HIT EXPANSION VIA LOCALIZED EXPLORATION OF SYNTHESIZABLE CHEMICAL SPACE

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative models for drug design which directly produce synthetic pathways have gained significant popularity due to their ability to constrain the search space to synthetically accessible molecules. However, existing methods have focused primarily on *de novo* molecular design, and rarely start the generation process from known binders. In this paper, we present HELiX: a template-based GFlowNet framework for localized exploration of chemical space. HELiX first learns to deconstruct a given hit by reversing selected reaction steps, and then performs forward synthesis in a manner that preserves synthetic tractability. Our approach demonstrates strong performance in efficiently identifying diverse, high-scoring analogs of known binders, and addresses the challenge of sample efficiency in GFlowNets by incorporating a Bayesian optimization loop which effectively balances exploration and exploitation. We also show that local exploration is inherently robust to noisy oracle evaluations, a common problem in drug development when using *in silico* predictors of binding affinity.

1 INTRODUCTION

In recent years, generative models for molecular design have attracted substantial attention in the machine learning community Tong et al. (2021); van den Broek et al. (2025), providing a powerful alternative to traditional supervised approaches for exploring vast chemical spaces. By learning to sample molecules from a target distribution, these models can directly propose candidates with desired properties, effectively acting as data-driven molecular search engines.

While early molecular generative models imposed few constraints on the structures they produced, the field has recently shifted toward incorporating molecular synthesizability to enable the direct translation of *in silico* designs to wet-lab experiments Koziarski et al. (2024); Cretu et al. (2025); Gaiński et al. (2025a), a consideration that has become increasingly important with the widespread adoption of iterative design–make–test–analyze (DMTA) pipelines in drug discovery Tom et al. (2024).

Despite their growing impact, most existing generative models—both unconstrained and synthesizability-aware—primarily focus on *de novo* molecular design, generating compounds from scratch without reference to an existing molecule. While *de novo* design is an important component of molecular discovery, later stages such as hit expansion and molecular optimization are equally critical, pose distinct challenges, and have received comparatively little attention in the generative modeling literature. In particular, these tasks require localized exploration around a given starting point and place strong demands on sample efficiency.

To address this gap, we introduce HELiX, a framework for **h**it **e**xpansion via **l**ocalized e**x**ploration of synthesizable chemical space. HELiX is, to our knowledge, the first generative model that enables localized exploration around a provided hit while explicitly incorporating synthesizability constraints. Our contributions are threefold:

- We propose the HELiX framework, an extension of template-based generative flow networks (GFlowNets) that reformulates the underlying Markov decision process to support reversible reaction steps, enabling controlled local exploration.

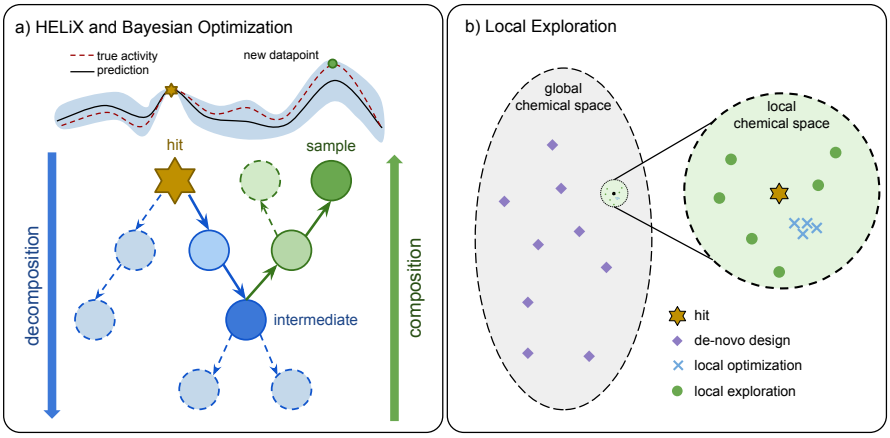

Figure 1: Figure (a) shows the generation process of HELiX. First, a given hit molecule is decomposed to an intermediate state by reversing the reactions. Then, the intermediate molecule is recomposed to a sample. The resulting samples can be evaluated using an expensive oracle (e.g., experimental assay) within a Bayesian optimization loop, and the observations are used to update the surrogate model guiding subsequent sampling. Figure (b) depicts local hit exploration. In contrast to *de novo* design, which performs global exploration of chemical space, HELiX concentrates sampling in the neighborhood of a given hit. This enables systematic exploration of the local chemical space around the hit.

- We improve sample efficiency of the framework by integrating Bayesian optimization, and demonstrate that the resulting method can achieve competitive performance in low oracle call regimes, a crucial property in molecular optimization.
- We present an extensive empirical evaluation demonstrating that HELiX achieves state-of-the-art performance in localized chemical space exploration, exhibiting high diversity and robustness to noisy objectives, while guaranteeing synthesizability out of the box.

2 METHODS

In this paper, we propose HELiX, the first framework for local exploration of chemical space using template-based GFlowNets. Unlike standard template-based generative models that construct molecules from scratch, HELiX initializes generation from a hit molecule, decomposes it to an intermediate molecule and then recomposes it into a new sample that partially shares a synthesis path with the hit. This process enforces a structural similarity to the hit, making the approach well-suited for local chemical space exploration. Similarly to other template-based GFlowNets, HELiX operates on a predefined set of reaction templates and building blocks.

2.1 FORWARD AND BACKWARD POLICIES

The generation process in HELiX starts from an initial hit molecule, which is sequentially decomposed to an intermediate state. From this intermediate, a composition phase reconstructs a new molecule. States are represented as a tuple (m, r, F), where m is the current molecule, r is a reaction template, and F is a list of fragments. The reaction template r may be unset (denoted by $\emptyset$), and the fragment list F may be empty or partially filled. Moreover, we track and limit the number of decomposition and composition steps; this information is omitted from the state representation for readability.

Decomposition Forward Action Space. Decomposition begins from the hit molecule m in the initial state $s_0 = (m, \emptyset, \emptyset)$. All single-step reactions that could have produced m are enumerated, and the policy selects one of them, leading to a state (m', r, F) such that applying reaction r to fragments F and molecule m' yields m. The fragments in F are then removed sequentially until

reaching a state of the form $(m', r, \emptyset)$ or a terminating action. At this point, the reaction template r can be removed, resulting in the state $(m', \emptyset, \emptyset)$, from which further decomposition actions may be taken. The decomposition process may terminate at any time, yielding a partially specified state of the form $(m, *, *)$.

Composition Forward Action Space. Starting from an intermediate state $(m, *, *)$, the composition process proceeds by sequentially selecting a reaction template r and its required fragments F, and applying the resulting reaction (m, r, F) to obtain a new state $(m', \emptyset, \emptyset)$. From this complete state, the agent may either continue composing by filling another template and fragments, or terminate the process, producing a final molecule.

Backward Action Spaces. Traversing the MDP backward introduces two main challenges. First, when reversing the composition phase, the trajectory must end at an intermediate molecule from which the initial hit can be recovered by reversing the decomposition phase. Second, when reversing decomposition, only molecules that lie on a valid synthesis path of the hit are admissible. Explicitly enforcing the validity of all backward actions along a trajectory is prohibitively expensive - more so than the retrosynthesis problem considered in Koziarski et al. (2024). Instead, following Cretu et al. (2025), we encourage backward policy P_B of HELiX to assign high probability only to actions that could eventually lead to a hit molecule by optimizing it direclty with REINFORCE algorithm Williams (1992):

$$J_B = E_{\tau \sim P_B}[R_B(\tau)] - \alpha H[P_B] \tag{1}$$

where $H[P_B]$ is the entropy term and the reward $R_B(\tau)$ is set to 1 if trajectory τ ends in the hit, and -1 otherwise.

3 Experiments

3.1 Local Hit Expansion

We compare HELiX against state-of-the-art baselines for local hit expansion, including SynFormer-ED Gao et al. (2024), REINVENT4 Loeffler et al. (2024) and TRACER Nakamura et al. (2025). To evaluate performance specifically in terms of local exploration, we use metrics computed over molecules that lie within a fixed Tanimoto similarity threshold to a given hit molecule, thereby restricting evaluation to the immediate chemical neighborhood of the hit. Initial hit molecules are generated by a SCENT model Gaiński et al. (2025a) (see Section D.1 for details).

We adapt local variants of standard metrics used in the literature: the number of high-reward molecules with pairwise Tanimoto similarity below 0.7 (number of modes), the number of high-reward Murcko scaffolds, the mean reward of the top-k molecules, and the top-1 molecule observed. Results for models trained with 10k oracle calls are shown in Table 1. We find that HELiX consistently outperforms REINVENT4 and TRACER across all metrics. Notably, SynFormer-ED approaches the performance of HELiX on most metrics evaluated at 10k oracle calls.

For comparison, we also include SCENT Gaiński et al. (2025a), a template-based GFlowNet designed for *de novo* drug design. For fairness, SCENT is evaluated within the same BO loop as HELiX. Its poor performance on local exploration metrics highlights the necessity of effective local exploration.

3.2 Sample Efficiency

While we previously compared the performance of each method at 10k oracle calls, we are similarly interested in the performance at earlier stages of optimization. To assess the sample efficiency of HELiX, we plot the number of high-reward local modes as a function of oracle calls in Figure 2, focusing on the first 5k observed molecules. We find that HELiX consistently outperforms all baselines in this regime. Notably, while SynFormer-ED achieves competitve performance at 10k oracle evaluations (Table 1), it exhibits markedly poorer sample efficiency at fewer oracle calls, failing to discover any high-reward modes until approximately 4,000 evaluations.

We further compare methods using the Area Under the Curve (AUC) top-10 metric evaluated at 10k oracle calls Gao et al. (2022a). AUC top-10 measures how quickly high-quality solutions accumulate

Table 1: Comparison to baselines on sEH proxy. Metrics are reported for molecules observed with > 0.4 ECFP4 Tanimoto similarity to the initial hit. Values are reported as the mean over 10 trials $\pm$ 1 standard deviation corresponding to different initial hits.

Method	Modes		Scaffolds		Top-10 Mean	Top-1
	(> 7.2)	(> 8)	(> 7.2)	(> 8)		
SCENT	174 $\pm$ 160	39 $\pm$ 40	613 $\pm$ 610	150 $\pm$ 172	7.40 $\pm$ 2.4	7.45 $\pm$ 2.5
SynFormer-ED	549 $\pm$ 126	49 $\pm$ 42	3417 $\pm$ 687	275 $\pm$ 188	**8.24** $\pm$ .08	8.30 $\pm$ .09
REINVENT4	203 $\pm$ 80	5 $\pm$ 4	418 $\pm$ 191	8 $\pm$ 8	7.99 $\pm$ .09	8.09 $\pm$ .08
TRACER	92 $\pm$ 64	3 $\pm$ 3	243 $\pm$ 170	4 $\pm$ 5	7.84 $\pm$ .27	8.04 $\pm$ .15
HELiX	**933** $\pm$ 227	**68** $\pm$ 29	**4559** $\pm$ 1221	**287** $\pm$ 129	8.22 $\pm$ .05	**8.31** $\pm$ .07

among the top-ranked candidates during an optimization procedure, capturing both solution quality and the speed of discovery under a fixed evaluation budget. Figure 2 demonstrates HELiX achieves the highest AUC top-10 among all baselines, further indicating strong sample efficiency.

Overall, these results show that HELiX is sample-efficient in terms of both exploration—identifying a larger number of high-reward local modes, and exploitation—recovering a stronger set of top-10 molecules compared to other methods.

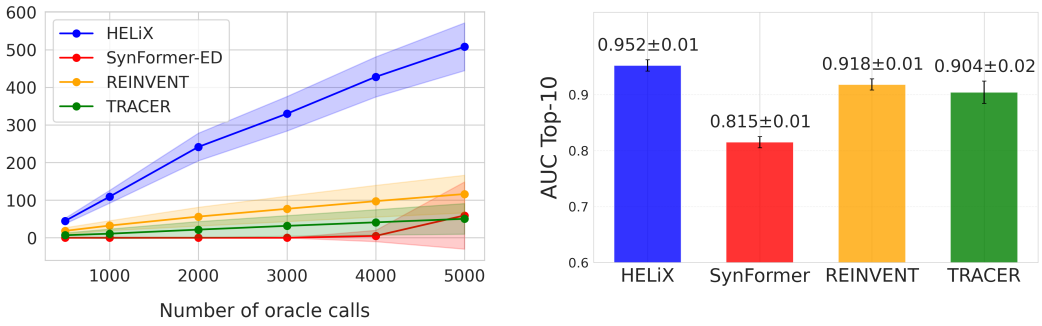

Figure 2: Comparison of sample efficiency in HELiX against other methods. (Left) The number of local modes found by each method in 5k oracle calls. (Right) The AUC top-10 of each method, computed on local molecules at 10k oracle calls.

3.3 Robustness to Noise

We further evaluate HELiX when oracle evaluations are corrupted by noise, simulating the use of imperfect binding affinity estimates such as molecular docking. We hypothesize that searching near known active compounds provides inherent robustness to noise. We test this by comparing HELiX to SCENT under increasing levels of additive Gaussian noise. Rather than constraining the metrics to a local neighborhood around the initial hit, we evaluate the *global* metrics in order to assess the overall quality of discovered molecules by each model, noting that SCENT typically performs better in this scenario due to its unconstrained search space. Metrics are then evaluated using the ground-truth (noise-free) proxy to assess how well each method recovers high-reward compounds despite noisy training signals.

Figure 3 shows that HELiX maintains nearly constant performance as the noise increases, while SCENT deteriorates rapidly. Since sEH proxy scores typically range from 0 to 8.5, Gaussian noise with $\sigma = 2.0$ represents a relatively high noise-to-signal ratio. This demonstrates that anchoring the search space in the local neighborhood of a known active binder provides robust optimization even under highly unreliable proxies.

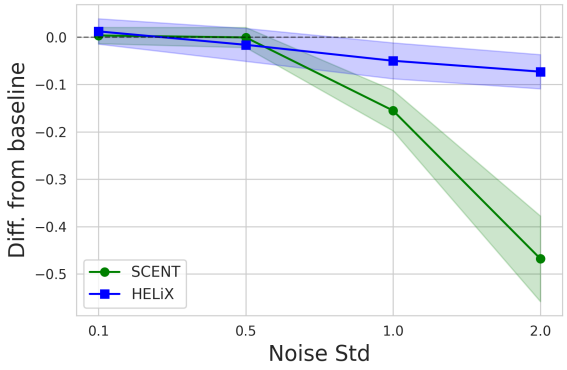

Figure 3: Performance degradation of HELiXand SCENT under increasing levels of additive Gaussian noise. Metric is reported as difference in global top-10 mean between noisy experiment and baseline (noise-free). Molecules are selected using the noisy oracle, while evaluation is performed using the ground-truth (noise-free) proxy.

3.4 Bayesian Optimization

We attribute a substantial amount of HELiX's performance, in particular its sample efficiency, to the use of BO in its training procedure. While BO has been combined with GFlowNets previously, prior results have shown limited improvements Gao et al. (2022a), suggesting that benefits depend on the underlying generative model and implementation.

Since SCENT is also a template-based GFlowNet, we incorporated BO into its training procedure to ensure fair comparison with HELiX. To understand BO's contributions to both methods, we perform an ablation comparing performance with and without BO. As shown in Table 2, BO consistently and substantially improves performance for both models in almost all settings, with an exception of SCENT in identifying local modes. We attribute this to variance in the results due to SCENT only identifying a small number of local modes. Overall, these results demonstrate that BO is a key contributor to the performance of HELiX.

Table 2: Comparison to SCENT. AUC Top-10 is computed *globally*.

Method	**Modes**		**AUC Top-10**
	Local	Global	
SCENT	296	1539	.906
SCENT + BO	174	**1962**	.942
HELiX	540	1255	.913
HELiX + BO	**933**	1486	**.952**

4 Conclusion

We propose HELiX, a synthesis-constrained GFlowNet framework for local hit expansion. We demonstrate the effectiveness of HELiXto identify many diverse, high-reward molecules in localized regions of chemical space around a given hit. Moreover, we show that by effectively incorporating BO into the training procedure, HELiX is able to be be highly sample efficient compared to baseline methods, a notable improvement over previous GFlowNet architectures. We also show that, by constraining the space of HELiX to high-reward regions, it is able to maintain robustness to noisy reward functions. Additionally, we can couple our model with a projection method that allows HELiX to expand hits which are incompatible with its chemical library (Appendix C.3).

References

Mila AI4Science, Alex Hernandez-Garcia, Alexandre Duval, Alexandra Volokhova, Yoshua Bengio, Divya Sharma, Pierre Luc Carrier, Yasmine Benabed, Michał Koziarski, and Victor Schmidt. Crystal-GFN: sampling crystals with desirable properties and constraints. *arXiv preprint arXiv:2310.04925*, 2023.

Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, 2021.

Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J. Hu, Mo Tiwari, and Emmanuel Bengio. GFlowNet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.

John Bradshaw, Brooks Paige, Matt J Kusner, Marwin Segler, and José Miguel Hernández-Lobato. A model to search for synthesizable molecules. *Advances in Neural Information Processing Systems*, 32, 2019.

John Bradshaw, Brooks Paige, Matt J Kusner, Marwin Segler, and José Miguel Hernández-Lobato. Barking up the right tree: an approach to search over molecule synthesis DAGs. *Advances in Neural Information Processing systems*, 33:6852–6866, 2020.

Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling*, 59(3): 1096–1108, 2019.

Yu-Chian Chen. Beware of docking! *Trends in pharmacological sciences*, 36(2):78–95, 2015.

Miruna Cretu, Charles Harris, Ilia Igashov, Arne Schneuing, Marwin Segler, Bruno Correia, Julien Roy, Emmanuel Bengio, and Pietro Lio. Synflownet: Design of diverse and novel molecules with synthesis constraints. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=uvHmnahyp1.

Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1:1–11, 2009.

Piotr Gaiński, Oussama Boussif, Andrei Rekesh, Dmytro Shevchuk, Ali Parviz, Mike Tyers, Robert A. Batey, and Michał Koziarski. Scalable and cost-efficient de novo template-based molecular generation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025a. URL https://openreview.net/forum?id=zssWxiiJZ1.

Piotr Gaiński, Michał Koziarski, Krzysztof Maziarz, Marwin Segler, Jacek Tabor, and Marek Śmieja. Diverse and feasible retrosynthesis using GFlowNets. *Information Sciences*, 714:122194, 2025b.

Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor W. Coley. Sample efficiency matters: A benchmark for practical molecular optimization, 2022a.

Wenhao Gao, Rocío Mercado, and Connor W Coley. Amortized tree generation for bottom-up synthesis planning and synthesizable molecular design. In *The Tenth International Conference on Learning Representations*, 2022b.

Wenhao Gao, Shitong Luo, and Connor W. Coley. Generative ai for navigating synthesizable chemical space. *Proceedings of the National Academy of Sciences*, 121:e2415665122, 2024.

Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist, and Esben Bjerrum. AiZynthFinder: a fast, robust and flexible open-source software for retrosynthetic planning. *Journal of cheminformatics*, 12(1):70, 2020.

Pouya M. Ghari, Alex M. Tseng, Gökcen Eraslan, Romain Lopez, Tommaso Biancalani, Gabriele Scalia, and Ehsan Hajiramezanali. Gflownet assisted biological sequence editing. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024.

Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.

Sai Krishna Gottipati, Boris Sattarov, Sufeng Niu, Yashaswi Pathak, Haoran Wei, Shengchao Liu, Simon Blackburn, Karam Thomas, Connor Coley, Jian Tang, et al. Learning to navigate the synthetically accessible chemical space using reinforcement learning. In *International conference on machine learning*, pp. 3668–3679. PMLR, 2020.

Jeff Guo and Philippe Schwaller. Directly optimizing for synthesizability in generative molecular design using retrosynthesis models. *Chemical Science*, 16(16):6943–6956, 2025.

Julien Horwood and Emmanuel Noutahi. Molecular design in synthetically accessible chemical space via deep reinforcement learning. *ACS omega*, 5(51):32984–32994, 2020.

Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics data commons: Machine learning datasets and tasks for drug discovery and development. *arXiv preprint arXiv:2102.09548*, 2021.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Hierarchical generation of molecular graphs using structural motifs. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.

Zygimantas Jocys, Henriette MG Willems, and Katayoun Farrahi. SynthFormer: Equivariant pharmacophore-based generation of molecules for ligand-based drug design. *arXiv preprint arXiv:2410.02718*, 2024.

Hyeonah Kim, Minsu Kim, Sanghyeok Choi, and Jinkyoo Park. Genetic-guided gflownets for sample efficient molecular optimization. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024.

Maksym Korablyov, Cheng-Hao Liu, Moksh Jain, Almer M van der Sloot, Eric Jolicoeur, Edward Ruediger, Andrei Cristian Nica, Emmanuel Bengio, Kostiantyn Lapchevskyi, Daniel St-Cyr, et al. Generative active learning for the search of small-molecule protein binders. *arXiv preprint arXiv:2405.01616*, 2024.

Michał Koziarski, Andrei Rekesh, Dmytro Shevchuk, Almer van der Sloot, Piotr Gaiński, Yoshua Bengio, Cheng-Hao Liu, Mike Tyers, and Robert A. Batey. Rgfn: Synthesizable molecular generation using gflownets. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024.

Cheng-Hao Liu, Maksym Korablyov, Stanisław Jastrzebski, Paweł Włodarczyk-Pruszynski, Yoshua Bengio, and Marwin Segler. RetroGNN: fast estimation of synthesizability for virtual screening and de novo design by learning from slow retrosynthesis software. *Journal of Chemical Information and Modeling*, 62(10):2293–2300, 2022.

Alston Lo, Connor W Coley, and Wojciech Matusik. A genetic algorithm for navigating synthesizable molecular spaces. *arXiv preprint arXiv:2509.20719*, 2025.

Hannes H. Loeffler, Jiazhen He, Alessandro Tibo, Jon Paul Janet, Alexey Voronov, Lewis H. Mervin, and Ola Engkvist. REINVENT 4: Modern ai–driven generative molecule design. *Journal of Cheminformatics*, 16(1):20, 2024.

Shitong Luo, Wenhao Gao, Zuofan Wu, Jian Peng, Connor W Coley, and Jianzhu Ma. Projecting molecules into synthesizable chemical spaces. In *International Conference on Machine Learning*, pp. 33289–33304. PMLR, 2024.

Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory balance: Improved credit assignment in gflownets. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, pp. 5955–5967, 2022.

Shogo Nakamura, Nobuaki Yasuo, and Masakazu Sekijima. Molecular optimization using a conditional transformer for reaction-aware compound exploration with reinforcement learning. *Communications Chemistry*, 8:40, 2025.

Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):48, 2017.

Andrei Rekesh, Miruna Cretu, Dmytro Shevchuk, Vignesh Ram Somnath, Pietro Liò, Robert A Batey, Mike Tyers, Michał Koziarski, and Cheng-Hao Liu. SynCoGen: Synthesizable 3D molecule generation via joint reaction and coordinate modeling. *arXiv preprint arXiv:2507.11818*, 2025.

Seonghwan Seo, Jaechang Lim, and Woo Youn Kim. Molecular generative model via retrosynthetically prepared chemical building block assembly. *Advanced Science*, 10(8):2206674, 2023.

Seonghwan Seo, Minsu Kim, Tony Shen, Martin Ester, Jinkyoo Park, Sungsoo Ahn, and Woo Youn Kim. Generative flows on synthetic pathway for drug design. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=pB1XSj2y4X.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175, 2015.

Tony Shen, Seonghwan Seo, Grayson Lee, Mohit Pandey, Jason R. Smith, Artem Cherkasov, Woo Youn Kim, and Martin Ester. TacoGFN: Target conditioned GFlowNet for structure-based drug design. *Transactions on Machine Learning Research*, 2024.

Niclas Ståhl, Goran Falkman, Alexander Karlsson, Gunnar Mathiason, and Jonas Bostrom. Deep reinforcement learning for multiparameter optimization in de novo drug design. *Journal of chemical information and modeling*, 59(7):3166–3176, 2019.

Gary Tom, Stefan P Schmid, Sterling G Baird, Yang Cao, Kourosh Darvish, Han Hao, Stanley Lo, Sergio Pablo-García, Ella M Rajaonson, Marta Skreta, et al. Self-driving laboratories for chemistry and materials science. *Chemical Reviews*, 124(16):9633–9732, 2024.

Xiaochu Tong, Xiaohong Liu, Xiaoqin Tan, Xutong Li, Jiaxin Jiang, Zhaoping Xiong, Tingyang Xu, Hualiang Jiang, Nan Qiao, and Mingyue Zheng. Generative models for de novo drug design. *Journal of Medicinal Chemistry*, 64(19):14011–14027, 2021.

Austin Tripp, Erik Daxberger, and José Miguel Hernández-Lobato. Sample-efficient optimization in the latent space of deep generative models via weighted retraining. *Advances in Neural Information Processing Systems*, 33:11259–11272, 2020.

Remco L van den Broek, Shivam Patel, Gerard JP van Westen, Willem Jespers, and Woody Sherman. In search of beautiful molecules: a perspective on generative modeling for drug design. *Journal of chemical information and modeling*, 65(18):9383–9397, 2025.

Alexandra Volokhova, Michał Koziarski, Alex Hernández-García, Cheng-Hao Liu, Santiago Miret, Pablo Lemos, Luca Thiede, Zichao Yan, Alán Aspuru-Guzik, and Yoshua Bengio. Towards equilibrium molecular conformation generation with GFlowNets. *Digital Discovery*, 3(5):1038–1047, 2024.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31, 2018.

A Background and Related Works

A.1 GFlowNets

Generative Flow Networks Bengio et al. (2021; 2023) are probabilistic models that learn a stochastic policy for constructing objects x through sequences of actions, such that the probability of sampling an object is proportional to a predefined reward $R(x)$. Generation is formulated as a trajectory $\tau \in \mathcal{T}$ that starts from an initial state s_0 and transition between states via actions a.

GFlowNets define a forward policy $P_F(a \mid s)$ over outgoing transitions and a backward policy $P_B(a \mid s)$ over reverse transitions. In this work, we train them using the Trajectory Balance objective Malkin et al. (2022):

$$F(s_0) \prod_{t=0}^{T-1} P_F(a_t \mid s_t) \;=\; R(s_f) \prod_{t=0}^{T-1} P_B(a_t \mid s_{t+1}), \tag{2}$$

where $F(s_0)$ is the flow going through the source state s_0.

The GFlowNet framework has been applied to molecular design, including synthesis-aware formulations such as Reaction-GFlowNets Koziarski et al. (2024), hybrid approaches that integrate evolutionary search to improve sample efficiency Kim et al. (2024), and structure-conditioned variants for protein–ligand design Shen et al. (2024). Beyond small molecule generation, GFlowNets have been applied to a plethora of scientific discovery problems, including biological sequence and peptide design Ghari et al. (2024), retrosynthesis Gaiński et al. (2025b), molecular conformation sampling Volokhova et al. (2024) and crystal generation AI4Science et al. (2023).

A.2 Synthesizability-aware Generative Models

Ensuring the synthesizability of generated molecules has been approached in several ways. One common strategy relies on synthesizability scoring functions, either as optimization objectives or post-hoc filters Korablyov et al. (2024). These include heuristic metrics Ertl & Schuffenhauer (2009); Genheden et al. (2020) and machine-learning–based predictors Liu et al. (2022); Guo & Schwaller (2025). However, such scores can be unreliable, particularly out of distribution, and often fail to capture practical considerations such as synthesis cost or reaction availability.

As an alternative, many methods explicitly constrain generation to chemically feasible spaces by constructing molecules from predefined building blocks and reaction templates, referred to as template-based approaches. Early work explored autoencoders Bradshaw et al. (2019; 2020), genetic algorithms Gao et al. (2022b), and reinforcement learning Gottipati et al. (2020); Horwood & Noutahi (2020). More recent approaches leverage autoregressive and transformer-based models Seo et al. (2023); Gao et al. (2024); Jocys et al. (2024) and flow matching Rekesh et al. (2025) to incorporate reaction-aware priors.

Template-based GFlowNets have recently emerged as a powerful paradigm for molecular generation in drug design, combining the diversity guarantees of flow-based objectives with explicit synthetic control Koziarski et al. (2024); Cretu et al. (2025); Seo et al. (2025); Gaiński et al. (2025a). By generating molecules through sequences of reaction templates applied to purchasable building blocks, these methods operate directly in synthesis space, producing candidates with valid, fully specified reaction routes.

In contrast to prior template-based GFlowNets, which primarily target *de novo* exploration of synthesis space, HELiX is the first framework explicitly designed for local hit exploration. We reformulate molecular generation as a coupled decompose–recompose process: a given hit molecule is partially decomposed into an intermediate molecule, which is then recomposed into a new sample. This construction enforces both structural proximity and shared synthetic pathways by design.

A.3 Molecular Optimization and Hit Expansion

Most computational drug design methods frame molecule generation as a global optimization problem, aiming to directly maximize predicted properties such as binding affinity, QED, or composite scores using deep generative models, reinforcement learning, or evolutionary search Gómez-Bombarelli et al. (2018); Olivecrona et al. (2017); You et al. (2018); Ståhl et al. (2019). While effective for *de novo* hit discovery, such approaches often collapse toward a small number of high-scoring candidates and are sensitive to the choice of proxy objective. This is particularly problematic in practice, as commonly used *in silico* affinity estimators, most notably docking and QSAR models, are weakly correlated with experimental binding, and susceptible to exploitation by optimization algorithms Chen (2015).

In contrast, hit expansion (or hit exploration) emphasizes the generation of structurally proximal analogs around known binders, aligning more closely with medicinal chemistry practice and SAR

development. Hit exploration seeks to map the local chemical neighborhood of an active scaffold, providing multiple viable candidates that hedge against proxy noise and experimental uncertainty.

This paradigm is realized through matched molecular pair analysis, analog enumeration, and conditional or translation-based generative models that improve properties while preserving core structure Brown et al. (2019). HierVAE performs supervised molecule-to-molecule translation with a hierarchical graph VAE, optimizing properties while preserving structural proximity Jin et al. (2020). REINVENT4 introduces a Transformer-based Mol2Mol generator that exhaustively samples analogs within a similarity radius, benefiting form large-scale pretraining Loeffler et al. (2024). TRACER predicts feasible reaction templates with a graph convolutional network and generates products using a conditional Transformer, exploring multi-step, synthetically valid modifications Nakamura et al. (2025). SynFormer models molecules as outcomes of reaction sequences over purchasable building blocks operating entirely in synthesis space Gao et al. (2024). SynGA similarly evolves synthesis routes using reaction-aware genetic operators, ensuring synthesizability and achieving efficient hit expansion by tightly coupling search to chemical transformations Lo et al. (2025).

A.4 Sample Efficiency of GFlowNets

Sample efficiency is a central challenge for GFlowNets in molecular design, as training typically requires many oracle evaluations Gao et al. (2022a). Prior work has improved efficiency through better objectives such as Trajectory Balance Malkin et al. (2022) or hybrid evolutionary schemes Kim et al. (2024). Bayesian optimization (BO) offers a complementary approach by adaptively selecting informative queries under limited budgets and has been explored in conjunction with generative models, including GFlowNets, to balance exploration and exploitation Gómez-Bombarelli et al. (2018); Shahriari et al. (2015); Tripp et al. (2020). Building on these ideas, HELiX couples BO with localized exploration around known hits in GFlowNet framework, enabling more efficient use of oracle evaluations in low-budget regimes, while proposing diverse analogs.

B HELiX

B.1 Parametrization of HELiX

The backbone of HELiX is a graph transformer $f(s) \in \mathbb{R}^D$ operating on states $s = (m, r, F)$, where m is a molecule, r is a reaction template, and F is a set of fragments. The model outputs a D-dimensional state embedding. The molecule m is represented as a molecular graph. Reaction templates r are encoded as one-hot vectors over the template library. Each fragment in F is represented by an ECFP fingerprint, linearly embedded and summed to obtain a fixed-size fragment-set representation.

Let $A_F(s)$ denote the set of forward actions available from a state s, and $A_B(s)$ the corresponding set of backward actions. The forward policy comprises three action types: reaction decomposition, fragment removal, and template removal. Symmetrically, the backward policy consists of reaction application, fragment addition, and template addition.

Both forward and backward policies are defined by a unified action-scoring formulation. For a given state s, actions are scored and normalized using a softmax operator σ^k over a length-k logit vector $\mathbf{c} \in \mathbb{R}^k$:

$$P_X(a \mid s) = \sigma^{|A_X(s)|}(\mathbf{c})_i,$$
$$c_i = \phi\Big(W_{A^*_X(s)} f(s)\Big)^\top e_{A^*_X(s)}(a_i), \tag{3}$$

where $X \in \{F, B\}$ denotes the forward or backward policy, ϕ is the GeLU activation function, and $W_{A^*_X(s)}$ is a linear projection specific to the type of action space $A_X(s)$. Each action $a_i \in A_X(s)$ is embedded using an action-type-dependent embedding function $e_{A^*_X(s)}(\cdot)$. For reaction decomposition actions, where a_i corresponds to a candidate predecessor tuple (m', r, F), the action embedding is obtained by applying the backbone transformer, i.e., $e(a_i) = f(a_i)$. For reaction template addition, actions are embedded as one-hot vectors over the reaction template set. For fragment addition, the action embedding is formed by concatenating a one-hot fragment identifier with a linearly transformed ECFP fingerprint.

B.2 TRAINING DETAILS

The training procedure of HELiX is outlined in Algorithm 1.

Algorithm 1 Bayesian Optimization with HELiX

Input: Hit molecule h, oracle $f : \mathcal{X} \to \mathbb{R}$, BO iterations N, candidates per iteration M, top-K selection K, acquisition function α, reward function R, GP model $\mathcal{F}$
Warm-start with 50 random samples from uniform forward policy
Evaluate on oracle and initialize dataset: $\mathcal{D} \leftarrow \{(x_j, f(x_j))\}$
for $i = 1$ to N **do**
  Fit GP surrogate on $\mathcal{D}$
  Train HELiX policy π_θ on acquisition function α
  Sample candidates: $\mathcal{C} \leftarrow \{x_j \sim \pi_\theta\}_{j=1}^{M}$
  Select top-K by acquisition: $\mathcal{C}_{\text{top}} \leftarrow \text{TopK}(\mathcal{C}, \alpha, K)$
  Evaluate on oracle and update dataset: $\mathcal{D} \leftarrow \{(x, f(x)) \mid x \in \mathcal{C}_{\text{top}}\}$
end for
return $\mathcal{D}$

GP Surrogate. We use a Gaussian Process with the Tanimoto (MinMax) kernel for molecular fingerprints:

$$k_{\text{MinMax}}(\mathbf{x}, \mathbf{x}') = \frac{\sum_i \min(x_i, x_i')}{\sum_i \max(x_i, x_i')}$$

where $\mathbf{x}$ and $\mathbf{x}'$ are 4096-dimensional ECFP4 fingerprints. The GP is fitted by maximizing the marginal log-likelihood. For efficiency, when $|\mathcal{D}| > 5000$, we subsample by taking the top 2500 molecules by score and 2500 random molecules.

Acquisition Function. We use the UCB acquisition function, defined as

$$\alpha(x) = m(x) + \beta\,\sigma(x),$$

where β is a tunable parameter controlling the greediness of the BO policy, and $m(x)$, $\sigma(x)$ are the posterior mean and variance of the GP model, respectively. During hyperparameter tuning, we found that a constant β equal to 1 was optimal.

Due to the large state space of GFlowNets, it is common to use a tunable exponential reward function in order to compensate for . Thus, we train the policy model with a reward function which takes the form

$$\begin{aligned} R(x) &= \exp(\gamma\,\alpha(x)) \\ &= \exp\left(\gamma\ m(x)\right) \exp\left(\gamma\,\beta\,\sigma(x)\right). \end{aligned}$$

We decoupled the parameters scaling the mean and variance by setting $\beta = \frac{1}{\gamma}$, allowing for independent control of the exploration-exploitation tradeoff in the model.

C EXTENDED RESULTS

C.1 ADDITIONAL ORACLES

In this section, we present the results of HELiX against the baseline methods on two additional oracle functions from the Therapeutic Data Commons (TDC; Huang et al. (2021)): GSK3β (Table 3 and JNK3 (Table 4). Similar trends can be observed as in Section 3.1.

C.2 PROPERTIES OF GENERATED MOLECULES

We analyze the properties of the top molecules produced by each model, including molecular weight, QED, SAScore, and AiZynthFinder success rate, which is defined as the proportion of molecules for which AiZynthFinder was able to successfully find a synthetic route via retrosynthesis. Figure 4 displays the distributions of each metric across the top 100 molecules generated from each initial hit, resulting in 1000 molecules in total.

Table 3: Comparison to baselines on GSK3β proxy. Values are reported as the mean over 10 trials $\pm$ 1 standard deviation corresponding to 10 different initial hits.

Method	Modes		Scaffolds		Top-10 Mean	Top-1
	(> 7.2)	(> 8)	(> 7.2)	(> 8)		
SynFormer-ED	52 $\pm$ 60	7 $\pm$ 9	34 $\pm$ 29	2 $\pm$ 3	.764 $\pm$ .07	.787 $\pm$ .065
REINVENT4	255 $\pm$ 111	78 $\pm$ 26	150 $\pm$ 56	45 $\pm$ 28	**.889** $\pm$ .021	.923 $\pm$ .02
TRACER	93 $\pm$ 44	20 $\pm$ 13	80 $\pm$ 58	19 $\pm$ 17	.843 $\pm$ .028	**.956** $\pm$ .029
HELiX	**747** $\pm$ 409	**298** $\pm$ 223	**556** $\pm$ 135	**238** $\pm$ 131	.88 $\pm$ .031	.891 $\pm$ .031

Table 4: Comparison to baselines on JNK3 proxy. Values are reported as the mean over 10 trials $\pm$ 1 standard deviation corresponding to 10 different initial hits.

Method	Modes		Scaffolds		Top-10 Mean	Top-1
	(> 7.2)	(> 8)	(> 7.2)	(> 8)		
SynFormer-ED	194 $\pm$ 37	74 $\pm$ 16	616 $\pm$ 215	236 $\pm$ 94	.884 $\pm$ .012	.903 $\pm$ .021
REINVENT4	515 $\pm$ 128	235 $\pm$ 93	879 $\pm$ 281	402 $\pm$ 195	**.928** $\pm$ .016	**.949** $\pm$ .022
TRACER	71 $\pm$ 28	21 $\pm$ 8	142 $\pm$ 64	35 $\pm$ 18	.856 $\pm$ .02	.873 $\pm$ .023
HELiX	**1005** $\pm$ 302	**645** $\pm$ 278	**4505** $\pm$ 655	**3063** $\pm$ 739	.911 $\pm$ .017	.92 $\pm$ .019

C.3 Non-synthesizable Reference Molecules

It is possible that the initial hit compound is not compatible with the building blocks and reaction templates of HELiX, making it impossible to decompose and perform local exploration. To demonstrate that HELiX can be extended to handle such cases, we combine it with ChemProjector Luo et al. (2024), which projects arbitrary molecules onto synthetic routes within a given chemical library.

Starting from a known sEH inhibitor, we generated 50 projections with ChemProjector, clustered the results, and selected the top 5 by proxy score as initial hits for exploration with HELiX. Table 5 shows the results at 1k oracle calls. While the absolute performance metrics are lower than those in Table 1, this reflects the lower quality of the projected starting points rather than a limitation of HELiX itself. Notably, HELiX achieves substantial improvements over the projected hits ($\Delta = 1.42$ compared to $\Delta = 0.23$ for hits from SCENT), demonstrating that local exploration remains effective even when initialized from suboptimal projections. SMILES strings and 2D structures of the sEH inhibitor and 5 projected analogs are shown in Section D.3.

These results establish the viability of coupling HELiX with projection methods to handle non-synthesizable reference molecules. Future work could optimize this pipeline by improving the quality of projected compounds or by directly integrating projection into the HELiX training procedure.

D Experimental Details

D.1 Generating hit compounds

We first trained a SCENT model that explored 320,000 molecules. From this set, we selected the top 50% unique molecules based on their proxy value that also satisfied the following criteria: $QED \geq 0.5$, $SA \leq 3.0$ and MW in the range $[300, 500]$ (i.e. $300 \leq MW \leq 500$). To ensure a similar size of chemical space across all reference molecules, we only used those constructed with ≥ 3 reactions to allow for 2 decompositions for each reference molecule. The resulting set clustered using greedy clustering in order of proxy value, where centroids where determined to have < 0.7 ECFP6 Tanimoto similarity to existing centroids. We then took the top 10 such centroids according to proxy value, resulting in a diverse set of initial hit compounds. This was repeated for each proxy (sEH, GSK3β, and JNK3). SMILES strings and properties of the resulting compounds can be seen in Table 6, and the 2D representations are shown in Figure 5.

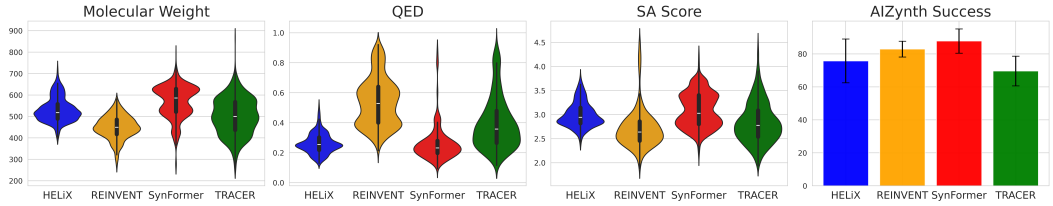

Figure 4: Distribution of molecular weight, QED, SA Score, and AiZynthFinder success rate for top-100 molecules from each method.

Table 5: Evaluation of HELiX combined with ChemProjector (HELiX + P). We compare to HELiX's performance on the baseline experiment. Metrics are evaluated at 1k oracle calls and are reported for molecules observed with > 0.4 ECFP4 Tanimoto similarity to the projected hit. Δ measures the improvement from the initial projected hit to the best molecule found.

Method	Modes > 7.2	Top-10 Mean	Top-1	Δ
HELiX	109	7.99	8.09	.227
HELiX+ P	27	7.73	7.79	1.42

D.2 Baseline methods

SynFormer-ED. We trained a SynFormer-ED model using approximately 220 million data points. This is less than the original model due to the more limited size of our chemical library. We employed a batch size of 512 instead of 1024, constrained by available computing power. To enable conditional generation, we slightly modified the RL fine-tuning script originally designed for SynFormer-D to re-incorporate the trained encoder. During this process, only the decoder parameters were updated. We used a batch size of 64, while all other settings remained at their default values, with both the prior and replay buffer enabled.

REINVENT4. Mol2Mol is the component of REINVENT4 dedicated to molecular optimization. We employed a prior trained to generate molecules with a Tanimoto similarity (ECFP4) of at least 0.5 to the input molecule. All other generation parameters were kept at their default settings.

TRACER. We utilized the pre-trained weights for the Transformer and Graph Convolutional Network (GCN) provided by the authors, which were trained on the USPTO dataset. This approach was adopted because constructing a comparable and fair dataset based on our specific library was not feasible. Consequently, we treated the model as a black box, similar to the REINVENT4 framework, where reaction templates and building blocks cannot be easily modified or swapped. The MTSC process was subsequently conducted using all default configurations.

D.3 Non-synthesizable Reference Molecule

Table 7 shows the SMILES strings and properties for the non-synthesizable reference molecule and projected analogs used in Section C.3. Figure 7 shows the corresponding 2D representations. The initial hit is a known sEH inhibitor that is not compatible with our building block library. We generated 50 projections using ChemProjector and selected the top 5 modes by proxy score as starting points for exploration with HELiX.

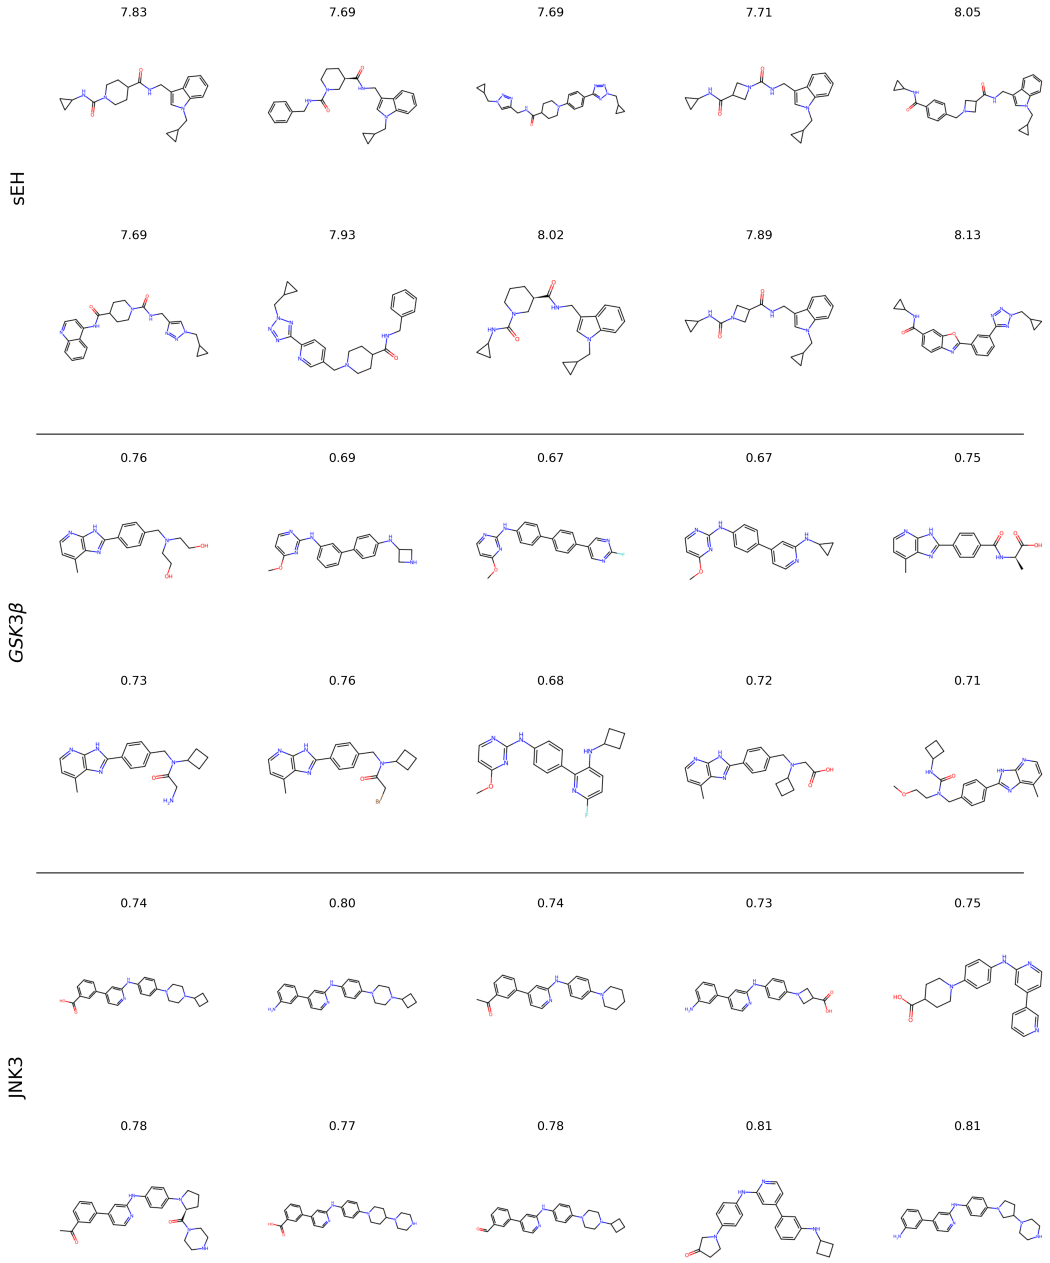

Figure 5: 2D representations of initial hits generated by SCENT. Corresponding proxy scores are shown.

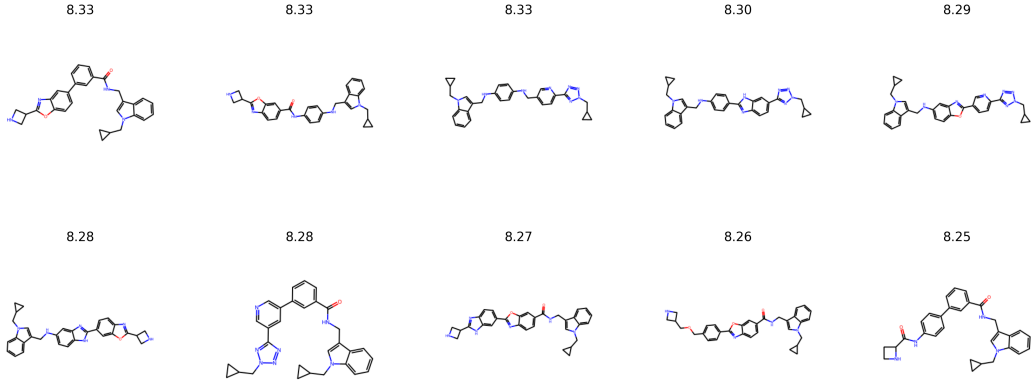

Figure 6: 10 modes generated by HELiX on sEH proxy. Corresponding sEH proxy scores are shown.

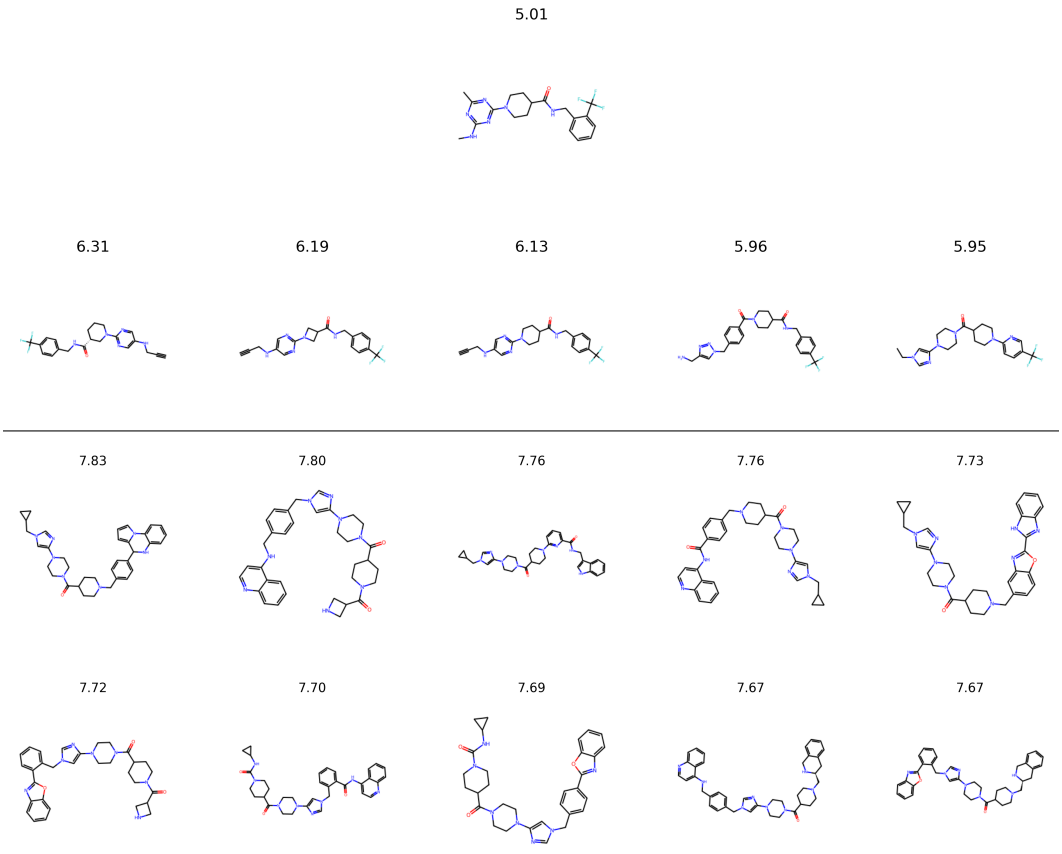

Figure 7: 2D representations of initial hit (first row), 5 projected analogs (second row) from ChemProjector, and top-10 modes generated from one run (rows 3-4). Corresponding sEH proxy scores are shown

Table 6: Properties of the selected reference molecules generated by SCENT for sEH, GSK3β, and JNK3, separated respectively into three sections.

Smiles	Proxy	QED	SA	MW
O=C(NCc1cn(CC2CC2)c2ccccc12)C1CCN(C(=O)NC2CC2)CC1	7.83	0.79	2.33	394
O=C(NCc1cn(CC2CC2)c2ccccc12)[C@@H]1CCCN(C(=O)NCc2ccccc2)C1	7.69	0.571	2.73	444
O=C(NCc1cn(CC2CC2)nn1)C1CCN(c2ccc(-c3nnn(CC4CC4)n3)cc2)CC1	7.69	0.521	2.7	461
O=C(NC1CC1)C1CN(C(=O)NCc2cn(CC3CC3)c3ccccc23)C1	7.71	0.825	2.41	366
O=C(NCc1cn(CC2CC2)nn1)C1CCN(c2ccc(-c3nnn(CC4CC4)n3)cc2)CC1	8.05	0.516	2.37	456
O=C(Nc1ccnc2ccccc12)C1CCN(C(=O)NCc2cn(CC3CC3)nn2)CC1	7.69	0.622	2.46	433
O=C(NCc1ccccc1)C1CCN(Cc2ccc(-c3nnn(CC4CC4)n3)nc2)CC1	7.93	0.59	2.4	431
O=C(NCc1cn(CC2CC2)c2ccccc12)[C@@H]1CCCN(C(=O)NC2CC2)C1	8.02	0.79	2.8	394
O=C(NCc1cn(CC2CC2)c2ccccc12)C1CN(C(=O)NC2CC2)C1	7.89	0.825	2.39	366
O=C(NC1CC1)c1ccc2nc(-c3cccc(-c4nnn(CC5CC5)n4)c3)oc2c1	8.13	0.532	2.49	400
Cc1ccnc2[nH]c(-c3ccc(CN(CCO)CCO)cc3)nc12	0.76	0.615	2.49	326
COc1ccnc(Nc2cccc(-c3ccc(NC4CNC4)cc3)c2)n1	0.69	0.636	2.27	347
COc1ccnc(Nc2ccc(-c3ccc(-c4cnc(F)nc4)cc3)cc2)n1	0.67	0.520	2.33	373
COc1ccnc(Nc2ccc(-c3ccnc(NC4CC4)c3)cc2)n1	0.67	0.713	2.24	333
Cc1ccnc2[nH]c(-c3ccc(C(=O)N[C@H](C)C(=O)O)cc3)nc12	0.75	0.681	2.71	324
Cc1ccnc2[nH]c(-c3ccc(CN(C(=O)CN)C4CCC4)cc3)nc12	0.73	0.741	2.54	349
Cc1ccnc2[nH]c(-c3ccc(CN(C(=O)CBr)C4CCC4)cc3)nc12	0.76	0.638	2.62	413
COc1ccnc(Nc2ccc(-c3nc(F)ccc3NC3CCC3)cc2)n1	0.68	0.631	2.47	365
Cc1ccnc2[nH]c(-c3ccc(CN(CC(=O)O)C4CCC4)cc3)nc12	0.72	0.712	2.49	350
COCCN(Cc1ccc(-c2nc3c(C)ccnc3[nH]2)cc1)C(=O)NC1CCC1	0.71	0.642	2.53	393
O=C(O)c1ccccc(-c2ccnc(Nc3ccc(N4CCN(C5CCC5)CC4)cc3)c2)c1	0.74	0.584	2.23	428
Nc1ccccc(-c2ccnc(Nc3ccc(N4CCN(C5CCC5)CC4)cc3)c2)c1	0.8	0.607	2.24	399
CC(=O)c1ccccc(-c2ccnc(Nc3ccc(N4CCCCC4)cc3)c2)c1	0.74	0.586	2.04	371
Nc1ccccc(-c2ccnc(Nc3ccc(N4CC(C(=O)O)C4)cc3)c2)c1	0.73	0.602	2.2	360
O=C(O)C1CCN(c2ccc(Nc3cc(-c4ccccn4)ccn3)cc2)CC1	0.75	0.697	2.16	374
CC(=O)c1ccccc(-c2ccnc(Nc3ccc(N4CCC[C@H]4C(=O)N4CCNCC4)cc3)c2)c1	0.78	0.53	2.95	469
O=C(O)c1ccccc(-c2ccnc(Nc3ccc(N4CCC(N5CCNCC5)CC4)cc3)c2)c1	0.77	0.515	2.43	457
O=Cc1ccccc(-c2ccnc(Nc3ccc(N4CCN(C5CCC5)CC4)cc3)c2)c1	0.78	0.579	2.32	412
O=C1CCN(c2ccc(Nc3cc(-c4ccccc(NC5CCC5)c4)ccn3)cc2)C1	0.81	0.595	2.39	398
Nc1ccccc(-c2ccnc(Nc3ccc(N4CCC(N5CCNCC5)C4)cc3)c2)c1	0.81	0.555	2.89	414

Table 7: Properties of the non-synthesizable initial hit and its top 5 projected analogs generated by ChemProjector. The projected analogs are synthesizable within our chemical library and serve as starting points for local exploration with HELiX.

Smiles	Proxy	QED	SA	MW
CNc1nc(C)nc(n1)N2CCC(CC2)C(=O)NCc3ccccc3C(F)(F)F	5.01	0.792	2.44	408
C#CCNc1cnc(N2CCC[C@@H](C(=O)NCc3ccc(C(F)(F)F)cc3)C2)nc1	6.31	0.707	3.01	417
C#CCNc1cnc(N2CC(C(=O)NCc3ccc(C(F)(F)F)cc3)C2)nc1	6.19	0.742	2.62	389
C#CCNc1cnc(N2CCC(C(=O)NCc3ccc(C(F)(F)F)cc3)CC2)nc1	6.13	0.707	2.54	417
NCc1cn(Cc2ccc(C(=O)N3CCC(C(=O)NCc4ccc(C(F)(F)F)cc4)CC3)cc2)nn1	5.96	0.519	2.41	501
CCn1cnc(N2CCN(C(=O)C3CCN(c4ccc(C(F)(F)F)cn4)CC3)CC2)c1	5.95	0.738	2.69	436