Systems & Control Transactions

Research Article - Peer Reviewed Conference Proceeding
ESCAPE 35 - European Symposium on Computer Aided Process Engineering
Ghent, Belgium. 6-9 July 2025
Jan F.M. Van Impe, Grégoire Léonard, Satyajeet S. Bhonsale,
Monika E. Polańska, Filip Logist (Eds.)

PSE PRESS

A data-driven hybrid multi-objective optimization framework for pressure swing adsorption systems

Siyang Ma[a], and Jie Li[a]*

[a] Centre for Process Integration, Department of Chemical Engineering, School of Engineering, The University of Manchester, Manchester, the United Kingdom
* Corresponding Author: jie.li-2 @manchester.ac.uk.

ABSTRACT

Pressure swing adsorption (PSA) is an energy-efficient technology for gas separation, while the multi-objective optimization of PSA is a challenging task. To tackle this, we propose a hybrid optimization framework, which integrates three steps. In the first step, we establish surrogate models for the constraints using Gaussian processes (GPs) and employ multi-objective Bayesian optimization to search for feasible points that satisfy the constraints. In the second step, we establish surrogate models for the objective function and constraints using GPs and utilize constrained multi-objective Bayesian optimization to search for an approximate Pareto front. In the third step, we perform a local search based on the approximate Pareto front. By employing the trust region filter method, we construct quadratic models for each constraint and objective function and refine the Pareto front to achieve local optimality. This framework demonstrates the efficiency of Bayesian optimization and the local optimality of the trust region method. A comparison with the popular evolutionary algorithm, Nondominated Sorting Genetic Algorithm II (NSGA-II), showed that this framework had a higher hypervolume of the Pareto front while halving the runtime and reducing the number of simulations by a factor of 20.

Keywords: Pressure swing adsorption, data-driven optimization, machine learning, multi-objective optimization

INTRODUCTION

Optimizing large-scale partial differential algebraic equation (PDAE) systems has always been a challenging task in chemical engineering. In recent years, there has been an increasing demand for multi-objective optimization of PDAE systems with coupled constraints. Therefore, it is necessary to develop an optimization framework that can be used to solve such problems. The main challenges in optimizing PDAE systems with coupled constraints include:

1. The coupled constraints with PDAEs make it difficult to obtain feasible solutions.
2. Solving PDAEs is very expensive, and satisfactory Pareto fronts need to be obtained within a limited number of solution iterations.
3. The gradient is difficult to obtain, which greatly restricts the use of gradient-based algorithms.

Pressure swing adsorption (PSA) for gas separation is a typical PDAE problem and the most common research subject. Researchers preferred to use Artificial Neural Networks to construct surrogate models and use Nondominated Sorting Genetic Algorithm II (NSGA-II) [1] for solving it [2]. While employing this approach can lead to an enhanced Pareto front, it necessitates the utilization of numerous PDAE problem solutions for acquiring training samples. Recently, Hao et al. [3] developed a hybrid optimization framework efficiently solving unconstrained PSA optimization problems using the Thompson sampling efficient multi-objective optimization (TSEMO) algorithm and DyOS, but it cannot handle PSA problems with coupled constraints. Pini et al. [4] used penalty function method to handle coupled constraints, but the selection of penalty function parameters becomes a new issue.

Therefore, we will develop a hybrid optimization framework that combines the advantages of multi-objective Bayesian optimization and local search to handle large PDAE problems with coupled constraints. This

framework can obtain a satisfactory Pareto front within a limited budget and handle challenging coupled constraints.

PROBLEM STATEMENTS

A PSA process usually operates cyclically with at least three steps: high-pressure adsorption, low-pressure desorption and pressurisation. Additional steps, such as pressure equalisation, rinse, and purge, are often incorporated to improve process performance. Figure 1 presents an example of the four-step PSA cycle for CO_2/N_2 separation from Haghpanah et al. [5], which comprises feed pressurisation, adsorption, blowdown and evacuation.

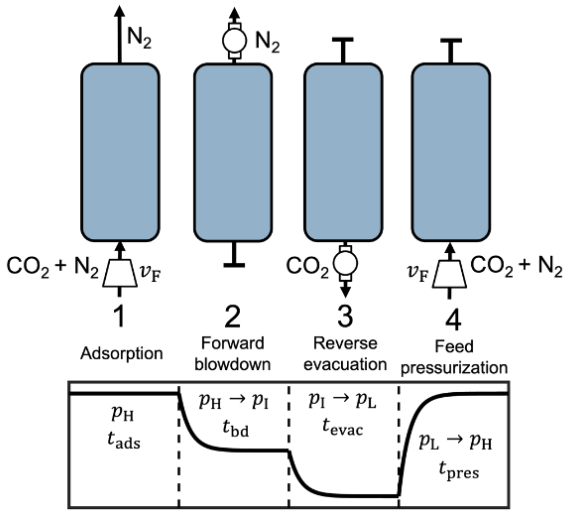

Figure 1. Schematic representation of the four-step PSA cycle for CO_2/N_2 separation

The problem can be stated as follows,

Given:

1. The number of adsorption bed, adsorbent layers.
2. Feed conditions such as flow rate, composition, temperature, and pressure.
3. Bed configuration, physical properties of the adsorbent and the corresponding adsorption isotherms.
4. Product purity requirement.

Determine:

1. The time for each step.
2. Operating conditions, such as operating pressures and flow velocity.

Assumptions:

- Gas phase follows ideal gas law.
- Non-isothermal adsorption process.
- The bulk flow pattern follows a plug flow model.
- The adsorption equilibrium is described by a dual-site Langmuir model with pure component isotherms dependent on gas phase concentration.
- The mass transfer is described by a linear driving force (LDF) model where the mass transfer resistance is controlled by the molecular diffusion in the macropores.
- The heat of adsorption of each gas component is constant.
- Instantaneous heat transfer between bulk fluid and adsorbent.
- Uniform bed with constant bed void fraction, particle porosity and particle size.
- No radial gradients in concentration, temperature, or pressure.

The objectives are to minimize energy cost and maximize productivity simultaneously while meeting cyclic steady-state (CSS) condition and requirements of product purity and recovery.

The dimensionless form of the PSA rigorous model equations from Liao et al. [6] is used in this work, which is a simplified version from that in Haghpanah et al. [5]. This is because the simplified model effectively captures the full dynamics of the PSA system without loss of accuracy whilst markedly improving simulation efficiency [6]. The PSA model equations are presented in Eqs(1-6). This model consists of PDEs that govern conservation of mass, momentum and bed energy, an ODE for calculating the fractional solid loading, complemented by the dual-site Langmuir adsorption isotherms and the Ergun equation that establishes the correlation between the pressure gradient and the local interstitial velocity. Two additional assumptions are made as follows in the model of Liao et al. [6]:

- The impact of axial dispersion is omitted.
- The column wall temperature is held constant at 298.15K.

$$\frac{\partial y_i}{\partial \tau} = -\frac{\bar{T}}{\bar{P}}\frac{\partial}{\partial Z}\left(\frac{y_i \bar{P}}{\bar{T}}\bar{v}\right) - \Psi\frac{\bar{T}}{\bar{P}}\frac{\partial x_i}{\partial \tau} - \frac{y_i}{\bar{P}}\frac{\partial \bar{P}}{\partial \tau} - \frac{y_i}{\bar{T}}\frac{\partial \bar{T}}{\partial \tau} \quad (1)$$

$$\frac{\partial \bar{P}}{\partial \tau} = -\bar{T}\frac{\partial}{\partial Z}\left(\frac{\bar{P}}{\bar{T}}\bar{v}\right) - \Psi\bar{T}\sum_{i=1}^{n_{comp}}\frac{\partial x_i}{\partial \tau} + \frac{\bar{P}}{\bar{T}}\frac{\partial \bar{T}}{\partial \tau} \quad (2)$$

$$\frac{\partial \bar{T}}{\partial \tau} = -\Omega_1\frac{\partial}{\partial Z}(\bar{v}\bar{P}) - \Omega_2\bar{T}\sum_{i=1}^{n_{comp}}\frac{\partial x_i}{\partial \tau} + \sum_{i=1}^{n_{comp}}(\sigma_i\frac{\partial x_i}{\partial \tau}) + -\Omega_3(\bar{T} - \bar{T}_w) - \Omega_1\frac{\partial \bar{P}}{\partial \tau} \quad (3)$$

$$\frac{\partial x_i}{\partial \tau} = \alpha_i(x_i^* - x_i) \quad (4)$$

$$q_i^* = \frac{q_{sb,i} b_i c_i}{1+\sum_{i=1}^{n_{comp}} b_i c_i} + \frac{q_{sd,i} d_i c_i}{1+\sum_{i=1}^{n_{comp}} d_i c_i} \quad (5)$$

$$\frac{\partial \bar{P}}{\partial Z} = -\frac{150}{4}\frac{1}{r_p^2}\left(\frac{1-\varepsilon}{\varepsilon}\right)^2 \frac{v_0 L}{P_0}\mu\bar{v} + \frac{1.75}{2}\frac{1}{r_p}\frac{1-\varepsilon}{\varepsilon}\frac{v_0^2 L}{P_0}\rho_g|\bar{v}|\bar{v} \quad (6)$$

All decision variables (denoted as x) must satisfy their respective lower and upper bounds (Table 1).

$$x_L \leq x \leq x_U \quad (7)$$

Table 1: Bounds of decision variables.

Variable	Lower bound	Upper bound
t_{ads} [s]	20	100
t_{bd} [s]	30	200
t_{evac} [s]	30	200
P_L [bar]	0.1	0.5
P_I [bar]	0.13	3
P_H [bar]	1	10
v_{feed} [m/s]	0.1	2

Some logical constraints may be imposed for the operating pressures to avoid unrealistic pressure values.

$$P_I \geq P_L + \Delta P \quad (8)$$

$$P_H \geq P_I + \Delta P \quad (9)$$

where ΔP is defined as the minimum allowable difference between P_L, P_I and P_H. In this work, ΔP =0.01bar. The objectives are to minimize energy cost and maximize productivity simultaneously while meeting CSS condition and requirements of product purity and recovery, which can be calculated below,

$$Purity = \frac{n_{CO_2,out}^{evac}}{n_{CO_2,out}^{evac}+n_{N_2,out}^{evac}} \times 100\% \quad (10)$$

where $n_{CO_2,out}^{evac}$ and $n_{N_2,out}^{evac}$ are the amounts of CO_2 and N_2 extracted during the evacuation step, respectively.

$$Recovery = \frac{n_{CO_2,out}^{evac}}{n_{CO_2,in}^{pres}+n_{CO_2,in}^{ads}} \times 100\% \quad (11)$$

where $n_{CO_2,in}^{pres}$ and $n_{CO_2,in}^{ads}$ are the amounts of CO_2 fed in the pressurization and adsorption steps.

$$CSS\ condition = st(t) - st(t+t_{cycle}) \leq \varepsilon \quad (12)$$

where $st(t)$ and $st(t+t_{cycle})$ is defined as the state variables at a given time t, and ε represents the tolerance vector.

$$Productivity = \frac{n_{CO_2,out}^{evac}}{V_{bed} \cdot t_{cycle}} \quad (13)$$

where V_{bed} is the volume of the adsorption bed, and t_{cycle} is the total amount of time required for one complete cycle.

$$Energy\ consumption = \frac{E_{ads}+E_{bd}+E_{evac}+E_{pres}}{m_{CO_2,out}^{evac}} \quad (14)$$

where E_{ads}, E_{bd}, E_{evac} and E_{pres} are the amounts of energy consumed by the four steps, and $m_{CO_2,out}^{evac}$ represents the mass of CO_2 extracted during the evacuation step.

The multi-objective optimization problem (denoted as MOOP) can be stated as follows,

$$\begin{aligned} (\mathbf{MOOP}) \quad & \max \quad Productivity \\ & \min \quad Energy\ consumption \\ \text{s.t.} \quad & \text{Eqs. (1)-(6)} \\ & \text{Eqs. (7)-(14)} \\ & \text{Purity} \geqslant 95\% \\ & \text{Recovery} \geqslant 90\% \\ & \text{Boundary conditions} \end{aligned}$$

The multi-objective optimization (MOOP) is difficult to solve due to the challenge of meeting two coupled constraints and the high cost of simulations. The two constraints, purity and recovery of CO_2 product, conflict with each other and are difficult to satisfy simultaneously. Moreover, the expensive PDAE solving process requires the use of efficient surrogate models to reduce optimization time. Therefore, it is necessary to develop a hybrid framework to efficiently optimize the PSA process design.

HYBRID OPTIMIZATION FRAMEWORK

A hybrid optimization framework is developed to solve the PSA problem. This framework combines multi-objective Bayesian optimization with a trust region algorithm, leveraging the strengths of both to solve expensive constrained multi-objective black-box problems within acceptable computational resources. Multi-objective Bayesian optimization (MOBO) is a global optimization algorithm that constructs surrogate models based on validated data points using a Gaussian process and updates the surrogate model by incorporating new validated points obtained in each iteration. In each iteration, an acquisition function is used as the evaluation criterion to determine the next sampled point for validation. Therefore, multi-objective Bayesian optimization can be used for efficient global search. Conversely, the trust region filter method (TRF) used in this work, constructs a quadratic surrogate model within a region defined by the trust radius to obtain a local optimal solution. The trust radius and center point are then adjusted based on the optimal solution, iteratively, until the hypervolume increment falls below a specified threshold. Therefore, this trust region algorithm is local, deterministic, and derivative-free. The hybrid framework combines these two approaches, executing a "rough to refined" search strategy to efficiently solve expensive black-box problems. Figure 2 shows the overall optimization process.

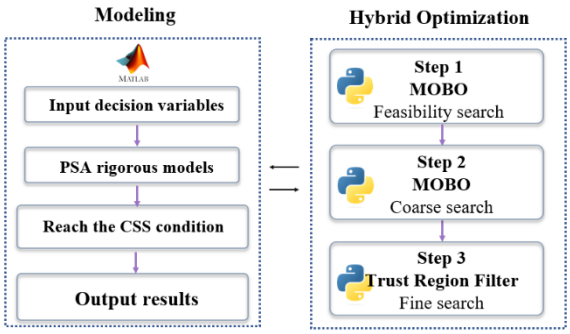

Figure 2. The overall PSA process optimization Schematic

The framework consists of three steps:

1. Searching for feasible points under unconstrained conditions using the multi-objective Bayesian optimization algorithm.
2. Searching for a rough Pareto front that satisfies the constraints based on the feasible points.
3. Obtaining an accurate Pareto front based on the rough Pareto front using the trust region method.

In the first step, our objective is to obtain feasible combinations of decision variables (feasible points) that satisfy the constraints. To achieve this, we treat the PSA problem as an unconstrained problem (in this phase, the decision variables remain unchanged, while the objective function becomes the two constraints: purity and recovery). In this way, we can use a multi-objective Bayesian optimization algorithm (e.g., TSEMO [7]) for solving. TSEMO first uses Latin hypercube sampling to randomly sample this black-box problem (in this experiment, 35 initial sampling points are extracted). After obtaining the starting point set, Gaussian surrogate models are constructed for each of the two constraints. After determining the acquisition function, NSGA-II is used for optimization. During the optimization process, we need to add two constraints (purity$\geqslant$95%, recovery$\geqslant$90%) to ensure that the obtained points satisfy the two constraints as much as possible. Then, new data points are validated, and the data set is updated. The new data set is used to construct a new Gaussian surrogate model, and this process repeats iteratively. After multiple iterations, we obtain a set of points that satisfy the constraints of purity and recovery.

In the second step, our objective is to obtain a rough Pareto front concerning energy consumption and productivity. We still use a specialized multi-objective Bayesian optimization (MOBO) to construct a surrogate model for solving this problem. First, we select a reference point P_{refer} (Productivity=0, Energy consumption=600kW·h/ton) and calculate the current hypervolume based on the feasible point set. Then, we use the entire data set to construct Gaussian surrogate models for purity, recovery, energy consumption, and productivity. The hypervolume improvement expectation with constraints is used as the acquisition function. This hypervolume improvement expectation can be expressed as:

$$EHVI_{con} = I_c \cdot EHVI \tag{15}$$

where I_c represents the probability of satisfying the purity and recovery constraints [8], and $EHVI$ represents the hypervolume improvement expectation under unconstrained conditions [9]. We use NSGA-II to optimize the acquisition function and determine the next evaluation point. After simulation, the data set, Gaussian surrogate model, and hypervolume are updated, and this process repeats iteratively. After multiple cycles, we obtain a rough Pareto front.

In the third step, our goal is to optimize the Pareto front. After the two steps of rough search, we have obtained points that satisfy the constraints and a rough Pareto front. However, there great gaps between the points on the Pareto front, and the hypervolume can improve further. Therefore, we need to use local search to further improve the Pareto front. We use the trust region method [10,11] to enhance the Pareto front. First, we select the points on the Pareto front and construct four local quadratic surrogate models (for purity, recovery, energy consumption, and productivity) within a predetermined trust radius. Based on these quadratic surrogate models, we construct a function for hypervolume increment and find the maximum value of hypervolume increment within the trust region. After validation, if the point significantly improves the hypervolume, the trust radius is increased; if the point slightly improves the hypervolume, the trust radius remains unchanged; if the point cannot improve the hypervolume, the trust radius is reduced. The termination conditions are set to an improvement less than 10^{-3}. After performing this process for each point in the trust region, the algorithm terminates, and we obtain the final Pareto front. Figure 3 shows the workflow of the hybrid optimization framework.

RESULTS AND DISCUSSION

As a comparison, we use NSGA-II to solve this problem to verify the computational efficiency and accuracy of this framework. In the hybrid optimization framework mentioned above, the first step employs the TSEMO algorithm implemented in MATLAB, while the second and third steps utilize optimization algorithms implemented in Python. The PSA numerical simulation model is solved in MATLAB, and the results are called through MALAB Engine API. All experiments were conducted on an Intel computer with 16 cores and 32GB memory. The Python version used was 3.10.0, and the MATLAB version was 2023a.

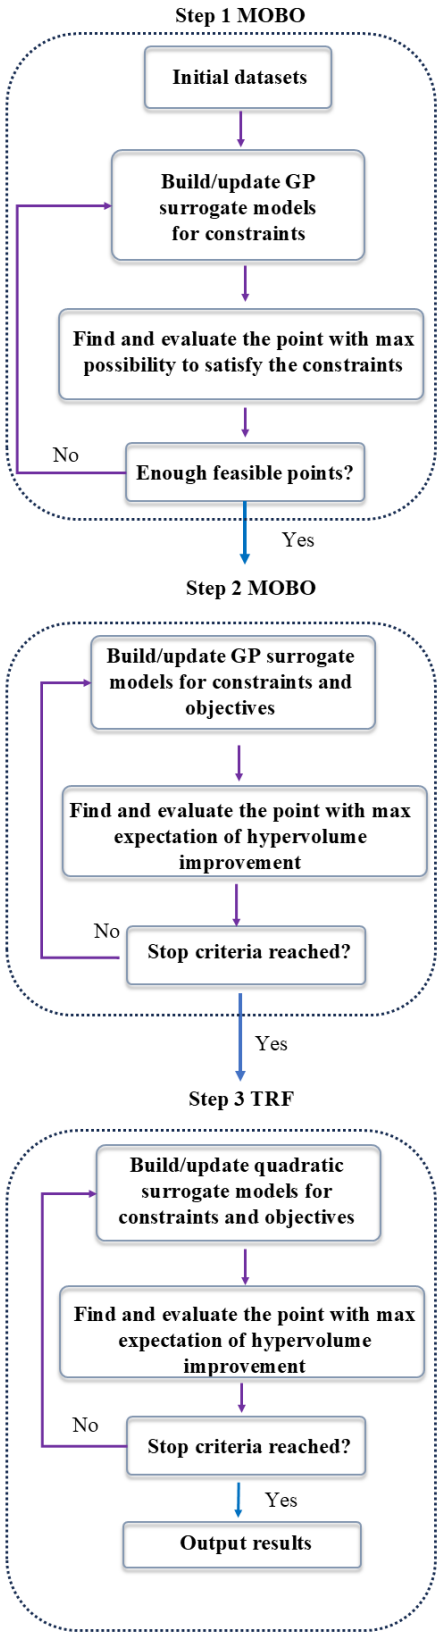

Figure 3. Workflow of the hybrid optimization framework

The first step of this hybrid optimization framework is to efficiently search for feasible points using multi-objective Bayesian optimization. In this step, we utilize the TSEMO algorithm as the external optimizer. To initialize the TSEMO algorithm, 35 sets of decision variables were randomly selected using Latin hypercube sampling and numerically simulated to obtain corresponding outputs. These data sets were used to train Gaussian process models. Based on these Gaussian processes, TSEMO evaluates the decision variable combinations that are most likely to satisfy the constraints and performs simulations. Subsequently, the simulation results are added to the entire data set to update the GP model for the next iteration. In this case, we conducted 300 simulations in the first step and recorded the increasing trend in the number of feasible points. When we conducted 300 simulations, we obtained 44 feasible points. However, conducting excessive simulations in the first step would prolong the time for building the Gaussian process model in the second step. Moreover, 44 feasible points are already sufficient to establish the hypervolume reference for productivity and energy consumption. To balance the efficiency of the second step's search, we only use 300 simulations in the first step before proceeding to the second step. In this step, with a wall time of 1.5 hours, an initial Pareto front with a hypervolume of 194.55 was obtained.

In the second step of this hybrid optimization framework, we optimize the original problem based on the feasible points obtained using multi-objective Bayesian optimization (MOBO). Firstly, we calculate the initial hypervolume reference value using the feasible points obtained in the first step. Then, we establish a multi-output Gaussian process model using all the data points. Using the $EHVI_{con}$ acquisition function, we search for the next set of decision variable combinations that are most likely to make progress. After performing simulations, we recalculate the hypervolume and update the data set and Gaussian process for the next iteration. In the second step, we conducted 600 simulations and obtained a rough Pareto front that cannot guarantee local optimality. As the number of simulations increases, the hypervolume also stabilizes and increases. However, after surpassing 500 simulations, the growth trend of the hypervolume noticeably slows down. It can be observed that when the number of simulations reaches 600, the efficiency of improving the hypervolume becomes extremely low. Therefore, we decided to stop the multi-objective Bayesian optimization at 600 simulations and proceed to the third step, where we use a local optimization algorithm to refine the Pareto front. In this step, with a wall time of 5.9 hours, a Pareto front with a hypervolume of 358.13 was obtained.

In the third step, the trust region filter method, a derivative-free local optimization method, will be employed

to enhance the hypervolume. The feasible points on the rough Pareto front obtained in the second step will serve as the starting points for the search. Sampling within the trust region of these points, local quadratic models will be constructed to optimize the hypervolume. After each iteration, the hypervolume will be recalculated, and the process will continue until the convergence condition is met. A total of 5737 simulations were conducted in the local search. In this step, with a wall time of 12.8 hours (8 cores), a Pareto front with a hypervolume of 429.28 was obtained.

To assess the performance of this hybrid optimization framework, we also employed NSGA-II to solve the problem. We utilized NSGA-II to solve the problem three times, selecting the best result from the obtained set. In the process of solving the problem with NSGA-II, a total of 600 iterations were executed. Each iteration consisted of a population of 250, calculated using 8 cores, taking a wall time of 44.3 hours. The resulting hypervolume of the Pareto front obtained was 373.73. Figure 4 shows the Pareto fronts of the hybrid optimization framework and NSGA-II. Table 3 presents a comparative analysis of the optimization performance between the hybrid optimization framework and NSGA-II. The results demonstrate that the hybrid approach yielded a 14% increase in hypervolume and a 62.5% augmentation in the number of Pareto front solutions relative to NSGA-II. The hybrid optimization framework demonstrates the ability to identify feasible points with an energy consumption close to 500 and a productivity of 5, showing superior global search capability compared to NSGA-II under limited wall time. Additionally, the search time of the hybrid optimization framework is only half of the wall time of NSGA-II, indicating significantly higher search efficiency. In summary, the hybrid optimization framework developed in this study effectively addresses the multi-objective optimization problem with coupled constraints.

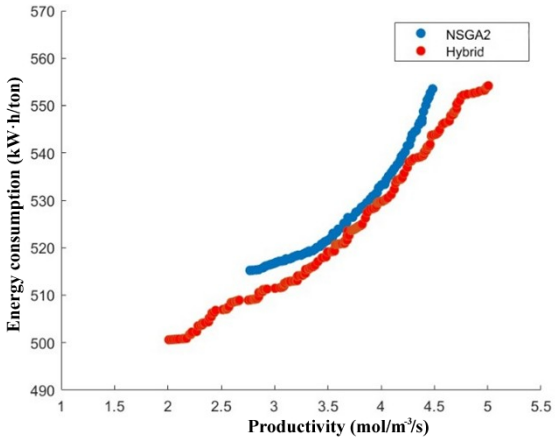

Figure 4. Hypervolume comparison between hybrid optimization framework and NSGA-II

Table 3: Performance Comparison

Method	Hybrid optimization framework	NSGA-II
Wall time (h)	20.2	44.3
CPU time (h)	109.8	364.4
Hypervolume	429.28	376.73
N_{PF}	143	88
N	6672	150000

N_{PF} represents the number of points on the Pareto front, and N is the number of rigorous model evaluations conducted.

Compared to the framework developed by Hao et al [3], this hybrid framework has two major advantages. Firstly, this method has improved the coarse search process, making it capable of handling multi-objective optimization problems with coupled constraints. On the other hand, Hao et al's framework [3] can only handle unconstrained multi-objective optimization problems. Secondly, this method uses the trust region method for fine search without using derivative information from dynamic systems, significantly improving computational efficiency. In contrast, Hao et al's framework [3] used the derivative-based DyOS algorithm for deterministic optimization. There are many limitations when using DyOS, including the reduction of decision variables and changes to CSS criteria, which may affect the results.

CONCLUSIONS

The main challenge in solving the design problems of chemical engineering processes, represented by the PSA process, is the high computational cost. Hence, it is imperative to develop an optimization framework capable of efficiently exploring the Pareto set for design problems. In this work, a hybrid optimization framework was developed, which can efficiently and quickly solve the CO_2 capture pressure swing adsorption problem. In the case study of PSA optimization design, this framework has achieved superior Pareto fronts while requiring only half the time compared to NSGA-II. In future research, we will focus on improving the coarse search step of this method to enhance the accuracy and efficiency of the coarse search while increasing compatibility with parallel computing.

ACKNOWLEDGEMENTS

Siyang Ma would like to thank Yangyanbing Liao and Andrew Wright for their assistance in constructing the PSA model. Jie Li would like to appreciate the financial support from UKRI Impact Acceleration Account (IAA 424).

REFERENCES

1. K. Deb et al. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* 6(2):182-197 (2002) https://doi.org/10.1109/4235.996017
2. N. S. Wilkins et al. Optimization of pressure-vacuum swing adsorption processes for nitrogen rejection from natural gas streams using a nitrogen selective metal organic framework. *Can. J. Chem. Eng.* 100(9):2374-93 (2022) https://10.1002/cjce.24469
3. Z. Hao et al. Efficient hybrid multiobjective optimization of pressure swing adsorption. *Chem. Eng. J.* 423:130248 (2021) https://doi.org/10.1016/j.cej.2021.130248
4. W. Adam & R. Pini. Efficient bayesian optimization of industrial-scale pressure-vacuum swing adsorption processes for CO2 capture. *Ind. Eng. Chem. Res.* 61(36):13650-13668 (2022). https://doi.org/10.1021/acs.iecr.2c02313
5. R. Haghpanah et al. Multiobjective optimization of a four-step adsorption process for postcombustion CO2 capture via finite volume simulation, *Ind. Eng. Chem. Res.* 52(11):4249–4265 (2013) https://doi.org/10.1021/ie302658y
6. Y. Liao et al. Simulation and optimisation of vacuum (pressure) swing adsorption with simultaneous consideration of real vacuum pump data and bed fluidisation. *Sep. Purif. Technol.* 358:130354 (2024) https://doi.org/10.1016/j.seppur.2024.130354
7. E. Bradford et al. Efficient multiobjective optimization employing Gaussian processes, spectral sampling and a genetic algorithm. *J. Glob. Optim.* 71(2): 407-438 (2018) https://doi.org/10.1007/s10898-018-0609-2
8. J.R. Gardner et al. Bayesian optimization with inequality constraints. *ICML* 2014:937-945 (2014) https://doi.org/10.5555/3044805.3044997
9. K. Yang et al. Multi-objective Bayesian global optimization using expected hypervolume improvement gradient. *Swarm Evol Comput.* 44:945-956 (2019) https://doi.org/10.1016/j.swevo.2018.10.007
10. B. Manuel et al. Multi-Objective Trust-Region Filter Method for Nonlinear Constraints using Inexact Gradients. arXiv preprint arXiv:2208.12094 (2022) https://doi.org/10.48550/arXiv.2208.12094
11. C. Cartis et al. Improving the flexibility and robustness of model-based derivative-free optimization solvers. *ACM Trans. Math. Software* 45(3):1-41 (2019) https://doi.org/10.48550/arXiv.1804.00154