
GENERATIVE ARTIFICIAL INTELLIGENCE FOR NAVIGATING SYNTHESIZABLE CHEMICAL SPACE

Wenhao Gao*

Department of Chemical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
whgao@mit.edu

Shitong Luo*

Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
luost@mit.edu

Connor W. Coley

Department of Chemical Engineering
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
ccoley@mit.edu

*Equal Contribution

ABSTRACT

We introduce SynFormer, a generative modeling framework designed to efficiently explore and navigate synthesizable chemical space. Unlike traditional molecular generation approaches, we generate synthetic pathways for molecules to ensure that designs are synthetically tractable. By incorporating a scalable transformer architecture and a diffusion module for building block selection, SynFormer surpasses existing models in synthesizable molecular design. We demonstrate SynFormer’s effectiveness in two key applications: (1) local chemical space exploration, where the model generates synthesizable analogs of a reference molecule, and (2) global chemical space exploration, where the model aims to identify optimal molecules according to a black-box property prediction oracle. Additionally, we demonstrate the scalability of our approach via the improvement in performance as more computational resources become available. With our code and trained models openly available, we hope that SynFormer will find use across applications in drug discovery and materials science.

1 Introduction

The discovery of novel functional molecules is a central challenge in chemical science and engineering and is crucial for addressing key societal and technological challenges, including those related to healthcare [1, 2, 3], energy [4, 5], and sustainability [6, 7, 8]. However, the process of discovery is often risky, complex, time-consuming, and resource-intensive [9, 10, 11]. Recent advances in artificial intelligence (AI) [12], particularly in generative modeling [13, 14, 15], have opened up new avenues for generative molecular design [16, 17, 18, 19, 20, 21, 22, 23]. The expressivity of deep learning models, coupled with automatic differentiation [24] and increasingly affordable computational resources, have made it possible to model the complex distribution of molecular structures directly. This capability allows for efficient exploration of broader virtual chemical space than traditional virtual screening approaches [25, 26, 27], with more fine-grained and controllable navigation compared to conventional combinatorial fragment-based methods [28, 29].

Despite the promise of generative design methods, their adoption has remained somewhat limited. One major challenge is that most generative models often produce synthetically intractable molecular structures when targeting specific design goals, such as optimizing property scores [30, 31, 32, 33]. When designed molecules cannot be synthesized and validated in the lab at a reasonable cost, their practical value is limited. In addition, the actionability of computer-aided molecular design is crucial for achieving rapid design cycles in early-stage drug discovery [34, 35] and for enabling

closed-loop autonomous discovery [36, 37, 38, 39, 40]. This issue, though long-standing and observed in earlier combinatorial algorithms, has become more pronounced with the advent of deep generative AI models.

There has been considerable interest in incorporating synthetic accessibility into molecular design methods [30, 33]. The most straightforward strategy is to incorporate a heuristic synthetic accessibility score as a design criterion [30, 41]. However, quantitatively estimating synthetic accessibility must account for factors such as regioselectivity, functional group compatibility, and building block availability, all of which contribute to a rugged structure-synthesizability landscape that makes the design of such scores an ongoing challenge. While notable efforts have been made to quantify synthetic accessibility [42, 43, 44], achieving reliable quantification through a simple heuristic remains a distant goal. Advances in the sample efficiency of molecular optimization [45, 46] have made it feasible to conduct explicit retrosynthesis analysis for each designed molecule as a means of evaluating synthetic accessibility [47]. However, the computational overhead remains significant, making this approach impractical for most existing generative design models. Furthermore, unsynthesizable molecules provide no learning signal, and when a model continues to propose such molecules, the sparse feedback hinders effective learning.

A more ideal and effective approach to synthesizable molecular design, in our opinion, involves constraining the design process to focus exclusively on synthesizable molecules by designing *synthetic pathways* rather than simply designing *structures*. With the growing role of massive make-on-demand libraries—whose sizes exceed what can be fully enumerated [27, 48, 49, 50, 51, 52, 53]—this class of methods has gained increasing interest, spanning both traditional Monte Carlo techniques [54, 55, 56, 57, 58, 59, 60] and more recent generative approaches [61, 62, 63, 64, 65, 66, 67]. Despite this progress, current synthesis-centric methods still fall short in terms of controllability and efficiency in navigating the chemical space, and as a result have not achieved widespread adoption. These methods that rely on property scores to guide the design process often exhibit lower efficiency in optimizing expensive oracles compared to structure-centric methods [45], as they model the synthetic action sequence-property landscape, which is inherently more complex than the structure-property landscape. Some recent models offer a more controlled exploration of local chemical spaces by using a reference input molecule [61, 64, 65, 67]. However, the low reconstruction rates for theoretically feasible molecules suggest a risk of partial collapse during decoding, suggesting that the practical chemical space accessible to these models is smaller than the theoretically accessible space. This behavior leads to certain regions of the chemical space becoming inaccessible, regardless of the input, ultimately hindering the overall design process.

Herein, we build upon our previous contributions [65, 67] to introduce SynFormer, a generative AI framework designed for efficient and controllable navigation within a synthesizable chemical space. Like earlier synthesizable design methods, SynFormer is a synthesis-centric approach, generating synthetic pathways by readily available building blocks through robust chemical transformations, ensuring synthetic traceability subject to the limitations of those transformation rules. SynFormer distinguishes itself from previous models with its compute-efficient and scalable transformer architecture [14], which empirically shows improvements in performance as the training dataset grows. To select suitable molecular building blocks from the large, discrete, and multimodal space of commercially-available options, SynFormer incorporates a denoising diffusion model as a token head module [15]. The framework is fully end-to-end differentiable, enabling effective training and optimization.

We demonstrate the utility and versatility of the SynFormer framework through two instantiations: (1) SynFormer-ED, an encoder-decoder model that generates synthetic pathways corresponding to a given input molecule for exact or approximate reconstruction of that input, and (2) SynFormer-D, a decoder-only model for generating synthetic pathways that is amenable to fine-tuning towards specific property goals. Both models are trained on a simulated chemical space derived from a curated set of 115 reaction templates and 223,244 commercially available building blocks, extending beyond Enamine’s REAL Space [50]. We validate SynFormer’s capabilities by demonstrating success on (a) reconstructing molecules within both the Enamine REAL and ChEMBL [68] chemical spaces; (b) local synthesizable chemical space exploration given a reference molecule; and (c) global synthesizable chemical space exploration guided by black-box property prediction model. These results not only highlight SynFormer’s ability to navigate synthesizable chemical space using varied control strategies but also demonstrate its practical applicability in real-world molecular design use cases. The scalability of the SynFormer framework with respect to both training data and model size suggests considerable potential for further performance enhancement. In conclusion, these findings underscore SynFormer’s flexibility and potential to impact molecular design across various domains, including drug development and materials science.

2 SynFormer: a generative framework for synthesizable molecular design

SynFormer is a generative framework designed for modeling synthesizable chemical spaces, as schematically illustrated in Figure 1. Although it is challenging to predict which molecules can be synthesized without experimental validation, the progress of make-on-demand libraries such as Enamine REAL Space [50], GalaXi [51], and eXplore [52] has

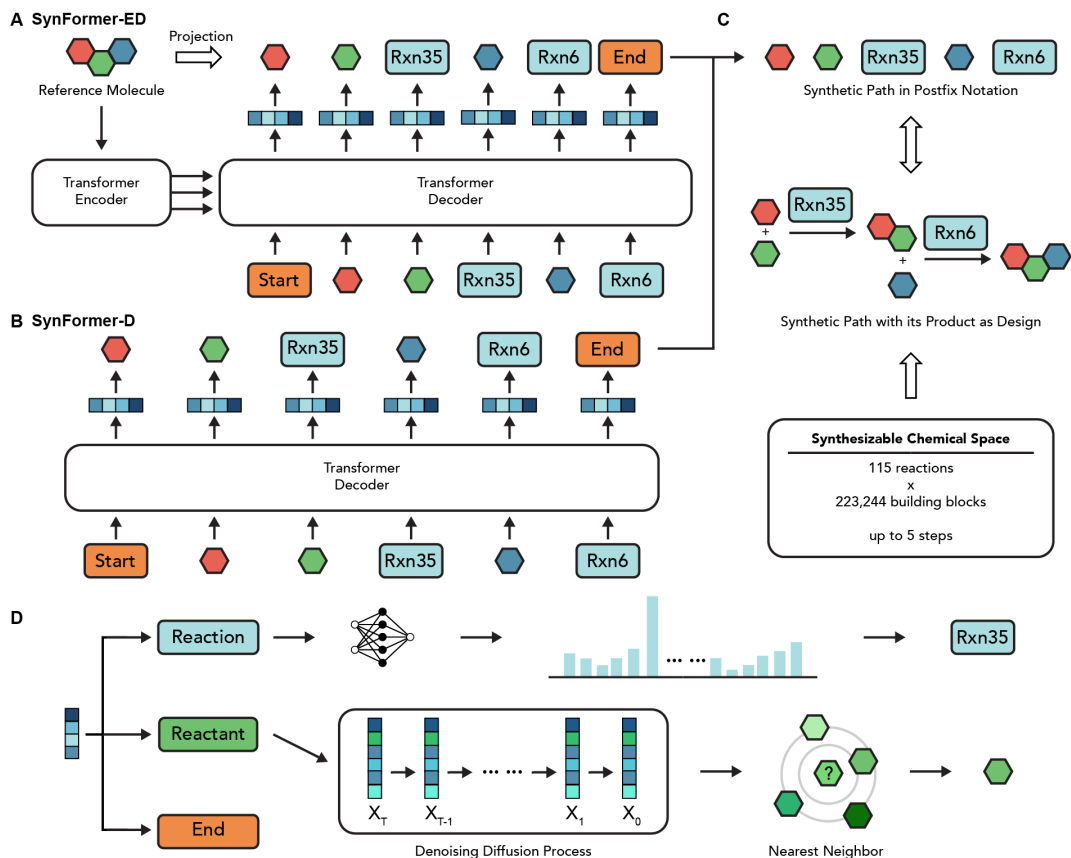


Figure 1: Schematic illustration of the SynFormer framework and architecture. (A) The SynFormer-ED architecture is an encoder-decoder that takes a molecule as input and outputs a synthetic route to the same or an analogous molecule. (B) SynFormer-D is a decoder-only framework designed to generate synthetic routes. (C) Synthetic routes are tokenized using a postfix notation to make them amenable to autoregressive generation. Routes are constructed by applying 115 reactions to a set of 223,244 molecular building blocks, covering a synthesizable chemical space estimated as $> 10^{60}$ molecules. (D) During generation, a token generated by the transformer is first classified by token type. If the token represents a reaction or reactant (building block), it undergoes an additional classification process to select the appropriate reactions or a denoising diffusion process followed by a nearest neighbor search to select the appropriate building block(s).

demonstrated that molecules constructed by connecting purchasable molecular building blocks through a series of curated, reliable reactions have a high likelihood of being synthesizable. Therefore, within the scope of this work, we define synthesizable chemical space as encompassing all molecules that can be formed within the reaction network by linking purchasable building blocks through up to five steps of known chemical transformations, though in principle, there is no limit to the length of pathways that can be generated, and building blocks and reactions beyond our selected list can also be utilized.

To effectively model chemically relevant and practically useful spaces, we adapted the reaction set used to construct REAL Space, which primarily focuses on bi- and tri-molecular reactions that couple multiple building blocks together. We further augment this set with additional common reactions, such as functional group interconversions, to better simulate a general organic synthesis endeavor, resulting in a total of 115 reaction templates. As the set of purchasable molecular building blocks, we use Enamine's U.S. stock catalog to ensure high availability, which also approximates those used to generate REAL Space. Although the proportion of synthesizable molecules may vary due to the broader focus of the chemical space, our model theoretically covers a chemical space broader than tens of billions currently represented in the Enamine REAL Space, making it suitable for extensive exploration during the molecular design stage. Both the choice of templates and the choice of building blocks can be straightforwardly modified prior to retraining and are not inherent to SynFormer's approach.

Generating synthetic pathways requires careful consideration of their data structure. Bottom-up synthesis planning involves sequential decision-making with branches and merges, encompassing data modalities that span reaction templates and molecular building blocks, which represent small and very large discrete design spaces, respectively. To manage this complexity, we adopt the postfix notation to represent synthetic pathways linearly, as described in Luo et al.'s [67] (Figure 1C). Specifically, we use four types of tokens to represent synthetic paths: a start token [START], an end token [END], reaction tokens [RXN], and building block tokens [BB]. Similar to the postfix notation of mathematical formulae, the reactions are placed after the reagents, allowing the sequence to be processed in a step-by-step manner without loss of generality in the synthetic pathways (e.g., the ability to accommodate any linear or convergent sequence). This linear notation also enables the use of autoregressive decoding via a transformer architecture [14], which is widely recognized as a scalable model backbone. A stack of standard transformer layers processes the sequence as it is decoded, taking the previous tokens and outputting the embedding for the next token. Each token embedding then passes through a multi-layer perceptron (MLP) classification head to determine the token type. If the token is [END], the generation process terminates. If the token is [RXN] or [BB], the embedding is used to select which building blocks to add and which reaction type to perform.

The selection of reactions can be relatively straightforward using an additional classification head; however, the number of purchasable building blocks can easily exceed hundreds of thousands, millions, or perhaps billions, and will continue to grow over time. Instead of relying on a static classification network, our past work [65, 67] has involved generating Morgan fingerprints [69] and then retrieving the nearest building block from a list of candidates to enable generalization to unseen building blocks. To further enhance this building block selection process, we adopt a denoising diffusion probabilistic module [15] to predict the posterior distribution of molecular fingerprints conditioned on the token embedding. Specifically, we model the molecular fingerprints using an n -dimensional joint Bernoulli distribution (where n is the length of the fingerprint) and employ a multinomial diffusion framework [70] to handle the discrete nature of the data. Using the denoising diffusion module as the building block token head offers several advantages: while maintaining the end-to-end differentiable nature, it allows for scaling to longer fingerprints (2048 in this work compared to 256 in previous studies), thereby avoiding bit collisions that can occur with shorter fingerprints and enabling generation at a finer resolution. Additionally, this approach more effectively handles multimodal distributions, where multiple "true" answers exist for a single decision—a particularly common scenario when modeling synthetic pathways.

Based on this generative modeling framework, we implemented two specific models: SynFormer-D, which straightforwardly implements a transformer decoder with the aforementioned reaction classification heads and building block diffusion heads (Figure 1B), and SynFormer-ED, which adds a standard transformer encoder to condition generation on an input SMILES string [71] (Figure 1A). To train the models, we uniformly sample synthetic pathways from the synthesizable chemical space simulated from the reaction templates and building block list as training data. Detailed descriptions of the model architectures, training protocols, and hyperparameter settings are provided in the Methods section and Supplementary Information.

3 Results

3.1 Molecule reconstruction and chemical space coverage

An important factor that affects a model's performance in molecular design applications is its coverage of chemical space [72]. Even though the theoretical synthesizable chemical space constructed from reaction templates and building blocks is vast, previous models have shown that a trained model can struggle to generate molecules that are known to be synthesizable [64, 65, 57, 67]. This limitation can hinder the model's effectiveness in molecular design if it cannot access the chemical space where the optimal design resides. To assess the coverage of chemical space, we tested the SynFormer-ED model on molecule reconstruction: how often SynFormer-ED can reconstruct a viable synthetic path given a large set of potential molecules of interest.

We evaluated reconstruction on a randomly selected subset of 1,000 molecules from ChEMBL Database and Enamine's REAL Diversity Set (Figure 2A and B). SynFormer-ED successfully recovers 66% of molecules from REAL, notably higher than that achieved by ChemProjector [67] and SynNet [30]. ChEMBL, by comparison, is a curated database of bioactive molecules with drug-like properties, which includes molecules whose syntheses may require building blocks or reactions beyond SynFormer's training set. This results in a reconstruction rate of 20%—lower than in REAL Space but still an improvement over previous models.

We further demonstrate that the model's performance, as measured by the binary cross-entropy (BCE) of the molecular fingerprints (FPs) [69], improves with increasing model size and training data (Figure 2C). Notably, we also observed that the performance of the largest model did not outperform the second-largest model when insufficient data was available. This finding suggests that estimating training performance based on early epochs may not be an effective

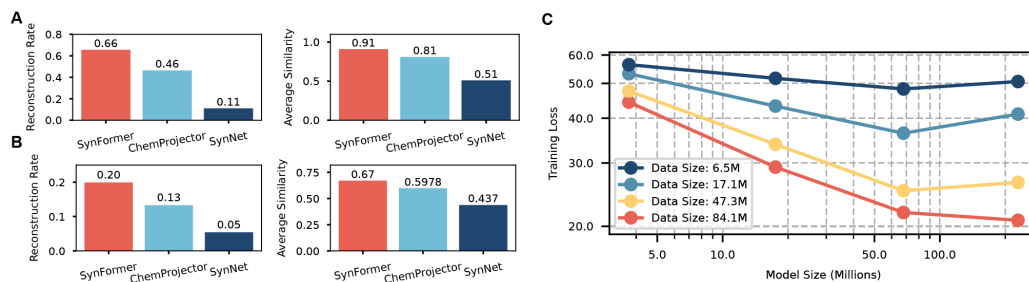


Figure 2: Model performance on molecular reconstruction. (A and B) Comparison of the reconstruction rate and average structural (Tanimoto) similarity between input and output molecules for SynFormer-ED, ChemProjector [67], and SynNet [65] on 1,000 randomly selected molecules from (A) REAL Diversity Set and (B) ChEMBL Database. (C) Scaling of model performance is measured by the training loss (binary cross-entropy (BCE) of the molecular fingerprint (FP) prediction) as model size and training data size increase.

strategy when the model size surpasses a certain threshold [73]. These results emphasize the scalability of the architecture [74] and indicate that further performance improvements can be achieved by increasing both the amount of data and computational resources. They also underscore the importance of scaling model size and training data in tandem to fully realize performance gains. As SynFormer is trained using simulated pathways according to its building blocks and reaction rules, there is virtually no limit to the size of training data available.

3.2 Local chemical space exploration with SynFormer

In this section, we present the application of SynFormer-ED in exploring the local chemical space around a given reference molecule, projecting it into a synthesizable chemical space. We demonstrate two use cases of our model in molecular design: (1) generating synthesizable analogs for unsynthesizable designs and (2) hit expansion within synthesizable chemical space.

3.2.1 Synthesizable analog generation

In molecule reconstruction experiments, we observed that SynFormer-ED is able to generate structurally similar outputs even when provided with unsynthesizable (according to its building blocks and reactions) molecules. This capability allows SynFormer-ED to be applied to generate synthesizable structural analogs of unsynthesizable designs, as illustrated in Figure 3A. The implicit goal is to preserve as much of the overall structure and key pharmacophores as possible, effectively transforming non-feasible molecules into synthesizable compounds, akin to “projecting” the input molecule onto the synthesizable space.

To evaluate the impact of synthesizable projection, we first followed the experimental setup of Luo et al. [67] and used SynFormer-ED to generate synthesizable analogs for molecules identified as unsynthesizable by ASKCOS [75] in [30]. The original molecules were optimized using *de novo* design methods targeting ten multi-objective scores, each ranging from 0 to 1, where higher values are preferable. These scores encompass structural features, physicochemical properties, and metrics related to similarity or dissimilarity to known drugs, as well as the presence of specific substructures [76]. After the generation of analogs, we evaluated them using the same objective scores that the original designs were optimized for. As shown in Figure 3B, the synthesizable projection effectively eliminated the peak in synthetic accessibility (SA) scores [42] at around 6-7 observed in the original designs, shifting the distribution toward more easily synthesizable molecules. Additionally, a significant fraction of the generated analogs exhibited objective scores comparable to those of the original molecules (Figure 3C and Figure S2 in the Supporting Information). We demonstrate two examples in Figure 3D. Starting with high-scoring but seemingly-unsynthesizable molecules, SynFormer-ED generated analogs that preserved key structural features while correcting the unsynthesizable fragments. The generated analogs retained objective scores comparable to those of the original molecules and are inherently synthesizable (Figure S4 in Supporting Information), demonstrating SynFormer-ED’s ability to transform an unsynthesizable molecule into a synthesizable analog while maintaining structural integrity.

To further validate the utility of SynFormer-ED in realistic drug discovery scenarios, we explored its application in structure-based drug design (SBDD) [77, 78, 79]. While numerous SBDD algorithms can generate ligands with promising predicted binding affinities (e.g., via 3D pocket-conditioned generation), these methods are often criticized for overlooking synthetic accessibility [80, 81]. SynFormer-ED addresses this gap by projecting the outputs of such

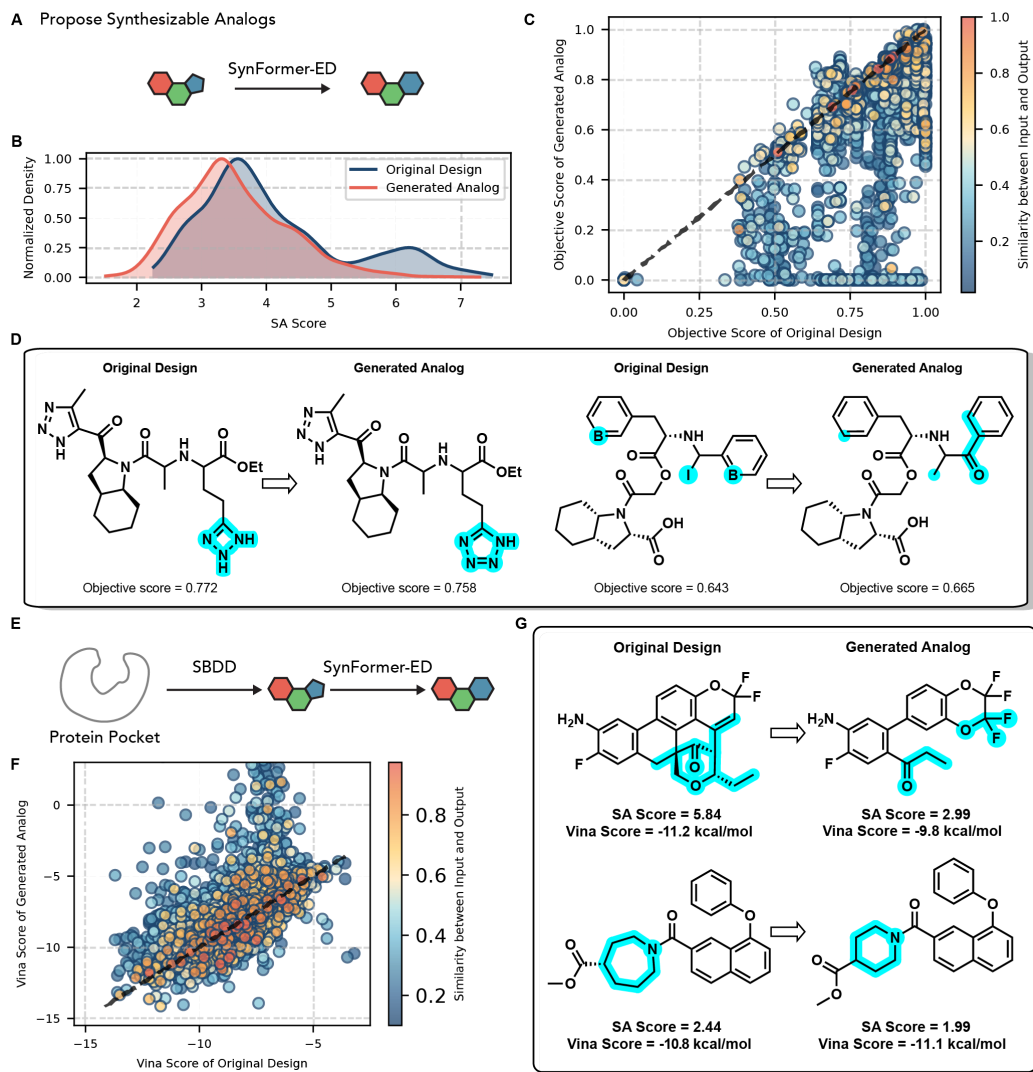


Figure 3: Application of SynFormer in projecting unsynthesizable design into synthesizable chemical space. (A) Schematic illustration of SynFormer-ED generating synthesizable structural analogs for unsynthesizable molecules. (B) Normalized distribution of SA scores for the originally designed molecules and their corresponding synthesizable analogs. Note that the distributions are normalized to a peak value of 1 for a clearer comparison. (C) Scatter plot comparing objective scores of originally designed molecules versus their generated analogs. Points are colored by the structural similarity between them, showing that structurally similar analogs tend to possess close properties. (D) Examples of originally designed unsynthesizable molecules and their SynFormer-generated synthesizable analogs. Objective scores are shown beneath each molecule, demonstrating comparable activities for the generated analogs. The modified fragments or atoms are highlighted in light blue. (E) The workflow shows SynFormer-ED generating synthesizable analogs for ligands generated by structure-based drug design. (F) Scatter plot comparing the Vina docking scores of originally designed ligands and their generated analogs. Points are colored based on the structural (Tanimoto) similarity between the input and output, showing strong agreement in general and an ability to generate analogs with comparable scores. (G) The original design and generated analog for Estrogen receptor alpha, with their SA Score and Vina score below. The modified fragments or atoms are highlighted in light blue.

models into synthesizable chemical space (3E). We demonstrate this approach using Pocket2Mol [82] as an exemplary SBDD generative model.

Targeting 15 protein targets from the LIT-PCBA dataset [83], we used SynFormer-ED to generate synthesizable analogs of molecules initially designed by Pocket2Mol. SynFormer-ED successfully produced analogs that preserve similar

Vina docking scores [84] while ensuring that each design has a synthetic pathway (Figure 3F). Two examples are shown in Figure 3G: in the first example, the original design had a Vina score of -11.2 kcal/mol but a high synthetic accessibility (SA) score of 5.84, indicating poor synthesizability. SynFormer-ED generated a structurally similar analog with an improved SA score of 2.99 while exhibiting a Vina score of -9.8 kcal/mol. In the second example, although the original design was likely synthesizable with access to the appropriate building blocks, SynFormer-ED further reduced the SA score from 2.44 to 1.99 with minor improvements in Vina score. These results demonstrate SynFormer-ED’s potential to complement existing SBDD algorithms by generating more practical and synthesizable drug candidates, suitable for experimental validation and further development.

3.2.2 Hit expansion within synthesizable chemical space

Building on SynFormer-ED’s ability to explore local synthesizable chemical space around reference molecules, we applied it to hit expansion within synthesizable chemical space. By allowing the model to select sub-optimal intermediate choices of building blocks or reactions via a beam search, SynFormer-ED can generate sets of close analogs around known hits, as illustrated in 4A. This approach aims to produce additional molecules that preserve the core motifs and overall structural integrity, potentially enhancing design objectives and expanding the pool of candidates for downstream processes, such as hit-to-lead optimization.

Following the setup in [85], we screened ZINC250k [86], a subset of ZINC database [87], using a predictive model for inhibition activity against c-Jun NH2-terminal kinase 3 (JNK3) [88, 89, 90]. The top 10 scoring molecules, with scores ranging from 0.49 to 0.68 (the higher, the better), were selected as hits and input into SynFormer-ED to generate analogs, which were then evaluated with the JNK3 inhibition predictor. As a baseline, a nearest neighbor search within Enamine REAL (approximately 7 billion molecules) was performed for the hits, retrieving approximately 100 of the most similar molecules for each hit and evaluating them for JNK3 inhibition. While both SynFormer-ED and the nearest neighbor search identified additional molecules with scores higher than those from a simple screening of ZINC250k, SynFormer-generated analogs exhibited a larger enrichment in the high-scoring region (Figure 4B). We present two sets of examples in Figure 4C. Though both methods produced analogs retaining the overall structure with higher predicted inhibition than the original hits, the SynFormer-ED analogs consistently achieved higher scores compared to those from the nearest neighbor search.

To further validate SynFormer-ED’s capability in a more realistic hit expansion scenario, we used experimentally verified ligands for Human Pyruvate Kinase M2 (PKM2, PDB ID: 3ME3) [91] and Lysine Acetyltransferase 2A (KAT2A, PDB ID: 5MLJ) [92] as reference molecules and evaluated their Vina docking scores. SynFormer-ED was used to expand the reported ligands, generating 191 analogs for PKM2, with 179 of them with Tanimoto similarities greater than 0.5 and Vina score changes lower than 1 kcal/mol. For KAT2A, SynFormer-ED generated 200 analogs, 111 of which met the criteria of Tanimoto similarity above 0.5 and Vina score changes lower than 1 kcal/mol. Figure 4D and E present three generated analogs for ligands targeting PKM2 and KAT2A, respectively, showing competitive Vina scores while preserving the overall structure, along with their synthetic pathways (see Figure S11 and 12 in the Supporting Information for more examples). Notably, in the case of KAT2A, the synthetic pathways of the generated analogs are diverse despite the structural similarity of the final analogs, highlighting that the SynFormer-based approach is not limited by the synthetic pathways of the input hits. This flexibility allows the generation of analogs with distinct synthetic routes, which traditional path-based enumeration strategies [93, 85] cannot achieve. Combined with these results, SynFormer-ED demonstrates significant potential in hit expansion, enabling the generation of practical, synthesizable analogs suitable for further drug development stages.

3.3 Global chemical space navigation with SynFormer

In this section, we demonstrate the application of SynFormer to global chemical space navigation, enabling the optimization of molecular properties treated as a black-box function across broad chemical spaces [45]. This process, also known as *de novo* molecular optimization, is key to exploring novel molecular designs. We first show that SynFormer-D can be fine-tuned using reinforcement learning to generate high-scoring molecules (Figure 5A). Furthermore, we demonstrate how SynFormer-ED can be integrated as a mutation step within an evolutionary algorithm (Figure 5C), achieving state-of-the-art sample efficiency while constraining the design space to synthesizable chemical space.

3.3.1 Molecular optimization by fine-tuning SynFormer-D with reinforcement learning

We fine-tuned SynFormer-D using reinforcement learning (RL) to guide its generation toward high-scoring molecules. Specifically, we adopted a variant of the REINFORCE algorithm [94], which iteratively generates a batch of molecules, evaluates their properties, and fine-tunes the model parameters with the goal of maximizing the desired properties (detailed methods can be found in the Methods section). In our experiments, we optimized predicted binding affinity

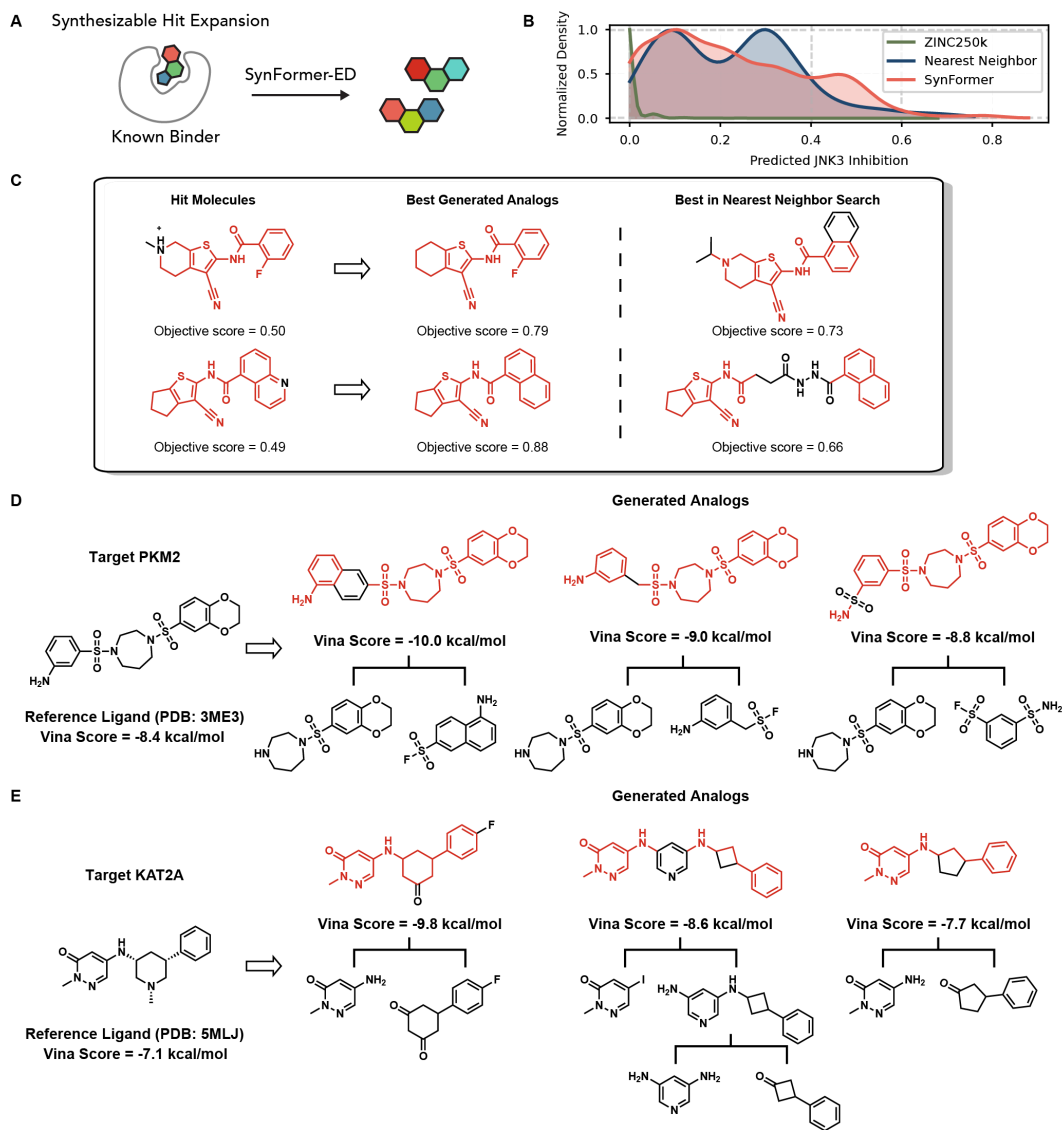


Figure 4: Application of SynFormer in hit expansion. (A) Schematic illustrating SynFormer-ED expanding a known hit compound into structurally similar, synthesizable analogs. (B) Normalized distribution of predicted JNK3 inhibition scores for screening ZINC250k, nearest neighbors in Enamine REAL of the hits, and SynFormer-generated analogs of the hits, highlighting the enrichment of high-scoring ligands among SynFormer analogs. Note that the distributions are normalized to a peak value of 1 for a clearer comparison. (C) Examples of hits from screening ZINC250k and corresponding best-scored generated analogs, alongside the best-scored molecules identified in the nearest neighbor search within Enamine REAL, demonstrate SynFormer-ED's ability to generate high scoring, synthesizable compounds. The motifs retained from the hits are colored in red. (D and E) Expanding experimentally validated ligands for the PKM2 target (PDB: 3ME3) (D) and the KAT2A target (PDB: 5MLJ) (E). Three representative analogs are shown for each target, along with their synthetic pathways. The structural motifs retained from the hits are colored in red.

against the dopamine receptor D_2 (DRD2) [95] as an exemplary oracle function. As shown in Figure 5B, our method (denoted as SF-RL) successfully biases generation toward high-scoring molecules and outperforms several popular optimization methods [96, 29]. Although SF-RL's sample efficiency is lower than the most efficient algorithms like GraphGA [29] and REINVENT [19], these results confirm the feasibility and effectiveness of fine-tuning SynFormer-D for *de novo* molecular optimization. Given that much of the algorithm's design space of RL is yet to be explored, the potential for additional algorithmic improvements in RL-based fine-tuning remains substantial.

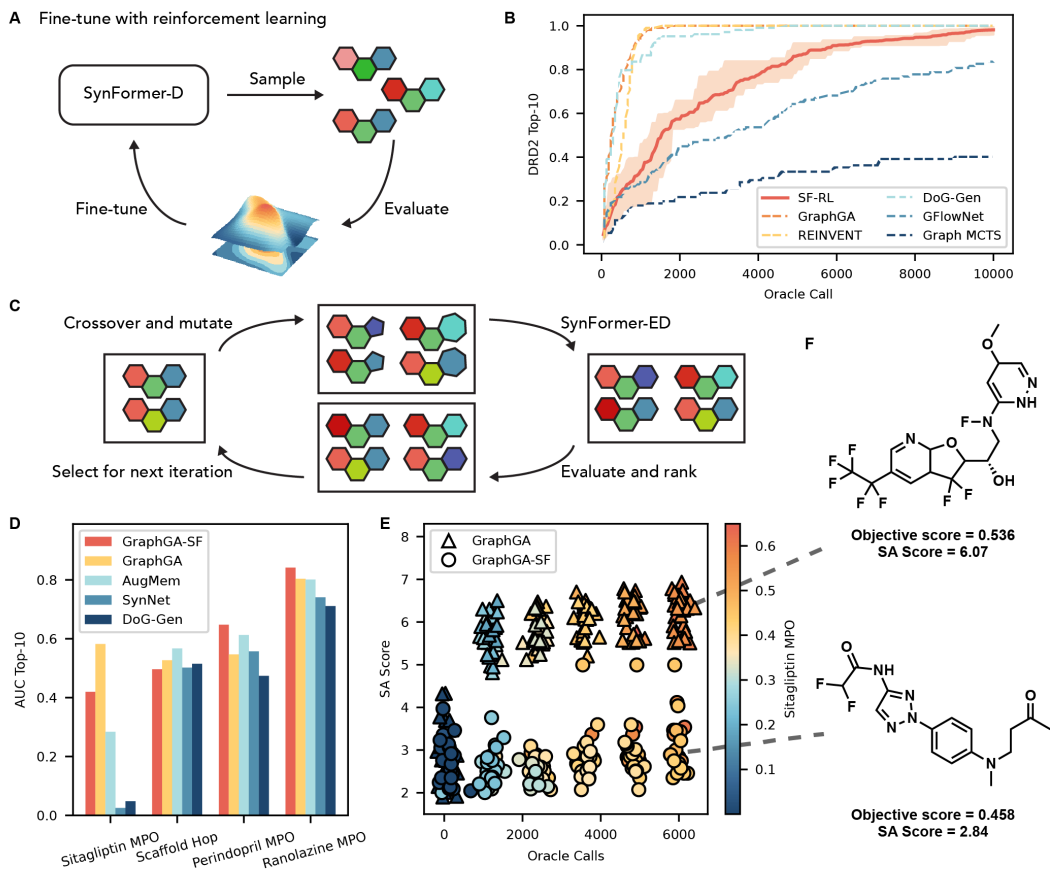


Figure 5: Application of SynFormer in global chemical space exploration. (A) Illustration of fine-tuning SynFormer-D with reinforcement learning. (B) Performance comparison of SynFormer fine-tuning with reinforcement learning (SF-RL) against other popular methods, showing the average top-10 molecule scores versus the number of oracle calls [45]. The plots represent the mean performance curves of 5 independent runs, with the shaded region for SF-RL indicating the range across these runs. (C) Illustration of a genetic algorithm with SynFormer-ED used for mutation steps. (D) AUC Top-10 performance comparison across different molecular design methods (GraphGA-SF, GraphGA [29], AugMem [97], SynNet [65], DoG-Gen [64]) for four tasks from GuacaMol [76]. (E) Distribution of SA Scores [42] for the top 25 molecules at various optimization steps, with colors representing objective scores, demonstrating how SynFormer effectively constraints its design space to synthesizable space exclusively. (F) The best molecules generated by GraphGA (top) and GraphGA-SF (bottom) show that GraphGA-SF identifies a more synthetically tractable candidate with a reduced SA score, albeit with a minor sacrifice in the objective score.

3.3.2 Molecular optimization by using SynFormer-ED as a mutation operator within a genetic algorithm

Lastly, we demonstrate a genetic algorithm framework incorporating SynFormer-ED projection as a mutation step. After the standard processes of crossover and mutation to diversify the candidate pool, SynFormer-ED projects all candidates into synthesizable chemical space. This step not only corrects unsynthesizable fragments and linkages, ensuring synthetic accessibility, but also introduces additional diversity into the candidate pool. Generated molecules are then scored and ranked, with the top-performing ones selected for further iterations (Figure 5C). For crossover and mutation rules, we follow the GraphGA algorithm [29], a well-established molecular optimization method that regularly performs well in molecular optimization benchmarks [45, 98] and real-world applications [41].

We evaluated the performance of this genetic algorithm, denoted as GraphGA-SF, across four tasks from the GuacaMol benchmark [76]: Sitagliptin MPO, Scaffold Hop, Perindopril MPO, and Ranolazine MPO. As shown in Figure 5D, GraphGA-SF demonstrated optimization efficiency comparable to the original GraphGA and the state-of-the-art RL method, Augmented Memory (denoted as AugMem) [97], while significantly outperforming previous methods for synthesizable molecular optimization. Notably, in the Sitagliptin MPO task, where the primary objective score often

avored unsynthesizable molecules, conventional methods without synthesizability constraints tended to generate exclusively unsynthesizable compounds (See SA Score in Figure 5E and structure in Figure 5F). On the other hand, synthesis-centric design methods, such as SynNet [65] and DoG-Gen [64], struggled to find meaningful optimization signals due to the sparse signal and their limited coverage of chemical space. In contrast, GraphGA-SF matched the performance of GraphGA while ensuring each design has a synthetic pathway like other synthesis-centric methods. We highlight the best designs from each genetic algorithm in Figure 5F, showing that the molecules generated by GraphGA-SF were significantly more synthesizable, with only a modest sacrifice in the main objective score.

Our results demonstrate that by combining the optimization capability of GraphGA with SynFormer-ED's ability to restrict the design space to synthesizable compounds, GraphGA-SF can effectively and efficiently optimize molecular properties while maintaining synthetic feasibility. Furthermore, SynFormer-ED's integration as a mutation operator exemplifies its potential as a generalizable component in molecular design algorithms, allowing these frameworks to confine their search to synthesizable chemical spaces. This adaptability enables SynFormer-ED to leverage ongoing advances in *de novo* molecular design algorithms, offering a versatile and robust solution to the problem of synthesizable molecular design.

4 Discussion

In this paper, we introduced SynFormer, a generative modeling framework designed to efficiently explore and navigate synthesizable chemical space. We curated a set of reaction rules that allow the model to construct a synthesizable chemical space beyond Enamine REAL Space, ensuring that the designed molecules are feasible for synthesis subject to the reliability of such rules. By integrating transformer architectures with a diffusion module, SynFormer enhances the ability to make beneficial choices of building blocks and reactions during autoregressive bottom-up pathway generation, surpassing previous models for synthesizable molecular design. We demonstrated two key applications of SynFormer: local chemical space exploration and global optimization within synthesizable chemical space, both of which yielded successful results. Additionally, we showed that SynFormer is a scalable architecture, with the potential for further performance improvements with increased resources. While our focus was on applications in drug discovery, the models are readily applicable to other domains of small organic molecule design [4, 7].

There are several areas where we expect further improvements can be made. First, the reconstruction rate of SynFormer-ED is not perfect, suggesting that certain regions of synthesizable chemical space may remain inaccessible, even though the model possess the reaction templates and building blocks that should make it possible to reach them. Additionally, RL-based fine-tuning in SynFormer-D currently requires a relatively large number of oracle calls to achieve competitive results. Another limitation lies in the quality and coverage of reaction templates and building blocks. These templates do not account for stereochemistry, meaning that chiral centers can only be inherited from building blocks, not installed during synthesis. Furthermore, the Enamine building blocks we choose are designed for combinatorial library construction and do not comprehensively cover commercially available materials. These constraints may limit the model's applicability to somewhat flat, linearly-constructed sp^2 -rich structures and restrict its use for more complex molecules like natural products [99]. Expanding the SynFormer framework to include more sophisticated reactions and a broader range of building blocks is possible, but will require further validation.

5 Methods

Synthesizable Chemical Space Construction

A synthesizable chemical space is defined as a set of molecules that can be accessed from purchasable starting materials (building blocks) through a series of reliable chemical transformations. One well-known example of such a chemical space is Enamine REAL Space (REadily Accessible) [50], which consists of 48 billion molecules. These molecules can be synthesized by applying 1-3-step synthetic transformations derived from approximately 170 well-validated parallel synthesis protocols to over 143,000 qualified building blocks. The synthesis and delivery of molecules within the REAL Space typically takes 3-4 weeks, with an average synthesis success rate of 80%, which attests to the validity of the way of constructing synthesizable chemical space. In this work, we aim to cover the REAL Space and extend it further by incorporating additional reactions and a greater number of reaction steps to mimic a general organic synthesis endeavor.

Reaction templates

Reaction rules are encoded as strings using the SMARTS grammar [100], known as reaction templates. We created a set of 86 reaction templates, which directly implement the 169 reactions used in the January 2022 version of the REAL

Space. This set includes both bi-molecular and tri-molecular reactions, with deprotection reactions treated as separate entities. Reactions that can be represented by the same transformation rule are encoded using a single SMARTS string. The combination of this reaction template set and the Enamine building block list should cover a space similar to the REAL Space.

In addition to linking building blocks, we also considered more general chemical synthesis processes. We manually selected 29 additional reaction templates from Hartenfeller et al.'s [101] and Button et al.'s [102] template collection, which were not covered by the initial 86 reactions. These reactions include well-known organic transformations such as the Wittig reaction [103], Buchwald-Hartwig reaction [104], and various functional group conversions. This expanded set of reactions complements the REAL reactions by providing a closer approximation to general laboratory organic synthesis.

We acknowledge certain limitations in our templates: no reaction conditions, regioselectivity, or functional group compatibility are currently considered. The templates represent potential transformations, meaning that the chemical space we constructed likely represents a superset of the actual synthesizable chemical space. Furthermore, stereochemistry is not considered in our reaction rules, so stereocenters can only be inherited from the building blocks and cannot be installed during synthesis.

Building blocks

The building block list comprises reagents or intermediates that are commercially available and can be directly purchased from the market. This list can include any set of purchasable molecules, including those found in Enamine REAL Space. For this study, we utilized Enamine's U.S. in-stock collection of building blocks, which was available for delivery as of October 1st, 2023. The building block list contains a total of 223,244 molecules. This selection was made to ensure the accessibility and immediate availability of the building blocks for synthesis.

Representing synthetic pathways using postfix notation

In SynFormer, synthetic pathways are represented using the *postfix notation of synthesis* [67]. In this format, the operators (chemical reactions) follow their operands (molecular building blocks), providing a structured and concise way to represent reaction sequences. Each synthesis pathway is encoded as a sequence of tokens, where each token represents either a reaction step or a building block. This linear sequence simplifies the representation of converging synthetic pathways by eliminating the need for "parentheses" to indicate the order of operations, making it particularly amenable to autoregressive modeling. This approach enables efficient handling of complex synthetic routes while ensuring that each step in the pathway is captured in the correct order.

SynFormer

The SynFormer framework is a generative modeling architecture designed to explore synthesizable chemical space by generating synthetic pathways for molecules rather than molecular graphs. SynFormer is based on a transformer architecture with a diffusion module to select molecular building blocks. We designed two variants of the model to handle distinct molecular design tasks:

- SynFormer-D, a decoder-only model designed for generating new molecules or optimizing molecules based on feedback from a black-box objective function.
- SynFormer-ED, an encoder-decoder model that generates synthetic pathways based on a given reference molecule, enabling local chemical space exploration.

Transformer decoder

SynFormer adopts the transformer decoder architecture as its backbone [14]. The decoder generates postfix notations of synthesis in an auto-regressive manner, which is, predicting the next token according to previous tokens. We assign embedding vectors to each reaction template and the [START] token, and convert building blocks fingerprint into continuous-valued vectors with a multi-layer perceptron (MLP) in the beginning. To indicate the position of each token within the sequence, we add positional encodings [14] to the embedding vectors.

The transformer decoder gathers information for each token from its preceding tokens via the multi-head attention mechanism [14], where each token's embedding vector is first projected to query, key, and value vectors, and then updated by aggregating the values of other tokens weighted by the softmax-ed inner product of the query vector and the other token's key vector.

Following the transformer decoder, we use a classifier based on an MLP to predict the next token’s type. If it is predicted as an [END] token, the decoding process will be terminated. If it is predicted as a reaction, another classifier is used to predict the probability of each reaction template being the next token. When the next token is predicted as a building block, a diffusion module is used to generate molecular fingerprints which are subsequently used to retrieve building blocks from the database.

Transformer encoder

The transformer decoder can be conditioned by a transformer encoder, i.e., given reference molecule represented by a SMILES string. The transformer encoder produces vector representations for each token in the SMILES string, which the decoder can attend to via the cross attention mechanism [14].

Diffusion module for building block fingerprints

We build a denoising diffusion model to learn the distribution of molecular fingerprints, represented as n -dimensional binary vectors, where n denotes the number of bits in the fingerprint. This distribution is modeled as a joint Bernoulli binary distribution of n dimensions.

During training, the forward diffusion process perturbs the fingerprint vector by randomly flipping each bit according to the following noise distribution:

$$q(x_t^{(j)} | x_0^{(j)}) = \text{Bernoulli}(x_t^{(j)} | \bar{\alpha}_t x_0^{(j)} + (1 - \bar{\alpha}_t) \cdot \frac{1}{2}), \quad (1)$$

where $x_0^{(j)}$ represents the j -th bit of the ground truth fingerprint vector, and $x_t^{(j)}$ represents the perturbed bit. Here, $\bar{\alpha}_t$ is a value that controls the noise level, which we define using a formulation similar to multinomial diffusion [70]. $\bar{\alpha}_t$ starts at 1 and monotonically decreases to 0. When $\bar{\alpha}_t = 0$, the distribution q becomes the uniform Bernoulli distribution, where the noise level reaches the maximum. The denoiser network learns the bitwise Bernoulli distribution of ground truth fingerprints conditioned on the perturb fingerprint and the embeddings from the transformer decoder. It is trained by minimizing the binary cross entropy loss that measures the discrepancy between the distribution and the ground truth fingerprint.

The inference process is known as the reverse diffusion process, where we start with a random bit vector drawn from the uniform Bernoulli distribution, and denoise it iteratively with the denoiser network. Finally, the denoised bit vector is used as a fingerprint to retrieve building block molecules from the database.

Reinforcement learning fine-tuning

For global optimization tasks, SynFormer-D was fine-tuned using a variant of the REINVENT algorithm [19] based on REINFORCE [94], a type of reinforcement learning (RL). Our RL setup was designed to minimize the following loss function, which incentivizes the model to bring the pseudo-log-likelihood (PLL) of generating a synthetic pathway closer to the objective score of the product molecules:

$$L(\theta) = \mathbb{E} \left[(\text{PLL}_\theta(p, m) - \sigma \cdot S(m))^2 \right], \quad (2)$$

where θ represents the trainable network parameters, p is the synthetic pathway, m is the product molecule, $S(m)$ is the objective score of molecule m , typically scaled between 0 and 1 (the higher, the better), σ is a constant scaling factor, $\text{PLL}_\theta(p, m)$ is the pseudo-log-likelihood of generating the pathway p leading to molecule m . The pseudo-log-likelihood, $\text{PLL}_\theta(p, m)$, is computed as the sum of three terms: the log-likelihood of selecting each token (reaction step or building block), the log-likelihood of selecting each reaction type, and the empirical log-likelihood value of the building block fingerprints:

$$\text{PLL}_\theta(p) = \sum_i \log P_\theta(x_i) + \log \hat{P}_\theta(x_{\text{building block}}), \quad (3)$$

where x_i represents each component of a synthetic path (token, reaction). The empirical likelihood of the building block fingerprints is approximated by:

$$\hat{P}_\theta(x_{\text{building block}}) = x_{\text{denoised}} \cdot x_{\text{target}} + (1 - x_{\text{denoised}}) \cdot (1 - x_{\text{target}}), \quad (4)$$

where x_{denoised} is the denoised fingerprint vector, x_{target} is the ground truth fingerprint.

Genetic algorithm with SynFormer-ED as a mutation operator

We integrated SynFormer-ED as a mutation step into an evolutionary framework for global chemical space exploration and optimization. The genetic algorithm follows a standard workflow: starting with a pool of candidates randomly sampled from the ZINC database, pairs of molecules are selected to undergo crossover, producing an offspring pool. Mutation is then applied to each offspring with a specific probability, followed by an additional mutation step where SynFormer-ED projects all offspring into synthesizable space. This step ensures that unsynthesizable fragments are corrected and adds additional diversity to the candidate pool. The offspring pool will be evaluated, and the highest-scored ones will form the next generation for the next iteration. The crossover and mutation rules were adopted from the GraphGA [29] algorithm; specifically, crossover involves graph matching and exchanging molecular halves, while mutation employs a set of hand-coded rules including both atom- and fragment-level modifications. We adopted the hyperparameters from [98].

Evaluation Details

This section describes the precise settings used in the evaluation of SynFormer-ED and SynFormer-D.

Reconstruction and local chemical space exploration

SynFormer-ED was employed for reconstruction, analog generation, and hit expansion. The workflows for these tasks are similar: SynFormer-ED encodes the input molecule and decodes it into synthetic pathways, using a search width of 24 and an exhaustiveness of 64 during decoding. For reconstruction and analog generation, Tanimoto (Jaccard) similarity [105] to the input molecule based on Morgan fingerprints [106] was evaluated for each generated molecule, with the most similar molecule selected as the output. In hit expansion, generated molecules were evaluated based on both the design objective score and Tanimoto similarity.

We tested the reconstruction performance on a random sample of 1,000 molecules from Enamine’s REAL Space [50] and the ChEMBL database [107]. Molecules were sampled from the REAL Diversity Set, which includes 48.2 million molecules and maximizes diversity to represent the broader REAL Space. These molecules comply with the Rule of 5 (Ro5) [108] and Veber criteria [109], making them suitable for drug-like properties. ChEMBL samples were taken from version 29 of the database, released in 2021.

The Synthetic Accessibility (SA) Score [42], a heuristic measure evaluating structural complexity based on fragment frequency in PubChem [110], was used to assess the synthesizability of the generated molecules. For synthesizable analog generation, 10 objective scores from the GuacaMol benchmark [76] were used, which involve optimizing various physicochemical properties, similarities or dissimilarities to known drugs, and specific structural motifs.

In hit expansion, we also used a predictive model [90] to assess inhibition against c-Jun N-terminal Kinase-3 (JNK3), a member of the mitogen-activated protein kinase family. This model is a random forest classifier utilizing ECFP6 fingerprints trained on the ExCAPE-DB dataset [89].

AutoDock Vina [84] was employed to evaluate binding affinities against protein targets in both structure-based drug design and hit expansion tasks. Protein structures and pocket information were sourced from the LIT-PCBA dataset [83], and an exhaustiveness setting of 16 was used for most calculations.

For baseline comparisons in analog finding, we conducted a nearest neighbor search within Enamine REAL. This was done by accessing the Enamine website (<https://new.enaminestore.com/draw-search>) and downloading approximately 100 nearest neighbors by adjusting the similarity threshold value.

Global chemical space exploration

For both SF-RL and GraphGA-SF, we employed these models to solve molecular optimization problems without predefined starting points [45]. In each experiment, we assumed the presence of an oracle, i.e., a function that evaluates the desired property of a molecule, providing the ground truth value. As property evaluation is often time- and resource-intensive in real-world scenarios, each experiment was limited to 10,000 oracle calls, meaning up to 10,000 molecules could be evaluated. During optimization, all evaluated molecules were recorded, and the primary performance metric was the area under the curve (AUC) of the top- K average property value versus the number of oracle calls (*AUC top- K*). The reported AUC values were min-max scaled to the range [0, 1].

We evaluated the global chemical space optimization capabilities of SF-RL and GraphGA-SF on optimizing DRD2 inhibition [111] and four multi-objective properties from the GuacaMol benchmark [76]: Sitagliptin MPO, Scaffold Hop, Perindopril MPO, and Ranolazine MPO. DRD2 inhibition was assessed using a classifier trained on the ExCAPE-DB dataset [89], employing a support vector machine (SVM) with a Gaussian kernel and ECFP6 fingerprints to distinguish

active from inactive compounds for the Dopamine Receptor D2 (DRD2) [111]. The multi-objective tasks from the GuacaMol benchmark are described above [76]. All oracle functions were accessed through the Therapeutic Data Commons (TDC) [112, 113].

The baseline methods compared in this study include:

- **GraphGA** [29] is a popular heuristic algorithm inspired by natural evolutionary processes, featuring crossover rules derived from graph matching and mutations applied at both the atom- and fragment-level.
- **Graph MCTS** [29] uses Monte Carlo Tree Search to explore molecular structures by locally searching each branch of the current state (molecule or partial molecule) and selecting the most promising candidates based on property scores for subsequent iterations.
- **REINVENT** [19] utilizes a policy-based reinforcement learning (RL) approach where agents take actions in an environment to maximize cumulative rewards, tuning recurrent neural networks (RNNs) to generate SMILES strings.
- **Augmented Memory (AugMem)** [97] integrates SMILES-based reinforcement learning with augmented training through experience replay and selective memory purge to prevent model collapse. This approach achieves state-of-the-art sample efficiency, demonstrating the highest reported performance in the PMO benchmark.
- **GFlowNet** [96] is a generative AI model that treats the generative process as a flow network and trains it with a temporal difference-like loss function, aligning the generation process with the target property distribution by matching the property of interest to the flow volume.
- **DoG-Gen** [64] models synthetic pathways as Directed Acyclic Graphs (DAGs) and uses an RNN-based generator, optimizing via an iterative learning method that incorporates high-scoring molecules into the training data for successive fine-tuning of the generative model.
- **SynNet** [65] is a synthesis-based method that applies a genetic algorithm to molecular fingerprints, decoding them into synthetic pathways using trained neural networks.

The implementations of these methods were adapted from their original publications with minor modifications, and hyperparameters were set according to the specifications in [45], except AugMem which is from the original publication.

Acknowledgments

This research was supported by the Office of Naval Research under grant number N00014-21-1-2195 and the AI2050 program at Schmidt Futures under grant number G-22-64475. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research. W.G. received additional funding from the Google Ph.D. fellowship. We thank Kaiming He and Yurii Moroz for their helpful discussions and comments on the manuscript.

Reproducibility Statement

All code and releasable data can be found at <https://github.com/wenhao-gao/synformer>, including instructions in a README file. The model weights can be accessed on <https://huggingface.co/whgao/synformer>. Appendix A provides more details about the experimental setup and implementation details.

References

- [1] W Patrick Walters, Jeremy Green, Jonathan R Weiss, and Mark A Murcko. What do medicinal chemists actually make? a 50-year retrospective. *Journal of medicinal chemistry*, 54(19):6405–6416, 2011.
- [2] Peng Wu, Thomas E Nielsen, and Mads H Clausen. Small-molecule kinase inhibitors: an analysis of fda-approved drugs. *Drug discovery today*, 21(1):5–10, 2016.
- [3] Jiankun Lyu, Sheng Wang, Trent E Balius, Isha Singh, Anat Levit, Yurii S Moroz, Matthew J O’Meara, Tao Che, Enkhjargal Alгаа, Kateryna Tolmachova, et al. Ultra-large library docking for discovering new chemotypes. *Nature*, 566(7743):224–229, 2019.
- [4] Zhiao Yu, Hansen Wang, Xian Kong, William Huang, Yuchi Tsao, David G Mackanic, Kecheng Wang, Xinchang Wang, Wenxiao Huang, Snehashis Choudhury, et al. Molecular design for electrolyte solvents enabling energy-dense and long-cycling lithium metal batteries. *Nature Energy*, 5(7):526–533, 2020.

- [5] Johannes Hachmann, Roberto Olivares-Amaya, Sule Atahan-Evrenk, Carlos Amador-Bedolla, Roel S Sánchez-Carrera, Aryeh Gold-Parker, Leslie Vogt, Anna M Brockway, and Alán Aspuru-Guzik. The harvard clean energy project: large-scale computational screening and design of organic photovoltaics on the world community grid. *The Journal of Physical Chemistry Letters*, 2(17):2241–2251, 2011.
- [6] Pawankumar Rai, Srishti Mehrotra, Smriti Priya, Edgard Gnansounou, and Sandeep K Sharma. Recent advances in the sustainable design and applications of biodegradable polymers. *Bioresource technology*, 325:124739, 2021.
- [7] Kyle M Diederichsen, Rezvan Sharifian, Jin Soo Kang, Yayuan Liu, Seoni Kim, Betar M Gallant, David Vermaas, and T Alan Hatton. Electrochemical methods for carbon dioxide separations. *Nature Reviews Methods Primers*, 2(1):68, 2022.
- [8] Jiayu Peng, Daniel Schwalbe-Koda, Karthik Akkiraju, Tian Xie, Livia Giordano, Yang Yu, C John Eom, Jaclyn R Lunger, Daniel J Zheng, Reshma R Rao, et al. Human-and machine-centred designs of molecules and materials for sustainability and decarbonization. *Nature Reviews Materials*, 7(12):991–1009, 2022.
- [9] Edward O Pyzer-Knapp, Changwon Suh, Rafael Gómez-Bombarelli, Jorge Aguilera-Iparraguirre, and Alán Aspuru-Guzik. What is high-throughput virtual screening? a perspective from organic materials discovery. *Annual Review of Materials Research*, 45(1):195–216, 2015.
- [10] Katarzyna Smietana, Marcin Siatkowski, and Martin Møller. Trends in clinical success rates. *Nat Rev Drug Discov*, 15(6):379–80, 2016.
- [11] Sudeep Pushpakom, Francesco Iorio, Patrick A Eyers, K Jane Escott, Shirley Hopper, Andrew Wells, Andrew Doig, Tim Guilliams, Joanna Latimer, Christine McNamee, et al. Drug repurposing: progress, challenges and recommendations. *Nature reviews Drug discovery*, 18(1):41–58, 2019.
- [12] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [13] Yixin Wang and David Blei. Variational bayes under model misspecification. *Advances in Neural Information Processing Systems*, 32, 2019.
- [14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [16] Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.
- [17] 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.
- [18] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pages 2323–2332. PMLR, 2018.
- [19] Thomas Blaschke, Josep Arús-Pous, Hongming Chen, Christian Margreitter, Christian Tyrchan, Ola Engkvist, Kostas Papadopoulos, and Atanas Patronov. Reinvent 2.0: an ai tool for de novo drug design. *Journal of chemical information and modeling*, 60(12):5918–5922, 2020.
- [20] Tianfan Fu, Wenhao Gao, Cao Xiao, Jacob Yasonik, Connor W Coley, and Jimeng Sun. Differentiable scaffolding tree for molecular optimization. *International Conference on Learning Representations*, 2021.
- [21] Joshua Meyers, Benedek Fabian, and Nathan Brown. De novo molecular design and generative models. *Drug discovery today*, 26(11):2707–2715, 2021.
- [22] Dylan M Anstine and Olexandr Isayev. Generative models as an emerging paradigm in the chemical sciences. *Journal of the American Chemical Society*, 145(16):8736–8750, 2023.
- [23] Akshay Subramanian, Wenhao Gao, Regina Barzilay, Jeffrey C Grossman, Tommi Jaakkola, Stefanie Jegelka, Mingda Li, Ju Li, Wojciech Matusik, Elsa Olivetti, et al. Closing the execution gap in generative ai for chemicals and materials: Freeways or safeguards. *An MIT Exploration of Generative AI*, 2024.
- [24] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. *Conference on Neural Information Processing Systems*, 2017.
- [25] W Patrick Walters, Matthew T Stahl, and Mark A Murcko. Virtual screening—an overview. *Drug discovery today*, 3(4):160–178, 1998.

- [26] Brian K Shoichet. Virtual screening of chemical libraries. *Nature*, 432(7019):862–865, 2004.
- [27] W Patrick Walters. Virtual chemical libraries: miniperspective. *Journal of medicinal chemistry*, 62(3):1116–1124, 2018.
- [28] Venkat Venkatasubramanian, King Chan, and James M Caruthers. Evolutionary design of molecules with desired properties using the genetic algorithm. *Journal of Chemical Information and Computer Sciences*, 35(2):188–195, 1995.
- [29] Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.
- [30] Wenhao Gao and Connor W Coley. The synthesizability of molecules proposed by generative models. *Journal of chemical information and modeling*, 60(12):5714–5723, 2020.
- [31] Philipp Renz, Dries Van Rompaey, Jörg Kurt Wegner, Sepp Hochreiter, and Günter Klambauer. On failure modes in molecule generation and optimization. *Drug Discovery Today: Technologies*, 32:55–63, 2019.
- [32] W Patrick Walters and Regina Barzilay. Critical assessment of ai in drug discovery. *Expert opinion on drug discovery*, 16(9):937–947, 2021.
- [33] Megan Stanley and Marwin Segler. Fake it until you make it? generative de novo design and virtual screening of synthesizable molecules. *Current Opinion in Structural Biology*, 82:102658, 2023.
- [34] Christos A Nicolaou, Christine Humblet, Hong Hu, Eva M Martin, Frank C Dorsey, Thomas M Castle, Keith Ian Burton, Haitao Hu, Jorg Hendle, Michael J Hickey, et al. Idea2data: toward a new paradigm for drug discovery. *ACS medicinal chemistry letters*, 10(3):278–286, 2019.
- [35] Cara E Brocklehurst, Eva Altmann, Corentin Bon, Holly Davis, David Dunstan, Peter Ertl, Carol Ginsburg-Moraff, Jonathan Grob, Daniel J Gosling, Guillaume Lapointe, et al. Microcycle: An integrated and automated platform to accelerate drug discovery. *Journal of medicinal chemistry*, 67(3):2118–2128, 2024.
- [36] Gisbert Schneider. Automating drug discovery. *Nature reviews drug discovery*, 17(2):97–113, 2018.
- [37] Connor W Coley, Natalie S Eyke, and Klavs F Jensen. Autonomous discovery in the chemical sciences part i: Progress. *Angewandte Chemie International Edition*, 59(51):22858–22893, 2020.
- [38] Connor W Coley, Natalie S Eyke, and Klavs F Jensen. Autonomous discovery in the chemical sciences part ii: outlook. *Angewandte Chemie International Edition*, 59(52):23414–23436, 2020.
- [39] Brent A Koscher, Richard B Canty, Matthew A McDonald, Kevin P Greenman, Charles J McGill, Camille L Bilodeau, Wengong Jin, Haoyang Wu, Florence H Vermeire, Brooke Jin, et al. Autonomous, multiproperty-driven molecular discovery: From predictions to measurements and back. *Science*, 382(6677):ead1407, 2023.
- [40] Felix Strieth-Kalthoff, Han Hao, Vandana Rathore, Joshua Derasp, Théophile Gaudin, Nicholas H Angello, Martin Seifrid, Ekaterina Trushina, Mason Guy, Junliang Liu, et al. Delocalized, asynchronous, closed-loop discovery of organic laser emitters. *Science*, 384(6697):eadk9227, 2024.
- [41] Julius Seumer, Jonathan Kirschner Solberg Hansen, Mogens Brøndsted Nielsen, and Jan H Jensen. Computational evolution of new catalysts for the morita–baylis–hillman reaction. *Angewandte Chemie International Edition*, 62(18):e202218565, 2023.
- [42] 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.
- [43] Amol Thakkar, Veronika Chadimová, Esben Jannik Bjerrum, Ola Engkvist, and Jean-Louis Reymond. Retrosynthetic accessibility score (rascore)—rapid machine learned synthesizability classification from ai driven retrosynthetic planning. *Chemical science*, 12(9):3339–3349, 2021.
- [44] Cheng-Hao Liu, Maksym Korablyov, Stanisław Jastrzebski, Paweł Włodarczyk-Pruszyński, 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.
- [45] Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a benchmark for practical molecular optimization. *Advances in neural information processing systems*, 35:21342–21357, 2022.
- [46] Jeff Guo and Philippe Schwaller. Saturn: Sample-efficient generative molecular design using memory manipulation. *arXiv preprint arXiv:2405.17066*, 2024.
- [47] Jeff Guo and Philippe Schwaller. Directly optimizing for synthesizability in generative molecular design using retrosynthesis models. *arXiv preprint arXiv:2407.12186*, 2024.
- [48] Torsten Hoffmann and Marcus Gastreich. The next level in chemical space navigation: going far beyond enumerable compound libraries. *Drug discovery today*, 24(5):1148–1156, 2019.

- [49] Hitesh Patel, Wolf-Dietrich Ihlenfeldt, Philip N Judson, Yurii S Moroz, Yuri Pevzner, Megan L Peach, Victorien Delannée, Nadya I Tarasova, and Marc C Nicklaus. Savi, in silico generation of billions of easily synthesizable compounds through expert-system type rules. *Scientific data*, 7(1):384, 2020.
- [50] Oleksandr O Grygorenko, Dmytro S Radchenko, Igor Dziuba, Alexander Chuprina, Kateryna E Gubina, and Yurii S Moroz. Generating multibillion chemical space of readily accessible screening compounds. *Isience*, 23(11), 2020.
- [51] WuXi AppTec. Virtual screening - hit finding and screening services, September 2024. Accessed: 2022-12-08.
- [52] Alexander Neumann, Lester Marrison, and Raphael Klein. Relevance of the trillion-sized chemical space “explore” as a source for drug discovery. *ACS Medicinal Chemistry Letters*, 14(4):466–472, 2023.
- [53] Corentin Bedart, Grace Shimokura, Frederick G West, Tabitha E Wood, Robert A Batey, John J Irwin, and Matthieu Schapira. The pan-canadian chemical library: A mechanism to open academic chemistry to high-throughput virtual screening. *Scientific Data*, 11(1):597, 2024.
- [54] H Maarten Vinkers, Marc R de Jonge, Frederik FD Daeyaert, Jan Heeres, Lucien MH Koymans, Joop H van Lenthe, Paul J Lewi, Henk Timmerman, Koen Van Aken, and Paul AJ Janssen. Synopsis: synthesize and optimize system in silico. *Journal of medicinal chemistry*, 46(13):2765–2773, 2003.
- [55] Markus Hartenfeller, Heiko Zettl, Miriam Walter, Matthias Rupp, Felix Reisen, Ewgenij Proschak, Sascha Weggen, Holger Stark, and Gisbert Schneider. Dogs: reaction-driven de novo design of bioactive compounds. *PLoS computational biology*, 8(2):e1002380, 2012.
- [56] Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos, Jeff Schneider, and Eric Xing. Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations. In *International Conference on Artificial Intelligence and Statistics*, pages 3393–3403. PMLR, 2020.
- [57] Dai Hai Nguyen and Koji Tsuda. A generative model for molecule generation based on chemical reaction trees. *arXiv preprint arXiv:2106.03394*, 2021.
- [58] Arman A Sadybekov, Anastasiia V Sadybekov, Yongfeng Liu, Christos Iliopoulos-Tsoutsouvas, Xi-Ping Huang, Julie Pickett, Blake Houser, Nilkanth Patel, Ngan K Tran, Fei Tong, et al. Synthon-based ligand discovery in virtual libraries of over 11 billion compounds. *Nature*, 601(7893):452–459, 2022.
- [59] Kyle Swanson, Gary Liu, Denise B Catacutan, Autumn Arnold, James Zou, and Jonathan M Stokes. Generative ai for designing and validating easily synthesizable and structurally novel antibiotics. *Nature Machine Intelligence*, 6(3):338–353, 2024.
- [60] Michael Sun, Alston Lo, Wenhao Gao, Minghao Guo, Veronika Thost, Jie Chen, Connor Coley, and Wojciech Matusik. Syntax-guided procedural synthesis of molecules. *arXiv preprint arXiv:2409.05873*, 2024.
- [61] 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.
- [62] 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*, pages 3668–3679. PMLR, 2020.
- [63] Julien Horwood and Emmanuel Noutahi. Molecular design in synthetically accessible chemical space via deep reinforcement learning. *ACS omega*, 5(51):32984–32994, 2020.
- [64] 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.
- [65] Wenhao Gao, Rocío Mercado, and Connor W Coley. Amortized tree generation for bottom-up synthesis planning and synthesizable molecular design. *International Conference on Learning Representations*, 2021.
- [66] Michał Koziarski, Andrei Rekish, 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. *arXiv preprint arXiv:2406.08506*, 2024.
- [67] Shitong Luo, Wenhao Gao, Zuofan Wu, Jian Peng, Connor W Coley, and Jianzhu Ma. Projecting molecules into synthesizable chemical spaces. *International Conference on Machine Learning*, 2024.
- [68] David Mendez, Anna Gaulton, A Patrícia Bento, Jon Chambers, Marleen De Veij, Eloy Félix, María Paula Magariños, Juan F Mosquera, Prudence Mutowo, Michał Nowotka, et al. ChEMBL: towards direct deposition of bioassay data. *Nucleic acids research*, 47(D1):D930–D940, 2019.

- [69] Harry L Morgan. The generation of a unique machine description for chemical structures—a technique developed at chemical abstracts service. *Journal of chemical documentation*, 5(2):107–113, 1965.
- [70] Emiel Hoogetboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34:12454–12465, 2021.
- [71] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- [72] Jie Zhang, Rocío Mercado, Ola Engkvist, and Hongming Chen. Comparative study of deep generative models on chemical space coverage. *Journal of Chemical Information and Modeling*, 61(6):2572–2581, 2021.
- [73] Nathan C Frey, Ryan Soklaski, Simon Axelrod, Siddharth Samsi, Rafael Gomez-Bombarelli, Connor W Coley, and Vijay Gadepally. Neural scaling of deep chemical models. *Nature Machine Intelligence*, 5(11):1297–1305, 2023.
- [74] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [75] ASKCOS Project. Askcos: Ai-driven synthesis planning. <https://askcos.mit.edu/>. Accessed: 2024-09-16.
- [76] 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.
- [77] Valerie J Gillet, William Newell, Paulina Mata, Glenn Myatt, Sandor Sike, Zsolt Zsoldos, and A Peter Johnson. Sprout: recent developments in the de novo design of molecules. *Journal of chemical information and computer sciences*, 34(1):207–217, 1994.
- [78] Regine S Bohacek, Colin McMartin, and Wayne C Guida. The art and practice of structure-based drug design: a molecular modeling perspective. *Medicinal research reviews*, 16(1):3–50, 1996.
- [79] Renxiao Wang, Ying Gao, and Luhua Lai. Ligbuilder: a multi-purpose program for structure-based drug design. *Molecular modeling annual*, 6:498–516, 2000.
- [80] Valerie J Gillet, Glenn Myatt, Zsolt Zsoldos, and A Peter Johnson. Sprout, hippo and caesa: Tools for de novo structure generation and estimation of synthetic accessibility. *Perspectives in Drug Discovery and Design*, 3:34–50, 1995.
- [81] Charles Harris, Kieran Didi, Arian R Jamasb, Chaitanya K Joshi, Simon V Mathis, Pietro Lio, and Tom Blundell. Benchmarking generated poses: How rational is structure-based drug design with generative models? *arXiv preprint arXiv:2308.07413*, 2023.
- [82] Shitong Luo, Jiaqi Guan, Jianzhu Ma, and Jian Peng. A 3d generative model for structure-based drug design. *Advances in Neural Information Processing Systems*, 34:6229–6239, 2021.
- [83] Viet-Khoa Tran-Nguyen, Célien Jacquemard, and Didier Rognan. Lit-pcba: an unbiased data set for machine learning and virtual screening. *Journal of chemical information and modeling*, 60(9):4263–4273, 2020.
- [84] Oleg Trott and Arthur J Olson. Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of computational chemistry*, 31(2):455–461, 2010.
- [85] Itai Levin, Michael E Fortunato, Kian L Tan, and Connor W Coley. Computer-aided evaluation and exploration of chemical spaces constrained by reaction pathways. *AIChE journal*, 69(12):e18234, 2023.
- [86] Matt J Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder. In *International conference on machine learning*, pages 1945–1954. PMLR, 2017.
- [87] Teague Sterling and John J Irwin. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337, 2015.
- [88] Chia-Yi Kuan, Alan J Whitmarsh, Derek D Yang, Guanghong Liao, Aryn J Schloemer, Chen Dong, Jue Bao, Kenneth J Banasiak, Gabriel G Haddad, Richard A Flavell, et al. A critical role of neural-specific jnk3 for ischemic apoptosis. *Proceedings of the National Academy of Sciences*, 100(25):15184–15189, 2003.
- [89] Jiangming Sun, Nina Jeliaskova, Vladimir Chupakhin, Jose-Felipe Golib-Dzib, Ola Engkvist, Lars Carlsson, Jörg Wegner, Hugo Ceulemans, Ivan Georgiev, Vedrin Jeliaskov, et al. Escape-db: an integrated large scale dataset facilitating big data analysis in chemogenomics. *Journal of cheminformatics*, 9:1–9, 2017.
- [90] Yibo Li, Liangren Zhang, and Zhenming Liu. Multi-objective de novo drug design with conditional graph generative model. *Journal of cheminformatics*, 10:1–24, 2018.

- [91] Dimitrios Anastasiou, Yimin Yu, William J Israelsen, Jian-Kang Jiang, Matthew B Boxer, Bum Soo Hong, Wolfram Tempel, Svetoslav Dimov, Min Shen, Abhishek Jha, et al. Pyruvate kinase m2 activators promote tetramer formation and suppress tumorigenesis. *Nature chemical biology*, 8(10):839–847, 2012.
- [92] Philip G Humphreys, Paul Bamborough, Chun-wa Chung, Peter D Craggs, Laurie Gordon, Paola Grandi, Thomas G Hayhow, Jameed Hussain, Katherine L Jones, Matthew Lindon, et al. Discovery of a potent, cell penetrant, and selective p300/cbp-associated factor (pcaf)/general control nonderepressible 5 (gcn5) bromodomain chemical probe. *Journal of medicinal chemistry*, 60(2):695–709, 2017.
- [93] Uschi Dolfus, Hans Briem, and Matthias Rarey. Synthesis-aware generation of structural analogues. *Journal of Chemical Information and Modeling*, 62(15):3565–3576, 2022.
- [94] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992.
- [95] Lucia Vallar and Jacopo Meldolesi. Mechanisms of signal transduction at the dopamine d2 receptor. *Trends in pharmacological sciences*, 10(2):74–77, 1989.
- [96] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394, 2021.
- [97] Jeff Guo and Philippe Schwaller. Augmented memory: Sample-efficient generative molecular design with reinforcement learning. *Jacs Au*, 2024.
- [98] Austin Tripp and José Miguel Hernández-Lobato. Genetic algorithms are strong baselines for molecule generation. *arXiv preprint arXiv:2310.09267*, 2023.
- [99] Michael W Mullowney, Katherine R Duncan, Somayah S Elsayed, Neha Garg, Justin JJ van der Hoof, Nathaniel I Martin, David Meijer, Barbara R Terlouw, Friederike Biermann, Kai Blin, et al. Artificial intelligence for natural product drug discovery. *Nature Reviews Drug Discovery*, 22(11):895–916, 2023.
- [100] Inc. Daylight Chemical Information Systems. Daylight theory: Smarts. <https://www.daylight.com/dayhtml/doc/theory/theory.smarts.html>, 2024. Accessed: 2024-09-05.
- [101] Markus Hartenfeller, Martin Eberle, Peter Meier, Cristina Nieto-Oberhuber, Karl-Heinz Altmann, Gisbert Schneider, Edgar Jacoby, and Steffen Renner. A collection of robust organic synthesis reactions for in silico molecule design. *Journal of chemical information and modeling*, 51(12):3093–3098, 2011.
- [102] Alexander Button, Daniel Merk, Jan A Hiss, and Gisbert Schneider. Automated de novo molecular design by hybrid machine intelligence and rule-driven chemical synthesis. *Nature machine intelligence*, 1(7):307–315, 2019.
- [103] Georg Wittig and Georg Geissler. Zur reaktionsweise des pentaphenyl-phosphors und einiger derivate. *Justus Liebigs Annalen der Chemie*, 580(1):44–57, 1953.
- [104] John F Hartwig. Transition metal catalyzed synthesis of arylamines and aryl ethers from aryl halides and triflates: Scope and mechanism. *Angewandte Chemie International Edition*, 37(15):2046–2067, 1998.
- [105] Paul Jaccard. Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bull Soc Vaudoise Sci Nat*, 37:547–579, 1901.
- [106] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- [107] Barbara Zdrazil, Eloy Felix, Fiona Hunter, Emma J Manners, James Blackshaw, Sybilla Corbett, Marleen de Veij, Harris Ioannidis, David Mendez Lopez, Juan F Mosquera, et al. The chembl database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. *Nucleic acids research*, 52(D1):D1180–D1192, 2024.
- [108] Christopher A Lipinski, Franco Lombardo, Beryl W Dominy, and Paul J Feeney. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced drug delivery reviews*, 23(1-3):3–25, 1997.
- [109] Daniel F Veber, Stephen R Johnson, Hung-Yuan Cheng, Brian R Smith, Keith W Ward, and Kenneth D Kopple. Molecular properties that influence the oral bioavailability of drug candidates. *Journal of medicinal chemistry*, 45(12):2615–2623, 2002.
- [110] Sunghwan Kim, Paul A Thiessen, Evan E Bolton, Jie Chen, Gang Fu, Asta Gindulyte, Liany Han, Jane He, Siqian He, Benjamin A Shoemaker, et al. Pubchem substance and compound databases. *Nucleic acids research*, 44(D1):D1202–D1213, 2016.

- [111] Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9:1–14, 2017.
- [112] 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. *Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021.
- [113] Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Artificial intelligence foundation for therapeutic science. *Nature chemical biology*, 18(10):1033–1036, 2022.
- [114] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021.
- [115] Greg Landrum. Colliding bits. <https://rdkit.blogspot.com/2014/02/colliding-bits.html>, 2014. Accessed: 2024-09-24.

A Additional method details

A.1 Model architectures

A.1.1 Transformer decoder

SynFormer adopts the standard Transformer decoder architecture [14] to generate postfix notation of synthesis in an auto-regressive manner. A postfix notation of synthesis may contain four types of tokens: a [START] token, [END] token, reaction token, and building block token. Each token, regardless of its type, is represented as a 768-dimensional embedding vector that is differentiable during training. Specifically, the embedding vectors of start token, end token, and reaction tokens are stored in a look-up table implemented by PyTorch’s `nn.Embedding` layer. Note that, as we use a fixed set of reactions consisting of only 169 reaction templates, it is practical to store the embedding vectors for each of them. However, the building block database, containing more than 200k molecules, is too large to assign embedding vectors to each building block separately. Thus, we use a multi-layer perceptron (MLP), to convert 2048-bit Morgan fingerprints of building blocks into 768-dimensional continuous vectors. This MLP consists of three layers, with hidden dimensions of 1,536 and ReLU activation.

To inform the model with the location of each token, we add sinusoidal positional encoding [14] to embedding vectors:

$$\mathbf{p}_i = \begin{cases} \sin\left(\frac{j}{10000^{\frac{2i}{d}}}\right), & i = 0, 2, \dots, 766 \\ \cos\left(\frac{j}{10000^{\frac{2i}{d}}}\right), & i = 1, 3, \dots, 767 \end{cases}, \quad (5)$$

where i denotes the dimension index, $d = 768$ denotes the embedding dimension, and j denotes the token position within the sequence.

The Transformer decoder network operates on the token embedding vectors. It consists of 10 attention layers with residual connection and layer normalization in between. In SynFormer-ED, the attention layer contains both self-attention and cross-attention that gathers information from the encoder, while in SynFormer-D, only self-attention is used. Each attention layer contains 16 attention heads. Each head projects the feature vectors into 256-dimensional query, key, and value vectors, which are then transformed into output vectors via the attention mechanism. Output vectors from different attention heads are concatenated into a 2,048-dimensional vector, which is then mapped into a 768-dimensional vector used to update the token embedding via residual connection and layer normalization.

A.1.2 Transformer encoder

SynFormer-ED contains a Transformer-based encoder network that conditions the decoder with a given SMILES string. The input SMILES string is split in to a list of tokens according to the following vocabulary:

H, He, Li, Be, B, C, N, O, F, Ne, Na, Mg, Al, Si, P, S, Cl, Ar, K, Ca, Sc, Ti, V, Cr, Mn, Fe, Co, Ni, Cu, Zn, Ga, Ge, As, Se, Br, Kr, Rb, Sr, Y, Zr, Nb, Mo, Tc, Ru, Rh, Pd, Ag, Cd, In, Sn, Sb, Te, I, Xe, Cs, Ba, La, Ce, Pr, Nd, Pm, Sm, Eu, Gd, Tb, Dy, Ho, Er, Tm, Yb, Lu, Hf, Ta, W, Re, Os, Ir, Pt, Au, Hg, Tl, Pb, Bi, Po, At, Rn, Fr, Ra, Ac, Th, Pa, U, Np, Pu, Am, Cm, Bk, Cf, Es, Fm, Md, No, Lr, Rf, Db, Sg, Bh, Hs, Mt, Ds, Rg, Cn, Nh, Fl, Mc, Lv, Ts, Og, b, c, n, o, s, p, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, [,], (,), ., =, #, -, +, \, /, :, , @, ?, >, *, \$, %

Each token in the list is represented by a 768-dimensional embedding vector, and position encodings (Eq. 5) are added to the embedding vector. The Transformer encoder contains 6 layers, each of which has 16 attention heads with 256-dimensional query, key, and value vectors.

A.1.3 Scaling of model size

Table 1 below shows the hyper-parameters for SynFormer-ED of different sizes measured by the number of trainable parameters. These hyper-parameter sets are used to investigate the scaling of model performance. For the remaining experiments, we use the model with 229M parameters.

A.1.4 Diffusion fingerprint head

The fingerprint of a building block token are modeled by an n -dimensional Bernoulli distribution. We use the denoising diffusion probabilistic model [15, 114] to model this distribution. The denoising diffusion probabilistic model generates data by iterative denoising an initial sample drawn from a simple prior distribution. To train a neural network capable of denoising, noisy samples based on ground truth data must be constructed via the forward diffusion process.

Table 1: Specific model hyper-parameters used in the scaling experiments.

Model Size	229M	68M	17.5M	3.7M
# Feature Dimensions	768	512	256	128
# Encoder Layers	8	6	6	4
# Encoder Attn. Heads	16	16	8	4
# Encoder Feed-forward Dimensions	4096	2048	1024	512
# Decoder Layers	10	8	8	6
# Decoder Attn. Heads	16	16	8	4
# Decoder Feed-forward Dimensions	4096	2048	1024	512

In SynFormer, we define the bit-wise forward diffusion process for fingerprint $\mathbf{x} = [x^{(1)}, \dots, x^{(n)}] \in \{0, 1\}^n$ as, following the formulation of multinomial diffusion [70]:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \prod_{j=1}^n q(x_t^{(j)} | x_{t-1}^{(j)}) \quad \text{where} \quad q(x_t^{(j)} | x_{t-1}^{(j)}) = \text{Bernoulli}(x_t^{(j)} | (1 - \beta_t)x_{t-1}^{(j)} + \beta_t \cdot \frac{1}{2}), \quad j = 1 \dots n. \quad (6)$$

Since the forward diffusion process is Markovian, it is possible to express the bit-wise distribution of \mathbf{x}_t directly with the ground truth fingerprint \mathbf{x}_0 :

$$q(\mathbf{x}_t | \mathbf{x}_0) = \prod_{j=1}^n q(x_t^{(j)} | x_0^{(j)}) \quad \text{where} \quad q(x_t^{(j)} | x_0^{(j)}) = \text{Bernoulli}(x_t^{(j)} | \bar{\alpha}_t x_0^{(j)} + (1 - \bar{\alpha}_t) \cdot \frac{1}{2}), \quad j = 1 \dots n. \quad (7)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{\tau=1}^t \alpha_\tau$. This expression enables efficient one-step generation of \mathbf{x}_t given \mathbf{x}_0 at training time, while it requires t steps if we follow Eq. 6. Intuitively, $1 - \bar{\alpha}_t$ determines the chance of the fingerprint bit being resampled uniformly at timestep t , and it should increase monotonically from $t = 0$ to T . We adopt the noise schedule in terms of $\bar{\alpha}_t$ proposed by Nichol and Dhariwal [114]:

$$\bar{\alpha}_t = \frac{f(t)}{f(0)}, \quad f(t) = \cos\left(\frac{\frac{t}{T} + s}{1 + s} \cdot \frac{\pi}{2}\right)^2. \quad (8)$$

From this definition, β_t and α_t can be obtained by $\beta_t = 1 - \frac{\bar{\alpha}_t}{\alpha_{t-1}}$ and $\alpha_t = 1 - \beta_t$ respectively. In practice, β_t is clipped to be no larger than 0.999 to prevent instability.

To denoise a noisy fingerprint \mathbf{x}_t into the less noisy \mathbf{x}_{t-1} which eventually leads to \mathbf{x}_0 , the denoiser network should learn the following posterior distribution derived from Eq.6 and Eq.7:

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \prod_{j=1}^n q(x_{t-1}^{(j)} | x_t^{(j)}, x_0^{(j)}) \quad \text{where} \quad q(x_{t-1}^{(j)} | x_t^{(j)}, x_0^{(j)}) = \text{Bernoulli}\left(x_{t-1}^{(j)} \left| \frac{\theta_1(x_t^{(j)}, x_0^{(j)})}{\theta_1(x_t^{(j)}, x_0^{(j)}) + \theta_0(x_t^{(j)}, x_0^{(j)})} \right.\right), \quad (9)$$

where

$$\theta_1(x_t^{(j)}, x_0^{(j)}) = \left(\alpha_t x_t^{(j)} + (1 - \alpha_t) \cdot \frac{1}{2}\right) \cdot \left(\bar{\alpha}_{t-1} x_0^{(j)} + (1 - \bar{\alpha}_{t-1}) \cdot \frac{1}{2}\right) \quad (10)$$

$$\theta_0(x_t^{(j)}, x_0^{(j)}) = \left(\alpha_t (1 - x_t^{(j)}) + (1 - \alpha_t) \cdot \frac{1}{2}\right) \cdot \left(\bar{\alpha}_{t-1} (1 - x_0^{(j)}) + (1 - \bar{\alpha}_{t-1}) \cdot \frac{1}{2}\right). \quad (11)$$

In this posterior distribution, \mathbf{x}_0 is the unknown variable to be predicted, which we parameterize with a neural network conditional on both the current noisy fingerprint \mathbf{x}_t and the embedding vector from the Transformer decoder \mathbf{h} , leading to the following reverse diffusion process:

$$p(\mathbf{x}_{t-1} | \mathbf{x}_t) = \prod_{j=1}^n \text{Bernoulli}\left(x_{t-1}^{(j)} \left| \frac{\theta_1(x_t^{(j)}, p_0^{(j)})}{\theta_1(x_t^{(j)}, p_0^{(j)}) + \theta_0(x_t^{(j)}, p_0^{(j)})} \right.\right) \quad \text{where} \quad p_0^{(j)} = \text{sigmoid}(\text{DenoiserNet}(\mathbf{x}_t, \mathbf{h})[j]). \quad (12)$$

At training time, we randomly sample a step index t from $\mathcal{U}\{1, T\}$, perturb the ground truth fingerprint according to Eq. 7, and finally compute and back-propagate through the following training loss to align the predicted distribution to the ground truth posterior:

$$L_{\text{fingerprint}} = \mathbb{E}_{t \sim \mathcal{U}\{1, T\}, q(\mathbf{x}_t | \mathbf{x}_0)} [\text{KL}(q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0), p(\mathbf{x}_{t-1} | \mathbf{x}_t))]. \quad (13)$$

During inference, we start with a random bit vector \mathbf{x}_T sampled from uniform Bernoulli distribution, and resample it following $p(\mathbf{x}_{t-1}|\mathbf{x}_t)$ for T steps.

SynFormer’s diffusion module has $T = 100$ timesteps and the hyper-parameter s controlling the noise schedule is set to 0.01. The DenoiserNet is a multi-layer perceptron, with input dimension 2,816 ($\dim(\mathbf{h}) + \dim(\mathbf{x})$), hidden dimension 4096, output dimension 2,048 (fingerprint size), and ReLU activation in between.

A.2 Training details

A.2.1 Training data generation

We generate random postfix notations of synthesis using the building blocks and reaction templates from the database to train SynFormer. To ensure the validity of generated postfix notations, we leverage the *synthesis stack* data structure introduced in our previous model [67] to store intermediate reaction products during synthesis. When training SynFormer-ED, the generated postfix notation is used to supervised the decoder part and the SMILES string of the final product is used as the input to the encoder part. For SynFormer-D, only postfix notations are used to train the model. The detail of this data generation procedure is described in Algorithm 1.

Algorithm 1: Training data generation [67]

Input: m_r , maximum number of reactions

Input: m_a , maximum number of atoms in the final product molecule

Input: $\mathcal{F} : M \mapsto \{R\}$, a function that returns available reactions for molecule M

Input: $\mathcal{H} : R \mapsto \{M\}$, a hash map that stores eligible molecules for each reaction

Output: an iterator of postfix notation and product molecule tuples $(P, S.\text{top}())$

```
1  $S \leftarrow []$ ; // Initialize empty stack
2  $P \leftarrow []$ ; // Initialize empty postfix notation
3  $B_0 \leftarrow$  a random building block;
4 Push( $B_0$  onto  $S$ );
5 Append( $B_0$  to  $P$ );
6 Yield( $P, S.\text{top}()$ );
7 while  $P.\text{countReactions}() < m_r$  and  $S.\text{top}().\text{countAtoms}() < m_a$  do
8    $R \sim \mathcal{R}(S.\text{top}())$ ; // Randomly choose a reaction according to the molecule on the stack
9   for  $i = 1$  to  $R.\text{numReactants} - S.\text{length}()$  do
10     $B_i \sim \mathcal{H}(R)$  subject to  $B_i$  and the molecules in  $S$  can react via  $R$ ;
11    Push( $B_i$  onto  $S$ );
12    Append( $B_i$  to  $P$ );
13   $\{X_i\} \leftarrow$  Pop( $R.\text{numReactants}$  elements from  $S$ );
14   $Y \leftarrow$  Call  $R(\{X_i\})$ ;
15  Push( $Y$  onto  $S$ );
16  Append( $R$  to  $P$ );
17  if  $S.\text{length}() = 1$  then
18    Yield( $P, S.\text{top}()$ );
```

A.2.2 Hyper-parameters

The SynFormer models are trained using the AdamW optimizer with a learning rate of 3×10^{-4} and a batch size of 1,024. The training process spans approximately 730,000 steps, equivalent to a total of 742M data points. Validation is performed every 5,000 steps. A ReduceLROnPlateau scheduler is applied, reducing the learning rate by 20% if the validation loss does not improve over the last 10 validation runs. The models are trained on 8 A100 GPUs, and the total training time is about 6 days and 1 hours.

A.3 Inference details

A.3.1 Sampling algorithm

The sampling process starts with an empty stack and an empty postfix notation. The stack and the postfix notation are updated iteratively as described in Algorithm 2. Specifically, the SynFormer decoder first outputs the embedding vector and predicts the type for the next token based on the current postfix notation.

If the next token is predicted to be a building block, the diffusion module will generate a set of $k_{\text{fingerprint}}$ fingerprints. For each fingerprint, the building block molecule with the most similar fingerprint measured by L2 distance is retrieved from the database. $k_{\text{fingerprint}}$ copies of the previous stack and postfix notation are made and each building block is added to the stack and the notation of each copy, leading to $k_{\text{fingerprint}}$ different next states.

If the next token is predicted to be a reaction, k_{reaction} reactions with the highest probability are selected. For each selected reaction, the required number of reactants are popped out of the stack. If there are too few molecules in the stack, the reaction is considered inapplicable. Then, the reactants are input to the reaction template to predict the product. If the reaction runs successfully, the reaction product will be pushed onto the stack and the reaction itself will be appended to the postfix notation to form a next state. Otherwise, the reaction will be ignored and no next state for this reaction is produced.

Finally, if the next token is an [END] token, we will check if the stack contains exactly one molecule. If it does, the sampling process is considered successful and the only molecule on the stack is the final product. Otherwise, more than one molecule in the stack means the synthesis process is incomplete and thus is considered failed.

A.3.2 Branched sampling

The sampling algorithm can lead to branched synthesis pathways at each building block and reaction step. This capability is particularly useful when a specific objective, for example similarity to a given molecule in analog generation, is set. We leverage this capability by formulating a beam-search-like algorithm as described in Algorithm 3. Specifically, we maintain a pool of states and each state is a tuple of synthesis stack and postfix notation. In the beginning, the pool contains only one empty state. At each state, we run the sampling step (Algorithm 2) for each state in the pool, remove the existing state, and add the set of next generated states back to the pool. If the pool size exceeds the limit, we sort the state according to the some given metric and then keep only the top-ranked states. If a state is marked successful, it will be moved out of the pool and added to the result buffer.

Algorithm 2: Sampling step

Input: P_{prev} , previous postfix notation of synthesis**Input:** S_{prev} , previous synthesis stack**Input:** \mathbf{c} , features from the encoder part of SynFormer, optional**Input:** k_{reaction} , branching factor of reaction**Input:** $k_{\text{fingerprint}}$, branching factor of building block fingerprint

// Global variables: SynFormer, BuildingBlockDatabase

Output: an iterator of next states in the form of postfix notation and stack tuple (P, S)

```
1  $\mathbf{h} \leftarrow \text{SynFormer.decoder}(\text{embed}(P_{\text{prev}}), \mathbf{c})$ ; // Get the embedding for the next token
2  $t \leftarrow \text{SynFormer.predictNextType}(\mathbf{h})$ ; // Predict the type of the next token
3 if  $t = \text{BuildingBlock}$  then
4   foreach  $\mathbf{x}$  in  $\text{SynFormer.generateFingerprints}(\mathbf{h}, k_{\text{fingerprint}})$  do
5      $P, S \leftarrow \text{copy}(P_{\text{prev}}, S_{\text{prev}})$ ;
6      $X \leftarrow \text{BuildingBlockDatabase.retrieve}(\mathbf{x})$ ;
7      $\text{Push}(X \text{ onto } S)$ ;
8      $\text{Append}(X \text{ to } P)$ ;
9      $\text{Yield}(P, S)$ ;
10 else if  $t = \text{Reaction}$  then
11   foreach  $R$  in  $\text{top } k_{\text{reaction}}$  of  $\text{SynFormer.predictNextReaction}(\mathbf{h})$  do
12      $P, S \leftarrow \text{copy}(P_{\text{prev}}, S_{\text{prev}})$ ;
13      $\{X_i\} \leftarrow \text{Pop}(R.\text{numReactants} \text{ elements from } S)$ ;
14      $Y \leftarrow \text{Call } R(\{X_i\})$ ;
15     if there are sufficient reactants in the stack and the reaction applied successfully then
16        $\text{Push}(Y \text{ onto } S)$ ;
17        $\text{Append}(R \text{ to } P)$ ;
18        $\text{Yield}(P, S)$ ;
19 else if  $t = [\text{END}]$  then
20   if  $S_{\text{prev}}.\text{length}() = 1$  then
21      $\text{Yield}(P_{\text{prev}}, S_{\text{prev}}, \text{EndFlag})$ ;
```

Algorithm 3: Branched sampling

Input: k_{reaction} , branching factor of reaction**Input:** $k_{\text{fingerprint}}$, branching factor of building block fingerprint**Input:** m , size limit of the state pool**Input:** n , expected number of outputs**Input:** $f(\cdot)$, scoring function**Data:** $\mathcal{M} \leftarrow \{(P_0, S_0)\}$ where $P_0 = \emptyset$ and $S_0 = \emptyset$, initial state pool**Data:** $\mathcal{N} \leftarrow \{\}$, output buffer

```
1 while  $|\mathcal{N}| < n$  do
2   foreach  $(P, S)$  in  $\mathcal{M}$  do
3      $\text{Remove}((P, S) \text{ from } \mathcal{M})$ ;
4     foreach  $(P_{\text{next}}, S_{\text{next}}, \text{hasEndFlag})$  in  $\text{SamplingStep}(P, S, \mathbf{c}, k_{\text{reaction}}, k_{\text{fingerprint}})$  do
5       if  $\text{hasEndFlag}$  then
6          $\text{Add}((P_{\text{next}}, S_{\text{next}}) \text{ to } \mathcal{N})$ ; // Save to output buffer
7       else
8          $\text{Add}((P_{\text{next}}, S_{\text{next}}) \text{ to } \mathcal{M})$ ; // Add to the state pool
9      $\text{Sort}(\mathcal{M} \text{ according to the scoring function } f(S))$ ;
10     $\mathcal{M} \leftarrow \mathcal{M}[1 \dots m]$ ;
11 Return}  $(\mathcal{N})$ 
```

B Additional results

B.1 Impact of enlarged fingerprint length on performance

The length of Morgan fingerprints is a critical parameter that determines the resolution of molecular fingerprinting. Shorter fingerprint lengths can lead to bit collisions, where two similar molecules become indistinguishable due to the limited bit space [115]. However, directly predicting larger fingerprint lengths has proven challenging to scale (see Figure 14 in the SI of [65]), often necessitating a compromise at shorter lengths, such as 256 bits.




Figure 6: Effect of fingerprint length (256, 512, 1024, and 2048) on reconstruction performance on molecules sampled from ChEMBL. The left panel shows the Reconstruction Rate, while the right panel displays the Average Similarity. Both metrics indicate that longer fingerprints enhance model performance. All models were trained for an identical wall time of 24 hours and used the same hyperparameters, except for the varying fingerprint lengths and the corresponding network widths.

The introduction of the diffusion module offers greater flexibility in handling longer fingerprint lengths. We trained models with varied fingerprint lengths (256, 512, 1024, and 2048 bits), all under identical conditions: a fixed training duration of 24 hours and consistent hyperparameters, except for adjustments in fingerprint length and the corresponding network widths. The models were evaluated on their reconstruction performance using ChEMBL samples, as shown in Figure 6. The results demonstrate that increasing the fingerprint length significantly enhances model performance, with a shift from 256 bits to 2048 bits resulting in an increase in the reconstruction rate from 8% to 12%, representing a 50% improvement. These findings indicate that the diffusion module not only strengthens the modeling of molecular fingerprints but also scales effectively to longer fingerprints, enabling more precise nearest neighbor searches with higher resolution.

B.2 Synthesizable analog generation

This section provides additional results for synthesizable analog generation, focusing on two sets of experiments: (1) finding synthesizable analogs for *de novo* design efforts using the GuacaMol benchmark [76] (where scores range from 0 to 1, with higher scores indicating better performance), and (2) generating molecules with Pocket2Mol [82], a structure-based design algorithm.

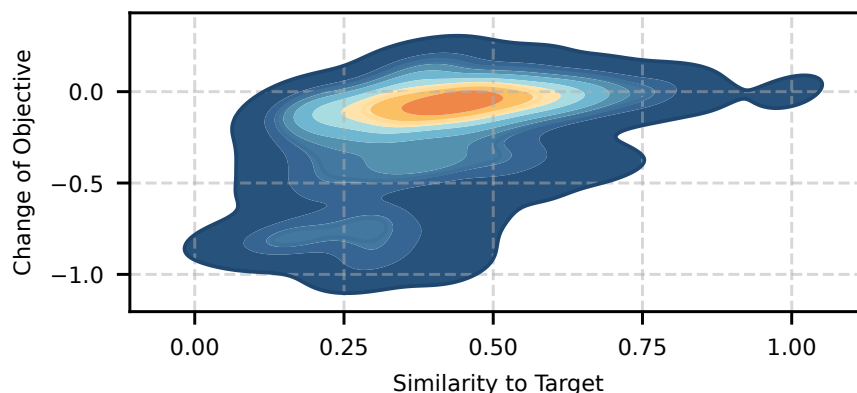


Figure 7: Relationship between similarity to the target molecule and the change in objective value. The contour plot shows the density of data points, indicating that higher similarity to the target molecule is generally associated with neutral changes in the objective value. This suggests that closer matches to the target tend to maintain the main objective score. The distribution highlights that a significant fraction of outputs exhibit neutral changes in objective scores.

Figure 7 presents the density plot of all molecules generated for unsynthesizable designs within the GuacaMol benchmark, plotted along the similarity to the target axis and the change in the main objective score axis. The results indicate that a substantial fraction of generated molecules exhibit relatively neutral changes in the main objective score, suggesting that the primary design objectives are largely preserved during analog generation. This is particularly evident for molecules with higher similarity to the target, which are clustered around zero change in the objective value, demonstrating that structurally similar analogs tend to retain their original objective scores.

However, the distribution also reveals a region characterized by low similarity and significant negative changes in objective scores, indicating instances where the model fails to generate meaningful analogs. These cases often result in outputs resembling random generation, highlighting the model’s limitations in some regions of the chemical space. This failure is partly due to the limited coverage of the explored chemical space compared to the theoretical coverage and partly due to the unsynthesizable input molecule being too distant from the synthesizable chemical space, making the “projection” unrealistic or irrelevant.

Figure 8 illustrates the distribution of Synthetic Accessibility (SA) Scores for molecules designed using Pocket2Mol (labeled as Target) and their corresponding generated analogs (labeled as Analog). The plot reveals a noticeable shift in the Analog distribution towards lower SA Scores, suggesting that the analog molecules are generally easier to synthesize compared to the original Pocket2Mol designs.

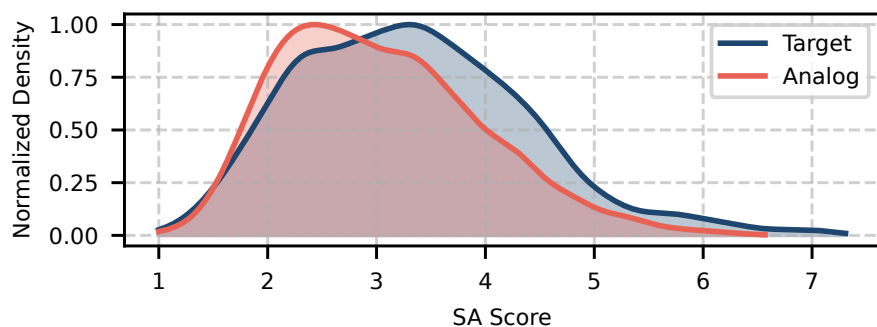


Figure 8: Distribution of SA Scores for molecules designed using Pocket2Mol (denoted as Target) and their corresponding analogs (denoted as Analog). The shift of the Analog distribution towards lower SA Scores indicates that the generated analog molecules are generally easier to synthesize compared to the original designs from Pocket2Mol.




Figure 9: Synthetic pathways for Perindopril MPO optimization examples as shown in the main text.




Figure 10: Additional examples of synthesizable analog generation for previously unsynthesizable designs. Each panel shows the transformation from the target molecule (left) to the product molecule (right), with objective scores indicating the performance of each molecule in the optimization task. The synthetic routes for the synthesizable analogs are shown below each panel, demonstrating the practical feasibility of the proposed transformations.

B.3 Hit expansion

This section provides additional results for hit expansion experiments, including (1) expanding hits from screening ZINC250k with activity predictors and (2) expanding from experimentally verified binders of disease targets.

In addition to JNK3, which is discussed in the main text, we also evaluated SynFormer on GSK3 β . Figure 11 displays contour plots comparing the density of molecules generated by SynFormer with those identified by the nearest neighbor approach for both JNK3 and GSK3 β . These plots illustrate the performance differences between the two methods in lead expansion. Additionally, Figure 12 presents the top-10 molecules identified from screening the ZINC250k database using each activity predictor.

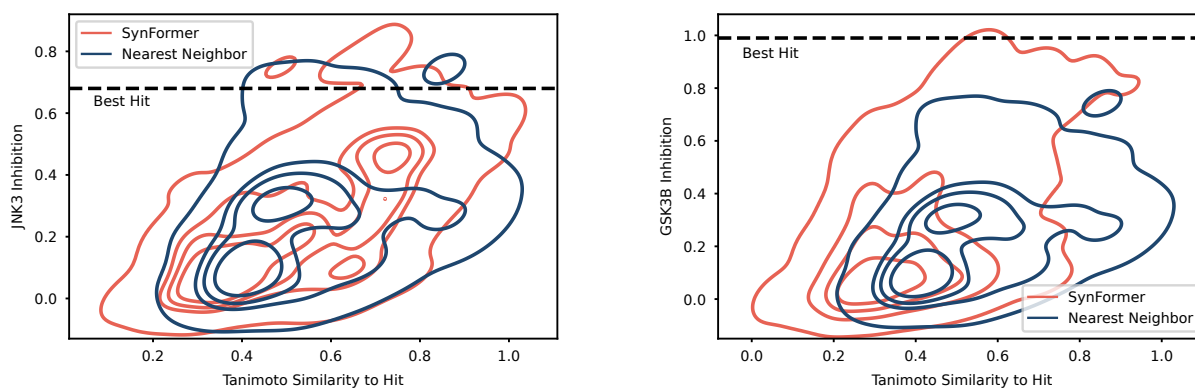




Figure 11: Comparison of hit expansion performance between SynFormer (red contours) and Nearest Neighbor (blue contours) approaches for JNK3 (left) and GSK3 β (right). The contour plots show inhibition scores versus Tanimoto similarity to the input hit compound. The dashed line represents the inhibition level of the best hit compound.

The density plots comparing analog discovery by nearest neighbor and SynFormer highlight that SynFormer consistently outperforms the nearest neighbor approach in identifying higher-scoring molecules in both tasks. As discussed in the main text, both SynFormer and nearest neighbor search identified molecules with predicted JNK3 inhibition scores higher than the best hits from initial screening, and SynFormer successfully found molecules with higher score than the highest one from the nearest neighbor search. For GSK3 β , the nearest neighbor approach failed to find molecules approaching the highest inhibition values of the original hits, whereas SynFormer succeeded. This is also demonstrated in Figure 13. Although nearest neighbor search showed enrichment in regions of higher similarity to the input references, molecules with scores exceeding those of the input references were often found at moderate similarity levels, ranging from 0.5 to 0.8. This observation suggests that SynFormer’s ability to explore beyond close structural analogs allows it to identify more potent molecules while maintaining a reasonable balance between structural similarity and enhanced activity.



(a) Hits for JNK3 identified from screening the ZINC250k database, with the predicted probability of being active labeled above each hit.



(b) Hits for GSK3 β identified from screening the ZINC250k database, with the predicted probability of being active labeled above each hit.

Figure 12: Hits identified from screening the ZINC250k database, which are subsequently used as inputs to the SynFormer for hit expansion.

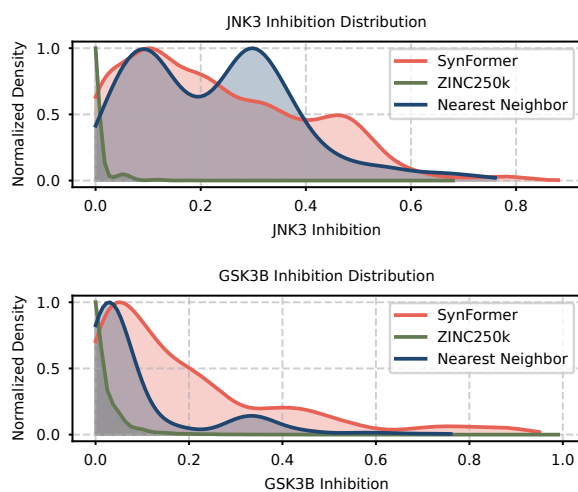


Figure 13: Normalized density plots of inhibition scores for JNK3 (top) and GSK3 β (bottom) showing the performance of SynFormer (red), nearest neighbor search (blue), and the initial ZINC250k screening (green). SynFormer consistently generates molecules with higher inhibition scores compared to the nearest neighbor and ZINC250k hits, demonstrating its superior capability in identifying potent analogs for both targets.




Figure 14: Additional examples of molecules generated from the hit against JNK3, along with their predicted inhibition scores between 0 and 1 and corresponding synthetic pathways.




Figure 15: Additional examples of molecules generated from the hit against $GSK3\beta$, along with their predicted inhibition scores between 0 and 1 and corresponding synthetic pathways.



Figure 16: Additional examples of molecules generated from the hit against PKM2.

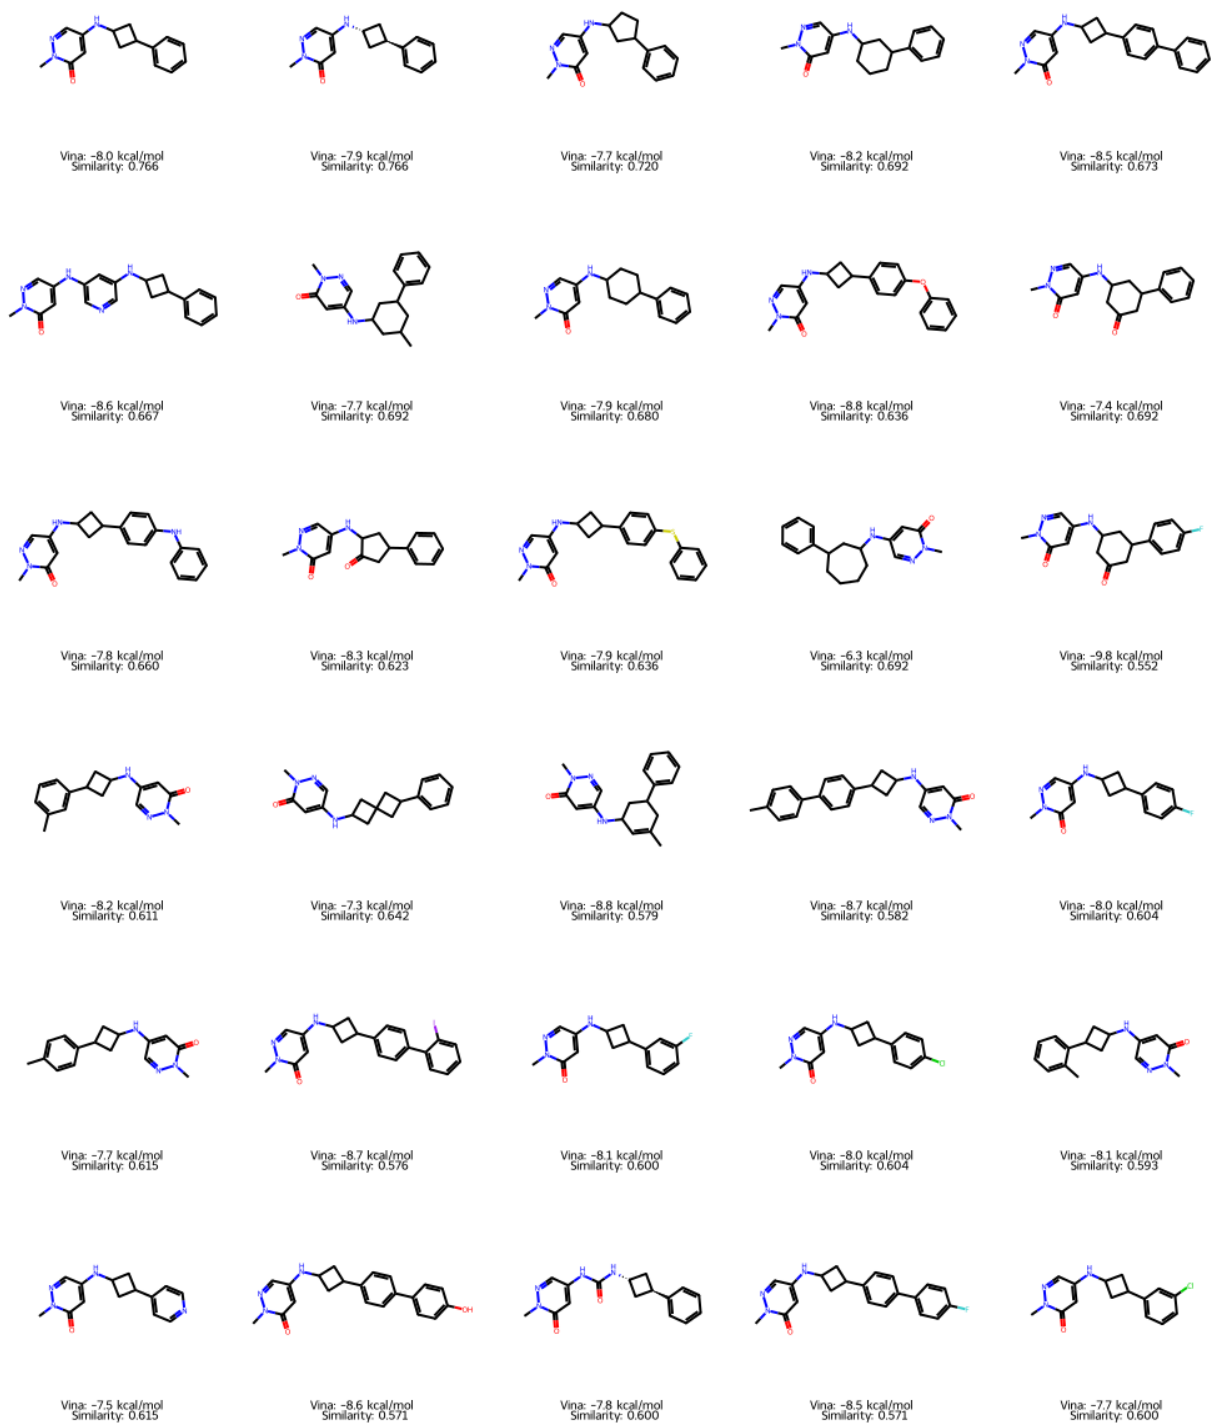



Figure 17: Additional examples of molecules generated from the hit against KAT2A.

B.4 Global chemical space exploration by *de novo* molecular optimization


B.4.1 Molecular optimization using reinforcement learning

This section describes additional results for using reinforcement learning to bias the generation of SynFormer-D toward high-scored molecules.

Figure 18a presents the score distribution of molecules generated by the SynFormer model after exposure to 5,000, 10,000, 15,000, 20,000, 25,000, and 30,000 labeled molecules and their corresponding scores. As the model processes more score-labeled data, the generation progressively shifts toward higher-scoring regions, demonstrating the effectiveness of the reinforcement learning (RL) algorithm in guiding the model toward optimized outputs. Figure 18b shows the accumulation of the top-100 molecules throughout the optimization process. The results indicate that the model’s performance saturates at the maximum score of around 1.0 after approximately 18,000 oracle evaluations, confirming the RL approach’s ability to consistently identify top-performing molecules.



(a) Score distribution of 200 molecules generated by SynFormer at different optimization steps, demonstrating that reinforcement learning effectively biases the generation distribution towards higher-scoring molecules.



(b) Score distribution of the accumulated top-100 molecules at various optimization steps, highlighting the progression and refinement of high-scoring molecules throughout the optimization process.

Figure 18: Additional results for SF-RL, the fine-tuned version of SynFormer-D using an approach adapted from REINVENT.

B.4.2 Molecular optimization using genetic algorithm

This section describes additional results for using a genetic algorithm with SynFormer-ED as a mutation.

Figure 19, 20, and 21 display the distribution of SA Scores [42] for the top 25 molecules at various optimization steps during the optimization of Perindopril MPO, Ranolazine MPO, and Scaffold Hop, respectively. In these tasks, since the main objective functions do not penalize synthetic accessibility, the differences between the standard GraphGA and GraphGA-SF are less pronounced compared to those observed in the Sitagliptin MPO task discussed in the main text.

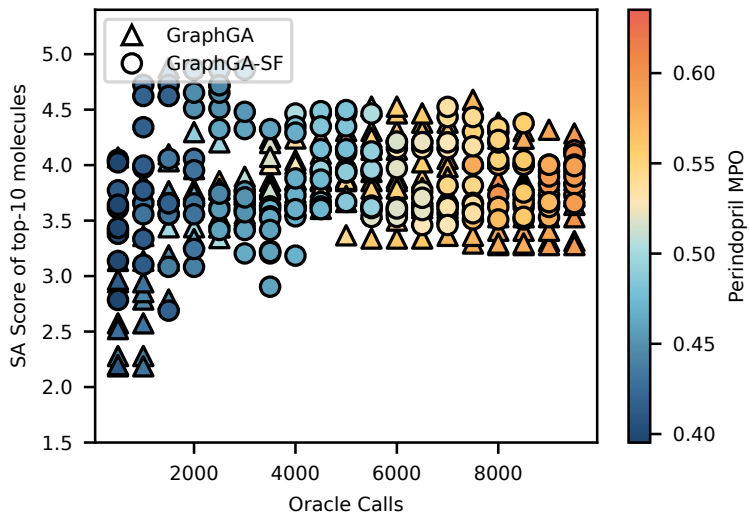


Figure 19: Distribution of SA Scores [42] for the top 25 molecules at various optimization steps, with colors representing objective scores.

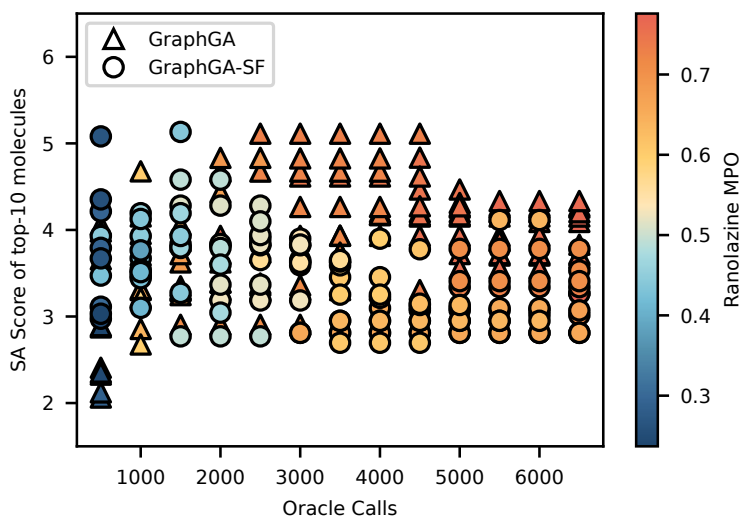


Figure 20: Distribution of SA Scores [42] for the top 25 molecules at various optimization steps, with colors representing objective scores.




Figure 21: Distribution of SA Scores [42] for the top 25 molecules at various optimization steps, with colors representing objective scores.