A Model Predictive Control-Inspired Framework for Generative Multi-Objective Chemical and Materials Design

Haralambos Sarimveis[1], Argyri Kardamaki[1], Vasileios Minadakis[1], Christos Paschalis[1], Philip Doganis[1]
[1]National Technical University of Athens, School of Chemical Engineering, 9 Iroon Polytechniou, 15772, Athens, Greece

hsarimv@mail.ntua.gr

Introduction

Generative artificial intelligence (GenAI) has made significant progress in enabling the exploration of vast chemical spaces and the proposal of novel molecular structures. Architectures such as variational autoencoders (VAEs), generative adversarial networks (GANs), reinforcement learning (RL)-based models, and transformer-based sequence generators have demonstrated strong potential for de novo molecular design by learning patterns from existing chemical datasets. In many discovery workflows, these approaches are coupled with property prediction models to guide generation toward candidates with desirable physicochemical properties through goal-directed optimization strategies. While effective, such methods often require modifying or fine-tuning the generative model itself and may be tightly coupled to specific model architectures, limiting flexibility when integrating different generators and predictive models [1]. In this work, we introduce a framework inspired by Model Predictive Control (MPC) that guides molecular design by steering the generative process toward sequences that satisfy predefined property objectives. Rather than adapting the generator's parameters, the proposed method directs exploration toward promising regions of chemical space through iterative evaluation and feedback. This strategy enables flexible integration of different sequential generative models and property prediction tools within the same framework, facilitating modular and scalable molecular design workflows.

Methods

The proposed algorithm is designed for sequential molecular generators that construct molecules through token-based representations. Drawing inspiration from the principles of MPC, each molecule is treated as a sequence of tokens, and the current system state corresponds to the molecular prefix (i.e., the partial sequence generated so far). The control action is defined as the selection of the next token, which extends the prefix and transitions the system to the next state.

At each iteration, a set of the most probable candidate next tokens is obtained by the generator. For each candidate, multiple complete molecules are sampled in the form of 'prefix + candidate token + remaining sequence', forming a rollout set that represents possible future design trajectories. These generated molecules are evaluated via a user-defined scoring function that aggregates multiple design objectives based on predictions from property models. The scoring function encodes the desired objective for each property (e.g., minimization, maximization, or target range), assigning rewards when objectives are satisfied and penalties otherwise. For minimization and maximization objectives, reward values are scaled to the interval [0,1] using statistics derived from the generator's training data, with thresholds relaxed by a tolerance proportional to the standard deviation of the training distribution. In contrast, range-based objectives assign maximum reward (1) within a specified interval, with rewards decaying outside the bounds and approaching zero further from the target region.

Each candidate token is evaluated based on the highest scoring molecule within its rollout set, and the token with the best score is selected. To stabilize the sequential decision process and retain high-quality solutions across iterations, the best design identified so far is propagated forward and included in subsequent evaluations. The selected token is then appended to the prefix, and the process is repeated for a predefined number of iterations or until sequence termination. Once the optimized prefix has been identified, a final batch of molecules is generated from that prefix, scored using the same objective function, and ranked to produce the final pool of top candidate designs.

Case Study

Generator:
For the implementation of the proposed framework, the S4 molecular generator introduced in [2] was employed. The publicly available implementation for de novo molecular design was used together with a pretrained model provided through the associated Zenodo repository. Specifically, the pretrained model trained on the ChEMBL v31 dataset, a large curated database of bioactive drug-like molecules widely used for training molecular generative models, was selected as the base generator.

Targeted properties:
The selected case study is motivated by the proposed restriction of octocrylene by the European Chemicals Agency (ECHA) [3]. The restriction report outlines key physicochemical characteristics of suitable alternatives to octocrylene. Based on these guidelines, three properties were selected and corresponding objectives were defined to characterize the desired chemical space:

1) Maximum absorbance wavelength (λ_{max}): target range of 290–310 nm.
2) Absorbance factor (*f*): objective to maximize UV absorption efficiency.
3) Partition coefficient n-octanol/water (log_{KOW}): minimum value of 5 to ensure adequate lipophilicity.

The targeted properties are predicted using graph neural network models trained on curated datasets of experimentally measured values [4,5].

Results and discussion

The proposed algorithm is evaluated on the three-objective design task described above, using scoring function weights of [0.4, 0.2, 0.4] for the respective objectives. For each run, the algorithm is executed for 12 rounds, followed by generation of 2000 designs from the optimized prefix. After duplicate removal, the top-50 designs are selected based on the scoring function ("Targeted"). For comparison, an unguided baseline is constructed by randomly sampling 2000 molecules, followed by duplicate removal, scoring, and selection of the top-50 designs ("Random"). The procedure is repeated over 600 independent runs, and the resulting top-50 sets are aggregated with duplicates removed. A training baseline is also constructed by selecting top-ranked molecules from the generator's training dataset to match the size of the random pool ("Training").

A design is considered a 'hit' if it satisfies both range constraints for λ_{max} and log_{KOW}. Figure 1 shows the hit rate of targeted designs grouped by the first token selected during generation. The labels on the x-axis report the corresponding hit rates (i.e., the proportion of designs satisfying said constraints). Among the most probable initial tokens, "C" consistently yields the highest hit rates, indicating that early token selection strongly influences the feasibility of the final designs.

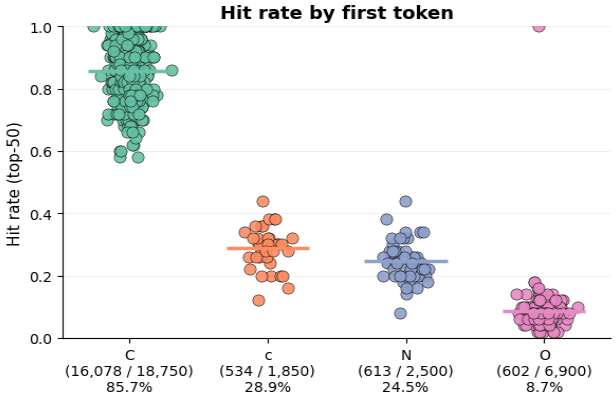

Figure 1. Hit rate of generated designs satisfying λ_{max} and log_{KOW} constraints, grouped by first token.

To further analyze this effect, designs originating from runs initialized with "C" are isolated. Given that two of the three objectives are enforced by restricting to hits, the remaining objective (*f*) is evaluated over the resulting subset. Figure 2 presents the distribution of *f* for the training, random, and targeted hits.

The targeted approach consistently produces higher *f* values compared to the other two baselines, while also yielding a substantially larger number of hits. This demonstrates that the proposed method not only improves the likelihood of satisfying the constraints but also effectively optimizes performance within the feasible region.

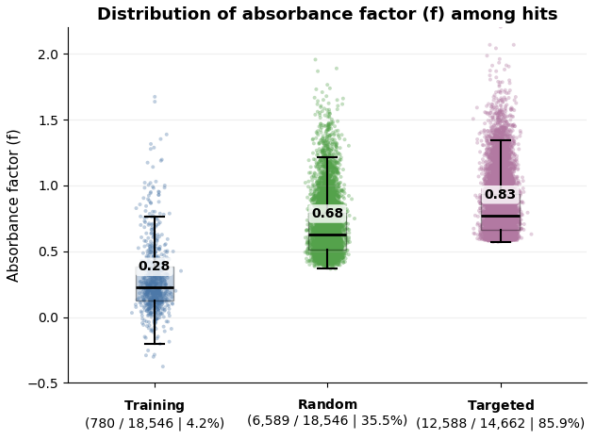

Figure 2. Distribution of absorbance factor (*f*) among hits (λ_{max} and log_{KOW} satisfied) for training baseline, random, and targeted generation.

Conclusions

This work introduces an MPC-inspired framework for guiding sequential molecular generation without retraining the underlying model. The proposed approach identifies molecule sequences that steer the generation process toward regions of chemical space where range constraints are satisfied, while simultaneously promoting other desired objectives. The proposed method can be extended to other types of materials for which a sequence-based generative model is available.

Acknowledgements: This work has been funded by the European Union's Horizon Europe Research and Innovation programme under project PINK (Grant Agreement No. 101137809).

References

[1] van den Broek R. L., Patel S., van Westen G. J. P., Jespers W., Sherman W., Journal of Chemical Information and Modeling, 18 (2025) 9383–9397.
[2] Özçelik, R., de Ruiter, S., Criscuolo, E., Grisoni, F., Nature Communications 15 (2024) 6176.
[3] European Chemicals Agency (ECHA), Annex XV Restriction Report: Proposal for a Restriction – Octocrilene, (2025).
[4] Beard, E.J., Sivaraman, G., Vázquez-Mayagoitia, Á., Vishwanath, V., Cole., J.M., Scientific Data 6 (2019) 307.
[5] Wang, H., Liu, W., Chen, Wang, Z., Data Science 44 (2023) 16906-16917.