Synthon Contrastive Learning for Synthesizable 3D Molecule Generation

Nahyun Kim [1 *] **Seul Lee** [1 *] **Sung Ju Hwang** [1 2]

Abstract

To consider synthesizability in structure-based drug design (SBDD), a recent approach has proposed to co-design synthetic pathways with molecular conformers of ligands. However, the intermediate conformers are predicted based solely on the current intermediate molecule, without information about the synthons that will be added later along the synthetic pathway. As a result, the intermediate conformers do not adequately reflect the upcoming substructures, creating the geometric mismatch that leads to cumulative errors along the synthetic trajectory. To this end, we propose SYnthon Contrastive LEarning (SYCLE), a synthesizable SBDD framework that injects future synthon information into intermediate conformer prediction via contrastive learning. SYCLE further utilizes this information in the pathway generation policy to guide synthon selection. SYCLE achieves state-of-the-art binding affinity across all 15 LIT-PCBA targets. On CrossDocked2020, SYCLE attains -9.80 kcal/mol and an AiZynthFinder success rate of 68.8%, nearly double that of the best baseline.

1. Introduction

Recent advances in deep learning are rapidly proliferating AI-based approaches across various stages of drug discovery, ranging from hit identification to lead optimization, establishing AI-driven drug discovery as a pivotal research paradigm. In particular, structure-based drug design (SBDD) (Anderson, 2003) is a crucial methodology that involves designing ligands using the structural information of the binding pocket of a target protein.

However, most previous studies on SBDD molecule generation (Luo & Ji, 2022; Liu et al., 2022; Peng et al., 2022; Guan et al., 2024; Schneuing et al., 2024) have a critical limitation that they do not take molecular synthesizability into account. As a result, most of the generated candidates frequently lack clear synthetic routes or fail to satisfy practical constraints on building block availability and reaction selectivity, which significantly limits the practical impact of these methods (Gao & Coley, 2020; Walters & Barzilay, 2021). While some approaches attempt to incorporate synthesizability scores as additional optimization objectives (Coley et al., 2018; Ertl & Schuffenhauer, 2009; Thakkar et al., 2021), these heuristic proxies are often insufficient to capture the underlying complexity of chemical synthesis.

Consequently, there have been attempts to consider synthetic pathways in molecular generative models. However, they are largely restricted to 2D molecule generation and fail to integrate the 3D structural and geometric constraints of the protein binding pocket (Koziarski et al., 2024; Cretu et al., 2024; Seo et al., 2024; Gao et al., 2024; Jocys et al., 2025; Lee et al., 2026). 3DSynthFlow (Shen et al., 2025) was recently introduced to bridge this gap through the co-design of 3D molecular conformations and synthetic pathways. 3DSynthFlow interleaves two processes: (i) Compositional Flow, which uses a GFlowNet (Bengio et al., 2023) to model the sequential addition of molecular building blocks or synthons based on synthetic pathways, and (ii) State Flow, which models the associated 3D coordinates of intermediate molecules using flow matching. Specifically, synthesis steps are sampled at discrete time intervals, while 3D conformers are refined through a continuous-time flow.

Despite these advances, 3DSynthFlow remains limited in fully incorporating chemical constraints into conformer generation. Because State Flow is conditioned only on the currently observed intermediate state, the conformers of the intermediate molecules do not account for the synthons that will be added subsequently. This results in neglecting the steric and electronic influences of upcoming molecular substructures. Even if the intermediate molecule is identical, the final product molecule can adopt a significantly different conformation depending on which synthons are added to it. Consequently, even though the intermediate molecule shares a substructure with the final product molecule, a gap arises between the coordinates of those substructures, as shown in the top panel of Figure 1(a). Furthermore, because this gap in the substructure conformations will also influence

[*]Equal contribution [1]KAIST, Seoul, South Korea [2]DeepAuto. Correspondence to: Nahyun Kim <hyun0414@kaist.ac.kr>, Seul Lee <seul.lee@kaist.ac.kr>, Sung Ju Hwang <sjhwang82@kaist.ac.kr>.

Accepted at the 2026 Workshop on Generative and Agentic AI for Biology (ICML 2026)

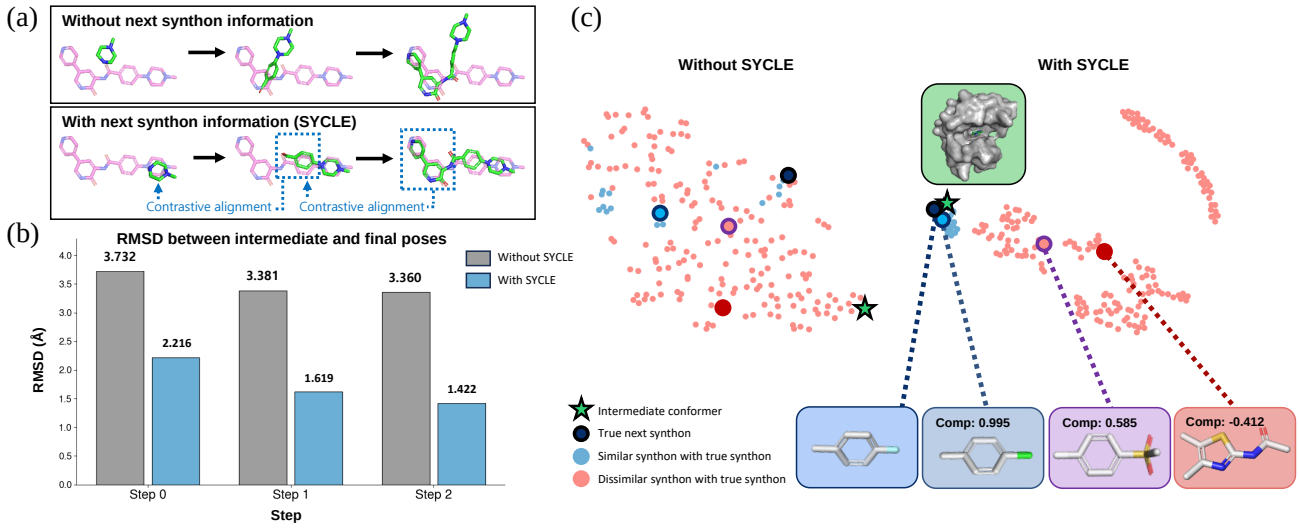

Figure 1. **Comparison of intermediate conformer prediction with and without SYCLE. (a)** Predicted docking poses at each synthesis step alongside the true binding pose. SYCLE generates intermediate and final conformers that are more consistent with the true binding pose. **(b)** RMSD to the true conformer at each synthesis step. SYCLE shows much lower RMSD across all steps, showing that the next synthon information improves conformer prediction. **(c)** t-SNE visualization of the shared embedding space with and without SYCLE. Synthons with a compatibility score above 0.99 to the true next synthon are denoted as similar synthons, and the other are denoted as dissimilar synthons. Without SYCLE, intermediate conformer embeddings are scattered with no clear clusters. With SYCLE, the intermediate conformer is pulled toward true and similar synthons and separated from dissimilar ones, demonstrating that SYCLE exhibits a chemically meaningful embedding space.

the selection of the next synthon, it can trigger cumulative errors that degrade both the subsequent pathway prediction and overall structural quality of the end product molecule. Therefore, a successful modeling of the conformers of intermediate molecules requires the knowledge of the reaction pathway and the synthons that will be added later.

To this end, rather than simply interleaving synthetic pathway generation and 3D molecule generation, we propose a synthetic pathway-aware 3D molecule generation approach. Specifically, we propose *SYnthon Contrastive LEarning* (SYCLE), a synthesizable 3D molecule generation framework that simultaneously considers the 3D geometry of protein pockets and the sequential constraints of reaction pathways. SYCLE leverages contrastive learning (Gao et al., 2023; Lohmann et al., 2025; Neeser et al., 2025), a self-supervised learning approach by aligning the representations of semantically similar pairs and contrasting them against negative samples. SYCLE extracts embeddings from partially generated intermediate 3D conformers within the pocket and computes embeddings for the candidate next synthons in a shared embedding space. By setting the positive set as the synthons that are attached in the next step to form the final ligand and controlling the degree of contrastive repulsion based on the synthon similarity (Figure 1(c) and Figure 2), the conformer predictor predicts intermediate conformers that are better aligned with the end product conformer, as shown in the bottom panel of Figure 1(a) and Figure 1(b). Furthermore, we propose to inject the *compatibility score*, which measures synthetic compatibility between the current intermediate conformer and a candidate next synthon, into the pathway generation policy to guide synthon selection. Overall, SYCLE enables effective construction of synthetic pathway-aware SBDD framework that reflects both protein-ligand interactions and synthetic compatibility between the intermediate synthon conformers.

Our contributions are summarized as follows:

- We introduce a synthon contrastive learning objective that learns the synthetic compatibility between a partial intermediate 3D conformer and the candidate next synthons to inject future synthon information into intermediate 3D conformer prediction.
- We propose the contrastive compatibility score that improves the synthon selection procedure to yield a high-quality final ligand under the pocket geometry.
- We empirically show that SYCLE achieves state-of-the-art performance in terms of both binding affinity and synthesizability, significantly outperforming previous baselines.

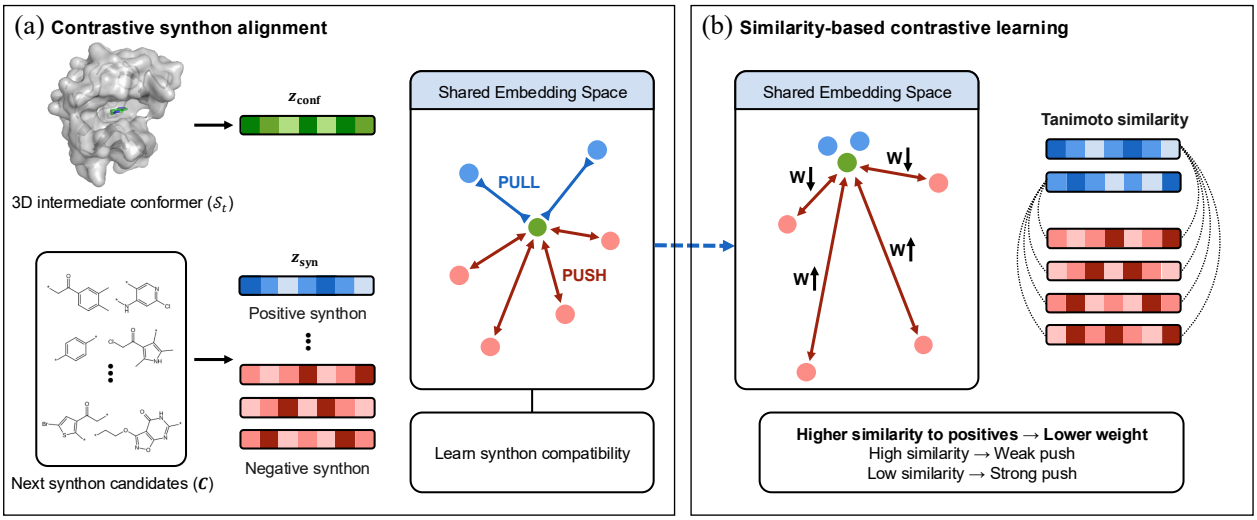

Figure 2. **The overall framework of SYCLE. (a)** The intermediate conformer and next synthon candidates are embedded in a shared embedding space. The contrastive objective pulls the embeddings of the true next synthons closer to the intermediate conformer embedding while pushing other synthon embeddings away. **(b)** Negative synthons are assigned weights inversely proportional to their similarity to the true next synthons, so that chemically closer negatives receive a weaker push.

2. Related Work

Structure-based drug design. Structure-based drug design (SBDD) leverages the structural information of target protein binding sites to design ligand molecules that satisfy geometric and chemical constraints while achieving high 3D structural stability and binding affinity (Guan et al., 2023). Early methods proposed to utilize only the shape of the binding site for ligand generation (Long et al., 2022; Adams & Coley, 2022), while others directly modeled target-ligand correspondence from protein-ligand complex data through diverse generative strategies, including autoregressive models that construct molecules atom-by-atom (Luo & Ji, 2022; Liu et al., 2022; Peng et al., 2022), and diffusion-based approaches that capture the continuous distribution of 3D molecular structures (Guan et al., 2024; Schneuing et al., 2024). Despite these advances, most SBDD methods do not account for molecular synthesizability, frequently producing candidates that lack clear synthetic routes or violate practical constraints on building block availability and reaction selectivity (Gao & Coley, 2020), which limits their practical utility in drug discovery pipelines.

Synthesizable molecule design. Ensuring synthetic accessibility in generative molecular design has been a central challenge in drug discovery. The most effective approach constrains the generative process to a space defined by reaction rules and purchasable building blocks. To navigate this constrained space, a range of optimization strategies have been explored in early work, including autoencoders (Bradshaw et al., 2019; 2020), genetic algorithms (Gao et al., 2021), and reinforcement learning (Gottipati et al., 2020; Horwood & Noutahi, 2020). More recently, GFlowNet-based (Koziarski et al., 2024; Cretu et al., 2024; Seo et al., 2024; Gaiński et al., 2025) and transformer-based (Gao et al., 2024; Jocys et al., 2025; Lee et al., 2026) approaches have demonstrated strong performance in synthesizable molecule generation. However, these methods are largely confined to 2D molecular representations, without accounting for 3D structural constraints. Recently, Rekesh et al. (2025) extended this to joint 3D coordinate generation, but operates unconditionally without protein pocket information. 3DSynthFlow (Shen et al., 2025) addressed this by co-designing synthesis pathways and 3D conformations within protein pockets. However, because its intermediate conformers lacks awareness of subsequent synthons along the synthetic pathways, a discrepancy arises between the geometric structures of the intermediate compounds and the final product. SYCLE addresses this gap with contrastive learning that models the compatibility between intermediate conformers and next synthon candidates, enabling pathway-aware 3D SBDD throughout the synthesis trajectory.

3. Method

We introduce SYCLE, a synthesizable SBDD framework with aligned intermediate conformations, by learning synthetic pathway-aware compatibility between the 3D conformations of partial ligands and the next synthon candidates. We first introduce the GFlowNet-based generation

framework and define the action space in Section 3.1. We then present the contrastive compatibility objective with similarity-based weighting in Section 3.2, followed by the forward logit augmentation for compatibility-guided generation in Section 3.3.

3.1. Preliminary: 3D Molecule and Synthesis Pathway Co-design

Following Shen et al. (2025), we model a synthesizable ligand as $\boldsymbol{x} = (\mathcal{C}, \mathcal{S})$, where $\mathcal{C}$ denotes the synthetic pathway and $\mathcal{S}$ denotes the corresponding 3D conformations. Here, $\mathcal{C} = (\boldsymbol{C}^{(i)})_{i=1}^{n}$ is a sequence of n synthons to be attached (one per synthesis step), where each synthon $\boldsymbol{C}^{(i)}$ contains m_i atoms. $\mathcal{S} = (\boldsymbol{S}^{(i)})_{i=1}^{n}$ denotes the 3D coordinates of synthons, where $\boldsymbol{S}^{(i)} \in \mathbb{R}^{m_i \times 3}$ is the conformation of $\boldsymbol{C}^{(i)}$. Generation proceeds by interleaving discrete synthon attachment with continuous 3D conformation refinement over a global continuous time $t \in [0, 1]$.

Synthon attachment. Let $k(t)$ be the number of synthons that have been attached by time t, so that the compositional structure is $\mathcal{C}_t = (\boldsymbol{C}^{(i)})_{i=1}^{k(t)}$. At discrete time $t = t^{(i)}_{\text{gen}}$, we sample the next synthon $\boldsymbol{C}^{(i)}$ from the synthon attachment policy π_ϕ:

$$\boldsymbol{C}^{(i)} \sim \pi_\phi(\cdot \mid \boldsymbol{x}_t, \hat{\boldsymbol{x}}_t^1), \tag{1}$$

where $\hat{\boldsymbol{x}}_t^1$ is the model's prediction of the clean 3D conformation. Then, a transition function $T(\boldsymbol{x}_t, \boldsymbol{C}^{(i)})$ merges the new synthon into the current object. T also samples an initial conformation $\boldsymbol{S}_0^{(i)}$ for the newly introduced synthon from an isotropic Gaussian and inserts it into $\mathcal{S}_t$.

Synthesizable molecule generation requires selecting discrete actions—choosing which synthon to attach and via which reaction—such that the resulting molecules are sampled in proportion to a terminal reward $R(x)$ combining docking affinity and drug-likeness. Each action is defined over a purchasable building block library $\mathcal{B}$, where $\mathcal{B}_s \subseteq \mathcal{B}$ denotes blocks with a single attachment site and $\mathcal{B}_b \subseteq \mathcal{B}$ denotes compatible reactant blocks for a bimolecular reaction template r. Each step in the synthesis trajectory corresponds to one of two action types (Seo et al., 2024). At the start of generation, a FIRSTSYNTHON action selects an initial synthon $\boldsymbol{C}^{(0)} \in \mathcal{B}_s$ conditioned on the protein pocket. The policy network π_ϕ computes a logit ℓ to select each action:

$$\ell = \pi_\phi(\boldsymbol{S}^{(0)}, \boldsymbol{C}^{(0)}). \tag{2}$$

All subsequent decisions are ADDSYNTHON actions, each defined as a $(r, \boldsymbol{C}^{(i)})$-pair consisting of a bimolecular reaction r compatible to $\mathcal{C}_t$ and a synthon $\boldsymbol{C}^{(i)} \in \mathcal{B}_b$. The corresponding logit ℓ is:

$$\ell = \pi_\phi(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})), \tag{3}$$

where $\mathcal{S}_t$ is the overall conformers of the attached synthons so far at the current time t. The trajectory terminates when a terminal synthon is selected or the maximum synthesis depth is reached. The forward policy $P_F(\boldsymbol{C}^{(i)} \mid \mathcal{S}_t; \phi)$ is obtained by normalizing ℓ over all available actions at each state.

3D conformer prediction. Each synthon $\boldsymbol{C}^{(i)}$ is generated at a different time, so its conformation should be refined from the moment it is attached. To encode this temporal bias, a component-wise local time is defined as:

$$t^{(i)}_{\text{loc}} = \text{clip}(t - t^{(i)}_{\text{gen}}) \in [0, 1], \tag{4}$$

where $t^{(i)}_{\text{gen}}$ is the attachment time of i-th synthon. An intermediate conformation is constructed via linear interpolation between the fully noisy initial conformation $\boldsymbol{S}_0^{(i)}$ and the clean target conformation $\boldsymbol{S}_1^{(i)}$, with Gaussian noise:

$$\boldsymbol{S}_t^{(i)} \sim \mathcal{N}\Big(t^{(i)}_{\text{loc}} \boldsymbol{S}_1^{(i)} + (1 - t^{(i)}_{\text{loc}}) \boldsymbol{S}_0^{(i)},\ \sigma^2 I\Big) \quad \text{for } t > t^{(i)}_{\text{gen}}. \tag{5}$$

The conformation prediction network $p^{\psi}_{1|t}$ consists of a POCKETENCODER and a LIGANDDECODER. The POCKETENCODER encodes the protein pocket using HierSemla (Shen et al., 2025), a hierarchical extension of Semla (Irwin et al., 2024), producing protein features $\boldsymbol{h}^{\text{pro}}$. The LIGANDDECODER processes the partially assembled ligand through Semla layers with protein-ligand cross-attention with $\boldsymbol{h}^{\text{pro}}$, producing ligand features $\boldsymbol{h}^{\text{lig}}$. The predicted clean coordinates $\hat{\boldsymbol{S}}_1^{(i)}$ are then obtained by normalizing and projecting the equivariant component of $\boldsymbol{h}^{\text{lig}}$.

Given the current object $\boldsymbol{x}_t = (\mathcal{C}_t, \mathcal{S}_t)$, $p^{\psi}_{1|t}$ predicts the clean conformations $\hat{\boldsymbol{S}}_1^{(i)} = p^{\psi}_{1|t}(\boldsymbol{x}_t)^{(i)}$ for all synthons attached so far. During sampling, we define a simple vector field $\hat{\boldsymbol{S}}_1^{(i)} - \boldsymbol{S}_t^{(i)}$ and update the conformation using an Euler step:

$$\boldsymbol{S}^{(i)}_{t+\Delta t} = \boldsymbol{S}^{(i)}_t + \big(\hat{\boldsymbol{S}}_1^{(i)} - \boldsymbol{S}_t^{(i)}\big) \cdot \kappa^{(i)} \Delta t, \tag{6}$$

where $\kappa^{(i)}$ scales the step size based on the remaining time until the interpolation for i-th synthon is complete.

Training objective. The conformer predictor $p^{\psi}_{1|t}$ is independently trained by minimizing

$$\mathcal{L}_{\text{conf}} = \mathbb{E}_{\boldsymbol{x}_1 \sim p_{\text{data}},\, t \sim \mathcal{U}(0,1)} \sum_{i=1}^{k(t)} \Big\| p^{\psi}_{1|t}(\boldsymbol{x}_t)^{(i)} - \boldsymbol{S}_1^{(i)} \Big\|_2^2. \tag{7}$$

After the conformer predictor is trained, we fix $p^{\psi}_{1|t}$ and optimize the synthon attachment policy π_ϕ using a trajectory-balance objective. We optimize the trajectory

balance (TB) objective of GFlowNets (Malkin et al., 2022), which is equivalent to optimizing the following objective (Shen et al., 2025):

$$\mathcal{L}_{\mathrm{TB}} = \left(\log \frac{Z_\phi \prod_{i=0}^{n-1} P_F\left(\boldsymbol{C}^{(i)} \mid \boldsymbol{x}_{t_{\mathrm{gen}}^{(i)}}, \hat{\boldsymbol{x}}^1_{t_{\mathrm{gen}}^{(i)}}; \phi\right)}{R(\boldsymbol{x}_1)} \right)^2, \tag{8}$$

where Z_ϕ is a learnable normalizer, $P_F(\boldsymbol{C}^{(i)} \mid \cdot\,; \phi)$ denotes the probability of selecting synthon $\boldsymbol{C}^{(i)}$ under the policy π_ϕ, and $R(\boldsymbol{x}_1)$ is a reward defined on the final generated object. At inference time, the learned policy and the conformer predictor are executed together by alternating discrete attachment and continuous refinement until $t = 1$.

3.2. Next Synthon Contrastive Learning

Intermediate 3D conformer predictor described in Section 3.1 is trained to refine a partially assembled ligand conformation without explicitly knowing which synthon will be attached next. However, the identity of the next synthon can strongly affect the final product geometry, and these geometric constraints should propagate back to earlier intermediate conformations. To incorporate this reaction pathway signal, we introduce next synthon contrastive compatibility learning, which aligns the current intermediate conformation with the synthons that can attach in the next step.

Construction of next synthon candidate set. We first define the next synthon candidate set via a synthon-level decomposition of each ligand. Specifically, each ligand is decomposed into synthons by applying reaction-aware bond-breaking rules following Enamine synthesis protocols (Grygorenko et al., 2020), yielding a set of synthon groups $(\boldsymbol{C}^{(i)})_{i=1}^n$. We further define a synthon connectivity graph $\mathcal{E}$, where $(\boldsymbol{C}^{(i)}, \boldsymbol{C}^{(j)}) \in \mathcal{E}$ if a chemical bond exists between $\boldsymbol{C}^{(i)}$ and $\boldsymbol{C}^{(j)}$ in the original ligand.

To determine the generation order, each synthon $\boldsymbol{C}^{(i)}$ is assigned a generation time $t_{\mathrm{gen}}^{(i)}$ via BFS-ordered traversal of $\mathcal{E}$. At global time t, the visible and masked synthon sets are defined as:

$$\mathcal{V}_t = \{\boldsymbol{C}^{(i)} \mid t_{\mathrm{gen}}^{(i)} \le t\}, \quad \mathcal{M}_t = \{\boldsymbol{C}^{(i)} \mid t_{\mathrm{gen}}^{(i)} > t\}. \tag{9}$$

Throughout Section 3.2, we use square-bracket superscripts $[\cdot]$ to index samples within a mini-batch, and use round-bracket superscripts $(\cdot)$ to index synthesis steps as in Section 3.1. The next synthon candidate set consists of all masked synthons that are adjacent to the current visible set:

$$\mathcal{C}_t^{[m]} = \{\boldsymbol{C}^{[m](i)} \in \mathcal{M}_t \mid \exists\, \boldsymbol{C}^{[m](j)} \in \mathcal{V}_t,\ (\boldsymbol{C}^{[m](i)}, \boldsymbol{C}^{[m](j)}) \in \mathcal{E}\}. \tag{10}$$

The positive set for the synthon contrastive compatibility learning is defined as $\mathcal{C}^{+,[m]} := \mathcal{C}^{[m]}$. For contrastive learning, the normalization pool is formed at the mini-batch level. For a mini-batch B, we aggregate the candidate sets across all samples:

$$\mathcal{C}_{\mathrm{mb}} = \bigcup_{[m] \in B} \mathcal{C}_t^{[m]}. \tag{11}$$

For each intermediate conformation $\mathcal{S}_t^{[m]}$, synthons in $\mathcal{C}^{+,[m]}$ are treated as positives, while in-batch negatives are defined as candidates in $\mathcal{C}^{-,[m]} := \mathcal{C}_{\mathrm{mb}} \setminus \mathcal{C}^{+,[m]}$.

Conformer and synthon embeddings. We extract atom-wise invariant and equivariant features $\boldsymbol{h}^{\mathrm{lig}}$ of the current intermediate ligand conformation $\mathcal{S}_t$ by passing it through the LIGANDDECODER conditioned on the pocket embedding obtained from the POCKETENCODER. Because the next reaction step is governed by the local geometry around the attachment site, we apply *attachment-aware mean pooling*:

- If attachment atoms exist, we compute the mean of features over attachment atoms.
- Otherwise, we fall back to mean pooling over all valid atoms.

Denote this pooled representation as $\boldsymbol{f}_{\mathrm{conf}} = \mathrm{AttPool}(\boldsymbol{h}^{\mathrm{lig}})$. We then map it to the shared contrastive space through a projection head $\mathrm{Proj}_{\mathrm{conf}}$:

$$\boldsymbol{z}_{\mathrm{conf}}(\mathcal{S}_t) = \mathrm{Proj}_{\mathrm{conf}}(\boldsymbol{f}_{\mathrm{conf}}) \in \mathbb{R}^d. \tag{12}$$

For each candidate next synthon $\boldsymbol{C} \in \mathcal{C}^{[m]}$, we first generate a 3D conformer via RDKit and represent it as a LIGANDBATCH consisting of atom types, formal charges, bond adjacency, 3D coordinates, and attachment point flags. This is encoded using a lightweight SYNTHONENCODER and produces atom-wise features $\boldsymbol{h}^{\mathrm{syn}}$, which we pool to obtain $\boldsymbol{f}_{\mathrm{syn}} = \mathrm{AttPool}(\boldsymbol{h}^{\mathrm{syn}})$. A symmetric projection head yields:

$$\boldsymbol{z}_{\mathrm{syn}}(\boldsymbol{C}) = \mathrm{Proj}_{\mathrm{syn}}(\boldsymbol{f}_{\mathrm{syn}}) \in \mathbb{R}^d. \tag{13}$$

Consequently, the current intermediate conformation and the next synthon candidates are embedded in a shared embedding space. The overall architecture of SYCLE is described in Figure 3 and Section C.

Contrastive learning objective. We define the *compatibility score* as the cosine similarity between a conformer $\mathcal{S}_t$ and a candidate next synthon $\boldsymbol{C}$ in the shared embedding space:

$$\mathrm{Comp}(\mathcal{S}_t, \boldsymbol{C}) = \frac{\langle \tilde{\boldsymbol{z}}_{\mathrm{conf}}(\mathcal{S}_t),\ \tilde{\boldsymbol{z}}_{\mathrm{syn}}(\boldsymbol{C}) \rangle}{\tau_c}, \tag{14}$$

where $\tilde{\boldsymbol{z}}$ denotes ℓ_2-normalized embeddings and τ_c is a temperature. We propose to use the multi-positive InfoNCE

objective (Khosla et al., 2020) to handle cases where multiple next synthons are valid:

$$\mathcal{L}_{\text{cont}}^{[m]} = -\log \frac{\sum_{\boldsymbol{C}\in\mathcal{C}^{+,[m]}} \exp(\text{Comp}(\mathcal{S}_t^{[m]}, \boldsymbol{C}))}{\sum_{\boldsymbol{C}\in\mathcal{C}_{\text{mb}}} \exp(\text{Comp}(\mathcal{S}_t^{[m]}, \boldsymbol{C}))}. \quad (15)$$

This objective pulls the next synthons along the synthetic pathway toward the current intermediate conformer embedding, while repelling other synthons, thereby incorporating the synthetic pathway information into the shared embedding space.

However, some negative synthons are chemically close to the true next synthon and can be alternative synthons to be attached. In this case, controlling the degree of contrastive repulsion would be more informative. Therefore, we propose to perform contrastive learning by setting *soft* negatives instead of hard ones. We introduce similarity-based weighting using the learned synthon features. We first binarize the synthon feature $\boldsymbol{f}_{\text{syn}}$ by thresholding each dimension at its batch-wise median, then compute a Tanimoto similarity $\text{Tan}(\boldsymbol{C}, \boldsymbol{C}') \in [0, 1]$ between candidate synthons. For each synthon $\boldsymbol{C} \in \mathcal{C}_{\text{mb}}$, we measure its maximum similarity to any true next synthon:

$$\text{sim}(\mathcal{C}^{+,[m]}, \boldsymbol{C}) = \max_{\boldsymbol{C}'\in\mathcal{C}^{+,[m]}} \text{Tan}(\boldsymbol{C}, \boldsymbol{C}'). \quad (16)$$

We then define a similarity-based weight:

$$w(\boldsymbol{C}) = \begin{cases} 1, & \text{if } \boldsymbol{C} \in \mathcal{C}^{+,[m]} \\ \exp\left(-\frac{\text{sim}(\mathcal{C}^{+,[m]}, \boldsymbol{C})}{\tau_w}\right), & \text{otherwise} \end{cases} \quad (17)$$

where τ_w is a temperature. We apply the weighting to Eq. (15):

$$\mathcal{L}_{\text{cont}}^{[m]} = -\log \frac{\sum_{\boldsymbol{C}\in\mathcal{C}^{+,[m]}} \exp(\text{Comp}(S_t^{[m]}, \boldsymbol{C}))}{\sum_{\boldsymbol{C}\in\mathcal{C}_{\text{mb}}} w(\boldsymbol{C}) \exp(\text{Comp}(S_t^{[m]}, \boldsymbol{C}))}. \quad (18)$$

Intuitively, synthons with a high structural similarity to the positive set receive a smaller weight $w(\boldsymbol{C})$, weakening the contrastive penalty imposed on them.

The final training objective for the conformer predictor is obtained by combining the conformation loss $\mathcal{L}_{\text{conf}}$ (Eq. (7)) and the contrastive learning loss $\mathcal{L}_{\text{cont}}$ (Eq. (18)):

$$\mathcal{L} = \mathcal{L}_{\text{conf}} + \lambda \cdot \mathcal{L}_{\text{cont}}, \quad (19)$$

where λ controls the strength of the contrastive learning signal. The contrastive learning objective guides the intermediate conformer generation with the knowledge about both the synthons in the next synthetic steps and the protein pocket structure.

3.3. Next Synthon Selection with Compatibility Score

The proposed learning scheme of the conformer and synthon embeddings is not only beneficial to the intermediate conformer generation, but also to the synthetic pathway generation itself. We propose to inject the compatibility score (Eq. (14)) into the pathway generation policy to guide synthon selection.

For AddSynthon actions, we augment the logits (Eq. (3)) with the compatibility score:

$$\ell_{\text{SYCLE}}(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})) = \ell(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})) + \beta \cdot \text{Comp}(\mathcal{S}_t, \boldsymbol{C}^{(i)}), \quad (20)$$

where $\beta > 0$ controls the strength of the compatibility score augmentation. The augmentation is not applied for FirstSynthon actions, as these actions select an initial synthon without an existing intermediate reaction context.

By utilizing the compatibility score both in the contrastive learning objective of the conformer predictor and the logits of the next synthon selection actions, SYCLE effectively couples (i) how intermediate 3D geometries are refined under the pocket structure and (ii) which synthons are selected along the synthesis pathway, yielding a synthetic pathway-aware 3D molecule generation pipeline.

In practice, synthon embeddings $\tilde{\boldsymbol{z}}_{\text{syn}}(\boldsymbol{C})$ are precomputed for all synthons in the block library prior to training, so that compatibility scoring introduces no additional computational overhead at train time. Conformer embeddings $\tilde{\boldsymbol{z}}_{\text{conf}}(\mathcal{S}_t)$ are computed once per visited state and then reused to evaluate multiple candidate synthons, making the compatibility augmentation computationally lightweight.

4. Experiments

4.1. Pocket-specific Optimization

Setup. Synthetic pathways are constructed using the 38 bimolecular Enamine reaction protocols from Gao et al. (2024). The pocket-specific optimization uses a 1.2M Enamine catalog as a building block libraries. Additionally, building blocks with a drug-likeness score under 60 are excluded following Lee et al. (2022). Molecular complexity is bounded at two synthesis steps per molecule, in line with the practical scope of Enamine REAL Space (Grygorenko et al., 2020). Synthesizability is evaluated via AiZynthFinder (Genheden et al., 2020), a retrosynthetic planning tool run against the full Enamine catalog as the stock database. Protein-ligand interactions are evaluated with PoseCheck (Harris et al., 2023), which counts hydrogen-bond acceptor/donor, van der Waals, and hydrophobic contacts between the ligand and pocket.

For reward design and evaluation, we follow the protocol of Seo et al. (2024) and Shen et al. (2025). Binding affinity

Table 1. **Pocket-specific optimization results on LIT-PCBA.** Average Vina docking scores for the top-100 diverse molecules across five representative targets, along with the average and median over all 15 LIT-PCBA targets. The results for the other 10 targets are in Table 10. For SynFlowNet, v2024.05[a] and v2024.10[b] are reported. The results for the baselines are taken from Shen et al. (2025). The best results are highlighted in bold.

Category	Method	Average Vina Docking Score (kcal/mol, ↓)						
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1	Avg.	Med.
Fragment	FragGFN	-10.19 (± 0.33)	-10.43 (± 0.29)	-9.81 (± 0.09)	-9.85 (± 0.13)	-7.67 (± 0.71)	-9.58	-9.85
	FragGFN+SA	-9.70 (± 0.61)	-9.83 (± 0.65)	-9.27 (± 0.95)	-10.06 (± 0.30)	-7.26 (± 0.10)	-9.22	-9.58
Reaction	SynNet	-8.03 (± 0.26)	-8.81 (± 0.21)	-8.88 (± 0.13)	-8.52 (± 0.16)	-6.36 (± 0.09)	-8.12	-8.52
	BBAR	-9.95 (± 0.04)	-10.06 (± 0.14)	-9.97 (± 0.03)	-9.92 (± 0.05)	-6.84 (± 0.07)	-9.35	-9.84
	SynFlowNet[a]	-10.85 (± 0.10)	-10.69 (± 0.09)	-10.44 (± 0.05)	-10.27 (± 0.04)	-7.47 (± 0.02)	-9.95	-10.34
	SynFlowNet[b]	-9.17 (± 0.68)	-9.37 (± 0.29)	-9.17 (± 0.12)	-9.05 (± 0.14)	-6.45 (± 0.13)	-8.78	-9.17
	RGFN	-9.84 (± 0.21)	-9.93 (± 0.11)	-9.99 (± 0.11)	-9.72 (± 0.14)	-6.92 (± 0.06)	-9.08	-9.91
	RxnFlow	-11.45 (± 0.05)	-11.26 (± 0.07)	-11.15 (± 0.02)	-10.77 (± 0.04)	-7.66 (± 0.02)	-10.46	-10.84
3D Reaction	3DSynthFlow	-11.96 (± 0.12)	-11.82 (± 0.03)	-11.58 (± 0.07)	-11.23 (± 0.08)	-7.79 (± 0.01)	-10.89	-11.26
	SYCLE	**-12.10** (± 0.09)	**-11.98** (± 0.08)	**-11.64** (± 0.21)	**-11.45** (± 0.05)	**-8.16** (± 0.10)	**-11.14**	**-11.50**

Table 2. **AiZynthFinder success rates** for the 15 LIT-PCBA protein targets. The results for the baselines are taken from Shen et al. (2025).

Method	Success Rate (%, ↑)	Synthesis Steps (↓)
FragGFN+SA	3.52	3.74
SynNet	47.50	3.45
BBAR	17.92	3.68
SynFlowNet[a]	54.60	2.55
SynFlowNet[b]	58.38	2.47
RGFN	47.43	2.46
RxnFlow	65.35	2.17
3DSynthFlow	68.58	2.39
SYCLE	**69.84**	**1.58**

Table 3. **Protein-ligand interaction.** PoseCheck scores averaged over 4 runs for the top-100 diverse molecules across the first 5 LIT-PCBA targets. For SynFlowNet, v2024.05[a] and v2024.10[b] are reported. The results for the baselines are taken from Shen et al. (2025). The best results are in bold.

Method	H-Bond Acceptors	H-Bond Donors	Van der Waals	Hydrophobic	Sum
SynFlowNet[a]	0.22 (± 0.01)	0.11 (± 0.01)	9.31 (± 0.05)	10.41 (± 0.05)	20.05
SynFlowNet[b]	0.22 (± 0.03)	0.10 (± 0.01)	8.38 (± 0.05)	9.44 (± 0.06)	18.14
RGFN	0.19 (± 0.01)	0.11 (± 0.01)	9.19 (± 0.28)	10.25 (± 0.12)	19.73
RxnFlow	0.22 (± 0.00)	0.10 (± 0.01)	9.63 (± 0.09)	10.67 (± 0.10)	20.62
3DSynthFlow	0.33 (± 0.06)	**0.17** (± 0.03)	11.04 (± 0.37)	10.79 (± 0.21)	22.32
SYCLE	**0.34** (± 0.02)	0.16 (± 0.03)	**11.22** (± 0.04)	**11.20** (± 0.01)	**22.91**

is estimated via Uni-Dock (Yu et al., 2023), a GPU-based docking engine, with a per-target budget of 64,000 generated molecules. QED (Bickerton et al., 2012) is jointly optimized as a co-objective, and molecules are generated with a maximum of 40 heavy atoms to discourage size-driven reward exploitation. The full GFlowNet reward function is provided in Section D.2. From these, we retain molecules with QED > 0.5 and select the top-100 by docking score, subject to pairwise Tanimoto similarity < 0.5 computed on RDKit fingerprints (Landrum et al., 2013) to maintain chemical diversity.

Baselines. For pocket-specific optimization, we compare SYCLE against previous methods spanning three categories: fragment-based, reaction-based, and 3D reaction-based generation. In the fragment category, we evaluate FragGFN and FragGFN+SA, which differ in whether synthetic accessibility (Ertl & Schuffenhauer, 2009) is explicitly incorporated into the objective. The reaction-based baselines include SynNet (Gao et al., 2021), BBAR (Seo et al., 2023), two variants of SynFlowNet (Cretu et al., 2024), RGFN (Koziarski et al., 2024), and RxnFlow (Seo et al., 2024), covering search-based, conditional generative, and GFlowNet-based methods. As the most closely related 3D reaction-based baseline, we include 3DSynthFlow as a baseline (Shen et al., 2025).

Results. Table 1 and Table 10 shows the docking scores for the 15 LIT-PCBA targets. SYCLE achieves the best scores on all targets, significantly outperforming the baselines in terms of the average docking score. These results demonstrate that SYCLE's strategy of incorporating future synthon information into the prediction of intermediate conformers generates high-quality ligands that bind stronger to the given protein pocket.

Synthesizability and the number of synthetic steps estimated by AiZynthFinder are shown in Table 2. SYCLE achieves the highest AiZynthFinder success rate, showing the best synthesizability that outperforms all the baselines. Furthermore, SYCLE requires significantly fewer synthesis steps compared to the baselines, indicating that the contrastive objective guides the synthon attachment policy toward more concise and feasible synthetic pathways. Fewer synthesis steps are generally preferred, as synthetic complexity is known to correlate with practical considerations such as yield and cost (Coley et al., 2018).

Table 4. **Pocket-conditional generation results on CrossDocked2020.** For each metric for valid molecules, the mean value (Avg.) and median value (Med.) are reported. References indicate known active molecules. Results for baselines other than SA scores and diversity are taken from Shen et al. (2025). SA scores and diversity for MolCRAFT, MolCRAFT-large, RxnFlow, and 3DSynthFlow are recomputed from author-provided SMILES; SA scores and diversity for the remaining baselines are taken from Shen et al. (2023). Best results are in bold.

Category	Method	Validity (%, ↑)	Vina (kcal/mol, ↓)		QED (↑)		SA (↑)		AiZynth. Succ Rate (%, ↑)		Div (↑)	Time (s, ↓)
		Validity	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Avg.
	Reference	–	-7.71	-7.80	0.48	0.47	0.73	0.74	36.1	–	–	–
Atom	Pocket2Mol (Peng et al., 2022)	98.3	-7.60	-7.16	0.57	0.58	0.74	0.75	29.1	22.0	0.69	2504
	TargetDiff (Guan et al., 2023)	91.5	-7.37	-7.56	0.49	0.49	0.58	0.58	9.9	3.2	0.72	3428
	DecompDiff (Guan et al., 2024)	66.0	-8.35	-8.25	0.37	0.35	0.61	0.60	0.9	0.0	0.68	6189
	DiffSBDD (Schneuing et al., 2024)	76.0	-6.95	-7.10	0.47	0.48	0.55	0.56	2.9	2.0	**0.76**	135
	MolCRAFT (Qu et al., 2024)	96.7	-8.05	-8.05	0.50	0.50	0.69	0.67	16.5	9.1	0.73	141
	MolCRAFT-large (Qu et al., 2024)	70.8	-9.25	-9.24	0.45	0.44	0.62	0.62	3.9	0.0	0.60	>141
Fragment	TacoGFN (Shen et al., 2023)	**100.0**	-8.24	-8.44	0.67	0.67	0.79	0.79	1.3	1.0	0.53	**4**
2D Reaction	RxnFlow (Seo et al., 2024)	**100.0**	-8.85	-9.03	0.67	0.67	0.54	0.54	34.8	34.5	0.54	**4**
3D Reaction	3DSynthFlow (low ρ) (Shen et al., 2025)	99.9	-9.14	-9.38	0.69	0.69	0.66	0.66	36.2	37.0	0.53	6
	3DSynthFlow (med ρ) (Shen et al., 2025)	99.9	-9.30	-9.62	0.72	0.71	0.66	0.66	35.1	36.0	0.49	6
	3DSynthFlow (high ρ) (Shen et al., 2025)	**100.0**	-9.42	-9.61	**0.73**	**0.73**	0.65	0.65	36.1	36.0	0.45	6
	SYCLE (low ρ)	**100.0**	-9.41 (± 0.00)	-9.51 (± 0.04)	0.69 (± 0.00)	0.69 (± 0.00)	0.77 (± 0.00)	0.77 (± 0.00)	58.4 (± 0.64)	59.5 (± 1.32)	0.58 (± 0.00)	9
	SYCLE (med ρ)	**100.0**	-9.63 (± 0.01)	-9.80 (± 0.01)	0.70 (± 0.00)	0.70 (± 0.00)	0.79 (± 0.00)	0.79 (± 0.00)	64.4 (± 0.59)	64.8 (± 1.13)	0.55 (± 0.00)	9
	SYCLE (high ρ)	**100.0**	**-9.80** (± 0.01)	**-9.96** (± 0.04)	0.70 (± 0.00)	0.70 (± 0.00)	**0.81** (± 0.00)	**0.81** (± 0.00)	**68.8** (± 1.72)	**69.8** (± 1.04)	0.52 (± 0.00)	9
Ablation	- similarity weighting	100.0	-9.68	-9.84	0.69	0.69	0.79	0.79	65.6	66.5	0.54	9
	- logit augmentation	100.0	-9.61	-9.87	0.69	0.69	0.78	0.78	57.3	57.0	0.55	9
	- sim. weighting & logit aug.	100.0	-9.50	-9.70	0.70	0.70	0.78	0.78	53.1	53.0	0.55	9

Table 3 summarizes PoseCheck metrics. SYCLE demonstrates higher interaction scores overall compared to the baseline models, suggesting that its performance improvement reflects an enhancement in binding quality thanks to its contrastive learning-based conformer prediction.

4.2. Pocket-conditional Generation

Setup. We also evaluate SYCLE in a pocket-conditional generation setting, where a single policy generates ligands for previously unseen protein pockets without relying on extra oracle evaluations (Peng et al., 2022; Guan et al., 2023; Schneuing et al., 2024). We use the same pocket-conditional protocol and reward definition as 3DSynthFlow (Shen et al., 2025), where the TacoGFN (Shen et al., 2023) proxy trained on CrossDocked2020 is used to estimate the reward. Synthetic pathway construction follows the same protocol as in Section 4.1, except that a 300k Enamine stock subset is used as the building block library.

Evaluation is conducted on the CrossDocked2020 benchmark (Francoeur et al., 2020) using 100 test pockets, and each method generates 100 molecules per pocket. In addition to the metrics used in the pocket-specific setting, we report the following metrics: **validity** measures the fraction of valid and unique molecules, **diversity** is the mean pairwise Tanimoto distance computed over RDKit fingerprints (Landrum et al., 2013), and **time** refers to the wall-clock time to generate 100 molecules.

Baselines. We compare 3DSynthFlow with a diverse set of baselines covering atom-based, fragment-based, and reaction-based molecular generation. The atom-based baselines include Pocket2Mol (Peng et al., 2022), TargetDiff (Guan et al., 2023), DecompDiff (Guan et al., 2024), DiffSBDD (Schneuing et al., 2024), MolCRAFT, and MolCRAFT-large (Qu et al., 2024). We include TacoGFN (Shen et al., 2023) as a fragment-based method and RxnFlow (Seo et al., 2024) as a 2D reaction-based method. Following 3DSynthFlow (Shen et al., 2025), SYCLE is tested with three reward regimes: low $\rho \sim U(1, 64)$, medium $\rho \sim U(16, 64)$, and high $\rho \sim U(32, 64)$, where ρ controls the exploration-exploitation trade-off during sampling (see Section D.5 for further details).

Results. Table 4 summarizes the pocket-conditional generation results on the CrossDocked2020 benchmark. SYCLE achieves state-of-the-art Vina docking scores across all ρ settings, with SYCLE (high ρ) outperforming all baselines. Notably, SYCLE outperforms the baselines in synthesizability by a significantly large margin. Not only SYCLE achieves the highest SA score, it also shows the AiZynthFinder success rate of 69.8%, significantly outperforming previous methods, which had a maximum success rate of only 36.1%. These results demonstrate that SYCLE's synthon contrastive learning significantly improves the quality of generated molecules in terms of both binding affinity and retrosynthetic feasibility.

4.3. Ablation Study

We conduct ablation studies of SYCLE on the pocket-conditional generation benchmark with the high ρ setting.

Effect of similarity weighting. To examine the effect of the proposed similarity-based weighting, we compare SYCLE against a variant that treats all negative synthons as hard negatives, removing the similarity-based weight w in Eq. (18). As shown in Table 4, removing the similarity weighting degrades both the SA score and the AiZyn-

thFinder success rate, verifying the importance of the proposed soft contrastive learning strategy. This result suggests that chemically close negatives may represent plausible alternative synthons and should receive a weaker contrastive push rather than being uniformly penalized.

Effect of logit augmentation. To examine the effect of the logit augmentation for synthon selection, we compare SYCLE against a variant that does not use the compatibility score in the action logits (i.e., $\beta = 0$ in Eq. (20)), and relies solely on the GFlowNet policy to select synthons. As shown in Table 4, removing the logit augmentation leads to a notable drop in the AiZynthFinder success rate, verifying that injecting the compatibility score into the synthon selection procedure is critical for generating synthesizable molecules. This result suggests that the proposed compatibility score provides useful inductive biases that cannot be captured by the GFlowNet policy alone, thereby guiding pathway generation to identify synthetically feasible candidates.

5. Conclusion

In this work, we present SYCLE, a synthesizable SBDD framework that addresses the limitation of existing synthesis-aware 3D generation methods by injecting future synthon information into intermediate conformer modeling. By aligning intermediate 3D conformers with compatible next synthons via contrastive learning and incorporating the compatibility score into the pathway generation policy, SYCLE enables more coherent co-design of 3D molecular structures and synthetic pathways. Experiments show that SYCLE achieves state-of-the-art performance in both binding affinity and synthesizability, demonstrating its potential as a promising tool for practical and synthesizable SBDD.

Impact Statement

This work has the potential to accelerate early-stage drug discovery by generating molecules with high binding affinity and feasible synthetic routes, reducing the time and cost of identifying viable drug candidates. As with generative molecular design broadly, misuse for harmful purposes remains a societal risk that warrants responsible deployment. We also caution that practical synthesis success depends on additional factors not captured computationally, such as reaction conditions, reagent availability, and cost.

References

Adams, K. and Coley, C. W. Equivariant shape-conditioned generation of 3d molecules for ligand-based drug design. *arXiv preprint arXiv:2210.04893*, 2022.

Anderson, A. C. The process of structure-based drug design. *Chemistry & biology*, 10(9):787–797, 2003.

Bengio, Y., Lahlou, S., Deleu, T., Hu, E. J., Tiwari, M., and Bengio, E. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.

Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S., and Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98, 2012.

Bradshaw, J., Paige, B., Kusner, M. J., Segler, M., and Hernández-Lobato, J. M. A model to search for synthesizable molecules. *Advances in Neural Information Processing Systems*, 32, 2019.

Bradshaw, J., Paige, B., Kusner, M. J., Segler, M., and Hernández-Lobato, J. M. Barking up the right tree: an approach to search over molecule synthesis dags. *Advances in neural information processing systems*, 33:6852–6866, 2020.

Coley, C. W., Rogers, L., Green, W. H., and Jensen, K. F. Scscore: synthetic complexity learned from a reaction corpus. *Journal of chemical information and modeling*, 58(2):252–261, 2018.

Cretu, M., Harris, C., Igashov, I., Schneuing, A., Segler, M., Correia, B., Roy, J., Bengio, E., and Liò, P. Synflownet: Design of diverse and novel molecules with synthesis constraints. *arXiv preprint arXiv:2405.01155*, 2024.

Ertl, P. and Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1(1):8, 2009.

Francoeur, P. G., Masuda, T., Sunseri, J., Jia, A., Iovanisci, R. B., Snyder, I., and Koes, D. R. Three-dimensional convolutional neural networks and a cross-docked data set for structure-based drug design. *Journal of chemical information and modeling*, 60(9):4200–4215, 2020.

Gaiński, P., Boussif, O., Rekesh, A., Shevchuk, D., Parviz, A., Tyers, M., Batey, R. A., and Koziarski, M. Scalable and cost-efficient de novo template-based molecular generation. *arXiv preprint arXiv:2506.19865*, 2025.

Gao, B., Qiang, B., Tan, H., Jia, Y., Ren, M., Lu, M., Liu, J., Ma, W.-Y., and Lan, Y. Drugclip: Contrastive protein-molecule representation learning for virtual screening. *Advances in Neural Information Processing Systems*, 36: 44595–44614, 2023.

Gao, W. and Coley, C. W. The synthesizability of molecules proposed by generative models. *Journal of chemical information and modeling*, 60(12):5714–5723, 2020.

Gao, W., Mercado, R., and Coley, C. W. Amortized tree generation for bottom-up synthesis planning and synthesizable molecular design. *arXiv preprint arXiv:2110.06389*, 2021.

Gao, W., Luo, S., and Coley, C. W. Generative artificial intelligence for navigating synthesizable chemical space. *arXiv preprint arXiv:2410.03494*, 2024.

Genheden, S., Thakkar, A., Chadimová, V., Reymond, J.-L., Engkvist, O., and Bjerrum, E. Aizynthfinder: a fast, robust and flexible open-source software for retrosynthetic planning. *Journal of cheminformatics*, 12(1):70, 2020.

Gottipati, S. K., Sattarov, B., Niu, S., Pathak, Y., Wei, H., Liu, S., Blackburn, S., Thomas, K., Coley, C., Tang, J., et al. Learning to navigate the synthetically accessible chemical space using reinforcement learning. In *International conference on machine learning*, pp. 3668–3679. PMLR, 2020.

Grygorenko, O. O., Radchenko, D. S., Dziuba, I., Chuprina, A., Gubina, K. E., and Moroz, Y. S. Generating multibillion chemical space of readily accessible screening compounds. *Iscience*, 23(11), 2020.

Guan, J., Qian, W. W., Peng, X., Su, Y., Peng, J., and Ma, J. 3d equivariant diffusion for target-aware molecule generation and affinity prediction. *arXiv preprint arXiv:2303.03543*, 2023.

Guan, J., Zhou, X., Yang, Y., Bao, Y., Peng, J., Ma, J., Liu, Q., Wang, L., and Gu, Q. Decompdiff: diffusion models with decomposed priors for structure-based drug design. *arXiv preprint arXiv:2403.07902*, 2024.

Harris, C., Didi, K., Jamasb, A. R., Joshi, C. K., Mathis, S. V., Lio, P., and Blundell, T. Benchmarking generated poses: How rational is structure-based drug design with generative models? *arXiv preprint arXiv:2308.07413*, 2023.

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

Irwin, R., Tibo, A., Janet, J. P., and Olsson, S. Efficient 3d molecular generation with flow matching and scale optimal transport. In *ICML 2024 AI for Science Workshop*, 2024.

Jain, M., Raparthy, S. C., Hernández-Garcıa, A., Rector-Brooks, J., Bengio, Y., Miret, S., and Bengio, E. Multi-objective gflownets. In *International conference on machine learning*, pp. 14631–14653. PMLR, 2023.

Jocys, Z., Zhu, Z., Willems, H. M., and Farrahi, K. Synthformer: Equivariant pharmacophore-based generation of synthesizable molecules for ligand-based drug design. *Artificial Intelligence in the Life Sciences*, pp. 100148, 2025.

Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., and Krishnan, D. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020.

Koziarski, M., Rekesh, A., Shevchuk, D., van der Sloot, A., Gaiński, P., Bengio, Y., Liu, C., Tyers, M., and Batey, R. Rgfn: Synthesizable molecular generation using gflownets. *Advances in Neural Information Processing Systems*, 37:46908–46955, 2024.

Landrum, G. et al. Rdkit: A software suite for cheminformatics, computational chemistry, and predictive modeling. *Greg Landrum*, 8(31.10):5281, 2013.

Lee, K., Jang, J., Seo, S., Lim, J., and Kim, W. Y. Drug-likeness scoring based on unsupervised learning. *Chemical Science*, 13(2):554–565, 2022.

Lee, S., Kreis, K., Veccham, S. P., Liu, M., Reidenbach, D., Paliwal, S., Nie, W., and Vahdat, A. Exploring synthesizable chemical space with iterative pathway refinements. *International conference on learning representations*, 2026.

Liu, M., Luo, Y., Uchino, K., Maruhashi, K., and Ji, S. Generating 3d molecules for target protein binding. *arXiv preprint arXiv:2204.09410*, 2022.

Liu, T., Naderi, M., Alvin, C., Mukhopadhyay, S., and Brylinski, M. Break down in order to build up: decomposing small molecules for fragment-based drug design with e molfrag. *Journal of chemical information and modeling*, 57(4):627–631, 2017.

Lohmann, F., Allenspach, S., Atz, K., Schiebroek, C. C., Hiss, J. A., and Schneider, G. Protein binding site representation in latent space. *Molecular Informatics*, 44(1): e202400205, 2025.

Long, S., Zhou, Y., Dai, X., and Zhou, H. Zero-shot 3d drug design by sketching and generating. *Advances in Neural Information Processing Systems*, 35:23894–23907, 2022.

Luo, S., Guan, J., Ma, J., and Peng, J. A 3d generative model for structure-based drug design. *Advances in Neural Information Processing Systems*, 34:6229–6239, 2021.

Luo, Y. and Ji, S. An autoregressive flow model for 3d molecular geometry generation from scratch. In *International conference on learning representations*, 2022.

Malkin, N., Jain, M., Bengio, E., Sun, C., and Bengio, Y. Trajectory balance: Improved credit assignment in

gflownets. *Advances in Neural Information Processing Systems*, 35:5955–5967, 2022.

Neeser, R. M., Igashov, I., Schneuing, A., Bronstein, M. M., Schwaller, P., and Correia, B. Flow-based fragment identification via contrastive learning of binding site-specific latent representations. In *AI for Accelerated Materials Design-ICLR 2025*, 2025.

Peng, X., Luo, S., Guan, J., Xie, Q., Peng, J., and Ma, J. Pocket2mol: Efficient molecular sampling based on 3d protein pockets. In *International conference on machine learning*, pp. 17644–17655. PMLR, 2022.

Qu, Y., Qiu, K., Song, Y., Gong, J., Han, J., Zheng, M., Zhou, H., and Ma, W.-Y. Molcraft: structure-based drug design in continuous parameter space. *arXiv preprint arXiv:2404.12141*, 2024.

Rekesh, A., Cretu, M., Shevchuk, D., Somnath, V. R., Liò, P., Batey, R. A., Tyers, M., Koziarski, M., and Liu, C.-H. Syncogen: Synthesizable 3d molecule generation via joint reaction and coordinate modeling. *arXiv preprint arXiv:2507.11818*, 2025.

Schneuing, A., Harris, C., Du, Y., Didi, K., Jamasb, A., Igashov, I., Du, W., Gomes, C., Blundell, T. L., Lio, P., et al. Structure-based drug design with equivariant diffusion models. *Nature Computational Science*, 4(12): 899–909, 2024.

Seo, S., Lim, J., and Kim, W. Y. Molecular generative model via retrosynthetically prepared chemical building block assembly. *Advanced Science*, 10(8):2206674, 2023.

Seo, S., Kim, M., Shen, T., Ester, M., Park, J., Ahn, S., and Kim, W. Y. Generative flows on synthetic pathway for drug design. *arXiv preprint arXiv:2410.04542*, 2024.

Shen, T., Seo, S., Lee, G., Pandey, M., Smith, J. R., Cherkasov, A., Kim, W. Y., and Ester, M. Tacogfn: Target-conditioned gflownet for structure-based drug design. *arXiv preprint arXiv:2310.03223*, 2023.

Shen, T., Seo, S., Irwin, R., Didi, K., Olsson, S., Kim, W. Y., and Ester, M. Compositional flows for 3d molecule and synthesis pathway co-design. *arXiv preprint arXiv:2504.08051*, 2025.

Thakkar, A., Chadimová, V., Bjerrum, E. J., Engkvist, O., and Reymond, J.-L. Retrosynthetic accessibility score (rascore)–rapid machine learned synthesizability classification from ai driven retrosynthetic planning. *Chemical science*, 12(9):3339–3349, 2021.

Tran-Nguyen, V.-K., Jacquemard, C., and Rognan, D. Litpcba: an unbiased data set for machine learning and virtual screening. *Journal of chemical information and modeling*, 60(9):4263–4273, 2020.

Walters, W. P. and Barzilay, R. Critical assessment of ai in drug discovery. *Expert opinion on drug discovery*, 16(9): 937–947, 2021.

Yu, Y., Cai, C., Wang, J., Bo, Z., Zhu, Z., and Zheng, H. Unidock: Gpu-accelerated docking enables ultralarge virtual screening. *Journal of chemical theory and computation*, 19(11):3336–3345, 2023.

A. Limitations

Although SYCLE shows strong empirical performance, there are several limitations. First, the contrastive objective considers only the immediately following synthon, leaving the framework partially myopic to longer-range synthetic dependencies. Finally, the generation order is fixed such that a FIRSTSYNTHON action always initiates the trajectory, which may not reflect the optimal order for some target pockets. Allowing pocket-conditioned adaptive ordering is left as a future work.

B. Theoretical Background

B.1. Trajectory Balance with Deterministic State Transitions

Following Shen et al. (2025), when the initial conformation $\boldsymbol{S}_0^{(i)}$ for each newly added synthon is sampled under a fixed random seed, the state $x_{t_{\text{gen}}^{(i)}}$ is uniquely determined by the sequence $(\boldsymbol{C}^{(1)}, \ldots, \boldsymbol{C}^{(k(t))})$. This determinism allows the forward transition probability to be simplified as:

$$P_F\left(x_{t_{\text{gen}}^{(i+1)}} \mid x_{t_{\text{gen}}^{(i)}}, \hat{\boldsymbol{x}}^1_{t_{\text{gen}}^{(i)}}; \phi, p^{\psi}_{1|t}\right) = P_F\left(\boldsymbol{C}^{(i)} \mid x_{t_{\text{gen}}^{(i)}}, \hat{\boldsymbol{x}}^1_{t_{\text{gen}}^{(i)}}; \phi\right), \tag{21}$$

and the TB objective in Eq. (8) decomposes into a tractable form over the sequence of synthon selections.

B.2. Convergence of Compatibility-Augmented Trajectory Balance

SYCLE augments the GFlowNet action logits with the compatibility score (Eq. 20):

$$\ell_{\text{SYCLE}}\left(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})\right) = \ell\left(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})\right) + \beta \cdot \text{Comp}\left(\mathcal{S}_t, \boldsymbol{C}^{(i)}\right), \tag{22}$$

where ℓ is the learnable GFlowNet logit (Eq. (29)) and $\text{Comp}(\mathcal{S}_t, \boldsymbol{C}^{(i)})$ is the compatibility score (Eq. (14)) computed from the frozen conformer predictor $p^{\psi}_{1|t}$. The resulting forward policy conditioned on pocket $\mathcal{P}$ is:

$$P_F^{\text{SYCLE}}\left(\boldsymbol{C}^{(i)} \,\middle|\, \mathcal{S}_t, \mathcal{P}\right) = \frac{\exp\Big(\ell\big(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})\big) + \beta \cdot \text{Comp}\big(\mathcal{S}_t, \boldsymbol{C}^{(i)}\big)\Big)}{\displaystyle\sum_{\boldsymbol{C}' \in \mathcal{A}(\mathcal{C}_t)} \exp\Big(\ell(\mathcal{S}_t, (r, \boldsymbol{C}')) + \beta \cdot \text{Comp}(\mathcal{S}_t, \boldsymbol{C}')\Big)}. \tag{23}$$

We now show that the TB objective with P_F^{SYCLE} retains a valid fixed point where $p(x_1) \propto R(x_1)$.

Proposition. *Let $c\big(\mathcal{S}_t, \boldsymbol{C}^{(i)}\big) = \beta \cdot \text{Comp}\big(\mathcal{S}_t, \boldsymbol{C}^{(i)}\big)$ be any fixed, bounded function of the current state and candidate synthon. Then the TB objective in Eq. (8) with P_F^{SYCLE} achieves $\mathcal{L}_{\text{TB}} = 0$ for all trajectories τ, with the optimal learnable logit given by:*

$$\ell^*\left(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})\right) = \log P^*\left(\boldsymbol{C}^{(i)} \,\middle|\, \mathcal{S}_t, \mathcal{P}\right) - c\left(\mathcal{S}_t, \boldsymbol{C}^{(i)}\right) + \text{const}, \tag{24}$$

where P^* is the reward-proportional forward policy.

proof Fix any trajectory $\tau = (\boldsymbol{C}^{(1)}, \ldots, \boldsymbol{C}^{(n)})$. By Section B.1, the states $\mathcal{S}_{t_{\text{gen}}^{(i)}}$ are uniquely determined by τ independently of ϕ. Therefore, $c(\mathcal{S}_t, \boldsymbol{C}^{(i)})$ is independent of ϕ along any fixed trajectory.

For any target distribution $P^*\big(\boldsymbol{C}^{(i)} \mid \mathcal{S}_t, \mathcal{P}\big)$ satisfying $\prod_i P^*\big(\boldsymbol{C}^{(i)}\big) \propto R(x_1)$, we can always find ℓ^* such that $\text{softmax}(\ell^* + c) = P^*$ by setting:

$$\ell^*\left(\mathcal{S}_t, (r, \boldsymbol{C}^{(i)})\right) = \log P^*\left(\boldsymbol{C}^{(i)} \mid \mathcal{S}_t, \mathcal{P}\right) - c\left(\mathcal{S}_t, \boldsymbol{C}^{(i)}\right) + \text{const},$$

since $c(\cdot)$ cancels in the numerator and denominator of the softmax:

$$\text{softmax}(\ell^* + c)_j = \frac{\exp(\log P^*_j - c_j + c_j)}{\sum_k \exp(\log P^*_k - c_k + c_k)} = \frac{P^*_j}{\sum_k P^*_k} = P^*_j.$$

Setting $Z_\phi = \sum_{x_1} R(x_1)$ then gives $Z_\phi \prod_i P_F^{\text{SYCLE}}(\boldsymbol{C}^{(i)}) = R(x_1)$ for all trajectories τ, and hence $\mathcal{L}_{\text{TB}} = 0$.

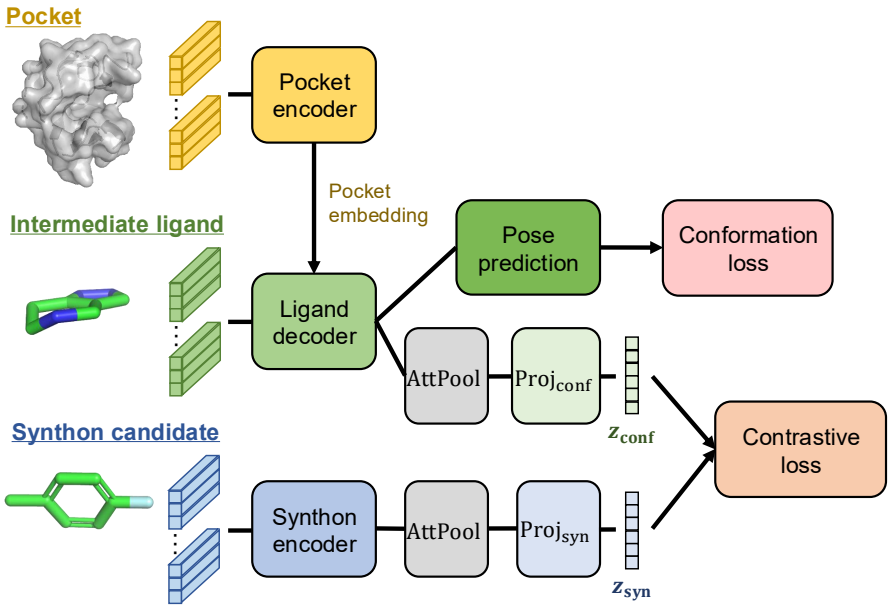

Figure 3. **The overall architecture of SYCLE.**

C. SYCLE Architecture

SYCLE consists of a conformer predictor $p^{\psi}_{1|t}$ and a pathway generation policy π_{ϕ}. The conformer predictor internally integrates a synthon encoder to learn the compatibility between intermediate conformers and candidate next synthons. Figure 3 illustrates the training architecture of the conformer predictor, which is jointly optimized with the conformation loss $\mathcal{L}_{\text{conf}}$ and the contrastive loss $\mathcal{L}_{\text{cont}}$ (Eq. (19)).

C.1. Conformer Predictor

The conformer predictor consists of a POCKETENCODER and a LIGANDDECODER. The POCKETENCODER produces pocket embeddings that condition the LIGANDDECODER, which in turn outputs both the predicted conformer and the conformer embedding z_{conf}(Eq. (12)) used for contrastive learning.

Pocket encoder. The protein pocket is encoded using *HierSemla*, a hierarchical extension of Semla (Irwin et al., 2024) introduced in Shen et al. (2025). Rather than processing all protein atoms jointly, HierSemla first applies intra-residue attention to obtain per-atom invariant and equivariant features, then collapses them into a single residue-level representation $(\boldsymbol{h}_i^{(\text{pro})}, \boldsymbol{x}_i^{(\text{pro})})$ through mean pooling. Limiting self-attention to atoms within the same residue makes the encoder scalable to full-atom protein pockets without prohibitive memory overhead.

Ligand decoder. The partial ligand is encoded by a sequence of Semla layers (Shen et al., 2025) that incorporate pocket context via cross-attention, yielding atom-level invariant and equivariant features $\boldsymbol{h}^{(\text{lig})}$. At each layer, cross-attention messages are computed over all ligand atom–pocket residue pairs $(i, j) \in \mathcal{V} \times \mathcal{P}$, where $\mathcal{V}$ and $\mathcal{P}$ denote the sets of ligand atoms and protein residues, respectively:

$$(\boldsymbol{m}_{i,j}^{\text{inv}}, \boldsymbol{m}_{i,j}^{\text{equi}}) = \Omega_{\theta}^{(\text{pro-lig})}\left(\tilde{h}_i, \tilde{h}_j, \tilde{x}_i \cdot \tilde{x}_j\right), \tag{25}$$

where $\tilde{h}_i$ and $\tilde{x}_i$ are the layer-normalized invariant and coordinate-normalized equivariant projections of node i, respectively. The cross-attention messages are combined with intra-ligand messages before the subsequent feature update, allowing each ligand atom to incorporate structural context from the surrounding binding site. Each atom is additionally annotated with a binary flag `IS_ATTACHMENT_POINT` indicating the reactive site at which the next synthon will be attached. This provides explicit geometric information for the conformer predictor to reserve sufficient space for the incoming synthon at the attachment site. The predicted clean coordinates $\hat{\boldsymbol{S}}_1^{(i)}$ are obtained by normalizing and projecting the equivariant

component of $\boldsymbol{h}^{(\text{lig})}$.

Synthon encoder. The SYNTHONENCODER processes a synthon $\boldsymbol{C}$ through Semla layers without pocket conditioning to produce atom-wise features. Given atom types, charges, adjacency, and 3D coordinates of the synthon, the encoder computes invariant and equivariant features as:

$$\left(\boldsymbol{h}^{(\text{syn})}, \boldsymbol{x}^{(\text{syn})}\right) = \text{SYNTHONENCODER}(\boldsymbol{C}), \tag{26}$$

where $\boldsymbol{h}^{(\text{syn})} \in \mathbb{R}^{L \times d_{\text{inv}}}$ and $\boldsymbol{x}^{(\text{syn})} \in \mathbb{R}^{L \times 3 \times d_{\text{equi}}}$ denote the invariant and equivariant features per atom respectively, with equivariant features defined over the three spatial directions. These dimensions are kept identical to those of the LIGANDDECODER to ensure compatibility in the shared embedding space. Two key differences distinguish it from the LIGANDDECODER. First, it uses no pocket conditioning, as synthons are encoded solely based on their intrinsic chemical structure. Second, it employs a lighter layer stack of 3 layers compared to the 8 layers of the LIGANDDECODER. The SYNTHONENCODER also includes an attachment point embedding $\boldsymbol{e}_{\text{att}}$ identical to that in the LIGANDDECODER, providing explicit information about the attachment site without time or size conditioning.

Projection heads. Both the conformer embedding $\boldsymbol{z}_{\text{conf}}$ and the synthon embedding $\boldsymbol{z}_{\text{syn}}$ are projected into the shared contrastive space from a common feature dimension. Before projection, atom-wise invariant features $\boldsymbol{h}^{(\cdot)} \in \mathbb{R}^{d_{\text{inv}}}$ and equivariant features $\boldsymbol{x}^{(\cdot)} \in \mathbb{R}^{3 \times d_{\text{equi}}}$ are pooled via ATTPOOL and concatenated into a single vector of dimension $d_{\text{inv}} + 3d_{\text{equi}}$. Both $\text{Proj}_{\text{conf}}$ and Proj_{syn} share the same architecture:

$$\text{Proj}(\boldsymbol{f}) = W_2 \, \text{ReLU}(\text{LayerNorm}(W_1 \boldsymbol{f})), \tag{27}$$

where $W_1 \in \mathbb{R}^{d_h \times (d_{\text{inv}} + 3d_{\text{equi}})}$ and $W_2 \in \mathbb{R}^{d \times d_h}$, with d_h denoting the hidden dimension of the projection head.

C.2. Pathway Generation Policy

The pathway generation policy π_ϕ reuses the architecture of 3DSynthFlow (Shen et al., 2025), with a graph transformer as backbone f_θ and an MLP g_θ for action embedding. The protein pocket is modeled as a residue-level graph whose nodes correspond to C_α atoms, with edges connecting every pair of residues. Each ligand atom attends to all protein nodes via cross-edges, and pairwise distances are projected into edge feature vectors to provide geometric context during message passing.

At each step of the synthesis trajectory, the policy chooses between two action types, conditioned on temperature β and multi-objective weights following Shen et al. (2025).

For FIRSTSYNTHON, the logit for selecting $C^{(0)} \in \mathcal{B}_s$ is computed as:

$$\ell = \text{MLP}_{\text{FIRSTSYNTHON}}\big(f_\theta(\mathcal{S}^{(0)}, \mathcal{P})\big) \odot g_\theta(C^{(0)}), \tag{28}$$

where f_θ encodes the current state and protein pocket $\mathcal{P}$, and $g_\theta(C^{(0)})$ is the synthon embedding.

For each subsequent ADDSYNTHON step, the logit for selecting $(r, C^{(i)}) \in \mathcal{R}(\mathcal{S}_t) \times \mathcal{B}_b$, where $\mathcal{R}(\mathcal{S}_t)$ denotes the set of reactions applicable to the current state $\mathcal{S}_t$, is:

$$\ell = \text{MLP}_{\text{ADDSYNTHON}}\big(f_\theta(\mathcal{S}_t, \mathcal{P})\big) \odot g_\theta(r, C^{(i)}), \tag{29}$$

where the reaction type r is jointly encoded with the synthon embedding by $g_\theta(r, C^{(i)})$.The full definition of the action space is provided in Appendix D.1.

The compatibility score augmentation (Eq. (20)) is applied only to ADDSYNTHON logits, as FIRSTSYNTHON operates without an existing intermediate conformation. The conformer embedding $\boldsymbol{z}_{\text{conf}}(\mathcal{S}_t)$ from the frozen conformer predictor is passed as an additional input to evaluate the compatibility score at each ADDSYNTHON step.

D. Experimental Details

D.1. Action Space

Following Shen et al. (2025), we model each synthesis step as a forward transition in the GFlowNet trajectory, so that a complete synthetic pathway corresponds to a full generation trajectory. Rather than treating each building block as a fully

formed molecule, we adopt the synthon representation of Liu et al. (2017), in which each fragment carries one or two reactive sites that define where the next reaction can occur. A fragment with one reactive site is a *brick*: once selected, no further extension is possible and the trajectory ends. A fragment with two reactive sites is a *linker*: it introduces an additional reactive site, allowing the assembly to continue. Valid connections between fragments are governed by the reaction rule set $\mathcal{R}$.

The building block library $\mathcal{B}$ is the complete synthon set, with $\mathcal{B}' \subseteq \mathcal{B}$ denoting the brick synthon set and $\mathcal{B}_r \subseteq \mathcal{B}$ the set of synthons compatible with reaction r. The allowable action space $\mathcal{A}(\mathcal{C})$ at compositional state $\mathcal{C}$ is defined as:

$$\mathcal{A}(\mathcal{C}) = \begin{cases} \mathcal{B}' & \text{if } t = 1, \\ \bigcup_{r \in \mathcal{R}(\mathcal{C})} \mathcal{B}_r & \text{otherwise,} \end{cases} \tag{30}$$

where $\mathcal{R}(\mathcal{C}) \subseteq \mathcal{R}$ denotes the subset of reactions applicable to the current state. Generation is capped at two synthesis steps, consistent with the tractable range of Enamine REAL Space (Grygorenko et al., 2020).

D.2. Reward Function

Pocket-specific optimization. Following Seo et al. (2024), the reward is a weighted sum of drug-likeness and binding affinity:

$$R(x) = w_1 \mathrm{QED}(x) + w_2 \widehat{\mathrm{Vina}}(x), \tag{31}$$

where $\widehat{\mathrm{Vina}}(x)$ denotes the normalized Vina score. The weights w_1 and w_2 act as multi-objective conditioning weights (Jain et al., 2023). Also, for baselines that do not use GFlowNet, it is set to 0.5.

Pocket-conditional generation. We build on the reward formulation of Seo et al. (2024), dropping the SA score term from Shen et al. (2023) because synthesizability is enforced structurally through the pathway generation process. The affinity component r_{affinity} rescales the proxy docking score via a piecewise linear mapping. r_{QED} saturates at a QED of 0.7. The final reward is divided by the cube root of heavy atom count to offset the size bias in docking scores, discouraging the policy from inflating molecular size as a shortcut to higher reward.

$$r_{\text{affinity}}(x) = \begin{cases} 0 & \text{if } 0 \leq \mathrm{Proxy}(x) \\ -0.04 \times \mathrm{Proxy}(x) & \text{if } -8 \leq \mathrm{Proxy}(x) \leq 0 \\ -0.2 \times \mathrm{Proxy}(x) - 1.28 & \text{if } -13 \leq \mathrm{Proxy}(x) \leq -8 \\ 1.32 & \text{if } \mathrm{Proxy}(x) \leq -13 \end{cases} \tag{32}$$

$$r_{\text{QED}}(x) = \begin{cases} \mathrm{QED}(x)/0.7 & \text{if } \mathrm{QED}(x) \leq 0.7 \\ 1 & \text{otherwise} \end{cases} \tag{33}$$

$$\text{SYCLE-Reward}(x) = \frac{r_{\text{affinity}}(x) \times r_{\text{QED}}(x)}{\sqrt[3]{\mathrm{HeavyAtomCounts}(x)}} \tag{34}$$

D.3. Datasets

Building block library. We use 38 bimolecular Enamine synthesis protocols for constructing synthetic pathways (Gao et al., 2024). For pocket-specific optimization, we use the Enamine comprehensive catalog (v2024.06.10) containing approximately 1.2M building blocks after filtering by drug-likeness score (Lee et al., 2022) and chemical validity. For pocket-conditional generation, we use a 300k Enamine stock set with the same filtering applied.

LIT-PCBA pockets. The LIT-PCBA benchmark (Tran-Nguyen et al., 2020), comprising 15 protein targets with experimentally validated active and inactive compounds, is used for pocket-specific optimization. Each binding pocket is constructed by collecting all residues containing heavy atoms within 10 Å of the co-crystallized ligand, consistent with the pocket extraction procedure used for CrossDocked2020. The PDB entries and target names are listed in Table 5.

Table 5. **Target information for the LIT-PCBA dataset.**

Target	PDB Id	Target name
ADRB2	4ldo	Beta2 adrenoceptor
ALDH1	5l2m	Aldehyde dehydrogenase 1
ESR_ago	2p15	Estrogen receptor α with agonist
ESR_antago	2iok	Estrogen receptor α with antagonist
FEN1	5fv7	FLAP Endonuclease 1
GBA	2v3d	Acid Beta-Glucocerebrosidase
IDH1	4umx	Isocitrate dehydrogenase 1
KAT2A	5h86	Histone acetyltransferase KAT2A
MAPK1	4zzn	Mitogen-activated protein kinase 1
MTORC1	4dri	PPIase domain of FKBP51, Rapamycin
OPRK1	6b73	Kappa opioid receptor
PKM2	4jpg	Pyruvate kinase muscle isoform M1/M2
PPARG	5y2t	Peroxisome proliferator-activated receptor γ
TP53	3zme	Cellular tumor antigen p53
VDR	3a2i	Vitamin D receptor

Table 6. **Hyperparameters for the conformer predictor.**

Hyperparameter	**Value**
Architecture	
Number of protein layers	4
Number of ligand layers	8
Noise prior	Gaussian
Time distribution	Beta(1.0, 1.0)
Time Scheduling	
Time per action	0.3
Interpolation window	1.0
Maximum decomposed parts	3
Minimum decomposed fragment size (atoms)	5
Minimum trajectory length	2
Maximum trajectory length	3

CrossDocked2020. CrossDocked2020 (Francoeur et al., 2020) is a large-scale dataset of protein-ligand complexes constructed by cross-docking ligands into non-native protein structures. Following the data split of Luo et al. (2021) and Peng et al. (2022), we used 99,000 complexes for training and 1,000 for validation, excluding all LIT-PCBA targets from the training set. The conformer predictor and the pocket-conditional policy are both trained on this split, with the 100 test pockets reserved exclusively for evaluating pocket-conditional generation.

D.4. Baselines

Pocket-specific optimization. We used the baseline results reported in Shen et al. (2025).

Pocket-conditional generation. We used the baseline results as detailed in the caption of Table 4.

D.5. Hyperparameters

Conformer predictor hyperparameters. We describe the hyperparameters used for the conformer predictor as below. All other architectural hyperparameters follow the default settings of SemlaFlow (Irwin et al., 2024).

Table 7. **Hyperparameters for the pathway generation policy.**

Hyperparameter	Value
GFlowNet Training	
GFN temperature (ρ)	Uniform(1, 64)
Sampling tau (pocket-specific)	0.9
Sampling tau (pocket-conditional)	0.0
Learning rate (Z)	10^{-2}
Learning rate (P_F)	10^{-4}
Model	
State embedding dimension	64
Action embedding dimension	64
Compatibility score coefficient β	10.0
Action Space	
Minimum trajectory length	2 (min. reaction steps: 1)
Maximum trajectory length	3 (max. reaction steps: 2)
Action space subsampling ratio	1%
Train random action probability	5%

1. **Number of protein layers**: Number of HierSemla layers for protein message passing within each residue.
2. **Number of ligand layers**: Number of Semla layers for ligand decoding with protein-ligand cross-attention.
3. **Noise prior**: Distribution from which the initial conformation $\mathbf{S}_0^{(i)}$ is sampled.
4. **Time distribution**: Distribution used to sample the global time t during training.
5. **Time per action**: Time interval per synthon addition, setting the attachment time as $t_{\text{gen}}^{(i)} = 0.3 \cdot (i-1)$, following 3DSynthFlow (Shen et al., 2025).
6. **Interpolation window**: Time window over which conformation refinement is performed. With $t_{\text{window}} = 1.0$, all synthon conformations are refined jointly through $t = 1$.
7. **Maximum decomposed parts**: Maximum number of synthons per molecule during data preparation.
8. **Minimum decomposed fragment size**: Minimum per-synthon atom count when fragmenting training ligands, filtering out trivially small fragments unlikely to appear in real synthesis steps.
9. **Minimum / Maximum trajectory length**: Lower and upper bounds on synthon attachment steps during generation, corresponding to 1–2 reaction steps.

Pathway generation policy hyperparameters. We define the hyperparameters used for the pathway generation policy below.

1. **GFN temperature (ρ)**: The reward exponentiation parameter (Jain et al., 2023) that reshapes the sampling distribution as $\pi(x) \propto R(x)^{\rho}$. Sampling ρ from a wider range (e.g., $\mathcal{U}(1, 64)$) encourages broader exploration, while restricting ρ to a narrower high-value range (e.g., $\mathcal{U}(32, 64)$) concentrates sampling toward high-reward molecules.
2. **Sampling tau (EMA factor)**: The exponential moving average factor applied to the policy parameters during data collection. For pocket-specific optimization, $\tau = 0.9$ is used to stabilize training over long optimization runs. For pocket-conditional generation, $\tau = 0.0$ is used to ensure the policy adapts rapidly across diverse pockets.
3. **Learning rate (Z)**: The learning rate for updating the partition function Z in the TB objective.

Table 8. **Hyperparameters for the contrastive learning objective.**

Hyperparameter	Value
Contrastive loss weight λ	0.1
Similarity temperature τ	0.3
Number of synthon layers	3
Projection dimension	256

4. **Learning rate (P_F)**: The learning rate for updating the forward policy P_F.
5. **State embedding dimension**: The hidden dimension of the graph transformer used to encode the protein-ligand complex state.
6. **Action embedding dimension**: The hidden dimension of the MLP used to embed each candidate synthon action.
7. **Compatibility score coefficient** β: The scaling coefficient applied to the compatibility score when augmenting the AddSynthon logits (Eq. (20)).
8. **Minimum / Maximum trajectory length**: The minimum and maximum number of synthon attachment steps during generation, corresponding to 1-2 reaction steps.
9. **Action space subsampling ratio**: The fraction of candidate synthons randomly subsampled for each FirstSynthon and bi-molecular reaction template at each step. A minimum subsampling size of 100 is enforced for reactions with small reactant block sets to avoid high variance.
10. **Train random action probability**: The probability of selecting a uniformly random action during training, ensuring that rare action categories are sufficiently explored.

Contrastive learning We define the hyperparameters used for the contrastive learning objective below.

1. **Contrastive loss weight** λ: The weighting coefficient for $\mathcal{L}_{\text{cont}}$ in the joint training objective of Eq. (19).
2. **Similarity temperature** τ: The temperature parameter in the contrastive loss that controls the sharpness of the similarity distribution. A lower temperature makes the loss more sensitive to differences between positive and negative pairs.
3. **Number of synthon encoder layers**: The number of layers in the SynthonEncoder used to compute synthon embeddings z_{syn}.
4. **Projection dimension**: The output dimension of the projection heads $\text{Proj}_{\text{conf}}$ and Proj_{syn}, which map conformer and synthon features into the shared contrastive embedding space.

D.6. Training Details

Conformer predictor training. The conformer predictor $p_\psi^{1|t}$ is initialized from the pretrained pose predictor of 3DSynthFlow (Shen et al., 2025) and fine-tuned on the CrossDocked dataset (Francoeur et al., 2020) for 30 epochs by minimizing the joint objective (Eq. (19)). All parameters, including the backbone, are jointly optimized during fine-tuning. Each CrossDocked ligand is fragmented into at most three synthons by applying 38 bimolecular Enamine reaction SMARTS, allowing the conformer predictor to train on intermediate ligand structures at different points along the synthesis trajectory. Training samples are constructed by drawing a random assembly order and a random global time t for each molecule, yielding snapshots of the ligand at different stages of incompletion. This curriculum trains the model to produce geometrically consistent binding poses even when only a subset of the final synthons has been attached.

Pocket-specific training. Following 3DSynthFlow (Shen et al., 2025), the conformer predictor is frozen after Stage 1, and the pathway generation policy π_ϕ is trained online for 1,000 iterations using the TB objective with a batch size of 64.

Table 9. **Summary of training time for SYCLE components.**

Model Component	Hardware	Batch Size	Iterations	Training Time
Conformer Predictor	4×RTX A5000 (24GB)	dynamic	30 epochs	∼80 hours
Policy (Pocket-Specific)	1×RTX A5000 (24GB)	64	1,000 steps	18–24 hours
Policy (Pocket-Conditional)	1×RTX A5000 (24GB)	64	–	up to 72 hours

Pocket-conditional training. For pocket-conditional generation, the policy is trained following the same protocol as 3DSynthFlow (Shen et al., 2025), the conformer predictor is frozen and the pathway generation policy π_ϕ is trained for pocket-conditional generation. The checkpoint is selected based on sampled reward and training loss monitored on 50 randomly held-out pockets from the training set, with training running for up to 72 hours on a single NVIDIA RTX A5000 (24GB) GPU.

Computational cost. Table 9 summarizes the training time for each component. The conformer predictor is fine-tuned on 4×NVIDIA RTX A5000 (24GB) GPUs, requiring approximately 80 hours. As the conformer predictor is trained once and shared across all targets, no per-target retraining is required. The policy is trained on a single NVIDIA RTX A5000 (24GB) GPU, taking 18–24 hours per target depending on the protein pocket.

E. Additional Experimental Results

E.1. Pocket-specific results

We report the full results across all 15 LIT-PCBA targets in Table 10, Table 11, Table 12, and Table 13. SYCLE consistently outperforms all baselines in docking score and ligand efficiency across most of the targets.

E.2. Visualization of sample ligand generation

Figure 4 illustrates how the predicted final binding conformer $\hat{x}_1$ evolves as synthesis progresses in SYCLE, evaluated against the ALDH1 target. At each autoregressive step, a new synthon is added to the growing molecule (shown in pink in the synthesis pathway), and the flow matching model re-estimates the final 3D binding conformer conditioned on all synthons assembled so far. The predicted conformer progressively refines with each additional synthon, converging to the final conformation at $t = 1.0$. The bottom row shows the corresponding synthesis pathway.

Table 10. **Average Vina docking score.** Results for top-100 diverse molecules across all 15 LIT-PCBA targets. The best results are in bold.

Category	Method	Average Vina Docking Score (kcal/mol, ↓)				
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Fragment	FragGFN	-10.19 (± 0.33)	-10.43 (± 0.29)	-9.81 (± 0.09)	-9.85 (± 0.13)	-7.67 (± 0.71)
	FragGFN+SA	-9.70 (± 0.61)	-9.83 (± 0.65)	-9.27 (± 0.95)	-10.06 (± 0.30)	-7.26 (± 0.10)
Reaction	SynNet	-8.03 (± 0.26)	-8.81 (± 0.21)	-8.88 (± 0.13)	-8.52 (± 0.16)	-6.36 (± 0.09)
	BBAR	-9.95 (± 0.04)	-10.06 (± 0.14)	-9.97 (± 0.03)	-9.92 (± 0.05)	-6.84 (± 0.07)
	SynFlowNet[a]	-10.85 (± 0.10)	-10.69 (± 0.09)	-10.44 (± 0.05)	-10.27 (± 0.04)	-7.47 (± 0.02)
	SynFlowNet[b]	-9.17 (± 0.68)	-9.37 (± 0.29)	-9.17 (± 0.12)	-9.05 (± 0.14)	-6.45 (± 0.13)
	RGFN	-9.84 (± 0.21)	-9.93 (± 0.11)	-9.99 (± 0.11)	-9.72 (± 0.14)	-6.92 (± 0.06)
	RxnFlow	-11.45 (± 0.05)	-11.26 (± 0.07)	-11.15 (± 0.02)	-10.77 (± 0.04)	-7.66 (± 0.02)
3D Reaction	3DSynthFlow	-11.97 (± 0.12)	-11.82 (± 0.04)	-11.58 (± 0.07)	-11.23 (± 0.08)	-7.79 (± 0.01)
	SYCLE	**-12.10** (± 0.09)	**-11.98** (± 0.08)	**-11.64** (± 0.21)	**-11.45** (± 0.05)	**-8.16** (± 0.10)
		GBA	IDH1	KAT2A	MAPK1	MTORC1
Fragment	FragGFN	-8.76 (± 0.46)	-9.91 (± 0.32)	-9.27 (± 0.20)	-8.93 (± 0.18)	-10.51 (± 0.31)
	FragGFN+SA	-8.92 (± 0.27)	-9.76 (± 0.64)	-9.14 (± 0.43)	-8.28 (± 0.40)	-10.14 (± 0.30)
Reaction	SynNet	-7.60 (± 0.09)	-8.74 (± 0.08)	-7.64 (± 0.38)	-7.33 (± 0.14)	-9.30 (± 0.45)
	BBAR	-8.70 (± 0.05)	-9.84 (± 0.09)	-8.54 (± 0.06)	-8.49 (± 0.07)	-10.07 (± 0.16)
	SynFlowNet[a]	-9.27 (± 0.06)	-10.40 (± 0.08)	-9.41 (± 0.04)	-8.92 (± 0.05)	-10.84 (± 0.03)
	SynFlowNet[b]	-8.28 (± 0.15)	-9.18 (± 0.35)	-8.06 (± 0.15)	-7.89 (± 0.13)	-9.60 (± 0.16)
	RGFN	-8.48 (± 0.06)	-9.49 (± 0.13)	-8.53 (± 0.11)	-8.22 (± 0.15)	-9.89 (± 0.06)
	RxnFlow	-9.62 (± 0.04)	-10.95 (± 0.05)	-9.73 (± 0.03)	-9.30 (± 0.01)	-11.39 (± 0.09)
3D Reaction	3DSynthFlow	-9.90 (± 0.14)	-11.28 (± 0.15)	-10.17 (± 0.37)	-9.61 (± 0.11)	-11.91 (± 0.01)
	SYCLE	**-10.05** (± 0.01)	**-11.69** (± 0.08)	**-10.43** (± 0.09)	**-9.76** (± 0.05)	**-12.16** (± 0.03)
		OPRK1	PKM2	PPARG	TP53	VDR
Fragment	FragGFN	-10.28 (± 0.15)	-11.24 (± 0.27)	-9.54 (± 0.12)	-7.90 (± 0.02)	-10.96 (± 0.06)
	FragGFN+SA	-9.58 (± 0.44)	-10.83 (± 0.34)	-9.19 (± 0.29)	-7.61 (± 0.27)	-10.66 (± 0.61)
Reaction	SynNet	-8.70 (± 0.36)	-9.55 (± 0.14)	-7.47 (± 0.34)	-5.34 (± 0.23)	-10.98 (± 0.57)
	BBAR	-9.84 (± 0.10)	-11.39 (± 0.08)	-8.69 (± 0.10)	-7.05 (± 0.09)	-11.07 (± 0.04)
	SynFlowNet[a]	-10.34 (± 0.07)	-11.98 (± 0.12)	-9.40 (± 0.05)	-7.90 (± 0.10)	-11.62 (± 0.13)
	SynFlowNet[b]	-9.36 (± 0.25)	-10.64 (± 0.19)	-8.25 (± 0.10)	-6.84 (± 0.06)	-10.32 (± 0.07)
	RGFN	-9.61 (± 0.11)	-10.96 (± 0.18)	-8.53 (± 0.07)	-7.07 (± 0.06)	-10.86 (± 0.11)
	RxnFlow	-10.84 (± 0.03)	-12.53 (± 0.02)	-9.73 (± 0.02)	-8.09 (± 0.06)	-12.30 (± 0.07)
3D Reaction	3DSynthFlow	-11.26 (± 0.41)	-13.36 (± 0.03)	-10.00 (± 0.04)	-8.41 (± 0.17)	-12.98 (± 0.10)
	SYCLE	**-11.58** (± 0.06)	**-13.55** (± 0.13)	**-10.21** (± 0.12)	**-9.09** (± 0.03)	**-13.29** (± 0.14)

Table 11. **Average ligand efficiency.** Results for top-100 diverse molecules across all 15 LIT-PCBA targets. The best results are in bold.

Category	Method	Average Ligand Efficiency (↓)				
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Fragment	FragGFN	0.410 (± 0.006)	0.368 (± 0.007)	0.347 (± 0.003)	0.358 (± 0.002)	0.246 (± 0.004)
	FragGFN+SA	0.406 (± 0.007)	0.374 (± 0.023)	0.369 (± 0.003)	0.345 (± 0.024)	0.210 (± 0.004)
Reaction	SynNet	0.274 (± 0.041)	0.272 (± 0.006)	0.317 (± 0.005)	0.289 (± 0.020)	0.196 (± 0.003)
	BBAR	0.412 (± 0.006)	0.401 (± 0.008)	0.380 (± 0.001)	0.387 (± 0.003)	0.257 (± 0.003)
	SynFlowNet[a]	0.401 (± 0.006)	0.380 (± 0.007)	0.361 (± 0.003)	0.361 (± 0.004)	0.247 (± 0.004)
	SynFlowNet[b]	0.380 (± 0.021)	0.362 (± 0.016)	0.351 (± 0.008)	0.349 (± 0.005)	0.234 (± 0.007)
	RGFN	0.393 (± 0.005)	0.357 (± 0.004)	0.346 (± 0.002)	0.344 (± 0.002)	0.241 (± 0.001)
	RxnFlow	0.412 (± 0.005)	0.396 (± 0.005)	0.375 (± 0.002)	0.380 (± 0.004)	0.246 (± 0.001)
3D Reaction	3DSynthFlow	0.448 (± 0.019)	0.395 (± 0.006)	0.391 (± 0.005)	0.398 (± 0.016)	**0.252** (± 0.005)
	SYCLE	**0.460** (± 0.010)	**0.409** (± 0.010)	**0.399** (± 0.006)	**0.410** (± 0.007)	0.247 (± 0.002)
		GBA	IDH1	KAT2A	MAPK1	MTORC1
Fragment	FragGFN	0.333 (± 0.018)	0.367 (± 0.009)	0.322 (± 0.008)	0.302 (± 0.002)	0.354 (± 0.005)
	FragGFN+SA	0.318 (± 0.005)	0.369 (± 0.020)	0.298 (± 0.020)	0.294 (± 0.015)	0.355 (± 0.027)
Reaction	SynNet	0.244 (± 0.013)	0.281 (± 0.016)	0.294 (± 0.042)	0.226 (± 0.004)	0.316 (± 0.035)
	BBAR	0.336 (± 0.002)	0.382 (± 0.005)	0.332 (± 0.003)	0.320 (± 0.002)	0.385 (± 0.004)
	SynFlowNet[a]	0.330 (± 0.004)	0.368 (± 0.002)	0.327 (± 0.003)	0.305 (± 0.002)	0.368 (± 0.002)
	SynFlowNet[b]	0.324 (± 0.007)	0.360 (± 0.013)	0.309 (± 0.004)	0.297 (± 0.011)	0.361 (± 0.009)
	RGFN	0.310 (± 0.002)	0.351 (± 0.003)	0.310 (± 0.003)	0.298 (± 0.002)	0.346 (± 0.004)
	RxnFlow	0.327 (± 0.004)	0.378 (± 0.001)	0.330 (± 0.001)	0.313 (± 0.001)	0.370 (± 0.001)
3D Reaction	3DSynthFlow	0.340 (± 0.018)	0.388 (± 0.007)	0.340 (± 0.016)	0.322 (± 0.009)	0.387 (± 0.008)
	SYCLE	**0.346** (± 0.002)	**0.411** (± 0.011)	**0.344** (± 0.007)	**0.324** (± 0.004)	**0.393** (± 0.007)
		OPRK1	PKM2	PPARG	TP53	VDR
Fragment	FragGFN	0.352 (± 0.004)	0.442 (± 0.008)	0.319 (± 0.007)	0.307 (± 0.005)	0.394 (± 0.006)
	FragGFN+SA	0.327 (± 0.014)	0.440 (± 0.009)	0.303 (± 0.013)	0.248 (± 0.025)	0.390 (± 0.020)
Reaction	SynNet	0.298 (± 0.039)	0.296 (± 0.005)	0.253 (± 0.031)	0.211 (± 0.031)	0.359 (± 0.015)
	BBAR	0.370 (± 0.006)	0.442 (± 0.004)	0.326 (± 0.007)	0.288 (± 0.005)	0.409 (± 0.002)
	SynFlowNet[a]	0.359 (± 0.004)	0.427 (± 0.003)	0.317 (± 0.002)	0.287 (± 0.008)	0.393 (± 0.003)
	SynFlowNet[b]	0.355 (± 0.008)	0.410 (± 0.018)	0.298 (± 0.010)	0.282 (± 0.005)	0.380 (± 0.005)
	RGFN	0.349 (± 0.001)	0.405 (± 0.002)	0.307 (± 0.002)	0.271 (± 0.001)	0.381 (± 0.002)
	RxnFlow	0.369 (± 0.007)	0.436 (± 0.005)	0.319 (± 0.002)	0.289 (± 0.003)	0.405 (± 0.002)
3D Reaction	3DSynthFlow	0.389 (± 0.006)	0.444 (± 0.010)	0.321 (± 0.005)	0.294 (± 0.006)	0.416 (± 0.009)
	SYCLE	**0.421** (± 0.005)	**0.460** (± 0.018)	**0.326** (± 0.011)	**0.329** (± 0.002)	**0.433** (± 0.002)

Table 12. **Average AiZynthFinder success rate.** Results for top-100 diverse molecules across all 15 LIT-PCBA targets. The best results are in bold.

Category	Method	AiZynthFinder Success Rate (%, ↑)				
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Fragment	FragGFN	4.00 (± 3.54)	3.75 (± 1.92)	1.00 (± 1.00)	3.75 (± 1.92)	0.25 (± 0.43)
	FragGFN+SA	5.75 (± 1.48)	6.00 (± 2.55)	4.00 (± 2.24)	1.00 (± 0.00)	0.00 (± 0.00)
Reaction	SynNet	54.17 (± 7.22)	50.00 (± 0.00)	50.00 (± 0.00)	25.00 (± 25.00)	50.00 (± 0.00)
	BBAR	21.25 (± 5.36)	19.50 (± 3.20)	17.50 (± 1.50)	19.50 (± 3.64)	20.00 (± 2.12)
	SynFlowNet[a]	52.75 (± 1.09)	57.00 (± 6.04)	53.75 (± 9.52)	56.50 (± 2.29)	53.00 (± 8.92)
	SynFlowNet[b]	56.50 (± 6.58)	56.00 (± 3.08)	61.00 (± 2.74)	64.50 (± 9.86)	60.75 (± 3.77)
	RGFN	46.75 (± 6.86)	47.50 (± 4.06)	50.25 (± 2.17)	49.25 (± 4.38)	48.50 (± 6.58)
	RxnFlow	60.25 (± 3.77)	63.25 (± 3.11)	71.25 (± 4.15)	66.50 (± 4.03)	**65.50** (± 4.09)
3D Reaction	3DSynthFlow	**85.00** (± 7.87)	69.33 (± 3.86)	**80.00** (± 4.24)	69.67 (± 14.06)	60.00 (± 15.64)
	SYCLE	81.00 (± 3.00)	**72.33** (± 3.51)	75.67 (± 3.51)	**74.00** (± 3.46)	58.67 (± 6.03)
		GBA	IDH1	KAT2A	MAPK1	MTORC1
Fragment	FragGFN	5.00 (± 4.24)	4.50 (± 1.66)	1.25 (± 0.83)	0.75 (± 0.83)	2.75 (± 1.30)
	FragGFN+SA	3.00 (± 1.00)	4.50 (± 4.97)	1.50 (± 0.50)	3.25 (± 1.48)	3.50 (± 2.50)
Reaction	SynNet	50.00 (± 0.00)	50.00 (± 0.00)	45.83 (± 27.32)	50.00 (± 0.00)	54.17 (± 7.22)
	BBAR	17.75 (± 2.28)	19.50 (± 1.50)	18.75 (± 1.92)	16.25 (± 3.49)	18.75 (± 3.90)
	SynFlowNet[a]	58.00 (± 4.64)	59.00 (± 4.06)	55.50 (± 10.23)	47.25 (± 6.61)	57.00 (± 7.58)
	SynFlowNet[b]	61.50 (± 3.84)	60.50 (± 3.91)	57.25 (± 4.97)	44.50 (± 9.29)	62.00 (± 1.22)
	RGFN	48.00 (± 1.22)	43.00 (± 2.74)	49.00 (± 1.22)	42.00 (± 3.00)	44.50 (± 4.03)
	RxnFlow	66.00 (± 1.58)	64.00 (± 5.05)	66.50 (± 2.06)	**63.00** (± 4.64)	**70.50** (± 2.87)
3D Reaction	3DSynthFlow	64.67 (± 3.09)	65.67 (± 13.91)	**68.00** (± 19.51)	54.67 (± 10.50)	64.67 (± 15.69)
	SYCLE	**68.33** (± 6.43)	**68.67** (± 8.39)	59.00 (± 3.46)	60.67 (± 6.81)	61.00 (± 5.57)
		OPRK1	PKM2	PPARG	TP53	VDR
Fragment	FragGFN	2.50 (± 2.29)	8.75 (± 3.11)	0.75 (± 0.43)	4.25 (± 1.64)	3.50 (± 2.18)
	FragGFN+SA	3.25 (± 1.79)	9.75 (± 2.28)	1.25 (± 1.09)	2.25 (± 1.92)	3.75 (± 2.77)
Reaction	SynNet	54.17 (± 7.22)	50.00 (± 0.00)	54.17 (± 7.22)	29.17 (± 18.16)	45.83 (± 7.22)
	BBAR	13.75 (± 3.11)	20.00 (± 0.71)	15.50 (± 2.29)	18.50 (± 3.28)	12.25 (± 3.34)
	SynFlowNet[a]	56.50 (± 7.63)	50.75 (± 1.09)	53.50 (± 5.68)	55.50 (± 9.94)	53.50 (± 1.80)
	SynFlowNet[b]	56.25 (± 2.49)	58.00 (± 7.00)	57.00 (± 5.74)	66.50 (± 6.80)	53.50 (± 3.84)
	RGFN	48.00 (± 2.55)	48.50 (± 3.20)	47.00 (± 5.83)	53.25 (± 3.63)	46.50 (± 2.69)
	RxnFlow	**72.25** (± 2.05)	62.00 (± 3.24)	65.50 (± 4.03)	67.50 (± 2.96)	66.75 (± 2.28)
3D Reaction	3DSynthFlow	67.33 (± 20.37)	**78.67** (± 8.73)	65.67 (± 8.18)	61.33 (± 11.73)	**74.00** (± 17.15)
	SYCLE	71.33 (± 9.87)	75.33 (± 3.79)	**69.33** (± 6.66)	**79.00** (± 7.55)	73.33 (± 6.51)

Table 13. **Average number of synthesis steps.** Results for top-100 diverse molecules across all 15 LIT-PCBA targets. The best results are in bold.

Category	Method	Average Number of Synthesis Steps (↓)				
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Fragment	FragGFN	3.60 (± 0.10)	3.83 (± 0.08)	3.76 (± 0.20)	3.76 (± 0.16)	3.34 (± 0.18)
	FragGFN+SA	3.73 (± 0.09)	3.66 (± 0.04)	3.66 (± 0.07)	3.67 (± 0.21)	3.79 (± 0.19)
Reaction	SynNet	3.29 (± 0.36)	3.50 (± 0.00)	3.00 (± 0.00)	4.13 (± 0.89)	3.50 (± 0.00)
	BBAR	3.60 (± 0.17)	3.62 (± 0.19)	3.76 (± 0.04)	3.72 (± 0.11)	3.59 (± 0.14)
	SynFlowNet[a]	2.64 (± 0.07)	2.48 (± 0.07)	2.60 (± 0.25)	2.45 (± 0.09)	2.56 (± 0.29)
	SynFlowNet[b]	2.42 (± 0.10)	2.48 (± 0.10)	2.38 (± 0.10)	2.34 (± 0.30)	2.41 (± 0.14)
	RGFN	2.88 (± 0.21)	2.65 (± 0.09)	2.78 (± 0.19)	2.91 (± 0.23)	2.76 (± 0.17)
	RxnFlow	2.42 (± 0.23)	2.1 (± 0.12)	1.95 (± 0.20)	2.15 (± 0.18)	2.23 (± 0.18)
3D Reaction	3DSynthFlow	1.77 (± 0.37)	2.23 (± 0.16)	2.02 (± 0.20)	2.43 (± 0.61)	2.73 (± 0.50)
	SYCLE	**1.27** (± 0.06)	**1.50** (± 0.27)	**1.54** (± 0.30)	**1.40** (± 0.12)	**1.97** (± 0.23)
		GBA	IDH1	KAT2A	MAPK1	MTORC1
Fragment	FragGFN	3.94 (± 0.11)	3.74 (± 0.10)	3.78 (± 0.09)	3.72 (± 0.18)	3.84 (± 0.18)
	FragGFN+SA	3.94 (± 0.15)	3.84 (± 0.23)	3.66 (± 0.18)	3.69 (± 0.21)	3.94 (± 0.08)
Reaction	SynNet	3.38 (± 0.22)	3.38 (± 0.22)	3.46 (± 0.95)	3.50 (± 0.00)	3.29 (± 0.36)
	BBAR	3.71 (± 0.12)	3.68 (± 0.02)	3.63 (± 0.05)	3.73 (± 0.05)	3.77 (± 0.09)
	SynFlowNet[a]	2.48 (± 0.18)	2.61 (± 0.13)	2.45 (± 0.37)	2.81 (± 0.24)	2.44 (± 0.27)
	SynFlowNet[b]	2.45 (± 0.08)	2.46 (± 0.12)	2.45 (± 0.12)	2.83 (± 0.27)	2.39 (± 0.17)
	RGFN	2.77 (± 0.20)	2.97 (± 0.15)	2.78 (± 0.10)	2.86 (± 0.19)	2.92 (± 0.06)
	RxnFlow	2.10 (± 0.08)	2.16 (± 0.11)	2.29 (± 0.05)	2.29 (± 0.11)	2.05 (± 0.09)
3D Reaction	3DSynthFlow	2.47 (± 0.24)	2.49 (± 0.42)	2.47 (± 0.61)	2.70 (± 0.33)	2.42 (± 0.61)
	SYCLE	**1.67** (± 0.05)	**1.35** (± 0.14)	**1.92** (± 0.44)	**1.65** (± 0.47)	**1.60** (± 0.25)
		OPRK1	PKM2	PPARG	TP53	VDR
Fragment	FragGFN	3.82 (± 0.13)	3.71 (± 0.12)	3.73 (± 0.24)	3.73 (± 0.23)	3.75 (± 0.06)
	FragGFN+SA	3.62 (± 0.12)	3.84 (± 0.21)	3.71 (± 0.04)	3.66 (± 0.05)	3.67 (± 0.25)
Reaction	SynNet	3.29 (± 0.36)	3.50 (± 0.00)	3.29 (± 0.36)	3.67 (± 0.91)	3.63 (± 0.22)
	BBAR	3.70 (± 0.17)	3.61 (± 0.05)	3.72 (± 0.13)	3.65 (± 0.05)	3.77 (± 0.16)
	SynFlowNet[a]	2.49 (± 0.33)	2.62 (± 0.10)	2.56 (± 0.12)	2.51 (± 0.27)	2.55 (± 0.09)
	SynFlowNet[b]	2.50 (± 0.11)	2.52 (± 0.20)	2.53 (± 0.06)	2.34 (± 0.10)	2.51 (± 0.17)
	RGFN	2.81 (± 0.12)	2.82 (± 0.10)	2.82 (± 0.18)	2.64 (± 0.10)	2.84 (± 0.18)
	RxnFlow	2.00 (± 0.09)	2.34 (± 0.19)	2.21 (± 0.06)	2.12 (± 0.12)	2.12 (± 0.12)
3D Reaction	3DSynthFlow	2.34 (± 0.53)	2.20 (± 0.41)	2.61 (± 0.39)	2.73 (± 0.47)	2.25 (± 0.78)
	SYCLE	**1.42** (± 0.13)	**1.55** (± 0.18)	**1.90** (± 0.28)	**1.35** (± 0.09)	**1.55** (± 0.20)

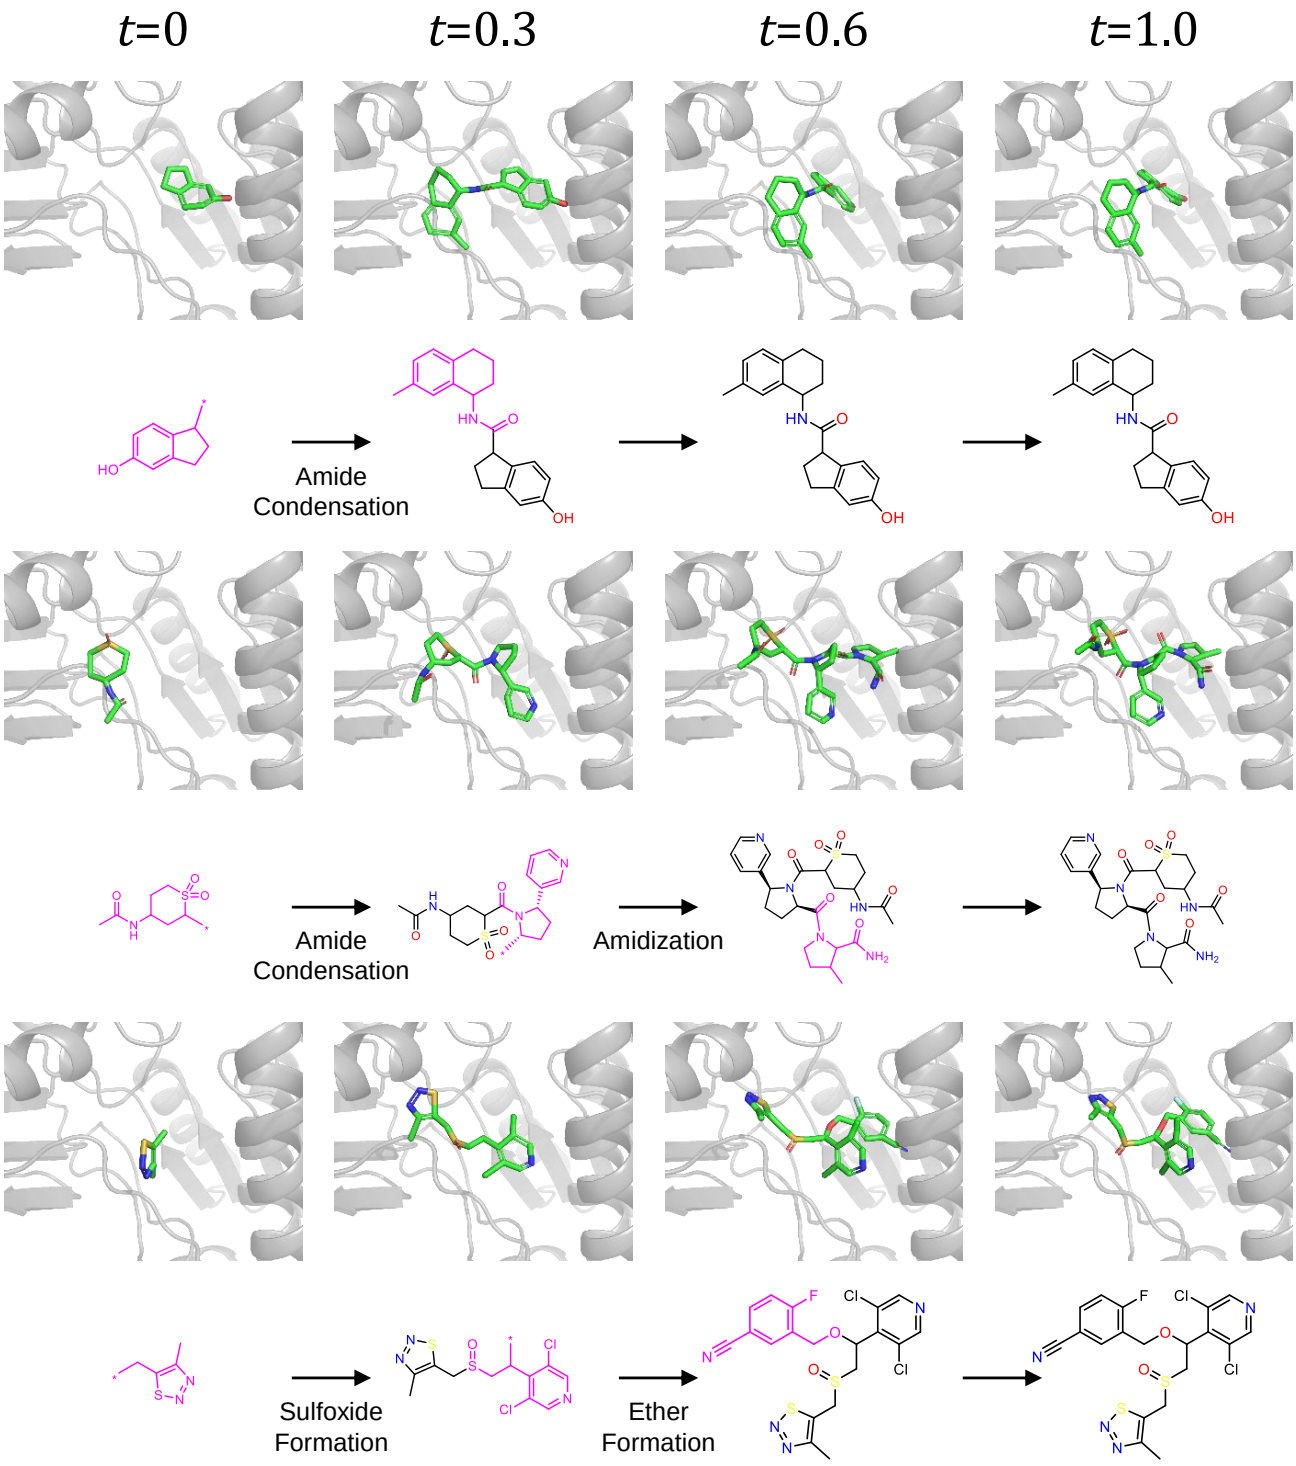

Figure 4. **Ligand generation trajectories of SYCLE.**