Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

LeadFlow learns activity-aligned molecular generation from chemical series for structure-based molecular optimization

Zhehuan Fan[1,2,†], Dan Teng[1†], Jie Yu[1,3,4†], Qi Huang[5,†], Keyue Qiu[6,†], Beijing Chen[5], Xia Sheng[1,2], Xinyu Jiang[1,2], Lehan Zhang[1,2], Xingyou Wang[1,2], Yaoyu Zheng[2,7], Mingan Chen[8], Runze Zhang[9], Jia Li[5,10], Hao Zhou[6,*], Mingliang Wang[2,5,11,*], Xiaomin Luo[1,2,*], Duanhua Cao[12,*], Mingyue Zheng[1,2,3,7,*]

[1]Drug Discovery and Design Center, State Key Laboratory of Drug Research, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai 201203, China

[2]University of Chinese Academy of Sciences, Beijing 100049, China

[3]School of Information Science and Technology, ShanghaiTech University, Shanghai, 201210, China

[4]Lingang Laboratory, Shanghai 200031, China

[5]Zhongshan Institute for Drug Discovery, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Zhongshan 528400, China

[6]Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China

[7]School of Pharmaceutical Science and Technology, Hangzhou Institute for Advanced Study, UCAS, Hangzhou, 330106, China

[8]iHuman institute, ShanghaiTech University, Shanghai, 201210, China

[9]Science for Life Laboratory, Department of Cell and Molecular Biology, Uppsala University, BMC, Uppsala, Sweden

[10]The State Key Laboratory of Chemical Biology, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai 201203, China

[11]Department of Medicinal Chemistry, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai 201203, China

[12]State Key Laboratory of Cardiovascular Diseases and Medical Innovation Center, Medical Artificial Intelligence Innovation Center, Shanghai East Hospital, Frontier Science Center for Stem Cell Research, School of Life Sciences and Technology, Tongji University, Shanghai, 200092, China

Corresponding Authors

*(Mingyue Zheng) E-mail: myzheng@simm.ac.cn

*(Duanhua Cao) E-mail: caodh@tongji.edu.cn

*(Xiaomin Luo) E-mail: xmluo@simm.ac.cn

*(Mingliang Wang) E-mail: wangmingliang@simm.ac.cn

*(Hao Zhou) E-mail: zhouhao@air.tsinghua.edu.cn

†These authors contributed equally to this work.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Abstract

Structure-based molecular optimization requires supervision that captures scaffold modifications associated with experimental activity. Here we introduce LeadFlow, combining maximum common substructure supervision from assay-consistent chemical series, Bayesian Flow Networks for joint generation of three-dimensional coordinates and molecular graphs, and experimental activity-based direct preference optimization (DPO). On MolGenBench, LeadFlow recovered reference active scaffolds in approximately 400 of 600 series and exact active molecules in approximately 120 series. Controlled scaffold-definition ablations supported broader active-chemotype recovery and fused-ring elaboration. Activity-based DPO increased activity-weighted recovery relative to pretraining and docking-score-based preferences. In a prospective ClpP campaign, generation followed by computational filtering and medicinal chemistry selection yielded five synthesized compounds with lower mean EC_{50} values than the starting compound A1. DC-ClpP-03 reached 38.7 nM, an approximately 156-fold potency improvement. These results support chemical-series supervision and experimental preferences for scaffold-conditioned molecular optimization.

1 Introduction

Generative models offer new approaches to molecular design for drug discovery[1-5]. In recent years, deep generative models have emerged as promising tools for *de novo* ligand design. However, empirical evidence suggests that these methods still struggle to rediscover active ligands for specific targets[6,7]. Furthermore, *de novo* generation approaches frequently produce molecules that lack feasible synthetic routes[8,9], making it difficult to validate whether the newly generated molecules are truly active. These problems limit the practical value of current *de novo* generation methods in drug discovery. In contrast, lead optimization based on lead scaffolds is a more practical scenario for current generative models, as it enables exploration of viable chemical space while preserving validated core structures and can reduce the risk of synthetically infeasible designs[10-12].

However, this strategy remains significantly under-explored. Ligand-based optimization methods often extract pharmacophore features or molecular surface properties from reference compounds to guide generation[13-15]. While these approaches can identify novel active molecules, they often fail to preserve the core scaffold during generation, which deviates substantially from the practical requirements of hit-to-lead optimization. Structure-based molecular optimization (SBMO) methods, on the other hand, largely reduce the optimization task to a geometric completion problem of elaborating a given scaffold within a protein

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

pocket[16-19]. Although such approaches can retain the given scaffold to some extent, their training supervision may inadequately represent activity-relevant scaffold modifications, resulting in limited practical optimization performance[10]. This limitation stems from bottlenecks at both the data and methodological levels. At the data level, existing methods predominantly rely on datasets such as CrossDocked[19,20], which were used for *de novo* generation tasks. Such datasets are organized as protein-ligand pairs rather than assay-consistent chemical series and do not explicitly encode relative activities among scaffold-sharing analogs. Methods trained on such data, including Delete[16] and DiffDec[17], are compelled to rely on simple hand-crafted rules to define scaffolds and modification regions[10,21-23]. This limits their ability to learn real-world hit-to-lead optimization patterns from experimentally validated scaffold modification examples, including complex transformations such as ring fusion and ring expansion. At the methodological level, DiffSBDD adapts a *de novo* generative model to scaffold optimization via the inpainting strategy[18], but this direct transfer was not designed with the specific requirements of lead optimization in mind and has limitations in the evaluated optimization setting. DiffDec[17], while specifically designed for scaffold decoration, requires explicit specification of modification sites at inference time, which restricts the flexibility of structural exploration. Recent works such as FLOWR[24] and PocketXMol[25] have begun to unify different molecular generation tasks and modalities within a single generative model, including *de novo* generation and fragment growing, allowing different tasks and molecular modalities to benefit from shared training[26]. However, their training objectives primarily focus on conditional molecular generation or reconstruction, without explicitly learning which scaffold modifications are more likely to improve activity. Diffleop[27] takes a further step toward activity-directed optimization by introducing gradient-based guidance from a learned affinity predictor during generation, with the guidance signal derived from predicted binding affinity. More broadly, although molecular generative models have explored various strategies for guided generation, including direct preference optimization[28] (DPO) and gradient-based guidance[29], how to define optimization signals that reflect real-world hit-to-lead activity improvement remains insufficiently explored.

To bridge this gap, we introduce LeadFlow, a framework that treats hit-to-lead optimization as an optimization-aligned conditional generation problem. The central idea is that the training signal, generative prior and preference objective should reflect experimentally observed structure-activity relationships[30] while preserving a functional core and exploring activity-relevant modifications within a target pocket[31-33]. LeadFlow therefore combines three components. First, we curate a Homologous Series Dataset in which compounds are grouped by target and assay, and the Maximum Common Substructure (MCS) within each series defines the conserved scaffold[34,35]. This construction exposes the model to experimentally observed modification patterns rather than artificial fragments from isolated protein-ligand pairs. Second, Bayesian

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Flow Networks (BFN)[36] provide a unified parameter-space formulation for jointly modeling continuous 3D coordinates and discrete atom and bond types in scaffold modifications. Third, DPO[37,38] uses within-assay experimental activity differences to bias generation toward higher-potency modifications without relying on docking scores as the preference signal.

We evaluate LeadFlow through benchmark comparisons, controlled scaffold-definition ablations, and a comparison of experimental activity-based and docking-score-based DPO. These analyses test whether supervision from real-world optimization patterns supports active-chemotype coverage and whether experimental preferences enrich recovery toward higher-activity reference molecules. An Interleukin-1 receptor-associated kinase 4 (IRAK4) retrospective case illustrates generation patterns consistent with published structure-activity relationships (SAR). A prospective Caseinolytic protease P (ClpP) agonist optimization campaign then evaluates the combined generation-and-prioritization workflow through synthesis and activity measurements: all five selected compounds showed lower mean EC_{50} values than the starting compound, and the most potent achieved an approximately 156-fold improvement.

2 Results

2.1 Homologous series data encode real-world optimization patterns

The foundation of our approach lies in extracting real-world hit-to-lead optimization patterns from chemical series to define scaffold-modification supervision, addressing a critical gap in the training data used by current SBMO methods. Many SBMO methods are trained on CrossDocked[16,17,19,20]. Since this dataset is composed of isolated protein-ligand pairs rather than homologous series, its standard training-task construction does not explicitly encode the relationships among analogs that share a conserved core. This limitation forces existing approaches to resort to rigid, rule-based fragmentation strategies to define scaffolds and modification regions. For example, Bemis-Murcko decomposition obtains the Bemis-Murcko scaffold (BMScaffold) by mechanically defining all terminal acyclic substituents as removable side chains, which limits its ability to represent scaffold-morphing situations such as ring fusion. Similarly, reaction-based slicing (RxnSlice) methods, such as RECAP[21], BRICS[23], and the strategy proposed by Arús-Pous et al.[22], rely on predefined synthetic rules to decompose molecules. These methods typically employ sets of handcrafted reaction rules to ensure that each split corresponds to a known, implementable chemical reaction. While this explicitly accounts for synthetic accessibility, these approaches are predominantly limited to cleaving acyclic bonds. Consequently, they cannot represent complex scaffold modifications such as ring fusion or expansion, which involve forming new cyclic systems rather than merely attaching fragments to the retained core. These fragmentation strategies therefore struggle to capture the successful

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

optimization patterns observed within homologous series around an empirically conserved core scaffold.

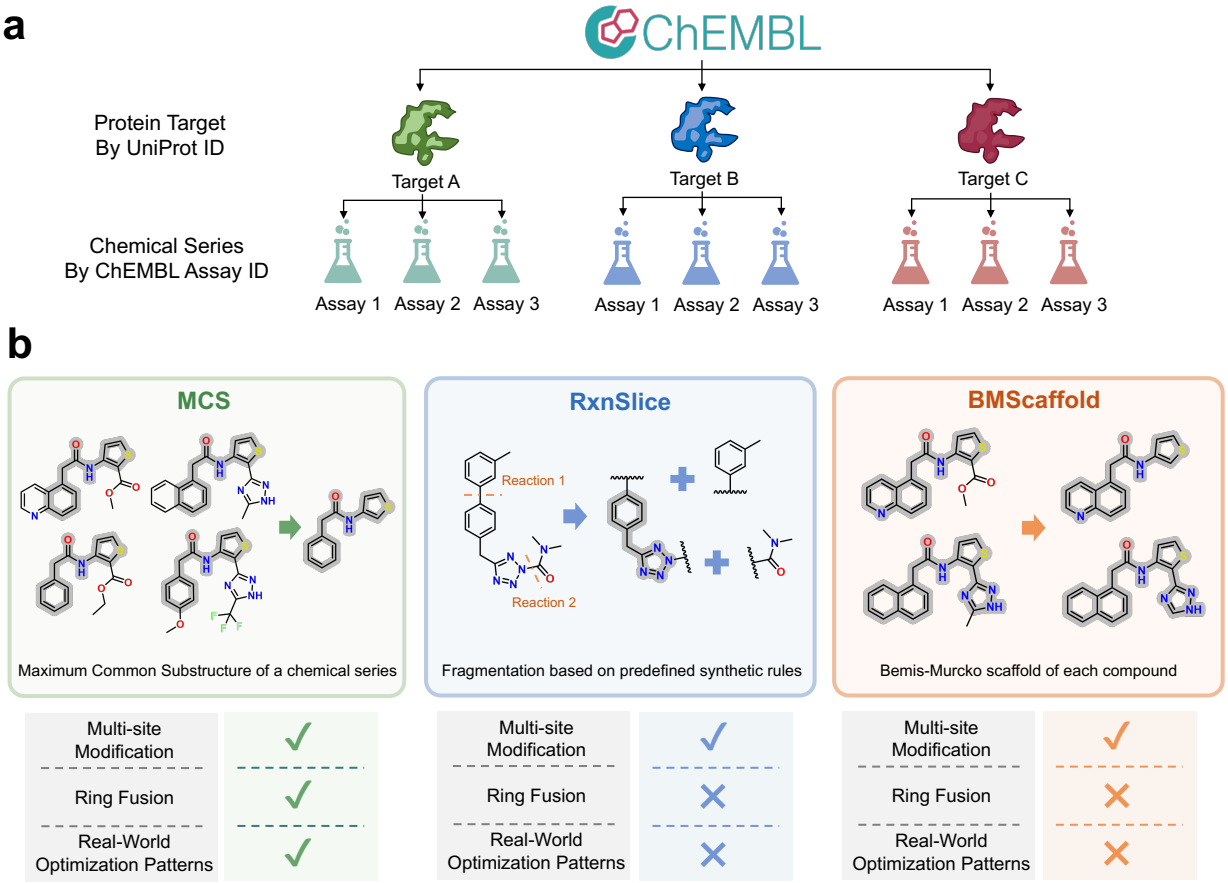


Fig. 1 Homologous-series data construction and MCS-based scaffold definition. a. Homologous series dataset construction from ChEMBL. Protein targets are identified by UniProt ID and compounds are grouped into chemical series by ChEMBL Assay ID, yielding homologous series for each target. **b.** Comparison of scaffold-modification definitions. MCS identifies a conserved scaffold shared across compounds within a series, whereas RxnSlice uses predefined reaction-based cleavage rules and BMScaffold assigns a Bemis-Murcko scaffold independently to each compound.

To overcome these limitations, we curated a dataset of homologous series from ChEMBL. As depicted in **Fig. 1a**, compounds were first grouped by target using UniProt IDs and then partitioned into chemical series according to their ChEMBL Assay IDs. Within each series, the compounds can share a common core structure and span structural modifications associated with measured activity differences against a protein target. Within these series, the MCS provides a series-derived definition of the conserved scaffold. **Fig. 1b** illustrates the extraction of the MCS from a chemical series and compares its scaffold–modification

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

definition with those of the other two strategies. Unlike artificial fragmentation rules, the MCS identifies the substructure shared across the measured analogs, providing an empirical scaffold context without requiring inference of the chemists' design decisions. This series-derived scaffold supervision exposes the model to experimentally characterized structural variations, including fused-ring elaborations that are difficult to represent with acyclic cleavage rules. Details of the MCS processing procedure are described in **Methods Section 4.1.1.**

2.2 LeadFlow formulates structure-based molecule optimization as Bayesian flow generation

To translate these data-driven insights into a design tool, we formulate the optimization task as a conditional generation problem. Given the 3D structure of the protein pocket and the conserved scaffold, the model aims to generate structural modifications by sampling atomic coordinates, atom types and bond types to construct a complete molecule that is chemically and geometrically compatible with the conserved scaffold and the protein pocket environment. To achieve this, we developed a generative model named LeadFlow based on BFN as illustrated in **Fig. 2**. A fundamental challenge in 3D molecular generation lies in the high sensitivity of molecular structure to noise, complicated by the heterogeneous nature of molecular data which comprise both continuous coordinates and discrete variables such as atom types and bond types[39,40]. Diffusion models perturb molecular variables in data space; jointly representing continuous coordinates and discrete chemical topology remains a challenge[39,41]. BFN provides an alternative generative framework[36]. By shifting the generative process from the data space to a continuous and unified parameter space, BFN iteratively refines the distribution parameters of data via a Bayesian update strategy. This formulation supports joint modeling of continuous and discrete variables and has been applied across scientific generation tasks[42-45].

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

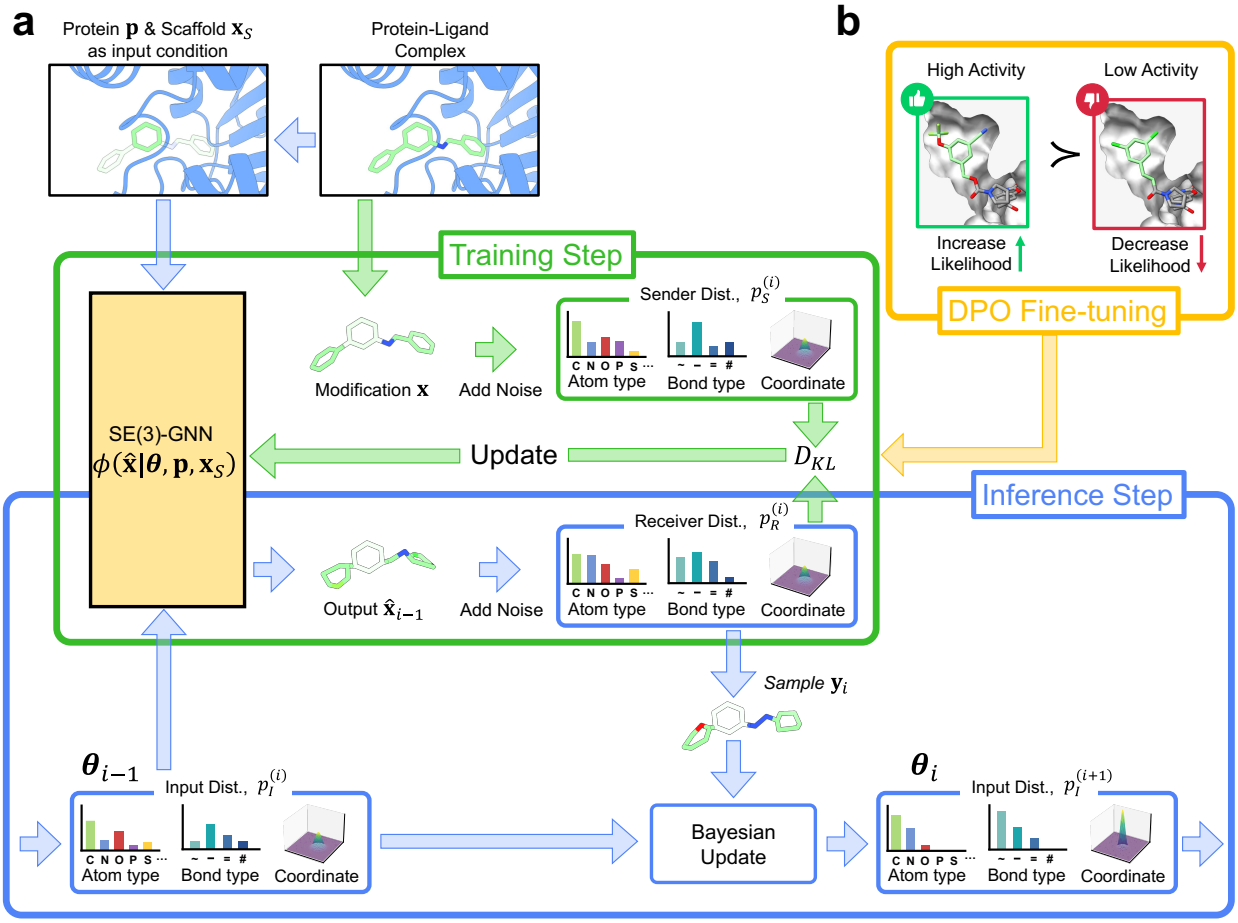


Fig. 2 An overview of the LeadFlow framework. a. The generative process of the LeadFlow for structure-based molecule optimization. It iteratively refines a set of parameters $\boldsymbol{\theta}$, using the SE(3)-equivariant graph transformer ϕ to predict the clean modification $\hat{\mathbf{x}}_{i-1}$ given current parameters $\boldsymbol{\theta}_{i-1}$ with protein $\mathbf{p}$ and scaffold conditions $\mathbf{x}_S$. The model is trained by minimizing the KL divergence between the sender distribution p_S over the noised observation $\mathbf{y}$ of ground truth modification $\mathbf{x}$, and the receiver distribution p_R over the noised estimation $\mathbf{y}$ of predicted modification $\hat{\mathbf{x}}$. Through Bayesian update of the sampled $\mathbf{y}$, the parameters are iteratively refined from $\boldsymbol{\theta}_{i-1}$ to $\boldsymbol{\theta}_i$, progressively enriching the information encoded in $\boldsymbol{\theta}$ and enabling the model to predict increasingly accurate modifications $\hat{\mathbf{x}}$. LeadFlow uses the noise-reduced sampling strategy in which $\hat{\mathbf{x}}_{i-1}$ is used in place of $\mathbf{y}_i$ to directly perform the Bayesian update. **b.** DPO fine-tuning. Within each series, high-activity and low-activity modifications are paired for comparison, guiding the fine-tuning process to increase the likelihood of generating high-activity modifications.

The generative process of LeadFlow is modeled as an iterative information transmission process between a **Sender Distribution** and a **Receiver Distribution**. The parameter $\boldsymbol{\theta}$ encodes the distribution parameters over the molecular modification variables, encompassing atomic coordinates, atom types, and bond types. During training, the **Sender Distribution** $p_S(\mathbf{y}_i \mid \mathbf{x}; \alpha_i)$ transforms the ground-truth molecular modification $\mathbf{x}$ into the parameter space with the noise scheduler α_i and sends the noisy structural modification $\mathbf{y}_i$ to the **Receiver Distribution** for Bayesian updates. Correspondingly, the **Receiver Distribution** $p_R(\mathbf{y}_i \mid \boldsymbol{\theta}, \phi, \alpha_i)$ employs a neural network ϕ to approximate the **Sender Distribution** at each

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

timestep i. In LeadFlow, we utilize an SE(3)-equivariant graph transformer as ϕ to model the structural modification with ensured rotational and translational equivariance. As depicted in **Fig. 2a**, the network takes the current parameters $\boldsymbol{\theta}_{i-1}$ and the context condition comprising the protein $\mathbf{p}$ and conserved scaffold $\mathbf{x}_S$ as input to predict the clean structural modification $\hat{\mathbf{x}}_{i-1}$, which is subsequently processed through the noise scheduler α_i to yield the **Receiver Distribution**. The final training objective is to minimize the Kullback-Leibler (KL) divergence between p_s and p_R across all timesteps. During the inference stage, the parameter update is driven by the model's own predictions. In the standard BFN formulation, an estimated noised observation $\mathbf{y}_i$ derived from $\hat{\mathbf{x}}_{i-1}$ is sampled from **Receiver Distribution** p_R to update the current parameters $\boldsymbol{\theta}_{i-1}$ into posterior $\boldsymbol{\theta}_i$ via the Bayesian updates. This iterative process gradually incorporates more information into $\boldsymbol{\theta}$ and increases certainty, thereby enabling the network ϕ to predict the structural modification $\hat{\mathbf{x}}$ with increasing precision. LeadFlow uses the noise-reduced sampling implementation described in **Methods Section 4.2.1**.

While generative models learn to generate structural modifications that are compatible with both the conserved scaffold and the protein pocket, they lack the capability to account for the impact of specific modifications on biological potency. To inject explicit knowledge regarding which structural modifications correspond to higher bioactivity, we employ DPO in the post-training stage to fine-tune the pre-trained LeadFlow. Existing structure-based generation methods often resort to docking scores (e.g., Vina) as proxies to guide optimization[28,29,46]. However, these scoring functions are approximations that often fail to correlate with true experimental activity[35,47,48]. We use experimental bioactivity measurements within homologous series to construct preference labels. Since our data curation strictly groups compounds by ChEMBL Assay ID, within-series comparisons reduce differences associated with assay conditions. As depicted in **Fig. 2b**, we fine-tune LeadFlow using DPO with these relative activity preferences as training signal. This aligns the model's generative probability specifically with high-activity modifications. Implementation details of the DPO fine-tuning procedure are provided in the **Methods Section 4.2.2**.

2.3 Molecular quality, pose validity and interaction recovery

We compared against seven baseline methods (detailed in **Methods Section 4.4**) on MolGenBench[7] hit-to-lead benchmark: two SBMO approaches based on different artificial fragmentation rules and generative paradigms (Delete[16] and DiffDec[17]), two DiffSBDD[18] inpainting models trained on different datasets (DiffSBDD-M and DiffSBDD-C), three ligand-based approaches guided by features from reference active molecules (ShEPhERD[13], PGMG[14], and ShapeMol[49]). Note that PGMG generates molecules as SMILES strings rather than 3D structures, and ShapeMol, while generating 3D molecules, does not place them directly within the binding pocket. These two methods are therefore excluded from

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

conformation-dependent evaluations where applicable.

Avoiding unstable or reactive substructures is a useful initial criterion for prioritizing generated molecules. Following MolGenBench[7], we employ ChemFilter to evaluate structural-alert filter compliance of generated molecules by filtering out those containing unstable or reactive substructures. As shown in **Fig. 3a**, LeadFlow achieves the highest ChemFilter pass rate among all evaluated methods, closely approaching the active reference level. In terms of BMScaffold diversity, LeadFlow outperforms most baselines and exceeds the reference level, only slightly behind PGMG. This is expected, as LeadFlow grows molecules upon existing scaffolds, which inherently constrains structural diversity, whereas PGMG generates molecules from scratch guided solely by ligand features without such constraints. This suggests that LeadFlow generates molecules passing structural-alert filters at a rate comparable to the reference actives but also explores diverse chemical space around the conditioned scaffold, which is useful for SAR exploration. Additional evaluations of basic molecular properties, such as validity, QED, synthesizability, and molecular motif distributions, are provided in **Supplementary Fig. 1**, where LeadFlow shows comparable performance to existing methods.

Physically plausible poses are a prerequisite for assessing generated designs in a protein pocket. For SBMO, it is necessary to examine whether the model generates physically realistic binding poses. This provides a direct test of whether the model has learned the geometric and energetic criteria of protein-ligand recognition. Recent studies have emphasized that pose evaluation cannot be omitted in structure-based generation, because generated complexes may contain substantial physical violations and pose artifacts[50]. We therefore evaluated the physical realism of LeadFlow generated poses. We first assessed overall pose quality using PoseBusters[51]. As shown in **Fig. 3b**, LeadFlow achieved the highest proportion of generated molecules passing all PoseBusters checks among all compared methods, with a PB-Valid rate of 77.57%. More detailed results for representative PoseBusters criteria are illustrated in **Fig. 3c**. LeadFlow achieved the best pass rates across intra- and intermolecular measurements such as the planarity of aromatic rings, standard bond lengths, and protein-ligand clashes. Beyond the Internal Energy check included in PoseBusters, we further quantified the conformational strain of the generated poses using PoseCheck[50]. Specifically, we calculated the strain energies of both the directly generated poses from models and the corresponding poses after redocking into the protein pocket with Vina[52]. As shown in **Fig. 3d**, LeadFlow exhibited lower strain energy than the competing methods across different quantile levels in both settings. Notably, the strain energies of LeadFlow-generated poses were also closer to those of the corresponding Vina-redocked poses, indicating that the conformations produced by LeadFlow were already closer to energetically relaxed binding poses before docking refinement.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

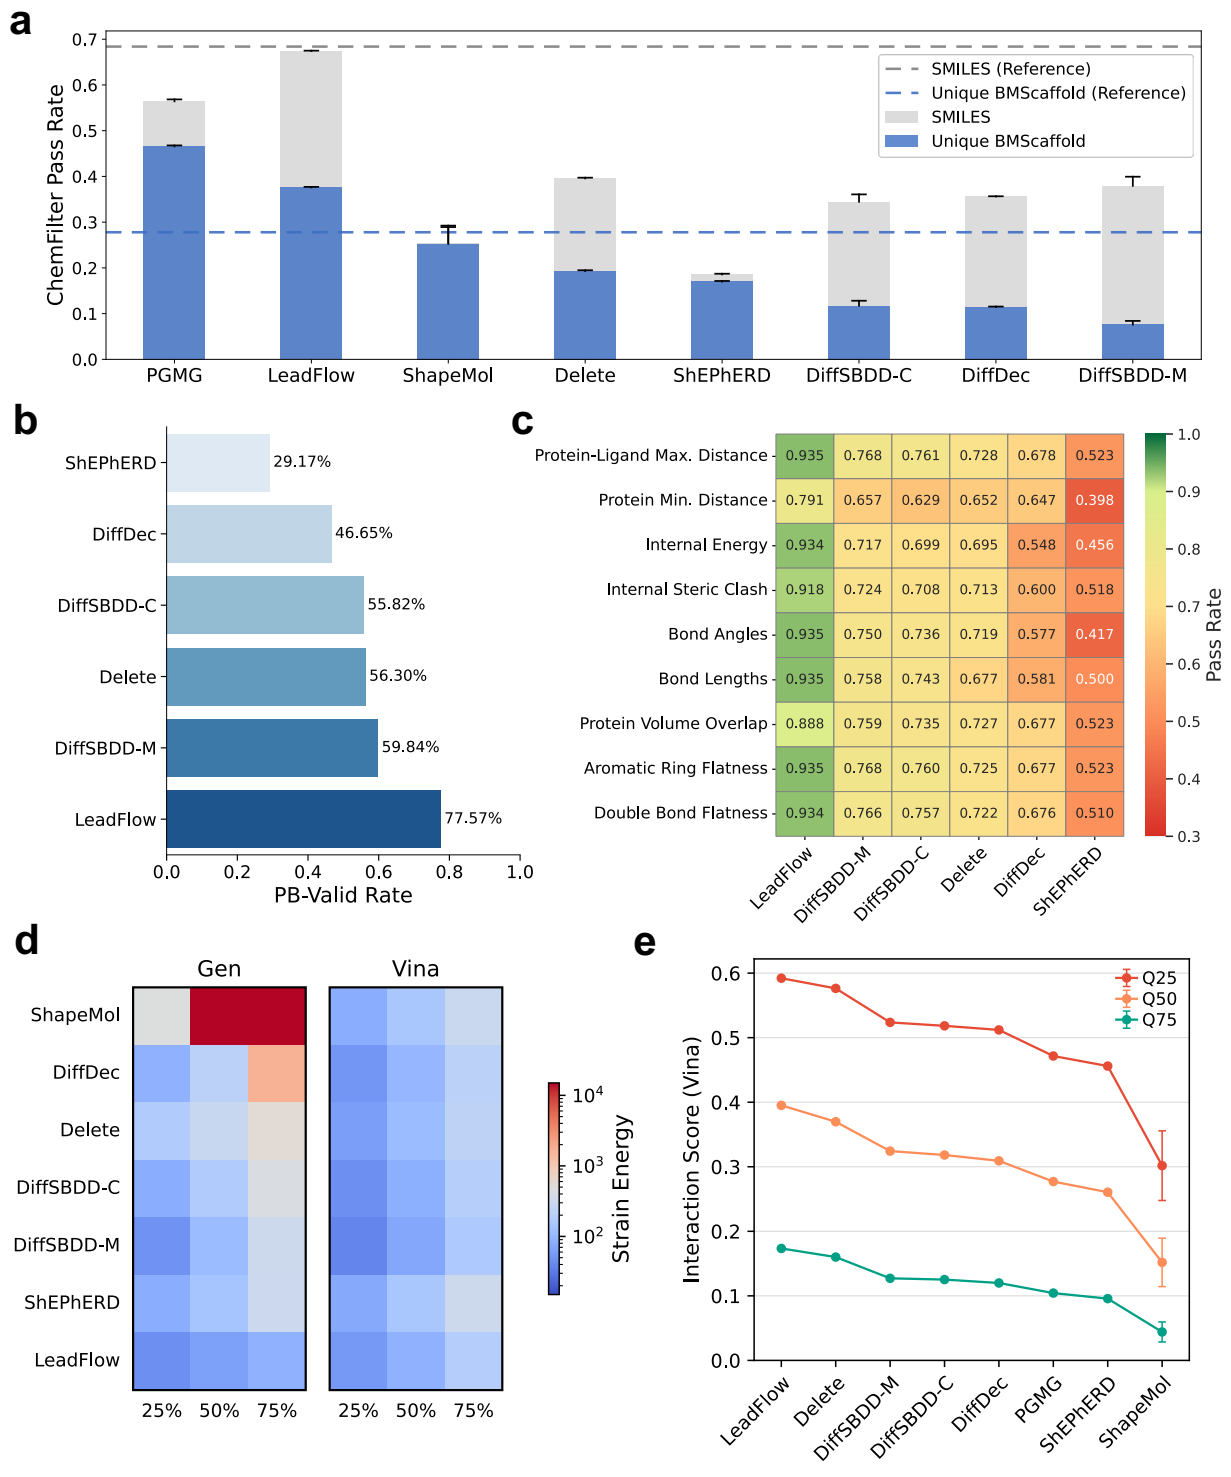


Fig. 3 | Evaluation of molecular quality, pose validity, and interaction recovery on the MolGenBench hit-to-lead benchmark. a. Structural-alert filter pass rates and Bemis-Murcko scaffold diversity of filter-passing molecules. Data are presented as mean ± s.d. (n = 3 sampling runs with different seeds from fixed trained checkpoints). **b.** Overall PoseBusters-Valid rate of generated poses across different models. **c.** Pass rates for representative PoseBusters criteria.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

d. Strain-energy distributions, summarized at different quantile levels, for directly generated poses (Gen) and Vina-docked poses (Vina). **e.** Proportion of Vina-redocked generated molecules whose interaction scores exceed the 25th, 50th, and 75th percentiles of the reference active distribution. Data are presented as mean ± s.d. (n = 3 sampling runs with different seeds from fixed trained checkpoints).

Geometric validity assesses the physical plausibility of a generated pose. However, this alone does not establish whether the model has captured biologically meaningful binding patterns. To address this, we further analyzed the interaction patterns of the generated molecules relative to those of the reference active molecules. Specifically, following the MolGenBench protocol, we redocked the generated molecules into the protein pocket with Vina and used ProLIF[53] to identify protein-ligand interaction types, and computed an interaction score for each molecule based on the interaction patterns of the reference active molecules. We then quantified the fraction of generated molecules whose scores exceeded the 25th, 50th, and 75th percentiles of the reference active distribution, denoted Q25, Q50, and Q75 in **Fig. 3e**. As shown in **Fig. 3e**, LeadFlow consistently achieved the highest proportions across all three thresholds. The corresponding results based on the directly generated poses are provided in **Supplementary Fig. 2**, where LeadFlow likewise remained the top-performing method. These results indicate that LeadFlow more frequently reproduces the conserved interaction patterns observed in the reference active molecules, supporting reference-like interaction recovery, without establishing the binding mode or activity of novel generated molecules.

The joint generation of topology and geometry in the Bayesian Flow Network formulation may contribute to this performance. Bayesian Flow Networks generate samples by iteratively updating distribution parameters, rather than directly perturbing molecular structures in data space. This parameter-space formulation is naturally compatible with joint modeling of discrete chemical topology and continuous 3D coordinates, and related BFN-based models have shown strong performance in both 3D molecular generation and broader scientific generation tasks. In the present setting, such a generation process may help reduce the accumulation of local structural errors during sampling, which is consistent with the observed geometric validity, conformational strain and reference-interaction recovery.

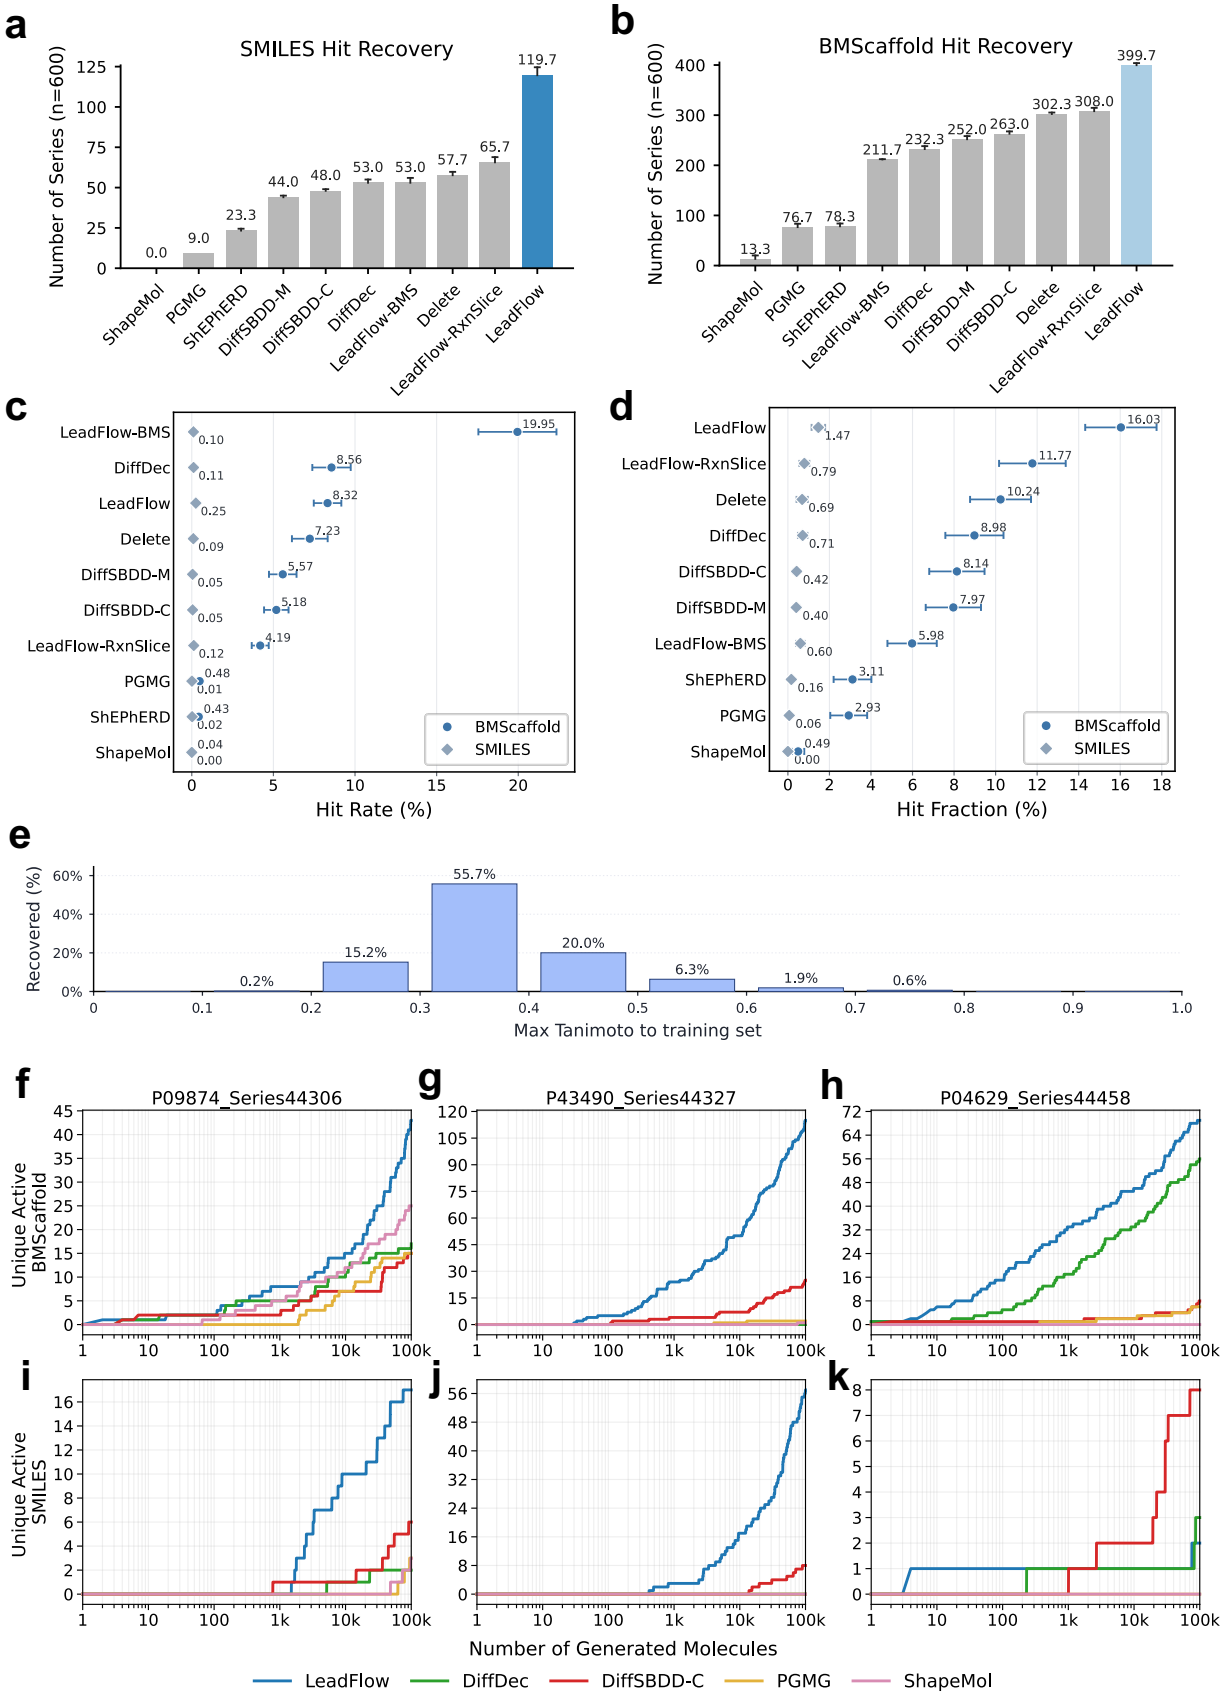
2.4 MCS-based supervision supports broad active-chemotype recovery

The primary objective of LeadFlow is to design optimized ligands from a given protein environment and core scaffold under realistic hit-to-lead settings. Building on the data construction strategy described before, we next examined whether supervision derived from homologous series and MCS-based decomposition translates into improved recovery of experimentally validated active chemotypes. To assess this, we evaluated LeadFlow on the hit-to-lead test set of MolGenBench, in which each test case provides

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

a target protein together with a starting scaffold derived from an active series. Models were asked to generate candidate ligands conditioned on the same protein-scaffold context and were evaluated by their ability to recover known active molecules or their corresponding BMScaffold. In addition to comparing LeadFlow with the existing baselines included in MolGenBench, we further conducted an ablation study by replacing the scaffold definition used in LeadFlow to isolate the contribution of our MCS-based scaffold-modification decomposition strategy. Specifically, we constructed two variants: LeadFlow-BMS, which uses BMScaffold to define the conserved core, and LeadFlow-RxnSlice, which adopts reaction-based slicing following the same strategy as DiffDec. All three LeadFlow variants shared the same training data split, model architecture and training protocol, differing only in how the core scaffold and modification regions were defined, with full-ligand masking applied when no valid core was identified to keep all raw data used across variants.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0



Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Fig. 4 | Evaluation of active molecule and BMScaffold recovery on the MolGenBench hit-to-lead benchmark. **a, b.** Number of benchmark series out of 600 in which each model recovered the exact active SMILES (**a**) and the active Bemis-Murcko scaffold (**b**). Data are presented as mean values ± s.d. (n = 3 sampling runs with different seeds from fixed trained checkpoints). **c, d.** Recovery performance measured by Hit Rate (**c**) and Hit Fraction (**d**) for BMScaffold and SMILES recovery (n = 600 series). Points show means across series after averaging three sampling runs; error bars indicate series-level 95% CIs calculated using Student's t-distribution. **e.** Distribution of the ECFP4 Tanimoto similarity between active molecules recovered by LeadFlow and their most similar molecules in the training set. **f-k.** Extended sampling analysis of active-space exploration on representative MolGenBench hit-to-lead series. (**f-h**)**,** Cumulative numbers of recovered unique active BMScaffold as a function of the number of generated molecules for three representative series. (**i-k**)**,** Corresponding cumulative numbers of recovered unique active SMILES.

We first quantified recovery performance at two levels: BMScaffold-level Hit Recovery and SMILES-level Hit Recovery (**Fig. 4a**, **b**). A test case was considered successful if the model recovered at least one known active SMILES or at least one corresponding BMScaffold. Because each of the 600 benchmark series represents a different protein-scaffold context, these metrics primarily assess the breadth of contexts in which a model can recover an experimentally validated active chemotype. Across the 600 series in benchmark, LeadFlow achieved the highest recovery under both criteria among the baseline methods and ablated variants. Specifically, LeadFlow recovered BMScaffold in approximately 400 series and SMILES in approximately 120 series, both markedly higher than those of the competing methods. Notably, its SMILES recovery was approximately 2-fold higher than the best-performing baseline Delete, supporting broader reference-active recovery across the evaluated protein-scaffold contexts.

The controlled ablation further linked this recovery advantage to the training supervision. Replacing the MCS-based decomposition with either BMScaffold-based decomposition or reaction-based slicing reduced recovery at both levels, with a particularly pronounced decrease in SMILES-level recovery. Because the three LeadFlow variants differed only in how the conserved core and modification regions were defined, these results indicate that MCS-based decomposition preserves more informative optimization relationships from real homologous series. This supervision enables LeadFlow to learn transferable scaffold-conditioned modification patterns, consistent with the broader reference-active recovery observed across the evaluated hit-to-lead contexts.

Although series-level recovery establishes whether a model can identify at least one correct hit within a series, it does not capture how many active molecules are recovered per test case. We therefore further evaluated Hit Rate and Hit Fraction (see **Evaluation Metrics Section 4.3**), which provide a more fine-grained view of recovery performance than binary success alone. Hit Rate measures the proportion of sampled molecules that correspond to unique actives and thus emphasizes discovery efficiency. Hit Fraction

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

measures the proportion of the known active space recovered for a given target and therefore emphasizes how broadly a model explores the reference active chemical space. As shown in **Fig. 4c, d,** LeadFlow remained among the strongest methods under both metrics. At the BMScaffold level, LeadFlow achieved a Hit Rate of 8.30%, comparable to DiffDec at 8.56% and higher than Delete at 7.23% and the remaining baselines. At the exact-SMILES level, LeadFlow achieved the highest Hit Rate of 0.25%, exceeding DiffDec at 0.11% and Delete at 0.09%. Under the Hit Fraction metric, LeadFlow also achieved the best performance at both the BMScaffold and exact-SMILES levels, reaching 16.09% and 1.44%, respectively.

The ablation results further showed that scaffold decomposition had a substantial effect on recovery performance. LeadFlow-BMS achieved the highest BMScaffold Hit Rate of 19.95%, but its BMScaffold Hit Fraction was only 5.98%, substantially lower than that of the MCS-based LeadFlow. It also showed limited exact-SMILES recovery, with a Hit Rate of 0.10% and a Hit Fraction of 0.60%. By contrast, LeadFlow-RxnSlice achieved higher Hit Fractions than most existing baselines. However, it remained below the MCS-based LeadFlow in both Hit Rate and Hit Fraction at the BMScaffold and exact-SMILES levels.

Notably, a similar contrast between BMScaffold Hit Rate and Hit Fraction was observed for both LeadFlow-BMS and DiffDec. Although LeadFlow-BMS achieved the highest BMScaffold Hit Rate and DiffDec showed a slightly higher value than LeadFlow, both methods exhibited substantially lower BMScaffold Hit Fractions. This difference is consistent with the information captured by the two metrics: BMScaffold Hit Rate counts recovered molecules bearing active BMScaffold, whereas BMScaffold Hit Fraction measures coverage of unique active BMScaffold. LeadFlow-BMS and DiffDec therefore tend to generate multiple distinct molecules around a comparatively limited set of active BMScaffolds, resulting in efficient local exploitation but narrower chemotype coverage. By contrast, LeadFlow maintained a comparable BMScaffold Hit Rate while recovering a substantially broader range of distinct active BMScaffolds, together with stronger exact-SMILES recovery. Collectively, these ablation and baseline comparisons indicate that MCS-based decomposition provides more informative supervision from real homologous series, enabling LeadFlow to achieve both broader active-chemotype coverage and stronger active-molecule recovery.

Fig. 4a-d reports recovery across all 600 benchmark series. **Supplementary Fig. 3** stratifies the same analysis by the seen/unseen protein subsets defined in MolGenBench relative to CrossDocked training exposure. These labels contextualize the baseline setting; all MolGenBench test targets were excluded from LeadFlow training by UniProt ID, together with identical test molecules and molecules sharing their BMScaffolds. LeadFlow maintained strong recovery across both subsets. To further assess whether the

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

recovered actives could be attributed to memorization of closely related training molecules, we calculated the maximum ECFP4 Tanimoto similarity between each active molecule recovered by LeadFlow and all molecules in its training set. The resulting distribution was concentrated in the low-similarity range, with 91.1% of the recovered active molecules exhibiting a maximum Tanimoto similarity below 0.5 (**Fig. 4e**). These results further support the robustness of LeadFlow across diverse protein-scaffold contexts and indicate that it captures transferable scaffold-conditioned optimization patterns.

Beyond the standard evaluation with 200 generated molecules per test case, MolGenBench also includes extended sampling tests that examine how recovery grows as the sampling budget increases. This setting provides a more stringent assessment of whether a model can continue to explore the active space of a given target series, rather than merely identifying a small number of hits under a fixed generation budget. On three series from this setting, we tracked the cumulative numbers of recovered unique active BMScaffold and unique active SMILES as the number of generated molecules increased (**Fig. 4f-k**). Across most series, LeadFlow showed substantially steeper recovery curves and reached markedly higher final counts than competing baselines. For BMScaffold recovery (**Fig. 4f-h**), LeadFlow consistently recovered more unique active scaffolds as sampling increased. A similar trend was observed for SMILES recovery in two of the three series (**Fig. 4i, j**), where LeadFlow recovered substantially more unique active SMILES than all baselines across nearly the entire sampling range. Although DiffSBDD-C outperformed LeadFlow in the third SMILES case (**Fig. 4k**), the overall trend across the three series remained in favor of LeadFlow for active scaffold and molecule recovery.

2.5 Series-derived supervision supports complex scaffold modifications

Having established the overall recovery advantage of LeadFlow and the contribution of MCS-based decomposition, we next examined how this advantage manifests across structurally diverse and increasingly challenging modification scenarios. More challenging scaffold modification cases provide a stricter test of whether the model has learned meaningful hit-to-lead optimization patterns, rather than merely recovering relatively simple substitutions. We therefore investigated LeadFlow's performance under increasingly demanding editing scenarios. We first compared the successful recovered active molecules across different models by stratifying them according to the number of newly generated atoms and the number of modified R-groups, two quantities that reflect the structural complexity of the required modification (**Fig. 5a-b**). Across a broad range of editing complexity, LeadFlow consistently recovered more actives than competing baselines. Its advantage remained evident for cases involving larger numbers of newly introduced atoms, indicating that LeadFlow can accommodate more substantial structural growth beyond small peripheral changes. A similar pattern was observed when recovery was grouped by the number of modified R-groups,

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

where LeadFlow retained the strongest performance in both single-site and multi-site modification settings, suggesting that its recovery advantage is not limited to simple edits but extends to more challenging combinatorial optimization patterns. Representative R-group modifications recovered by LeadFlow are shown in **Supplementary Fig. 4**. These results suggest that the recovery advantage of LeadFlow extends to more demanding scaffold modification scenarios rather than being limited to relatively simple edits.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

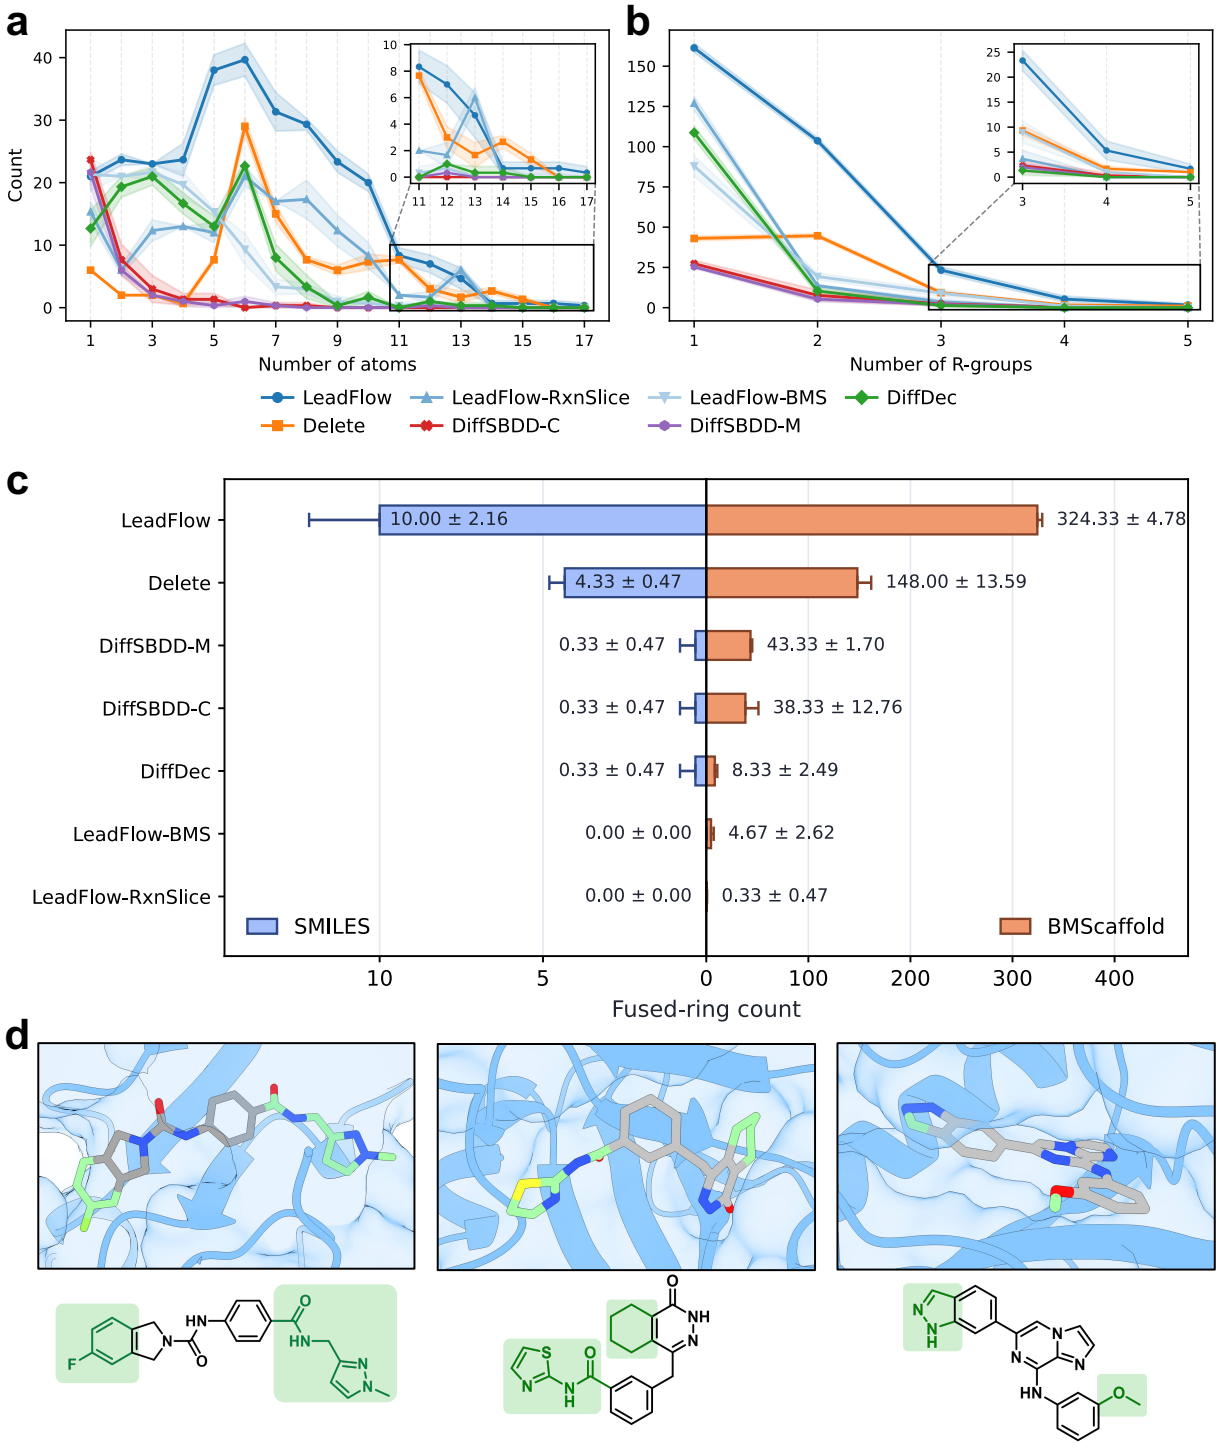


Fig. 5 | Active recovery across modification complexity and fused-ring elaboration. a, b. Distribution of recovered active molecules stratified by structural modification complexity, measured by the number of newly generated atoms (**a**) and the number of modified R-groups (**b**). Curves show the mean values across three rounds, and shaded regions indicate the standard deviations across rounds. Insets highlight the low-count region for more difficult cases. **c.** Comparison of different models' performance in recovering active molecules/BMscaffold via ring-fusion elaboration.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Fused-ring counts are deduplicated per series using recovered SMILES (SMILES recovery) or generated SMILES (BMScaffold recovery). Data are reported as mean ± s.d. (n = 3 sampling runs with different seeds from fixed trained checkpoints). **d.** Representative example of an experimentally active molecule recovered by LeadFlow through fused-ring elaboration. The blue region denotes the protein environment, and the gray region denotes the starting scaffold, both provided as model inputs, whereas the green region represents the structural modification generated by LeadFlow.

As mentioned in **Section 2.1**, the MCS-based data construction strategy of LeadFlow naturally allows the model to learn scaffold optimization patterns involving fused-ring formation and other nontrivial topological modifications. **Fig. 5d** shows representative cases in which LeadFlow successfully recovers active molecules via fused-ring elaboration. In these examples, the blue region denotes the protein environment and the gray region denotes the starting scaffold, which are provided as model inputs, whereas the green region represents the structural modification generated by LeadFlow. Additional fused-ring recovery examples are provided in **Supplementary Fig. 5**.

The recovery results for fused-ring across the full test set are summarized in **Fig. 5c**. Here, the reported values correspond to the numbers of recovered exact SMILES and recovered BMScaffolds. On this subset, LeadFlow achieved 10.00 ± 2.16 recovered exact SMILES and 324.33 ± 4.78 recovered BMScaffolds, substantially exceeding Delete (4.33 ± 0.47 and 148.00 ± 13.59), DiffDec (0.33 ± 0.47 and 8.33 ± 2.49), DiffSBDD-C (0.33 ± 0.47 and 38.33 ± 12.76), and DiffSBDD-M (0.33 ± 0.47 and 43.33 ± 1.70). These results support recovery in the evaluated fused-ring elaboration setting.

Comparison with the two ablated variants of LeadFlow further clarifies the origin of these advantages. Across the atom-count and R-group-count stratifications (**Fig. 5a-b**), LeadFlow-BMS recovered actives almost exclusively within the low-complexity regime, with its recovery counts dropping sharply as the number of newly generated atoms or modified R-groups increased. This pattern is consistent with the nature of BMScaffold-based decomposition, which predominantly excises small peripheral substituents at terminal positions of the molecule and therefore provides limited training signal for more substantial structural changes. LeadFlow-RxnSlice exhibited a similarly restricted recovery profile across both stratifications, indicating that reaction-based slicing alone, when used in place of series-derived MCS supervision, is likewise insufficient to capture more complex editing patterns. This limitation became even more pronounced on the fused-ring subset (**Fig. 5c**), where both ablated variants recovered almost no exact SMILES (0.00 ± 0.00 for both) and only a small number of BMScaffolds (4.67 ± 2.62 for LeadFlow-BMS and 0.33 ± 0.47 for LeadFlow-RxnSlice), performing markedly worse than the full LeadFlow model across all three tasks. Together, these results indicate that neither BMScaffold-based fragmentation nor reaction-based slicing alone provides adequate training signal for learning fused-ring formation or other complex topological modifications, since both strategies predominantly generate training pairs involving simple

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

peripheral substitutions rather than nontrivial changes to the scaffold core. In contrast, the full LeadFlow model maintained substantially higher recovery counts across nearly the entire range of editing complexity, from atom- and R-group-level stratifications to fused-ring elaboration, corroborating that its MCS-based data construction strategy is a key contributor to its ability to handle complex, large-scale scaffold modifications.

We attribute this strong rediscovery ability primarily to the MCS-based data construction strategy. By defining the optimization context through the MCS shared within homologous active series, LeadFlow is trained on transformations that better reflect real hit-to-lead campaigns. Compared with heuristic fragmentation, the MCS-based setup provides a more faithful and chemically meaningful supervision signal, enabling the model to learn how active series evolve under target-specific constraints.

2.6 Experimental activity preferences align generation toward potent analogs

The preceding results show that LeadFlow can generate scaffold-preserving modifications compatible with the target pocket and achieve strong active-molecule recovery. However, in practical hit-to-lead optimization, the central objective is to optimize molecules toward higher potency rather than generate active molecules in general. The pre-trained LeadFlow has learned pocket-compatible, activity-conferring modifications from the data, but it does not explicitly learn which modifications are more likely to confer higher activity in a given context. We therefore constructed paired preference data from experimental bioactivity measurements and applied DPO for post-training of LeadFlow, with the goal of shifting the generative distribution toward modifications with greater potential to enhance binding potency.

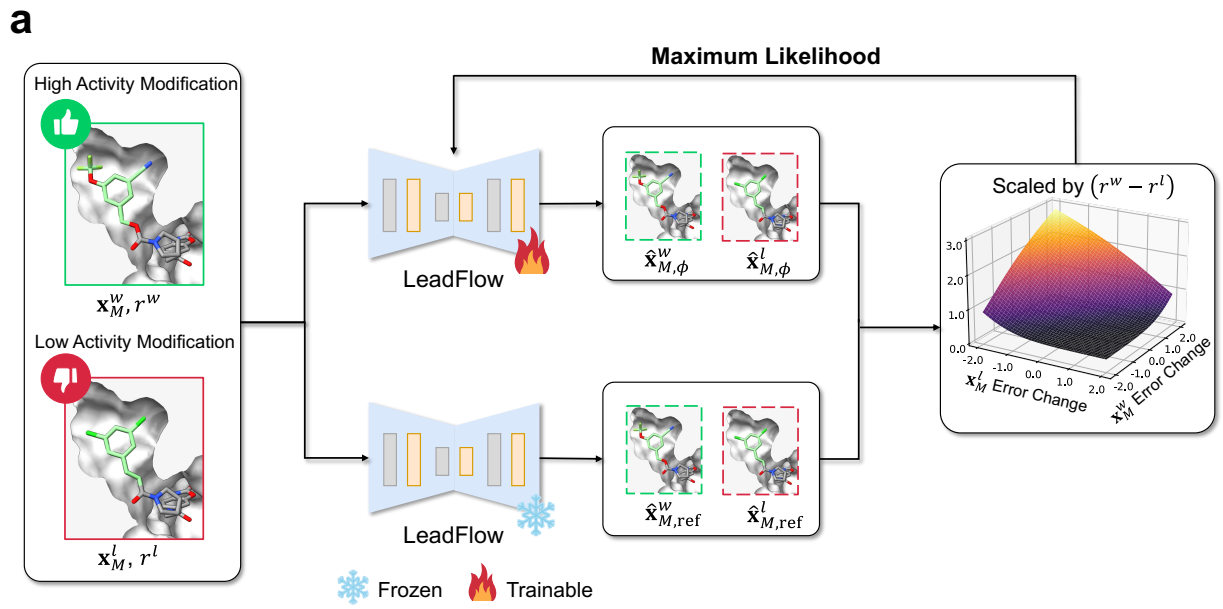
The DPO training data were constructed from the training set following the preference alignment procedure described in **Methods Section 4.1.2**. Briefly, within each eligible chemical series, molecules sharing the same core scaffold were paired according to their experimental activities. A representative preference pair is shown in **Fig. 6a**, where the higher-activity modification is treated as the preferred sample and the lower-activity modification as the dispreferred sample. To assess whether the improvement arises from experimentally grounded preference labels rather than from DPO fine-tuning alone, we further constructed a docking score-based DPO ablation, denoted as LeadFlow-DPO-DS. This model used the same molecular pairs as LeadFlow-DPO, but the preferred and dispreferred samples within each pair were reassigned according to Glide docking scores, with the molecule receiving the more favorable docking score treated as the preferred sample. Correlation analysis in **Supplementary Fig. 6** showed that docking scores had weak within-series correlations with experimentally measured activities. After reassignment, the docking score-derived preference labels agreed with the experimental-activity-derived labels for only

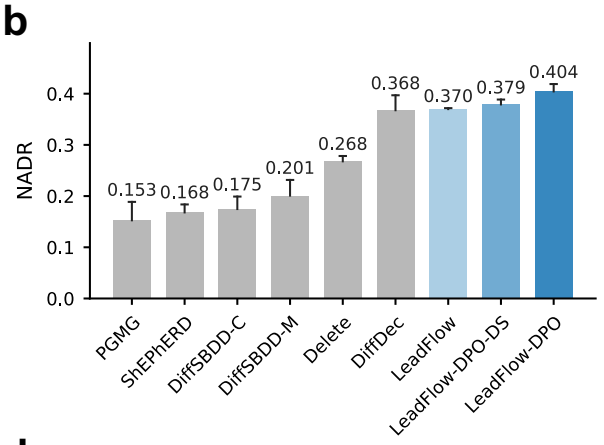
Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

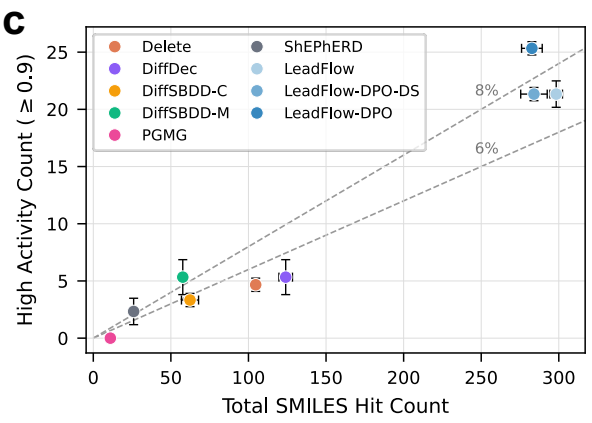
467 approximately 65.5% of the pairs (**Supplementary Fig. 7**), indicating limited consistency between docking
468 scores and measured activities and suggesting that docking score-based preferences introduce substantial
469 label noise.

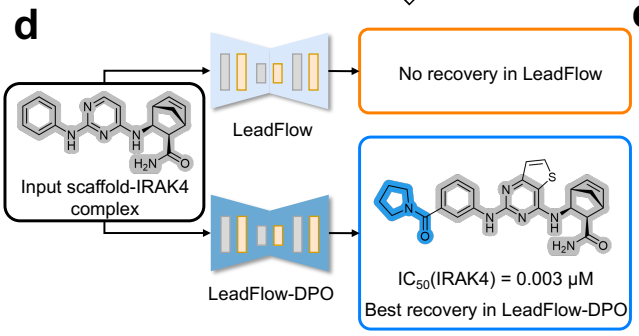
470 DPO fine-tuning was then performed by optimizing a trainable LeadFlow against a frozen reference
471 model, both initialized from the same pre-trained LeadFlow. During fine-tuning, the preference pairs
472 encourage the trainable model to assign higher relative likelihood to preferred modifications than to
473 dispreferred ones, while the frozen reference model serves as a stable anchor that constrains the trainable
474 model from drifting too far from the original generative distribution. In addition, the contribution of each
475 preference pair was modulated by the activity difference between the preferred and dispreferred
476 modifications, such that pairs with larger potency gaps provided a stronger optimization signal. The detailed
477 formulation of DPO fine-tuning is described in **Methods Section 4.2.2**.

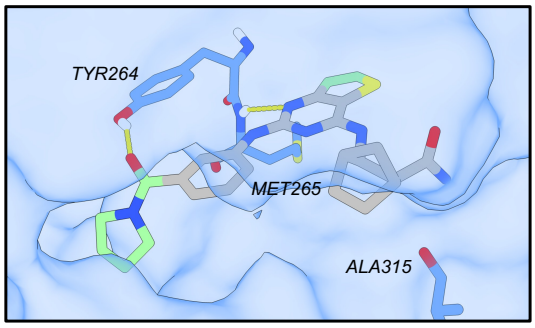
Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

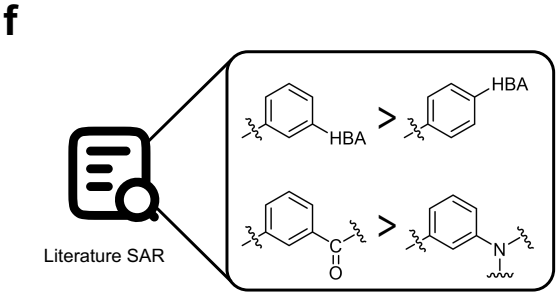


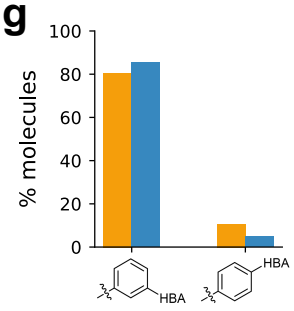


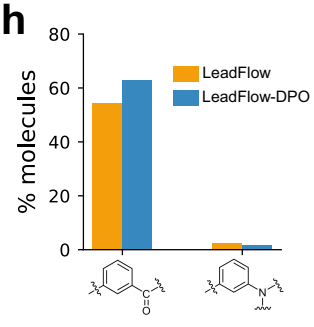












Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Fig. 6 | Experimental activity-based DPO enriches recovery toward higher-activity molecules. a. Schematic illustration of DPO fine-tuning. Higher-activity and lower-activity modifications measured within the same assay are used to construct preference pairs. A trainable LeadFlow model is optimized against a frozen reference model, and the preference signal is weighted by the experimental activity difference, $(r^w - r^l)$. **b.** NADR scores across different models on the MolGenBench hit-to-lead benchmark. Data are presented as mean ± s.d. (n = 3 sampling runs with different seeds from fixed trained checkpoints). **c.** Comparison of total and high activity SMILES hit counts across different models. Data are presented as mean ± s.d. (n = 3 sampling runs with different seeds from fixed trained checkpoints). **d.** Comparison of active molecule recovery between LeadFlow and LeadFlow-DPO given an input scaffold-IRAK4 complex. **e.** Interaction pattern of the active molecule recovered by LeadFlow-DPO binding to IRAK4. The blue region denotes the protein, and the gray region denotes the starting scaffold, both provided as model inputs, whereas the green region represents the structural modification generated by LeadFlow-DPO. Yellow dashed lines indicate hydrogen bond interactions. **f.** Scaffold-specific literature SAR rules for IRAK4 inhibitors. **g.** Proportion of molecules generated by LeadFlow and LeadFlow-DPO bearing HBA at the meta-position and the para-position. **h.** Proportion of molecules generated by LeadFlow and LeadFlow-DPO bearing meta-substituted carbonyl groups or tertiary amine substituents.

To compare different models in their ability to generate higher-activity molecules, we used NADR, a metric designed to quantify whether the recovered hits are enriched for more potent compounds within each reference chemical series. By assigning greater weight to molecules with higher activity, NADR rewards overall recovery while giving larger contributions to the most active recovered molecules. The detailed definition and calculation of NADR are provided in **Methods Section 4.3**. As shown in **Fig. 6b**, the pre-trained LeadFlow already achieved a strong NADR score of 0.370, outperforming most baseline methods, including Delete (0.268), DiffSBDD-M (0.201), DiffSBDD-C (0.175), ShEPhERD (0.168), and PGMG (0.153), and slightly exceeding DiffDec (0.368). LeadFlow-DPO-DS achieved a comparable NADR score of 0.379, remaining close to that of the pre-trained LeadFlow. LeadFlow-DPO increased the NADR to 0.404, achieving the highest score among the evaluated models and giving higher scores than LeadFlow and LeadFlow-DPO-DS, suggesting that assay-controlled activity pairs provide a more effective DPO signal for shifting the generative distribution toward higher-activity molecules.

More direct evidence is provided in **Fig. 6c**, where we counted both the total number of recovered molecules and the number of high-activity recovered molecules. Here, high-activity molecules were defined as those with min-max normalized activity greater than 0.9 within each chemical series. Compared with the pre-trained LeadFlow baseline, LeadFlow-DPO reduced the total number of recovered SMILES hits while increasing the number of high-activity recoveries. The proportion of high-activity molecules among all recovered hits increased from 7.15% to 8.96%, indicating enhanced enrichment for high-activity candidates. By contrast, LeadFlow-DPO-DS did not produce a clear increase in the number of high-activity recovered

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

molecules relative to LeadFlow. This result suggests that the fidelity of preference signal is critical. Experimentally measured activities provide a directly comparable signal, whereas docking-score-derived preferences are only partially aligned with true potency and can dilute the benefit of DPO fine-tuning. The same trend is observed across a range of evaluated normalized-activity thresholds (**Supplementary Fig. 8**). This pattern is consistent with a shift toward higher-activity reference molecules under the evaluated sampling budget. Importantly, DPO fine-tuning largely preserved the overall quality of generated molecules, with only slight decreases in uniqueness and diversity (**Supplementary Table 1**).

To illustrate the observed activity-enrichment trend, we performed a retrospective analysis on an IRAK4 inhibitor scaffold to provide a concrete illustration of how DPO fine-tuning shapes the generative behavior of LeadFlow. IRAK4 plays a critical role in transduction of IL-1R/TLR signaling underlying innate immunity, and aberrant activation of this pathway is implicated in inflammatory and autoimmune diseases as well as in hematological malignancies[54,55]. Targeting IRAK-4 to develop inhibitors and degraders is an important direction for the treatment of related diseases, and several IRAK4-targeting agents are currently being evaluated in preclinical and clinical studies[56].

Chen et al. identified an IRAK4 inhibitor through high throughput screening and conducted SAR exploration around its scaffold[57]. We therefore applied both models to this scaffold-IRAK4 complex and compared their recovery of active molecules, with 200 molecules generated per round across three independent rounds for each model (**Fig. 6d**). While LeadFlow failed to recover any active molecules, LeadFlow-DPO successfully recovered a molecule with an IC_{50} of 0.003 μM against IRAK4. Notably, IRAK4 was absent from both the training datasets, and the recovered molecule exhibited low structural similarity to the training set, with a maximum ECFP4 Tanimoto similarity of 0.347. In the predicted pose of the recovered molecule, the carbonyl group is positioned to accept a hydrogen bond from TYR264 within the ATP-binding pocket of IRAK4 (**Fig. 6e**). Redocking using Glide yielded an RMSD of 0.400 Å between the generated and redocked conformations (**Supplementary Fig. 9**), confirming that the generated binding mode is stable and reliable.

This prompted us to revisit their SAR findings conducted by Chen et al. (**Fig. 6f**), from which two key rules can be summarized: meta-substituted hydrogen bond acceptors (HBA) on the phenyl ring enhance binding potency against IRAK4 compared to para-substituted ones, as the meta-position places the substituent in a more favorable spatial orientation for hydrogen bond formation with TYR264. Additionally, meta-substituted carbonyl groups confer greater potency than tertiary amine substituents. This can be attributed to two factors: the meta-substituted carbonyl group is geometrically closer to TYR264, and the lone pair of amine nitrogen is delocalized through conjugation with the aromatic ring, weakening its

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

hydrogen bond accepting ability, whereas the carbonyl group is less susceptible to such electronic attenuation. The corresponding activity data is provided in **Supplementary Table 2**.

We therefore examined the differences in generation behavior between the two models for molecules built around this scaffold with respect to these SAR patterns. HBA groups on the phenyl ring were identified by SMARTS-based substructure matching, with the full list of patterns and raw counts of SAR-compliant molecules provided in **Supplementary Tables 3, 4**. As shown in **Fig. 6g**, LeadFlow-DPO showed a higher proportion of molecules bearing HBA at the meta-position and LeadFlow showed a higher proportion at the para-position. At the meta-position, LeadFlow-DPO showed a higher tendency to generate carbonyl substituents, whereas LeadFlow generated a relatively higher proportion of tertiary amine substituents (**Fig. 6h**). These results suggest that pre-training enables the model to learn general activity-relevant modification patterns. DPO fine-tuning further refined the model's understanding of structure-activity relationships.

The retrospective results indicate complementary roles for the two training stages: scaffold-conditioned pretraining supports reference-active coverage, whereas experimental preference alignment concentrates recovery toward higher-activity compounds. We next evaluated the practical utility of the LeadFlow-DPO-guided workflow in a prospective design-synthesis-test campaign.

2.7 A LeadFlow-DPO-guided workflow yields a ClpP agonist with 156-fold improved potency

The preceding benchmark and retrospective analyses demonstrated that LeadFlow can generate chemically and geometrically plausible molecules, recover experimentally validated active chemotypes, and learn activity-relevant scaffold modification patterns. To determine whether these capabilities could translate into experimentally validated potency improvements, we conducted a design-synthesis-test study targeting ClpP. ClpP is a mitochondrial serine protease that maintains mitochondrial proteostasis and respiratory function by degrading damaged or misfolded proteins[58-61]. It is frequently upregulated in hematologic malignancies and solid tumors, where cancer cells can become dependent on its proteolytic activity[62-66], making ClpP an attractive anticancer target. We selected A1, a ClpP agonist chemotype previously derived from ONC201 through ring-opening simplification[67]. A1 showed modest agonistic activity (EC_{50} = 6,052.0 ± 1,743.1 nM) in a fluorogenic cleavage assay using purified recombinant human ClpP (residues 57–277) and Ac-WLA-AMC as the substrate. ClpP activation was quantified from the increase in fluorescence resulting from substrate cleavage, providing an experimentally validated starting point with substantial optimization potential (**Fig. 7a**).

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

The A1-derived core scaffold and the activator-binding pocket extracted from the ONC201-bound ClpP structure (PDB 6DL7) were used as conditioning inputs for LeadFlow-DPO. Neither ClpP nor the A1 scaffold was present in the training datasets, and the maximum ECFP4-Tanimoto similarity between A1 scaffold and any training molecule was 0.39. A single generation run produced 10,000 molecular structures, corresponding to 3,147 unique candidates after largest-fragment retention and deduplication. Seven substitution sites were identified on the input scaffold, with approximately 71% of substitution events occurring on the two terminal aryl arms (**Fig. 7b, c**). This distribution was consistent with the predicted U-shaped binding mode of A1, in which arms A and B occupy two distinct subpockets. Pocket-surface analysis further revealed unoccupied volumes adjacent to both arms, suggesting these regions as plausible modification vectors for subsequent prioritization (**Fig. 7d**).

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

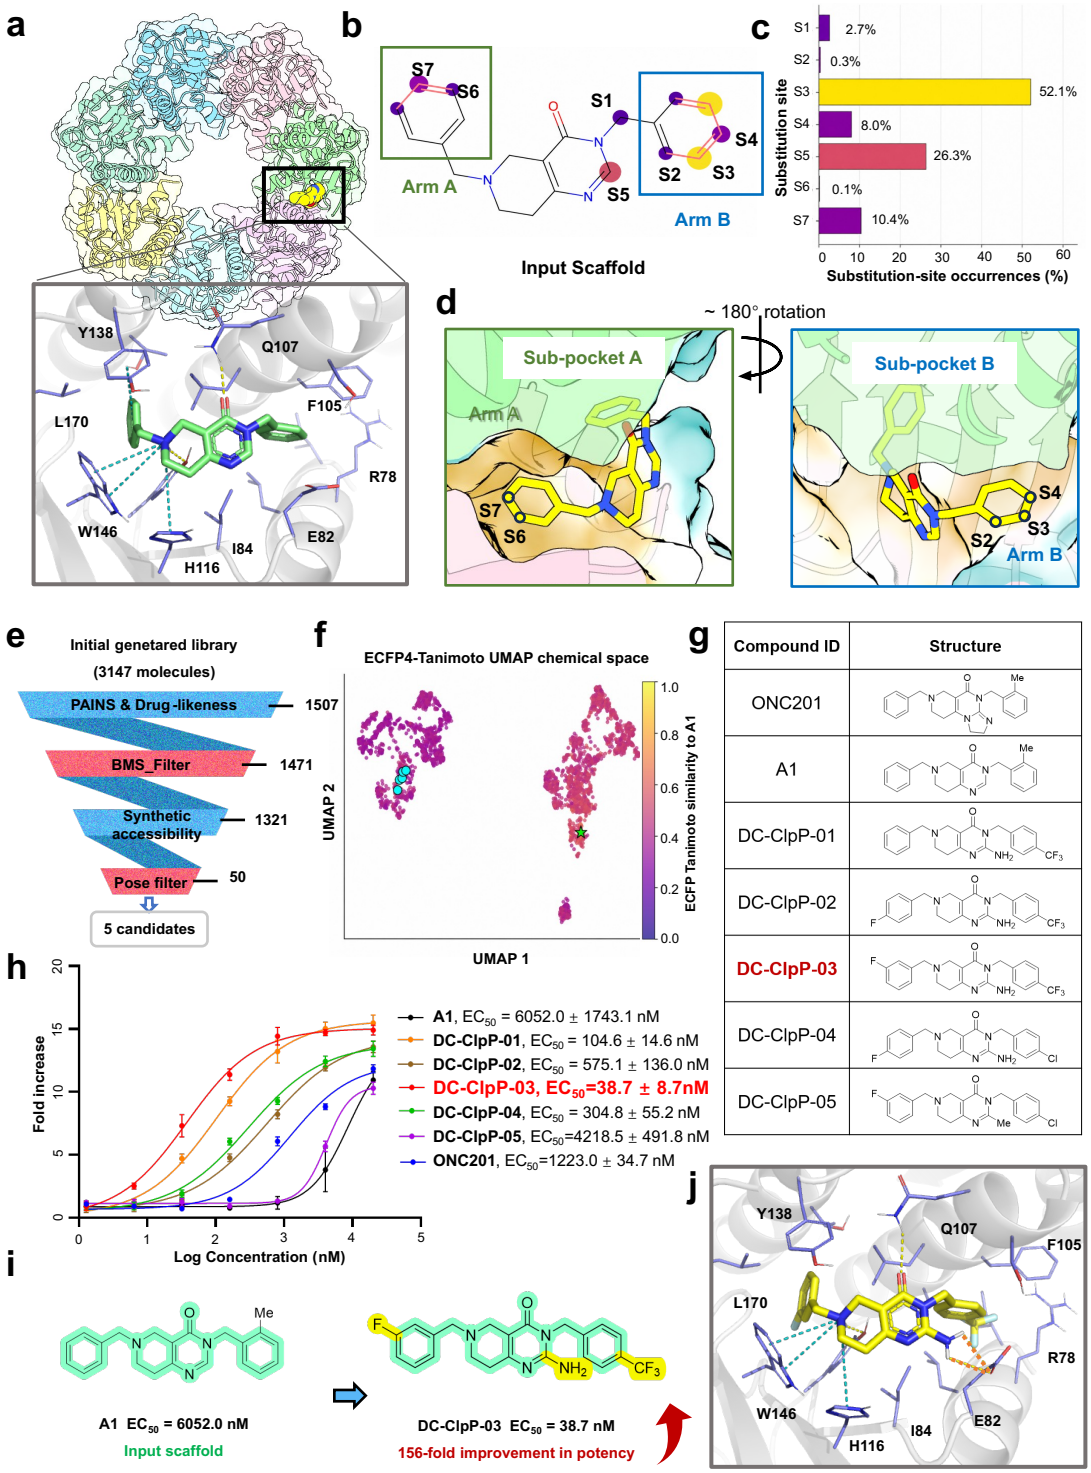


Fig. 7 | A LeadFlow-DPO-guided generation-and-prioritization workflow yields a ClpP agonist with 156-fold improved potency. a. Heptameric ClpP assembly and predicted binding mode of A1 in the activator-binding pocket. **b.** Seven substitution sites mapped onto the input scaffold, with aryl arms A and B indicated. **c.** Proportion of generated molecules associated with each substitution site. **d.** Complementary surface views of subpockets A and B

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

accommodating the two aryl arms. **e.** Sequential filtering and candidate-prioritization workflow. **f.** UMAP projection of the generated chemical space based on ECFP4-Tanimoto distances. Points are colored by similarity to A1; A1 and the selected candidates are indicated by a green star and cyan circles, respectively. **g.** Structures of ONC201, A1 and the five synthesized candidates. **h.** Concentration-response curves for ClpP activation. Data are mean ± s.d. (n = 4 technical replicates). **i.** Structural evolution from A1 to DC-ClpP-03 and the corresponding 156-fold potency improvement. Introduced substituents are highlighted in yellow. **j.** Predicted binding mode of DC-ClpP-03 in ClpP. Hydrogen bonds to Q107 and E82 are shown as yellow and orange dashed lines, respectively; aromatic contacts are shown as cyan dashed lines.

The generated molecules were subsequently prioritized through a sequential computational filtering and medicinal chemistry selection workflow (**Fig. 7e; Methods Section 4.5**). This procedure retained 50 candidates, of which five were prioritized based on known ClpP SAR and successfully synthesized and tested (**Fig. 7g**). UMAP visualization based on ECFP4-Tanimoto distances showed that the generated molecules occupied multiple regions beyond the immediate neighborhood of A1, indicating structural diversification around the conditioned scaffold (**Fig. 7f**).

Experimental testing showed that all five prioritized compounds exhibited lower mean EC_{50} values than A1 under identical assay conditions, and four also showed lower EC_{50} values than ONC201. DC-ClpP-01 and DC-ClpP-03 were the two most potent compounds, with EC_{50} values of 104.6 ± 14.6 and 38.7 ± 8.7 nM, respectively. DC-ClpP-03 was the most potent compound, corresponding to an approximately 156-fold potency improvement over A1 and a 32-fold improvement over ONC201 (**Fig. 7g-i**). Analysis of the matched analogs revealed a pronounced positional effect. Based on the mean EC_{50} values, the S6 meta-fluoro analog DC-ClpP-03 was 2.7-fold more potent than nonfluorinated DC-ClpP-01, whereas the S7 para-fluoro analog DC-ClpP-02 (EC_{50} = 575.1 ± 136.0 nM) was 5.5-fold less potent than DC-ClpP-01, resulting in a 14.9-fold difference between the two fluoro analogs. Although the predicted poses did not provide a clear structural explanation for this effect, docking suggested that DC-ClpP-03 maintained key contacts with H116 and W146 and a putative hydrogen bond with Q107. The introduced amino group was oriented toward E82, consistent with an additional potential hydrogen bond, whereas the trifluoromethyl group occupied an adjacent hydrophobic region (**Fig. 7j**). The LeadFlow-DPO-guided workflow yielded experimentally measured potency improvements in this ClpP campaign. These outcomes reflect the combined contributions of generation, computational filtering and candidate selection.

3 Discussion

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

LeadFlow highlights the importance of aligning generative training-task construction and subsequent preference optimization with real-world molecular optimization. The scaffold-definition ablations support the contribution of series-derived MCS contexts to broad active-chemotype recovery. Building on this capacity for exploration, experimental activity-based DPO shifted recovery toward higher-activity reference molecules and outperformed the docking-score-based counterpart on activity-weighted recovery. Lower total exact-SMILES recovery and modest reductions in uniqueness and diversity accompanied this gain, indicating a trade-off between coverage and activity enrichment. The IRAK4 case provides a scaffold-specific illustration of the resulting shift toward published SAR patterns. These findings identify training-task construction as an important component of molecular optimization, alongside the generative architecture. Reference-active recovery measures coverage of known chemical space but does not establish the activity of novel outputs. Moreover, the benchmark supplies scaffolds derived from known active series, whereas prospective optimization requires selecting a scaffold from the starting hit.

The prospective ClpP campaign tested whether this approach could contribute to experimentally useful modifications. Generation followed by structural and synthetic prioritization yielded five compounds with lower mean EC_{50} values than the starting compound A1, with DC-ClpP-03 showing an approximately 156-fold improvement. This outcome supports the practical utility of the combined workflow in the evaluated campaign. Its attribution remains at the workflow level because candidate selection contributed to the result and the prospective study did not isolate the effect of DPO. The campaign also exposes a limit of the computational interpretation. Meta-fluorination improved potency relative to the corresponding unsubstituted analog, whereas para-fluorination reduced it, producing a difference that neither site-level generation frequencies nor inspection of the predicted poses readily explained. Local packing, desolvation, water organization and conformational free-energy costs could contribute despite similar predicted poses[68]. Because EC_{50} measures functional activation, differential stabilization of activation-competent ClpP conformations is another possibility[69]. These remain untested explanations. The positional effect therefore illustrates that proposing productive modifications and resolving fine-grained SAR are distinct capabilities.

Within this scope, LeadFlow offers a route to integrating scaffold-conditioned generation into design-make-test-analyze cycles. Series-derived supervision supports exploration, while historical within-assay activity preferences can steer generation. The ClpP study evaluated one generation-and-selection campaign, iterative retraining on its new measurements remains a future extension. Extending this framework to measured ADMET preferences and additional prospective campaigns could address multi-objective optimization and experimental efficiency.

4 Methods

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

4.1 Dataset

4.1.1 Homologous Series Dataset

The homologous series dataset was collected from the ChEMBL[70] database (version 34) and contains protein-ligand records with experimental activity measurements (IC_{50}, EC_{50}, K_i, and K_d). Ligands failing RDKit molecular sanitization and activity records containing null values, 0, or containing inequality symbols '<' or '>' were excluded. Entries were initially partitioned by UniProt ID to group ligands by protein target and further stratified into chemical series based on ChEMBL Assay ID. Grouping by assay reduces differences in experimental context when constructing activity preferences for DPO training. To prevent potential data leakage, we excluded all protein targets included in the MolGenBench test set based on UniProt IDs. Furthermore, any molecules identical to those in the MolGenBench test set or sharing the same BMScaffold were removed from the dataset. This filtering process resulted in a final curated dataset comprising 31,917 data entries distributed across 2,922 chemical series. From the chemical series whose protein targets were not shared with any other series, 50 series containing 399 data entries were randomly selected as the validation set, and the remaining series in the curated dataset were used as the training set. Representative PDB structures were selected for each chemical series by identifying co-crystallized ligands with the highest Tanimoto similarity to the series compounds. Protein and ligand structures were prepared using Schrödinger's Protein Preparation Wizard and LigPrep with default parameters. Receptor grids were centered on the native co-crystallized ligand, and docking was performed via Glide[29] with default parameters to generate a docked pose for each ligand. All Schrödinger-related operations were performed using version 2023-1.

To construct the dataset for the molecular structural modification generative model, we utilized RDKit to identify the MCS as the scaffold to be modified for each chemical series. A valid MCS was required to contain more than one-third of the atoms of each retained molecule in a chemical series. A unified training framework was established by applying specific masking strategies. For molecules with a valid MCS, the shared substructure is retained while the remaining portions are masked, requiring the model to elaborate the existing scaffold to reconstruct the full ligand structure. In cases where no valid MCS is identified, or the ligand is identical to the MCS, the entire ligand was masked to support full-ligand reconstruction training. This strategy includes both scaffold-conditioned and full-ligand reconstruction examples under the same training objective. An overview of the curated homologous chemical series dataset is provided in **Supplementary Fig. 10**.

4.1.2 Preference alignment dataset

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

We constructed the preference alignment dataset for DPO training based on the training data described above. Specifically, within each chemical series in the training set, we paired molecules that share the same core scaffold but exhibit an activity difference greater than 5-fold. This procedure yields a dataset containing 546,231 preference pairs from 1,748 eligible chemical series. This dataset is defined as $\mathcal{D} = \{(\mathbf{p}, \mathbf{x}_S, \mathbf{x}_M^w, \mathbf{x}_M^l)\}$, where $\mathbf{p}$ denotes the protein target, $\mathbf{x}_S$ represents the starting scaffold, $\mathbf{x}_M^w$ denotes the preferred scaffold modification which exhibits higher activity, and $\mathbf{x}_M^l$ denotes the rejected modification which exhibits lower activity. By fine-tuning on these pairwise comparisons, the model learns to distinguish high-activity modification from low-activity ones given the same core structure and protein context. This encourages the model to preferentially generate high-activity modification when conditioned on $(\mathbf{p}, \mathbf{x}_S)$.

4.1.3 MolGenBench hit-to-lead test set

MolGenBench[7] is a comprehensive benchmark specifically designed to evaluate SBDD models on both *de novo* molecular generation and hit-to-lead optimization tasks. The hit-to-lead dataset contains 120 protein targets, with the top five representative chemical series selected for each target, resulting in a total of 600 series. Constructed from the ChEMBL database, each series comprises experimentally verified active molecules from the same assay to ensure bioactivity comparability by mitigating inter-assay discrepancies. For each series, the generative input is derived from the docked pose of a representative active molecule by retaining only the 3D conformation of the shared MCS. The model is required to perform structural modification based on this conserved core and the protein pocket to recover reference active ligands, particularly those exhibiting higher bioactivity. For each series, the model generates 200 molecules per round across three independent sampling rounds initialized with different seeds.

4.2 Model implementation

Notation

SBMO aims to model the conditional probability $p(\mathbf{x}_M | \mathbf{p}, \mathbf{x}_S)$, where the generated modification $\mathbf{x}_M$ combined with the invariant scaffold $\mathbf{x}_S$ constitutes the complete ligand $\mathbf{x}_L = \mathbf{x}_S \cup \mathbf{x}_M$. Let N_L, N_S, and N_M denote the numbers of atoms in $\mathbf{x}_L$, $\mathbf{x}_S$, and $\mathbf{x}_M$, respectively, with $N_L = N_S + N_M$. The modification is represented as $\mathbf{x}_M = (\mathbf{r}_M, \mathbf{h}_M, \mathbf{A}_M)$. Here, $\mathbf{r}_M \in \mathbb{R}^{N_M \times 3}$ denotes the continuous coordinates of modification atoms, and $\mathbf{h}_M \in \mathbb{R}^{N_M \times K_h}$ denotes their atom types in one-hot form, where K_h is the number of atom types. For bond types, let $\mathcal{E}_{\mathrm{fc}}(N_L)$ denote the directed fully connected atom-pair set of the complete ligand with self-loops excluded. Let $E_M \subseteq \mathcal{E}_{\mathrm{fc}}(N_L)$ be the subset of atom pairs involving modification atoms. We represent bond types on this subset as $\mathbf{A}_M \in \mathbb{R}^{|E_M| \times K_A}$, where K_A is the number of bond types.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

The protein target is represented as $\mathbf{p} = (\mathbf{r}_P, \mathbf{h}_P)$. To simplify the generative formulation, the modification data is partitioned into continuous variables $\mathbf{x}_c := \mathbf{r}_M$ and discrete variables $\mathbf{x}_d := (\mathbf{h}_M, \mathbf{A}_M)$. The model generates $\mathbf{x}_M$ conditioned on the 3D environment provided by the protein $\mathbf{p}$ and invariant scaffold $\mathbf{x}_S$. When no valid MCS is identified or the ligand is identical to the MCS, we mask the entire ligand and treat it as *de novo* generation, in which case $\mathbf{x}_S = \emptyset$. For brevity, these conditions are omitted from notation where the context is unambiguous. Detailed featurization of the protein and ligand inputs is provided in the **Supplementary Table 5, 6**.

4.2.1 Scaffold Modification with Bayesian Flow Networks

Modeling the synergistic relationship between continuous 3D atomic geometry and discrete chemical topology requires a generative framework capable of integrating these different modalities within a unified formulation. We employ BFN, which provide a robust mathematical basis for handling hybrid modality data by operating within a continuous parameter space $\boldsymbol{\theta}$. Unlike conventional diffusion models that directly perturb discrete or continuous variables, BFN model the evolution of distribution parameters through iterative Bayesian updates.

The generative process of BFN initiates with an uninformative prior $\boldsymbol{\theta}_0$ and progressively refines the distribution parameters to produce a sequence of posteriors $\boldsymbol{\theta}_i$ that incorporate increasing information. Specifically, a standard multivariate normal distribution is employed as the prior for the continuous atomic positions $\mathbf{x}_c$, while a uniform categorical distribution is used for the discrete variables $\mathbf{x}_d$. The refinement is governed by a message-passing process between the sender and receiver distributions, where the sender's visibility is confined to the sample space, while the receiver infers messages based on its interpretation of samples and parameters.

Within each communication round of the BFN, the sender distribution $p_S(\mathbf{y} \mid \mathbf{x}; \alpha)$ perturbs the molecular modification data $\mathbf{x} = (\mathbf{x}_c, \mathbf{x}_d)$ depending on modality-specific noise schedules and sends the observed noisy latent variables $\mathbf{y}$ to the receiver distribution. We discretize $t \in [0,1]$ into T steps with indices $i \in \{1, \dots, T\}$, where $t_i = i/T$ and $\alpha_i \triangleq \alpha(t_i)$. In the following, we use t interchangeably with the discrete index i for notational simplicity. Then the sender distribution p_S can be formulated as follows:

$$p_S(\mathbf{y} \mid \mathbf{x}; \alpha) = \begin{cases} \mathcal{N}(\mathbf{x}_c, \alpha_c(t)^{-1}\mathbf{I}), & \text{for continuous data } \mathbf{x}_c \\ \mathcal{N}\left(\alpha_d(t)\left(K\mathbf{e}_{\mathbf{x}_d} - \mathbf{1}\right), \alpha_d(t)K\mathbf{I}\right), & \text{for discrete data } \mathbf{x}_d \end{cases} \quad (1)$$

where K is the number of possible types. For the discrete data $\mathbf{x}_d \in \{1, \dots, K\}$, $\mathbf{e}_{\mathbf{x}_d}$ is the one-hot encoding of $\mathbf{x}_d$, and the corresponding noisy observation $\mathbf{y}_d \in \mathbb{R}^K$ is also represented as K-dimensional vector. The

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

noise levels $\alpha_c(t)$ and $\alpha_d(t)$ are derived from monotonically increasing, differentiable noise schedule functions $\beta_c(t)$ and $\beta_d(t)$:

$$\beta_c(t) = \sigma_1^{-2t} - 1,\ \beta_d(t) = \beta_1 t^2, \tag{2}$$

where $\sigma_1 \in (0,1)$ and $\beta_1 \in \mathbb{R}^+$ are hyperparameters. The per-timestep noise is defined by $\alpha(t) = \beta'(t) := d\beta(t)/dt$

Unlike diffusion models that operate by iteratively denoising noisy latent variables, the BFN receiver distribution employs a neural network ϕ to predict an estimate of the clean data $\hat{\mathbf{x}}_{i-1} = \tilde{\mathbf{x}}_\phi(\boldsymbol{\theta}_{i-1}, t_i, \mathbf{p}, \mathbf{x}_S)$ with its current belief $\boldsymbol{\theta}_{i-1}$. The receiver distribution is formulated by integrating $\hat{\mathbf{x}}_{i-1}$ with the noise scheduler α_i, thereby enabling the sampling of a noisy estimation $\mathbf{y}_i \sim q(\cdot\,|\hat{\mathbf{x}}_{i-1}; \alpha_i)$:

$$p_R(\mathbf{y}_i \mid \boldsymbol{\theta}_{i-1}, \phi; \alpha_i) = q\big(\mathbf{y}_i \mid \tilde{\mathbf{x}}_\phi(\boldsymbol{\theta}_{i-1}, t_i, \mathbf{p}, \mathbf{x}_S); \alpha_i\big) \tag{3}$$

The information gathering process is modeled via a Bayesian update function h, which iteratively refines the prior $\boldsymbol{\theta}_{i-1}$ into the data-informed posterior $\boldsymbol{\theta}_i$ given the noisy estimation $\mathbf{y}_i$:

$$\boldsymbol{\theta}_i = h(\boldsymbol{\theta}_{i-1}, \mathbf{y}_i, \alpha_i) \tag{4}$$

The parameter update for $\boldsymbol{\theta}_i$ given $\boldsymbol{\theta}_{i-1}$ and a data point $\mathbf{x}$ is denoted by the Bayesian update distribution $p_U\big(\boldsymbol{\theta}_i|\boldsymbol{\theta}_{i-1}, \mathbf{x}; \alpha(t_i)\big)$. Exploiting the property of additive accuracy, BFNs define a simulation-free Bayesian flow distribution $p_F(\boldsymbol{\theta}|\mathbf{x}; t)$. This allows for direct computation of the updated parameter at any timestep during training process, bypassing iterative noising:

$$p_F(\boldsymbol{\theta}|\mathbf{x}; t) = \begin{cases} \mathcal{N}\big(\gamma(t)\mathbf{x}_c, \gamma(t)\big(1-\gamma(t)\big)\mathbf{I}\big), & \text{for continuous data } \mathbf{x}_c \\ \delta\Big(\boldsymbol{\theta}_d(t) - \text{softmax}\big(\mathbf{y}_d(t)\big)\Big), & \text{for discrete data } \mathbf{x}_d \end{cases} \tag{5}$$

where $\gamma(t) := \frac{\beta_c}{1+\beta_c}$, and $\mathbf{y}_d(t)$ follows the discrete case in **Eq. 1**

In practice, we follow the noise reduced sampling strategy of MolCRAFT[41] to mitigate the variance introduced by sampling $\mathbf{y}_i$. Specifically, we directly use the network prediction $\hat{\mathbf{x}}_{i-1}$ derived from the previous parameters $\boldsymbol{\theta}_{i-1}$ and obtain $\boldsymbol{\theta}_i$ via Bayesian flow distribution $p_F(\cdot\,|\hat{\mathbf{x}}_{i-1}; t_i)$, thereby bypassing sampling noisy $\mathbf{y}_i$ and improving both sampling efficiency and generation quality.

The model ϕ is trained by minimizing the KL divergence between sender and receiver distributions across timesteps and trajectories, yielding the following loss:

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

$$\mathcal{L}_{\mathrm{BFN}}(\phi) = \mathbb{E}_{\mathbf{x}\sim p_{\mathrm{data}}} \sum_{i=1}^{T} D_{KL}\left(p_S(\mathbf{y}_i \mid \mathbf{x}; \alpha_i) \,||\, p_R(\mathbf{y}_i \mid \boldsymbol{\theta}_{i-1}, \phi; \alpha_i)\right) \tag{6}$$

According to the BFN framework[36], these losses are computed at each timestep to minimize the divergence between the sender and receiver distributions, with the continuous-variable loss $\mathcal{L}_t^c$ and discrete-variable loss $\mathcal{L}_t^d$ defined as follows.

$$\mathcal{L}_t^c = \frac{\alpha_c(t_i)}{2} \left\| \mathbf{x}_c - \tilde{\mathbf{x}}_{\phi,c}(\boldsymbol{\theta}_{i-1}, t_i, \mathbf{p}, \mathbf{x}_S) \right\|^2 \tag{7}$$

$$\mathcal{L}_t^d = \ln p_S\left(\mathbf{y}_d^{(i)} \mid \mathbf{x}_d; \alpha_d(t_i)\right) - \ln p_R\left(\mathbf{y}_d^{(i)} \mid \tilde{\mathbf{x}}_{\phi,d}(\boldsymbol{\theta}_{i-1}, t_i, \mathbf{p}, \mathbf{x}_S); \alpha_d(t_i)\right) \tag{8}$$

where $\tilde{\mathbf{x}}_{\phi,c}(\cdot)$ and $\tilde{\mathbf{x}}_{\phi,d}(\cdot)$ represent the network's estimations of the continuous and discrete variables, respectively. The neural network ϕ is parameterized with an SE(3)-equivariant graph transformer. The detailed architecture of ϕ is described in the **Supplementary Information**.

4.2.2 Direct Preference Optimization

Reinforcement learning from human feedback (RLHF) aligns generative models with human preferences and has been widely used in large language models by first fitting a reward model and then optimizing the generative model with a KL-regularized reinforcement learning algorithm such as Proximal Policy Optimization (PPO)[71]. Compared to PPO, Direct Preference Optimization (DPO) provides an alternative that bypasses explicit reinforcement learning by directly optimizing the model on preference pairs[38]. It can be derived as a closed-form reformulation of the KL-regularized RLHF objective under the Bradley-Terry preference model[72], without training an explicit reward model. More recently, DPO-style preference optimization has been extended beyond language to a broader range of generative models, including diffusion-based and flow-based formulations, and has been explored in applications such as text-to-image generation and biomolecular design[28,37,73]. In our scaffold modification setting, experimental activity measurements provide a natural preference signal, where candidate modifications under the same protein target and starting scaffold are paired by their measured scores to form winning-losing pairs for DPO training. We shift the conditional distribution of scaffold modifications toward higher activity conditioned on the protein target and starting scaffold.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Using our preference alignment dataset $\mathcal{D}$, we treat $(\mathbf{p}, \mathbf{x}_S)$ as the conditioning context and $(\mathbf{x}_M^w, \mathbf{x}_M^l)$ as the preferred and dispreferred scaffold modifications, respectively. The DPO objective is then given by

$$\mathcal{L}_{\mathrm{DPO}}(\phi) = -\mathbb{E}_{(\mathbf{p},\mathbf{x}_S,\mathbf{x}_M^w,\mathbf{x}_M^l)\sim\mathcal{D}}\left[\log\sigma\left(\beta\log\frac{p_\phi(\mathbf{x}_M^w \mid \mathbf{p},\mathbf{x}_S)}{p_{\mathrm{ref}}(\mathbf{x}_M^w \mid \mathbf{p},\mathbf{x}_S)} - \beta\log\frac{p_\phi(\mathbf{x}_M^l \mid \mathbf{p},\mathbf{x}_S)}{p_{\mathrm{ref}}(\mathbf{x}_M^l \mid \mathbf{p},\mathbf{x}_S)}\right)\right] \tag{9}$$

where p_ϕ and p_{ref} denote the conditional distributions induced by the DPO fine-tuned and frozen pre-trained LeadFlow models, respectively. The hyperparameter β controls the strength of preference optimization, and $\sigma(\cdot)$ is the logistic sigmoid function.

Minimizing this loss encourages the fine-tuned model p_ϕ to assign higher relative likelihood to the winning samples than to the losing samples, while remaining close to p_{ref}. In our implementation, we adapted the BFN-DPO framework originally proposed by Lee et al.[74] to the task of scaffold modification. The resulting objective is formulated as follows:

$$\mathcal{L}_{\mathrm{BFN-DPO}}(\phi) = -\mathbb{E}\left[\log\sigma\left(\beta\log\frac{p_\phi(\mathbf{y}_{1:T}^w \mid \boldsymbol{\theta}_{0:T-1}^w)}{p_{\mathrm{ref}}(\mathbf{y}_{1:T}^w \mid \boldsymbol{\theta}_{0:T-1}^w)} - \beta\log\frac{p_\phi(\mathbf{y}_{1:T}^l \mid \boldsymbol{\theta}_{0:T-1}^l)}{p_{\mathrm{ref}}(\mathbf{y}_{1:T}^l \mid \boldsymbol{\theta}_{0:T-1}^l)}\right)\right] \tag{10}$$

where we omit the condition $(\mathbf{p}, \mathbf{x}_S)$ for brevity. Here, $\mathbf{y}_t$ is a noisy sample drawn at time step t from the receiver output, governed by the BFN's sender/receiver distributions, and $\boldsymbol{\theta}_t$ are parameters sampled from the Bayesian flow distribution p_F in **Eq. 5**.

Expanding the ratio into a sum over time steps with Jensen's inequality, one obtains an upper bound involving differences of KL terms at each time step. By approximating the receiver distribution in expectation with the sender distribution, we can obtain the following loss:

$$\mathcal{L}_{\mathrm{BFN-DPO}}(\phi) = -\mathbb{E}\left[\mathcal{L}_t^{c'} + \mathcal{L}_t^{d'}\right] \tag{11}$$

where $\mathcal{L}_t^{c'}$ and $\mathcal{L}_t^{d'}$ are the continuous-variable loss and discrete-variable loss under the DPO setting, respectively.

Let $\mathbf{x}_c^w$ and $\mathbf{x}_c^l$ be the clean coordinates for the winning and losing modifications. The continuous-variable loss is defined as

$$\mathcal{L}_t^{c'} = -\mathbb{E}\left[\log\sigma\left(-\beta T\frac{\alpha_c(t)}{2}\left[\Delta_{\mathbf{x}_c}^w - \Delta_{\mathbf{x}_c}^l\right]\right)\right] \tag{12}$$

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

where

$$\Delta_{\mathbf{x}_c}^{w} = \left\|\mathbf{x}_c^{w} - \tilde{\mathbf{x}}_{\phi,c}(\boldsymbol{\theta}_{i-1}^{w}, t_i, \mathbf{p}, \mathbf{x}_S)\right\|^2 - \left\|\mathbf{x}_c^{w} - \tilde{\mathbf{x}}_{\mathrm{ref},c}(\boldsymbol{\theta}_{i-1}^{w}, t_i, \mathbf{p}, \mathbf{x}_S)\right\|^2 \tag{13}$$

and similarly for $\Delta_{\mathbf{x}_c}^{l}$. Here, $\tilde{\mathbf{x}}_{\phi,c}(\cdot)$ and $\tilde{\mathbf{x}}_{\mathrm{ref},c}(\cdot)$ generate coordinate predictions under the fine-tuned and frozen pre-trained LeadFlow models, respectively.

For discrete variables, we have a similar construction based on KL divergences. The discrete-variable loss has the form of

$$\mathcal{L}_t^{d'} = -\mathbb{E}\left[\log\sigma\left(-\beta T\left[\Delta_{\mathbf{x}_d}^{w} - \Delta_{\mathbf{x}_d}^{l}\right]\right)\right] \tag{14}$$

where

$$\begin{aligned}\Delta_{\mathbf{x}_d}^{w} = &\left[\ln p_S\left(\mathbf{y}_d^{(i,w)} \mid \mathbf{x}_d^{w}; \alpha_d(t_i)\right) - \ln p_R\left(\mathbf{y}_d^{(i,w)} \mid \tilde{\mathbf{x}}_{\phi,d}(\boldsymbol{\theta}_{i-1}^{w}, t_i, \mathbf{p}, \mathbf{x}_S); \alpha_d(t_i)\right)\right] \\ &- \left[\ln p_S\left(\mathbf{y}_d^{(i,w)} \mid \mathbf{x}_d^{w}; \alpha_d(t_i)\right) - \ln p_R\left(\mathbf{y}_d^{(i,w)} \mid \tilde{\mathbf{x}}_{\mathrm{ref},d}(\boldsymbol{\theta}_{i-1}^{w}, t_i, \mathbf{p}, \mathbf{x}_S); \alpha_d(t_i)\right)\right]\end{aligned} \tag{15}$$

and similarly for $\Delta_{\mathbf{x}_d}^{l}$.

To alleviate overfitting to winning samples, we adopt E^{2}PO[28], following its use in prior BFN-DPO work[74]. In our setting, the BFN-DPO surrogate decomposes into three components: one continuous term for coordinates and two discrete terms for atom types and bond types. Therefore, the complementary construction in E^{2}PO yields the constant $K = 3$. We replace the BFN-DPO loss with the following E^{2}PO-modified objective:

$$\mathcal{L}_{\mathrm{BFN-E^2PO}} = -\mathbb{E}\left[\sigma\left(r^w - r^l\right)\mathcal{L}_{\mathrm{BFN-DPO}} + \left(\left(1 - \sigma\left(r^w - r^l\right)\right)\left(K - \mathcal{L}_{\mathrm{BFN-DPO}}\right)\right)\right], K = 3 \tag{16}$$

where r^w and r^l denote the reward values assigned to the winning and losing modifications, respectively. For LeadFlow-DPO, we define pActivity as $r = -\log_{10}(x)$, where x is the experimentally measured IC_{50}, EC_{50}, K_i or K_d value in molar units, corresponding to pIC_{50}, pEC_{50}, pK_i, or pK_d. For LeadFlow-DPO-DS, r was defined as the negative Glide docking score. The reward-dependent weighting in **Eq.16** modulates the strength of preference optimization according to the reward difference between winning and losing modifications. When $\left(r^w - r^l\right)$ is large, $1 - \sigma(r^w - r^l)$ approaches zero, and the objective

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

approaches the original BFN-DPO objective. Conversely, smaller reward differences result in weaker preference optimization.

4.2.3 Protein Pocket Extraction

The protein pocket was defined as the set of residues whose center of mass fell within a radius of 10 Å from any heavy atom of the input MCS scaffold. The center of mass of each residue was computed as the mass-weighted average position of all its constituent atoms, where atomic masses were used as weights.

4.2.4 Atom Number Sampling

To determine the number of atoms to generate for structural modification, we first analyzed the joint distribution of ligand atom counts and pocket sizes across the training set. Pocket size was quantified as the median of the ten largest pairwise distances among all heavy atoms in the pocket, which captures the spatial extent of the binding site without being sensitive to outliers. Pockets were grouped into 10 bins using equal-frequency binning, where bin boundaries were determined by the deciles of the pocket size distribution in the training set. For each bin, we computed the empirical probability distribution over ligand heavy atom counts observed in the training set.

During model inference, the pocket size of the input structure is assigned to the corresponding bin, with the largest bin being open-ended. The total ligand atom count is then sampled from the associated distribution. If the sampled count is less than or equal to the number of scaffold heavy atoms, it is resampled. The number of atoms to be generated is then obtained by subtracting the number of scaffold heavy atoms from the sampled total.

4.2.5 Molecule Construction

Since our model directly generates atomic coordinates, atom types, and bond connectivity for the structural modification, the complete molecule can be assembled by combining the generated atoms with the input scaffold using RDKit.

4.3 Evaluation metrics

We follow the evaluation protocol of MolGenBench[7]. Specifically, the medchem[75] toolkit is employed to compute ChemFilter pass rate, while PoseBusters[51] and PoseCheck[50] are used to evaluate the conformational metrics. Protein-ligand interaction types are analyzed using ProLIF[53]. Hit Recovery, Hit

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Rate, and Hit Fraction are used to evaluate the ability of models to recover experimentally validated active chemotypes, following the same calculation protocol as MolGenBench, with their definitions provided below.

Hit Recovery

Hit Recovery is defined as the successful recovery of at least one known active SMILES or one corresponding BMScaffold among the generated molecules for a test case.

Hit Rate

Hit Rate is defined as the ratio of unique active entities recovered to the total number of molecules sampled, as formalized in **Eq.17**:

$$\text{Hit Rate}_{\text{label}} = \frac{\mathcal{M}_{\text{active}}}{\mathcal{M}_{\text{sampled}}} \; ; \text{label} \in \{\text{SMILES}, \text{BMScaffold}\} \tag{17}$$

Here, $\mathcal{M}_{\text{active}}$ denotes the number of unique active SMILES or molecules bearing active BMScaffolds identified in the sampled set, while $\mathcal{M}_{\text{sampled}}$ represents the total number of generated molecules sampled for evaluation.

Hit Fraction

Hit Fraction is defined as the ratio of unique active entities recovered by a model to the total number of reference active entities for that series, as formalized in **Eq.18**:

$$\text{Hit Fraction}_{\text{label}} = \frac{F_{\text{active}}}{F_{\text{ref}}} \; ; \text{label} \in \{\text{SMILES}, \text{BMScaffold}\} \tag{18}$$

Here, F_{active} denotes the number of unique active SMILES or BMScaffold identified in the model's sampled outputs, while F_{ref} represents the total number of unique reference active SMILES or BMScaffold for the specific series. This metric directly reflects a model's ability to cover the known bioactive space, complementing the hit rate's focus on efficiency.

NADR

To evaluate the model's capacity to prioritize the generation of high-activity candidates, we utilized the Normalized Activity Discounted Rank (NADR). Inspired by the Normalized Discounted Cumulative Gain frequently employed in information retrieval, NADR is employed to evaluate the recovery of high-

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

activity candidates from the reference chemical series. This metric is defined as:

$$\mathrm{NADR} = \sum_{i}^{n} \frac{\mathrm{pActivity}_i^{\mathrm{series}} - \mathrm{pActivity}_{\mathrm{min}}^{\mathrm{series}}}{\mathrm{pActivity}_{\mathrm{max}}^{\mathrm{series}} - \mathrm{pActivity}_{\mathrm{min}}^{\mathrm{series}}} \times \frac{1}{\log_2(1 + \mathrm{Rank}_i)} \tag{19}$$

where n is the number of unique active molecules recovered within the corresponding chemical series in each sampling round, with duplicate recoveries counted only once based on canonical isomeric SMILES. $\mathrm{pActivity}_i^{\mathrm{series}}$ is the activity of the i-th successfully recovered unique active molecule, while $\mathrm{pActivity}_{\mathrm{max}}^{\mathrm{series}}$ and $\mathrm{pActivity}_{\mathrm{min}}^{\mathrm{series}}$ denote the maximum and minimum activities within the corresponding chemical series, respectively. Higher pActivity values indicate stronger activity. Rank_i represents the activity-based rank of the i-th recovered molecule among all reference molecules in the corresponding chemical series, with rank 1 assigned to the most active reference molecule. In this formulation, the activity of each recovered molecule is min-max normalized within its chemical series. A logarithmic rank penalty is applied based on the activity rank, ensuring that high-activity compounds contribute more significantly to the final score than those with lower activity. Higher NADR values indicate stronger activity-weighted recovery of known reference molecules under the evaluated sampling budget. NADR was calculated and averaged across chemical series with at least one recovered active molecule.

4.4 Baseline models

To comprehensively evaluate the performance of LeadFlow, we compared it against a diverse set of baseline models from MolGenBench. We directly adopted the reported performance metrics in the MolGenBench[7] study for the following methods. Based on their input modalities, these baselines could be categorized into structure-based molecular optimization methods and ligand-based methods.

Structure-based molecular optimization methods explicitly utilize the 3D structures of both the scaffold and the protein pocket as inputs to condition the generation process. These methods include:

- **Delete**[16]. An SE(3)-equivariant autoregressive model for lead optimization that utilizes the extracted protein surface to guide geometric generation. The model is trained on the CrossDocked dataset via a masked reconstruction objective for different lead optimization tasks. The version trained with the fragment mask is evaluated on the scaffold modification task.

- **DiffDec**[17]. An SE(3)-equivariant 3D diffusion-based generative model specializing in structure-based scaffold modification. Trained on the CrossDocked dataset using a reaction-based slicing method to define scaffold and R-group pairs. Crucially, the generation process requires explicit modification

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

anchors on the scaffold.

- **DiffSBDD**[18]. An SE(3)-equivariant 3D conditional diffusion model originally developed for structure-based de novo molecule generation. By leveraging its inpainting strategy, this model can be adapted for the scaffold modification task, generating ligands conditioned on both the scaffold and the protein pocket. We include results for versions trained on CrossDocked (DiffSBDD-C) and Binding MOAD (DiffSBDD-M).

Ligand-based methods utilize features extracted from reference active molecules as inputs to guide the generation. These methods include:

- **ShEPhERD**[13]. A diffusion-based generative model that learns the joint distribution of molecular structures and their 3D interaction profiles, generating 3D molecules conditioned on the shapes, electrostatics, and pharmacophores from the reference molecule.

- **PGMG**[14]. A pharmacophore-guided model trained on ChEMBL dataset that employs a GNN encoder and a Transformer decoder to generate molecular SMILES conditioned on pharmacophores extracted from the reference molecule.

- **ShapeMol**[49]. An SE(3)-equivariant diffusion model that utilizes the extracted surface of a small molecule as a conditioning input, aiming to generate 3D molecules with matching geometric shapes.

4.5 LeadFlow-DPO-guided ClpP optimization and candidate prioritization

The A1-derived core scaffold and the activator-binding pocket extracted from the ONC201-bound ClpP structure (PDB 6DL7) were used as conditioning inputs for LeadFlow-DPO. A single sampling run with a random seed of 3407 generated 10,000 molecular structures. Generated structures were processed by retaining the largest molecular fragment, and duplicates were removed based on canonical isomeric SMILES, yielding 3,147 unique candidates. Substitution events were assigned to sites S1-S7 according to their attachment positions on the input scaffold, and site frequencies were calculated over all mapped events. Chemical-space distributions were visualized by UMAP using ECFP4-Tanimoto distances.

The unique candidates were sequentially filtered using rd_filters (https://github.com/PatWalters/rd_filters). The permitted ranges were molecular weight ≤500 Da, $-5 \leq$ LogP ≤ 5, hydrogen-bond donors ≤5, hydrogen-bond acceptors ≤10, rotatable bonds ≤10 and topological polar surface area ≤200 Å^2. Compounds matching PAINS[76] or Bristol-Myers Squibb structural alerts[77] were excluded. Synthetic-accessibility scores were calculated using the RDKit SA_Score implementation, with

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

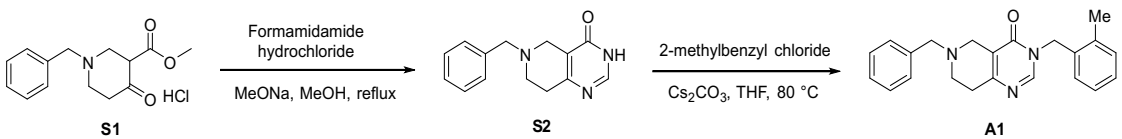
lower values indicating greater predicted synthetic accessibility, and compounds with scores ≤3.0 were retained.

The remaining candidates were redocked into the ClpP activator-binding pocket using Glide in Schrödinger Release 2023-1 with standard-precision mode and otherwise default settings. The receptor grid was centered on the ONC201-binding site. Pose consistency was assessed using the heavy-atom RMSD between the generated and redocked poses. Candidates with RMSD values below 2.0 Å were further inspected for pocket complementarity, steric clashes and preservation of key contacts involving H116, W146 and Q107. This procedure retained 50 candidates (**Supplementary Table 7**), from which five were prioritized through medicinal chemistry review informed by known ClpP structure-activity relationships. All five compounds were direct model outputs without manual structural modification and were successfully synthesized and evaluated using the ClpP activation assay described below.

4.6 Chemical synthesis

General information: starting materials, reagents and solvents were purchased from Bide Pharmatech (Shanghai, China), Adamas-beta (Shanghai, China), Energy Chemical (Shanghai, China), and J&K (Beijing, China), and were used without further purification. Flash Chromatography was performed on a Biotage Isolera One (Biotage AB publ, Uppsala, Sweden) using a Sfär Bio C18 column (Biotage AB publ, Uppsala, Sweden). Nuclear magnetic resonance (NMR) spectroscopy was performed on a Bruker Ascend 500 or 600 MHz NMR instrument with tetramethylsilane (TMS) as the internal standard. Chemical shifts were reported in parts per million (ppm, δ) downfield from tetramethylsilane. Proton coupling patterns were described as singlet (s), doublet (d), triplet (t), quartet (q), multiplet (m), and broad (br). High-resolution mass spectra (HRMS) and low-resolution mass spectra (LRMS) were obtained by electrospray ionization (ESI) using Thermo Exactive Plus (Thermo Fisher Scientific, Waltham, MA, USA) and Agilent 6125C (Agilent Technologies Inc., Santa Clara, CA, USA), respectively. HPLC data analysis of compounds was performed on an Agilent 1290 (Agilent Technologies Inc., Santa Clara, CA, USA) with a quaternary pump and diode-array detector (DAD), the peak purity was verified by UV spectra. All tested compounds had HPLC purity >95%. NMR spectra, HRMS data and HPLC chromatograms are provided in **Supplementary Information**.

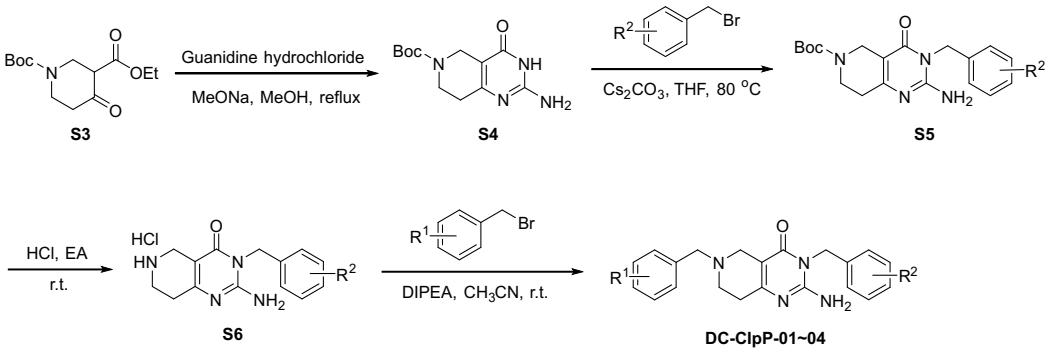
6-benzyl-3-(2-methylbenzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (**A1**).



Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

A mixture of **S1** (2.0 g, 7.1 mmol), formamidine hydrochloride (0.9 g, 11.6 mmol), and MeONa (2.0 g, 35.5 mmol) in MeOH (30 mL) was stirred overnight at 80 °C. The residue was diluted with H_2O (20 mL) and extracted with EA (3 × 30 mL). The combined organic layers were washed with brine, dried over Na_2SO_4, and concentrated. The residue was purified by silica gel chromatography (PE: EA = 2/1 to 1/1) to afford intermediate **S2** as a yellow solid (0.6 g, 36% yield). LC/MS (ESI): m/z 242.2 $[M + H]^+$. Then, a solution of **S2** (0.6 g, 2.5 mmol), 2-methylbenzyl chloride (0.4 mL, 3.3 mmol), and Cs_2CO_3 (1.2 g, 3.7 mmol) in THF (15 mL) was stirred overnight at 80 °C. The mixture was diluted with H_2O (20 mL) and extracted with EA (3 × 20 mL). The combined organic layers were washed with brine, dried over Na_2SO_4, and concentrated. The residue was purified via silica gel chromatography (PE: EA = 3/1 to 1/1) to afford the final compound **A1** as a yellow powder (0.3 g, 35% yield). ^{1}H NMR (500 MHz, DMSO-d_6) δ 8.37 (s, 1H), 7.37 – 7.31 (m, 4H), 7.30 – 7.24 (m, 1H), 7.21 – 7.12 (m, 3H), 6.82 (d, J = 7.6 Hz, 1H), 5.06 (d, J = 2.6 Hz, 2H), 3.68 (s, 2H), 3.24 (s, 2H), 2.75 – 2.63 (m, 4H), 2.33 (s, 3H). ^{13}C NMR (151 MHz, DMSO-d_6) δ 163.76, 159.54, 158.48, 149.87, 138.31, 136.05, 134.80, 130.69, 129.34, 128.73, 127.88, 126.81, 126.55, 119.43, 61.88, 56.56, 49.43, 46.97, 31.44, 19.23; HRMS (ESI) calcd for $C_{22}H_{24}N_3O^+$ 346.1913 $[M + H]^+$, found 346.1912.

General method for the preparation of compounds **DC-ClpP-01~04**.



Typical procedures: Compound **S3** (5.4 g, 20.0 mmol) and guanidine hydrochloride (2.9 g, 30 mmol) were dissolved in methanol (100 mL), MeONa (5.5 g, 100.0 mmol) was added, and the mixture was stirred at 80 °C for 6 h. Upon completion of the reaction, the mixture was cooled to room temperature and concentrated under reduced pressure. The reaction mixture was diluted with DCM (100 mL) and water (50 mL) was added. The aqueous layer was extracted with DCM (3 × 30 mL). The combined organic layers were dried over Na_2SO_4, filtered and concentrated under reduced pressure. The resulting residue was used for the next step without further purification.

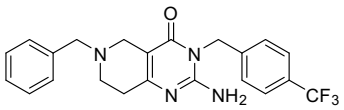
The crude **S4** (1.3 g, 5.0 mmol) from the above step was dissolved in THF (20 mL), and 1.3 equiv substituted benzyl bromide, 1.3 equiv Cs_2CO_3 were added. The mixture was stirred at 80 °C until the starting

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

material was consumed. The reaction was quenched with water (50 mL), and the aqueous layer was extracted with EA (3 × 50 mL). The combined organic layers were dried over Na_2SO_4 and concentrated under reduced pressure. The residue was purified by silica gel chromatography (PE: EA = 3/1 to 1/1) to afford intermediate **S5** (35%~38% yield).

To a solution of the compound **S5** in DCM (c = 0.1 M) was added HCl (20 equiv., 2.0 mol/L in EA) and the mixture was stirred at room temperature. After completion of the reaction, the solvent was removed in vacuo to afford intermediate **S6**. Then, **S6** and DIPEA (3.0 equiv.) were dissolved in MeCN (c = 0.2 M), substituted benzyl bromide (1.5 equiv.) was added, and the mixture was stirred at room temperature. Upon completion of the reaction, the mixture was diluted with H_2O (20 mL) and extracted with EA (3 × 30 mL). The combined organic layers were washed with brine, dried over Na_2SO_4, and concentrated. The residue was purified by silica gel chromatography (PE: EA = 3/1 to 1/1) to afford the final compound **DC-ClpP-01~04** (63%~74% yield).

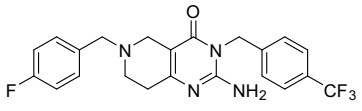
2-amino-6-benzyl-3-(4-(trifluoromethyl)benzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (**DC-ClpP-01**)



DC-ClpP-01

Yield = 66%. White solid. ^{1}H NMR (600 MHz, $CDCl_3$) δ 7.62 (d, *J* = 8.0 Hz, 2H), 7.41 – 7.32 (m, 6H), 7.31 – 7.25 (m, 1H), 5.24 (s, 2H), 5.09 (s, 2H), 3.73 (s, 2H), 3.46 (s, 2H), 2.71 (t, *J* = 5.8 Hz, 2H), 2.56 (t, *J* = 5.9 Hz, 2H). ^{13}C NMR (151 MHz, $CDCl_3$) δ 161.00, 159.04, 152.92, 138.75, 137.93, 130.61 (q, *J* = 32.7 Hz), 129.22, 128.39, 127.29, 127.02, 126.27 (q, *J* = 3.7 Hz), 123.80 (q, *J* = 272.2 Hz), 109.45, 62.64, 50.17, 49.23, 44.44, 31.69. ^{19}F NMR (471 MHz, $CDCl_3$) δ -62.70. HRMS (ESI) calcd for $C_{22}H_{22}F_3N_4O^+$ 415.1741 $[M + H]^+$, found 415.1727.

2-amino-6-(4-fluorobenzyl)-3-(4-(trifluoromethyl)benzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (**DC-ClpP-02**)



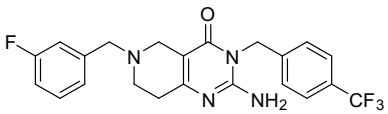
DC-ClpP-02

Yield = 63%. White solid. ^{1}H NMR (500 MHz, $CDCl_3$) δ 7.64 (d, *J* = 8.0 Hz, 2H), 7.42 – 7.30 (m, 4H), 7.03 (t, *J* = 8.6 Hz, 2H), 5.27 (s, 2H), 4.83 (brs, 2H), 3.70 (s, 2H), 3.45 (s, 2H), 2.72 (t, *J* = 5.7 Hz, 2H), 2.59 (t, *J* = 5.8 Hz, 2H). ^{13}C NMR (151 MHz, $CDCl_3$) δ 162.16 (d, *J* = 245.1 Hz), 160.97, 159.02, 152.79, 138.70, 133.60, 130.70 (q, *J* = 32.7 Hz), 130.69 (d, *J* = 8.1 Hz), 127.04, 126.33 (q, *J* = 3.7 Hz), 123.78 (q, *J* = 272.2 Hz), 115.21 (d, *J* = 21.2 Hz), 109.48, 61.81, 50.02, 49.26, 44.52, 31.70. ^{19}F NMR (471 MHz, $CDCl_3$) δ -62.71, -115.62. HRMS (ESI) calcd for $C_{22}H_{21}F_4N_4O^+$ 433.1646 $[M + H]^+$, found 433.1635.

2-amino-6-(3-fluorobenzyl)-3-(4-(trifluoromethyl)benzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

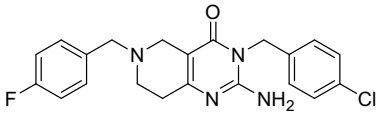
4(3H)-one (**DC-ClpP-03**)



DC-ClpP-03

Yield = 74%. White solid. ^{1}H NMR (500 MHz, $CDCl_3$) δ 7.63 (d, *J* = 8.1 Hz, 2H), 7.37 (d, *J* = 8.0 Hz, 2H), 7.33 – 7.26 (m, 1H), 7.16 – 7.10 (m, 2H), 6.97 (td, *J* = 8.5, 2.6 Hz, 1H), 5.25 (s, 2H), 5.06 (brs, 2H), 3.72 (s, 2H), 3.45 (d, *J* = 2.2 Hz, 2H), 2.72 (t, *J* = 5.8 Hz, 2H), 2.57 (t, *J* = 5.8 Hz, 2H). ^{13}C NMR (126 MHz, $CDCl_3$) δ 163.02 (d, *J* = 245.8 Hz), 160.97, 158.96, 152.92, 140.83 (d, *J* = 7.0 Hz), 138.71, 130.66 (q, *J* = 32.8 Hz), 129.79 (d, *J* = 8.2 Hz), 127.02, 126.29 (q, *J* = 3.8 Hz), 124.55 (d, *J* = 2.7 Hz), 123.79 (q, *J* = 272.1 Hz), 115.74 (d, *J* = 21.3 Hz), 114.16 (d, *J* = 21.3 Hz), 109.35, 62.03, 50.10, 49.36, 44.48, 31.67. ^{19}F NMR (471 MHz, $CDCl_3$) δ -62.70, -113.56. HRMS (ESI) calcd for $C_{22}H_{21}F_4N_4O^+$ 433.1646 $[M + H]^+$, found 433.1636.

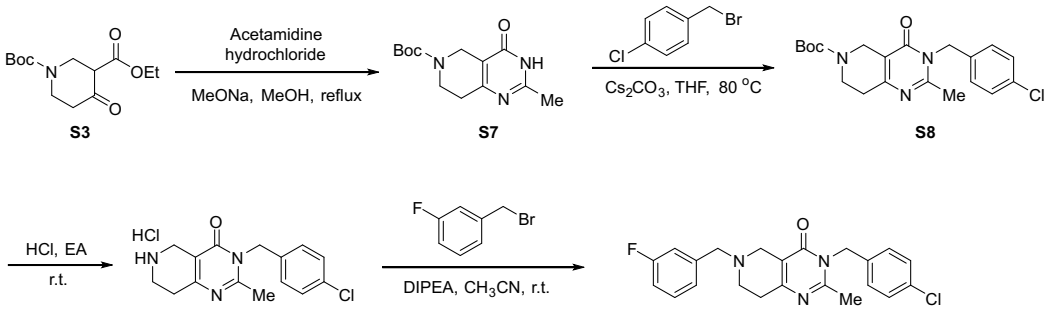
2-amino-3-(4-chlorobenzyl)-6-(4-fluorobenzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (**DC-ClpP-04**)



DC-ClpP-04

Yield = 71%. White solid. ^{1}H NMR (500 MHz, $CDCl_3$) δ 7.38 – 7.30 (m, 4H), 7.20 (d, *J* = 8.4 Hz, 2H), 7.03 (t, *J* = 8.7 Hz, 2H), 5.18 (s, 2H), 4.85 (brs, 2H), 3.69 (s, 2H), 3.44 (s, 2H), 2.70 (t, *J* = 5.8 Hz, 2H), 2.57 (t, *J* = 5.8 Hz, 2H). ^{13}C NMR (151 MHz, $CDCl_3$) δ 162.13 (d, *J* = 245.0 Hz), 161.06, 158.88, 152.95, 134.28, 133.72 (d, *J* = 3.2 Hz), 133.14, 130.67 (d, *J* = 8.0 Hz), 129.55, 128.14, 115.18 (d, *J* = 21.2 Hz), 109.47, 61.83, 50.07, 49.27, 44.31, 31.70. ^{19}F NMR (471 MHz, $CDCl_3$) δ -115.67. HRMS (ESI) calcd for $C_{21}H_{21}ClFN_4O^+$ 399.1383 $[M + H]^+$, found 399.1372.

3-(4-chlorobenzyl)-6-(3-fluorobenzyl)-2-methyl-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (**DC-ClpP-05**)



To a solution of the compound **S3** (2.7 g, 10.0 mmol) and acetamidine hydrochloride (1.4 g, 15.0 mmol) in MeOH (30 mL) was added MeONa (2.8 g, 50.0 mmol), and the mixture was stirred at 80 °C. Upon completion of the reaction, the mixture was cooled to room temperature and concentrated under reduced pressure. The reaction mixture was diluted with DCM (50 mL), and water (30 mL) was added, the aqueous

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

layer was extracted with DCM (3 × 20 mL). The combined organic layers were dried over Na_2SO_4, filtered and concentrated under reduced pressure. The resulting residue was used for the next step without further purification.

A mixture of **S7** from the above step, 1-(bromomethyl)-4-chlorobenzene (2.7 g, 13.0 mmol), Cs_2CO_3 (4.2 g, 13.0 mmol) in THF (30 mL) was stirred at 80 °C until the starting material was consumed. The reaction was quenched with water and the organic layer was separated, dried over Na_2SO_4 and concentrated under reduced pressure. The residue was purified by silica gel chromatography (PE: EA = 3/1 to 1/1) to afford intermediate **S8** as a yellow solid (1.3 g, 34% yield over two steps).

To a solution of the compound **S8** (1.3 g, 3.4 mmol) in DCM (15 mL) was added HCl (17 mL, 2.0 mol/L in EA) and the mixture was stirred at room temperature. After completion of the reaction, the solvent was removed in vacuo to afford intermediate **S9**. Then, a mixture of **S9** and DIPEA (1.8 mL, 10.2 mmol), 1-(bromomethyl)-3-fluorobenzene (1.0 g, 5.1 mmol) in MeCN (20 mL) was stirred overnight at room temperature until the starting material was consumed. Upon completion of the reaction, the mixture was diluted with H_2O (20 mL) and extracted with EA (3 × 30 mL). The combined organic layers were washed with brine, dried over Na_2SO_4, and concentrated. The residue was purified by silica gel chromatography (PE: EA = 3/1 to 1/1) to afford the final compound **DC-ClpP-05** as a white solid (0.9 g, 64% yield over two steps). ^{1}H NMR (600 MHz, $CDCl_3$) δ 7.36 – 7.29 (m, 3H), 7.19 – 7.10 (m, 4H), 6.98 (td, J = 8.4, 2.6 Hz, 1H), 5.24 (s, 2H), 3.74 (s, 2H), 3.50 (s, 2H), 2.78 – 2.72 (m, 4H), 2.44 (s, 3H). ^{13}C NMR (151 MHz, $CDCl_3$) δ 163.02 (d, J = 245.7 Hz), 161.04, 157.56, 156.33, 140.66 (d, J = 7.2 Hz), 133.95, 133.74, 129.83 (d, J = 8.2 Hz), 129.14, 128.23, 124.57, 117.59, 115.76 (d, J = 21.4 Hz), 114.22 (d, J = 21.2 Hz), 61.96, 49.95, 49.28, 46.47, 31.66, 22.91. ^{19}F NMR (471 MHz, $CDCl_3$) δ -113.51. HRMS (ESI) calcd for $C_{22}H_{22}ClFN_3O^+$ 398.1430 $[M + H]^+$, found 398.1420.

4.7 Measurement of ClpP agonistic activity

4.7.1 Protein Preparation for In Vitro Enzyme Assay

ClpP protein was prepared using a protocol adapted from Ishizawa et al.[62]. ClpP sequence 57−277 with an *N*-terminal SUMO tag was cloned and transformed in Escherichia coli. The bacteria were cultured with LB medium in the presence of 50 μg/mL ampicillin at 37 °C. The bacteria were harvested after treatment with 0.5 mM isopropyl β D-1 thiogalactopyranoside for 4 h and sequentially lysed in buffer (25 mM Tris-HCl pH 7.5, 0.5 M NaCl, 5 mM β-mercaptoethanol, 10 mM imidazole, 10% glycerol, 1% Triton X-100) by sonication. The pellet was removed by centrifugation, the supernatant was combined with Ni sepharose 6 Fast Flow for 2 h, and then purified using a gravity column, first with washing buffer (25 mM Tris-HCl pH 7.5, 0.5 M NaCl, 5 mM β-mercaptoethanol, 10 mM imidazole, 10% glycerol) for the first

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

impurity protein elution and finally with 400 mM imidazole for the protein of interest. The eluted protein was dialyzed against 4 L of dialysate solution (25 mM Tris-HCl pH 7.5, 0.5 M NaCl, 5 mM β-mercaptoethanol, 10% glycerol) for 6 h. Subsequently, the His-Sumo tag fused to ClpP was cleaved via overnight incubation with Ulp1 protease at 16 °C. Protein concentration was determined using the BCA method, and the purified protein was finally preserved at −80 °C in storage buffer supplemented with 50% glycerol.

4.7.2 In Vitro Enzyme Assay

The fluorogenic ClpP assay was adapted from Ishizawa et al.[62]. ClpP and Ac-WLA-AMC were diluted in assay buffer (50 mM Tris-HCl, pH 8, 10 mM $MgCl_2$, 100 mM KCl, 1 mM DTT, 5% glycerol and 0.02% Triton X-100) so that their final concentrations in the completed 25 µL reaction were 0.07 µM and 100 µM, respectively. Compound stocks (10 mM) were serially diluted fivefold in DMSO to prepare seven concentrations and then diluted 100-fold in ddH_2O. Each well of a black 384-well plate received 10 µL of ClpP solution and 5 µL of compound (final concentration 20 µM, 4 µM, 0.8 µM, 0.16 µM, 32 nM, 6.4 nM, 1.28 nM) solution. Vehicle-control wells contained 10 µL of ClpP solution, 10 µL of substrate solution and 5 µL of 1% DMSO; blank wells contained 10 µL of assay buffer, 10 µL of substrate solution and 5 µL of 1% DMSO. The final DMSO concentration was 0.2%. After a 10-min incubation at room temperature, 10 µL of Ac-WLA-AMC solution was added to the compound-treated wells, giving a total reaction volume of 25 µL. Initial fluorescence was measured using a microplate reader with excitation at 355 nm and emission at 460 nm (PerkinElmer, EnVision® 2105 Multimode Plate). The plate was then sealed and incubated at 37 °C for 6 h before fluorescence was measured again. Fold activation was calculated using the expression below, and EC_{50} values were obtained by fitting a variable-slope four-parameter concentration-response model in GraphPad Prism 8.3.0. Data are mean ± s.d. from four technical replicates.

$$\text{Fold increase} = \frac{\left[\left(\mathrm{FI}_{2,\text{compound}} - \mathrm{FI}_{1,\text{compound}}\right) - \left(\mathrm{MFI}_{2,\text{blank}} - \mathrm{MFI}_{1,\text{blank}}\right)\right]/h}{\left[\left(\mathrm{MFI}_{2,\text{DMSO}} - \mathrm{MFI}_{1,\text{DMSO}}\right) - \left(\mathrm{MFI}_{2,\text{blank}} - \mathrm{MFI}_{1,\text{blank}}\right)\right]/h} \quad (20)$$

Data Availability

The datasets used in this study are available via Zenodo at https://zenodo.org/records/22116036.

Code Availability

The source code of this study is provided at https://github.com/HShokaku/LeadFlow.

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Author Contributions

M.Y.Z., D.H.C., and Z.H.F. designed the research study. Z.H.F., J.Y., D.T. and K.Y.Q. developed the methodology and algorithms. Z.H.F., J.Y., K.Y.Q., X.S., X.Y.J., L.H.Z., X.Y.W., M.A.C., R.Z.Z., Y.Y.Z., and D.T. conducted the experiments, curated the data, and performed the data analysis. Q.H. and B.J.C. conducted compound synthesis and characterization. Z.H.F., M.Y.Z., D.H.C., D.T. and K.Y.Q. prepared the figures and wrote the manuscript. M.Y.Z., D.H.C., X.M.L., M.L.W., H.Z., and J.L. supervised the project and provided scientific guidance and resources. All authors read and approved the final manuscript.

Competing Interests

The authors declare no competing interests.

Acknowledgements

We gratefully acknowledge financial support from National Natural Science Foundation of China (T2225002, 82404513 and 82273855), Strategic Priority Research Program of the Chinese Academy of Sciences (XDB0830000 and XDA0580000), National Key Research and Development Program of China (2022YFC3400504), Innovative Drug Research and Development National Science and Technology Major Project (2025ZD1800700). Lingang Laboratory (LGL-8888).

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

References

1. Krishnan, A.*, et al.* A generative deep learning approach to *de novo* antibiotic design. *Cell* **188**, 5962–5979.e5922 (2025).

2. Zhang, K.*, et al.* Artificial intelligence in drug development. *Nature Medicine* **31**, 45–59 (2025).

3. Xu, Z.*, et al.* A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomized phase 2a trial. *Nature Medicine* **31**, 2602–2610 (2025).

4. Watson, J.L.*, et al.* De novo design of protein structure and function with RFdiffusion. *Nature* **620**, 1089–1100 (2023).

5. Abramson, J.*, et al.* Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **630**, 493–500 (2024).

6. Liu, H.*, et al.* How Good are Current Pocket-Based 3D Generative Models?: The Benchmark Set and Evaluation of Protein Pocket-Based 3D Molecular Generative Models. *Journal of Chemical Information and Modeling* **64**, 9260–9275 (2024).

7. Cao, D.*, et al.* Evaluating generative models for structure-based drug discovery from de novo design to lead optimization. *Nature Communications* (2026).

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

9. Yang, B., Xiang, C., Li, T., Xu, Y. & Li, J. Structure-Based Generation of 3D Small-Molecule Drugs: Are We There Yet? *Journal of Medicinal Chemistry* **68**, 22756–22768 (2025).

10. Zhang, O.*, et al.* Deep Lead Optimization: Leveraging Generative AI for Structural Modification. *Journal of the American Chemical Society* **146**, 31357–31370 (2024).

11. Langevin, M., Minoux, H., Levesque, M. & Bianciotto, M. Scaffold-Constrained Molecular Generation. *Journal of Chemical Information and Modeling* **60**, 5637–5646 (2020).

12. Liu, X., Ye, K., van Vlijmen, H.W.T., Ijzerman, A.P. & van Westen, G.J.P. DrugEx v3: scaffold-constrained drug design with graph transformer-based reinforcement learning. *Journal of Cheminformatics* **15**, 24 (2023).

13. Adams, K., Abeywardane, K., Fromer, J.C. & Coley, C.W. ShEPhERD: Diffusing shape, electrostatics, and pharmacophores for bioisosteric drug design. in *ICLR* (2025).

14. Zhu, H., Zhou, R., Cao, D., Tang, J. & Li, M. A pharmacophore-guided deep learning approach for bioactive molecular generation. *Nature Communications* **14**, 6234 (2023).

15. Xie, W.*, et al.* Accelerating discovery of bioactive ligands with pharmacophore-informed generative models. *Nature Communications* **16**, 2391 (2025).

16. Chen, S.*, et al.* Deep lead optimization enveloped in protein pocket and its application in

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

designing potent and selective ligands targeting LTK protein. *Nature Machine Intelligence* **7**, 448–458 (2025).

17.Xie, J., Chen, S., Lei, J. & Yang, Y. DiffDec: structure-aware scaffold decoration with an end-to-end diffusion model. *Journal of Chemical Information and Modeling* **64**, 2554–2564 (2024).

18.Schneuing, A.*, et al.* Structure-based drug design with equivariant diffusion models. *Nature Computational Science* **4**, 899–909 (2024).

19.Lin, H.*, et al.* CBGBench: Fill in the blank of protein-molecule complex binding graph. *arXiv preprint arXiv:2406.10840* (2024).

20.Francoeur, P.G.*, et al.* Three-Dimensional Convolutional Neural Networks and a Cross-Docked Data Set for Structure-Based Drug Design. *Journal of Chemical Information and Modeling* **60**, 4200–4215 (2020).

21.Lewell, X.Q., Judd, D.B., Watson, S.P. & Hann, M.M. Recap retrosynthetic combinatorial analysis procedure: a powerful new technique for identifying privileged molecular fragments with useful applications in combinatorial chemistry. *Journal of chemical information and computer sciences* **38**, 511–522 (1998).

22.Arús-Pous, J.*, et al.* SMILES-based deep generative scaffold decorator for de-novo drug design. *Journal of cheminformatics* **12**, 38 (2020).

23.Degen, J., Wegscheid-Gerlach, C., Zaliani, A. & Rarey, M. On the art of compiling and using'drug-like'chemical fragment spaces. *ChemMedChem* **3**, 1503 (2008).

24.Cremer, J.*, et al.* FLOWR: flow matching for structure-aware de novo, interaction- and fragment-based ligand generation. *Nature Computational Science* **6**, 565–574 (2026).

25.Peng, X.*, et al.* Unified modeling of 3D molecular generation via atomic interactions with PocketXMol. *Cell* **189**, 1904–1922.e1928 (2026).

26.Kong, X.*, et al.* Programming Biomolecular Interactions with All-Atom Generative Model. *bioRxiv*, 2026.2003.2012.711044 (2026).

27.Qiao, A.*, et al.* A 3D pocket-aware lead optimization model with knowledge guidance and its application for discovery of new glutaminyl cyclase inhibitors. *Briefings in Bioinformatics* **26**, bbaf345 (2025).

28.Gu, S.*, et al.* Aligning target-aware molecule diffusion models with exact energy optimization. *Advances in Neural Information Processing Systems* **37**, 44040–44063 (2024).

29.Qiu, K.*, et al.* Empower structure-based molecule optimization with gradient guided bayesian flow networks. *arXiv preprint arXiv:2411.13280* (2024).

30.Black, K., Janner, M., Du, Y., Kostrikov, I. & Levine, S. Training diffusion models with reinforcement learning. in *International Conference on Learning Representations*, Vol. 2024 4965–4987 (2024).

31.Stacey, J., Canault, B., Pickett, S.D. & Gillet, V.J. Visualising lead optimisation series

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

using reduced graphs. *Journal of Cheminformatics* **17**, 60 (2025).

32. Takeuchi, K., Kunimoto, R. & Bajorath, J. Systematic mapping of R-group space enables the generation of an R-group replacement system for medicinal chemistry. *European Journal of Medicinal Chemistry* **225**, 113771 (2021).

33. Hoffer, L., Muller, C., Roche, P. & Morelli, X. Chemistry-driven Hit-to-lead Optimization Guided by Structure-based Approaches. *Molecular Informatics* **37**, 1800059 (2018).

34. Yu, J.*, et al.* Atomic-level protein–ligand recognition with PBCNet2.0 for probe discovery. *Nature Chemical Biology* **22**, 1299–1310 (2026).

35. Yu, J.*, et al.* Computing the relative binding affinity of ligands based on a pairwise binding comparison network. *Nature computational science* **3**, 860–872 (2023).

36. Graves, A., Srivastava, R.K., Atkinson, T. & Gomez, F. Bayesian flow networks. *arXiv preprint arXiv:2308.07037* (2023).

37. Wallace, B.*, et al.* Diffusion model alignment using direct preference optimization. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 8228–8238 (2024).

38. Rafailov, R.*, et al.* Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems* **36**, 53728–53741 (2023).

39. Song, Y.*, et al.* Unified generative modeling of 3d molecules with bayesian flow networks. in *The Twelfth International Conference on Learning Representations* (2024).

40. Qiu, K.*, et al.* Piloting structure-based drug design via modality-specific optimal schedule. *arXiv preprint arXiv:2505.07286* (2025).

41. Qu, Y.*, et al.* MolCRAFT: structure-based drug design in continuous parameter space. *arXiv preprint arXiv:2404.12141* (2024).

42. Atkinson, T.*, et al.* Protein sequence modelling with Bayesian flow networks. *Nature Communications* **16**, 3197 (2025).

43. Ruple, L., Torresi, L., Schopmans, H. & Friederich, P. Symmetry-aware Bayesian flow networks for crystal generation. *npj Computational Materials* **12**, 182 (2026).

44. Zeng, C.*, et al.* Structure-free generation of transition state conformations with Bayesian flow networks. *ChemRxiv* **2026**(2026).

45. Jiao, R.*, et al.* Mof-bfn: Metal-organic frameworks structure prediction via bayesian flow networks. *Advances in Neural Information Processing Systems* **38**, 161989–162015 (2026).

46. Kadan, A., Ryczko, K., Lloyd, E., Roitberg, A. & Yamazaki, T. Guided multi-objective generative AI to enhance structure-based drug design. *Chemical Science* **16**, 13196–13210 (2025).

47. Schindler, C.E.*, et al.* Large-scale assessment of binding free energy calculations in active drug discovery projects. *Journal of Chemical Information and Modeling* **60**, 5457–5474 (2020).

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

48. Wang, L., *et al.* Accurate and Reliable Prediction of Relative Ligand Binding Potency in Prospective Drug Discovery by Way of a Modern Free-Energy Calculation Protocol and Force Field. *Journal of the American Chemical Society* **137**, 2695–2703 (2015).

49. Chen, Z., Peng, B., Parthasarathy, S. & Ning, X. Shape-conditioned 3D molecule generation via equivariant diffusion models. *arXiv preprint arXiv:2308.11890* (2023).

50. Harris, C., *et al.* Posecheck: Generative models for 3d structure-based drug design produce unrealistic poses. in *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop* (2023).

51. Buttenschoen, M., Morris, G.M. & Deane, C.M. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical Science* **15**, 3130–3139 (2024).

52. Trott, O. & Olson, A.J. AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of computational chemistry* **31**, 455–461 (2010).

53. Bouysset, C. & Fiorucci, S. ProLIF: a library to encode molecular interactions as fingerprints. *Journal of cheminformatics* **13**, 72 (2021).

54. Bothe, U., *et al.* Discovery of IRAK4 Inhibitors BAY1834845 (Zabedosertib) and BAY1830839. *Journal of Medicinal Chemistry* **67**, 1225–1242 (2024).

55. Ngo, V.N., *et al.* Oncogenically active MYD88 mutations in human lymphoma. *Nature* **470**, 115–119 (2011).

56. Bai, Y.-R., *et al.* The recent advance of Interleukin-1 receptor associated kinase 4 inhibitors for the treatment of inflammation and related diseases. *European Journal of Medicinal Chemistry* **258**, 115606 (2023).

57. Chen, Y., *et al.* Bicyclic pyrimidine compounds as potent IRAK4 inhibitors. *Bioorganic & Medicinal Chemistry Letters* **73**, 128900 (2022).

58. Haynes, C.M., Petrova, K., Benedetti, C., Yang, Y. & Ron, D. ClpP Mediates Activation of a Mitochondrial Unfolded Protein Response in *C. elegans*. *Developmental Cell* **13**, 467–480 (2007).

59. Annesley, S.J. & Fisher, P.R. Mitochondria in Health and Disease. *Cells* **8**, 680 (2019).

60. Topf, U., Wrobel, L. & Chacinska, A. Chatty Mitochondria: Keeping Balance in Cellular Protein Homeostasis. *Trends in Cell Biology* **26**, 577–586 (2016).

61. Vyas, S., Zaganjor, E. & Haigis, M.C. Mitochondria and Cancer. *Cell* **166**, 555–566 (2016).

62. Ishizawa, J., *et al.* Mitochondrial ClpP-Mediated Proteolysis Induces Selective Cancer Cell Lethality. *Cancer Cell* **35**, 721–737.e729 (2019).

63. Suomalainen, A. & Battersby, B.J. Mitochondrial diseases: the contribution of organelle stress responses to pathology. *Nature Reviews Molecular Cell Biology* **19**, 77–92 (2018).

64. Cole, A., *et al.* Inhibition of the Mitochondrial Protease ClpP as a Therapeutic Strategy for

Langtaosha (LTS) Preprint doi: https://doi.org/10.65215/LTSpreprints.2026.10.29.000370 This version posted October 29, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder. Creative Commons license: CC Attribution-NonCommercial-NoDerivatives 4.0 https://creativecommons.org/licenses/by-nc-nd/4.0

Human Acute Myeloid Leukemia. *Cancer Cell* **27**, 864–876 (2015).

65.Uhlen, M.*, et al.* A pathology atlas of the human cancer transcriptome. *Science* **357**, eaan2507 (2017).

66.Mirali, S. & Schimmer, A.D. The role of mitochondrial proteases in leukemic cells and leukemic stem cells. *Stem Cells Translational Medicine* **9**, 1481–1487 (2020).

67.Chen, S.*, et al.* Design, Synthesis, and Structure-Activity Relationships of Novel Piperidine-Fused Imidazolone ClpP Activators as Potential Anti-Cancer Agents. *Molecules* **31**, 2916 (2026).

68.Kumar, R.*, et al.* Structure and Energetics of Ligand–Fluorine Interactions with Galectin-3 Backbone and Side-Chain Amides: Insight into Solvation Effects and Multipolar Interactions. *ChemMedChem* **14**, 1528–1536 (2019).

69.Mabanglo, M.F.*, et al.* Potent ClpP agonists with anticancer properties bind with improved structural complementarity and alter the mitochondrial N-terminome. *Structure* **31**, 185–200.e110 (2023).

70.Zdrazil, B.*, et al.* The ChEMBL Database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. *Nucleic acids research* **52**, D1180–D1192 (2024).

71.Ouyang, L.*, et al.* Training language models to follow instructions with human feedback. *Advances in neural information processing systems* **35**, 27730–27744 (2022).

72.Bradley, R.A. & Terry, M.E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika* **39**, 324–345 (1952).

73.Zhou, X.*, et al.* Antigen-specific antibody design via direct energy-based preference optimization. *Advances in Neural Information Processing Systems* **37**, 120861–120891 (2024).

74.Lee, S., Jo, M., Ok, J. & Kim, D. Enhancing Ligand Validity and Affinity in Structure-Based Drug Design with Multi-Reward Optimization. in *Forty-second International Conference on Machine Learning* (2025).

75.Noutahi, E.*, et al.* datamol-io/medchem: Molecular filtering for drug discovery. (Zenodo, 2025).

76.Baell, J.B. & Holloway, G.A. New Substructure Filters for Removal of Pan Assay Interference Compounds (PAINS) from Screening Libraries and for Their Exclusion in Bioassays. *Journal of Medicinal Chemistry* **53**, 2719–2740 (2010).

77.Pearce, B.C., Sofia, M.J., Good, A.C., Drexler, D.M. & Stock, D.A. An Empirical Process for the Design of High-Throughput Screening Deck Filters. *Journal of Chemical Information and Modeling* **46**, 1060–1068 (2006).

LeadFlow learns activity-aligned molecular generation from chemical series for structure-based molecular optimization

Zhehuan Fan[1,2,†], Dan Teng[1†], Jie Yu[1,3,4†], Qi Huang[5,†], Keyue Qiu[6,†], Beijing Chen[5], Xia Sheng[1,2], Xinyu Jiang[1,2], Lehan Zhang[1,2], Xingyou Wang[1,2], Yaoyu Zheng[2,7], Mingan Chen[8], Runze Zhang[9], Jia Li[5,10], Hao Zhou[6,*], Mingliang Wang[2,5,11,*], Xiaomin Luo[1,2,*], Duanhua Cao[12,*], Mingyue Zheng[1,2,3,7,*]

[1]Drug Discovery and Design Center, State Key Laboratory of Drug Research, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai 201203, China

[2]University of Chinese Academy of Sciences, Beijing 100049, China

[3]School of Information Science and Technology, ShanghaiTech University, Shanghai, 201210, China

[4]Lingang Laboratory, Shanghai 200031, China

[5]Zhongshan Institute for Drug Discovery, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Zhongshan 528400, China

[6]Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China

[7]School of Pharmaceutical Science and Technology, Hangzhou Institute for Advanced Study, UCAS, Hangzhou, 330106, China

[8]iHuman institute, ShanghaiTech University, Shanghai, 201210, China

[9]Science for Life Laboratory, Department of Cell and Molecular Biology, Uppsala University, BMC, Uppsala, Sweden

[10]The State Key Laboratory of Chemical Biology, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai 201203, China

[11]Department of Medicinal Chemistry, Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai 201203, China

[12]State Key Laboratory of Cardiovascular Diseases and Medical Innovation Center, Medical Artificial Intelligence Innovation Center, Shanghai East Hospital, Frontier Science Center for Stem Cell Research, School of Life Sciences and Technology, Tongji University, Shanghai, 200092, China

Corresponding Authors

*(Mingyue Zheng) E-mail: myzheng@simm.ac.cn

*(Duanhua Cao) E-mail: caodh@tongji.edu.cn

*(Xiaomin Luo) E-mail: xmluo@simm.ac.cn

*(Mingliang Wang) E-mail: wangmingliang@simm.ac.cn

*(Hao Zhou) E-mail: zhouhao@air.tsinghua.edu.cn

†These authors contributed equally to this work.

Supplementary Information

Table of Contents

Supplementary Results

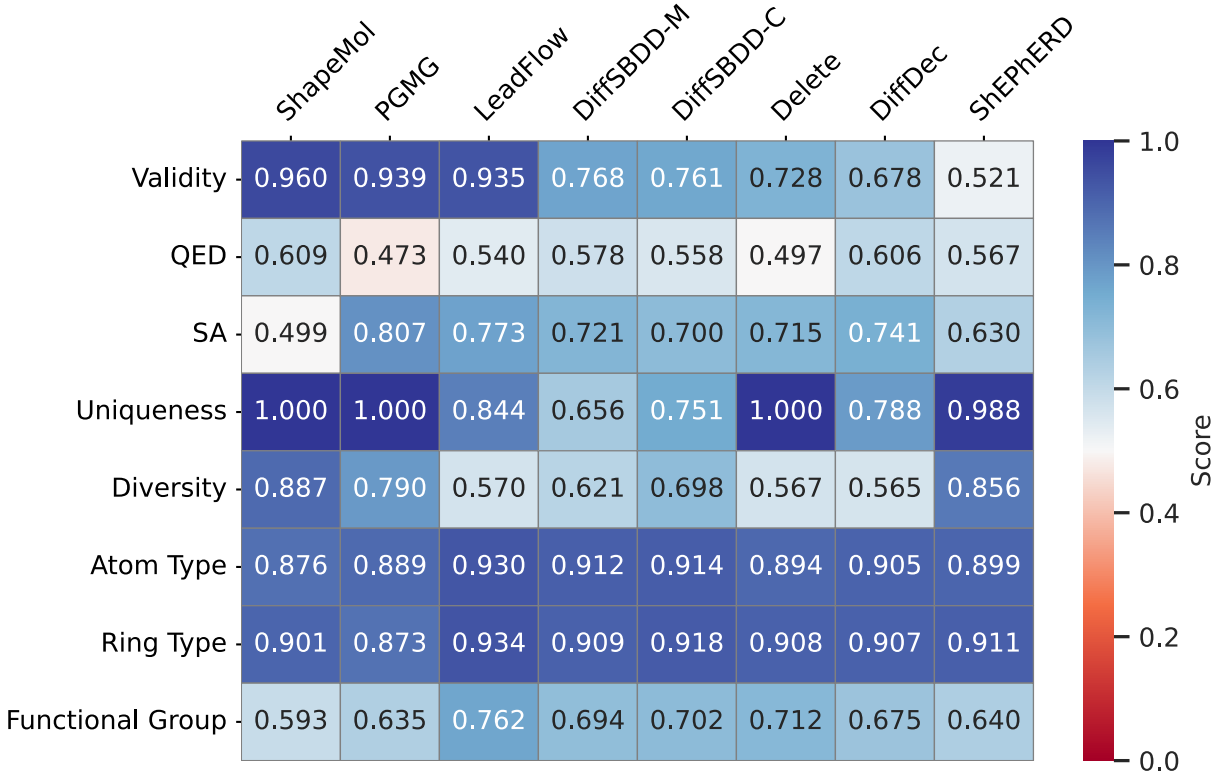
Implementation Details

Architecture of SE(3)-Graph Transformer

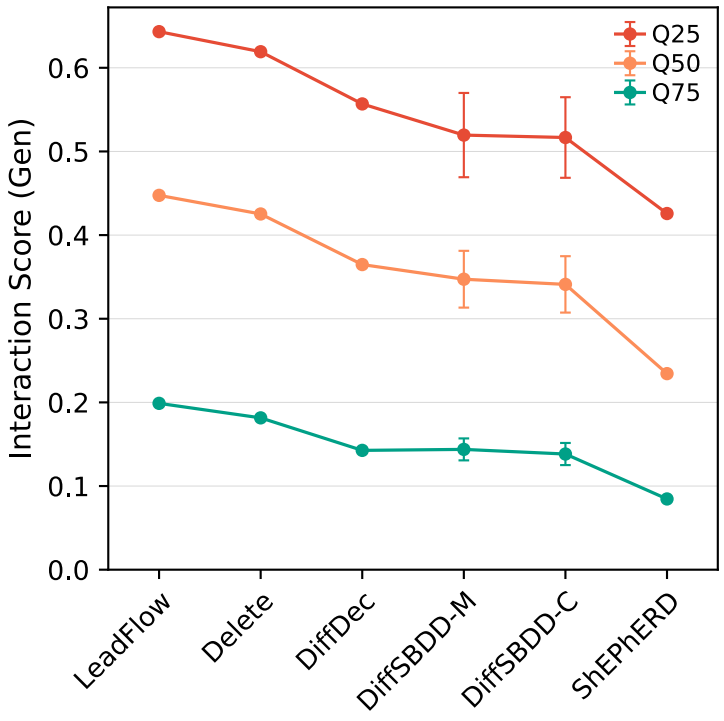
1H/ 13C/ 19F NMR, HRMS spectra and HPLC spectra of compounds

ClpP activator candidates

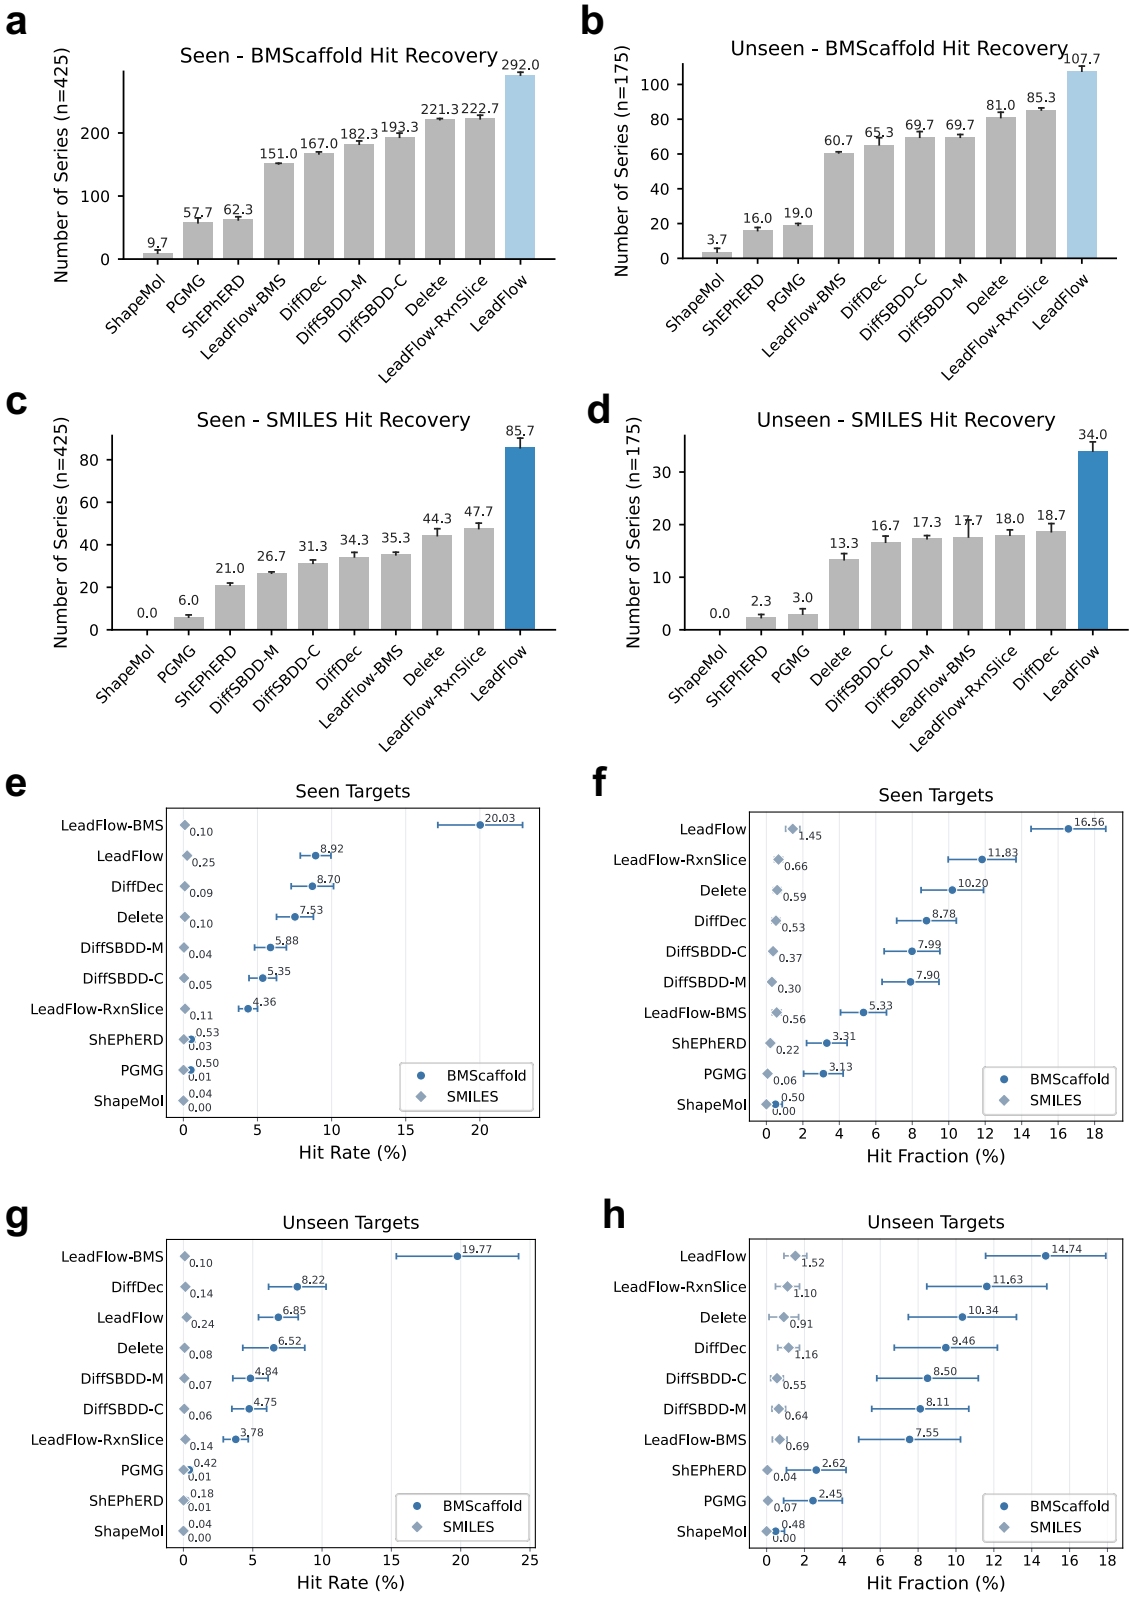
Supplementary Results



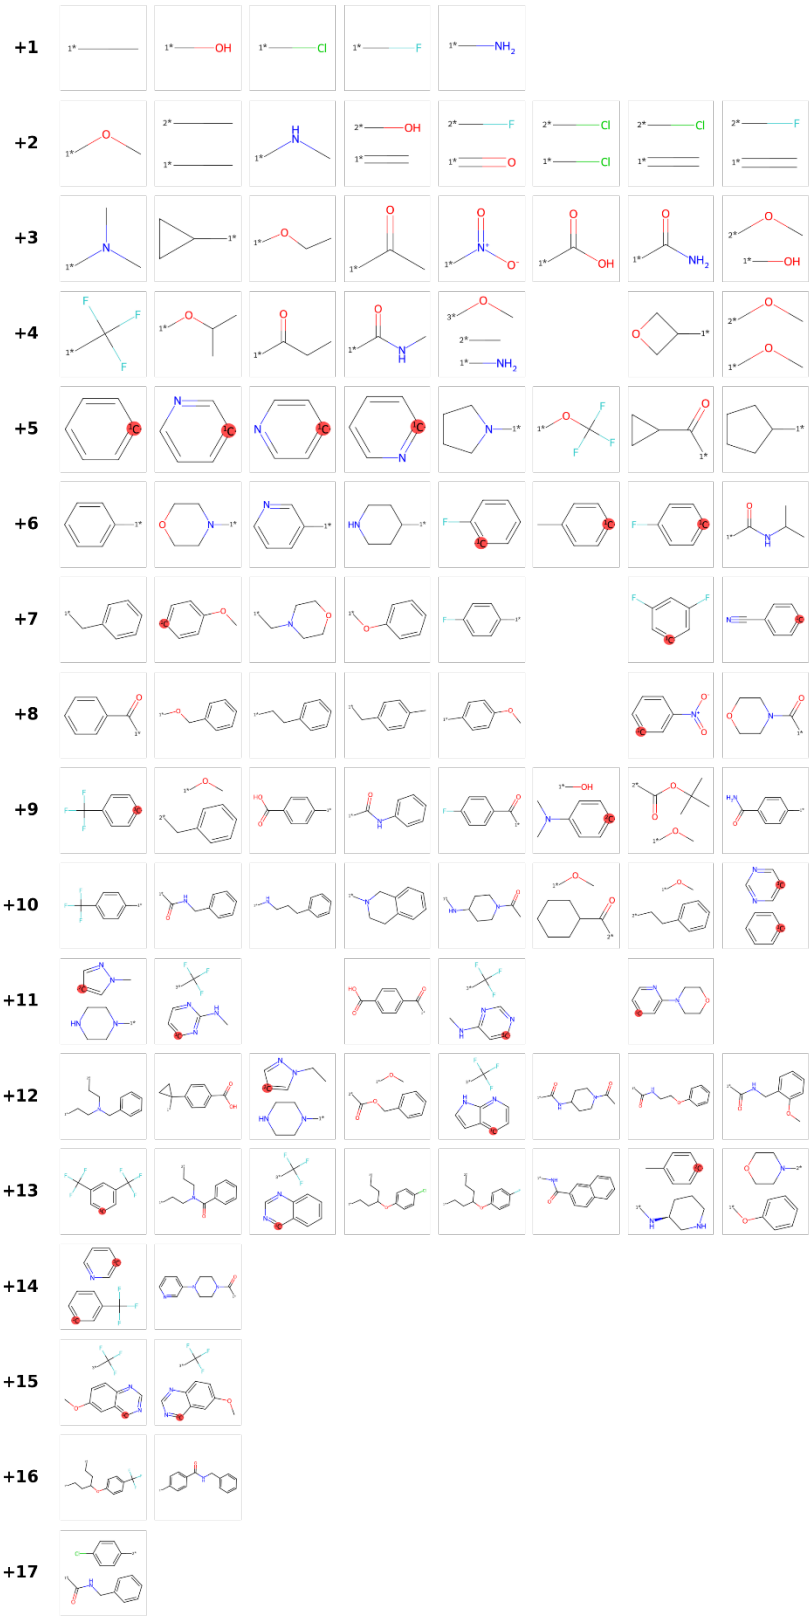
Supplementary Fig. 1 | Evaluation results of basic molecular properties. Metrics include validity, QED, SA score, uniqueness, diversity, and motif distributions (atom types, ring types, functional groups) with reference actives. Values represent the mean from three sampling runs with different seeds. All metrics are normalized to a 0–1 scale, where higher values indicate better performance.



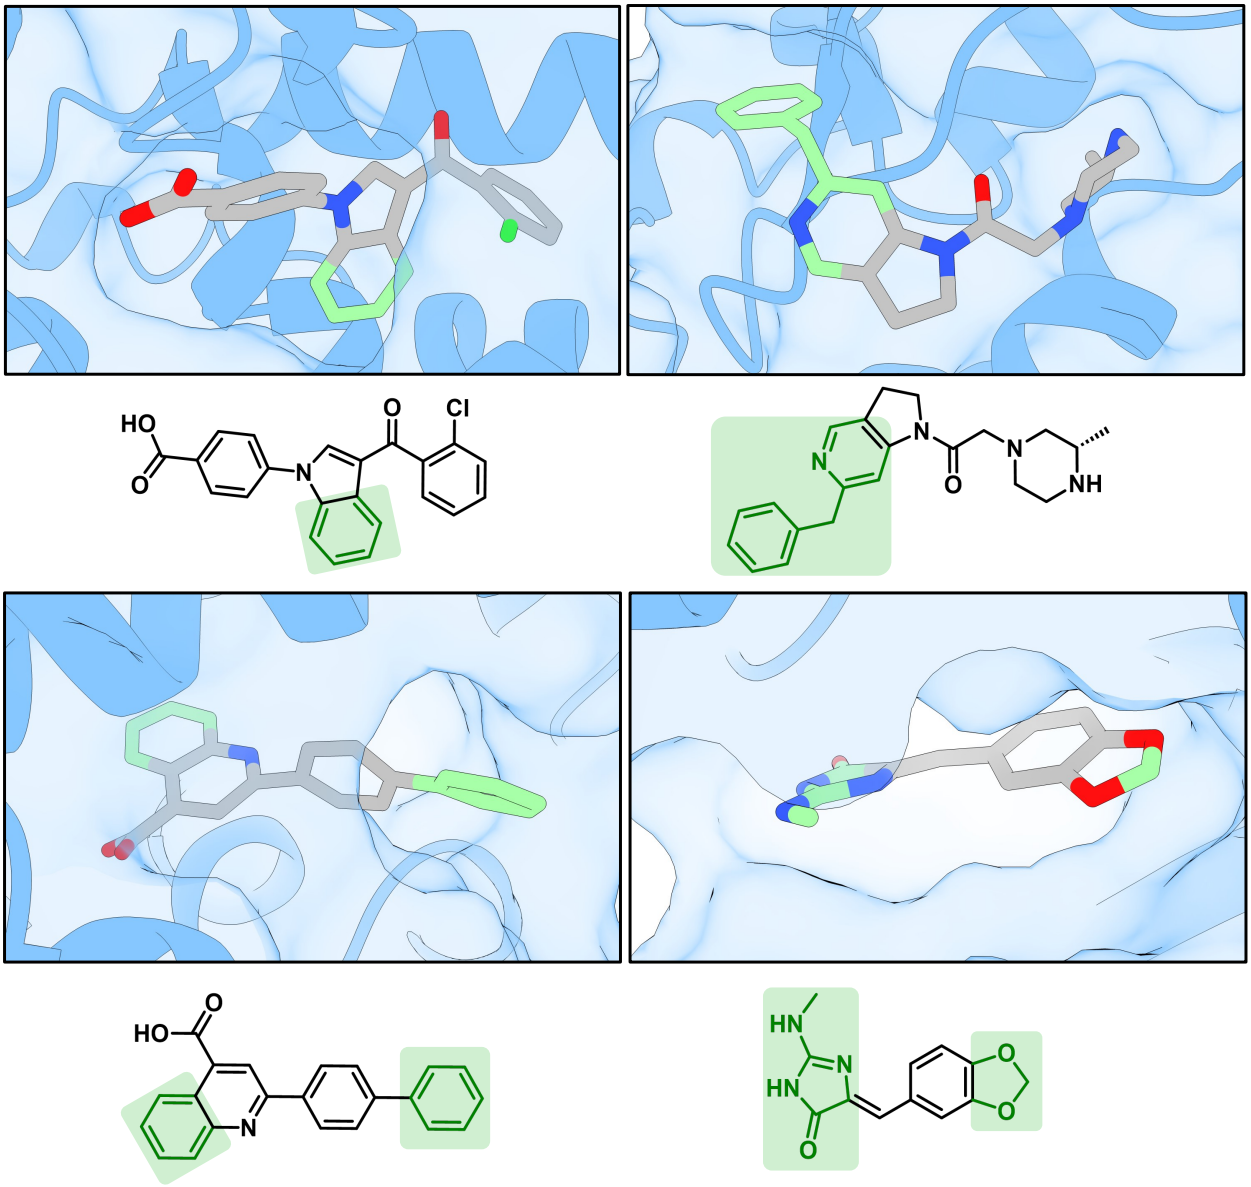
Supplementary Fig. 2 | Proportion of directly generated poses whose interaction scores exceed the 25th, 50th, and 75th percentiles of the reference active distribution. Data are presented as mean ± s.d. (n = 3 sampling runs with different seeds from fixed trained checkpoints)


a
Seen - BMScaffold Hit Recovery
Number of Series (n=425)
ShapeMol
PGMG
ShEPhERD
LeadFlow-BMS
DiffDec
DiffSBDD-M
DiffSBDD-C
Delete
LeadFlow-RxnSlice
LeadFlow
b
Unseen - BMScaffold Hit Recovery
Number of Series (n=175)
c
Seen - SMILES Hit Recovery
Number of Series (n=425)
d
Unseen - SMILES Hit Recovery
Number of Series (n=175)
e
Seen Targets
BMScaffold
SMILES
Hit Rate (%)
f
Seen Targets
Hit Fraction (%)
g
Unseen Targets
Hit Rate (%)
h
Unseen Targets
Hit Fraction (%)

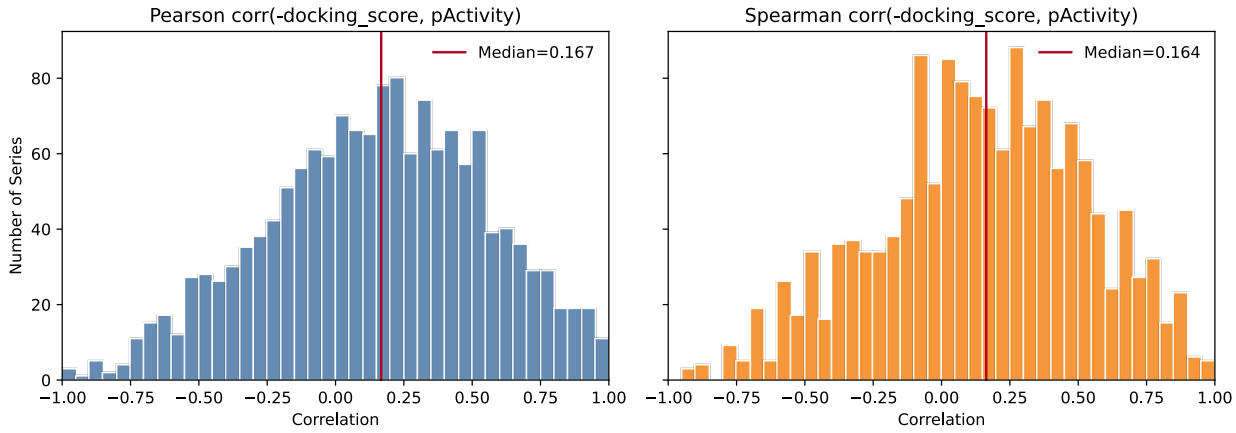
Supplementary Fig. 3 | Recovery performance stratified by the MolGenBench seen/unseen protein subsets. a, b. BMScaffold-level Hit Recovery counts on the seen (**a**) and unseen (**b**) protein subsets of MolGenBench. **c, d.** SMILES-level Hit Recovery counts on the seen (**c**) and unseen (**d**) protein subsets of MolGenBench; data are presented as mean values ± s.d. **e, f**. Recovery performance on the seen subset (n = 425 series) summarized as Hit Rate (**e**) and Hit Fraction (**f**), reported separately for BMScaffold and exact SMILES recovery. **g, h.** Corresponding Hit Rate (**g**) and Hit Fraction (**h**) results on the unseen subset (n = 175 series). Points show means across series after averaging three sampling runs; error bars indicate series-level 95% CIs calculated using Student's t-distribution. The subsets denote CrossDocked training exposure. All benchmark test targets were excluded from LeadFlow training by UniProt ID.


+1
+2
+3
+4
+5
+6
+7
+8
+9
+10
+11
+12
+13
+14
+15
+16
+17

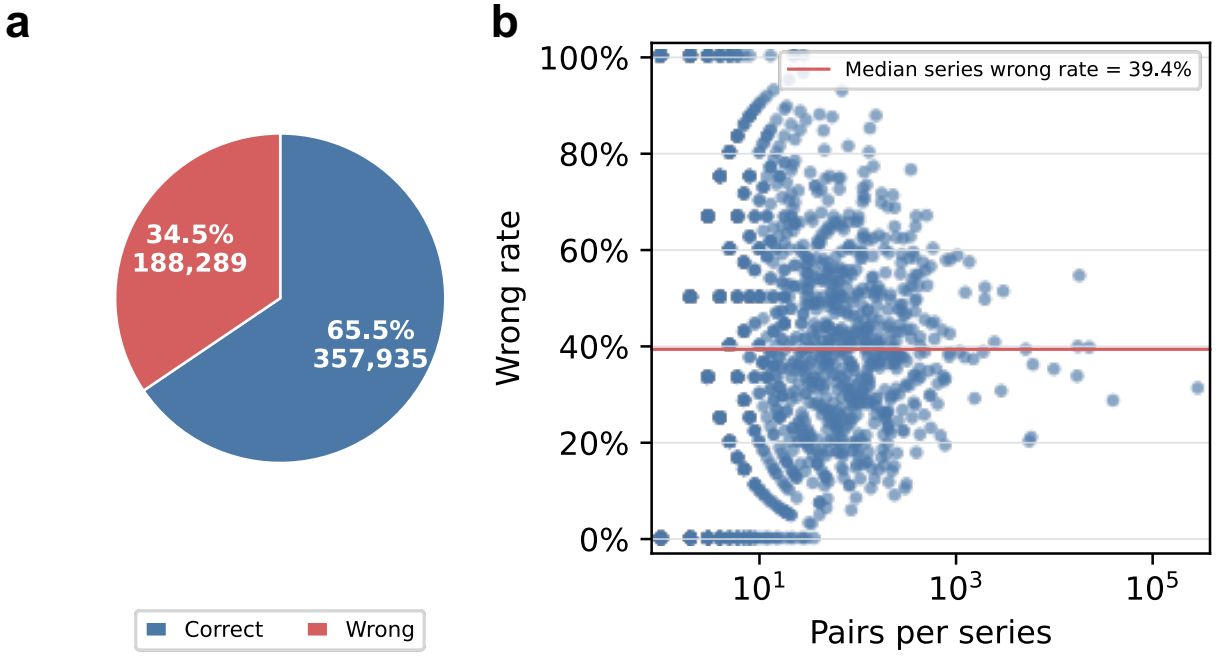
Supplementary Fig. 4 | Representative R-group fragments stratified by recovered atom count for LeadFlow. The number at the start of each row indicates how many more heavy atoms the molecule has than the scaffold. Atoms marked with * are the attachment points to the scaffold; a red "C" indicates the group is fused to the scaffold, with the red "C" being the shared atom of the fused ring. Multiple disconnected fragments in one cell mean the molecule has different groups attached at multiple positions of the scaffold simultaneously.



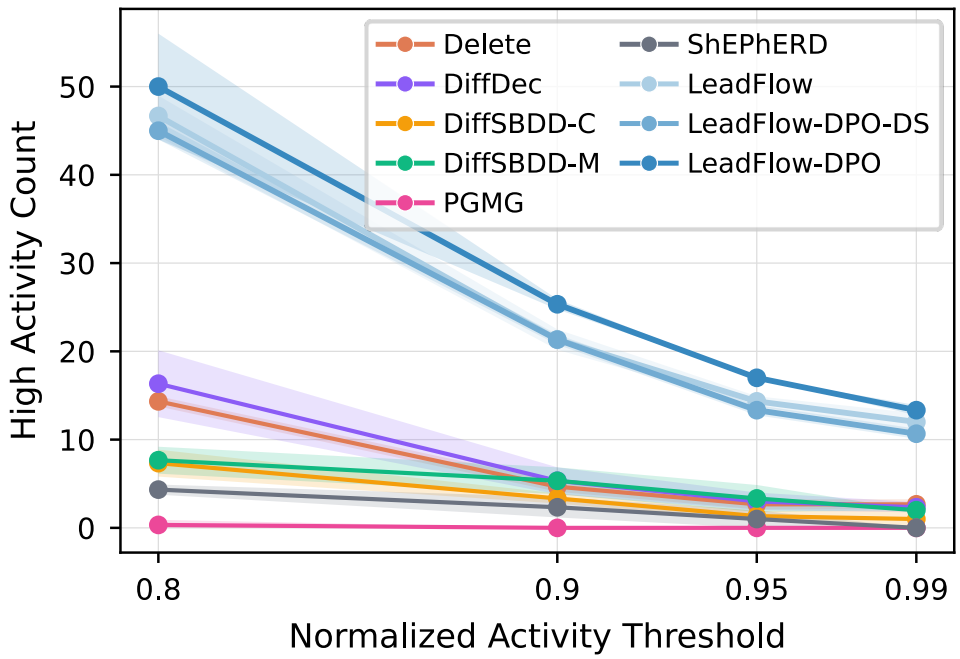
Supplementary Fig. 5 | Representative examples of experimentally active molecules recovered by LeadFlow through fused-ring elaboration.



Supplementary Fig. 6 | Distribution of correlations between docking scores and experimentally measured activities within each chemical series. Docking scores were obtained using Glide and sign-inverted so that larger values indicate more favorable predicted binding. Experimental activities were converted to pActivity values, such as pIC_{50}, so that larger values indicate higher potency. Red vertical lines indicate the median values.



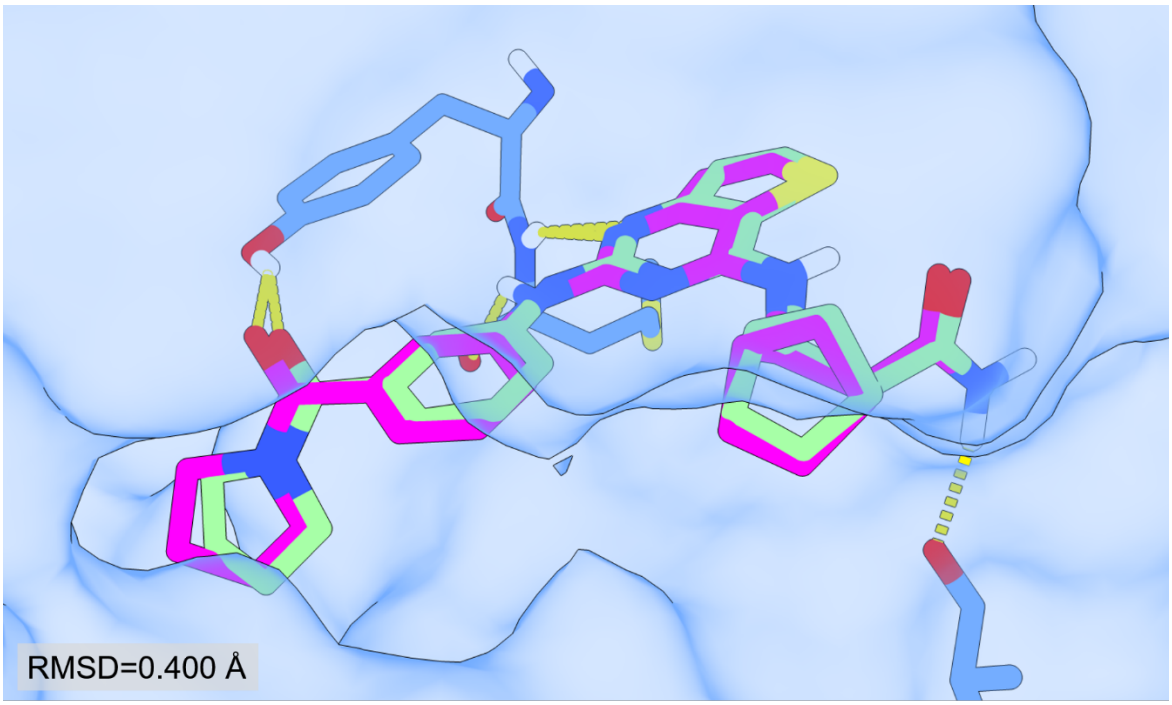
Supplementary Fig. 7 | Correctness of docking score-based reassignment of experimentally defined DPO pairs. a. The fraction of non-tied pairs for which docking-score-based labels agreed or disagreed with the experimental activity labels; seven tied pairs were excluded from this comparison. **b.** Chemical series level distribution of docking-based wrong rates, stratified by the number of pairs per series**.** The red horizontal line indicates the median per-series wrong rate.



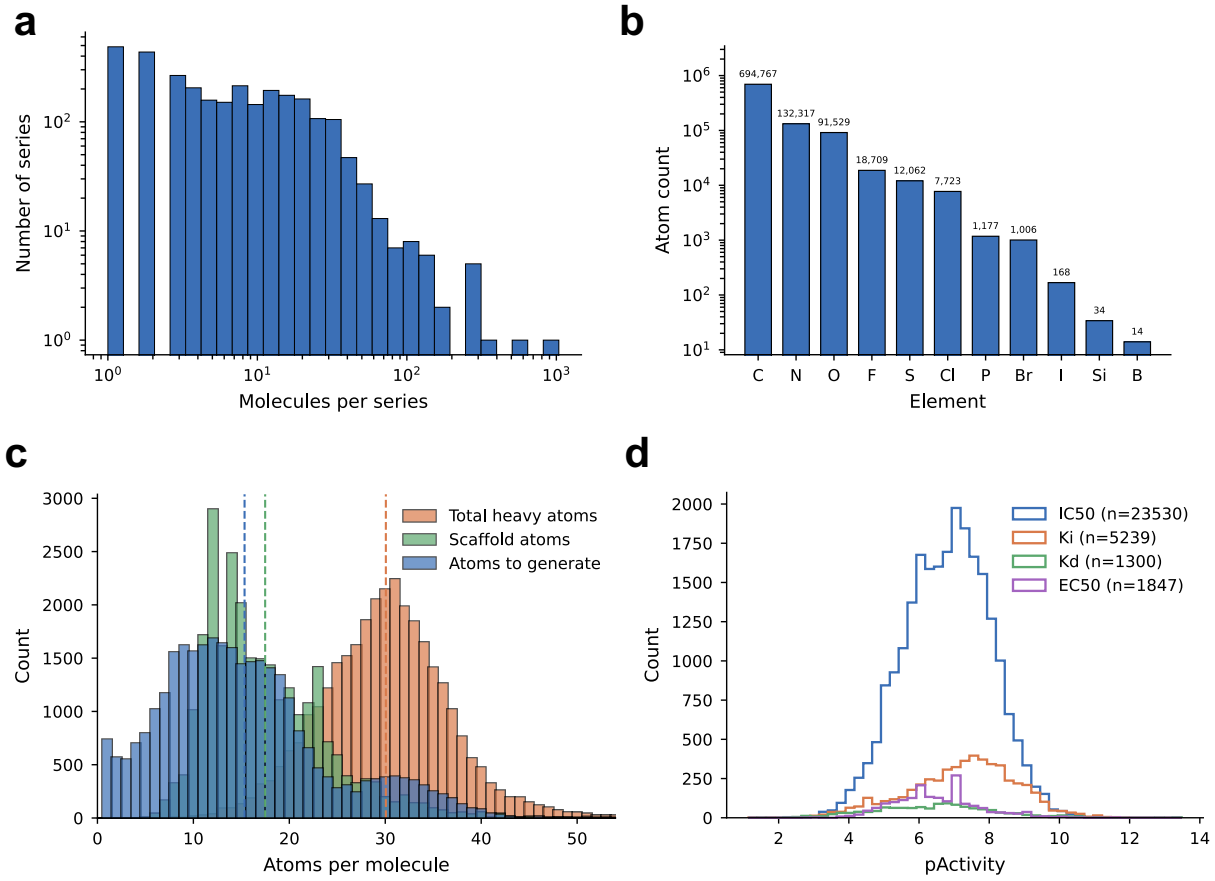
Supplementary Fig. 8 | High-activity reference SMILES recovery across models at varying activity thresholds. The threshold is the normalized activity (per-series min-max scaled, 1.0 is the strongest reference active in that series). Lines indicate the mean over three independent rounds; shaded bands represent ± standard deviation.

Supplementary Table 1. Comparison of molecular quality between LeadFlow and LeadFlow-DPO generations on the MolGenBench test set.

Model	PB-Valid	QED	SA	Uniqueness	Diversity
LeadFlow	77.57%	0.540	0.773	0.844	0.570
LeadFlow-DPO	78.16%	0.546	0.773	0.816	0.562

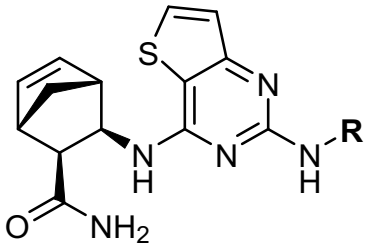


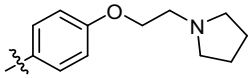
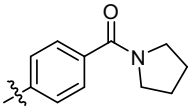
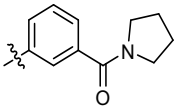
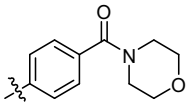
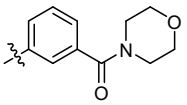
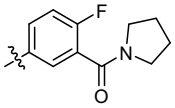
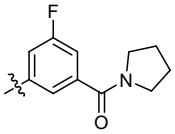
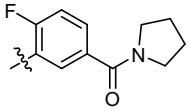
Supplementary Fig. 9 | Comparison of binding poses of the active molecule recovered by LeadFlow-DPO in the IRAK4 binding site. The green pose represents the conformation directly generated by LeadFlow-DPO, and the magenta pose represents the conformation obtained after redocking.

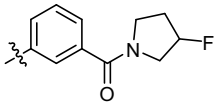
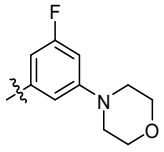
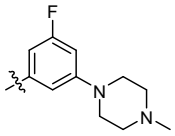
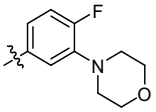
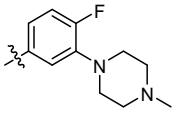


Supplementary Fig. 10 | Overview of the homologous series dataset. a. Distribution of the number of molecules per chemical series. **b.** Element composition across all ligand atoms in the dataset. **c.** Distribution of the total heavy-atom count per molecule (orange), decomposed into the MCS-derived scaffold portion (green) and the remaining atoms to be generated by the model (blue). Dashed vertical lines indicate the mean of each distribution. **d.** Distribution of pActivity values stratified by assay type, including IC_{50}, K_i, K_d, and EC_{50},

Supplementary Table 2. Compounds by varying substituents from Chen et al.[1]



Cpds	R	IRAK4 (IC_{50}, μM)
1		0.049
2		0.010
3		0.003
4		0.014
5		0.006
6		0.003
7		0.002
8		0.711

9		0.015
13		0.012
14		0.012
15		0.095
16		0.050

Supplementary Table 3. SMARTS patterns used for HBA functional group matching.

Functional Group	SMARTS
Carbonyl	[CX3](=[OX1])
Alkoxy/ether	[OX2H0][#6]
Phenol	[OX2H]
Tertiary amine	[NX3;+0;H0;!$([NX3]-[C,S]=[O,S,N])]
Nitrile	[CX2]#[NX1]
Sulfonyl	[SX4](=[OX1])(=[OX1])
Sulfoxide	[SX3](=[OX1])

Supplementary Table 4. SAR compliance of molecules generated by LeadFlow and LeadFlow-DPO for IRAK4 inhibitors.

Model	meta-HBA	para-HBA	meta-carbonyl	meta-tertiary amine	Total valid molecules generated
LeadFlow	476	63	322	14	591
LeadFlow-DPO	509	28	374	9	596

Implementation Details

Training and Inference

For LeadFlow pre-training, we use Adam optimizer with a learning rate of 5e-4 and a batch size of 6 on eight NVIDIA H800 80GB GPUs. We set $\beta_1 = 1.5$ for discrete atom types and bond types, $\sigma_1 = 0.03$ for atom coordinates. The model converged in approximately 350 epochs.

For LeadFlow-DPO, we fine-tuned the pretrained LeadFlow model using Adam optimizer with a learning rate of 1e-6 and a batch size of 3. The hyperparameter β in the DPO loss below was set to 5. Fine-tuning was performed for one epoch. All other settings followed the pretraining configuration.

$$\mathcal{L}_{\mathrm{BFN-DPO}}(\phi) = -\mathbb{E}\left[\log\sigma\left(\beta\log\frac{p_\phi(\mathbf{y}^w_{1:T} \mid \boldsymbol{\theta}^w_{0:T-1})}{p_{\mathrm{ref}}(\mathbf{y}^w_{1:T} \mid \boldsymbol{\theta}^w_{0:T-1})} - \beta\log\frac{p_\phi(\mathbf{y}^l_{1:T} \mid \boldsymbol{\theta}^l_{0:T-1})}{p_{\mathrm{ref}}(\mathbf{y}^l_{1:T} \mid \boldsymbol{\theta}^l_{0:T-1})}\right)\right]$$

For inference, we use the exponential moving average of the weights from training or fine-tuning that is updated at every optimization step with a decay factor of 0.999. We run inference with 100 sampling steps.

Hyperparameters for Network

We set the network to be kNN graphs with $k = 32$, $N = 9$ layers with $d = 128$ hidden dimension, and 16-headed attention.

Featurization

We used separate atom-level and edge-level featurization schemes for protein–ligand complexes. Protein atoms were represented by element type, amino acid type, and a backbone flag. Ligand atoms were represented using a one-hot encoding over predefined element–aromaticity combinations, such that aromaticity was incorporated directly into the ligand atom-type vocabulary (**Supplementary Table 5**).

For edge-level featurization, the heterogeneous protein–ligand graph used edge-type indicators to distinguish ligand–ligand, protein–protein, ligand–protein, and protein–ligand edges, while ligand bonds were encoded by bond type after RDKit kekulization of aromatic bonds (**Supplementary Table 6**).

After these initial features were embedded, we appended a 3-dimensional node-type indicator to distinguish protein atoms, preserved scaffold atoms, and generated ligand atoms.

Supplementary Table 5. Atom-level featurization for protein–ligand complexes

Node type	Feature category	Description
Protein atom	Element type	One-hot encoding over H, C, N, O, S, Se
Protein atom	Amino acid type	20-dimensional one-hot encoding of residue identity
Protein atom	Backbone flag	Binary indicator of whether the atom belongs to the protein backbone
Ligand atom	Atom type	One-hot encoding over predefined element–aromaticity pairs: Si; C (non-aromatic/aromatic); O (non-aromatic/aromatic); N (non-aromatic/aromatic); F; S (non-aromatic/aromatic); Cl; Br; I; P (non-aromatic/aromatic); and B

Supplementary Table 6. Edge-level featurization for heterogeneous protein–ligand graphs

Graph	Edge feature	Description
Heterogeneous protein-ligand graph	Edge type	4-dimensional one-hot encoding indicating ligand–ligand, protein–protein, ligand–protein, or protein–ligand edge type
Ligand graph	Bond type	4-dimensional one-hot encoding over non-bond, single, double, and triple bonds

Architecture of SE(3)-Graph Transformer

We employ the equivariant Graph Transformer architecture similar to Guan et al.[2] Denote node positions as $\mathbf{r} \in \mathbb{R}^{N\times 3}$, one-hot atom types as $\mathbf{h} \in \mathbb{R}^{N\times K_h}$, and one-hot bond types as $\mathbf{A} \in \mathbb{R}^{N\times(N-1)\times K_A}$, where K_h and K_A are the numbers of atom and bond types, respectively. We describe the model architecture below.

Heterogeneous Message Passing for Node Update

The heterogeneous update for node representation consists of two parts: (1) a K-nearest neighbors graph $\mathcal{G}_K$ built for interactions within the protein-ligand complex, and (2) a fully connected ligand graph $\mathcal{G}_L$ over ligand atoms to capture intra-ligand pairwise interactions.

For the complex graph $\mathcal{G}_K$, the node message is computed as

$$\Delta\mathbf{h}_{K,i} \leftarrow \sum_{j\in\mathcal{N}_K(i)} \phi_{h_K}\left(\boldsymbol{X}_{kv}^K, \boldsymbol{X}_q^K\right),$$

where

$$\boldsymbol{X}_{kv}^K = \text{concat}\left(\mathbf{h}_i, \mathbf{h}_j, \mathbf{e}_{ij}^K\right), \qquad \boldsymbol{X}_q^K = \mathbf{h}_i.$$

Here,

$$\mathbf{e}_{ij}^K = \text{concat}\left(\mathbf{h}_{ij}^{\text{edge}} \otimes \sigma\left(\mathbf{d}_{ij}\right), \mathbf{h}_{ij}^{\text{edge}}\right),$$

$\mathbf{d}_{ij} = \left\|\mathbf{r}_i - \mathbf{r}_j\right\|$ is the pairwise atom distance, $\sigma(\cdot)$ is the Gaussian smearing function applied to the pairwise distance, and $\mathbf{h}_{ij}^{\text{edge}}$ is the one-hot encoding for edge types between protein-ligand atoms (protein-protein, protein-ligand, ligand-protein or ligand-ligand).

For the ligand graph $\mathcal{G}_L$, the node message is computed as

$$\Delta\mathbf{h}_{L,i} \leftarrow \sum_{j\in\mathcal{N}_L(i)} \phi_{h_L}\left(\boldsymbol{X}_{kv}^L, \boldsymbol{X}_q^L\right),$$

with

$$\boldsymbol{X}_{kv}^L = \text{concat}\left(\mathbf{h}_i, \mathbf{h}_j, \mathbf{h}_{ij}^{\text{bond}}\right), \qquad \boldsymbol{X}_q^L = \mathbf{h}_i.$$

where $\mathbf{h}_{ij}^{\text{bond}}$ denotes the current bond hidden representation on the ligand graph edge.

The node-level representation is updated by aggregating the heterogeneous messages:

$$\mathbf{h}_i \leftarrow \mathbf{h}_i + \boldsymbol{W}_h\big(\Delta\mathbf{h}_{K,i} + \Delta\mathbf{h}_{L,i}\big).$$

Intra-Ligand Message Passing for Bond Update

The bond representation is updated by a directional message passing:

$$\mathbf{h}_{ji}^{\text{bond}} \leftarrow \sum_{k\in\mathcal{N}_L(j)\setminus\{i\}} \phi_E\big(\boldsymbol{X}_{kv}^E, \boldsymbol{X}_q^E\big),$$

where

$$\boldsymbol{X}_{kv}^E = \text{concat}\big(\mathbf{h}_{kj}^{\text{bond}}, \sigma(\mathbf{d}_{kj}), \sigma(\mathbf{d}_{ji}), \mathbf{a}_{kji}, \mathbf{h}_k, \mathbf{h}_j\big), \qquad \boldsymbol{X}_q^E = \text{concat}\big(\mathbf{h}_{ji}^{\text{bond}}, \mathbf{h}_i\big).$$

Here, $\mathbf{a}$ is the angular encoding for the bond triplets (k, j, i).

Heterogeneous Message Passing for Position Updates

Similarly, the positions are updated from heterogeneous messages:

$$\Delta\mathbf{r}_{K,i} \leftarrow \sum_{j\in\mathcal{N}_K(i)} \big(\mathbf{r}_i - \mathbf{r}_j\big)\, \phi_{r_K}\big(\boldsymbol{X}_{kv}^K, \boldsymbol{X}_q^K\big),$$

$$\Delta\mathbf{r}_{L,i} \leftarrow \sum_{j\in\mathcal{N}_L(i)} \big(\mathbf{r}_i - \mathbf{r}_j\big)\, \phi_{r_L}\big(\boldsymbol{X}_{kv}^L, \boldsymbol{X}_q^L\big),$$

given the updated node features $\mathbf{h}_i$ and bond features $\mathbf{h}_{ij}^{\text{bond}}$ for input feature $\boldsymbol{X}$. The final position is updated by

$$\mathbf{r}_i \leftarrow \mathbf{r}_i + \big(\Delta\mathbf{r}_{K,i} + \Delta\mathbf{r}_{L,i}\big) \cdot \mathbf{1}_{\text{mod}}$$

where $\mathbf{1}_{\text{mod}}$ is the Boolean mask for the atoms of modification, as the protein atoms and scaffold atoms remain fixed.

Attention Mechanism

The function ϕ is implemented as a multi-headed attention mechanism. For each branch, the query, key, and value are obtained from the corresponding input features:

$$\mathbf{Q} = \psi_q(\boldsymbol{X}_q), \qquad \mathbf{K} = \psi_k(\boldsymbol{X}_{kv}), \qquad \mathbf{V} = \psi_v(\boldsymbol{X}_{kv}).$$

where ψ_q, ψ_k, and ψ_v are learnable projections implemented as multilayer perceptrons.

The attention scores are computed by scaled dot-product attention:

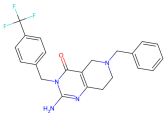
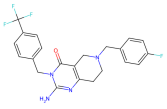
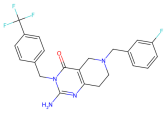
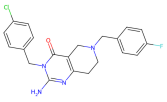
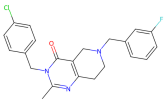
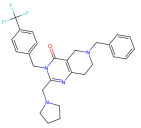
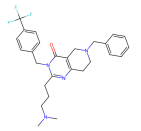
$$\boldsymbol{\alpha}_{ij} = \mathrm{softmax}\left(\frac{\mathbf{q}_i \mathbf{k}_j}{\sqrt{d_h}}\right),$$

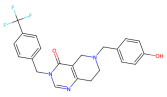
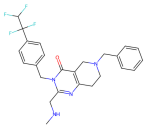
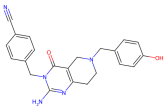
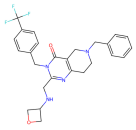
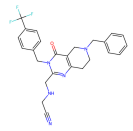
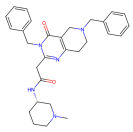
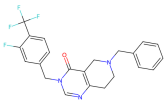
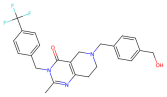
where d_h is the head dimension. The output of ϕ is the attention-weighted message:

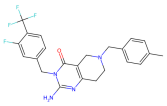
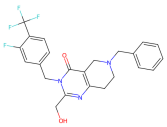
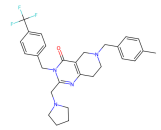
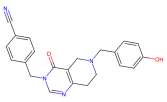
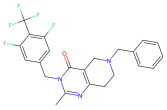
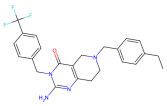
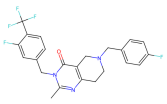
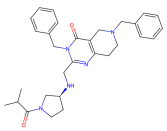
$$\phi(\boldsymbol{X}_{kv}, \boldsymbol{X}_q) = \boldsymbol{\alpha}_{ij} \mathbf{v}_j.$$

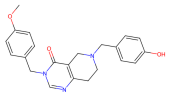
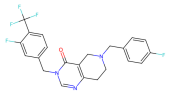
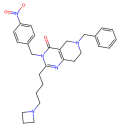
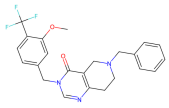
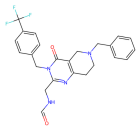
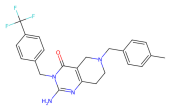
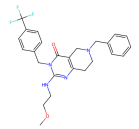
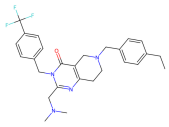
ClpP activator candidates

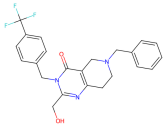
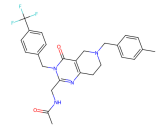
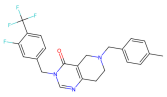
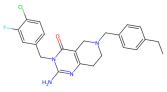
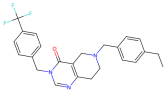
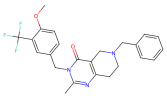
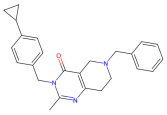
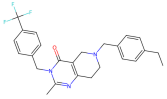
Supplementary Table 7. Computational evaluation of 50 selected ClpP activator candidates

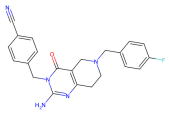
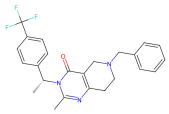
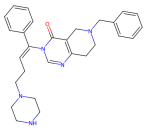
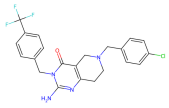
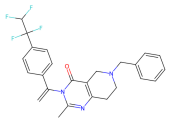
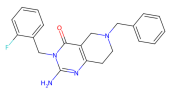
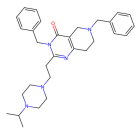
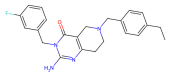
Original Output ID	Structure	Synthesized Compound ID	SMILES	SA score	Docking score (kcal/mol)	RMSD (Å)
mol_1		DC-ClpP-01	Nc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccccc1)CC2	2.373	-10.43	0.202
mol_200		DC-ClpP-02	Nc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccc(F)cc1)CC2	2.439	-10.44	0.515
mol_44		DC-ClpP-03	Nc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1cccc(F)c1)CC2	2.480	-11.00	0.220
mol_407		DC-ClpP-04	Nc1nc2c(c(=O)n1Cc1ccc(Cl)cc1)CN(Cc1ccc(F)cc1)CC2	2.297	-9.99	0.152
mol_960		DC-ClpP-05	Cc1nc2c(c(=O)n1Cc1ccc(Cl)cc1)CN(Cc1cccc(F)c1)CC2	2.243	-10.40	0.169
mol_1881			O=c1c2c(nc(CN3CCCC3)n1Cc1ccc(C(F)(F)F)cc1)CCN(Cc1ccccc1)C2	2.444	-11.04	0.230
mol_2094			CN(C)CCCc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccccc1)CC2	2.493	-10.99	0.465

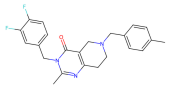
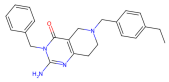
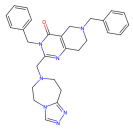
Original Output ID	Structure	Synthesized Compound ID	SMILES	SA score	Docking score (kcal/mol)	RMSD (Å)
mol_1546			O=c1c2c(ncn1Cc1ccc(C(F)(F)F)cc1)CCN(Cc1ccc(O)cc1)C2	2.401	-10.98	0.245
mol_211			CNCc1nc2c(c(=O)n1Cc1ccc(C(F)(F)C(F)F)cc1)CN(Cc1ccccc1)CC2	2.750	-10.96	0.325
mol_2817			N#Cc1ccc(Cn2c(N)nc3c(c2=O)CN(Cc2ccc(O)cc2)CC3)cc1	2.479	-10.93	0.143
mol_1745			O=c1c2c(nc(CNC3COC3)n1Cc1ccc(C(F)(F)F)cc1)CCN(Cc1ccccc1)C2	2.630	-10.84	0.381
mol_591			N#CCNCc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccccc1)CC2	2.633	-10.78	0.206
mol_2892			CN1CCC[C@H](NC(=O)Cc2nc3c(c(=O)n2Cc2ccccc2)CN(Cc2ccccc2)CC3)C1	2.944	-10.78	0.165
mol_1078			O=c1c2c(ncn1Cc1ccc(C(F)(F)F)c(F)c1)CCN(Cc1ccccc1)C2	2.445	-10.73	0.222
mol_469			Cc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccc(CO)cc1)CC2	2.417	-10.72	0.369

Original Output ID	Structure	Synthesized Compound ID	SMILES	SA score	Docking score (kcal/mol)	RMSD (Å)
mol_486			Cc1ccc(CN2CCc3nc(N)n(Cc4ccc(C(F)(F)F)c(F)c4)c(=O)c3C2)cc1	2.583	-10.71	0.269
mol_3118			O=c1c2c(nc(CO)n1Cc1ccc(C(F)(F)F)c(F)c1)CCN(Cc1ccccc1)C2	2.547	-10.66	0.186
mol_1870			Cc1ccc(CN2CCc3nc(CN4CCCC4)n(Cc4ccc(C(F)(F)F)cc4)c(=O)c3C2)cc1	2.495	-10.65	0.276
mol_763			N#Cc1ccc(Cn2cnc3c(c2=O)CN(Cc2ccc(O)cc2)CC3)cc1	2.388	-10.65	0.179
mol_1999			Cc1nc2c(c(=O)n1Cc1cc(F)c(C(F)(F)F)c(F)c1)CN(Cc1ccccc1)CC2	2.602	-10.64	0.197
mol_961			CCc1ccc(CN2CCc3nc(N)n(Cc4ccc(C(F)(F)F)cc4)c(=O)c3C2)cc1	2.461	-10.62	0.392
mol_1725			Cc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)c(F)c1)CN(Cc1ccc(F)cc1)CC2	2.510	-10.60	0.504
mol_2179			CC(C)C(=O)N1CC[C@H](NCc2nc3c(c(=O)n2Cc2ccccc2)CN(Cc2ccccc2)CC3)C1	2.999	-10.60	0.159

Original Output ID	Structure	Synthesized Compound ID	SMILES	SA score	Docking score (kcal/mol)	RMSD (Å)
mol_1566			COc1ccc(Cn2cnc3c(c2=O)CN(Cc2ccc(O)cc2)CC3)cc1	2.247	-10.59	0.279
mol_428			O=c1c2c(ncn1Cc1ccc(C(F)(F)F)c(F)c1)CCN(Cc1ccc(F)cc1)C2	2.510	-10.59	0.509
mol_2608			O=c1c2c(nc(CCCCN3CCC3)n1Cc1ccc([N+](=O)[O-])cc1)CCN(Cc1ccccc1)C2	2.486	-10.57	0.438
mol_1220			COc1cc(Cn2cnc3c(c2=O)CN(Cc2ccccc2)CC3)ccc1C(F)(F)F	2.438	-10.57	0.471
mol_574			O=CNCc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccccc1)CC2	2.591	-10.56	0.200
mol_229			Cc1ccc(CN2CCc3nc(N)n(Cc4ccc(C(F)(F)F)cc4)c(=O)c3C2)cc1	2.424	-10.56	0.391
mol_1054			COCCNc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccccc1)CC2	2.511	-10.53	0.220
mol_1801			CCc1ccc(CN2CCc3nc(CN(C)C)n(Cc4ccc(C(F)(F)F)cc4)c(=O)c3C2)cc1	2.538	-10.52	0.374

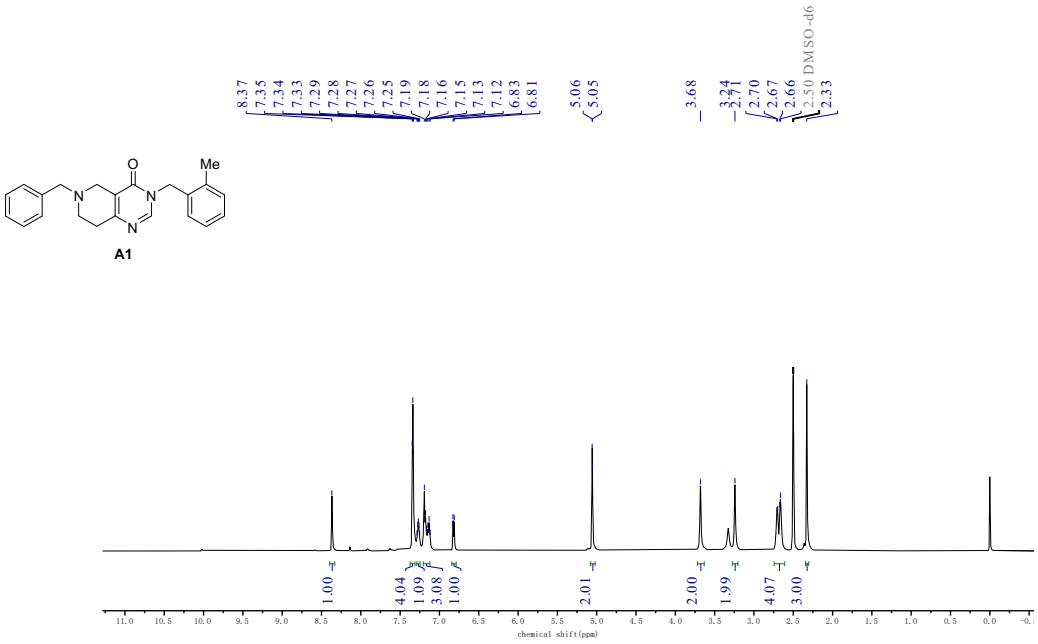
Original Output ID	Structure	Synthesized Compound ID	SMILES	SA score	Docking score (kcal/mol)	RMSD (Å)
mol_2103			O=c1c2c(nc(CO)n1Cc1ccc(C(F)(F)F)cc1)CCN(Cc1ccccc1)C2	2.389	-10.50	0.203
mol_2183			CC(=O)NCc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccc(C)cc1)CC2	2.499	-10.46	0.471
mol_2447			Cc1ccc(CN2CCc3ncn(Cc4ccc(C(F)(F)F)c(F)c4)c(=O)c3C2)cc1	2.495	-10.43	0.256
mol_171			CCc1ccc(CN2CCc3nc(N)n(Cc4ccc(Cl)c(F)c4)c(=O)c3C2)cc1	2.462	-10.41	0.325
mol_2384			CCc1ccc(CN2CCc3ncn(Cc4ccc(C(F)(F)F)cc4)c(=O)c3C2)cc1	2.371	-10.40	0.370
mol_655			COc1ccc(Cn2c(C)nc3c(c2=O)CN(Cc2ccccc2)CC3)cc1C(F)(F)F	2.420	-10.39	0.172
mol_1626			Cc1nc2c(c(=O)n1Cc1ccc(C3CC3)cc1)CN(Cc1ccccc1)CC2	2.242	-10.38	0.148
mol_881			CCc1ccc(CN2CCc3nc(C)n(Cc4ccc(C(F)(F)F)cc4)c(=O)c3C2)cc1	2.374	-10.38	0.426

Original Output ID	Structure	Synthesized Compound ID	SMILES	SA score	Docking score (kcal/mol)	RMSD (Å)
mol_2601			N#Cc1ccc(Cn2c(N)nc3c(c2=O)CN(Cc2ccc(F)cc2)CC3)cc1	2.424	-10.32	0.137
mol_1525			Cc1nc2c(c(=O)n1[C@H](C)c1ccc(C(F)(F)F)cc1)CN(Cc1ccccc1)CC2	2.877	-10.30	0.385
mol_1787			O=c1c2c(ncn1/C(=C\CCN1CCNCC1)c1ccccc1)CCN(Cc1ccccc1)C2	2.873	-10.29	0.210
mol_543			Nc1nc2c(c(=O)n1Cc1ccc(C(F)(F)F)cc1)CN(Cc1ccc(Cl)cc1)CC2	2.435	-10.27	0.411
mol_588			C=C(c1ccc(C(F)(F)C(F)F)cc1)n1c(C)nc2c(c1=O)CN(Cc1ccccc1)CC2	2.971	-10.20	0.176
mol_1043			Nc1nc2c(c(=O)n1Cc1ccccc1F)CN(Cc1ccccc1)CC2	2.286	-10.13	0.075
mol_1659			CC(C)N1CCN(CCc2nc3c(c(=O)n2Cc2ccccc2)CN(Cc2ccccc2)CC3)CC1	2.455	-10.07	0.203
mol_1328			CCc1ccc(CN2CCc3nc(N)n(Cc4cccc(F)c4)c(=O)c3C2)cc1	2.366	-10.06	0.339

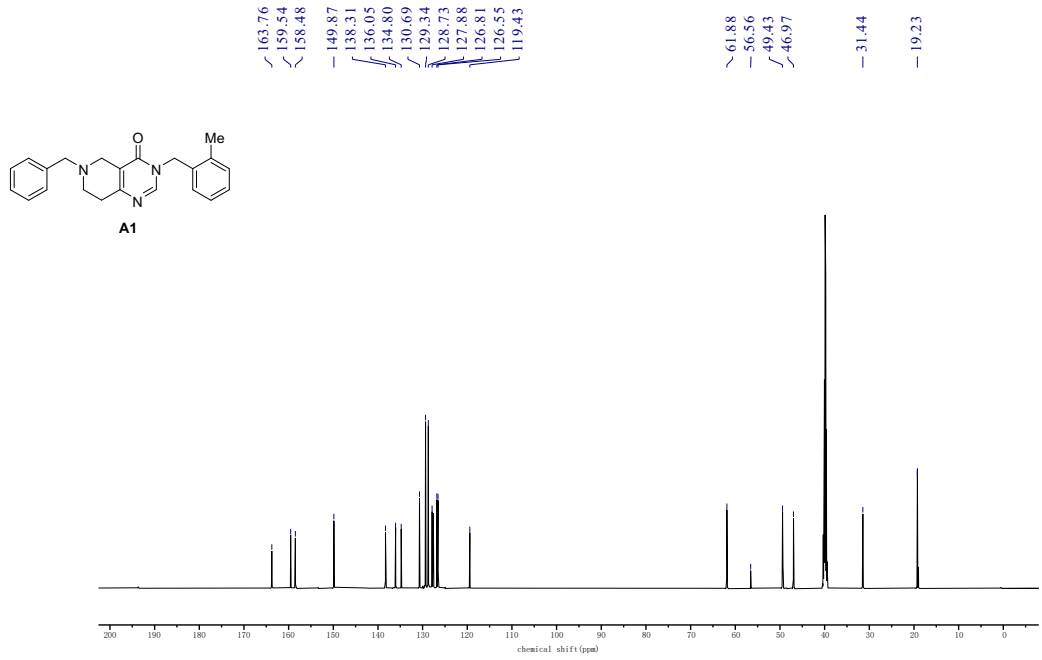
Original Output ID	Structure	Synthesized Compound ID	SMILES	SA score	Docking score (kcal/mol)	RMSD (Å)
mol_1811			Cc1ccc(CN2CCc3nc(C)n(Cc4ccc(F)c(F)c4)c(=O)c3C2)cc1	2.308	-9.97	0.382
mol_138			CCc1ccc(CN2CCc3nc(N)n(Cc4ccccc4)c(=O)c3C2)cc1	2.254	-9.96	0.248
mol_313			O=c1c2c(nc(CN3CCc4nncn4CC3)n1Cc1ccccc1)CCN(Cc1ccccc1)C2	2.675	-9.83	0.138

1H/ 13C/ 19F NMR, HRMS spectra and HPLC spectra of compounds

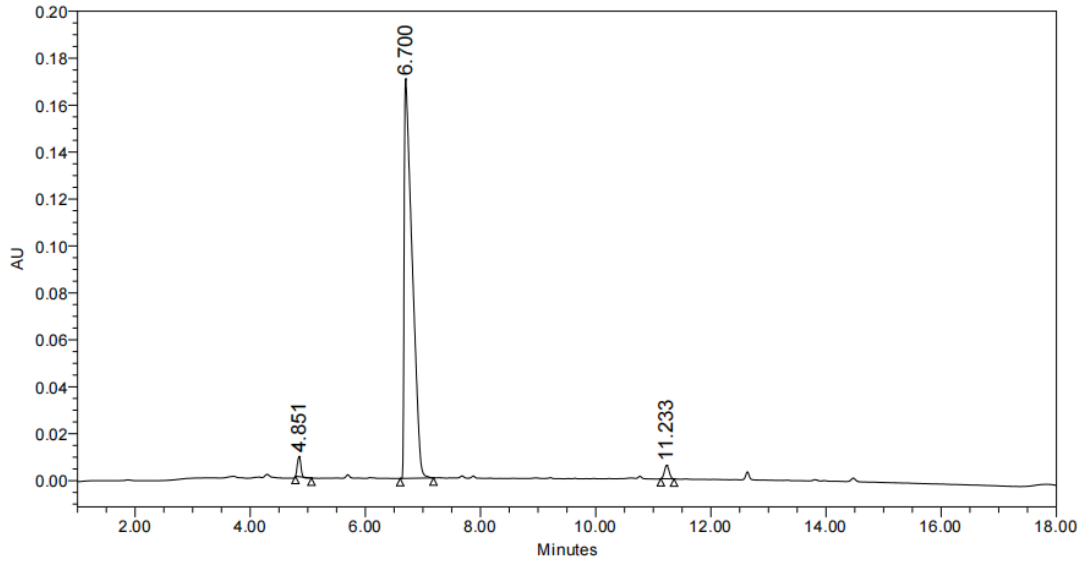
6-benzyl-3-(2-methylbenzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (A1)



^{1}H NMR of **A1**



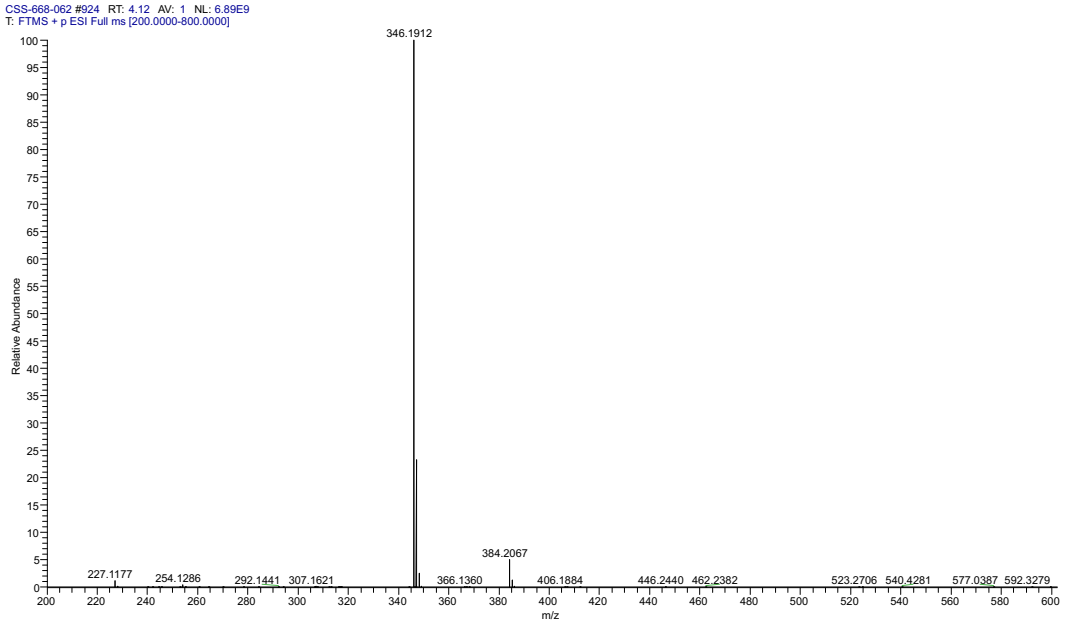
^{13}C NMR of **A1**



Peak Results

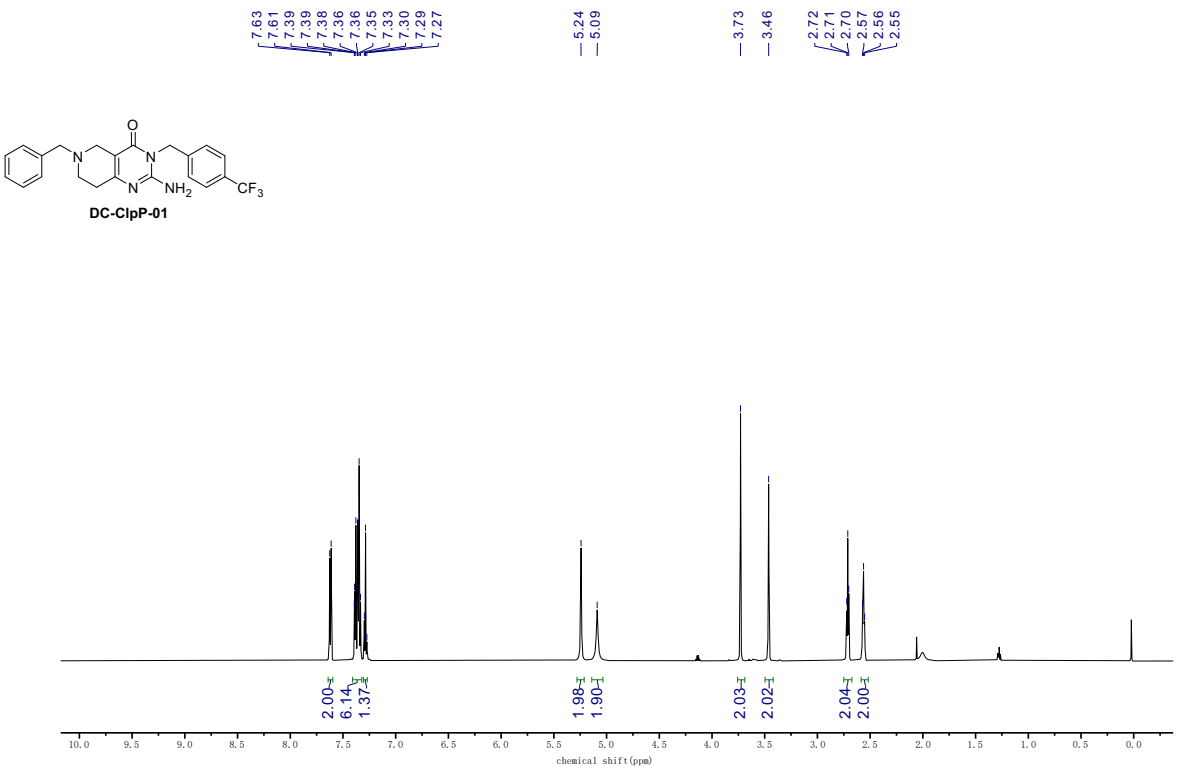
	RT	Area	Height	% Area
1	4.851	36871	8635	2.16
2	6.700	1638953	170264	95.96
3	11.233	32197	5771	1.89

HPLC of **A1**

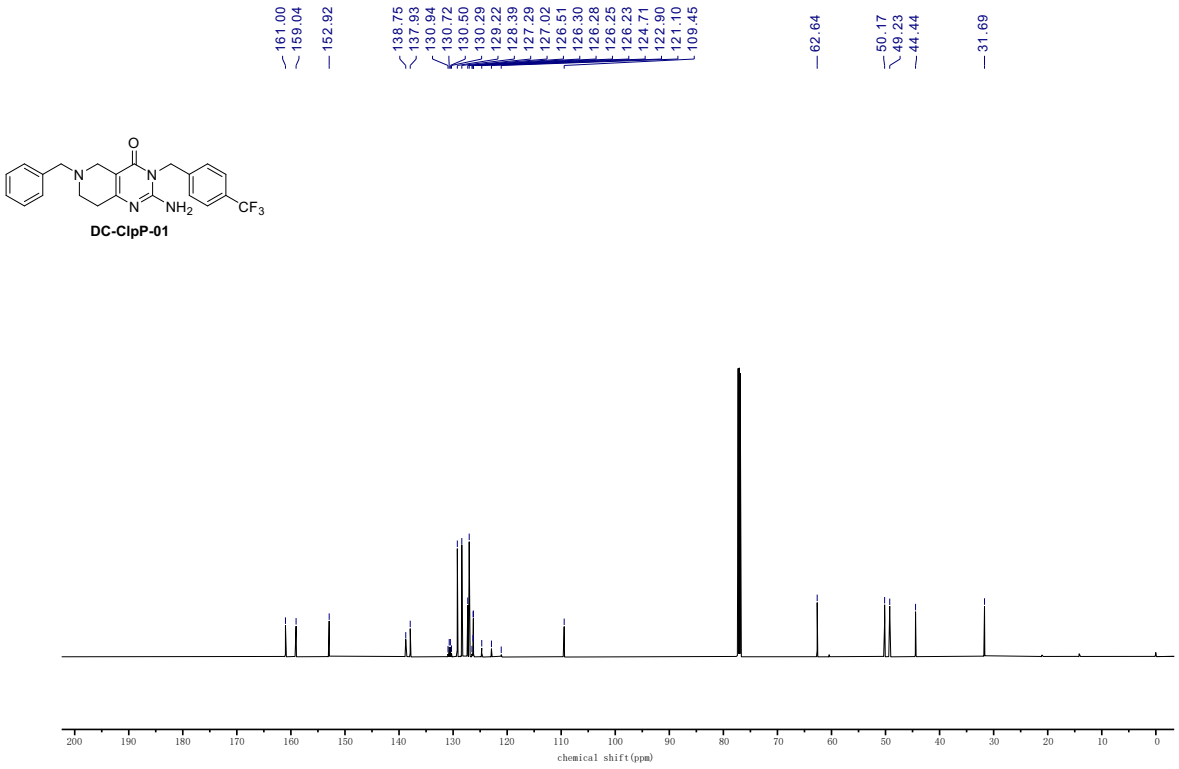


ESI-HRMS of **A1**

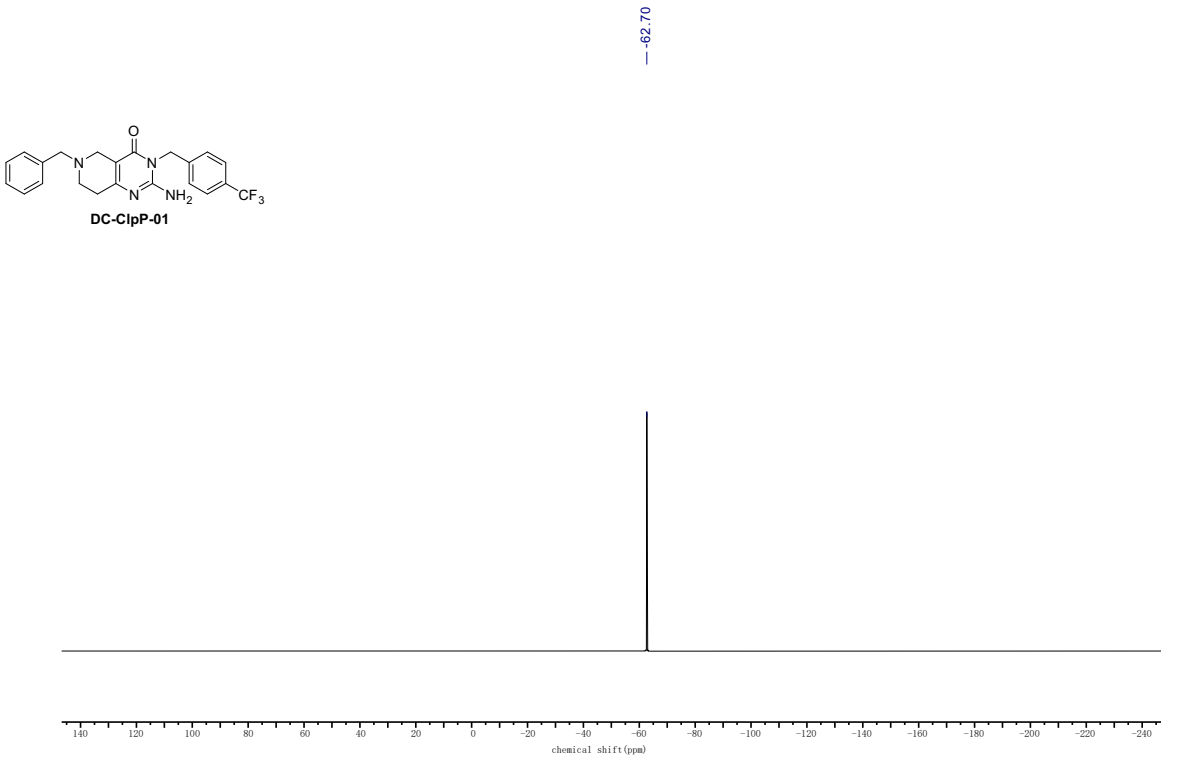
2-amino-6-benzyl-3-(4-(trifluoromethyl)benzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (DC-ClpP-01)



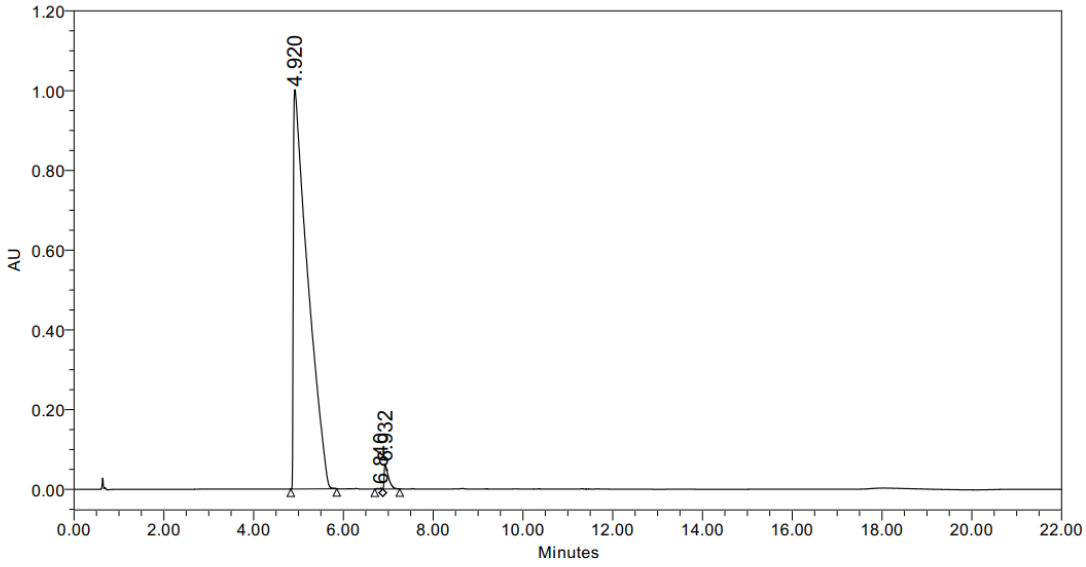
^{1}H NMR of **DC-ClpP-01**



^{13}C NMR of **DC-ClpP-01**



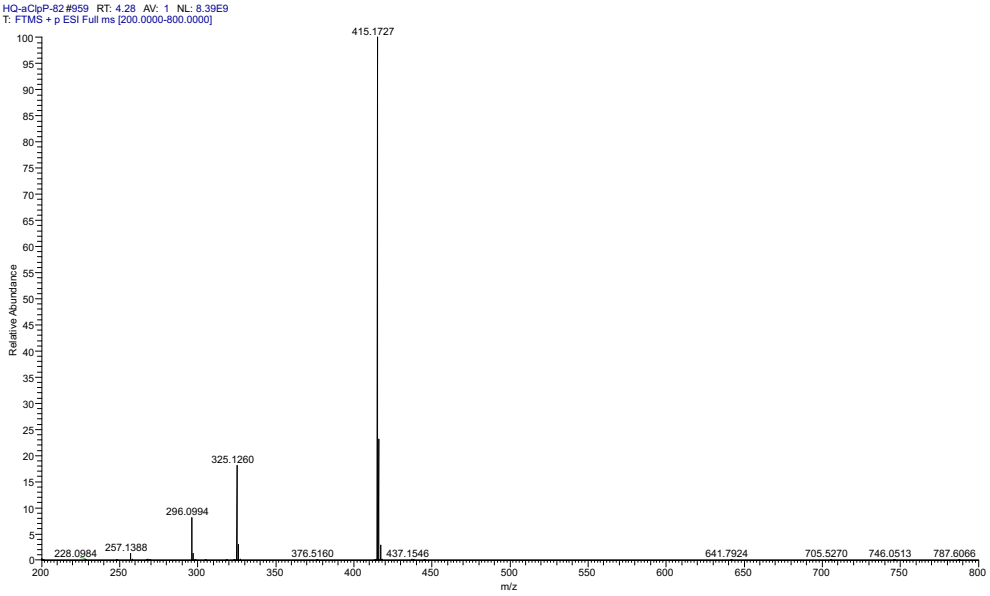
^{19}F NMR of **DC-ClpP-01**



Peak Results

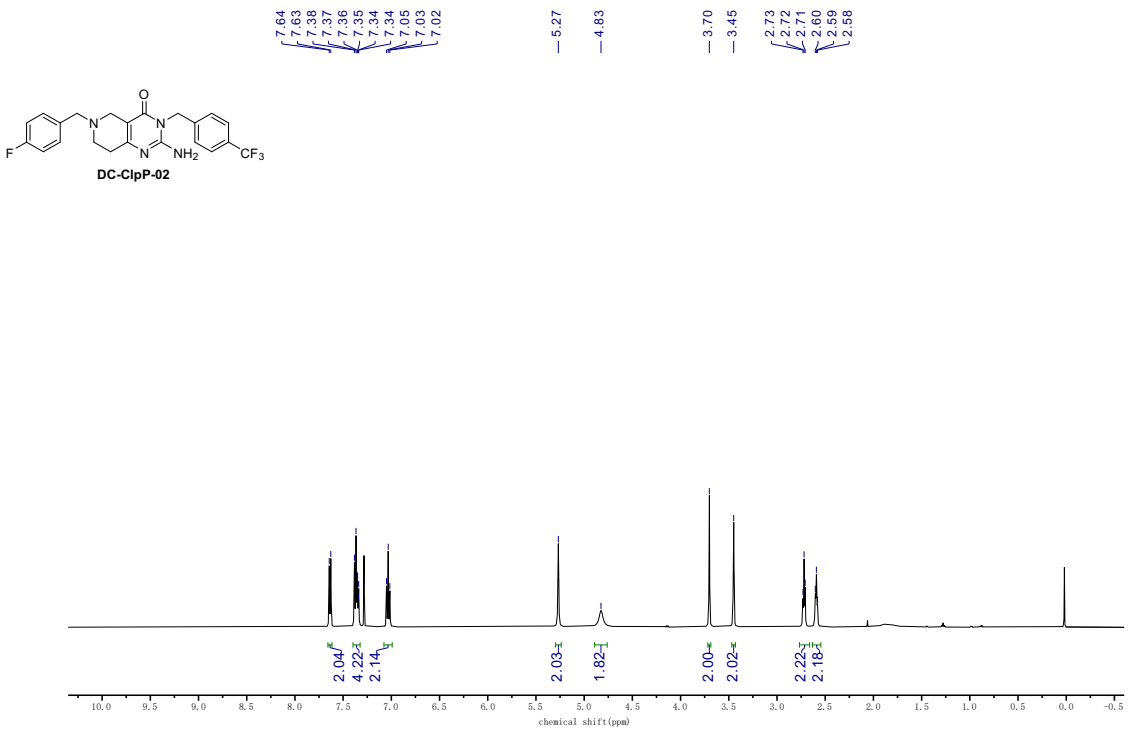
	RT	Area	Height	% Area
1	4.920	22314453	1001509	98.20
2	6.840	11485	2548	0.05
3	6.932	397272	60956	1.75

HPLC of **DC-ClpP-01**

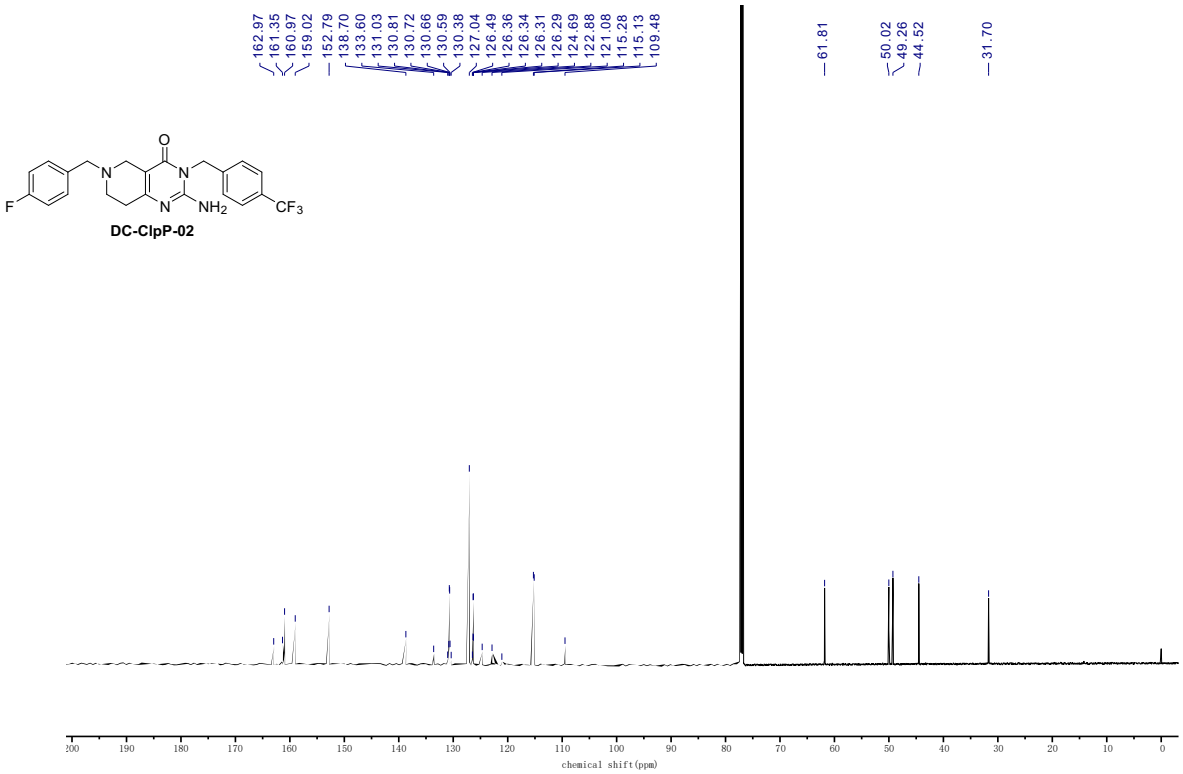


ESI-HRMS of **DC-ClpP-01**

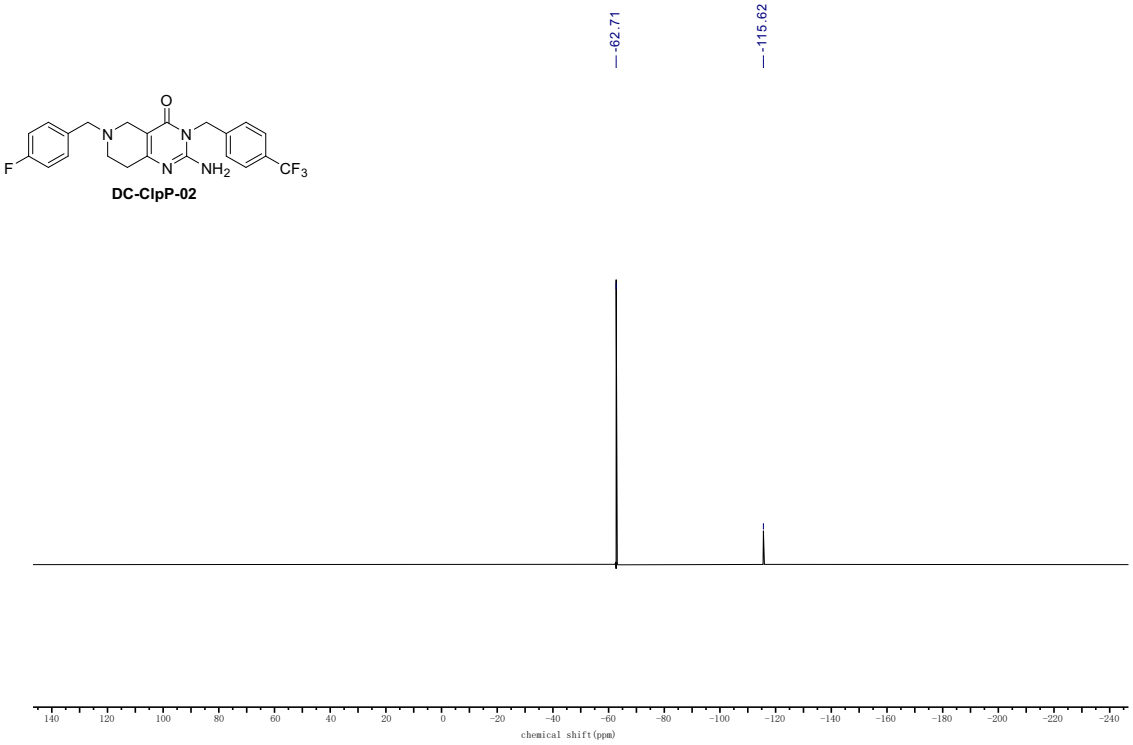
2-amino-6-(4-fluorobenzyl)-3-(4-(trifluoromethyl)benzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (DC-ClpP-02)



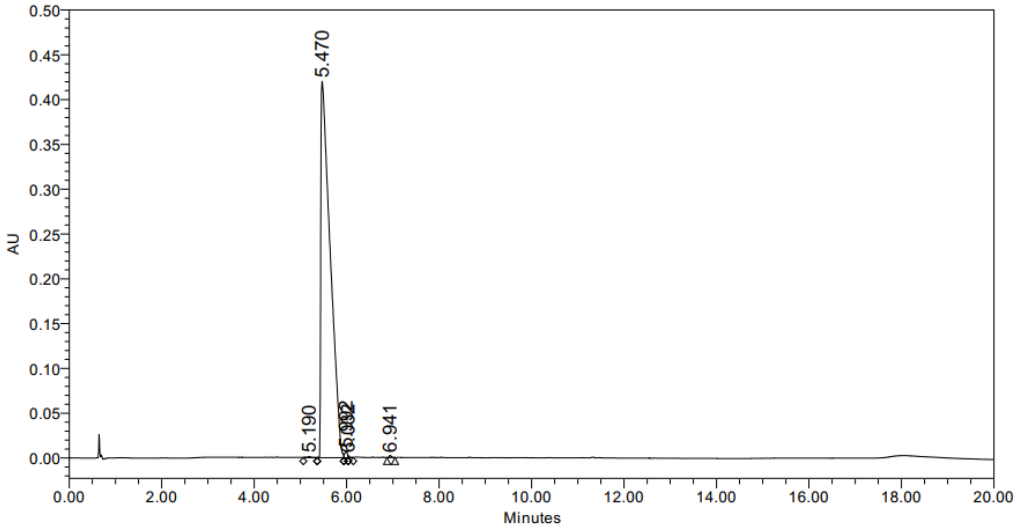
^{1}H NMR of **DC-ClpP-02**



^{13}C NMR of **DC-ClpP-02**



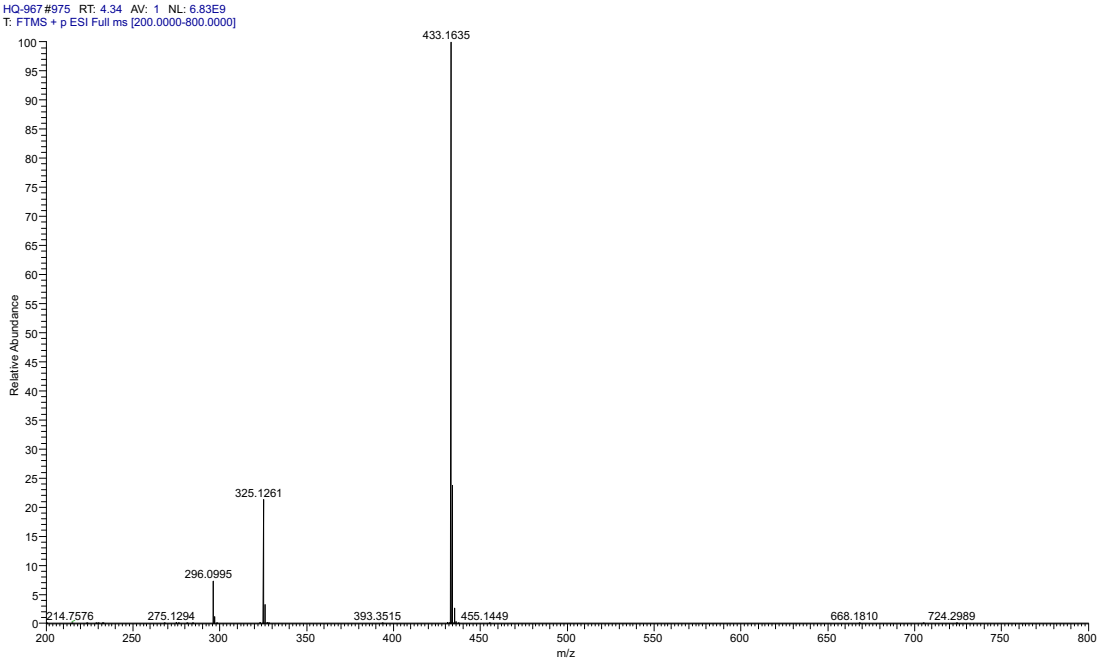
^{19}F NMR of **DC-ClpP-02**



Peak Results

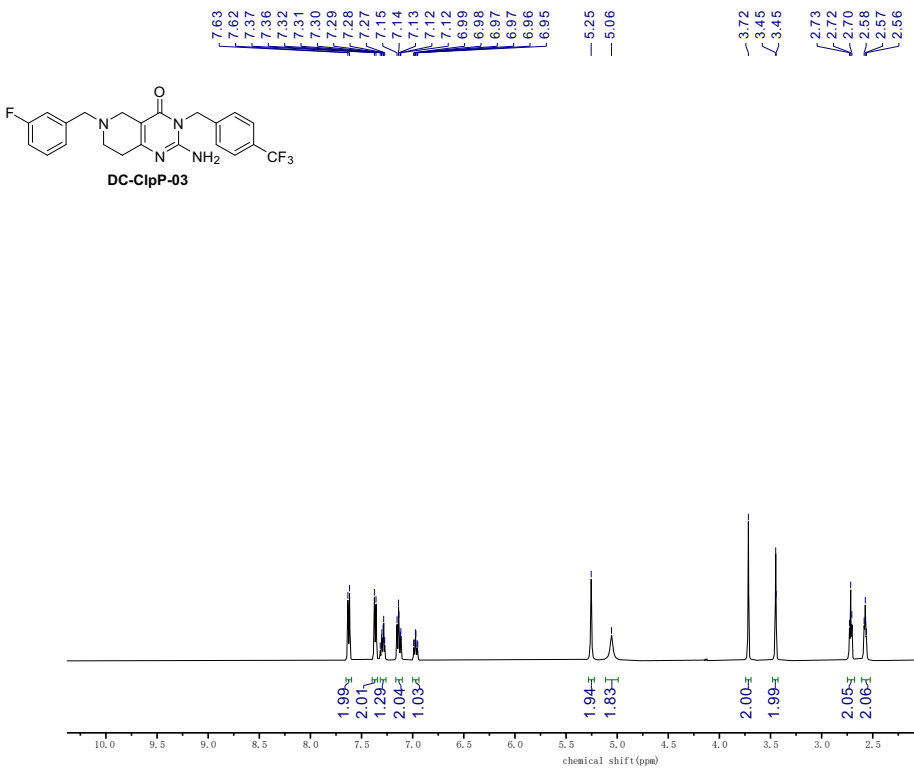
	RT	Area	Height	% Area
1	5.190	10063	1364	0.17
2	5.470	5879795	420004	99.19
3	5.992	27776	6714	0.47
4	6.032	3898	2206	0.07
5	6.941	6534	2339	0.11

HPLC of **DC-ClpP-02**

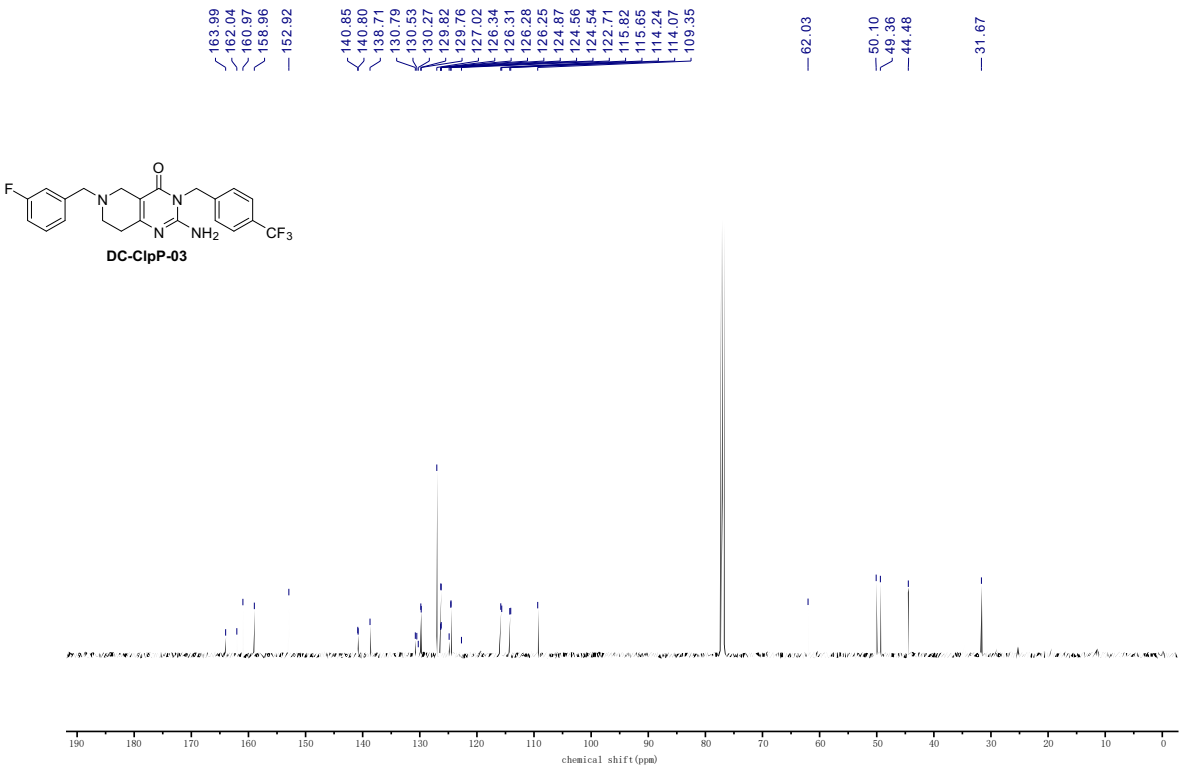


ESI-HRMS of **DC-ClpP-02**

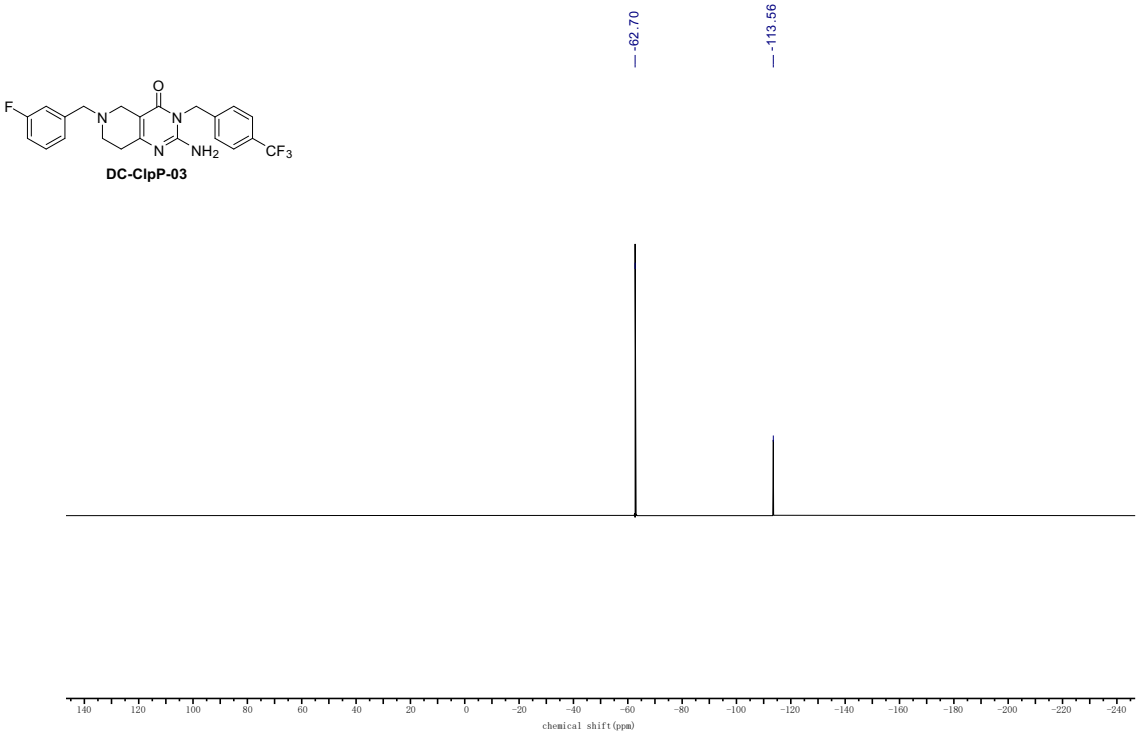
2-amino-6-(3-fluorobenzyl)-3-(4-(trifluoromethyl)benzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (DC-ClpP-03)



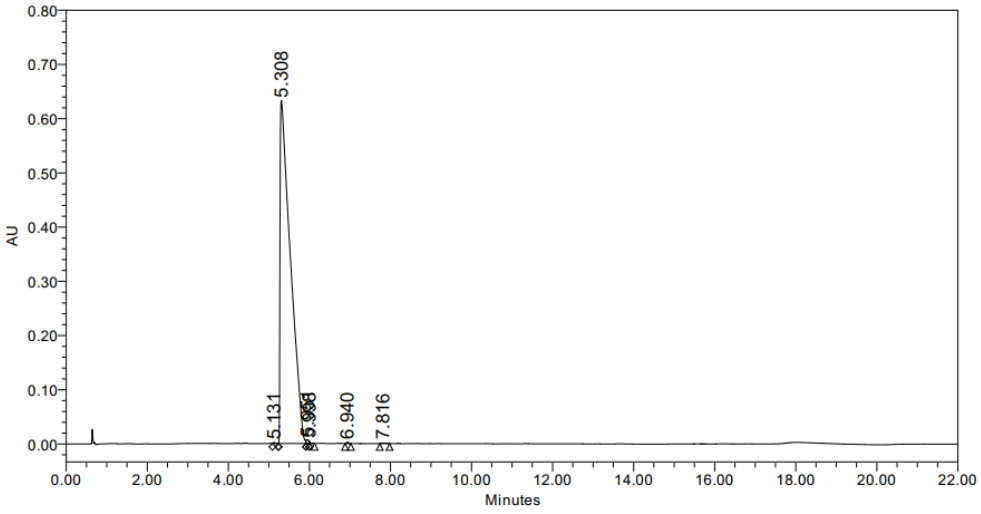
^{1}H NMR of **DC-ClpP-03**



^{13}C NMR of **DC-ClpP-03**



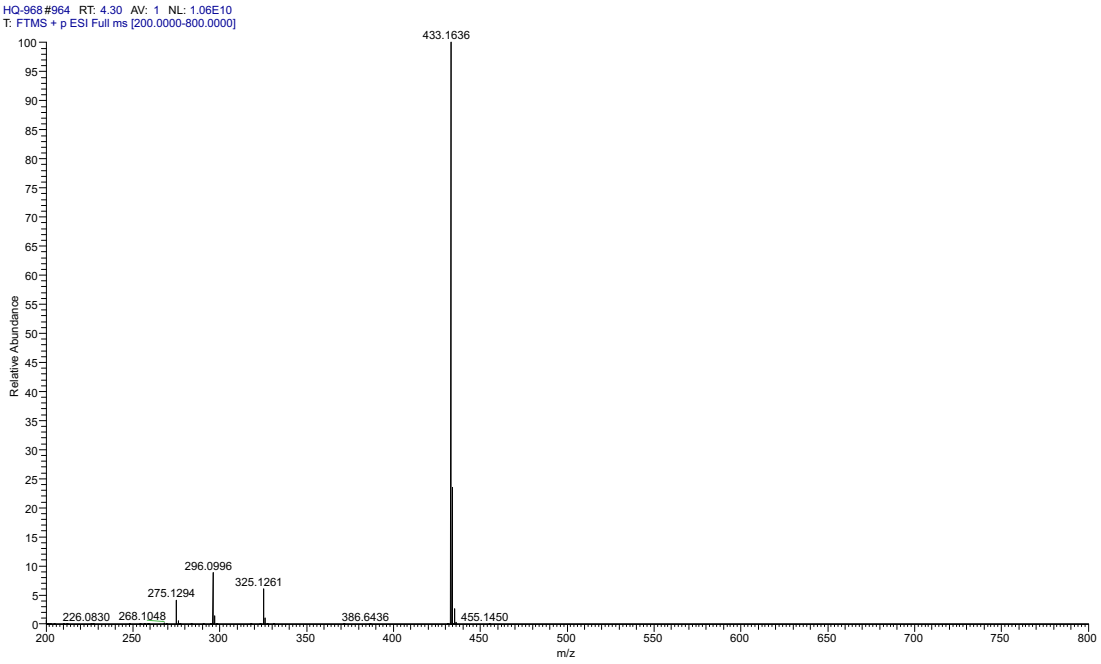
^{19}F NMR of **DC-ClpP-03**



Peak Results

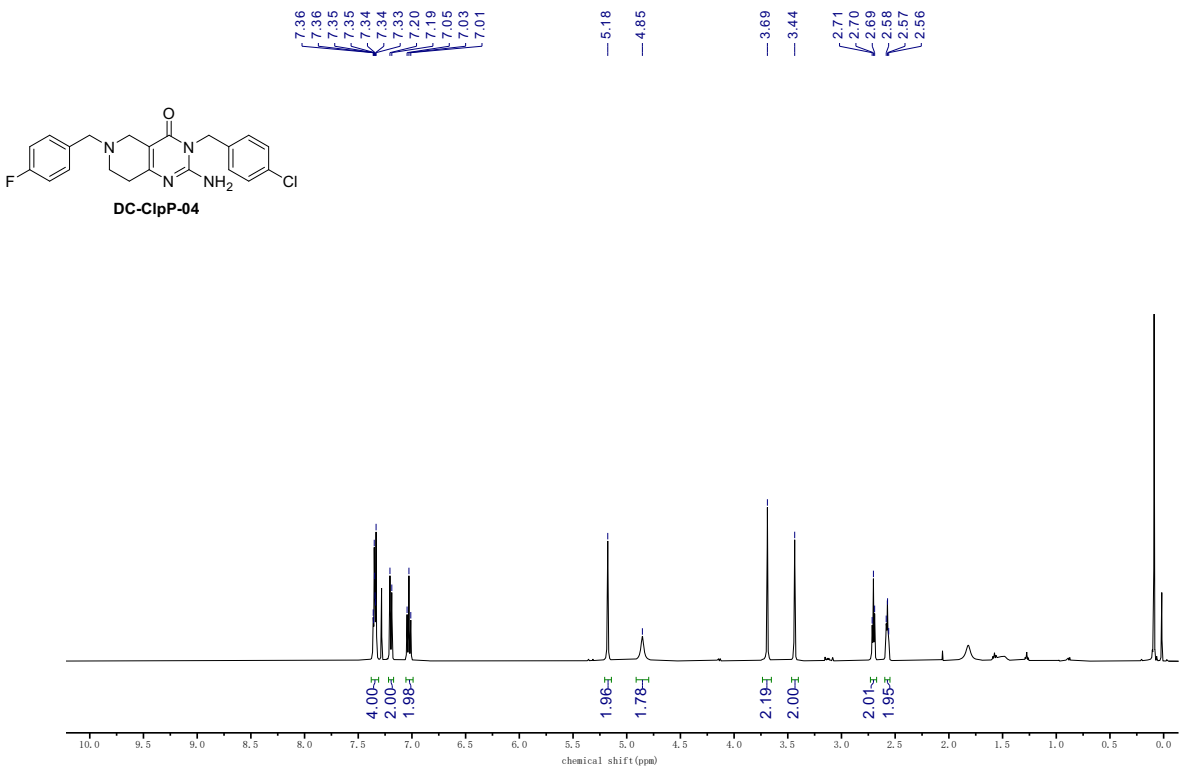
	RT	Area	Height	% Area
1	5.131	6246	2061	0.06
2	5.308	11116813	632367	99.49
3	5.951	29761	6790	0.27
4	5.998	5246	3812	0.05
5	6.940	7732	2745	0.07
6	7.816	7460	1232	0.07

HPLC of **DC-ClpP-03**

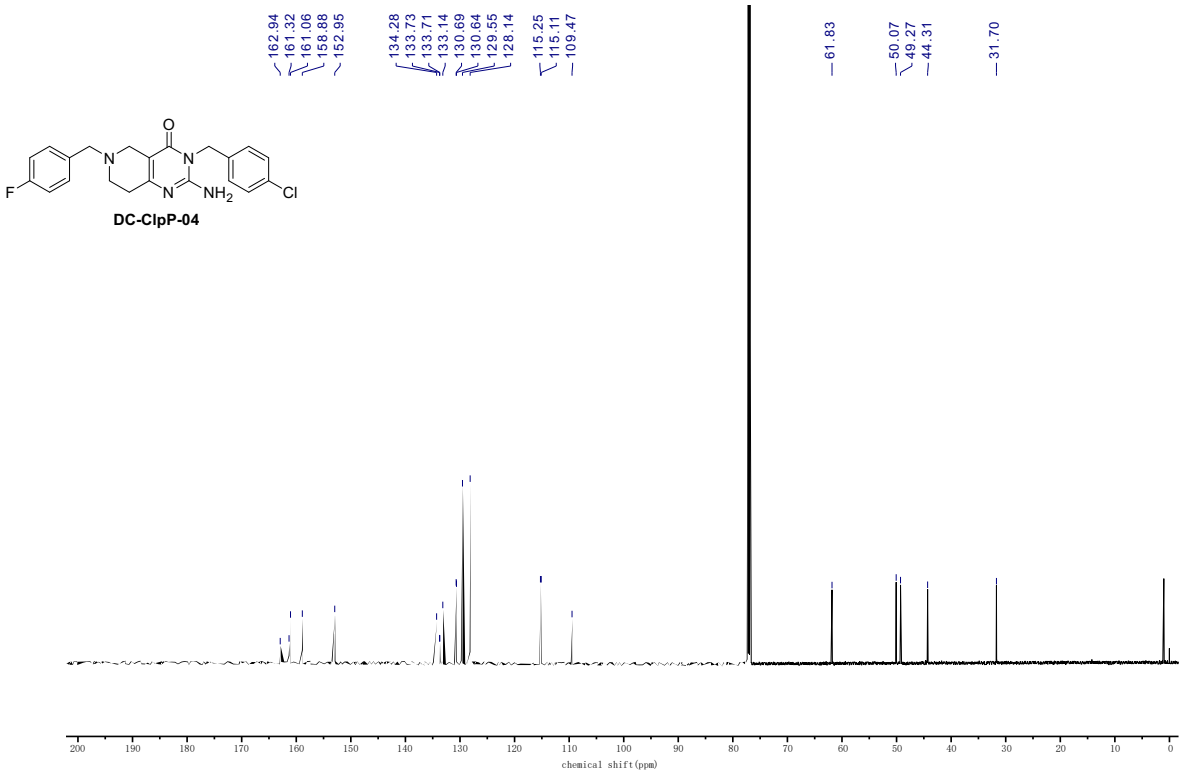


ESI-HRMS of **DC-ClpP-03**

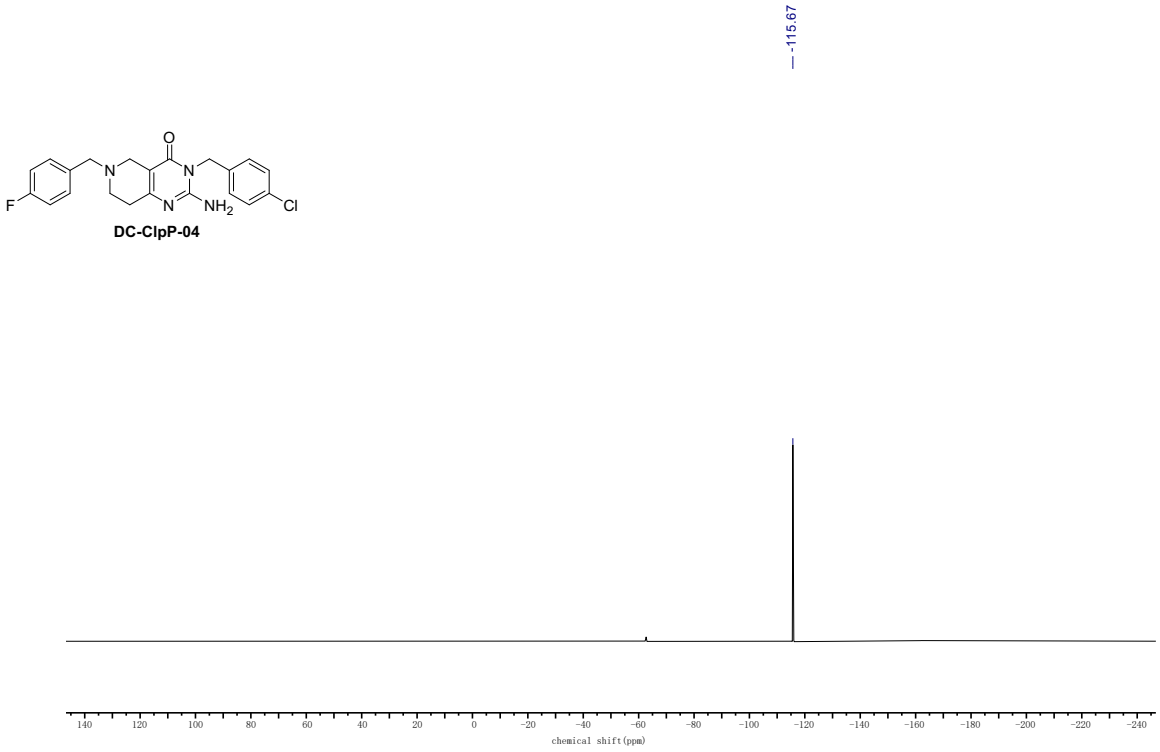
2-amino-3-(4-chlorobenzyl)-6-(4-fluorobenzyl)-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (DC-ClpP-04)



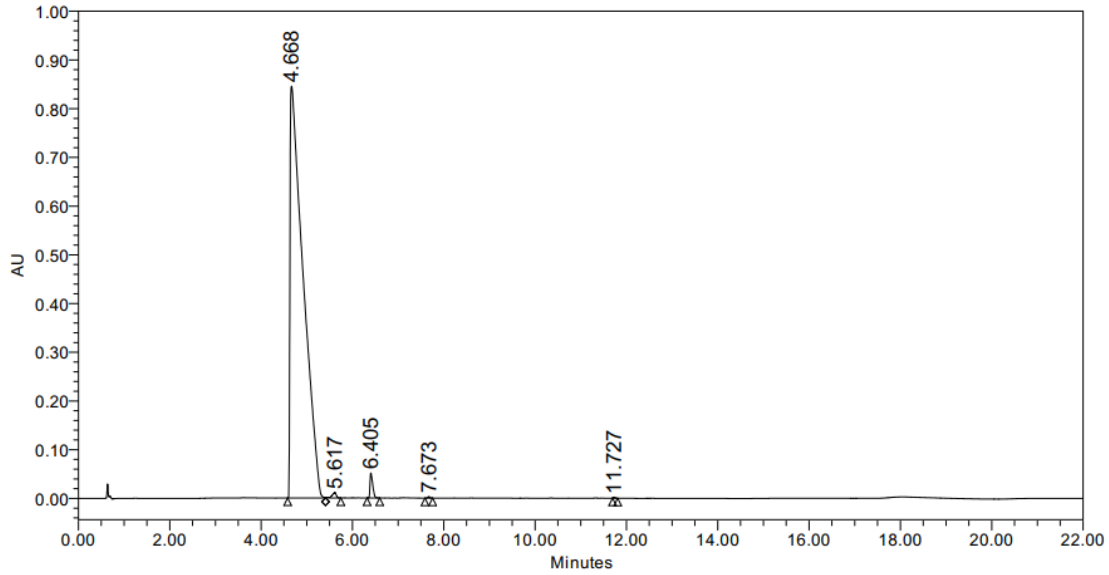
^{1}H NMR of **DC-ClpP-04**



^{13}C NMR of **DC-ClpP-04**



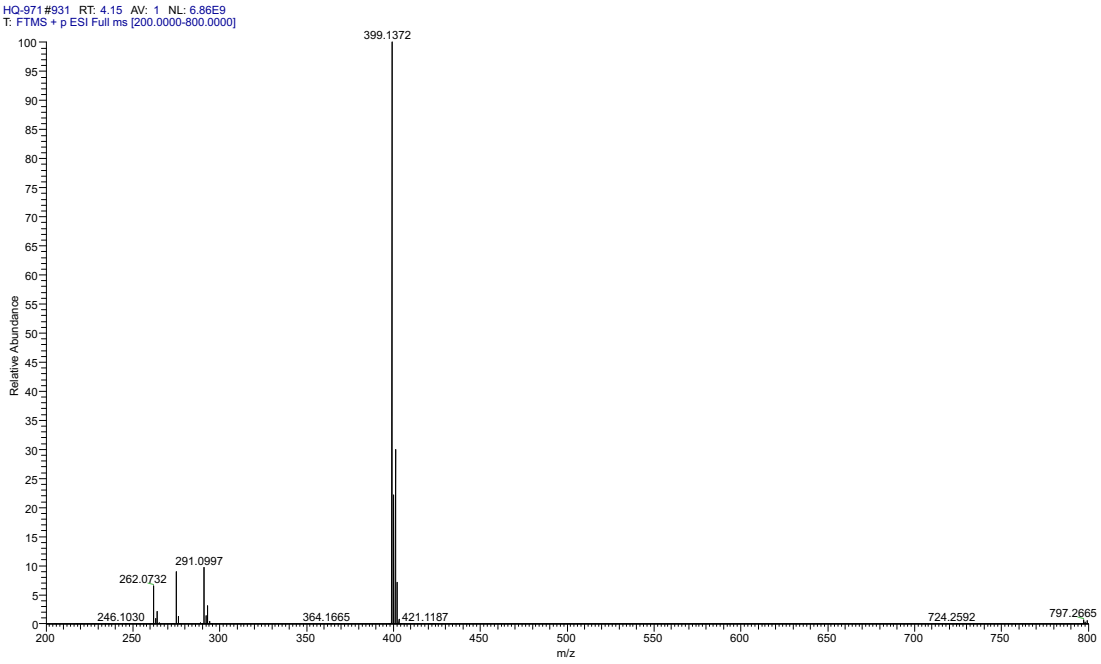
^{19}F NMR of **DC-ClpP-04**



Peak Results

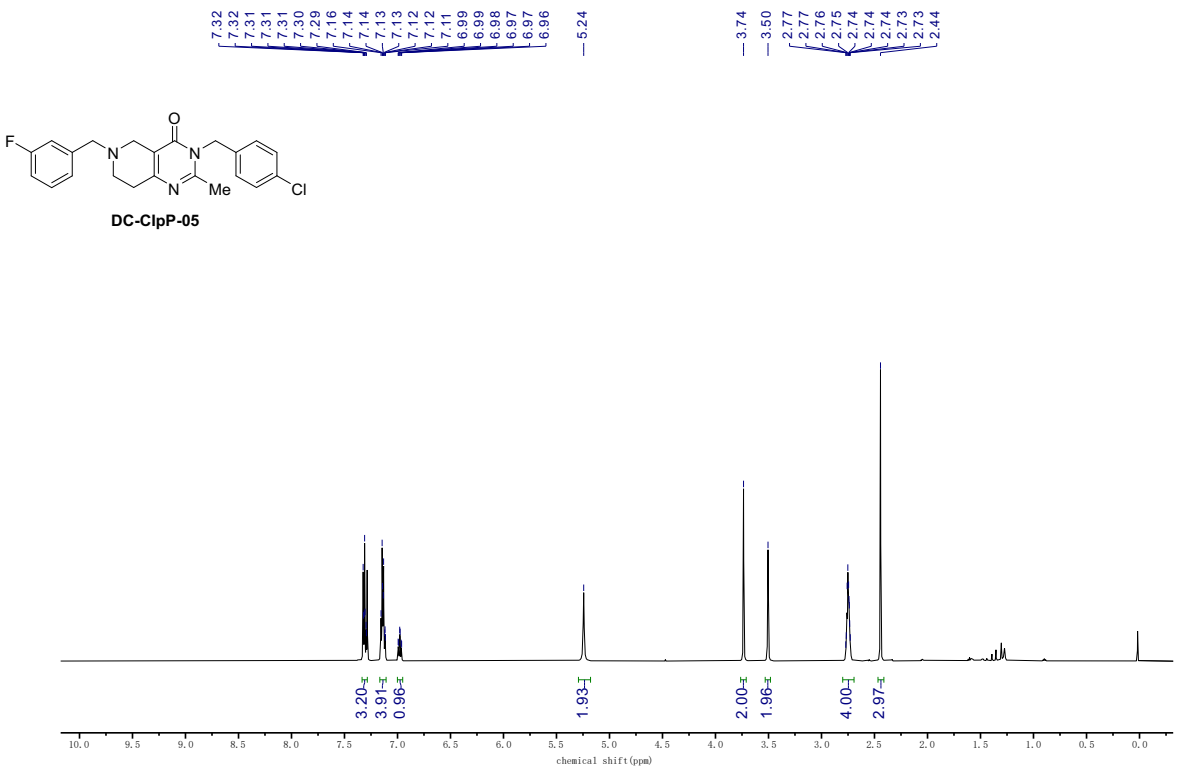
	RT	Area	Height	% Area
1	4.668	16735647	845285	98.27
2	5.617	72625	11363	0.43
3	6.405	207607	51108	1.22
4	7.673	8735	2831	0.05
5	11.727	4831	1684	0.03

HPLC of **DC-ClpP-04**

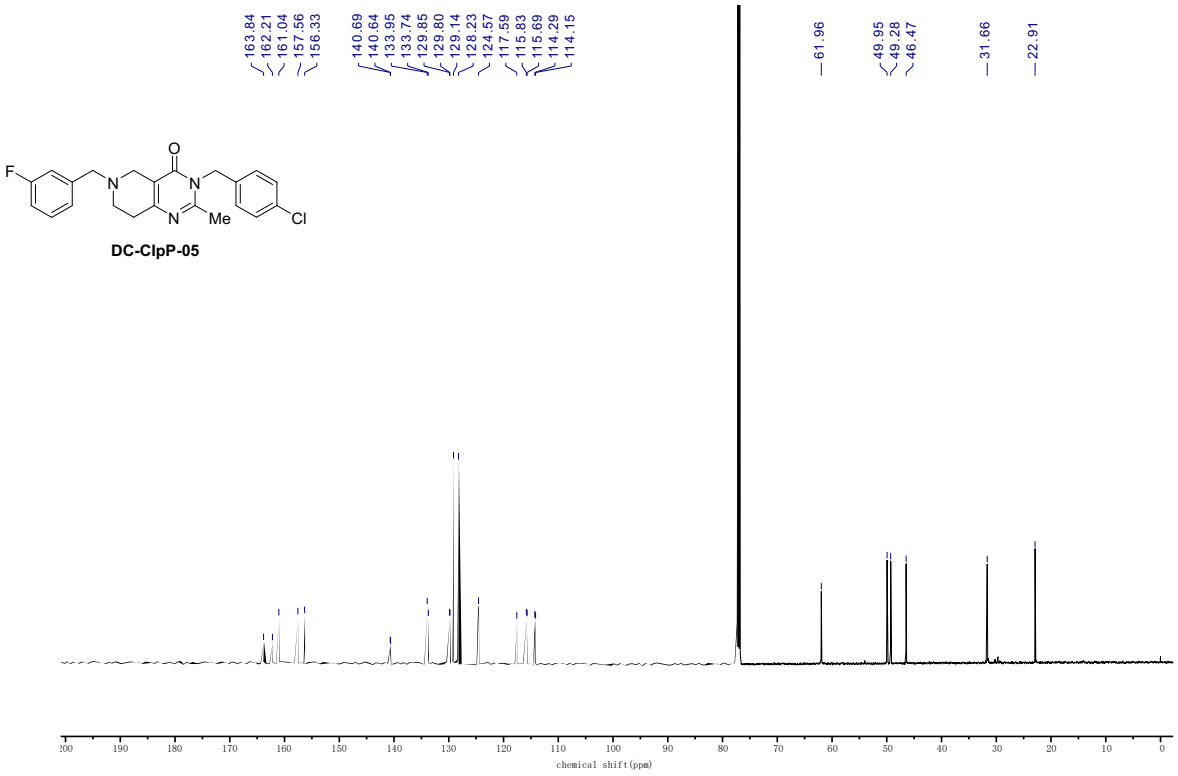


ESI-HRMS of **DC-ClpP-04**

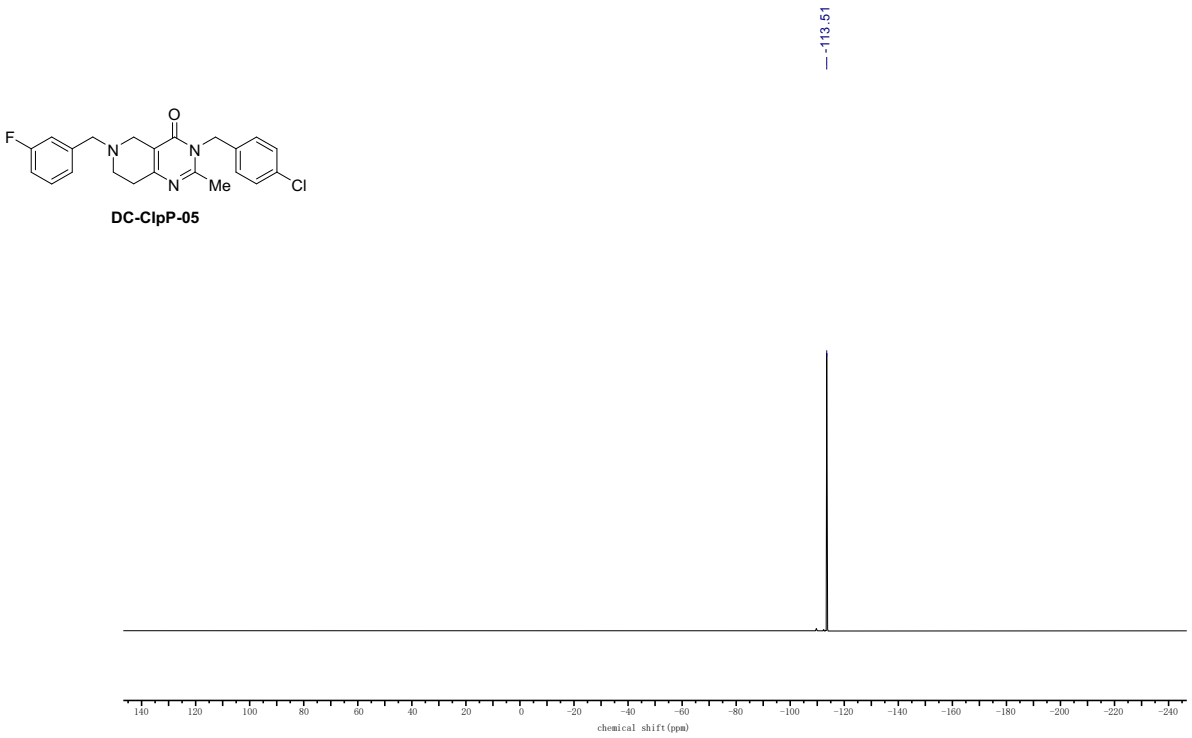
3-(4-chlorobenzyl)-6-(3-fluorobenzyl)-2-methyl-5,6,7,8-tetrahydropyrido[4,3-d]pyrimidin-4(3H)-one (DC-ClpP-05)



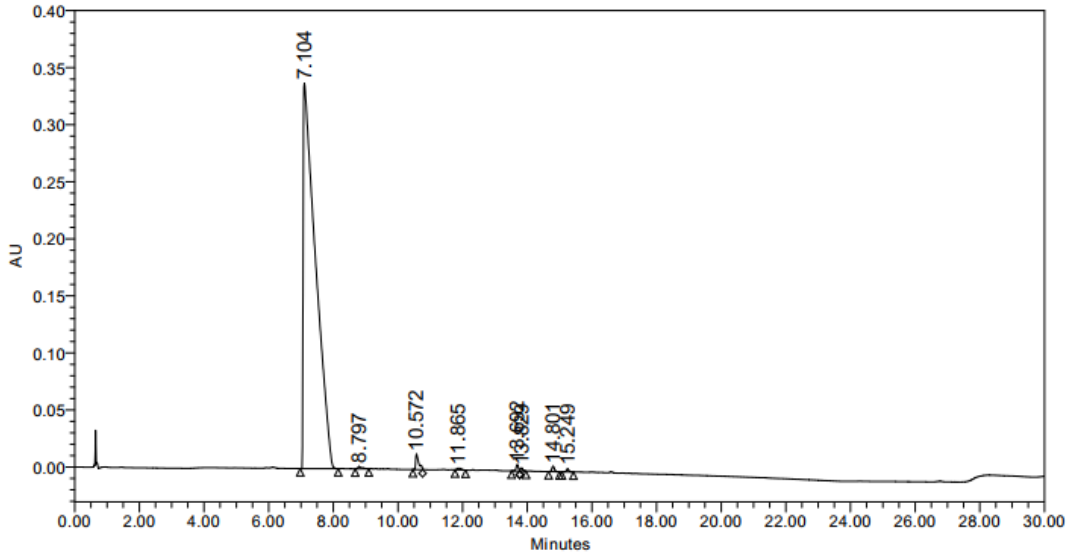
^{1}H NMR of **DC-ClpP-05**



^{13}C NMR of **DC-ClpP-05**



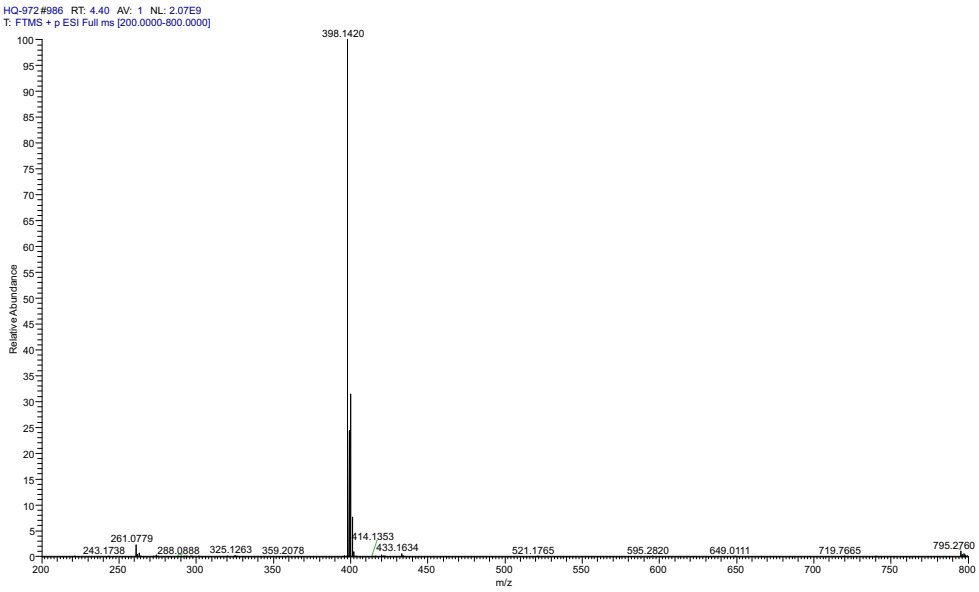
^{19}F NMR of **DC-ClpP-05**



Peak Results

	RT	Area	Height	% Area
1	7.104	8900336	337467	97.79
2	8.797	17314	2103	0.19
3	10.572	89268	13653	0.98
4	11.865	13780	1402	0.15
5	13.692	28340	5542	0.31
6	13.829	12735	2636	0.14
7	14.801	26268	4750	0.29
8	15.249	13693	2652	0.15

HPLC of **DC-ClpP-05**



ESI-HRMS of **DC-ClpP-05**

References

1 Chen, Y. *et al.* Bicyclic pyrimidine compounds as potent IRAK4 inhibitors. *Bioorganic & Medicinal Chemistry Letters* **73**, 128900 (2022). https://doi.org/10.1016/j.bmcl.2022.128900

2 Guan, J. *et al.* DecompDiff: diffusion models with decomposed priors for structure-based drug design. *arXiv preprint arXiv:2403.07902* (2024).