

COMPOSITIONAL FLOWS FOR 3D MOLECULE AND SYNTHESIS PATHWAY CO-DESIGN

Tony Shen^{1*}, Seonghwan Seo^{2*}, Ross Irwin^{3,4}, Kieran Didi^{5,6}

Simon Olsson⁴, Woo Youn Kim², Martin Ester¹

¹ Simon Fraser University, ² KAIST, ³ AstraZeneca, ⁴ Chalmers University,

⁵ University of Oxford, ⁶ NVIDIA

ABSTRACT

Many generative applications, such as synthesis-based 3D molecular design, involve constructing compositional objects with continuous features. Here, we introduce Compositional Generative Flows (CGFlow), a novel framework that extends flow matching to generate objects in compositional steps while modeling continuous states. Our key insight is that modeling compositional state transitions can be formulated as a straightforward extension of the flow matching interpolation process. We further build upon the theoretical foundations of generative flow networks (GFlowNets), enabling reward-guided sampling of compositional structures. We apply CGFlow to synthesizable drug design by jointly designing the molecule’s synthetic pathway with its 3D binding pose. Our approach achieves state-of-the-art binding affinity on all 15 targets from the LIT-PCBA benchmark, and 5.8 \times improvement in sampling efficiency compared to 2D synthesis-based baseline. To our best knowledge, our method is also the first to achieve state-of-art-performance in both Vina Dock (-9.38 kcal/mol) and AiZynthFinder success rate (62.2%) on the CrossDocked benchmark.

1 INTRODUCTION

Sampling objects through *compositional* steps while modeling *continuous* state is essential for a wide range of scientific applications Jain et al. (2023a); Wang et al. (2023). One such important application is synthesizable target-based drug design, which aims to jointly generate molecules through a sequence of compositional reaction steps and predict their continuous 3D conformations relative to a protein target Li et al. (2022). To this end, we propose a flow-based generative framework that jointly models the compositional structure and continuous state of objects.

Diffusion models Sohl-Dickstein et al. (2015); Ho et al. (2020); Song et al. (2021) and flow matching models Lipman et al. (2023) have achieved state-of-the-art performance in high-dimensional modeling tasks such as 3D molecule generation and protein structure design Hoogetboom et al. (2022); Campbell et al. (2024); Schneuing et al. (2024a). However, standard diffusion and flow matching are restricted to modeling all the dimensions of the object at once (Chen et al., 2024). This results in an inability to model the compositional structure of objects through sequential construction steps. As a consequence, two main limitations arise: (1) the validity of compositional objects cannot be ensured, as invalid generative actions cannot be masked, and (2) the potential for efficient reward credit assignment in the compositional space is restricted Bengio et al. (2021); Hansen et al. (2022); Yao et al. (2023). In drug design, where synthesizability is crucial for wet-lab validation, molecules can be naturally viewed as compositional objects constructed through sequential synthesis steps. Unfortunately, existing diffusion and flow matching models lack the ability to effectively model and respect the compositional nature of synthesis constraints when generating molecules.

Sequential models are a natural fit for generating composite objects. For instance, autoregressive models have been applied for 3D molecular design Peng et al. (2023); Gebauer et al. (2020). However, current autoregressive models lack mechanisms to correct errors from earlier steps, causing slight errors in early position predictions to cascade Jin et al. (2022). Generative flow networks

*Equal contribution. Correspondence to tsa87@sfu.ca




Figure 1: Overview of 3DSynthFlow generation. 3D Molecule $x = (\mathcal{C}, \mathcal{S})$ consisting of its synthesis pathway \mathcal{C} and 3D conformation \mathcal{S} (visualized using x). CGFlow generation interleaves 1). the continuous process for modeling position \mathcal{S} , and 2). the sequential sampling of synthesis steps at discrete intervals ($t=\{0, 0.25, 0.5, 0.75\}$). The modeling of synthesis pathways and position are both dependent on the object $x_t = \{\mathcal{S}_t, \mathcal{C}_t\}$, ensuring the interplay between the two processes.

(GFlowNets; Bengio et al., 2021) have recently shown success in sampling compositional structure for synthesis-based molecule design Koziarski et al. (2024); Cretu et al. (2024); Seo et al. (2024), but remain limited to 2D molecules.

In this paper, we identify a gap in standard *flow matching* and *sequential models* for generating compositional objects with continuous properties. To address this, we introduce **Compositional Generative Flows (CGFlow)**, a framework that combines flow matching for high-dimensional data with a respect for compositional structure. CGFlow interleaves two flow processes: the *Compositional Flow* gradually dismantles the structure from the data distribution to an empty state, while the *State Flow* transports the corresponding state variables from data to noise, assigning higher noise levels to components removed earlier (Ruhe et al., 2024). Constructive compositional steps are then sampled via a generative policy, and the conditional flow matching (CFM) objective (Lipman et al., 2023) estimates the vector field for state generation, ensuring that structure and state remain interdependent.

As an application of the CGFlow framework, we present 3DSynthFlow, the method for target-based drug design ensuring synthesizability. We combine flow matching-based 3D structure generation with the GFlowNet-based synthesis-aware molecular generative model developed by Seo et al. (2024). Fig. 1 illustrates how 3DSynthFlow jointly generates the synthesis pathway (compositional structure) and 3D conformation (continuous state) of molecules. Previous flow-based generative models focused on generating either the 3D molecular structure or the synthesis pathway, but not both (see Sec. E). 3DSynthFlow can jointly generate synthesis pathways and 3D molecular structures. This enables effective modeling of protein-ligand interactions and ensures synthesizability, both of which are essential for target-based drug discovery.

As an application, we present 3DSynthFlow for target-based drug design, which jointly generates synthesis pathways (compositional structure) and 3D conformations (continuous state). This joint modeling improves protein-ligand interaction predictions and ensures synthesizability. 3DSynthFlow achieves 5.8x sampling efficiency improvement, and state-of-the-art performance across all 15 targets in the LIT-PCBA benchmark (Tran-Nguyen et al., 2020) for binding affinity. 3DSynthFlow can be extended to pocket-conditional setting and achieve state-of-the-art performance in both Vina Dock (-9.38) and AiZynth success rate (62.2%) on the CrossDocked benchmark.

Our contributions are summarized as follows: (1) We propose Compositional Generative Flows (CGFlow), a flow-based framework that enables the generation of compositional objects while modeling continuous states. (2) We incorporate GFlowNets in CGFlow to efficiently explore the compositional state-space for high-reward samples. (3) We apply this framework to develop 3DSynthFlow for 3D molecule and synthesis pathway co-design, achieving *state-of-the-art* results on both LIT-PCBA and CrossDocked benchmark.

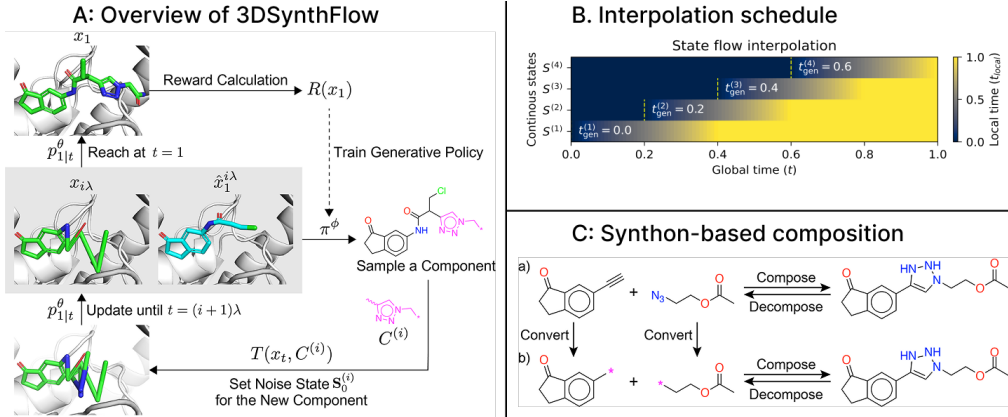


Figure 2: (A.) Overview of sampling and training using the pre-trained state flow model $p_{1|t}^\theta$. At $t = i\lambda$, compositional flow model π^ϕ samples a component $C^{(i)}$ based on the state x_t and its previous prediction \hat{x}_1^t . Then, the transition function T incorporates the component $C^{(i)}$ into the object. At $t = 1$, the compositional flow model is trained based on the reward of the generated object x_1 . (B.) Local time for each component over time, with $n = 4$, $\lambda = 0.2$, and $t_{\text{window}} = 0.4$. $t_{\text{local}}^{(i)} = 0$ indicates the $S^{(i)}$ has not yet been initialized. (C.) Illustration of synthesis-based composition rules: (a.) reactant unit-based and (b.) synthon unit-based (ours).

2 COMPOSITIONAL GENERATIVE FLOWS

CGFlow is a generative framework for modeling compositional objects with continuous states. It consists of two interleaved flows: the *Compositional Flow* for compositional structures and the *State Flow* for continuous states.

Data representation. An object x is represented as $(\mathcal{C}, \mathcal{S})$. The compositional structure $\mathcal{C} = (C^{(i)})_{i=1}^n$ is an ordered sequence of components (e.g., molecular building blocks) generated along a trajectory τ (e.g., synthesis pathway). The continuous states $\mathcal{S} = (S^{(i)})_{i=1}^n$ correspond to these components. Each component $C^{(i)}$ consists of m_i points (e.g., atoms) with states $S^{(i)} \in \mathbb{R}^{m_i \times d}$. Standard flow matching models only state variables. In contrast, CGFlow jointly models structure and states to ensure valid compositions.

Joint conditional flow process. We define a joint conditional flow (Lipman et al., 2023) that interpolates an object x from an initial state $x_0 = (\emptyset, [\])$ to a final state $x_1 = (C_1, S_1)$ via two flows: the *compositional flow* and the *state flow*. The process $\mathcal{P}_{t|1}(\cdot|x_1)$ satisfies the boundary conditions $\mathcal{P}_{t|1}(x_t|x_1) = \delta(x_t = x_0)$ at $t = 0$ and $\mathcal{P}_{t|1}(x_t|x_1) = \delta(x_t = x_1)$ at $t = 1$, ensuring a transition from x_0 to x_1 .

The *compositional flow* defines a conditional probability flow over the structure \mathcal{C} , transitioning from an empty graph $C_0 = \emptyset$ at $t = 0$ to a full structure C_1 at $t = 1$. Components are added sequentially in a fixed order. The number of components added by time t is given by

$$k(t) = \begin{cases} 0, & t = 0, \\ \min(\lfloor t/\lambda \rfloor + 1, n), & t > 0, \end{cases}$$

with $k(0) = 0$ and $k(1) = n$. Each component $C^{(i)}$ is generated at time $t_{\text{gen}}^{(i)} = \lambda(i - 1)$, and the constraint $\lambda \leq 1/n$ ensures that all components are added by $t = 1$. The structure at time t is then $\mathcal{C}_t = (C^{(i)})_{i=1}^{k(t)}$. For further details on scheduling and parameter choices, see Sec. A.4.

The *state flow* defines a conditional path over the continuous states $\mathcal{S} = (S^{(i)})_{i=1}^n$. Each state $S^{(i)}$ is introduced once its corresponding component is generated at $t_{\text{gen}}^{(i)}$. To capture varying uncertainty over time, we define a local time $t_{\text{local}}^{(i)} = \text{clip}\left(\frac{t - t_{\text{gen}}^{(i)}}{t_{\text{window}}}\right)$, where $\text{clip}(\cdot)$ restricts the value to $[0, 1]$ and t_{window} is the interpolation window. The state is updated via linear interpolation with Gaussian noise:

$\mathbf{S}_t^{(i)} = \mathcal{N}(t_{\text{local}}^{(i)} \mathbf{S}_1^{(i)} + (1 - t_{\text{local}}^{(i)}) \mathbf{S}_0^{(i)}, \sigma^2)$ for $t > t_{\text{gen}}^{(i)}$, and $\mathbf{S}_t^{(i)} = []$ otherwise. This formulation gradually refines each continuous state as time progresses (see Sec. A.5 for details).

Sampling. During the sampling process, objects are generated by interleaving the integration of the *state flow* model $p_{1|t}^\theta$ and sampling actions using the *compositional flow* policy π^θ . The process alternates between refining continuous states \mathcal{S}_t and sequentially constructing the compositional structure \mathcal{C}_t at fixed time points. We provide pseudocode in Algorithm 1 and details in Sec. A.3.

Training Objectives. CGFlow consist of two models: (1) a state flow model $p_{1|t}^\theta$ that updates continuous states, and (2) a compositional flow policy π^ϕ that sequentially samples compositional components. Both models are conditioned on the current object $\mathbf{x}_t = (\mathcal{C}_t, \mathcal{S}_t)$ and employ self-conditioning $\hat{\mathbf{x}}_1^t = (\mathcal{C}_t, \hat{\mathcal{S}}_1^t)$, where $\hat{\mathcal{S}}_1^t$ is obtained from the previous prediction (Chen et al., 2022).

State Flow Loss. The state flow model is trained independently in a simulation-free manner. For a given object \mathbf{x} , we sample a time t and generate \mathbf{x}_t via the joint interpolation process (see Section 2). The denoiser $p_{1|t}^\theta$ then predicts the refined continuous states $\hat{\mathcal{S}}_1^{t+\Delta t}$. The loss is computed as the mean squared error over the states corresponding to the generated components:

$$\mathcal{L}_{\text{state}} = \mathbb{E}_{p_{\text{data}}(\mathbf{x}_1), t \sim \mathcal{U}(0,1)} \sum_{i=1}^{k(t)} \left\| p_{1|t}^\theta(\mathbf{x}_t)^{(i)} - \mathbf{S}_1^{(i)} \right\|_2^2.$$

Compositional Flow Loss. Given a fixed state flow model $p_{1|t}^\theta$, the compositional flow policy π^ϕ samples a trajectory τ by sequentially adding components at times $t = i\lambda$, conditioned on \mathbf{x}_t and the self-conditioning $\hat{\mathbf{x}}_1^t$. The goal is to assign a sampling probability proportional to the reward $R(\mathbf{x}_1)$ for the final object \mathbf{x}_1 . Using the a special case of trajectory balance (TB) objective (Malkin et al., 2023), the loss is given by

$$\mathcal{L}_{\text{TB}}(\tau) = \left(\log \frac{Z_\phi \prod_{i=0}^{n-1} P_F(\mathbf{C}^{(i)} \mid \mathbf{x}_{i\lambda}, \hat{\mathbf{x}}_1^{i\lambda}; \phi)}{R(\mathbf{x}_1)} \right)^2.$$

where Z_ϕ is a normalizing constant and $P_F(\mathbf{C}^{(i)} \mid \mathbf{x}_{i\lambda}, \hat{\mathbf{x}}_1^{i\lambda}; \phi)$ is the forward probability of adding component $\mathbf{C}^{(i)}$ at time $i\lambda$. To ensure $P_B(-|-) = 1$, a deterministic transition is enforced by fixing the random seed when sampling the initial state $\mathbf{S}_0^{(i)}$ so that \mathbf{x}_t is uniquely determined by τ . We refer to Appendix B.2 for further theoretical background.

We refer readers to Sec. A.1 for a summary of the key steps in applying CGFlow to a new data domain. In the next section, we demonstrate CGFlow’s application to 3D molecular generation and synthesis pathway design.

3 POCKET-BASED 3D MOLECULE AND SYNTHESIS PATHWAY CO-DESIGN

Synthesizability is a critical factor for ensuring molecules can be readily made for wet lab validation. Recent works have incorporated combinatorial chemistry principles into generative models to respond to this challenge (Gao et al., 2022). Despite these advancements, conventional methods are restricted to 2D molecular graphs, limiting their ability to capture the protein-ligand interactions that are essential for biological efficacy. To address this limitation, we introduce 3DSynthFlow, a generative method based on CGFlow, enabling the synthesis pathways and binding poses co-design.

State flow model predicts the docking pose of a molecule within a target protein pocket. Unlike typical docking, where the full 2D molecule is provided at $t = 0$, the state flow model predicts the binding pose for the ligand that is sequentially constructed at discrete time intervals. The state flow model is trained on the protein-ligand complex dataset, independent of the compositional flow model. We modify the architecture from SemlaFlow (Irwin et al., 2024), originally designed for 3D single-body molecular generation, to enable protein-ligand modeling (See Sec. D.2.1).

Compositional flow model generates a synthesis pathway to construct a molecular structure. The action space for the composition flow model consists of all valid synthetic steps for a given compositional structure \mathcal{C}_t (details in App. D.1). We modify the architecture from RxnFlow (Seo et al., 2024) to model the sampling policy for synthesis steps using 3D protein-ligand complexes as input (see Sec. D.2.2). The compositional flow model is trained online.

Table 1: **Average Vina docking score of Top-100 diverse modes.** We report two versions of SynFlowNet (v2024.05^a and v2024.10^b). Avg. and Med. are the average and median values over the average docking scores for all 15 LIT-PCBA protein targets. The results for the remaining 10 target proteins are reported in Appendix. 16 The best results are in bold.

Category	Method	Average Vina Docking Score (kcal/mol, ↓)					Avg.	Med.
		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)	-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	-11.97 (± 0.43)	-12.25 (± 0.07)	-11.31 (± 0.12)	-11.25 (± 0.07)	-7.92 (± 0.22)	-10.97	-11.25




Figure 3: Number of discovered modes (satisfying Vina <-10 kcal/mol, QED >0.5, Sim <0.5) as a function of sampling budget for the first four LIT-PCBA targets for 3DSynthFlow (3D) vs RxnFlow (2D) across 2 seeds. Lower is better.

4 EXPERIMENTS

Overview We evaluate 3DSynthFlow for synthesizable target-based drug design in two common settings: pocket-specific optimization (Sec. 5) and pocket-conditional generation (Sec. 5). Our experiments aim to address three main key questions: (1) Does 3DSynthFlow generate molecules with improved binding affinity and ligand efficiency compared to existing synthesis-based baselines? (2) Does co-designing 3D structures improve the sampling efficiency in discovering diverse high-reward modes? (3) How does 3DSynthFlow generalize to the pocket-conditional setting, and how does it compare to existing SBDD baselines?

To answer these questions, We first apply 3DSynthFlow to optimize for targets in the LIT-PCBA benchmark (Tran-Nguyen et al., 2020). We evaluate the affinity, ligand efficiency, synthesis success rate and protein-ligand interactions of generated molecules. Then, we compare the sampling efficiency of 3DSynthFlow against 2D-based baseline. Lastly, we compare 3DSynthFlow against SBDD methods in the pocket-conditional setting on the CrossDocked dataset (Francoeur et al., 2020). The setup and baseline details are provided in Sec. F.1.

5 RESULTS

Pocket-specific optimization results. Table 1 presents the Vina results for the first five targets, while the full property results for all targets are in Sec. G.8. 3DSynthFlow consistently outperforms all baselines across LIT-PCBA targets in affinity and ligand efficiency, demonstrating that co-designing 3D molecular structures alongside synthesis pathways enhances the discovery of high-affinity molecules. Furthermore, 3DSynthFlow achieves 5.8x sampling efficiency to 2D baselines, as detailed in Sec. 5. Molecules generated by 3DSynthFlow also exhibit a higher number of protein-ligand interactions, as quantified by PoseCheck, highlighting the benefits of 3D-aware modeling for target-based drug design (see App. G.5).

Table 2: **Benchmark Results for Generative Methods.** We report the average (Avg.) and median (Med.) values for each metric when available. Reference denotes known actives. For methods where only one value is available, the median is indicated as “-”.

Category	Method	Validity (\uparrow)	Vina (\downarrow)		QED (\uparrow)		AiZynth. Succ. (\uparrow)		Div (\uparrow)	Time (\downarrow)
		Validity	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Avg.
	Reference	-	-7.71	-7.80	0.48	0.47	36.1%	-	-	-
Atom	Pocket2Mol	98.3%	-7.60	-7.16	0.57	0.58	29.1%	22.0%	0.83	2504
	TargetDiff	91.5%	-7.37	-7.56	0.49	0.49	9.9%	3.2%	0.87	3428
	DecompDiff	66.0%	-8.35	-8.25	0.37	0.35	0.9%	0.0%	0.84	6189
	DiffSBDD	76.0%	-6.95	-7.10	0.47	0.48	2.9%	2.0%	0.88	135
	MolCRAFT	96.7%	-8.05	-8.05	0.50	0.50	16.5%	9.1%	0.84	141
	MolCRAFT-large	70.8%	-9.25	-9.24	0.45	0.44	3.9%	0.0%	0.82	>141
Fragment	TacoGFN	100.0%	-8.24	-8.44	0.67	0.67	1.3%	1.0%	0.67	4
2D Reaction	RxnFlow	100.0%	-8.85	-9.03	0.67	0.67	34.8%	34.5%	0.81	4
3D Reaction	3DSynthFlow (low β)	100.0%	-9.00	-9.27	0.72	0.72	55.0%	54.5%	0.79	24
	3DSynthFlow (med β)	100.0%	-9.16	-9.41	0.73	0.74	<u>56.6%</u>	<u>56.0%</u>	0.76	24
	3DSynthFlow (high β)	100.0%	-9.38	-9.62	0.74	0.74	62.2%	63.0%	0.66	24

Finally, we conduct extensive ablation studies on the various technical decisions regarding: flow matching steps (App. G.2), time scheduling (App. G.3), and use of pose-based rewards (App. G.1). Analysis on training efficiency can be found in App. F.2.3.

3DSynthFlow improves sampling efficiency. We evaluate sample efficiency on the first 5 LIT-PCBA targets and report results in Fig. 3 and Table. 15. Diverse high-scoring modes were defined by QED > 0.5, Vina < -10 kcal/mol¹, and Tanimoto similarity < 0.5 to any other mode. After sampling 10,000 molecules, 3DSynthFlow identified 5.8x number of diverse high-scoring modes compared to RxnFlow (316 vs 54.6). This enhanced sampling efficiency validates the effectiveness of 3DSynthFlow framework, and suggests a higher probability of experimental success in downstream drug discovery applications.

Pocket-conditional generation results. As shown in Table. 2, 3DSynthFlow achieves significant improvements in pocket-conditional generation, particularly in synthesizability and docking scores. In particular, 3DSynthFlow- high β attains an average docking score of -9.38 kcal/mol, outperforming RxnFlow (-8.85) and state-of-the-art diffusion-based methods like MolCRAFT-large (-9.25) and DecompDiff (-8.35). We attribute this improvement largely to our explicit consideration of 3D co-design in the reaction-based generation framework.

3DSynthFlow achieves competitive synthesizability. 3DSynthFlow also ensures high AiZynth success rate via synthesis-based generation, achieving success rates of 55.0%-62.2% across different β values in the pocket-conditioned setting. This surpasses synthesis-based RxnFlow (34.8%), and TacoGFN (1.3%), which optimize for SA score. SBDD methods such as MolCRAFT-large can attain strong binding affinity (-9.25) similar to 3DSynthFlow (-9.38), however their synthesis success rate is much lower (3.9% vs 62.2%). The main contribution of 3DSynthFlow is representing both the compositional nature for synthesis constraints and modeling 3D poses for binding. Both synthesizability and binding are both necessary requirement for experimental validation.

6 CONCLUSION

In this work, we introduce Compositional Generative Flows (CGFlow), a flexible generative framework for jointly modeling compositional structures and continuous states. We introduce a simple extension to the flow matching interpolation process for handling compositional state transition. CGFlow enables the integration of GFlowNets for efficient exploration of compositional state spaces with flow matching for continuous state modeling. We apply CGFlow to 3D molecule and synthesis pathway co-design and develop 3DSynthFlow, which achieves state-of-the-art performance on both LIT-PCBA and CrossDocked benchmark. Future work includes using a more expressive model for pose prediction in 3DSynthFlow and developing more application-specific methods using CGFlow.

¹Except for FEN1, where we use Vina threshold below -7 kcal/mol to maintain similar number of modes compared to the other four targets.

REFERENCES

- Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 27381–27394. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/e614f646836aaed9f89ce58e837e2310-Paper.pdf.
- G Richard Bickerton, Gaia V Paolini, J  r  my Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98, 2012.
- Caterina Bissantz, Bernd Kuhn, and Martin Stahl. A medicinal chemist’s guide to molecular interactions. *Journal of Medicinal Chemistry*, 53(14):5061–5084, 2010. doi: 10.1021/jm100112j. URL <https://doi.org/10.1021/jm100112j>. PMID: 20345171.
- I. D. Brown. On the geometry of O–H · · · O hydrogen bonds. *Acta Crystallographica Section A*, 32(1):24–31, Jan 1976. doi: 10.1107/S0567739476000041.
- Andrew Campbell, William Harvey, Christian Weilbach, Valentin De Bortoli, Tom Rainforth, and Arnaud Doucet. Trans-dimensional generative modeling via jump diffusion models, 2023. URL <https://arxiv.org/abs/2305.16261>.
- Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design, 2024. URL <https://arxiv.org/abs/2402.04997>.
- Boyuan Chen, Diego Marti Monso, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion, 2024. URL <https://arxiv.org/abs/2407.01392>.
- Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022.
- Connor W Coley, Luke Rogers, William H Green, and Klavs F Jensen. Scscore: synthetic complexity learned from a reaction corpus. *Journal of chemical information and modeling*, 58(2): 252–261, 2018.
- Miruna Cretu, Charles Harris, Julien Roy, Emmanuel Bengio, and Pietro Lio. Synflownet: Towards molecule design with guaranteed synthesis pathways. In *ICLR 2024 Workshop on Generative and Experimental Perspectives for Biomolecular Design*, 2024. URL <https://openreview.net/forum?id=kjcBA2I2My>.
- 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.
- Paul G. Francoeur, Tomohide Masuda, Jocelyn Sunseri, Andrew Jia, Richard B. Iovanisci, Ian Snyder, and David R. Koes. 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. doi: 10.1021/acs.jcim.0c00411. PMID: 32865404.
- Richard A. Friesner, Jay L. Banks, Robert B. Murphy, Thomas A. Halgren, Jasna J. Klicic, Daniel T. Mainz, Matthew P. Repasky, Eric H. Knoll, Mee Shelley, Jason K. Perry, David E. Shaw, Perry Francis, and Peter S. Shenkin. Glide: A new approach for rapid, accurate docking and scoring. 1. method and assessment of docking accuracy. *Journal of Medicinal Chemistry*, 47(7):1739–1749, 2004. doi: 10.1021/jm0306430. URL <https://doi.org/10.1021/jm0306430>. PMID: 15027865.
- Wenhao Gao and Connor W. Coley. The synthesizability of molecules proposed by generative models, 2020. URL <https://arxiv.org/abs/2002.07007>.

- Wenhao Gao, Rocío Mercado, and Connor W. Coley. Amortized tree generation for bottom-up synthesis planning and synthesizable molecular design. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=FRxhHdnxt1>.
- Wenhao Gao, Shitong Luo, and Connor W. Coley. Generative artificial intelligence for navigating synthesizable chemical space. *arXiv preprint arXiv:2410.03494*, 2024.
- Niklas W. A. Gebauer, Michael Gastegger, and Kristof T. Schütt. Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules, 2020. URL <https://arxiv.org/abs/1906.00957>.
- Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist, and Esben Bjerrum. Aizynthfinder: a fast, robust and flexible open-source software for retrosynthetic planning. *Journal of cheminformatics*, 12(1):70, 2020.
- Mahdi Ghorbani, Leo Gendele, Paul Beroza, and Michael J. Keiser. Autoregressive fragment-based diffusion for pocket-aware ligand design, 2023. URL <https://arxiv.org/abs/2401.05370>.
- 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. *Iscience*, 23(11), 2020.
- Jiaqi Guan, Wesley Wei Qian, Xingang Peng, Yufeng Su, Jian Peng, and Jianzhu Ma. 3d equivariant diffusion for target-aware molecule generation and affinity prediction, 2023. URL <https://arxiv.org/abs/2303.03543>.
- Jiaqi Guan, Xiangxin Zhou, Yuwei Yang, Yu Bao, Jian Peng, Jianzhu Ma, Qiang Liu, Liang Wang, and Quanquan Gu. Decompdiff: Diffusion models with decomposed priors for structure-based drug design, 2024. URL <https://arxiv.org/abs/2403.07902>.
- Jeff Guo and Philippe Schwaller. It takes two to tango: Directly optimizing for constrained synthesizability in generative molecular design. *arXiv preprint arXiv:2410.11527*, 2024.
- Nicklas Hansen, Xiaolong Wang, and Hao Su. Temporal difference learning for model predictive control, 2022. URL <https://arxiv.org/abs/2203.04955>.
- 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?, 2023. URL <https://arxiv.org/abs/2308.07413>.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. URL <https://arxiv.org/abs/2006.11239>.
- Emiel Hoogetboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d, 2022. URL <https://arxiv.org/abs/2203.17003>.
- Ross Irwin, Alessandro Tibo, Jon Paul Janet, and Simon Olsson. Efficient 3d molecular generation with flow matching and scale optimal transport, 2024. URL <https://arxiv.org/abs/2406.07266>.
- Moksh Jain, Tristan Deleu, Jason Hartford, Cheng-Hao Liu, Alex Hernandez-Garcia, and Yoshua Bengio. Gflownets for ai-driven scientific discovery. *Digital Discovery*, 2:557–577, 2023a. doi: 10.1039/D3DD00002H. URL <http://dx.doi.org/10.1039/D3DD00002H>.
- Moksh Jain, Sharath Chandra Rapparthi, Alex Hernández-García, Jarrid Rector-Brooks, Yoshua Bengio, Santiago Miret, and Emmanuel Bengio. Multi-objective gflownets. In *International conference on machine learning*, pp. 14631–14653. PMLR, 2023b.
- Wengong Jin, Jeremy Wohlwend, Regina Barzilay, and Tommi Jaakkola. Iterative refinement graph neural network for antibody sequence-structure co-design, 2022. URL <https://arxiv.org/abs/2110.04624>.

- Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael J. L. Townshend, and Ron Dror. Learning from protein structure with geometric vector perceptrons, 2021. URL <https://arxiv.org/abs/2009.01411>.
- Bowen Jing, Ezra Erives, Peter Pao-Huang, Gabriele Corso, Bonnie Berger, and Tommi Jaakkola. Eigenfold: Generative protein structure prediction with diffusion models, 2023. URL <https://arxiv.org/abs/2304.02198>.
- Hyeonwoo Kim, Kyunghoon Lee, Chansu Kim, Jaechang Lim, and Woo Youn Kim. Dfrscore: deep learning-based scoring of synthetic complexity with drug-focused retrosynthetic analysis for high-throughput virtual screening. *Journal of Chemical Information and Modeling*, 64(7):2432–2444, 2023.
- Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching, 2023. URL <https://arxiv.org/abs/2306.15030>.
- 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.
- Tuan Le, Julian Cremer, Frank Noé, Djork-Arné Clevert, and Kristof Schütt. Navigating the design space of equivariant diffusion-based generative models for de novo 3d molecule generation, 2023. URL <https://arxiv.org/abs/2309.17296>.
- Xinze Li, Penglei Wang, Tianfan Fu, Wenhao Gao, Chengtao Li, Leilei Shi, and Junhong Liu. Autodiff: Autoregressive diffusion modeling for structure-based drug design, 2024. URL <https://arxiv.org/abs/2404.02003>.
- Yibo Li, Jianfeng Pei, and Luhua Lai. Synthesis-driven design of 3d molecules for structure-based drug discovery using geometric transformers. *arXiv preprint arXiv:2301.00167*, 2022.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling, 2023. URL <https://arxiv.org/abs/2210.02747>.
- Tairan Liu, Misagh Naderi, Chris Alvin, Supratik Mukhopadhyay, and Michal Brylinski. 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.
- Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution, 2024. URL <https://arxiv.org/abs/2310.16834>.
- Shitong Luo, Jiaqi Guan, Jianzhu Ma, and Jian Peng. A 3D generative model for structure-based drug design. In *NeurIPS*, 2021.
- Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory balance: Improved credit assignment in gflownets, 2023. URL <https://arxiv.org/abs/2201.13259>.
- 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.
- Rebecca M Neeser, Bruno Correia, and Philippe Schwaller. Fsscore: A personalized machine learning-based synthetic feasibility score. *Chemistry-Methods*, 4(11):e202400024, 2024.
- Mohit Pandey, Gopeshh Subbaraj, Artem Cherkasov, Martin Ester, and Emmanuel Bengio. Pre-training generative flow networks with inexpensive rewards for molecular graph generation, 2025. URL <https://arxiv.org/abs/2503.06337>.
- Xingang Peng, Shitong Luo, Jiaqi Guan, Qi Xie, Jian Peng, and Jianzhu Ma. Pocket2mol: Efficient molecular sampling based on 3d protein pockets. In *ICML*, pp. 17644–17655. PMLR, 2022.
- Xingang Peng, Jiaqi Guan, Jian Peng, and Jianzhu Ma. Pocket-specific 3d molecule generation by fragment-based autoregressive diffusion models, 2023. URL <https://openreview.net/forum?id=HGsoelwmRW5>.

- Yanru Qu, Keyue Qiu, Yuxuan Song, Jingjing Gong, Jiawei Han, Mingyue Zheng, Hao Zhou, and Wei-Ying Ma. Molcraft: Structure-based drug design in continuous parameter space, 2024. URL <https://arxiv.org/abs/2404.12141>.
- Matthew Ragoza, Tomohide Masuda, and David Ryan Koes. Generating 3d molecules conditional on receptor binding sites with deep generative models. *Chemical science*, 13(9):2701–2713, 2022.
- Danny Reidenbach. EvoSBDD: Latent evolution for accurate and efficient structure-based drug design. In *ICLR 2024 Workshop on Machine Learning for Genomics Explorations*, 2024. URL <https://openreview.net/forum?id=sLhUNz0uTz>.
- David Ruhe, Jonathan Heek, Tim Salimans, and Emiel Hoogetboom. Rolling diffusion models, 2024. URL <https://arxiv.org/abs/2402.09470>.
- Arne Schneuing, Charles Harris, Yuanqi Du, Kieran Didi, Arian Jamasb, Ilia Igashov, Weitao Du, Carla Gomes, Tom Blundell, Pietro Lio, Max Welling, Michael Bronstein, and Bruno Correia. Structure-based drug design with equivariant diffusion models, 2024a. URL <https://arxiv.org/abs/2210.13695>.
- Arne Schneuing, Charles Harris, Yuanqi Du, Kieran Didi, Arian Jamasb, Ilia Igashov, Weitao Du, Carla Gomes, Tom Blundell, Pietro Lio, Max Welling, Michael Bronstein, and Bruno Correia. Structure-based drug design with equivariant diffusion models, 2024b. URL <https://arxiv.org/abs/2210.13695>.
- Seonghwan Seo and Woo Youn Kim. Pharmaconet: deep learning-guided pharmacophore modeling for ultra-large-scale virtual screening. *Chem. Sci.*, 15:19473–19487, 2024. doi: 10.1039/D4SC04854G. URL <http://dx.doi.org/10.1039/D4SC04854G>.
- Seonghwan Seo, Jaechang Lim, and Woo Youn Kim. Molecular generative model via retrosynthetically prepared chemical building block assembly. *Advanced Science*, 10(8):2206674, 2023.
- Seonghwan Seo, Minsu Kim, Tony Shen, Martin Ester, Jinkyoo Park, Sungsoo Ahn, and Woo Youn Kim. Generative flows on synthetic pathway for drug design, 2024. URL <https://arxiv.org/abs/2410.04542>.
- Tony Shen, Seonghwan Seo, Grayson Lee, Mohit Pandey, Jason R Smith, Artem Cherkasov, Woo Youn Kim, and Martin Ester. Tacogfn: Target-conditioned gflownet for structure-based drug design, 2024. URL <https://arxiv.org/abs/2310.03223>.
- Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics, 2015. URL <https://arxiv.org/abs/1503.03585>.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations, 2021. URL <https://arxiv.org/abs/2011.13456>.
- Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou, and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport, 2023. URL <https://arxiv.org/abs/2312.07168>.
- Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design, 2024. URL <https://arxiv.org/abs/2402.05841>.
- Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Harmonic self-conditioned flow matching for multi-ligand docking and binding site design, 2024. URL <https://arxiv.org/abs/2310.05764>.
- 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.

- 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.
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, Anima Anandkumar, Karianne Bergen, Carla P Gomes, Shirley Ho, Pushmeet Kohli, Joan Lasenby, Jure Leskovec, Tie-Yan Liu, Arjun Manrai, Debora Marks, Bharath Ramsundar, Le Song, Jimeng Sun, Jian Tang, Petar Veličković, Max Welling, Linfeng Zhang, Connor W Coley, Yoshua Bengio, and Marinka Zitnik. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, August 2023.
- Tong Wu, Zhihao Fan, Xiao Liu, Yeyun Gong, Yelong Shen, Jian Jiao, Hai-Tao Zheng, Juntao Li, Zhongyu Wei, Jian Guo, Nan Duan, and Weizhu Chen. Ar-diffusion: Auto-regressive diffusion model for text generation, 2023. URL <https://arxiv.org/abs/2305.09515>.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models, 2023. URL <https://arxiv.org/abs/2305.10601>.
- Yuejiang Yu, Chun Cai, Jiayue Wang, Zonghua Bo, Zhengdan Zhu, and Hang Zheng. Uni-dock: Gpu-accelerated docking enables ultralarge virtual screening. *Journal of chemical theory and computation*, 19(11):3336–3345, 2023.
- Seongjun Yun, Minbyul Jeong, Sungdong Yoo, Seunghun Lee, S Yi Sean, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. Graph transformer networks: Learning meta-path graphs to improve gnn. *Neural Networks*, 153:104–119, 2022.
- Zihan Zhang, Richard Liu, Kfir Aberman, and Rana Hanocka. Tedi: Temporally-entangled diffusion for long-term motion synthesis, 2023. URL <https://arxiv.org/abs/2307.15042>.
- Alex Zhavoronkov, Yan A Ivanenkov, Alex Aliper, Mark S Veselov, Vladimir A Aladinskiy, Anastasiya V Aladinskaya, Victor A Terentiev, Daniil A Polykovskiy, Maksim D Kuznetsov, Arip Asadulaev, et al. Deep learning enables rapid identification of potent ddr1 kinase inhibitors. *Nature biotechnology*, 37(9):1038–1040, 2019.

A CGFLOW DETAILS

A.1 CGFLOW RECIPE

We now summarize the key steps for implementing CGFlow for generative tasks which constructs compositional structures with continuous states.

1. Decompose the object into its compositional and continuous part (e.g., 2D structure and 3D coordinates) (See Sec. 2).
2. Define action space for additive composition steps.
3. Define a function for sampling action sequence for constructing the compositional structure (e.g. synthesis order).
4. Train State Flow model on existing datasets (e.g. 3D protein-ligand complexes) (Sec. 2).
5. Train the compositional flow model using GFlowNets-based Compositional Flow loss, if reward function $R(x)$ (docking score) is available, else use cross-entropy loss (Sec. 2).
6. Run sampling using trained state flow and compositional flow model (Algorithm 1).

A.2 SAMPLING ALGORITHM

Algorithm 1 Compositional Flow Sampling with CGFlow

Require: step size Δt , interval λ , time window t_{window}

```

1: init  $t = 0, i = 0, \mathcal{C}_t = \mathcal{C}_0, \mathcal{S}_0 = \mathcal{S}_0, \hat{\mathcal{S}}_1^t = \mathcal{S}_0$ 
2: while  $t < 1$  do
3:    $\mathbf{x}_t \leftarrow (\mathcal{C}_t, \mathcal{S}_t, \hat{\mathcal{S}}_1^t)$   $\triangleright$  Use self-conditioning.
4:   if  $t \bmod \lambda = 0$  then
5:      $i \leftarrow i + 1$ 
6:      $\mathbf{C}^{(i)} \sim \pi_\theta(\mathbf{x}_t)$   $\triangleright$  Compositional flow model: Sample the next compositional component  $i$ .
7:      $\mathcal{S}_t^{(i)} \sim \mathcal{N}(0, 1)$   $\triangleright$  Initialize new state values for component  $i$  under the fixed seed
8:      $\mathbf{x}_t \leftarrow T(\mathbf{x}_t, \mathbf{C}^{(i)}, \mathcal{S}_t^{(i)})$   $\triangleright$  Transition with sampled composition and state value.
9:      $t_{\text{gen}}^{(i)} \leftarrow t$ 
10:    end if
11:     $t_{\text{local}}^{(j)} \leftarrow \text{clip}\left(\frac{t - t_{\text{gen}}^{(j)}}{t_{\text{window}}}\right), \forall j \leq i$ 
12:     $\hat{\mathcal{S}}_1^{(j)} \leftarrow p_{1|t}^\theta(\mathbf{x}_t, t_{\text{local}}^{(j)}), \forall j \leq i$   $\triangleright$  State flow model: Predict final state values.
13:     $\mathcal{S}_{t+\Delta t}^{(j)} \leftarrow \mathcal{S}_t^{(j)} + (\hat{\mathcal{S}}_1^{(j)} - \mathcal{S}_t^{(j)}) \cdot \kappa^{(j)} \Delta t, \forall j \leq i$   $\triangleright$  Step according to the predicted vector field.
14:     $t \leftarrow t + \Delta t$ 
15:  end while
16: return  $\mathbf{x}_1$ 

```

A.3 SAMPLING DETAILS

For state flow, the vector field governing the continuous states for the i -th component $\mathcal{S}^{(i)}$ is defined as $\hat{\mathcal{S}}_1^{(i)} - \mathcal{S}_t^{(i)}$, where $\hat{\mathcal{S}}_1^{(i)}$ is the predicted clean state by $p_{1|t}^\theta$, as formulated in previous works (Le et al., 2023). The rate at which we step in this vector field $\kappa^{(i)}$ is determined by the time remaining in the interpolation process for the state $\mathcal{S}^{(i)}$:

$$\kappa^{(i)} = \frac{\min(t_{\text{end}}^{(i)} - t, \Delta t)}{t_{\text{window}}}, \quad (1)$$

where $t_{\text{end}}^{(i)} = t_{\text{gen}}^{(i)} + t_{\text{window}}$ is the time at which the interpolation for component i is completed, and t_{window} is the interpolation window. The state values $\mathcal{S}^{(i)}$ are updated using Euler’s method:

$$\mathcal{S}_{t+\Delta t}^{(i)} = \mathcal{S}_t^{(i)} + (\hat{\mathcal{S}}_1^{(i)} - \mathcal{S}_t^{(i)}) \cdot \kappa^{(i)} \Delta t, \quad (2)$$

Intuitively, if $t \geq t_{\text{end}}^{(i)}$, the state $\mathbf{S}_t^{(i)}$ is directly set to the predicted clean value $\hat{\mathbf{S}}_1^{(i)}$, ensuring the coherence of the continuous state.

For compositional structure generation, new components $\mathbf{C}^{(i)}$ are sampled from the compositional flow policy π^θ at discrete time intervals separated by λ . Transition function $T(\mathbf{x}_t, \mathbf{C}^{(i)})$ incorporates newly sampled component $\mathbf{C}^{(i)}$ into the object. T also incorporates the new component’s associated state $\mathbf{S}_0^{(i)}$ by sampling its value from a noise distribution (typically Gaussian).

A.4 COMPOSITIONAL FLOW DETAILS

In the compositional flow, we define:

$$k(t) = \begin{cases} 0, & t = 0, \\ \min(\lfloor t/\lambda \rfloor + 1, n), & t > 0, \end{cases}$$

which determines the number of components added at time t . The key points are:

- **Initial and final states:** $k(0) = 0$ implies $\mathcal{C}_0 = \emptyset$ and $k(1) = n$ implies that the full structure \mathcal{C}_1 is generated.
- **Generation times:** Each component $\mathbf{C}^{(i)}$ is generated at $t_{\text{gen}}^{(i)} = \lambda(i - 1)$. For example, the first component is generated at $t = 0$, the second at $t = \lambda$, and so on.
- **Scheduling constraint:** The requirement $\lambda \leq 1/n$ guarantees that $t_{\text{gen}}^{(n)} \leq 1 - \lambda$ so that all n components are generated by $t = 1$.

This design allows for a gradual, discrete build-up of the structure in fixed time intervals.

A.5 STATE FLOW DETAILS

In the state flow, each continuous state $\mathbf{S}^{(i)}$ is associated with a local time variable defined as

$$t_{\text{local}}^{(i)} = \text{clip}\left(\frac{t - t_{\text{gen}}^{(i)}}{t_{\text{window}}}\right),$$

with the following properties:

- **Local time range:** The clipping operation ensures that $t_{\text{local}}^{(i)}$ lies in the interval $[0, 1]$. It starts at 0 when $t = t_{\text{gen}}^{(i)}$ and reaches 1 once $t \geq t_{\text{gen}}^{(i)} + t_{\text{window}}$.
- **Interpolation:** When $t > t_{\text{gen}}^{(i)}$, the state $\mathbf{S}^{(i)}$ is updated using a linear interpolation between the initial noisy state $\mathbf{S}_0^{(i)}$ and the refined final state $\mathbf{S}_1^{(i)}$. Gaussian noise with variance σ^2 is added to this interpolation:

$$\mathbf{S}_t^{(i)} = \mathcal{N}\left(t_{\text{local}}^{(i)} \mathbf{S}_1^{(i)} + (1 - t_{\text{local}}^{(i)}) \mathbf{S}_0^{(i)}, \sigma^2\right).$$

- **Conditional state existence:** If $t \leq t_{\text{gen}}^{(i)}$, the state $\mathbf{S}_t^{(i)}$ does not exist (denoted by $[]$).

This method allows the uncertainty in each component’s state to decrease as time progresses, mirroring processes in diffusion-based video generation where uncertainty is reduced gradually.

A.6 ALTERNATIVE TRAINING OBJECTIVES

There may be cases where reward function $R(x)$ is not available, or the goal is to model the data distribution instead. In these settings, the compositional flow model can adopt cross-entropy loss to maximize the log-likelihood of sampling $\mathbf{C}^{(\sigma_i)}$ at each step i , aligned with the valid generation order σ . The loss is:

$$\mathcal{L}_{\text{comp}} = -\mathbb{E}_{p_t(\mathbf{x}_t)} \sum_{i=1}^n \log \pi^\theta(\mathbf{C}^{(\sigma_i)} | \mathbf{x}_{t=t_{\text{gen}}^{(\sigma_i)}}), \quad (3)$$

where $\pi^\theta(\mathbf{C}^{(\sigma_i)} | \mathbf{x}_t)$ represents the policy model for generating the next compositional component $\mathbf{C}^{(\sigma_i)}$, conditioned on the entire object \mathbf{x}_t .

B THEORETICAL BACKGROUND

B.1 GFlowNet PRELIMINARY

Generative Flow Networks (GFlowNets; Bengio et al., 2021) are a family of probabilistic models that learn a stochastic policy to construct compositional objects $x \in \mathcal{X}$ proportional to the reward of terminate state $R(x)$, i.e., $p(x) \propto R(x)$. Each object x is constructed through a trajectory $\tau = (s_0 \rightarrow \dots \rightarrow s_n = x) \in \mathcal{T}$ from the initial state s_0 and a series of state transitions $s \rightarrow s'$, where the terminate state is the object $s_n = x \in \mathcal{X}$.

A GFlowNet models a flow F as an unnormalized density function along a directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{S}, \mathcal{A})$, where \mathcal{S} denotes the state space and \mathcal{A} represents transitions. We define the *trajectory flow* $F(\tau)$ as a flow through the trajectory τ . The *node flow* $F(s)$ is defined as the sum of trajectory flows through the node s , i.e., $F(s) = \sum_{\tau \in \mathcal{T}} F(\tau)$, and the *edge flow* $F(s \rightarrow s')$ is defined as the total flow along the edge $s \rightarrow s'$, i.e., $F(s \rightarrow s') = \sum_{(\tau \in \mathcal{T})} F(\tau)$.

From the flow network, we define two policy distributions. The *forward policy* $P_F(s'|s)$ executes the state transition $s \rightarrow s'$ from the flow distribution, i.e., $P_F(s'|s) = F(s \rightarrow s')/F(s)$. Similarly, the *backward policy* $P_B(s|s')$ distributes the node flow $F(s)$ to reverse transitions $s \dashrightarrow s'$, i.e., $P_B(s|s') = F(s' \rightarrow s)/F(s')$.

To match the likelihood of generating $x \in \mathcal{X}$ with the reward function R , two boundary conditions must be achieved. First, the node flow of each terminal state x , which represents the unnormalized probability of sampling the object x , must equal its reward, i.e., $F(x) = R(x)$. Second, the initial node flow s_0 , which represents the *partition function* Z , must equal the sum of all rewards. i.e., $Z = \sum_{x \in \mathcal{X}} R(x)$. One such objective to satisfy these conditions is *trajectory balance* (TB; Malkin et al., 2023), defined as follows:

$$\mathcal{L}_{\text{TB}}(\tau) = \left(\log \frac{Z_\theta \prod_{t=1}^n P_F(s_t|s_{t-1}; \theta)}{R(x) \prod_{t=1}^n P_B(s_{t-1}|s_t; \theta)} \right)^2, \quad (4)$$

where the P_F , P_B , and Z are directly parameterized to minimize the TB objective.

B.2 TRAINING OBJECTIVE OF COMPOSITIONAL FLOW MODEL

The object $x = (\mathcal{C}, \mathcal{S})$ is sequence data where $\mathcal{C} = (\mathbf{C}^{(i)})_{i=1}^n$ and $\mathcal{S} = (\mathbf{S}^{(i)})_{i=1}^n$. Therefore, we formulate the generative process as an auto-regressive process, i.e., $P_B(-|-) = 1.0$.

To train the compositional flow model with the trajectory balance (TB) objective, we must estimate forward transition probabilities along the trajectory $\tau = (\mathbf{x}_0 \rightarrow \mathbf{x}_{i\lambda} \rightarrow \dots \rightarrow \mathbf{x}_{n\lambda} \rightarrow \mathbf{x}_1)$, where λ is the time interval between successive components $\mathbf{C}^{(\cdot)}$, and $n < 1/\lambda$ is the total number of components.

Since the state flow model $p_{1|t}^\theta$ introduces randomness in the initial state $\mathbf{S}_0^{(i)} \sim \mathcal{N}(0, \sigma^2)$, estimating the forward transition probability involves integrating over this noise distribution as follows:

$$\begin{aligned} & P_F \left(\mathbf{x}_{(i+1)\lambda} \middle| \mathbf{x}_{i\lambda}, \hat{\mathbf{x}}_1^t; \phi, p_{1|t}^\theta \right) \\ &= P_F \left(\mathbf{C}^{(i)} \middle| \mathbf{x}_{i\lambda}, \hat{\mathbf{x}}_1^t; \phi \right) \int P \left(\mathbf{x}_{(i+1)\lambda} \middle| T(\mathbf{x}_{i\lambda}, \mathbf{C}^{(i)}, \mathbf{S}_0^{(i)}); p_{1|t}^\theta \right) p(\mathbf{S}_0^{(i)}) d\mathbf{S}_0^{(i)} \\ &= P_F \left(\mathbf{C}^{(i)} \middle| \mathbf{x}_{i\lambda}, \hat{\mathbf{x}}_1^t; \phi \right) \int \delta \left(\Phi^\theta \left(T(\mathbf{x}_{i\lambda}, \mathbf{C}^{(i)}, \mathbf{S}_0^{(i)}), \lambda, \Delta t \right) - \mathbf{x}_{(i+1)\lambda} \right) p(\mathbf{S}_0^{(i)}) d\mathbf{S}_0^{(i)}, \quad (5) \end{aligned}$$

where Φ^θ is the ODE solver of the state flow model $p_{1|t}^\theta$. This integration induces a Dirac delta term, making direct probability estimation challenging.

Proposition B.1. *Let the initial state $\mathbf{S}_0^{(i)}$ be fixed². Then, the object $\mathbf{x}_t = (\mathcal{C}_t, \mathcal{S}_t)$ is uniquely determined by a given trajectory τ sampled from the compositional flow model π^ϕ .*

²To minimize the influence of the fixed initial state value, we sample the initial state from the noise distribution using the manual random seed which is equal to the size of the current object, such as the number of atoms.

Proof. When $S_0^{(i)}$ is fixed, the ODE solver fully specifies $S_{(i+1)\lambda}$ based on $S_{i\lambda}$, the self-conditioning \hat{x}_1^t , and the sampled component $C^{(i)}$. Consequently, the sequence $(C^{(1)}, \dots, C^{(k(t))})$ directly determines S_t . Given that the generative progress is auto-regressive, there is a one-to-one correspondence between x_t and $(C^{(1)}, \dots, C^{(k(t))})$ for every t .

Building on this determinism, we simplify the forward transition probability $P_F(x_{(i+1)\lambda} | x_{i\lambda}; \phi; p_{1|t}^\theta)$ as:

$$P_F(x_{(i+1)\lambda} | x_{i\lambda}, \hat{x}_1^t; \phi, p_{1|t}^\theta) = P_F(C^{(i)} | x_{i\lambda}, \hat{x}_1^t; \phi) \quad (6)$$

This allows us to write the TB objective in a more tractable form:

$$\begin{aligned} \mathcal{L}_{TB}(\tau) &= \left(\log \frac{Z_\phi \prod_{i=0}^{n-1} P_F(x_{(i+1)\lambda} | x_{i\lambda}; \phi, p_{1|t}^\theta)}{R(x_1) \prod_{i=0}^{n-1} P_B(x_{i\lambda} | x_{(i+1)\lambda}; \phi, p_{1|t}^\theta)} \times \frac{P(x_1 | x_{n\lambda}; p_{1|t}^\theta)}{P(x_{n\lambda} | x_1; p_{1|t}^\theta)} \right)^2 \\ &= \left(\log \frac{Z_\phi \prod_{i=0}^{n-1} P_F(C^{(i)} | x_{i\lambda}; \phi)}{R(x_1)} \right)^2. \end{aligned} \quad (7)$$

By enforcing trajectory balance under these conditions, we ensure that the likelihood of sampling a final object x_1 is proportional to its reward $R(x_1)$, while sidestepping the complexities introduced by the continuous noise integration.

C RELATED WORKS

Synthesis-based generative models. Generative models have emerged as the key paradigm for discovering candidates by bypassing the expensive virtual screening. However, most generative models often render molecules outside the bounds of synthesizable chemical space, limiting their practical use in real-world applications (Gao & Coley, 2020). To address this limitation, several studies (Shen et al., 2024; Guo & Schwaller, 2024) have employed various synthetic complexity estimation methods (Ertl & Schuffenhauer, 2009; Coley et al., 2018; Kim et al., 2023; Neeser et al., 2024; Genheden et al., 2020) as the reward function.

Another promising direction is to design molecules by assembling purchasable building blocks under predefined synthesis protocols. (Gao et al., 2022; Li et al., 2022; Seo et al., 2023; Swanson et al., 2024; Gao et al., 2024). This strategy explicitly constrains the sample space to synthesizable chemical space. More recently, Koziarski et al. (2024); Cretu et al. (2024); Seo et al. (2024) have extended this strategy using GFlowNets, formulating synthesis pathway generation as trajectories of GFlowNets. This effectively explores the chemical space to discover diverse candidate molecules while balancing exploration and exploitation.

Diffusion for sequential data. Diffusion models have recently gained traction for generative modeling of sequential structures in diverse domains, spanning biological sequences Campbell et al. (2024); Stark et al. (2024), videos Ruhe et al. (2024), and language modeling Lou et al. (2024). Campbell et al. (2023) propose a jump process for transitioning between different dimensional spaces to address the variable-dimension nature of data. To exploit the temporal causal dependency in sequences, Ruhe et al. (2024); Zhang et al. (2023) explore frame-level noise schedules for diffusion-based video generation for arbitrary-length frame rollout. Wu et al. (2023); Chen et al. (2024) apply similar ideas of training next-token prediction models while diffusing past ones for applications in planning and language modeling. Most similar to our work for molecule generation using diffusion models are methods that use a separate diffusion process for each sequentially added fragment Peng et al. (2023); Ghorbani et al. (2023); Li et al. (2024).

We provide an extended related works for sequential diffusion for molecular generation and structure-based drug design in App. C.

Sequential diffusion for molecular generation Peng et al. (2023); Ghorbani et al. (2023); Li et al. (2024) sequentially generate molecules fragment by fragment using diffusion models. These methods use a separate diffusion process to generate each 3D fragment graph, with atomic positions

fixed post-generation. They also lack the ability to enforce compositional synthesis constraints during the generation process. Instead, 3DSynthFlow formulates a joint flow process: state flow refines all atomic positions throughout, mitigating the issue with cascading error in position prediction; composition flow sequentially constructs the synthesis pathway, effectively enforcing the synthesis constraint.

Structure-based drug design We categorize structure-based drug design (SBDD) in two main categories: pocket-specific and pocket-conditional following Seo et al. (2024).

The first approach optimizes docking scores for a target. Methods include evolutionary algorithms (Reidenbach, 2024), reinforcement learning (RL) (Zhavoronkov et al., 2019), and GFlowNets (Bengio et al., 2021; Pandey et al., 2025; Koziarski et al., 2024). The drawback of this approach is that each pocket must be optimized individually, which can restrict scalability.

In contrast, the pocket-conditional generation approach produces molecules tailored to any given pocket without the need for extra training. This strategy leverages distribution-based generative models (Ragoza et al., 2022; Peng et al., 2022; Guan et al., 2024; Schneuing et al., 2024b; Qu et al., 2024) that are trained on protein-ligand complex datasets to learn the distribution of ligands suitable for different pockets. Recently, Shen et al. (2024); Seo et al. (2024) adopted pocket-conditioned policy for GFlowNets that generates samples from reward-biased distributions in a zero-shot setting.

D 3DSYNTHFLOW DETAILS

D.1 ACTION SPACE

Following Cretu et al. (2024); Koziarski et al. (2024); Seo et al. (2024), we treat chemical reactions as forward transitions and synthetic pathways as trajectories for molecular generation.

Compared to previous methods, we represent a building block as *synthon*, which is not a complete molecule. The synthons can be connected at the *attachment point* according to the pre-defined connection rules, i.e., reactions \mathcal{R} (See Fig. 2). To prevent a generation trajectory terminating at an incomplete structure, we define two types of synthon according to Liu et al. (2017): *brick* is the synthon including one attachment point, and *linker* is the synthon including two attachment points.




Figure 4: Examples of brick and linker synthons.

We define \mathcal{B} as the entire synthon set, $\mathcal{B}' \subseteq \mathcal{B}$ as the brick synthon set, and $\mathcal{B}_r \subseteq \mathcal{B}$ as the allowable synthon set for reaction $r \in \mathcal{R}$.

At the initial state \mathcal{C}_0 , the model always selects `FirstSynthon`, which samples a brick synthon b from the entire brick synthon set \mathcal{B}' to serve as the starting molecule. For subsequent states \mathcal{C}_t , the model chooses `AddSynthon`, which firstly identifies the available reaction set $\mathcal{R}(\mathcal{C}_t) \subseteq \mathcal{R}$ and then samples the synthon from the available synthon set $\cup_{r \in \mathcal{R}(\mathcal{C}_t)} \mathcal{B}_r$. If the brick synthon is selected, the trajectory is terminated. When the trajectory reaches the maximum length, the model always selects a brick synthon.

In summary, the allowable action space $\mathcal{A}(\mathcal{C})$ for a given compositional state \mathcal{C} is:

$$\mathcal{A}(\mathcal{C}) = \begin{cases} \mathcal{B}' & \text{if } t = 1, \\ \cup_{r \in \mathcal{R}(\mathcal{C})} \mathcal{B}_r & \text{otherwise,} \end{cases} \quad (8)$$

D.2 3DSYNTHFLOW MODEL ARCHITECTURE

D.2.1 STATE FLOW MODEL

To model state flow for predicting ligand docking poses in 3DSynthFlow, we extend the *Semla* architecture introduced by Irwin et al. (2024). *Semla* is a scalable, E(3)-equivariant model originally

designed for 3D molecular generation; We refer reader to the original paper for full details. To adapt it for modeling 3D protein-ligand complexes, we incorporate protein encoding layers and protein-ligand message passing to capture interactions critical for binding, making the following modifications:

Let $P = \{1, \dots, M\}$ index protein residues. Specifically, we use GVP (Jing et al., 2021) to encode the protein pocket, producing an invariant feature vector for each residue $\mathbf{h}_i^{(\text{pro})} \in \mathbb{R}^d$, $\forall i \in P$. The protein residue embeddings, along with their original positions $\mathbf{x}_i^{(\text{pro})}$, are combined with the ligand atom embeddings $(\mathbf{h}_j^{(\text{lig})}, \mathbf{x}_j^{(\text{lig})})$ for pairwise message passing at each modified *Semla* layer $(i, j) \in V \times P$, defined as:

$$(\mathbf{m}_{i,j}^{\text{inv}}, \mathbf{m}_{i,j}^{\text{equi}}) = \Omega_{\theta}^{(\text{pro-lig})}(\tilde{\mathbf{h}}_i, \tilde{\mathbf{h}}_j, \tilde{\mathbf{x}}_i \cdot \tilde{\mathbf{x}}_j), \quad (9)$$

where $\tilde{\mathbf{h}}_i$ and $\tilde{\mathbf{x}}_i$ are normalized linearly projected features. The protein-ligand messages are aggregated with the original ligand-ligand messages and used for attention-based updates. This formulation enables message exchange between residue nodes and atom nodes, ensuring that the model effectively learns protein-ligand interactions critical to binding.

Since the atom type and bond connections are generated by the compositional flow model instead of the state flow model, therefore they remain fixed throughout the process. This means we cannot use techniques such as equivariant-OT to reduce the transport cost (Klein et al., 2023; Song et al., 2023), since they assume interchangeability between newly initialized nodes. We have also tried using harmonic prior from Jing et al. (2023); Stärk et al. (2024) to initialize position with bonding information as prior. We observed improved pose prediction with Gaussian noise during initial training, but no advantage of the harmonic prior over the Gaussian prior at convergence.

D.2.2 COMPOSITIONAL FLOW MODEL

We adopt a model architecture inspired by Bengio et al. (2021); Cretu et al. (2024); Seo et al. (2024), using a graph transformer (Yun et al., 2022) as the backbone f_{θ} , and a multi-layer perceptron (MLP) g_{θ} for action embedding. The graph embedding dimension is d_1 and the synthon embedding dimension is d_2 . The GFlowNet condition such as temperature β or multi-objective weights are encoded in condition vector c .

To capture spatial relationships between the protein and ligand, the model uses a 3D protein-ligand complex graph for each state. The protein is represented as a residual graph, where each node corresponds to a C_{α} atom and connects to its K-nearest neighbors. Each ligand atom also connects to its K-nearest protein nodes. All pairwise distances between nodes are encoded in the graph edges, enabling the model to reason about the 3D structure. We represent each state s as a molecular graph s and include a GFlowNet condition vector c .

Initial synthon selection. For the initial state $s = s_0$, the model always selects `FirstSynthon` to sample brick synthon $b \in \mathcal{B}'$ for the starting molecule with $\text{MLP}_{\text{AddSynthon}} : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2}$.

$$F_{\theta}(s_0, b, c) = \text{MLP}_{\text{FirstSynthon}}(f_{\theta}(s_0, c)) \odot g_{\theta}(b) \quad (10)$$

Adding synthon selection. For the later states $s \neq s_0$, the model selects `AddSynthon` to sample allowable brick or linker synthon $b \in \cup_{r \in \mathcal{R}(s)} \mathcal{B}_r$ with $\text{MLP}_{\text{AddSynthon}} : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2}$. We note that the reaction type information is included in synthon embedding:

$$F_{\theta}(s, (r, b), c) = \text{MLP}_{\text{AddSynthon}}(f_{\theta}(s, c)) \odot g_{\theta}(r, b). \quad (11)$$

Synthon masking. The state flow model is trained on the data distribution of active ligands. This may sometimes not work well for out-of-distribution molecules that are much larger than the training data. Therefore, we restrict the actions which make the state to have more than 40 atoms. This process can be performed simply in a synthon-based action space.

E RELATED WORKS

Synthesis-based generative models. Generative models have emerged as the key paradigm for discovering candidates by bypassing the expensive virtual screening. However, most generative

models often render molecules outside the bounds of synthesizable chemical space, limiting their practical use in real-world applications (Gao & Coley, 2020). To address this limitation, several studies (Shen et al., 2024; Guo & Schwaller, 2024) have employed various synthetic complexity estimation methods (Ertl & Schuffenhauer, 2009; Coley et al., 2018; Kim et al., 2023; Neeser et al., 2024; Genheden et al., 2020) as the reward function.

Another promising direction is to design molecules by assembling purchasable building blocks under predefined synthesis protocols. (Gao et al., 2022; Li et al., 2022; Seo et al., 2023; Swanson et al., 2024; Gao et al., 2024). This strategy explicitly constrains the sample space to synthesizable chemical space. More recently, Koziarski et al. (2024); Cretu et al. (2024); Seo et al. (2024) have extended this strategy using GFlowNets, formulating synthesis pathway generation as trajectories of GFlowNets. This effectively explores the chemical space to discover diverse candidate molecules while balancing exploration and exploitation.

Diffusion for sequential data. Diffusion models have recently gained traction for generative modeling of sequential structures in diverse domains, spanning biological sequences Campbell et al. (2024); Stark et al. (2024), videos Ruhe et al. (2024), and language modeling Lou et al. (2024). Campbell et al. (2023) propose a jump process for transitioning between different dimensional spaces to address the variable-dimension nature of data. To exploit the temporal causal dependency in sequences, Ruhe et al. (2024); Zhang et al. (2023) explore frame-level noise schedules for diffusion-based video generation for arbitrary-length frame rollout. Wu et al. (2023); Chen et al. (2024) apply similar ideas of training next-token prediction models while diffusing past ones for applications in planning and language modeling. Most similar to our work for molecule generation using diffusion models are methods that use a separate diffusion process for each sequentially added fragment Peng et al. (2023); Ghorbani et al. (2023); Li et al. (2024).

We provide an extended related works for sequential diffusion for molecular generation and structure-based drug design in App. C.

F EXPERIMENTAL DETAILS

F.1 PROBLEM SETUP

Overall setup. We utilize 1.2M molecular fragments from the Enamine Catalog and 38 bimolecular Enamine synthesis protocols from Gao et al. (2024). To ensure tractability, molecular generation is limited to two synthesis steps, consistent with Enamine REAL Space guidelines (Grygorenko et al., 2020). To estimate synthesizability, we employ the retrosynthetic analysis tool AiZynthFinder (Genheden et al., 2020).

We pre-train the state flow model for 3D pose prediction on the CrossDocked dataset Francoeur et al. (2020) with LIT-PCBA targets removed (see App. F.3.2). We sample a synthesis pathway for each ligand using reaction rules to fragment the ligand (see Fig. 2). This exposes the state flow model to intermediate states - allowing it to predict their poses for partial structures during compositional flow model training (see App. F.2.2 for motivation).

F.1.1 POCKET-SPECIFIC OPTIMIZATION SETUP

Setup. We follow Seo et al. (2024) in both the reward function and evaluation protocol. To estimate the binding affinity, we use the GPU-accelerated docking tool UniDock (Yu et al., 2023). Each method generates up to 64,000 molecules for each protein target. To prevent *reward hacking* by increasing molecular size to improve the docking score, QED (Bickerton et al., 2012) is jointly optimized, and we set a generation of 40 heavy atoms. This ensure that the generated molecules have drug-like properties. Specifically, we optimize $R(x) = w_1 \text{QED}(x) + w_2 \widehat{\text{Vina}}(x)$, where $\widehat{\text{Vina}}(x)$ is a normalized docking score. The parameters w_1 and w_2 serve as conditions of multi-objective GFlowNets (Jain et al., 2023b), and are set to 0.5 for non-GFlowNet baselines.

Generated molecules are filtered based on property constraints ($\text{QED} > 0.5$), and the top 100 diverse modes are selected according to docking scores, ensuring structural diversity with a Tanimoto distance threshold of 0.5³. Finally, we evaluated the average **Vina Docking Score** and **AiZynthFinder**

³Since we select top 100 modes filtering for similarity, diversity is not reported in this section.

Success Rate of the selected molecules. **Ligand Efficiency** as computed by (Vina / number of heavy atoms) is also reported to confirm the docking score improvement do not arise from simply increase in molecular size.

Baselines. We compare our approach against several synthesis-based methods, including a genetic algorithm **SynNet** (Gao et al., 2022), a conditional generative model **BBAR** (Seo et al., 2023), and multiple GFlowNets. We consider two settings of fragment-based GFlowNets, with and without synthetic accessibility score (SA; Ertl & Schuffenhauer, 2009) objective (**FragGFN**, **FragGFN+SA**), and three different synthesis-based GFlowNets (**RGFN**, **SynFlowNet**, **RxnFlow**) (Kozlarski et al., 2024; Cretu et al., 2024; Seo et al., 2024).

F.1.2 POCKET-CONDITIONAL GENERATION SETUP

Setup. Our method generalizes to pocket-conditional generation problem setting (Peng et al., 2022; Guan et al., 2023; Schneuing et al., 2024a), enabling the design of binders for unseen pockets with a single model and no additional oracle calls. We follow the same pocket-conditional experimental setup used in TacoGFN (Shen et al., 2024) and RxnFlow (Seo et al., 2024). 3DSynthFlow adopt the pre-trained proxy from TacoGFN trained on CrossDock2020 train set, which leverages pharmacophore representation (Seo & Kim, 2024), to compute rewards for training a pocket-conditional policy.

Each method generates 100 molecules for each of the 100 test pockets in the CrossDocked2020 benchmark (Francoeur et al., 2020) and are evaluated on these additional metrics compared to the pocket-specific setting: **Diversity** represents the average pairwise Tanimoto distance computed from ECFP4 fingerprints (Morgan, 1965). We also report **Validity (%)**, the proportion of unique molecules that can be parsed by RDKit, and **Time (sec.)**, the average duration required to sample 100 molecules.

Baselines. We compare 3DSynthFlow with state-of-the-art distribution learning-based models trained on a synthesizable drug set, including the autoregressive model **Pocket2Mol** (Peng et al., 2022), and diffusion-based methods: **TargetDiff** (Guan et al., 2023), **DiffSBDD** (Schneuing et al., 2024b), **DecompDiff** (Guan et al., 2024), and **MolCRAFT** (Qu et al., 2024). We further include comparisons with optimization-based approaches: the fragment-based method **TacoGFN** (Shen et al., 2024) and the reaction-based method **RxnFlow** (Seo et al., 2024). To ensure a fair evaluation of distribution learning-based methods, we use the docking proxy trained on CrossDocked training set. To balance exploration and exploitation during sampling, we vary the reward exponentiation parameter β of **3DSynthFlow**: **low** β (sampled from Uniform(1, 64)), **medium** β (Uniform(32, 64)), and **high** β (Uniform(48, 64)).

F.2 TRAINING DETAILS

F.2.1 3DSYNTHFLOW HYPERPARAMETERS

Table 3: **Default hyperparameters** used for compositional structure.

Hyperparameters	Values
Time per action λ	0.3
Interpolation window t_{window}	0.4
Maximum decomposed part	3
Minimum decomposed fragment size	5
Minimum trajectory length	2 (minimum reaction steps: 1)
Maximum trajectory length	3 (maximum reaction steps: 2)

Here we define the hyperparameters used for 3DSynthFlow:

1. **Time per action (λ):** Defines the time interval between adding each additional synthon.
2. **Interpolation window (t_{window}):** Specifies the fixed time window that affects the noise scheduling of the continuous states of each component.
3. **Maximum decomposed part:** Determines the maximum number of synthons a molecule can be decomposed into, preventing molecules from being associated with excessive synthesis steps.

4. **Minimum decomposed fragment size:** Specifies the minimum number of atoms that each synthon product must contain when a molecule is decomposed according to reaction rules. This ensures that synthons are of realistic size when decomposing CrossDocked ligands for training a pose prediction model.
5. **Minimum trajectory length:** Defines the minimum trajectory length for sampling composition steps in 3DSynthFlow.
6. **Maximum trajectory length:** Specifies the maximum trajectory length that can be reached.

We further exam the effect of performance for different setting of time scheduling, varying λ and t_{window} in Appendix G.3.

Table 4: **Default hyperparameters** used for State flow model.

Hyperparameters	Values
Number of protein layers	6
Noise prior	Gaussian
Time alpha α	1.0
Time beta β	1.0
Time step Δt	0.0125

For parameterizing the state flow model for pose prediction, we use the same hyperparameter from the official repository ⁴. We show the additional hyperparameters in Table. 4 and provide an explanation for each below:

1. **Number of protein layers:** the number of layers with protein-ligand message passing. This must be less or equal to the total number of layers.
2. **Noise prior:** Defines the noise distribution which initial state values are drawn from.
3. **Time alpha:** Defines the alpha parameter of the Beta distribution which time is sampled from.
4. **Time beta:** Defines the beta parameter of the Beta distribution which time is sampled from.
5. **Time step:** Defines the time step for inference.

Composition flow model In this work, we set most parameters to the default values⁵ for all GFlowNet baselines (Some of the parameters are in Table. 5). However, since 3DSynthFlow is built from RxnFlow, it follows some important parameters of RxnFlow except for maximum trajectory length and state embedding size. We note that our state embedding dimension is smaller than other GFlowNet baselines.

Table 5: Default hyperparameters used in all GFlowNets. The settings are from `seh_frag.py` (Bengio et al., 2021)

Hyperparameters	Values
GFN temperature β	Uniform(0, 64)
Sampling tau (EMA factor)	0.9
Learning rate (Z)	10^{-3}
Learning rate (P_F, P_B)	10^{-4}
State embedding dimension	128

For action space subsampling, we randomly subsample 1% actions for `AddFirstReactant` and each bi-molecular reaction template $r \in \mathcal{R}_2$. However, for bi-molecular reactions with small possible reactant block sets $\mathcal{B}_r \in \mathcal{B}$, the memory benefit from the action space subsampling is small while a variance penalty is large. Therefore, we set the minimum subsampling size to 100 for each bi-molecular reaction, and the action space subsampling is not performed when the number of actions is smaller than 100.

⁴Follow SemlaFlow’s hyperparameter settings in `train.py` at <https://github.com/rssrwn/sem-la-flow/blob/main/sem-la-flow/train.py>

⁵Follow default hyperparameter settings in `seh_frag.py` and `seh_frag_moo.py` at <https://github.com/recursionpharma/gflownet/blob/trunk/src/gflownet/tasks>

Table 6: Specific hyperparameters used in 3DSynthFlow training. The parameters are from RxnFlow (Seo et al., 2024), except for the maximum trajectory length (4 for RxnFlow) and state embedding size (128 for RxnFlow and other GFlowNets).

Hyperparameters	Values
Minimum trajectory length	2 (minimum reaction steps: 1)
Maximum trajectory length	3 (maximum reaction steps: 2)
State embedding dimension	64
Action embedding dimension	64
Action space subsampling ratio	1%
Train random action probability	5%

The number of actions for each action type is imbalanced, and the number of reactant blocks (\mathcal{B}_r) for each bi-molecular reaction template r is also imbalanced. This can make some rare action categories not being sampled during training. We empirically found that `ReactBi` action were only sampled during 20,000 iterations (1.28M samples) in a toy experiment that uses one bi-molecular reaction template and 10,000 building blocks in some random seeds. Therefore, we set the random action probability as the default of 5%, and the model uniformly samples each action category in the random action sampling. This prevents incorrect predictions by ensuring that the model experiences trajectories including rare actions. We note that this random selection is only performed during model training.

F.2.2 DETAILS OF STATE FLOW MODEL TRAINING

To expose the state flow model to realistic partial ligand structures, we decompose each CrossDocked ligand into up to three fragments using 38 bimolecular Enamine synthesis protocols, defined by reaction SMARTS patterns. For each molecule, we randomly select a fragment ordering and sample a time step t , so the model observes the molecule at varying stages of assembly (e.g., only fragment A at $t = 1$, fragments A and B at $t = 2$, etc.). This mimics the compositional assembly process used during generation and teaches the model to predict physically plausible docking conformations for incomplete ligands.

Importantly, the decomposition is not intended to yield commercially purchasable Enamine synthons. Instead, the use of Enamine protocols ensures the resulting fragments are chemically meaningful and synthetically relevant, even if they do not match existing catalog synthons. This allows us to use any CrossDocked molecule for training while staying aligned with the synthetic priors used at generation time.

Across the CrossDocked dataset, we find that over 93% of molecules can be decomposed into at least two fragments using this protocol, and over 87% can be decomposed into three. This high decomposition rate enables robust learning from partially constructed molecules and ensures strong coverage of fragment combinations encountered during inference.

F.2.3 TRAINING TIME AND COMPUTATIONAL EFFICIENCY

The state flow model (i.e., the pocket-conditional pose predictor) was trained for 100 epochs with a batch size of 32 on 4 L40 GPUs (48GB), requiring a total of 18.4 hours. Since the state flow model is trained on the CrossDocked dataset and is reused across different test pockets, it incurs only a one-time computational cost. We plan to release the pretrained weights, so that users will need to train only the composition flow model tailored to their custom reward function and target.

In contrast, the composition flow model is trained individually for each pocket for 1,000 steps with a batch size of 64 and 80 flow matching steps. Training with GPU-accelerated docking takes between 12 and 20 hours (depending on the target) on a single A4000 GPU (16GB), making the computational requirement accessible for most practical drug discovery campaigns. Moreover, the composition flow model can be trained in a pocket-conditioned manner to enable zero-shot molecule sampling for any target pocket, thereby converting the training into a one-time cost in this setting.

Table 7: Summary of Training Time for 3DSynthFlow Components.

Model Component	Hardware	Batch Size	Number of Iterations	Training Time
State Flow Model	4 × L40 (48GB)	32	100 epochs	18.4 hours
Composition Flow Model	1 × A4000 (16GB)	64	1,000 steps	12–20 hours

F.3 DATASETS

F.3.1 LIT-PCBA POCKETS

Table. 8 describes the protein information used in pocket-specific optimization with UniDock, which is performed on Sec. 4. We follow the same procedure used in pocket extraction for the CrossDock dataset: taking all residue of the protein within 10 Å radius to the reference ligand as the binding pocket.

Table 8: **The basic target information** of the LIT-PCBA dataset and PDB entry used in this work.

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

F.3.2 CROSSDOCKED2020

We train the State flow model of 3DSynthFlow on the commonly used **CrossDocked** dataset (Francoeur et al., 2020). We apply the splitting and processing protocol on the CrossDocked dataset to obtain the same training split of protein-ligand pairs as previous methods (Luo et al., 2021; Peng et al., 2022). To ensure fairness against baseline methods on the LIT-PCBA benchmark, we remove all PDB ids associated with LIT-PCBA proteins from the training set. Since we co-design 3D binding pose and synthesis pathway, unlike the pervious 2D-based methods, we can leverage this dataset for training the auxiliary State flow pose prediction model.

F.4 BASELINES

SynNet, BBAR We reused the values reported in Seo et al. (2024).

FragGFN, RGFN, RxnFlow. All GFlowNet baselines share the same training parameters under the multi-objective GFlowNet (Jain et al., 2023b) setting. We also reused the values reported in Seo et al. (2024).

SynFlowNet. There are two versions for SynFlowNet (Cretu et al., 2024): 2024.3 and 2024.8. For version 2024.3, we reused the values in Seo et al. (2024). For version 2024.8, we followed the processes and settings according to the original paper and official code repository ⁶. To construct the action space, we randomly selected 10,000 building blocks from Enamine Global Stock

⁶Follow SynFlowNet’s hyperparameter settings in `reactions_task.py` at <https://github.com/mirunacrt/synflownet/tree/46a4acabd2255eb964c317ffbb86b743a13a4685>

(v2025.01.11) with 105 reaction templates. We trained SynFlowNet using backward policy learning maximum likelihood, maximum trajectory length to 3, and action embedding with Morgan fingerprint (Morgan, 1965). Finally, we set the training parameters used in other GFlowNet baselines (See Table. 5). The training code and data used in the benchmark study are included in the supplementary materials.

G ADDITIONAL RESULTS

G.1 REWARD COMPUTATION WITHOUT FULL DOCKING

Table 9: Ablation study on reward setting for Vina docking scores. Evaluated on the ADH1 pocket. Results are averaged over 3 runs.

Method	Reward setting	Docking score (\downarrow)
RxnFlow	Full docking	-11.26 (\pm 0.07)
3DSynthFlow	Full docking	-11.97 (\pm 0.07)
3DSynthFlow	Local opt.	-11.44 (\pm 0.06)
3DSynthFlow _{finetuned}	Local opt.	-11.62 (\pm 0.02)

The most computational and time-intensive process in the training process is often docking the generated candidates with molecular docking for evaluation. Glide, a standard docking software in the industry, takes around 6 minutes per compound in its most accurate setting (Friesner et al., 2004). This constraint necessitates either reducing the number of Oracle queries or trading speed for less accurate docking settings. Docking scores are typically computed via **Full docking**, which performs a full search for the optimal binding pose at a high computational cost. However, it can also be computed via **Local optimization** for significantly faster computation, if an existing binding pose is available.⁷

In this setting, we investigate directly using the final predicted pose from 3DSynthFlow to compute reward using local optimization compared to using the full docking score. In the 3DSynthFlow_{finetuned} setting for local opt., we first finetune the pose predictor on re-docked poses from the first 9,600 sampled molecules to improve binding pose prediction, thereby enabling accurate local optimization docking scores.

From Table. 9, we see that regardless of the scoring function used or whether we perform fine tuning, 3DSynthFlow outperformed the best 2D baseline RxnFlow, confirming the benefit of 3D structure co-design. Interestingly, 3DSynthFlow using local opt. do not perform as well as 3DSynthFlow trained with signals from full docking in both settings. While fine tuning the pose prediction module helps, the gap between full docking and local opt. is not fully closed. Empirically, we find the poses predicted from 3DSynthFlow sometimes contain steric clashes, leading to inaccurate estimation of the reward signal. Accordingly, improving the pose prediction module emerges as a key next step to reduce steric clashes and enhance the accuracy of local optimization docking scores.

By co-designing binding pose and molecule, 3DSynthFlow using local opt. can bypass the computational burden of full docking which 2D generative methods are subjected to, while surpassing their performance. This is a key advantage for 3DSynthFlow in cases where accurate full docking is prohibitively expensive for a large number of ligands.

G.2 EFFECT OF FLOW MATCHING STEPS

We further analyzed how the number of flow matching steps impacts performance using the ALDH1 target with the Vina reward. As shown in Table 10, performance slightly improves with increasing flow matching steps and appears to saturate around 40–60 steps. This marginal improvement may stem from the fact that the pose prediction module’s primary role is to provide a spatial context between intermediate molecules and the pocket; hence, extremely precise pose predictions have limited additional impact on model decisions.

⁷Scoring a predicted pose using the local-optimization setting in AutoDock Vina is 7 \times faster than performing full docking. This speedup is even greater for more accurate docking programs, since the vast majority of time is spent on pose searching rather than scoring.

Table 10: Effect of flow matching steps on performance for the ALDH1 target with the Vina reward. We report both the average docking score (Avg Vina) over all generated molecules and the top 100 docking scores (Top 100 Vina), with lower values indicating better binding. Training time is reported in seconds per iteration, and sampling time is reported in seconds per molecule.

Flow Matching Steps	Avg Vina (\downarrow)	Top 100 Vina (\downarrow)	Training Time (sec/iter)	Sampling Time (sec/mol)
10	-10.28 ± 0.32	-14.27 ± 0.59	33	0.053
20	-10.24 ± 0.18	-14.38 ± 0.22	34	0.080
40	-10.40 ± 0.13	-14.51 ± 0.27	39	0.123
60	-10.50 ± 0.14	-14.57 ± 0.24	44	0.160
80	-10.44 ± 0.18	-14.53 ± 0.16	49	0.199

G.3 ABLATION STUDY ON TIME SCHEDULING OF 3DSYNTHFLOW

We conducted an ablation study to assess the effect of time scheduling in state flow training. Specifically, we compared three scheduling strategies:

- **No overlap**: strictly autoregressive denoising - each synthon denoised after the previous one is completed. ($\lambda = 0.33, t_{window} = 0.33$).
- **Overlapping**: partial overlap of synthon denoising ($\lambda = 0.3, t_{window} = 0.4$).

To highlight the impact of noise scheduling on the final pose, we report the average local-optimized Vina docking scores across different training iterations for the ALDH1 target over 3 seeds:

# of mol explored	10,000	20,000	30,000
No overlap	-5.68 ± 0.29	-6.33 ± 0.26	-7.02 ± 0.34
Overlapping	-6.28 ± 0.22	-7.28 ± 0.21	-7.22 ± 0.12

Table 11: Ablation study on time scheduling for state flow training on the ALDH1 target.

3DSynthFlow’s **overlapping** strategy, where positions are refined as synthons are added, clearly outperforms conventional autoregressive approaches - **No overlap**, adopted in previous works (See Appendix C).

G.4 STATE SPACE SIZE ESTIMATION

We estimate the sample space size based on the number of synthetic steps: 10^{11} molecules with a single reaction step, 10^{17} molecules with two reaction steps, and 10^{23} molecules with three reaction steps. In our experiments, we employed up to two reaction steps according to Enamine REAL. The resulting state space size is comparable to RGFN (up to 4 steps with 8,350 blocks) and SynFlowNet (up to 3 steps with 200k blocks).

To assess the breadth of building block (BB) exploration, we analyzed the number of unique BBs encountered during training across the first five LIT-PCBA targets. Our model explored an average of approximately 55,000 unique BBs within 1,000 training iterations using a batch size of 64. This indicates significantly broader exploration compared to SynFlowNet, which reported around 15,000 unique BBs over 8,000 iterations with a batch size of 8.

Table 12: Number of unique building blocks explored during training across 5 LIT-PCBA targets.

Target	ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Number of Unique BBs	45520 ± 7876	48644 ± 1983	55211 ± 5611	58097 ± 8529	69400 ± 5259

G.5 POSECHECK PROTEIN-LIGAND INTERACTIONS

Table 13: **PoseCheck protein-ligand interactions.** We report two versions of SynFlowNet (v2024.05^a and v2024.10^b). Averages and standard deviations over 3 runs, for the top 100 diverse modes per LIT-PCBA pocket. The best results in each column are in **bold**.

Method	PoseCheck Metrics (\uparrow)				Sum
	H-Bond Acceptors	H-Bond Donors	Van der Waals	Hydrophobic	
SynFlowNet ^a	0.22 (\pm 0.01)	0.11 (\pm 0.01)	9.31 (\pm 0.05)	10.41 (\pm 0.05)	20.05
SynFlowNet ^b	0.22 (\pm 0.03)	0.10 (\pm 0.01)	8.38 (\pm 0.05)	9.44 (\pm 0.06)	18.14
RGFN	0.19 (\pm 0.01)	0.11 (\pm 0.01)	9.19 (\pm 0.28)	10.25 (\pm 0.12)	19.73
RxnFlow	0.22 (\pm 0.00)	0.10 (\pm 0.01)	9.63 (\pm 0.09)	10.67 (\pm 0.10)	20.62
3DSynthFlow	0.27 (\pm 0.03)	0.12 (\pm 0.01)	10.33 (\pm 0.13)	11.15 (\pm 0.05)	21.87

Lastly, we evaluate the protein-ligand interactions of the top 100 generated molecules using PoseCheck Harris et al. (2023). Since baselines do not generate 3D poses, we assess all methods using their redocked structures for consistency. PoseCheck evaluates four key protein-ligand interactions: H-bond acceptors, H-bond donors, van der Waals contacts, and hydrophobic effects.

Hydrogen bonds (H-bonds) are the most important specific interactions in protein-ligand recognition Bissantz et al. (2010) and require precise geometric alignment to form Brown (1976). Unlike hydrophobic and van der Waals interactions, which are non-directional and broadly applicable, H-bonds require strict atomic alignment, reinforcing the need for accurate 3D molecular modeling.

Notably, 3DSynthFlow achieves the highest H-bond acceptor and donor counts, outperforming all baselines. This improvement suggests that co-designing 3D structure and synthesis pathways enhances the geometric alignment of polar functional groups, leading to more stable and specific protein-ligand interactions. In addition, 3DSynthFlow-generated molecules also result in more hydrophobic and van der Waals interactions to baselines, as shown in Table. 13, further enhancing binding stability.

G.6 SYNTHESIS SUCCESS RATE RESULT ON LIT-PCBA

Table 14: The average success rate and synthetic steps estimated from AiZynthFinder for all 15 LIT-PCBA protein targets.

Method	Success Rate (% , \uparrow)	Synthesis Steps (\downarrow)
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	48.04	3.15

Table. 14 shows that 3DSynthFlow attains comparable synthesizability to 2D baselines, demonstrating its generated molecules remain likely to be synthesizable. While the main contribution of our work is 3D co-design rather than improving synthesizability, these results confirm that our synthon-based 3D generation remains effective, with minimal trade-offs in synthetic feasibility.

G.7 FULL SAMPLING EFFICIENCY RESULTS

Table 15: Average number of diverse modes discovered versus the number of molecules explored. The best results are in bold.

Target	Method	Number of Molecules Explored				
		1,000	5,000	10,000	30,000	64,000
ADRB2	RxnFlow	3.0 (± 1.4)	21.5 (± 4.9)	52.0 (± 24.0)	500.0 (± 137.2)	1282.5 (± 234.1)
	3DSynthFlow	16.0 (± 5.7)	124.0 (± 2.8)	318.5 (± 60.1)	1609.5 (± 939.7)	5378.0 (± 1492.0)
ALDH1	RxnFlow	4.5 (± 2.1)	26.5 (± 7.8)	73.5 (± 33.2)	472.5 (± 99.7)	1240.0 (± 75.0)
	3DSynthFlow	10.0 (± 5.7)	198.0 (± 83.4)	623.0 (± 291.3)	3069.5 (± 1885.9)	8838.0 (± 5152.0)
ESR_ago	RxnFlow	5.0 (± 4.2)	17.0 (± 4.2)	44.0 (± 8.5)	440.5 (± 60.1)	1174.0 (± 100.4)
	3DSynthFlow	8.0 (± 4.2)	51.0 (± 2.8)	234.5 (± 174.7)	1498.0 (± 1306.7)	3738.0 (± 1861.1)
ESR_antago	RxnFlow	3.0 (± 2.8)	14.5 (± 0.7)	28.0 (± 0.0)	218.0 (± 21.2)	559.0 (± 50.9)
	3DSynthFlow	6.5 (± 7.8)	53.5 (± 9.2)	184.5 (± 40.3)	1083.5 (± 113.8)	3010.0 (± 264.5)
FEN1	RxnFlow	3.5 (± 0.7)	29.0 (± 18.4)	75.5 (± 57.3)	372.0 (± 46.7)	1057.5 (± 38.9)
	3DSynthFlow	10.0 (± 4.2)	92.0 (± 97.6)	219.5 (± 195.9)	774.5 (± 323.1)	3276.5 (± 451.8)
Average	RxnFlow	3.8	21.7	54.6	400.6	1062.6
	3DSynthFlow	10.1	103.7	316.0	1607.0	4848.1

G.8 FULL RESULTS FOR LIT-PCBA TARGET

Table 16: Average Vina docking score for top-100 diverse modes generated against 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 (± 0.17)	-12.25 (± 0.43)	-11.31 (± 0.07)	-11.25 (± 0.12)	-7.92 (± 0.07)
Category	Method	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	-9.89 (± 0.22)	-11.42 (± 0.26)	-10.19 (± 0.19)	-9.69 (± 0.15)	-12.03 (± 0.27)
Category	Method	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	-11.22 (± 0.12)	-13.73 (± 0.25)	-10.26 (± 0.14)	-8.49 (± 0.27)	-12.97 (± 0.20)

Table 17: Average ligand efficiency for top-100 diverse modes generated against 15 LIT-PCBA targets. The best results are in bold.

Category	Method	Average Ligand Efficiency (\downarrow)				
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Fragment	FragGFN	0.410 (\pm 0.006)	0.368 (\pm 0.007)	0.347 (\pm 0.003)	0.358 (\pm 0.002)	0.246 (\pm 0.004)
	FragGFN+SA	0.406 (\pm 0.007)	0.374 (\pm 0.023)	0.369 (\pm 0.003)	0.345 (\pm 0.024)	0.210 (\pm 0.004)
Reaction	SynNet	0.274 (\pm 0.041)	0.272 (\pm 0.006)	0.317 (\pm 0.005)	0.289 (\pm 0.020)	0.196 (\pm 0.003)
	BBAR	0.412 (\pm 0.006)	0.401 (\pm 0.008)	0.380 (\pm 0.001)	0.387 (\pm 0.003)	0.257 (\pm 0.003)
	SynFlowNet ^a	0.401 (\pm 0.006)	0.380 (\pm 0.007)	0.361 (\pm 0.003)	0.361 (\pm 0.004)	0.247 (\pm 0.004)
	Synflownet ^b	0.380 (\pm 0.021)	0.362 (\pm 0.016)	0.351 (\pm 0.008)	0.349 (\pm 0.005)	0.234 (\pm 0.007)
	RGFN	0.393 (\pm 0.005)	0.357 (\pm 0.004)	0.346 (\pm 0.002)	0.344 (\pm 0.002)	0.241 (\pm 0.001)
	RxnFlow	0.412 (\pm 0.005)	0.396 (\pm 0.005)	0.375 (\pm 0.002)	0.380 (\pm 0.004)	0.246 (\pm 0.001)
3D Reaction	3DSynthFlow	0.438 (\pm 0.015)	0.409 (\pm 0.020)	0.385 (\pm 0.006)	0.396 (\pm 0.010)	0.247 (\pm 0.000)
Category	Method	GBA	IDH1	KAT2A	MAPK1	MTORC1
Fragment	FragGFN	0.333 (\pm 0.018)	0.367 (\pm 0.009)	0.322 (\pm 0.008)	0.302 (\pm 0.002)	0.354 (\pm 0.005)
	FragGFN+SA	0.318 (\pm 0.005)	0.369 (\pm 0.020)	0.298 (\pm 0.020)	0.294 (\pm 0.015)	0.355 (\pm 0.027)
Reaction	SynNet	0.244 (\pm 0.013)	0.281 (\pm 0.016)	0.294 (\pm 0.042)	0.226 (\pm 0.004)	0.316 (\pm 0.035)
	BBAR	0.336 (\pm 0.002)	0.382 (\pm 0.005)	0.332 (\pm 0.003)	0.320 (\pm 0.002)	0.385 (\pm 0.004)
	SynFlowNet ^a	0.330 (\pm 0.004)	0.368 (\pm 0.002)	0.327 (\pm 0.003)	0.305 (\pm 0.002)	0.368 (\pm 0.002)
	SynFlowNet ^b	0.324 (\pm 0.007)	0.360 (\pm 0.013)	0.309 (\pm 0.004)	0.297 (\pm 0.011)	0.361 (\pm 0.009)
	RGFN	0.310 (\pm 0.002)	0.351 (\pm 0.003)	0.310 (\pm 0.003)	0.298 (\pm 0.002)	0.346 (\pm 0.004)
	RxnFlow	0.327 (\pm 0.004)	0.378 (\pm 0.001)	0.330 (\pm 0.001)	0.313 (\pm 0.001)	0.370 (\pm 0.001)
3D Reaction	3DSynthFlow	0.344 (\pm 0.021)	0.394 (\pm 0.012)	0.322 (\pm 0.010)	0.314 (\pm 0.012)	0.376 (\pm 0.009)
Category	Method	OPRK1	PKM2	PPARG	TP53	VDR
Fragment	FragGFN	0.352 (\pm 0.004)	0.442 (\pm 0.008)	0.319 (\pm 0.007)	0.307 (\pm 0.005)	0.394 (\pm 0.006)
	FragGFN+SA	0.327 (\pm 0.014)	0.440 (\pm 0.009)	0.303 (\pm 0.013)	0.248 (\pm 0.025)	0.390 (\pm 0.020)
Reaction	SynNet	0.298 (\pm 0.039)	0.296 (\pm 0.005)	0.253 (\pm 0.031)	0.211 (\pm 0.031)	0.359 (\pm 0.015)
	BBAR	0.370 (\pm 0.006)	0.442 (\pm 0.004)	0.326 (\pm 0.007)	0.288 (\pm 0.005)	0.409 (\pm 0.002)
	SynFlowNet ^a	0.359 (\pm 0.004)	0.427 (\pm 0.003)	0.317 (\pm 0.002)	0.287 (\pm 0.008)	0.393 (\pm 0.003)
	SynFlowNet ^b	0.355 (\pm 0.008)	0.410 (\pm 0.018)	0.296 (\pm 0.010)	0.282 (\pm 0.005)	0.380 (\pm 0.005)
	RGFN	0.349 (\pm 0.001)	0.405 (\pm 0.002)	0.307 (\pm 0.002)	0.271 (\pm 0.001)	0.381 (\pm 0.002)
	RxnFlow	0.369 (\pm 0.007)	0.436 (\pm 0.005)	0.319 (\pm 0.002)	0.289 (\pm 0.003)	0.405 (\pm 0.002)
3D Reaction	3DSynthFlow	0.373 (\pm 0.015)	0.460 (\pm 0.019)	0.328 (\pm 0.007)	0.293 (\pm 0.013)	0.408 (\pm 0.012)

Table 18: Average success rate of AiZynthFinder for top-100 diverse modes generated against 15 LIT-PCBA targets. The best results are in bold.

Category	Method	AiZynthFinder Success Rate (% , \uparrow)				
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Fragment	FragGFN	4.00 (\pm 3.54)	3.75 (\pm 1.92)	1.00 (\pm 1.00)	3.75 (\pm 1.92)	0.25 (\pm 0.43)
	FragGFN+SA	5.75 (\pm 1.48)	6.00 (\pm 2.55)	4.00 (\pm 2.24)	1.00 (\pm 0.00)	0.00 (\pm 0.00)
Reaction	SynNet	54.17 (\pm 7.22)	50.00 (\pm 0.00)	50.00 (\pm 0.00)	25.00 (\pm 25.00)	50.00 (\pm 0.00)
	BBAR	21.25 (\pm 5.36)	19.50 (\pm 3.20)	17.50 (\pm 1.50)	19.50 (\pm 3.64)	20.00 (\pm 2.12)
	SynFlowNet ^a	52.75 (\pm 1.09)	57.00 (\pm 6.04)	53.75 (\pm 9.52)	56.50 (\pm 2.29)	53.00 (\pm 8.92)
	SynFlowNet ^b	56.50 (\pm 6.58)	56.00 (\pm 3.08)	61.00 (\pm 2.74)	64.50 (\pm 9.86)	60.75 (\pm 3.77)
	RGFN	46.75 (\pm 6.86)	47.50 (\pm 4.06)	50.25 (\pm 2.17)	49.25 (\pm 4.38)	48.50 (\pm 6.58)
	RxnFlow	60.25 (\pm 3.77)	63.25 (\pm 3.11)	71.25 (\pm 4.15)	66.50 (\pm 4.03)	65.50 (\pm 4.09)
3D Reaction	3DSynthFlow	47.67 (\pm 19.15)	49.33 (\pm 3.40)	50.33 (\pm 5.73)	61.33 (\pm 5.56)	52.67 (\pm 5.79)
Category	Method	GBA	IDH1	KAT2A	MAPK1	MTORC1
Fragment	FragGFN	5.00 (\pm 4.24)	4.50 (\pm 1.66)	1.25 (\pm 0.83)	0.75 (\pm 0.83)	2.75 (\pm 1.30)
	FragGFN+SA	3.00 (\pm 1.00)	4.50 (\pm 4.97)	1.50 (\pm 0.50)	3.25 (\pm 1.48)	3.50 (\pm 2.50)
Reaction	SynNet	50.00 (\pm 0.00)	50.00 (\pm 0.00)	45.83 (\pm 27.32)	50.00 (\pm 0.00)	54.17 (\pm 7.22)
	BBAR	17.75 (\pm 2.28)	19.50 (\pm 1.50)	18.75 (\pm 1.92)	16.25 (\pm 3.49)	18.75 (\pm 3.90)
	SynFlowNet ^a	58.00 (\pm 4.64)	59.00 (\pm 4.06)	55.50 (\pm 10.23)	47.25 (\pm 6.61)	57.00 (\pm 7.58)
	SynFlowNet ^b	61.50 (\pm 3.84)	60.50 (\pm 3.91)	57.25 (\pm 4.97)	44.50 (\pm 9.29)	62.00 (\pm 1.22)
	RGFN	48.00 (\pm 1.22)	43.00 (\pm 2.74)	49.00 (\pm 1.22)	42.00 (\pm 3.00)	44.50 (\pm 4.03)
	RxnFlow	66.00 (\pm 1.58)	64.00 (\pm 5.05)	66.50 (\pm 2.06)	63.00 (\pm 4.64)	70.50 (\pm 2.87)
3D Reaction	3DSynthFlow	51.50 (\pm 12.08)	50.50 (\pm 7.79)	38.67 (\pm 9.10)	48.00 (\pm 18.24)	34.00 (\pm 7.87)
Category	Method	OPRK1	PKM2	PPARG	TP53	VDR
Fragment	FragGFN	2.50 (\pm 2.29)	8.75 (\pm 3.11)	0.75 (\pm 0.43)	4.25 (\pm 1.64)	3.50 (\pm 2.18)
	FragGFN+SA	3.25 (\pm 1.79)	9.75 (\pm 2.28)	1.25 (\pm 1.09)	2.25 (\pm 1.92)	3.75 (\pm 2.77)
Reaction	SynNet	54.17 (\pm 7.22)	50.00 (\pm 0.00)	54.17 (\pm 7.22)	29.17 (\pm 18.16)	45.83 (\pm 7.22)
	BBAR	13.75 (\pm 3.11)	20.00 (\pm 0.71)	15.50 (\pm 2.29)	18.50 (\pm 3.28)	12.25 (\pm 3.34)
	SynFlowNet ^a	56.50 (\pm 7.63)	50.75 (\pm 1.09)	53.50 (\pm 5.68)	55.50 (\pm 9.94)	53.50 (\pm 1.80)
	SynFlowNet ^b	56.25 (\pm 2.49)	58.00 (\pm 7.00)	57.00 (\pm 5.74)	66.50 (\pm 6.80)	53.50 (\pm 3.84)
	RGFN	48.00 (\pm 2.55)	48.50 (\pm 3.20)	47.00 (\pm 5.83)	53.25 (\pm 3.63)	46.50 (\pm 2.69)
	RxnFlow	72.25 (\pm 2.05)	62.00 (\pm 3.24)	65.50 (\pm 4.03)	67.50 (\pm 2.96)	66.75 (\pm 2.28)
3D Reaction	3DSynthFlow	43.67 (\pm 2.36)	54.67 (\pm 13.42)	41.00 (\pm 2.94)	49.33 (\pm 13.42)	49.00 (\pm 19.87)

Table 19: Average synthesis steps estimated from AiZynthFinder for top-100 diverse modes generated against 15 LIT-PCBA targets. The best results are in bold.

Category	Method	Average Number of Synthesis Steps (\downarrow)				
		ADRB2	ALDH1	ESR_ago	ESR_antago	FEN1
Fragment	FragGFN	3.60 (\pm 0.10)	3.83 (\pm 0.08)	3.76 (\pm 0.20)	3.76 (\pm 0.16)	3.34 (\pm 0.18)
	FragGFN+SA	3.73 (\pm 0.09)	3.66 (\pm 0.04)	3.66 (\pm 0.07)	3.67 (\pm 0.21)	3.79 (\pm 0.19)
Reaction	SynNet	3.29 (\pm 0.36)	3.50 (\pm 0.00)	3.00 (\pm 0.00)	4.13 (\pm 0.89)	3.50 (\pm 0.00)
	BBAR	3.60 (\pm 0.17)	3.62 (\pm 0.19)	3.76 (\pm 0.04)	3.72 (\pm 0.11)	3.59 (\pm 0.14)
	SynFlowNet ^a	2.64 (\pm 0.07)	2.48 (\pm 0.07)	2.60 (\pm 0.25)	2.45 (\pm 0.09)	2.56 (\pm 0.29)
	SynFlowNet ^b	2.42 (\pm 0.10)	2.48 (\pm 0.10)	2.38 (\pm 0.10)	2.34 (\pm 0.30)	2.41 (\pm 0.14)
	RGFN	2.88 (\pm 0.21)	2.65 (\pm 0.09)	2.78 (\pm 0.19)	2.91 (\pm 0.23)	2.76 (\pm 0.17)
	RxnFlow	2.42 (\pm 0.23)	2.19 (\pm 0.12)	1.95 (\pm 0.20)	2.15 (\pm 0.18)	2.23 (\pm 0.18)
3D Reaction	3DSynthFlow	2.95 (\pm 0.42)	3.17 (\pm 0.21)	3.09 (\pm 0.23)	2.72 (\pm 0.24)	3.11 (\pm 0.24)
Category	Method	GBA	IDH1	KAT2A	MAPK1	MTORC1
Fragment	FragGFN	3.94 (\pm 0.11)	3.74 (\pm 0.10)	3.78 (\pm 0.09)	3.72 (\pm 0.18)	3.84 (\pm 0.18)
	FragGFN+SA	3.94 (\pm 0.15)	3.84 (\pm 0.23)	3.66 (\pm 0.18)	3.69 (\pm 0.21)	3.94 (\pm 0.08)
Reaction	SynNet	3.38 (\pm 0.22)	3.38 (\pm 0.22)	3.46 (\pm 0.95)	3.50 (\pm 0.00)	3.29 (\pm 0.36)
	BBAR	3.71 (\pm 0.12)	3.68 (\pm 0.02)	3.63 (\pm 0.05)	3.73 (\pm 0.05)	3.77 (\pm 0.09)
	SynFlowNet ^a	2.48 (\pm 0.18)	2.61 (\pm 0.13)	2.45 (\pm 0.37)	2.81 (\pm 0.24)	2.44 (\pm 0.27)
	SynFlowNet ^b	2.45 (\pm 0.08)	2.46 (\pm 0.12)	2.45 (\pm 0.12)	2.83 (\pm 0.27)	2.39 (\pm 0.17)
	RGFN	2.77 (\pm 0.20)	2.97 (\pm 0.15)	2.78 (\pm 0.10)	2.86 (\pm 0.19)	2.92 (\pm 0.06)
	RxnFlow	2.10 (\pm 0.08)	2.16 (\pm 0.11)	2.29 (\pm 0.05)	2.29 (\pm 0.11)	2.05 (\pm 0.09)
3D Reaction	3DSynthFlow	3.05 (\pm 0.38)	3.05 (\pm 0.31)	3.43 (\pm 0.23)	3.20 (\pm 0.39)	3.45 (\pm 0.22)
Category	Method	OPRK1	PKM2	PPARG	TP53	VDR
Fragment	FragGFN	3.82 (\pm 0.13)	3.71 (\pm 0.12)	3.73 (\pm 0.24)	3.73 (\pm 0.23)	3.75 (\pm 0.06)
	FragGFN+SA	3.62 (\pm 0.12)	3.84 (\pm 0.21)	3.71 (\pm 0.04)	3.66 (\pm 0.05)	3.67 (\pm 0.25)
Reaction	SynNet	3.29 (\pm 0.36)	3.50 (\pm 0.00)	3.29 (\pm 0.36)	3.67 (\pm 0.91)	3.63 (\pm 0.22)
	BBAR	3.70 (\pm 0.17)	3.61 (\pm 0.05)	3.72 (\pm 0.13)	3.65 (\pm 0.05)	3.77 (\pm 0.16)
	SynFlowNet ^a	2.49 (\pm 0.33)	2.62 (\pm 0.10)	2.56 (\pm 0.12)	2.51 (\pm 0.27)	2.55 (\pm 0.09)
	SynFlowNet ^b	2.50 (\pm 0.11)	2.52 (\pm 0.20)	2.53 (\pm 0.06)	2.34 (\pm 0.10)	2.51 (\pm 0.17)
	RGFN	2.81 (\pm 0.12)	2.82 (\pm 0.10)	2.82 (\pm 0.18)	2.64 (\pm 0.10)	2.84 (\pm 0.18)
	RxnFlow	2.00 (\pm 0.09)	2.34 (\pm 0.19)	2.21 (\pm 0.06)	2.12 (\pm 0.12)	2.12 (\pm 0.12)
3D Reaction	3DSynthFlow	3.34 (\pm 0.07)	3.10 (\pm 0.49)	3.45 (\pm 0.12)	3.08 (\pm 0.39)	3.06 (\pm 0.62)

G.9 EXAMPLE GENERATION TRAJECTORIES

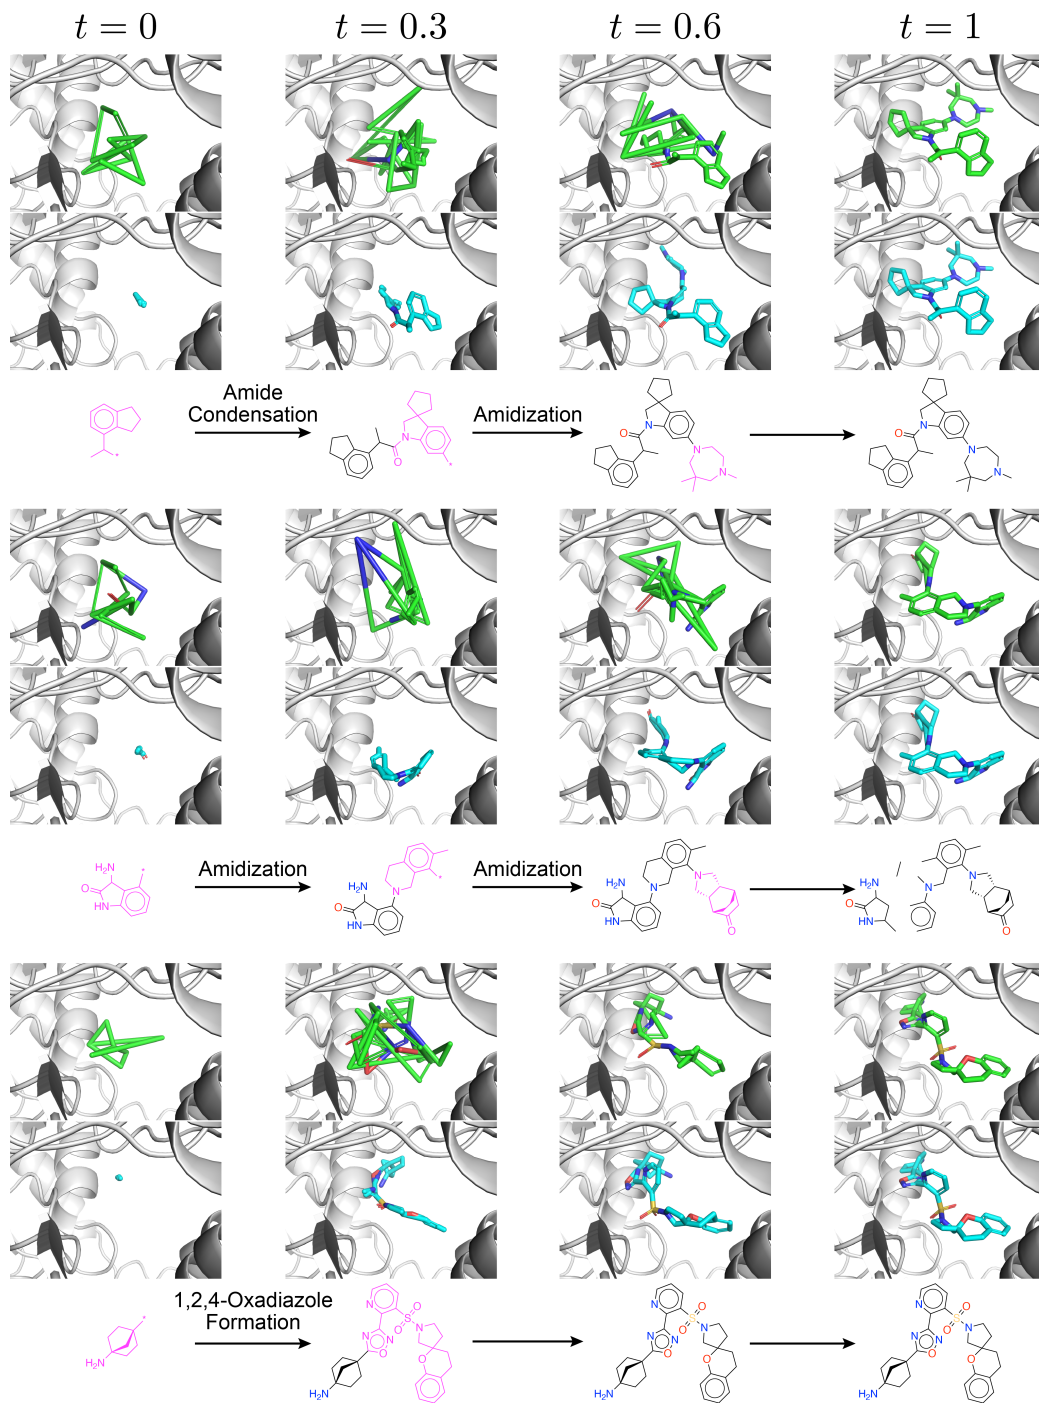


Figure 5: The example generation trajectory against ALDH1 target. The top row shows the 3D molecule x_t (green). The mid row shows the predicted final pose at each time step (cyan). The bottom row shows the synthesis pathway.