Systems & Control Transactions

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

PSE PRESS

A Generative AI Approach to Inverse Design for Continuous Pharmaceutical Manufacturing

Consuelo Del Pilar Vega-Zambrano, Vassilis M. Charitopoulos*

Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, University College London, Torrington Place, London WC1E 7JE, UK
* Corresponding Author: v.charitopoulos@ucl.ac.uk

ABSTRACT

Continuous pharmaceutical manufacturing (CM) offers improved quality assurance, operational agility, and supply resilience, yet process development remains dominated by expensive trial-and-error experimentation and high-dimensional space exploration. Motivated by ICH Q13, we develop a generative inverse-design framework that maps target product quality to feasible process recipes for an integrated twin-screw wet granulation and segmented fluidized-bed drying line. The framework integrates three components: (i) a Conditional Variational Autoencoder (CVAE) generator that proposes process parameter sets conditioned on desired Critical Quality Attributes (CQAs), (ii) a Gaussian Process (GP) surrogate validator that screens candidates for manufacturing feasibility, and (iii) SHapley Additive exPlanations (SHAP) to interpret the generated designs. Training data were produced from a validated gPROMS digital twin of the Diamond Pilot Plant (DiPP) ConsiGma-25 line, covering liquid -to-solid ratio, drying temperature, drying time and air flowrate, with CQAs including granule moisture content, average particle size and porosity. The trained CVAE generated ~50, 000 candidate recipes and learned constrained feasible regions of the design space. Across seven operating scenarios, generated recipes achieved low deviations from targets. For a target-center case, deviations were 2.0% (moisture), 0.4% (particle size) and 0.5% (porosity). Edge cases remained acceptable, with the largest deviation observed for moisture in a high-moisture scenario (8.8%). SHAP analysis highlighted drying time and liquid-to-solid ratio as the dominant drivers of moisture, while liquid-to-solid ratio governed particle size and porosity. Overall, the approach enables rapid, explainable exploration of CM design spaces, reducing experimental burden and supporting QbD-aligned development for integrated continuous processes.

Keywords: Design space, Generative Artificial Intelligence, Pharmaceutical manufacturing, Inverse Design, Quality by Digital Design, Conditional Variational Autoencoder

INTRODUCTION

Continuous manufacturing (CM) is increasingly adopted in the pharmaceutical sector to improve product quality, manufacturing agility and supply resilience. Unlike traditional batch processing, CM operates closer to a steady state and can enable tighter coupling between process analytical technology, advanced process control and real-time release testing. This evolution is supported by the broader shift toward digital applications (online/real-time, model-based decision support and control) and by the growing role of AI across the pharmaceutical lifecycle [1-4]

Regulatory agencies have actively encouraged CM adoption; in particular, ICH Q13 consolidates scientific and regulatory considerations for CM development, implementation, operation and lifecycle management, and emphasizes systematic approaches to define and maintain a state of control [5] Within the Quality-by-Design (QbD) paradigm (e.g., ICH Q8), the central engineering task is to connect Critical Process Parameters (CPPs) to Critical Quality Attributes (CQAs) and to identify a design space (DS) that guarantee product quality [6]

Despite these advantages, CM process development and scale-up remain resource intensive. Integrated drug-product flowsheets frequently involve several unit operations. As a result, CPP-CQA relationships are highly nonlinear and can exhibit strong cross-unit interactions.

A representative example is the integrated twin-screw wet granulation and segmented fluidized-bed drying line used in continuous tableting: granulation conditions govern nucleation, growth and consolidation through liquid bridge formation and shear, while drying conditions determine moisture removal and influence granule structure and porosity [7]. Because these mechanisms interact across units, traditional design-of-experiments (DoE) campaigns can become prohibitively expensive, requiring many experiments (or high-fidelity simulations) to explore multidimensional spaces and capture cross-unit effects.

Inverse design reframes this workflow by treating product specifications as the starting point. Instead of asking "what CQAs result from these CPPs?", the inverse problem asks "which CPPs can achieve a target set of CQAs?". This aligns naturally with QbD because target CQAs and acceptable ranges define the desired product profile, and the process recipe should be selected from within a feasible and validated design space. However, inverse design in CM is challenging: several distinct CPP combinations can yield similar CQAs (a one-to-many mapping), feasible regions can be non-convex, and manufacturing constraints (equipment limits, residence-time windows, safe operating envelopes) are difficult to encode explicitly.

Deep generative models offer an alternative route to solve inverse problems. Variational autoencoders (VAEs) learn probabilistic latent representations enabling stable training and sampling of new candidates through latent-space exploration [8].Conditional variants (CVAEs) extend VAEs by learning conditional distributions given labels or targets, providing direct controllability for inverse design [9]. In materials discovery, deep generative models have been widely explored to accelerate inverse design by encoding structure-property relationships into latent spaces that can be navigated to propose candidates with targeted characteristics [10]. In pharmaceutical sciences, generative models have achieved notable success in drug discovery and are now being explored for digital formulation optimization[11-12]. Related work in particle engineering has shown that CVAEs can generate particle geometries to match target sphericity and packing fraction, illustrating the suitability of conditional generation for constrained inverse tasks [13]

Despite this progress, the application of generative inverse design to integrated, industrial-style CM unit-operation networks remain limited. Existing CM studies often focus on data augmentation. For example, Ma et al. combined GAN-based data augmentation with surrogate models and multi-objective optimization for an integrated synthesis-crystallization case study [14]. In GMP-relevant CM, however, robustness, feasibility screening and interpretability are essential: candidate recipes must remain within validated operating envelopes and the rationale for recipe selection should be explainable to support process understanding and lifecycle management.

In this work, we propose a three-component framework tailored to CM inverse design. First, a CVAE acts as a generator that learns the conditional distribution of CPP recipes given target CQAs, enabling controllable, one-to-many generation. Second, an independent Gaussian Process (GP) surrogate acts as a validator to screen generated candidates for feasibility and to quantify predictive uncertainty; this reduces the risk of selecting recipes that exploit model artefacts or extrapolate beyond the validated region. Third, we apply SHapley Additive exPlanations (SHAP) [15] to the surrogate validator (GP model) that maps CPPs to CQAs. The resulting feature attributions explain how changes in CPPs drive the validator's predicted CQAs and can be used to interpret candidate CPP sets generated by the CVAE.

We demonstrate the framework using a validated gPROMS flowsheet model representative of the Diamond Pilot Plant (DiPP) ConsiGma-25 continuous tableting line, integrating twin-screw wet granulation with segmented fluidized-bed drying [7-16] .The case study covers four CPPs (liquid-to-solid ratio, drying temperature, drying time and air flowrate) and three CQAs (granule moisture content, average particle size and porosity). Results show that the CVAE can generate large sets of candidate recipes that cluster around desired CQAs while maintaining operational diversity, and that feasibility and interpretability layers provide robust screening and actionable process insight. Together, these capabilities support data-efficient CM development aligned with ICH Q13 and QbD principles.

METHODOLOGY

Figure 1 summarizes the Conditional Variational Autoencoder (CVAE) architecture adopted for inverse design. The encoder maps process parameters together with target CQAs into a low-dimensional latent representation, and the decoder reconstructs process parameters conditioned on the target CQAs. Sampling in latent space enables generation of multiple diverse recipes for the same target product profile.

Data generation and preprocessing

Training data were generated using a validated gPROMS FormulatedProducts flowsheet model of the integrated twin-screw wet granulation and segmented fluidized-bed drying line [7-16]. A space-filling Sobol sequence was used to sample the CPP design space efficiently and to avoid clustering typical of purely random sampling [17]. A dataset of 300 formulations was used with 210 for training and 90 for validation. The same dataset was used to train the GP and CVAE models. Training on the same data as the CVAE ensures consistency in the

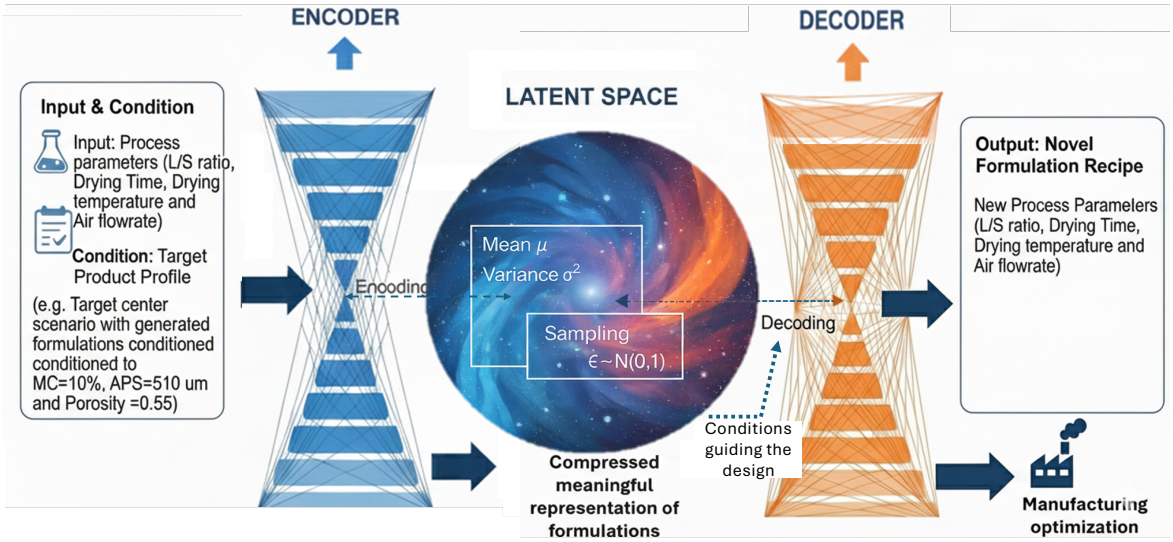

Figure 1. Conditional Variational Autoencoder (CVAE) architecture adapted for inverse design.

modeling framework and leverages the full available training information for robust surrogate predictions. For the GP model, the validation dataset was used to evaluate the GP's generalization performance on unseen data.

Inputs and outputs were standardized prior to training; CQAs were scaled to facilitate conditioning and multi-output learning.

CVAE generator for inverse design

We formulate inverse design as learning the conditional distribution $p(\boldsymbol{u}|\boldsymbol{y})$, where $\boldsymbol{u}$ denotes CPPs and $\boldsymbol{y}$ denotes target CQAs. The CVAE is trained by minimizing a composite loss function (Eq. 1) consisting of an MSE reconstruction term and a Kullback-Leibler divergence regularizer [8], [9]

$$L_{CVAE}(\theta,\ \varphi;\ x,\ y) = L_{recon}(x,\ \hat{x}) + \beta(D_{KL}(q_{\varphi}(z \mid x,\ y) \mid\mid p(z))) \quad (1)$$

Where:

- $\boldsymbol{x}$: vector of critical process parameters (CPPs), scaled to the range [0, 1].
- $\hat{\boldsymbol{x}}$: reconstructed vector of critical process parameters (CPPs).
- $L_{recon}\ (\boldsymbol{x},\hat{\boldsymbol{x}})$: reconstruction loss that enforces similarity between the original CPPs and the reconstructed CPPs. In this work, the mean squared error (MSE) is used.
- $\boldsymbol{y}$: vector of conditioning critical quality attributes (CQAs), representing the desired product quality targets.
- $q_{\varphi}(\mathbf{z} \mid \mathbf{x}, \mathbf{y})$: approximate posterior distribution learned by the encoder network. It is assumed to follow a multivariate Gaussian distribution with diagonal covariance
- p(**z**): prior distribution over the latent variables, defined as a standard normal distribution N(0, 1).
- β: weighting coefficient controlling the trade-off between reconstruction accuracy and latent space smoothness/continuity of the latent space.

Latent space and regularisation

The CVAE introduces a low-dimensional latent variable z that captures the unobserved factors and one-to-many variability in feasible CPP recipes that can produce the same target CQAs. This is important for inverse design because multiple CPP combinations can satisfy a single product specification; sampling different z values allows the model to generate diverse candidate recipes for the same y, rather than collapsing to one deterministic solution.

To make latent-space sampling meaningful, the latent distribution must be well-structured. The Kullback–Leibler (KL) divergence term regularises the encoder's approximate posterior $q\varphi(z|x,\ y)$ toward a simple prior $p(z)=N(0,\ 1)$. This discourages the encoder from memorising training samples (overfitting) and helps avoid irregular latent regions ("holes") where sampling would decode to unrealistic CPPs. As a result, the decoder can reliably generate feasible recipes by sampling $z\sim p(z)$ and conditioning on the desired CQAs y.

A key practical requirement in CM is recipe diversity under constraints. By sampling the latent space, the CVAE generates a distribution of recipes rather than a single optimum. This flexibility can support operational preferences, robustness to disturbances and equipment

limitations while still meeting target CQAs.

GP validator for feasibility assessment

To screen AI-generated recipes, we train a Gaussian Process (GP) surrogate model that predicts CQAs from CPPs and provides predictive uncertainty. Candidate recipes are accepted when predicted CQAs fall within specification windows and uncertainty remains below a threshold. This validator step reduces the risk of selecting recipes outside the validated region and provides an additional safety layer for GMP-relevant deployment.

Explainability with SHAP

To support interpretability, we apply SHapley Additive exPlanations (SHAP) to the surrogate validator (GP) used to predict CQAs from CPPs. SHAP decomposes each predicted CQA into additive contributions from individual CPPs, providing global trends (feature importance across the design space) and local explanations for specific CPP candidates. In the continuous manufacturing context, this supports "Quality by explainability" [18] by identifying which CPPs most strongly drive each predicted CQA and therefore should be prioritized for monitoring and control.

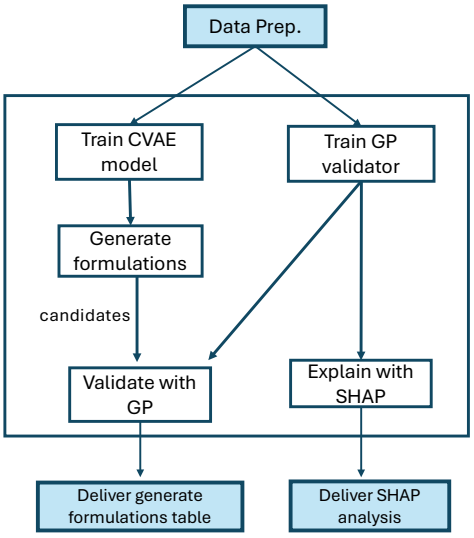

Figure 2. Generative inverse design workflow: CVAE generation.

CASE STUDY: INTEGRATED WET GRANULATION AND DRYING

The case study considers an integrated twin-screw granulator coupled with a segmented fluidized-bed dryer from the ConsiGma-25 continuous tableting line at the Diamond Pilot Plant (DiPP) [7]. The process model, implemented in gPROMS, captures nonlinearities and cross-unit couplings relevant to granule formation and moisture removal.

Four CPPs are considered: liquid-to-solid ratio (L/S), drying temperature (Tem), drying time (DT) and air flowrate (AF). Three CQAs are targeted: granule moisture content (MC), average particle size (APS) and porosity (Poro).

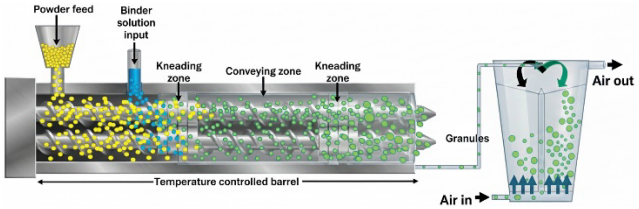

Figure 3. Integrated twin-screw wet granulator and segmented fluidized-bed dryer (ConsiGma-25 line, DiPP).

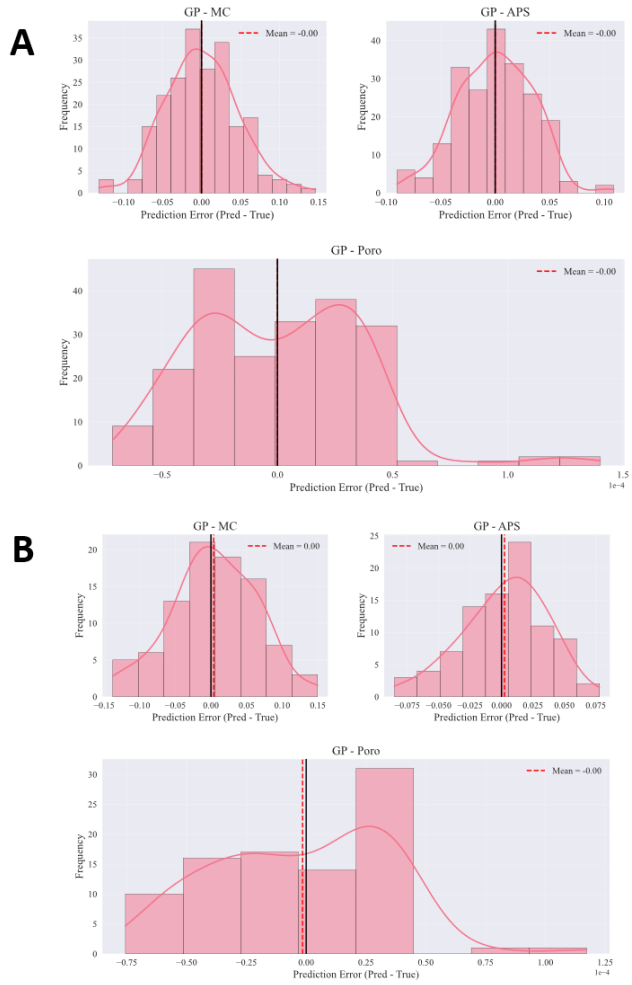

Figure 4. Prediction error distribution for the Gaussian Process validator. A. Training B. Validation

The inverse design task is to generate CPP recipes that meet specified targets (e.g., MC≈10%, APS≈510 µm, Poro≈0.55) while remaining operable for the integrated process.

RESULTS AND DISCUSSION

Validator accuracy

The GP validator achieved high accuracy on training and validation data, with prediction errors tightly centered around zero for all CQAs (Fig. 4).

Explainability and process insight

SHAP analysis (Fig. 5) indicates that drying time (DT) and liquid-to-solid ratio (L/S) exert the strongest influence on moisture content, consistent with heat and mass transfer fundamentals in drying.

In contrast, APS and porosity are primarily governed by L/S, aligning with granulation theory where liquid availability controls nucleation, coalescence and consolidation. These explanations help translate data-driven recipes into actionable engineering understanding and support monitoring and control strategies centered on the most influential CPPs.

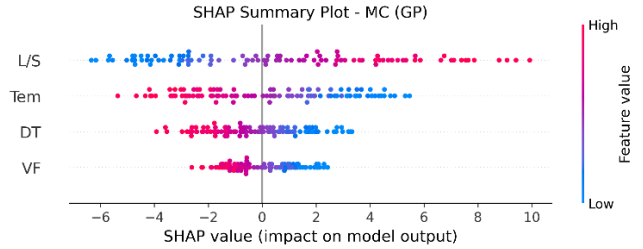

Figure 5. SHAP summary plot highlighting dominant CPP drivers for the CQAs.

Design-space coverage and conditional generation

The trained CVAE generated approximately 50, 000 candidate recipes. Figure 6A shows that generated recipes explore the broader CPP design space while clustering into feasible regions that correspond to target formulations. The generator naturally produced multiple distinct CPP combinations that meet the same target CQAs, illustrating the one-to-many nature of the inverse design

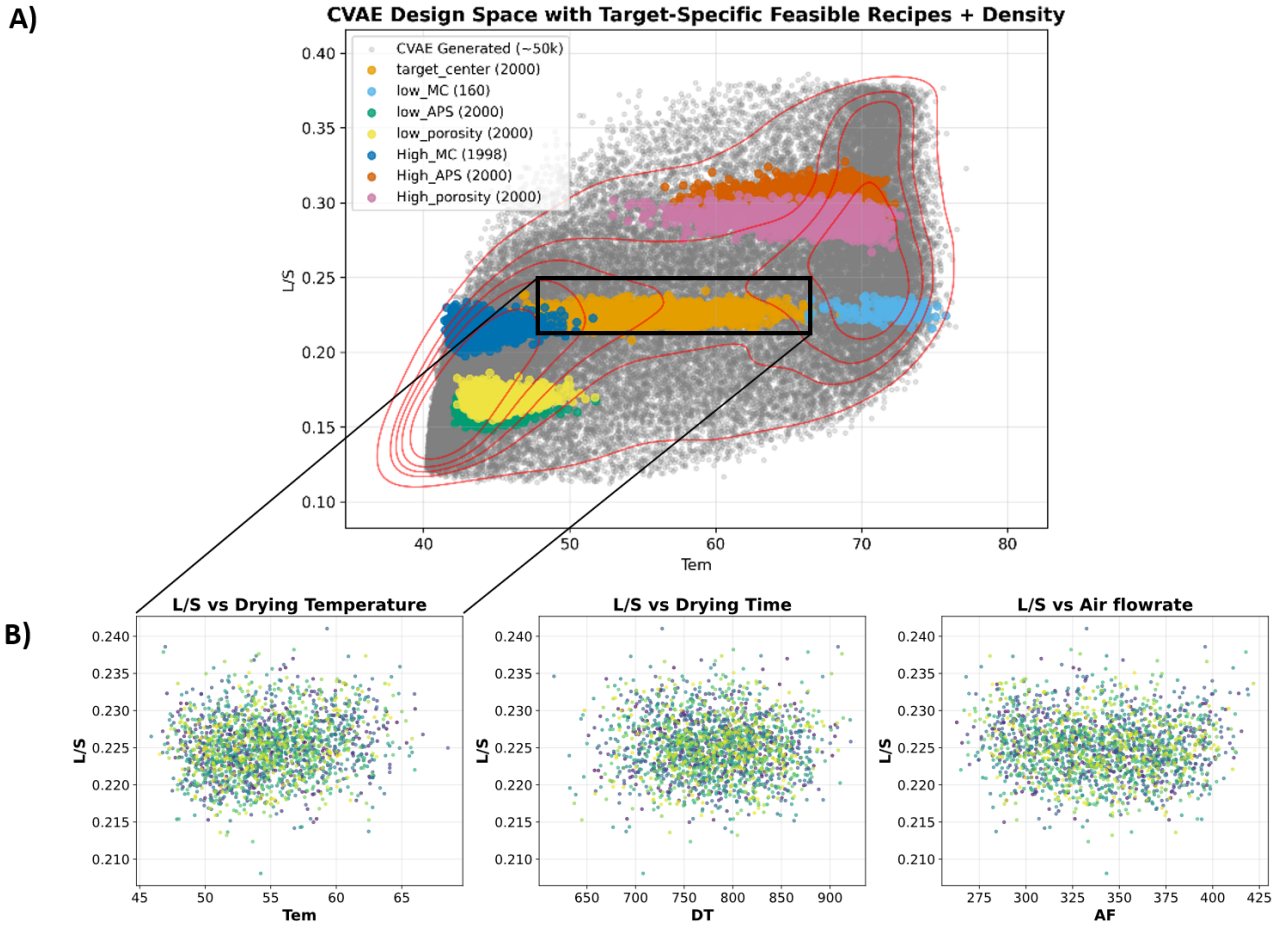

Figure 6. A. CVAE design-space exploration with clustered target formulation regions (colored) within the broader feasible set (grey). B. Close-up of Target center scenario with generated formulations conditioned to MC=10%, APS=510 um and Porosity =0.55.

mapping.

This inverse-design workflow complements established DS determination and uncertainty-aware DS methodologies by providing a generative mechanism to propose diverse CPP candidates that can then be validated within a DS framework [19-23]

Figure 6B illustrates example generated recipes conditioned on a representative target profile (MC=10%, APS=510 µm, Poro=0.55). Recipes remain diverse in CPP space while producing similar CQAs after feasibility screening, enabling operational flexibility.

Robustness across operating scenarios

Robustness was tested across seven operating scenarios that probe challenging specification demands (including low/high moisture, low/high APS, and low/high porosity).

Figure 7 summarizes percentage deviations from targets across scenarios. The target-center scenario achieved deviations of 2.0% (MC), 0.4% (APS) and 0.5% (Poro). The low-moisture scenario also performed well (3.1%, 0.3%, 0.7%).

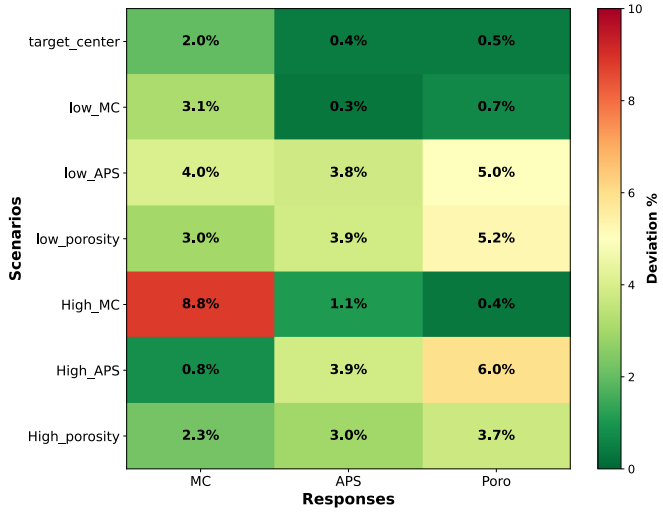

Figure 6. Deviation study for generated recipes across seven operating scenarios.

The most challenging case was high moisture content, where MC deviation increased to 8.8% while APS and porosity remained close to targets (1.1% and 0.4%). Overall, results indicate that the CVAE can learn constrained feasible regions and produce recipes that satisfy tight specifications across a range of industrially relevant conditions.

CONCLUSIONS AND FUTURE WORK

We presented a generative inverse-design framework for integrated continuous pharmaceutical manufacturing that combines a CVAE generator, a GP feasibility validator and SHAP-based explainability. Using a realistic wet granulation-drying case study, the framework generated large sets of candidate recipes (~50k) and achieved low deviations from targets across seven scenarios (best-case: target-center: 2.0% MC, 0.4% APS, 0.5% porosity; worst-case: 8.8% MC in the high-moisture scenario). Explainability analysis provided physically consistent insights, highlighting drying time and liquid-to-solid ratio as key levers for moisture control and L/S as the dominant driver for particle size and porosity.

Future work will focus on hybrid generative architectures, as well as experimental validation of selected recipes on pilot-scale equipment.

ACKNOWLEDGEMENTS

Financial support from EPSRC grants EP/V051008/1, EP/V034723/1, EP/W003317/1 and Peruvian Government through PRONABEC is gratefully acknowledged

AUTHOR IDENTIFIERS

Author ORCIDs:
Vega-Zambrano CDP: 0000-0003-0155-2088
Charitopoulos VM: 0000-0001-9051-917X

REFERENCES

1. Pantelides CC, Pereira FE. The future of digital applications in pharmaceutical operations. Current Opinion in Chemical Engineering 45:101038 (2024). https://doi.org/10.1016/j.coche.2024.101038
2. Jayakrishnan A, Phang HC, Ying Tan VX, Kee PE, Loke YH, Mod Razif MRF, Yee KM, Md Noh SM, Ming LC, Gan SH, Liew KB. Towards pharmaceutical industry 5.0: impact of artificial intelligence in drug discovery and development. CPB 27: (2025). https://doi.org/10.2174/0113892010396035250710120501
3. Nagy B, Galata DL, Farkas A, Nagy ZK. Application of artificial neural networks in the process analytical technology of pharmaceutical manufacturing—a review. AAPS J 24: (2022). https://doi.org/10.1208/s12248-022-00706-0
4. Destro F, Barolo M. A review on the modernization of pharmaceutical development and manufacturing – trends, perspectives, and the role of mathematical modeling. International Journal of Pharmaceutics 620:121715 (2022). https://doi.org/10.1016/j.ijpharm.2022.121715
5. International Conference on Harmonisation (ICH). Continuous Manufacturing of Drug Substances and Drug Products, Q13 (2022).
6. International Council for Harmonisation (ICH). ICH

Q8(R2): Pharmaceutical Development (2009).
7. Wang LG, Omar C, Litster J, Slade D, Li J, Salman A, Bellinghausen S, Barrasso D, Mitchell N. Model driven design for integrated twin screw granulator and fluid bed dryer via flowsheet modelling. International Journal of Pharmaceutics 628:122186 (2022). https://doi.org/10.1016/j.ijpharm.2022.122186
8. Kingma DP, Welling M. Auto-encoding variational bayes. arXiv:1312.6114 (2013) https://arxiv.org/abs/1312.6114
9. Sohn K, Lee H, Yan X. Learning structured output representation using deep conditional generative models. Adv Neural Inf Process Syst 28 (2015).
10. Fuhr AS, Sumpter BG. Deep generative models for materials discovery and machine learning-accelerated innovation. Front. Mater. 9: (2022). https://doi.org/10.3389/fmats.2022.865270
11. Tong X, Liu X, Tan X, Li X, Jiang J, Xiong Z, Xu T, Jiang H, Qiao N, Zheng M. Generative models for de novo drug design. J. Med. Chem. 64:14011-14027 (2021). https://doi.org/10.1021/acs.jmedchem.1c00927
12. Hornick T, Mao C, Koynov A, Yawman P, Thool P, Salish K, Giles M, Nagapudi K, Zhang S. In silico formulation optimization and particle engineering of pharmaceutical products using a generative artificial intelligence structure synthesis method. Nat Commun 15: (2024). https://doi.org/10.1038/s41467-024-54011-9
13. Qian Y, Li S. Inverse design of particle shapes with target sphericity and packing fraction using variational autoencoders. Engineering Applications of Artificial Intelligence 162:112509 (2025). https://doi.org/10.1016/j.engappai.2025.112509
14. Ma Y, Li W, Liu J, Shang G, Yang H, Gong J, Nagy ZK, Benyahia B. Digital design and optimization of the integrated synthesis and crystallization process using data?driven approaches. AIChE Journal 71: (2025). https://doi.org/10.1002/aic.18931
15. Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. Adv Neural Inf Process Syst 30 (2017). https://doi.org/10.48550/arXiv.1705.07874
16. Siemens Process Systems Engineering, 2025. gPROMS.https://www.psenterprise.com/products/gproms.
17. Bratley P, Fox BL. Algorithm 659. ACM Trans. Math. Softw. 14:88-100 (1988). https://doi.org/10.1145/42288.214372
18. Vega-Zambrano C, Diangelakis NA, Charitopoulos VM. Data-driven model predictive control for continuous pharmaceutical manufacturing. International Journal of Pharmaceutics 672:125322 (2025). https://doi.org/10.1016/j.ijpharm.2025.125322
19. Geremia M, Bezzo F, Ierapetritou MG. Design space determination of pharmaceutical processes: effects of control strategies and uncertainty. European Journal of Pharmaceutics and Biopharmaceutics 194:159-169 (2024). https://doi.org/10.1016/j.ejpb.2023.12.008
20. Bano G, Facco P, Ierapetritou M, Bezzo F, Barolo M. Design space maintenance by online model adaptation in pharmaceutical manufacturing. Computers & Chemical Engineering 127:254-271 (2019). https://doi.org/10.1016/j.compchemeng.2019.05.019
21. Meng Q, Bogle D, Charitopoulos VM. Probabilistic design space exploration and optimization via bayesian approach for a fluid bed drying process. European Journal of Pharmaceutical Sciences 210:107116 (2025). https://doi.org/10.1016/j.ejps.2025.107116
22. Chiplunkar R, Pauzi SM, Sachio S, Papathanasiou MM, Kontoravdi C. Probabilistic design space identification for upstream bioprocesses under limited data availability?. Systems and Control Transactions 4:2561-2567 (2025). https://doi.org/10.69997/sct.166359
23. Boukouvala F, Muzzio FJ, Ierapetritou MG. Design space of pharmaceutical processes using data-driven-based methods. J Pharm Innov 5:119-137 (2010). https://doi.org/10.1007/s12247-010-9086-y

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