
Advancing Ligand-based Virtual Screening and Molecular Generation with Pretrained Molecular Embedding Distance

Shiyun Wa¹ Yifei Wang¹ Simone Sciabola¹ Ye Wang¹

Abstract

Molecular similarity plays a central role in ligand-based drug discovery, such as virtual screening, analog searching, and goal-directed molecular generation. However, traditional similarity measures, ranging from fingerprint-based Tanimoto coefficients to 3D shape overlays, are often computationally expensive at scale or rely on handcrafted molecular descriptors. Meanwhile, many deep learning approaches to similarity-aware design still depend on similarity-specific supervision or costly data curation, limiting their generality across targets. In this work, we propose pretrained embedding distance (PED) as an effective alternative, computed directly from pretrained molecular models without task-specific training. Experimental results show that PED exhibits distinct correlations with traditional similarity metrics, and performs effectively in both ranking molecules for virtual screening and guiding molecular generation via reward design. These findings suggest that pretrained molecular embeddings capture rich structural information and can serve as a promising and scalable similarity measurement for modern AI-aided drug discovery.

1. Introduction

Molecular similarity occupies a central position in ligand-based drug discovery. This approach assumes that molecules with similar structural and pharmacophoric arrangements are likely to occupy the same binding site and elicit similar biological effects (Hendrickson, 1991; Maggiora et al., 2014). This concept serves as the critical computational engine for a broad class of modern drug discovery workflows such as virtual screening and generative molecular design. Virtual screening serves as a high-throughput retrieval task,

¹Biogen, Cambridge, MA, USA. Correspondence to: Ye Wang <ye.wang@biogen.com>.

identifying promising binders against the therapeutic target of interest from expansive chemical libraries by ranking candidates against known active templates (Klebe, 2006; Gao et al., 2023). Generative design focuses on the *de novo* construction of molecules that satisfy complex multi-objective property and structure constraints (Haddad et al., 2025; Edwards et al., 2026). Despite their different formulations, both tasks are fundamentally governed by the evaluation of molecular similarity. In ligand-based virtual screening, similarity serves as the primary ranking heuristic; in reinforcement learning-based molecular generation, the similarity-derived rewards guide the exploration toward bioactive regions of chemical space. Consequently, the development of an accurate, computationally efficient, and general-purpose similarity function remains a vital research direction in modern drug discovery as it enables the navigation of chemical space at an unprecedented scale.

Traditional measurement requires an appropriate selection of molecular descriptors, similarity coefficients, and weighting schemes (Maldonado et al., 2006). These descriptors serve as the mathematical representation of chemical entities, translating constitutional, configurational, and conformational features into a format suitable for computational analysis (Nikolova & Jaworska, 2003). They span a hierarchy of complexity ranging from 1D physicochemical properties (Raevsky, 2004) and 2D topological fingerprints (Rogers & Hahn, 2010), to sophisticated 3D shape and color¹ alignment such as ROCS (Hawkins et al., 2007) and ROSHAMBO2 (Atwi et al., 2025). However, these approaches face a fundamental trade-off between speed and accuracy. Handcrafted descriptors and 2D fingerprints are often rigid and low-dimensional to capture complex biological mechanisms. 3D approaches, while considered the industry gold standard, require computationally intensive conformer generation and spatial alignment procedures that hinder their application to large-scale molecular libraries.

Recently, machine learning or deep learning approaches have offered a more flexible and efficient definition of molecular similarity, shifting from manual feature engineering to automated representation learning. To achieve this, many

¹Chemical color represents the spatial arrangement of pharmacophore features like hydrogen bond donors and acceptors.

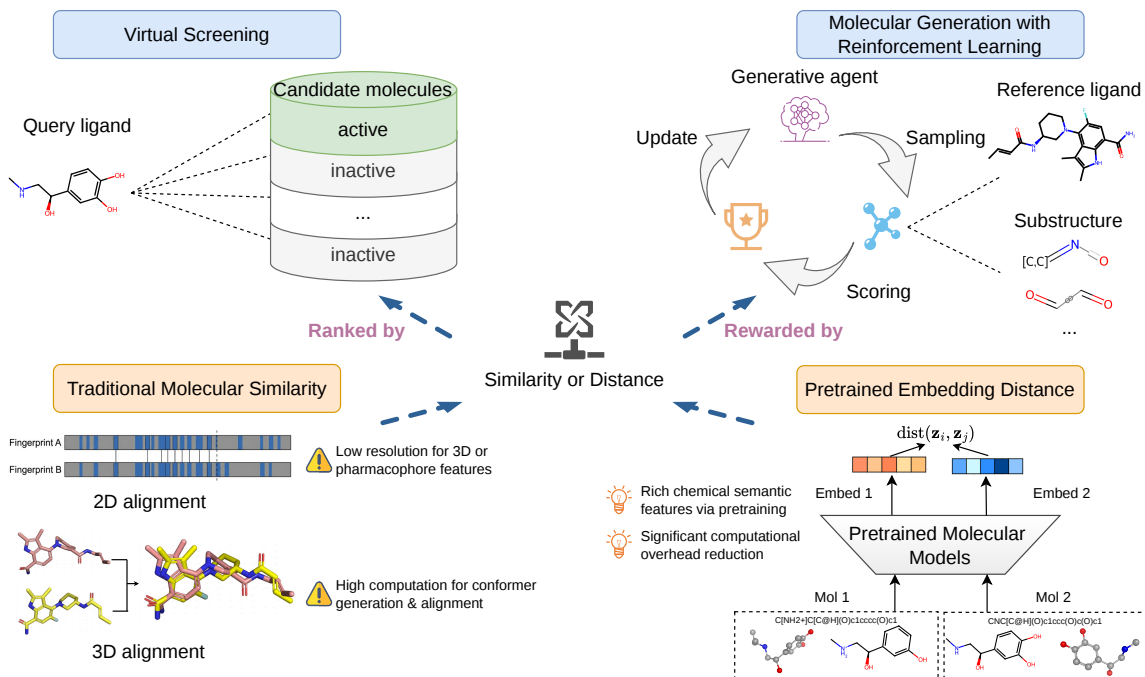


Figure 1. PED-Driven Conceptual Framework for Scalable Ligand-based Drug Discovery. Modern ligand-based drug discovery relies on robust similarity or distance signals to function as scoring functions for prioritizing compounds in virtual screening and as reward for guiding reinforcement learning-based molecular generation. Traditional 2D/3D alignment methods often face inherent representation limitation or prohibitive computational bottlenecks. Instead, this work proposes utilizing Pretrained Embedding Distance (PED) as a highly scalable and computationally efficient alternative. By leveraging embeddings from pretrained models like GeoDiff or MoLFormer, PED captures advanced chemical semantics without explicit alignment, enabling high-resolution discovery at a significantly reduced cost.

methods utilize supervised or semi-supervised frameworks to learn task-specific similarity functions from experimental data. As a general bottleneck of data-driven methods, these models often fail to generalize well due to high cost and scarcity of high-quality molecular annotations. For example, supervised approaches remain constrained by small, target-specific datasets (McNutt et al., 2025; Shah et al., 2025), while semi-supervised or unsupervised approaches such as contrastive learning (Atsango et al., 2022; Gao et al., 2023; Zhou et al., 2024; Wang et al., 2025a) depend on constructing paired data using predefined similarity measures. Consequently, these methods either rely on costly annotations or are still tied to existing similarity tools.

Given these constraints, a compelling alternative is to leverage pretrained molecular representations that have already learned broad chemical knowledge from massive, unlabeled databases. By utilizing these foundational models, we can bypass the need for task-specific labels and instead extract general-purpose embeddings that capture the intrinsic structural features of the chemical space. While this concept has demonstrated good utility through finetuning for downstream tasks such as quantitative structure–activity relationship (QSAR) modeling (Li & Fourches, 2020; Ross et al., 2022; Singh et al., 2026; Burns et al., 2026) and ADMET (Absorption, Distribution, Metabolism, Excretion, and Toxi-

city) prediction (Gu et al., 2024; MacDermott-Opeskin et al., 2026; Xu et al., 2026), its application as a direct similarity function remains underexplored. Specifically, it lacks a comprehensive study of these embedding-derived similarities in guiding ligand-based virtual screening and generative molecular design. Moreover, it is critical to investigate whether these embedding-derived similarities can match or even exceed the utility of computationally intensive 3D alignment methods, offering a more scalable and efficient framework for navigating chemical space. To fill this gap, as shown in Figure 1, this work investigates Pretrained Embedding Distance (PED) derived from two architecturally distinct molecular models: GeoDiff (Xu et al., 2022), a diffusion-based model trained for molecular conformation generation, and MoLFormer (Ross et al., 2022), Transformer-based model trained on Simplified Molecular Input Line Entry System (SMILES) strings (Weininger, 1988) using masked language modeling. Without any task-specific fine-tuning or explicit similarity supervision, we evaluate the utility of PED across three important tasks:

- Correlation analysis between PED and traditional 3D alignment methods.
- Large-scale virtual screening on the challenging LIT-PCBA benchmark (Tran-Nguyen et al., 2020).

- Reinforcement learning-based molecular generation including SMILES-based generation like REINVENT (Li & Tabor, 2023) and synthesizable generation like SynFormer (Gao et al., 2025).

Our results show that PED provides a competitive and flexible similarity signal across both virtual screening and molecular generation, suggesting that pretrained molecular embeddings can serve as a strong and general-purpose alternative to traditional similarity measures.

2. Related Work

Molecular Embedding-derived Similarity. Several studies use learned embedding space to design similarity or distance metric for molecular analog search and virtual screening. For example, DrugCLIP (Gao et al., 2023) reformulates virtual screening as a similarity matching problem between proteins and molecules, aligning protein pocket and molecule representations in a contrastive learning manner, which enables efficient large-scale screening without relying on explicit docking or affinity labels. CHEESE (CHEMical Embeddings Search Engine) (Lžičar & Gamouh, 2024) accelerates virtual screening by embedding complex 3D and electrostatic similarities into a vector space, enabling billion-scale hit identification through rapid approximate nearest neighbor searches. S-MolSearch (Zhou et al., 2024) is the first semi-supervised contrastive learning framework that leverages binding affinity signals for ligand-based virtual screening. It trains a 3D encoder on affinity-paired molecules, then uses it to generate similarity-based soft labels for unlabeled pairs, and train another encoder for large-scale retrieval. In addition, a prior pretrained molecular language model, MoLFormer (Ross et al., 2022), has shown that its embedding space is informative of chemical structural similarity. Specifically, based on the randomly sampled molecule pairs, their embedding similarities correlate with traditional fingerprint-based similarities, and the embedding Euclidean distances correlate with the atom numbers in maximum common subgraph. These studies suggest that embedding similarity can act as an operational measure for retrieval and structure-aware molecular comparison. However, most existing approaches rely on models that are explicitly trained for similarity learning. It remains unclear whether distances in embeddings from non-similarity-specific pretrained molecular models, can reliably capture chemically meaningful similarity, and whether they can further serve as a unified measure for drug discovery tasks.

RL-guided Molecular Generation. Reinforcement learning (RL) has been widely used for molecular generation, where the generator is iteratively updated to maximize the reward function, such as molecular properties (Olivecrona et al., 2017; Li & Tabor, 2023), binding affinity, and

structure-aware similarity (Langevin et al., 2020) to target molecules. In this way, the model updates its policy to favor regions of chemical space that satisfy these complex requirements. For example, REINVENT (Li & Tabor, 2023) is a SMILES-based RL framework that trains a generative agent from a pretrained prior and updates it through iterative sampling and scoring using an augmented likelihood that balances reward and prior constraints. LIMO (Eckmann et al., 2022) are graph-based generative models using RL to improve molecular properties in generation. MOFF (Wang et al., 2025b) is a fragment-based reinforcement learning framework that uses docking score as rewards to generate both covalent and non-covalent inhibitors. In the context of ligand-based generative design, the selection of similarity score function directly decides the quality of generation. Fingerprint-based similarity may fail to capture higher-level structural or functional relationships, while docking-based and 3D similarity rewards are computationally expensive and difficult to scale. This challenge motivates customized reward design like embedding-based similarity as an alternative way for fast and approximated evaluation.

3. Method

3.1. Pretrained Embedding Distance (PED)

We define Pretrained Embedding Distance (PED) as a general similarity measure computed in the representation space of frozen pretrained molecular models. Given a pretrained molecular encoder $f(\cdot)$, a molecule represented by SMILES string s is mapped to a vector embedding:

$$\mathbf{z} = f(s) \in \mathbb{R}^d \quad (1)$$

Given two molecules s_i and s_j (e.g., a molecule and a reference ligand), their similarity is measured via a distance function dist , such as cosine distance or Euclidean distance:

$$D(s_i, s_j) = \text{dist}(f(s_i), f(s_j)) \quad (2)$$

We consider multiple representation modes: **2D**: embeddings derived from molecular graphs or SMILES. **3D**: embeddings derived from 3D conformations. **Concat**: concatenation after ℓ_2 normalization. These variants allow us to analyze how different structural information contributes to embedding-space similarity.

3.2. Instantiations with Pretrained Molecular Models

We compute PED using two representative pretrained models: GeoDiff and MoLFormer, which provide 2D, 3D and sequence-based representations. The model settings are detailed in Appendix A.

GeoDiff. GeoDiff (Xu et al., 2022) is a diffusion-based model for molecular conformation generation. It models

the distribution of molecular conformations by learning to reverse a noise perturbation process. Given a noisy conformation C_t , GeoDiff predicts the noise term $\hat{\epsilon}_\theta$ using a dual-encoder architecture with two corresponding graph field networks (GFN). The dual encoder consists of: a 2D Graph Isomorphism Network (GIN) (Xu et al., 2019) encoding molecular topology, and a 3D equivariant SchNet (Schütt et al., 2017) encoding spatial structure. Overall, the model is trained with a denoising objective:

$$\mathcal{L}_{\text{GeoDiff}} = \sum_{t=1}^T \gamma_t \mathbb{E}_{C_0, \epsilon} \|\epsilon - \hat{\epsilon}_\theta(C_t, t)\|^2 \quad (3)$$

where t is the index for diffusion steps, C_0 is the ground truth conformation, γ_t is a weight.

We use the pretrained model to extract embeddings from the dual-encoder components: (i) 2D embedding from the GIN encoder, (ii) 3D embedding from the SchNet encoder, and (iii) the concat embedding by concatenating both ℓ_2 normalized representations. These embeddings are mean-pooled over atoms from the last layer to obtain fixed-length molecular representations.

MoLFormer. MoLFormer (Ross et al., 2022) is a large-scale chemical language model pretrained on 1.1 billion SMILES sequences using a masked language modeling objective. For a given SMILES string, the model’s Transformer architecture generates a hidden representation for each individual token (atom or symbol). To derive a fixed-length molecular embedding from these token-level outputs, mean-pooling is applied across the last hidden states. This process aggregates the contextual information from the entire sequence into a global representation, capturing the latent chemical patterns and structural motifs learned from large-scale unlabeled data. The resulting high-dimensional vector serves as the basis for computing the corresponding PED.

3.3. Applications of PED in Virtual Screening and Molecular Generation

We evaluate PED on two representative drug discovery tasks: large-scale virtual screening and ligand-based generative molecular design.

3.3.1. VIRTUAL SCREENING

Given a reference ligand and a library of candidate molecules, we compute embeddings for both and rank candidates based on PED in ascending order. Molecules with smaller distances are considered more similar to the reference and are prioritized. We evaluate PED on the LIT-PCBA benchmark, which consists of 15 targets, each with multiple reference ligands and a library of active and inactive candidate molecules. We follow the practice of (Tran-Nguyen

et al., 2020; Atwi et al., 2025) to use the enrichment factor at 1% (EF1%) as the metric, which measures the concentration of retrieved active compounds in the top-ranked subset:

$$\text{EF1\%} = \frac{N_{\text{actives}}^{1\%}}{N^{1\%}} / \frac{N_{\text{actives}}}{N_{\text{total}}} \quad (4)$$

To account for multiple reference ligands per target, we adopt a best-pooled evaluation strategy. For each candidate molecule s_i and a set of reference ligands $\{s_r^{(j)}\}$, we compute the PED with respect to each reference and take the minimum distance:

$$D_{\text{pool}}(s_i) = \min_j D(s_i, s_r^{(j)}) \quad (5)$$

If the approach relies on similarity measures, a maximum similarity would be taken for best-pooling. The pooled distances are then used to rank all candidates, and EF1% is computed based on the resulting ranking using Equation 4. This strategy allows each candidate molecule to give its best chance to match any of the known ligands for the target, better reflecting practical virtual screening settings where multiple reference ligands are available to search a chemical library (Atwi et al., 2025).

3.3.2. MOLECULAR GENERATION

Generative molecular design has emerged as a powerful paradigm for accelerating drug discovery, enabling *in silico* exploration of vast chemical spaces to find molecules with specific structural and functional properties. In ligand-based design, this methodology is typically formulated as an iterative optimization loop centered on a reference template. With the goal of maximizing the desired properties, in each step, the base generative model samples a batch of candidate molecules, which are evaluated by a similarity or distance-based scoring function as reward signals. These rewards are then used to update the model, iteratively shifting the generative distribution toward the chemical space defined by the reference compound. In what follows, we introduce two representative generative models and evaluate how PED serves as an effective scoring function to guide the generation.

REINVENT. REINVENT (Li & Tabor, 2023) utilizes a SMILES-based chemical language model, pretrained on expansive molecular databases, as the foundation for generative design. In ligand-based molecular generation, it is as a reinforcement learning task where the pretrained model acts as a prior to ensure chemical validity. This agent model is finetuned to generate molecules toward the specific reference by optimizing an augmented objective, which increases the probability of high-scoring molecules while regularizing the agent toward the prior distribution. The agent network is

then updated by minimizing the squared difference between its likelihood and this augmented likelihood.

We integrate PED into the REINVENT framework as a scoring component to guide molecular generation toward a reference ligand. A reverse sigmoid function is used to transform PED to a bounded reward in $[0, 1]$, where smaller distance yields higher reward (Equation S1). Compared to cosine distance, raw Euclidean distance is unbounded and exhibits a larger dynamic range, providing a more informative reward signal. Therefore, we prioritize Euclidean PED in our generation experiments.

In practice, the final scoring function combines the PED-based score with additional constraint term that penalizes molecules containing undesirable substructures. Both components are weighted equally, ensuring that the model balances similarity optimization with chemical validity. Detailed experimental settings and hyperparameters are provided in Appendix A.

SynFormer. Besides SMILES-based molecular generation, synthesizable molecular generation is another critical task in AI-aid drug discovery with real-world utility. It reformulates molecular design by first generating its sequential synthetic pathway and then assembling the final molecule. SynFormer (Gao et al., 2025) is a representative architecture, which utilizes a hybrid Transformer and diffusion framework and operates over a joint vocabulary of building block tokens and reaction tokens, generating a discrete sequence that encodes both the components and the assembly order of a synthesizable molecule. It is pretrained on a large collection of synthetic trajectories, learning a prior over chemically valid construction sequences before any task-specific optimization. During generation, SynFormer takes a reference SMILES as input and autoregressively decodes a building block-reaction sequence conditioned on the encoded input. It is further finetuned using a variant of the REINFORCE algorithm (Williams, 1992), a type of RL approach, to guide its generation toward high-scoring molecules.

4. Experiments and Results

4.1. Evaluating PED for Virtual Screening

4.1.1. PED UNDER DIFFERENT MODES CORRELATES WITH CORRESPONDING 3D SIMILARITIES

To validate whether pretrained embeddings capture meaningful chemical and spatial features, we first assess the correlation between PED and traditional 3D score that characterizes molecular similarity through a combination of shape (volumetric overlap) and color (pharmacophore matching). As for the pretrained model used in PED, GeoDiff learns both topological information via a 2D GIN encoder and

spatial structure via a 3D SchNet encoder, we hypothesize that its latent space implicitly aligns with these structural and chemical features. We choose AmpC dataset (Cao et al., 2024) for this correction study, as pre-computed ROCS 3D scores for all molecules are public available. For this analysis, we sample 200k molecules from this dataset, uniformly across ROCS 3D combination similarity bins from 0 to 2 with an interval of 0.2. For each molecule, Color Tanimoto similarity, Shape Tanimoto similarity, and their balanced 50:50 combination score relative to a reference ligand are provided. We compute PED using MoLFormer and GeoDiff under different modes, and compare them with the corresponding ROCS similarities.

Figure 2 shows the Pearson correlations between each pair of all measures. The binned correlation plots for specific pairs can be found in Appendix B (Figure S1 to S7). We observe that GeoDiff 2D PED highly correlates with ROCS color similarity ($r = -0.60$), and GeoDiff 3D PED shows comparable correlation with ROCS shape similarity ($r = -0.60$). It is important to note that these strong negative values indicate high alignment: as the traditional similarity score increases, the PED decreases. When considering mismatched pairs, correlations generally decrease, particularly for 2D PED (2D vs. shape: $r = -0.45$), whereas 3D PED remains relatively stable (3D vs. color: $r = -0.59$). Motivated by the fact that combined similarity often provides stronger signals, we concatenate normalized 2D and 3D embeddings and compute the joint PED. This leads to improved alignment with ROCS combination similarity ($r = -0.67$). Notably, GeoDiff 3D PED consistently demonstrates superior correlation across all ROCS metrics compared to its 2D counterpart. Besides, MoLFormer PED also exhibits relatively high alignment with ROCS combination similarity ($r = -0.63$), and exhibits higher correlation with color than shape ($r = -0.64$ vs. -0.48), which is consistent with its SMILES-based 2D characteristic.

4.1.2. PED OUTPERFORMS TRADITIONAL MEASURES ON VIRTUAL SCREENING BENCHMARKS

We benchmark PED against traditional 3D similarity measures on the LIT-PCBA virtual screening dataset. Figure 3 presents the EF1% distributions for three modes of 3D similarities computed by ROSHAMBO2² (Atwi et al., 2025), three modes of GeoDiff cosine PEDs, and MoLFormer cosine PEDs. Details are reported in Tables S2 and S3.

For targets highlighted in blue (identified as favorable for 3D shape-based screening with EF1% > 2 using ROCS (Hawkins et al., 2007; Tran-Nguyen et al., 2020)), most PED variants achieve strong performance: 7 out of 8 targets

²ROSHAMBO2 provides a robust, open-source alternative for 3D molecular similarity where commercial licensing of ROCS is restricted.

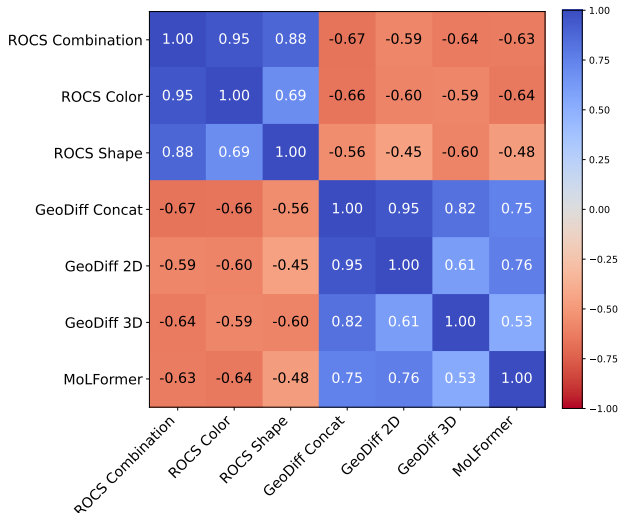


Figure 2. Pearson correlation matrix across ROCS similarities, GeoDiff variant and MoLFormer cosine PEDs. We focus on the top right corner to analyze the correlations between ROCS similarities and PEDs. As smaller distances correspond to higher similarity, the correlation sign is negative.

obtain best-pooled EF1% greater than 2. Notably, PED also performs well on less favorable targets, with 6 out of the remaining 7 targets exceeding the same threshold. Across all targets, MoLFormer cosine PED achieves the highest average mean and best-pooled EF1% (4.53 ± 2.79 and 6.15), followed by 2D ECFP4 similarity (3.94 ± 2.43 and 4.83), as shown in Table S3.

Using Euclidean distance (Figure S8, Tables S4 and S5), the overall trends remain consistent. PED achieves best-pooled EF1% greater than 2 on 7 out of 8 highlighted targets and 6 out of 7 remaining targets. MoLFormer Euclidean PED attains the highest average mean EF1% (4.31 ± 3.11), while GeoDiff 3D Euclidean PED achieves the highest average best-pooled EF1% (5.47), followed by MoLFormer (5.08).

Comparing different modes, ROSHAMBO2 achieves its highest average mean EF1% under the color-only setting, while GeoDiff performs best under the 2D PED setting for both cosine and Euclidean distances. In contrast, for average best-pooled EF1%, ROSHAMBO2 performs best under the combination mode, whereas GeoDiff achieves its highest values under the 3D PED setting.

4.2. PED-Guided Molecular Generation

To evaluate the impact of different embedding-based similarities in guiding molecular generation, we conduct a targeted case study using two generative frameworks, REINVENT and SynFormer, focused on the generation of novel analogs using a reference compound BMS-986195, a known Bruton’s tyrosine kinase (BTK) inhibitor (Watterson et al.,

2019). We compare three scoring regimes: ROSHAMBO2, which utilizes the experimental 3D coordinates from the reference’s crystallographic structure (PDB: 6O8I); transformed GeoDiff and MoLFormer Euclidean PEDs.

Training dynamics. Figure S9 shows that REINVENT with all methods achieves stable convergence with high valid SMILES rates and exhibits distinct optimization behaviors. ROSHAMBO2 converges more slowly, whereas GeoDiff and MoLFormer-based PEDs lead to faster convergence of the reward. Meanwhile, SynFormer converges fast with a strong pretrained encoder for the reference compound (Figure S10), so we only train it for fewer steps, resulting in a substantially smaller number of generated molecules compared to REINVENT (Table S6, Appendix E). Therefore, we focus on comparing different scoring methods within each framework. We also observe significant differences in computational efficiency. 3D-based scoring via ROSHAMBO2 exhibits the slowest convergence and highest computational overhead due to required conformer generation. In contrast, PED-based methods provide a substantial speed-up. According to Table S6, in the REINVENT experiments, GeoDiff and MoLFormer-based PEDs achieve a $1.5\times$ and $3.3\times$ speed-up, respectively, compared to the ROSHAMBO2 baseline. These trends are consistently observed in the SynFormer experiments, where both GeoDiff and MoLFormer approximately double the sampling speed ($\sim 2\times$ faster) relative to ROSHAMBO2. These results underscore the computational advantage of leveraging learned representations over explicit spatial optimization for speeding up molecular generation.

Scaffold diversity of generated molecules. Figure 4 shows that scaffold diversity of REINVENT-generated molecules consistently decreases as total score increases across all methods, indicating that higher-scoring molecules are more structurally concentrated. This trend is most pronounced for MoLFormer and GeoDiff Concat, suggesting stronger convergence. MoLFormer and GeoDiff 3D maintain comparatively higher diversity in the early stages, exhibiting strong exploration in low-score regions, while MoLFormer shifts toward more focused exploitation in high-score regions. In contrast, SynFormer exhibits a different pattern. GeoDiff 2D achieves the highest scaffold diversity across most quantiles, particularly in Q2, indicating strong exploration behavior. The scaffold diversity of GeoDiff Concat increases toward higher-score regions. However, other methods show a decrease trend of scaffold diversity along with the total score, which are similar to REINVENT.

Table 1 further summarizes the unique scaffold count and ratio within the top-5,000 molecules. This corresponds to the highest score quantile (Q4) of REINVENT, while it covers most of the SynFormer-generated unique molecules due




Figure 3. LIT-PCBA virtual screening performance (EF1%) obtained by ROSHAMBO2 similarities, GeoDiff and MoLFormer cosine PEDs. Different modes of ROSHAMBO2 similarities are in green, while those for GeoDiff PEDs are in blue. Boxplots illustrate the distribution of EF1% values across query ligands for each target. Circular markers indicate the best-pooled EF1%, while cross markers denote the outliers of the boxplots. In case where there is only one reference ligand, the boxes would be a single line and the best-pooled circle overlaps with it. Targets highlighted in a blue background are those identified as favorable for 3D shape-based screening (EF1% > 2 with ROCS) in the original LIT-PCBA benchmark. Dashed horizontal red and blue lines correspond to EF1% of 1 and 2, respectively.

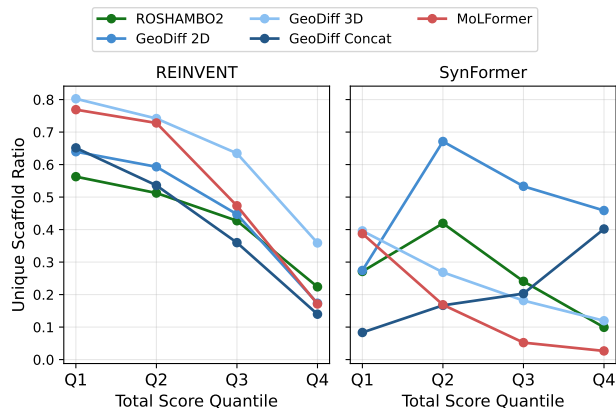


Figure 4. Scaffold diversity across total score quantiles. For each method, molecules are grouped into four equal-frequency bins based on total score (Q1–Q4, from low to high). The y-axis shows the ratio of unique ring-linker scaffold to the total number of molecules within each bin.

to the smaller sample size. Within REINVENT, GeoDiff 3D achieves the highest scaffold diversity (35.16%), followed by MoLFormer (12.84%), while other methods remain lower. In SynFormer, GeoDiff 2D shows the highest scaffold ratio (46.36%), substantially exceeding other methods, but MoLFormer produces the lowest diversity (5.04%). Despite differences in scale, both frameworks indicate that embedding-based scoring can significantly affect the diversity of high-scoring molecules.

Drug-likeness of top-5,000 generated molecules. Table S7 and S8 in Appendix F show drug-likeness statistics of

Table 1. Unique ring-linker scaffold and its ratio over top-5,000 high-score generated molecules.

Model	Method	Unique scaffold	Unique scaffold ratio (†)
REINVENT	ROSHAMBO2	397	7.94%
	GeoDiff 2D	379	7.58%
	GeoDiff 3D	1,758	35.16%
	GeoDiff Concat	301	6.02%
	MoLFormer	642	12.84%
SynFormer	ROSHAMBO2	1,041	20.82%
	GeoDiff 2D	2,318	46.36%
	GeoDiff 3D	624	12.48%
	GeoDiff Concat	1,046	20.92%
	MoLFormer	252	5.04%

top-5,000 high-score molecules. REINVENT consistently achieves strong compliance with good drug-like property ranges across all methods, which can be attributed to the unfavorable substructure filters in the reward function. In contrast, SynFormer exhibits much larger variation in physicochemical properties due to its building-block-based generation process, which allows more substantial structural modifications. This effect is particularly evident for methods with higher scaffold diversity, such as GeoDiff 2D and GeoDiff Concat, where molecules frequently fall outside desirable ranges in molecular weight, TPSA, LogP, and QED. These results highlight a trade-off between structural diversity and drug-likeness across different generative frameworks.

Binding affinity prediction of top-score generated molecules by Boltz-2. To validate the biological relevance of our generative results, we select the top-1,000 molecules ranked by the total score for each method to assess their predicted potency against the BTK target using Boltz-2 (Pas-

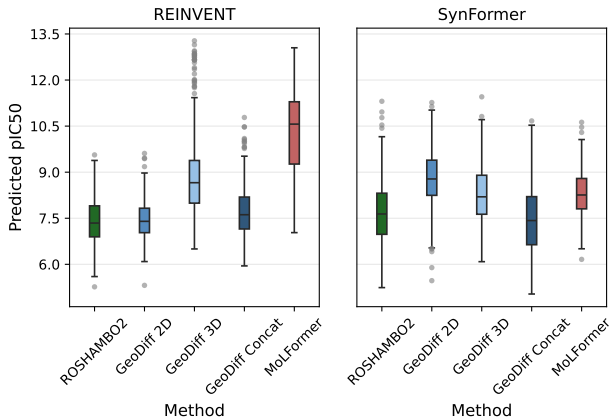


Figure 5. Predicted pIC50 distribution of scaffold-balanced top-ranked generated molecules. For each method, the top-1,000 molecules are first grouped by scaffold. We then uniformly sample 500 molecules across scaffolds for each method.

saro et al., 2025). Boltz-2 is a particularly reliable evaluator for this task as it has demonstrated high fidelity on the kinase family (Passaro et al., 2025; Cecchini & Sinenka, 2026). We focus on the pIC50 metric, where a higher value indicates lower IC50 concentration and thus higher potency, to assess whether molecules favored by different reward signals are also predicted to exhibit strong predicted inhibition. To reduce bias from over-represented scaffolds, we further construct a scaffold-balanced subset by grouping the top-1,000 molecules by ring-linker scaffold and uniformly sampling up to 500 molecules per method.

As shown in Figure 5 and Table 2, the ranking of reward signals differs across the two generative frameworks. For REINVENT, MoLFormer yields the highest predicted pIC50 distribution (10.27 ± 1.34 , effect size $\Delta = 0.92$), followed by GeoDiff 3D (8.83 ± 1.29 , $\Delta = 0.71$). GeoDiff Concat shows only moderate improvement (7.70 ± 0.79 , $\Delta = 0.20$), while ROSHAMBO2 and GeoDiff 2D remain lower and close to each other (7.40 ± 0.69 vs. 7.44 ± 0.60 , $\Delta = 0.04$). For SynFormer, the overall differences are less significant, but embedding-based rewards still outperform ROSHAMBO2 in most cases. GeoDiff 2D achieves the highest predicted pIC50 (8.81 ± 0.87 , $\Delta = 0.61$), followed by MoLFormer (8.31 ± 0.74 , $\Delta = 0.39$) and GeoDiff 3D (8.28 ± 0.90 , $\Delta = 0.34$), whereas GeoDiff Concat performs slightly worse than ROSHAMBO2.

Overall, these results suggest that molecules generated with the guidance of PED-based reward signals are more likely to bind the target protein pocket, although the strongest PED mode varies with the generative framework. Specifically, REINVENT benefits most from MoLFormer and GeoDiff 3D, whereas SynFormer performs well with GeoDiff 2D and MoLFormer. The top-5 generated molecules ranked by total score with predicted pIC50 are shown in Appendix

Table 2. Predicted pIC50 statistics of scaffold-balanced top-500 molecules for REINVENT and SynFormer. Cliff’s delta is reported as an effect size, measuring how often the predicted pIC50 values from a PED-based method exceed those from ROSHAMBO2.

Model	Method	Predicted pIC50 (\uparrow)	Effect size
REINVENT	ROSHAMBO2	7.40 ± 0.69	—
	GeoDiff 2D	7.44 ± 0.60	0.04
	GeoDiff 3D	8.83 ± 1.29	0.71
	GeoDiff Concat	7.70 ± 0.79	0.20
	MoLFormer	10.27 ± 1.34	0.92
SynFormer	ROSHAMBO2	7.69 ± 1.04	—
	GeoDiff 2D	8.81 ± 0.87	0.61
	GeoDiff 3D	8.28 ± 0.90	0.34
	GeoDiff Concat	7.48 ± 1.06	-0.11
	MoLFormer	8.31 ± 0.74	0.39

Figure S11 and S12.

5. Conclusion

This study establishes a connection between similarity-based virtual screening and molecular generation via pretrained embedding distance (PED). Our findings show that PEDs derived from models like GeoDiff and MoLFormer exhibit strong correlations with industry-standard 3D similarity metrics, suggesting that pretrained molecular encoders implicitly capture rich spatial and pharmacophoric information without requiring explicit 3D alignment. This signal performs effectively across both virtual screening and reinforcement learning-based generation, providing a reliable ranking function for the challenging LIT-PCBA benchmark and a continuous reward signal for frameworks like REINVENT and SynFormer. Most notably, PED offers a substantial leap in computational efficiency, achieving up to a 3.3 \times speed-up in generative sampling by bypassing the need for expensive conformer generation and iterative spatial optimization.

Impact Statement

This paper demonstrates that pretrained molecular embeddings can be repurposed as efficient similarity measurement to significantly accelerate early-stage drug discovery. While this work focuses on a BTK inhibitor case study in molecular generation task and evaluates the generation compound using Boltz-2, the biological efficacy and safety of the generated compounds are still not wet-lab validated. Future work involving experimental synthesis and assays is required to confirm the therapeutic potential of the identified leads.

References

- Atsango, A., Diamant, N. L., Lu, Z., Biancalani, T., Scalia, G., and Chuang, K. V. A 3d-shape similarity-based contrastive approach to molecular representation learning. In *NeurIPS 2022 AI for Science: Progress and Promises*, 2022. URL <https://openreview.net/forum?id=1T3jXy8hPKI>.
- Atwi, R., Farr, S., Wang, Y., Antoszewski, A., and Sciabola, S. Roshambo2: Accelerating molecular alignment for large chemical libraries with gpu optimization and algorithmic advances. *Journal of Chemical Information and Modeling*, 65(19):9842–9849, 2025. doi: 10.1021/acs.jcim.5c01322. URL <https://doi.org/10.1021/acs.jcim.5c01322>. PMID: 40990368.
- Axelrod, S. and Gómez-Bombarelli, R. Geom, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data*, 9(1):185, Apr 2022. ISSN 2052-4463. doi: 10.1038/s41597-022-01288-4. URL <https://doi.org/10.1038/s41597-022-01288-4>.
- Burns, J. W., Zalte, A. S., Abreu, C. R. A., Sieg, J., Feldmann, C., Mathea, M., and Green, W. H. Deep learning foundation models from classical molecular descriptors, 2026. URL <https://arxiv.org/abs/2506.15792>.
- Cao, Z., Sciabola, S., and Wang, Y. Large-scale pre-training improves sample efficiency of active learning-based virtual screening. *Journal of Chemical Information and Modeling*, 64(6):1882–1891, 2024. doi: 10.1021/acs.jcim.3c01938. URL <https://doi.org/10.1021/acs.jcim.3c01938>. PMID: 38442000.
- Cecchini, M. and Sinenka, H. Rise of ai technologies in virtual screening. *Journal of Chemical Information and Modeling*, 2026. doi: 10.1021/acs.jcim.6c00877. URL <https://doi.org/10.1021/acs.jcim.6c00877>.
- Eckmann, P., Sun, K., Zhao, B., Feng, M., Gilson, M. K., and Yu, R. Limo: Latent inceptionism for targeted molecule generation, 2022. URL <https://arxiv.org/abs/2206.09010>.
- Edwards, C., Han, C., Lee, G., Nguyen, T., Szymkuć, S., Prasad, C. K., Jin, B., Han, J., Diao, Y., Liu, G., Peng, H., Grzybowski, B. A., Burke, M. D., and Ji, H. mCLM: A modular chemical language model that generates functional and makeable molecules. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=r2HG3xOMJI>.
- Gao, B., Qiang, B., Tan, H., Jia, Y., Ren, M., Lu, M., Liu, J., Ma, W.-Y., and Lan, Y. DrugCLIP: Contrastive protein-molecule representation learning for virtual screening. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=1AbCgNcxm7>.
- Gao, W., Luo, S., and Coley, C. W. Generative ai for navigating synthesizable chemical space. *Proceedings of the National Academy of Sciences*, 122(41):e2415665122, 2025. doi: 10.1073/pnas.2415665122. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2415665122>.
- Grygorenko, O. O., Radchenko, D. S., Dziuba, I., Chuprina, A., Gubina, K. E., and Moroz, Y. S. Generating multi-billion chemical space of readily accessible screening compounds. *iScience*, 23(11):101681, 2020. ISSN 2589-0042. doi: <https://doi.org/10.1016/j.isci.2020.101681>. URL <https://www.sciencedirect.com/science/article/pii/S2589004220308737>.
- Gu, Y., Yu, Z., Wang, Y., Chen, L., Lou, C., Yang, C., Li, W., Liu, G., and Tang, Y. admetsar3.0: a comprehensive platform for exploration, prediction and optimization of chemical admet properties. *Nucleic Acids Research*, 52(W1):W432–W438, 07 2024. ISSN 0305-1048. doi: 10.1093/nar/gkae298. URL <https://doi.org/10.1093/nar/gkae298>.
- Haddad, R., Litsa, E. E., Liu, Z., Yu, X., Burkhardt, D., and Bhisetti, G. Targeted molecular generation with latent reinforcement learning. *Scientific Reports*, 15(1):15202, Apr 2025. ISSN 2045-2322. doi: 10.1038/s41598-025-99785-0. URL <https://doi.org/10.1038/s41598-025-99785-0>.
- Hawkins, P. C. D., Skillman, A. G., and Nicholls, A. Comparison of shape-matching and docking as virtual screening tools. *Journal of Medicinal Chemistry*, 50(1):74–82, 2007. doi: 10.1021/jm0603365. URL <https://doi.org/10.1021/jm0603365>. PMID: 17201411.
- Hendrickson, J. B. Concepts and applications of molecular similarity. *Science*, 252(5009):1189–1190, 1991.
- Huang, A., Knight, I. S., and Naprienko, S. Data leakage and redundancy in the lit-pcba benchmark, 2025. URL <https://arxiv.org/abs/2507.21404>.
- Klebe, G. Virtual ligand screening: strategies, perspectives and limitations. *Drug Discovery Today*, 11(13):580–594, 2006. ISSN 1359-6446. doi: <https://doi.org/10.1016/j.drudis.2006.05.012>. URL <https://www.sciencedirect.com/science/article/pii/S1359644606001784>.

- Langevin, M., Minoux, H., Levesque, M., and Bianciotto, M. Scaffold-constrained molecular generation. *Journal of Chemical Information and Modeling*, 60(12):5637–5646, 2020. doi: 10.1021/acs.jcim.0c01015. URL <https://doi.org/10.1021/acs.jcim.0c01015>. PMID: 33301333.
- Li, C.-H. and Tabor, D. P. Generative organic electronic molecular design informed by quantum chemistry. *Chem. Sci.*, 14:11045–11055, 2023. doi: 10.1039/D3SC03781A. URL <http://dx.doi.org/10.1039/D3SC03781A>.
- Li, X. and Fourches, D. Inductive transfer learning for molecular activity prediction: Next-gen qsar models with molpmpfit. *Journal of Cheminformatics*, 12(1):27, Apr 2020. ISSN 1758-2946. doi: 10.1186/s13321-020-00430-x. URL <https://doi.org/10.1186/s13321-020-00430-x>.
- Lžičař, M. and Gamouh, H. Cheese: 3d shape and electrostatic virtual screening in a vector space. *ChemRxiv*, 2024(1105), 2024. doi: 10.26434/chemrxiv-2024-cswth. URL <https://chemrxiv.org/doi/abs/10.26434/chemrxiv-2024-cswth>.
- MacDermott-Opeskin, H., Scheen, J., Wognum, C., Horton, J. T., West, D., Payne, A. M., Castellanos, M. A., Colby, S., Griffen, E., Cousins, D., Stacey, J., Reid, L., Aschenbrenner, J. C., Fearon, D., Balcomb, B., Marples, P., Tomlinson, C. W. E., Lithgo, R., Godoy, A. S., Winokan, M., Barr, H., Lahav, N., Lavi, M., Duberstein, S., Cohen, G., Fate, G., Lefker, B., Robinson, R., Szommer, T., Lynch, N., Minh, D. D. L., La, V. N. T., Kang, L., Huddleston, K., Renslow, R., Tollefson, M., Walters, W. P., Xu, C., Hsu, J., St-Laurent, J., Etsmoberg, H., Zhu, L., Quirke, A., Abdul Haleem, M. I., Alibay, I., Baid, G., Birnbaum, B., Bishop, K. P., Bohorquez, H., Bose, A., Brown, C. J., Burns, J., Cai, L., Cedeno, R., de Cesco, S., Chupakhin, V., Clark, F., Cole, D. J., Corbi-Verge, C., Darnial, M., Davi, A., Dehaen, W., Doering, N. P., Dougha, A., Dréanic, M.-P., Eakin, B., Ehrlich, A., Elijosius, R., Fülöp, J., Gitter, A., Goossens, K., Gu, Y., Head-Gordon, T., Hoffer, L., Hofmans, J., Jiang, E., Kaminow, B., Khosravi, S., Khoualdi, A. F., Lenselink, E. B., Liu, Z., Liu, Y., Liu, S., Ma, Y., Maher, P., Mayer, I., Mendez-Lucio, O., Mey, A. S. J. S., Michel, J., Montanari, F., Niu, T., Ogino, R., Palaniappan, A., Pan, X., Patnaik, A., Pham, L.-H., Pinto, L., Purnomo, J., Rich, A., Schaaf, L., Schran, C., Singh, R. K., Srilakshmi, M., Srivastava, S. P., Sun, K., Sun, Z., Talagayev, V., Thirukonda Subramanian Balakrishnan, B., Titus, I., Tkatchenko, A., Treyde, W., Tricarico, G., Tripp, A., Vithayapalart, N., Wang, Y., Wasi, A. T., Wedig, S., Wolber, G., Xu, B., Zhou, W., von Delft, F., Lee, A., Kirkegaard, K., Sjö, P., Fraser, J. S., and Chodera, J. D. A computational community blind challenge on pan-coronavirus drug discovery data. *Journal of Chemical Information and Modeling*, 66(6):3129–3149, 2026. doi: 10.1021/acs.jcim.5c02106. URL <https://doi.org/10.1021/acs.jcim.5c02106>. PMID: 41749419.
- Maggiora, G., Vogt, M., Stumpfe, D., and Bajorath, J. Molecular similarity in medicinal chemistry. *Journal of Medicinal Chemistry*, 57(8):3186–3204, 2014. doi: 10.1021/jm401411z. URL <https://doi.org/10.1021/jm401411z>. PMID: 24151987.
- Maldonado, A. G., Doucet, J. P., Petitjean, M., and Fan, B.-T. Molecular similarity and diversity in chemoinformatics: From theory to applications. *Molecular Diversity*, 10(1):39–79, Feb 2006. ISSN 1573-501X. doi: 10.1007/s11030-006-8697-1. URL <https://doi.org/10.1007/s11030-006-8697-1>.
- McNutt, A. T., Li, Y., Meli, R., Aggarwal, R., and Koes, D. R. Gnina 1.3: the next increment in molecular docking with deep learning. *Journal of Cheminformatics*, 17(1):28, Mar 2025. ISSN 1758-2946. doi: 10.1186/s13321-025-00973-x. URL <https://doi.org/10.1186/s13321-025-00973-x>.
- Nikolova, N. and Jaworska, J. Approaches to measure chemical similarity – a review. *QSAR & Combinatorial Science*, 22(9-10):1006–1026, 2003. doi: <https://doi.org/10.1002/qsar.200330831>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/qsar.200330831>.
- Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9(1):48, Sep 2017. ISSN 1758-2946. doi: 10.1186/s13321-017-0235-x. URL <https://doi.org/10.1186/s13321-017-0235-x>.
- Passaro, S., Corso, G., Wohlwend, J., Reveiz, M., Thaler, S., Somnath, V. R., Getz, N., Portnoi, T., Roy, J., Stark, H., Kwabi-Addo, D., Beaini, D., Jaakkola, T., and Barzilay, R. Boltz-2: Towards accurate and efficient binding affinity prediction. *bioRxiv*, 2025. doi: 10.1101/2025.06.14.659707. URL <https://www.biorxiv.org/content/early/2025/06/18/2025.06.14.659707>.
- Raevsky, O. A. Physicochemical descriptors in property-based drug design. *Mini-Reviews in Medicinal Chemistry*, 4(10):1041–1052, 2004. ISSN 1389-5575/1875-5607. doi: 10.2174/1389557043402964. URL <https://www.eurekaselect.com/article/25208>.

- Rogers, D. and Hahn, M. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010. doi: 10.1021/ci100050t. URL <https://doi.org/10.1021/ci100050t>. PMID: 20426451.
- Ross, J., Belgodere, B., Chenthamarakshan, V., Padhi, I., Mroueh, Y., and Das, P. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(12):1256–1264, Dec 2022. ISSN 2522-5839. doi: 10.1038/s42256-022-00580-7. URL <https://doi.org/10.1038/s42256-022-00580-7>.
- Schütt, K., Kindermans, P.-J., Saucedo Felix, H. E., Chmiela, S., Tkatchenko, A., and Müller, K.-R. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/303ed4c69846ab36c2904d3ba8573050-Paper.pdf.
- Shah, P. M., Zhu, H., Lu, Z., Wang, K., Tang, J., and Li, M. Deepdtagen: a multitask deep learning framework for drug-target affinity prediction and target-aware drugs generation. *Nature Communications*, 16(1):5021, May 2025. ISSN 2041-1723. doi: 10.1038/s41467-025-59917-6. URL <https://doi.org/10.1038/s41467-025-59917-6>.
- Singh, R., Barsainyan, A. A., Irfan, R., Amorin, C. J., He, S., Davis, T., Thiagarajan, A., Sankaran, S., Chithrananda, S., Ahmad, W., Jones, D., McLoughlin, K., Kim, H., Bhutani, A., Sathyanarayana, S. V., Viswanathan, V., Allen, J. E., and Ramsundar, B. Chemberta-3: an open source training framework for chemical foundation models. *Digital Discovery*, 5:662–685, 2026. doi: 10.1039/D5DD00348B. URL <http://dx.doi.org/10.1039/D5DD00348B>.
- Tran-Nguyen, V.-K., Jacquemard, C., and Rognan, D. Litpcba: An unbiased data set for machine learning and virtual screening. *Journal of Chemical Information and Modeling*, 60(9):4263–4273, 2020. doi: 10.1021/acs.jcim.0c00155. URL <https://doi.org/10.1021/acs.jcim.0c00155>. PMID: 32282202.
- Wang, Y., Zhao, H., Sciabola, S., and Wang, W. cmolpvt: A conditional generative pre-trained transformer for target-specific de novo molecular generation. *Molecules*, 28(11), 2023. ISSN 1420-3049. doi: 10.3390/molecules28114430. URL <https://www.mdpi.com/1420-3049/28/11/4430>.
- Wang, Y., Li, Y., Liu, L., Hong, P., and Xu, H. Advancing drug discovery with enhanced chemical understanding via asymmetric contrastive multimodal learning. *Journal of Chemical Information and Modeling*, 65(13):6547–6557, 2025a. doi: 10.1021/acs.jcim.5c00430. URL <https://doi.org/10.1021/acs.jcim.5c00430>. PMID: 40548496.
- Wang, Y., Wang, Z., Li, Y., Yan, P., and Li, X. Improving covalent and noncovalent molecule generation via reinforcement learning with functional fragments. *Journal of Chemical Information and Modeling*, 65(12):5934–5944, 2025b. doi: 10.1021/acs.jcim.5c00944. URL <https://doi.org/10.1021/acs.jcim.5c00944>. PMID: 40471710.
- Watterson, S. H., Liu, Q., Beaudoin Bertrand, M., Batt, D. G., Li, L., Pattoli, M. A., Skala, S., Cheng, L., Obermeier, M. T., Moore, R., Yang, Z., Vickery, R., Elzinga, P. A., Discenza, L., D’Arienzo, C., Gillooly, K. M., Taylor, T. L., Pulicicchio, C., Zhang, Y., Heimrich, E., McIntyre, K. W., Ruan, Q., Westhouse, R. A., Catlett, I. M., Zheng, N., Chaudhry, C., Dai, J., Galella, M. A., Tebben, A. J., Pokross, M., Li, J., Zhao, R., Smith, D., Rampulla, R., Allentoff, A., Wallace, M. A., Mathur, A., Salter-Cid, L., Macor, J. E., Carter, P. H., Fura, A., Burke, J. R., and Tino, J. A. Discovery of branebrutinib (bms-986195): A strategy for identifying a highly potent and selective covalent inhibitor providing rapid in vivo inactivation of bruton’s tyrosine kinase (btk), 2019. URL <https://doi.org/10.1021/acs.jmedchem.9b00167>. PMID: 30893553.
- Weininger, D. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988. doi: 10.1021/ci00057a005. URL <https://doi.org/10.1021/ci00057a005>.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, May 1992. ISSN 1573-0565. doi: 10.1007/BF00992696. URL <https://doi.org/10.1007/BF00992696>.
- Xu, B., Lu, Y., Li, C., Yue, L., Wang, X., Fu, T., Shen, M., and Chen, L. Smiles-mamba: Chemical mamba foundation models for drug admet prediction, 2026. URL <https://arxiv.org/abs/2408.05696>.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks?, 2019. URL <https://arxiv.org/abs/1810.00826>.
- Xu, M., Yu, L., Song, Y., Shi, C., Ermon, S., and Tang, J. Geodiff: A geometric diffusion model for molecular

conformation generation. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=PzcvxEMzvQC>.

Zhou, G., Wang, Z., Yu, F., Ke, G., Wei, Z., and Gao, Z. S-molsearch: 3d semi-supervised contrastive learning for bioactive molecule search. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=wJAF8TGVUG>.

A. Experiment Setting and Hyperparameters

Pretrained Molecular Models. We used GeoDiff (Xu et al., 2022) and MoLFormer (Ross et al., 2022) models as two representative pretrained molecular models to compute PED. GeoDiff was pretrained on the GEOM-Drugs (Axelrod & Gómez-Bombarelli, 2022) dataset, and MoLFormer was the XL version (<https://huggingface.co/ibm-research/MoLFormer-XL-both-10pct>) pretrained on 1.1B molecules from ZINC and PubChem (10% of both datasets). Both 2D GIN and 3D SchNet encoders of GeoDiff generate 128-length embeddings, the size of the concatenated embedding from the dual encoder is 256. The size of MoLFormer’s embedding is 768.

SMILES-based Generative Model. We used REINVENT 3.2 (<https://github.com/MolecularAI/Reinvent>) with a pretrained prior and agent initialized from the same checkpoint, and optimized the agent for 10,000 RL steps with batch size 128 and learning rate 1×10^{-4} . Each run generates 1,280,000 molecules (including duplicates). The scoring function was a equally weighted sum of two components—custom substructure alerts and sigmoid-transformed PED. An inception mechanism (memory size 100, sample size 10) was used to replay high-scoring molecules. The reward scaling parameter was set to 128, and a margin threshold of 50 was used during training.

Synthesizable Generative Model. For SynFormer, we used the default model architecture and pretrained weights provided by the authors’ GitHub repository (<https://github.com/wenhao-gao/synformer>). This method operated over a synthesizable chemical space defined by 115 reaction templates and 223,244 commercially available building blocks sourced from the Enamine US in stock (Grygorenko et al., 2020), covering an estimated space of $> 10^{60}$ molecules. For RL finetuning, we also followed the original paper with minimal modification. Specifically, we retained the frozen prior model and employed a replay buffer that sampled 128 saved molecules at each step using a weighting mechanism over accumulated scores to stabilize the gradient updates. The augmented likelihood scaling factor was set to $\sigma = 500$ and the Adam optimizer was used with a learning rate of 1×10^{-4} , both matching the original paper’s settings. We used a batch size of 128 and the maximal training step of 100, resulting in a total of 12,800 generated molecules per run (including duplicates).

Reverse Sigmoid Function. Given a generated molecule s_i and a reference molecule s_r , the distance between them in embedding space is transformed into a reward via a reverse sigmoid function:

$$R(s_i) = \frac{1}{1 + 10^{10k \cdot \frac{D(s_i, s_r) - m}{h - l}}} \quad (\text{S1})$$

where k adjusts the sharpness of the reward, $m = \frac{h+l}{2}$ is the midpoint, and l and h denote the lower and upper bounds of the distance range, respectively. Specifically, Table S1 displays the hyperparameters of PED sigmoid transformation in different distance modes, which are applied to the above two generative models as a scoring function component.

Table S1. Hyperparameters for the sigmoid transformation used to convert different Euclidean PEDs into reward signals.

Distance mode	l (low)	h (high)	k
GeoDiff 2D	0.1	1.2	0.25
GeoDiff 3D	1	10	0.25
GeoDiff Concat	0.1	0.58	0.25
MoLFormer	5	17	0.25

Experiments were conducted in an HPC server by using an NVIDIA A100 GPU with 80GB of memory. The datasets we used are public and available in <https://drugdesign.unistra.fr/LIT-PCBA/> and <https://zenodo.org/records/10277786>. The code repository is available in GitHub: <https://github.com/molecularinformatics/PED>.

B. Correlation Results of AmpC-ROCS Similarities and Embedding Distance

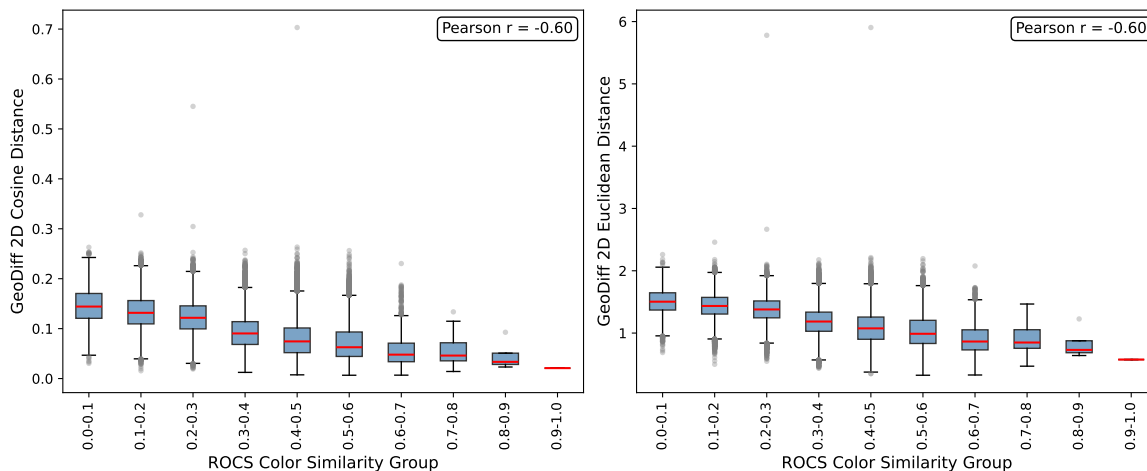


Figure S1. GeoDiff 2D cosine and Euclidean PED vs. ROCS ColorTanimoto.




Figure S2. GeoDiff 3D cosine and Euclidean PED vs. ROCS ShapeTanimoto.




Figure S3. GeoDiff 2D cosine and Euclidean PED vs. ROCS ShapeTanimoto.




Figure S4. GeoDiff 3D cosine and Euclidean PED vs. ROCS ColorTanimoto.




Figure S5. GeoDiff concatenated cosine and Euclidean PED vs. ROCS combination.




Figure S6. MolFormer cosine and Euclidean PED vs. ROCS ColorTanimoto.




Figure S7. MolFormer cosine and Euclidean PED vs. ROCS ShapeTanimoto.

C. LIT-PCBA Benchmark and Virtual Screening Results on EF1%

We used a challenging virtual screening dataset, LIT-PCBA (Tran-Nguyen et al., 2020), to benchmark PED. However, Huang et al. (2025) found that the provided original reference ligands in the .mol2 files often lack bond orders, impeding accurate structure reconstruction and yielding incorrect SMILES. We then used their corrected canonical SMILES for all the reference ligands.

The EF1% results of ROSHAMBO2 are provided by its original paper (Atwi et al., 2025). 2D ECFP4 and ROCS 3D similarities and molecular docking results are from the original LIT-PCBA benchmark (Tran-Nguyen et al., 2020).




Figure S8. LIT-PCBA virtual screening performance (EF1%) obtained by ROSHAMBO2 similarities, GeoDiff and MoLFormer Euclidean PEDs. Different modes of ROSHAMBO2 similarities are in green, while those for GeoDiff PEDs are in blue. Boxplots illustrate the distribution of EF1% values across query ligands for each target. Circular markers indicate the best-pooled EF1%, while cross markers denote the outliers of the boxplots. In case where there is only one reference ligand, the boxes would be a single line and the best-pooled circle overlaps with it. Targets highlighted in a blue background are those identified as favorable for 3D shape-based screening (EF1% > 2 with ROCS) in the original LIT-PCBA benchmark. Dashed horizontal red and blue lines correspond to EF1% of 1 and 2, respectively.

Table S2. Complete virtual screening performance (EF1%) obtained by different modes of ROSHAMBO2 similarity and GeoDiff cosine PED. The best average mean and best-pooled EF1% values for each approach are shown in bold.

Target	ROSHAMBO2 color		ROSHAMBO2 shape		ROSHAMBO2 combination		GeoDiff 2D		GeoDiff 3D		GeoDiff concat	
	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled
ADRB2	10.94 \pm 5.54	12.50	0.78 \pm 2.21	6.25	3.91 \pm 3.24	12.50	14.70 \pm 4.45	11.76	2.21 \pm 3.04	5.88	6.62 \pm 3.77	11.76
ALDH1	0.98 \pm 0.04	1.02	1.23 \pm 0.30	0.88	1.18 \pm 0.52	0.68	0.93 \pm 0.18	0.85	1.20 \pm 0.27	1.09	1.08 \pm 0.25	1.29
ESR1_ago	3.03 \pm 4.21	0.00	1.82 \pm 4.91	8.47	1.82 \pm 3.61	8.47	4.02 \pm 6.28	0.00	3.01 \pm 3.82	15.06	4.02 \pm 4.82	15.06
ESR1_ant	0.73 \pm 0.88	1.00	1.67 \pm 1.34	1.00	2.53 \pm 2.06	4.00	1.11 \pm 0.98	2.23	2.39 \pm 1.85	5.57	2.71 \pm 1.56	2.23
FEN1	5.16	5.16	0.54	0.54	2.17	2.17	2.78	2.78	0.83	0.83	2.50	2.50
GBA	2.30 \pm 0.66	3.64	0.36 \pm 0.54	1.21	1.58 \pm 0.54	1.21	3.20 \pm 1.35	3.05	1.37 \pm 0.30	3.05	2.13 \pm 0.35	1.83
IDH1	0.20 \pm 0.73	0.00	1.01 \pm 1.71	0.00	1.01 \pm 1.71	0.00	2.20 \pm 2.22	7.69	0.73 \pm 1.20	2.56	1.28 \pm 1.67	2.56
KAT2A	1.04 \pm 0.73	1.55	1.30 \pm 1.10	2.07	1.55 \pm 0.73	1.04	0.00 \pm 0.00	0.00	0.77 \pm 0.36	1.03	0.26 \pm 0.36	0.52
MAPK1	0.60 \pm 0.48	0.65	1.56 \pm 1.31	1.30	1.35 \pm 0.67	2.28	1.42 \pm 1.28	1.62	1.95 \pm 0.86	1.95	1.95 \pm 0.89	2.27
MTORC1	2.37 \pm 0.82	1.04	0.28 \pm 0.67	2.09	1.04 \pm 0.47	1.04	2.47 \pm 0.92	3.09	1.65 \pm 0.92	1.03	4.74 \pm 2.69	0.00
OPRK1	0.00	0.00	4.35	4.35	0.00	0.00	4.17	4.17	4.17	4.17	4.17	4.17
PKM2	2.70 \pm 2.06	0.37	3.83 \pm 2.55	1.65	4.74 \pm 3.58	2.39	3.03 \pm 1.83	4.76	3.42 \pm 1.95	5.49	4.01 \pm 2.53	7.14
PPARG	2.49 \pm 4.19	0.00	5.81 \pm 6.03	7.74	6.64 \pm 5.35	15.49	5.16 \pm 4.98	8.15	7.33 \pm 3.51	20.37	5.97 \pm 5.52	12.22
TP53	0.43 \pm 0.67	0.00	2.16 \pm 1.95	5.18	0.86 \pm 1.34	2.59	0.25 \pm 0.62	0.00	0.00 \pm 0.00	0.00	0.00 \pm 0.00	0.00
VDR	8.55 \pm 0.08	8.49	0.34 \pm 0.00	0.45	3.34 \pm 0.08	3.17	2.07	2.07	0.69	0.69	0.69	0.69
Average	2.77 \pm 1.62	2.36	1.80 \pm 1.89	2.88	2.25 \pm 1.84	3.80	3.17 \pm 2.09	3.48	2.11 \pm 1.51	4.58	2.81 \pm 2.03	4.28

Table S3. Virtual screening performance (EF1%) obtained by different methods. Each method reports per-reference Mean \pm SD and Best-pooled EF1%. ROSHAMBO2 score is based on the combination mode. GeoDiff uses 3D mode with cosine PED, and MoLFormer uses cosine PED. The best average values are shown in **bold**, and the second best are underlined.

Target	2D ECFP4 similarity		ROCS 3D similarity		Molecular docking		ROSHAMBO2 score		GeoDiff cosine distance		MoLFormer cosine distance	
	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled
ADRB2	18.38 \pm 6.62	17.65	9.76 \pm 5.35	20.00	2.94 \pm 3.14	5.88	3.91 \pm 3.24	12.50	2.21 \pm 3.04	5.88	16.17 \pm 4.16	17.64
ALDH1	1.40 \pm 0.50	2.61	1.20 \pm 0.39	1.85	1.19 \pm 0.10	0.96	1.18 \pm 0.52	0.68	1.20 \pm 0.27	1.09	1.67 \pm 0.32	1.95
ESR1_ago	7.52 \pm 6.09	7.69	4.61 \pm 4.86	7.69	0.00	0.00	1.82 \pm 3.61	8.47	3.01 \pm 3.82	15.06	13.05 \pm 8.76	7.53
ESR1_ant	1.39 \pm 1.06	0.98	1.66 \pm 1.55	3.92	0.91 \pm 0.78	0.98	2.53 \pm 2.06	4.00	2.39 \pm 1.85	5.57	1.11 \pm 1.75	0.00
FEN1	1.08	1.08	0.27	0.27	4.34	4.34	2.17	2.17	0.83	0.83	1.39	1.39
GBA	2.01 \pm 1.41	3.61	0.72 \pm 0.35	0.86	5.72 \pm 1.97	8.43	1.58 \pm 0.54	1.21	1.37 \pm 0.30	3.05	4.27 \pm 0.50	4.27
IDH1	1.83 \pm 3.09	5.13	0.86 \pm 2.32	0.00	0.73 \pm 1.20	0.00	1.01 \pm 1.71	0.00	0.73 \pm 1.20	2.56	1.65 \pm 1.91	0.00
KAT2A	0.17 \pm 0.30	0.52	1.48 \pm 0.07	1.55	1.89 \pm 0.59	2.58	1.55 \pm 0.73	1.04	0.77 \pm 0.36	1.03	2.83 \pm 1.82	4.12
MAPK1	1.20 \pm 0.74	1.95	1.96 \pm 1.01	5.19	1.32 \pm 0.53	1.62	1.35 \pm 0.67	2.28	1.95 \pm 0.86	1.95	3.90 \pm 1.98	4.55
MTORC1	0.28 \pm 0.48	0.00	1.38 \pm 1.29	3.23	1.03 \pm 0.65	1.03	1.04 \pm 0.47	1.04	1.65 \pm 0.92	1.03	0.82 \pm 0.86	2.06
OPRK1	12.50	12.50	0.00	0.00	4.17	4.17	0.00	0.00	4.17	4.17	4.17	4.17
PKM2	3.30 \pm 2.73	7.95	4.86 \pm 4.30	6.22	0.57 \pm 0.37	0.18	4.74 \pm 3.58	2.39	3.42 \pm 1.95	5.49	4.17 \pm 3.85	10.25
PPARG	3.65 \pm 5.14	7.41	8.79 \pm 5.24	18.52	5.92 \pm 5.20	7.41	6.64 \pm 5.35	15.49	7.33 \pm 3.51	20.37	5.98 \pm 5.73	28.52
TP53	1.06 \pm 0.95	0.00	0.88 \pm 1.36	2.53	0.00	0.00	0.86 \pm 1.34	2.59	0.00 \pm 0.00	0.00	1.02 \pm 1.85	0.00
VDR	3.37	3.37	0.35	0.35	0.34 \pm 0.48	0.00	3.34 \pm 0.08	3.17	0.69	0.69	5.74	5.74
Average	3.94 \pm 2.43	<u>4.83</u>	2.59 \pm 2.34	4.81	2.07 \pm 1.36	2.51	2.25 \pm 1.84	3.80	2.11 \pm 1.51	4.58	4.53 \pm 2.79	<u>6.15</u>

Table S4. Complete virtual screening performance (EF1%) obtained by different modes of ROSHAMBO2 similarity and GeoDiff Euclidean PED. The best average mean and best-pooled EF1% values for each approach are shown in **bold**.

Target	ROSHAMBO2 color		ROSHAMBO2 shape		ROSHAMBO2 combination		GeoDiff 2D		GeoDiff 3D		GeoDiff concat	
	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled
ADRB2	10.94 \pm 5.54	12.50	0.78 \pm 2.21	6.25	3.91 \pm 3.24	12.50	14.70 \pm 4.45	11.76	2.21 \pm 3.04	5.88	6.62 \pm 3.77	11.76
ALDH1	0.98 \pm 0.04	1.02	1.23 \pm 0.30	0.88	1.18 \pm 0.52	0.68	0.92 \pm 0.17	0.92	1.23 \pm 0.29	1.09	1.08 \pm 0.25	1.29
ESR1_ago	3.03 \pm 4.21	0.00	1.82 \pm 4.91	8.47	1.82 \pm 3.61	8.47	4.02 \pm 6.28	0.00	3.01 \pm 3.82	22.59	4.02 \pm 4.82	15.06
ESR1_ant	0.73 \pm 0.88	1.00	1.67 \pm 1.34	1.00	2.53 \pm 2.06	4.00	1.19 \pm 1.02	2.23	2.31 \pm 1.88	5.57	2.71 \pm 1.56	2.23
FEN1	5.16	5.16	0.54	0.54	2.17	2.17	2.78	2.78	0.83	0.83	2.50	2.50
GBA	2.30 \pm 0.66	3.64	0.36 \pm 0.54	1.21	1.58 \pm 0.54	1.21	3.20 \pm 1.35	3.05	1.37 \pm 0.30	3.05	2.13 \pm 0.35	1.83
IDH1	0.20 \pm 0.73	0.00	1.01 \pm 1.71	0.00	1.01 \pm 1.71	0.00	2.20 \pm 2.22	7.69	0.73 \pm 1.20	2.56	1.28 \pm 1.67	2.56
KAT2A	1.04 \pm 0.73	1.55	1.30 \pm 1.10	2.07	1.55 \pm 0.73	1.04	0.00 \pm 0.00	0.00	1.03 \pm 0.73	1.55	0.26 \pm 0.36	0.52
MAPK1	0.60 \pm 0.48	0.65	1.35 \pm 1.31	1.30	1.35 \pm 0.67	2.28	1.42 \pm 1.26	1.30	1.92 \pm 0.83	1.95	1.95 \pm 0.89	2.27
MTORC1	2.37 \pm 0.82	1.04	0.28 \pm 0.67	2.09	1.04 \pm 0.47	1.04	2.47 \pm 0.92	2.06	1.65 \pm 0.92	1.03	4.74 \pm 2.69	0.00
OPRK1	0.00	0.00	4.35	4.35	0.00	0.00	4.17	4.17	8.33	8.33	4.17	4.17
PKM2	2.70 \pm 2.06	0.37	3.83 \pm 2.55	1.65	4.74 \pm 3.58	2.39	3.05 \pm 1.79	5.13	3.58 \pm 2.06	6.59	4.01 \pm 2.53	7.14
PPARG	2.49 \pm 4.19	0.00	5.81 \pm 6.03	7.74	6.64 \pm 5.35	15.49	4.89 \pm 5.15	8.15	6.52 \pm 3.71	20.37	5.97 \pm 5.52	12.22
TP53	0.43 \pm 0.67	0.00	2.16 \pm 1.95	5.18	0.86 \pm 1.34	2.59	0.25 \pm 0.62	0.00	0.00 \pm 0.00	0.00	0.00 \pm 0.00	0.00
VDR	8.55 \pm 0.08	8.49	0.34 \pm 0.00	0.45	3.34 \pm 0.08	3.17	2.07	2.07	0.69	0.69	0.69	0.69
Average	2.77 \pm 1.62	2.36	1.80 \pm 1.89	2.88	2.25 \pm 1.84	3.80	3.16 \pm 2.10	3.42	2.36 \pm 1.57	5.47	2.81 \pm 2.03	4.28

Table S5. Virtual screening performance (EF1%) obtained by different methods. Each method reports per-reference Mean \pm SD and Best-pooled EF1%. ROSHAMBO2 score is based on the combination mode. GeoDiff uses 3D mode with Euclidean PED, and MoLFormer uses Euclidean PED. The best average values are shown in **bold**, and the second best are underlined.

Target	2D ECFP4 similarity		ROCS 3D similarity		Molecular docking		ROSHAMBO2 score		GeoDiff Euclidean distance		MoLFormer Euclidean distance	
	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled	Mean \pm SD	Best-pooled
ADRB2	18.38 \pm 6.62	17.65	9.76 \pm 5.35	20.00	2.94 \pm 3.14	5.88	3.91 \pm 3.24	12.50	2.21 \pm 3.04	5.88	11.76 \pm 6.29	17.64
ALDH1	1.40 \pm 0.50	2.61	1.20 \pm 0.39	1.85	1.19 \pm 0.10	0.96	1.18 \pm 0.52	0.68	1.23 \pm 0.29	1.09	1.69 \pm 0.29	1.99
ESR1_ago	7.52 \pm 6.09	7.69	4.61 \pm 4.86	7.69	0.00	0.00	1.82 \pm 3.61	8.47	3.01 \pm 3.82	22.59	13.05 \pm 9.64	0.00
ESR1_ant	1.39 \pm 1.06	0.98	1.66 \pm 1.55	3.92	0.91 \pm 0.78	0.98	2.53 \pm 2.06	4.00	2.31 \pm 1.88	5.57	0.96 \pm 1.06	0.00
FEN1	1.08	1.08	0.27	0.27	4.34	4.34	2.17	2.17	0.83	0.83	1.11	1.11
GBA	2.01 \pm 1.41	3.61	0.72 \pm 0.35	0.86	5.72 \pm 1.97	8.43	1.58 \pm 0.54	1.21	1.37 \pm 0.30	3.05	4.11 \pm 0.58	3.66
IDH1	1.83 \pm 3.09	5.13	0.86 \pm 2.32	0.00	0.73 \pm 1.20	0.00	1.01 \pm 1.71	0.00	0.73 \pm 1.20	2.56	2.20 \pm 2.82	0.00
KAT2A	0.17 \pm 0.30	0.52	1.48 \pm 0.07	1.55	1.89 \pm 0.59	2.58	1.55 \pm 0.73	1.04	1.03 \pm 0.73	1.55	2.32 \pm 0.36	2.58
MAPK1	1.20 \pm 0.74	1.95	1.96 \pm 1.01	5.19	1.32 \pm 0.53	1.62	1.35 \pm 0.67	2.28	1.92 \pm 0.83	1.95	3.69 \pm 2.23	4.22
MTORC1	0.28 \pm 0.48	0.00	1.38 \pm 1.29	3.23	1.03 \pm 0.65	1.03	1.04 \pm 0.47	1.04	1.65 \pm 0.92	1.03	2.27 \pm 0.86	1.03
OPRK1	12.50	12.50	0.00	0.00	4.17	4.17	0.00	0.00	8.33	8.33	4.17	4.17
PKM2	3.30 \pm 2.73	7.95	4.86 \pm 4.30	6.22	0.57 \pm 0.37	0.18	4.74 \pm 3.58	2.39	3.58 \pm 2.06	6.59	4.39 \pm 3.88	10.44
PPARG	3.65 \pm 5.14	7.41	8.79 \pm 5.24	18.52	5.92 \pm 5.20	7.41	6.64 \pm 5.35	15.49	6.52 \pm 3.71	20.37	6.52 \pm 7.35	24.45
TP53	1.06 \pm 0.95	0.00	0.88 \pm 1.36	2.53	0.00	0.00	0.86 \pm 1.34	2.59	0.00 \pm 0.00	0.00	1.53 \pm 1.93	0.00
VDR	3.37	3.37	0.35	0.35	0.34 \pm 0.48	0.00	3.34 \pm 0.08	3.17	0.69	0.69	4.94	4.94
Average	3.94 \pm 2.43	4.83	2.59 \pm 2.34	4.81	2.07 \pm 1.36	2.51	2.25 \pm 1.84	3.80	2.36 \pm 1.57	5.47	4.31 \pm 3.11	<u>5.08</u>

D. Training Dynamics of Generative Frameworks with different Rewards

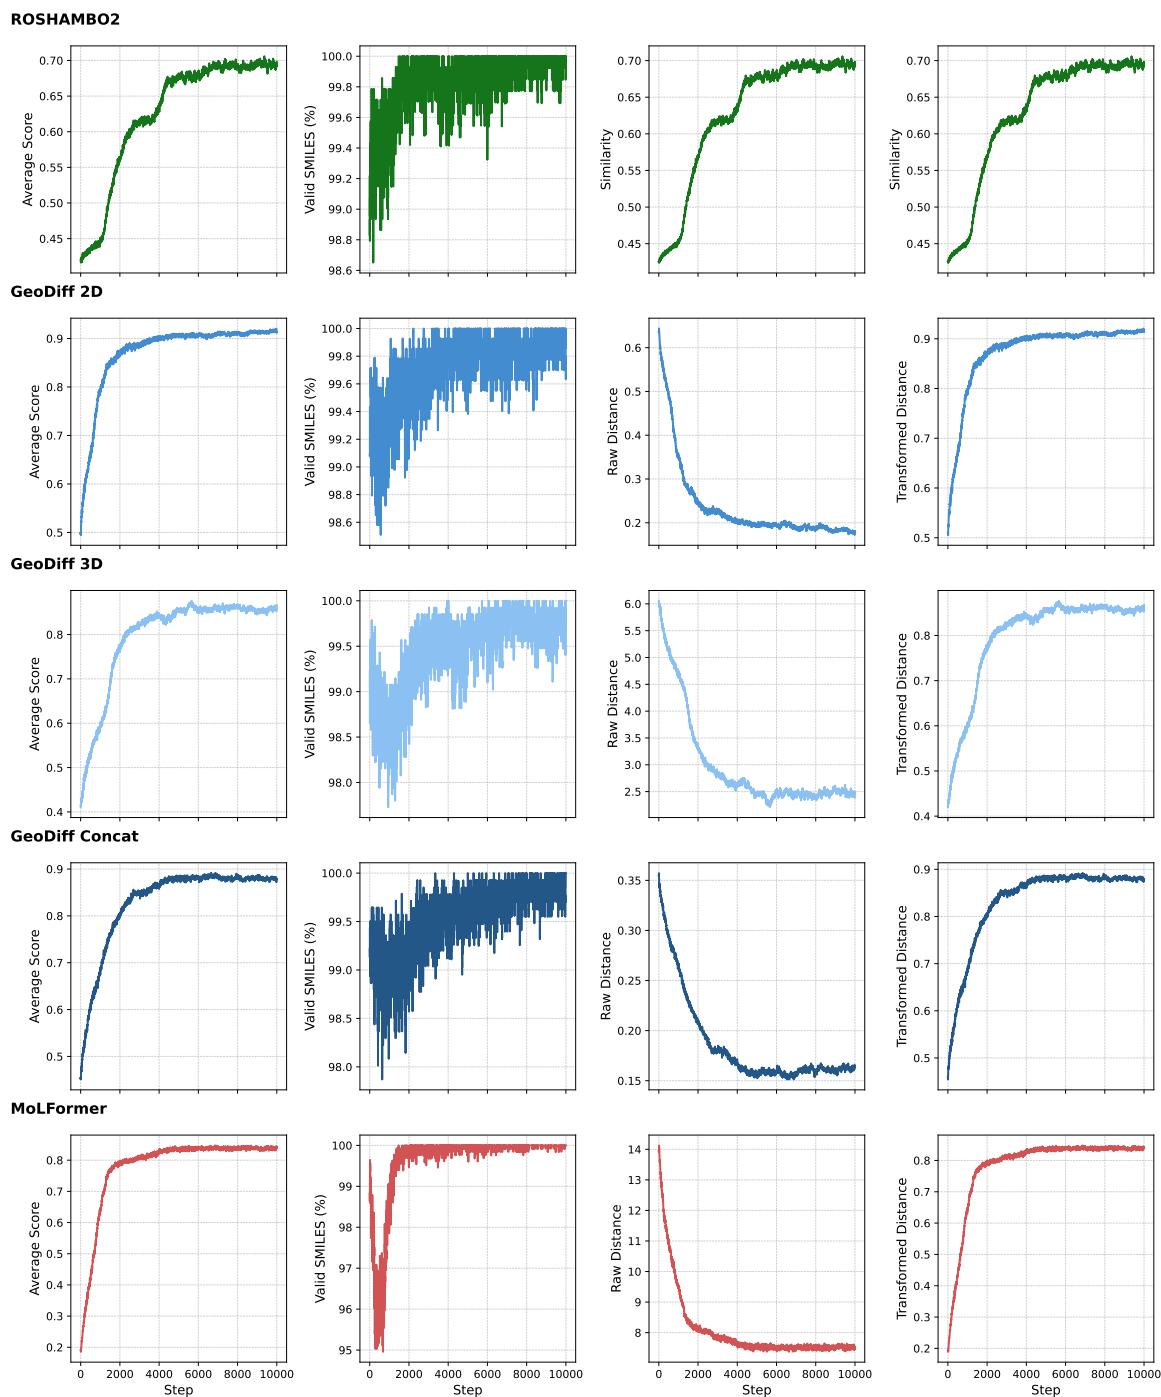


Figure S9. Training dynamics of REINVENT with different reward components. Average score, valid SMILES rate, and similarity or distance metrics are tracked during training for ROSHAMBO2, GeoDiff PED variants, and MoLFormer PED. Curves are computed by averaging scores within each training step and applying a moving average smoothing to reduce stochastic variation from sampling.

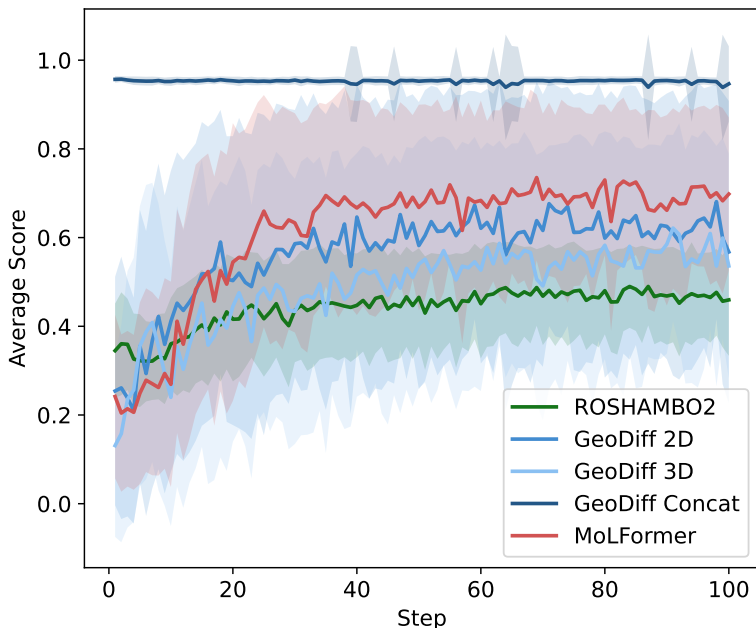


Figure S10. Training dynamics of SynFormer with different reward components. Curves are computed by averaging scores within each training step.

E. Statistics of Generated Molecules

As mentioned in Appendix A, in each run, REINVENT generates 1,280,000 molecules and SynFormer generates 12,800 molecules. Table S6 shows all the unique molecules and scaffolds generated by REINVENT and SynFormer.

Table S6. REINVENT and SynFormer generated unique molecules, unique ring-linker scaffolds, and the unique scaffold ratio.

Model	Method	Unique molecules	Unique scaffold	Unique scaffold ratio (\uparrow)	Running time (h) (\downarrow)
REINVENT	ROSHAMBO2	277,116	100,983	36.44%	47
	GeoDiff 2D	211,835	84,764	40.01%	31
	GeoDiff 3D	421,879	230,346	54.60%	31
	GeoDiff Concat	338,333	119,441	35.30%	31
	MoLFormer	179,598	86,091	47.94%	14
SynFormer	ROSHAMBO2	6,813	1,459	21.41%	1.33
	GeoDiff 2D	8,853	3,827	43.23%	0.70
	GeoDiff 3D	7,999	1,532	19.15%	0.70
	GeoDiff Concat	6,089	1,081	17.75%	0.70
	MoLFormer	7,357	1,003	13.63%	0.65

F. Drug-likeness of Top-5,000 Generated Molecules

We followed the evaluation protocol of Wang et al. (2023) to assess the physicochemical properties of the generated compounds. Specifically, for the top 5,000 generated molecules, we computed five physicochemical properties, including molecular weight (MW), topological polar surface area (TPSA), partition coefficient (LogP, indicating hydrophobicity), number of hydrogen bond donors (HBD), and number of hydrogen bond acceptors (HBA), along with quantitative estimate of drug-likeness (QED) and synthetic accessibility (SA) scores. The minimum and maximum values of these physicochemical properties are reported in Table S7, showing that the generated molecules cover a broad range of physicochemical space. In addition, we evaluated the proportion of molecules that fall within desirable drug-like ranges for each property (Table S8).

Table S7. Physical-chemical properties of top-5,000 generated molecules by REINVENT and SynFormer.


Model	Method	MW _{min}	MW _{max}	TPSA _{min}	TPSA _{max}	LogP _{min}	LogP _{max}	HBD _{min}	HBD _{max}	HBA _{min}	HBA _{max}
REINVENT	ROSHAMBO2	245.12	469.04	41.57	136.82	-0.26	5.31	0.00	4.00	2.00	7.00
	GeoDiff 2D	258.15	478.24	45.48	134.31	0.96	4.56	1.00	5.00	1.00	5.00
	GeoDiff 3D	325.19	466.11	38.13	165.65	0.09	5.23	0.00	5.00	2.00	8.00
	GeoDiff Concat	301.18	433.19	54.34	147.20	0.39	4.17	1.00	4.00	2.00	5.00
	MoLFormer	307.15	496.25	39.34	151.55	0.11	5.14	0.00	5.00	2.00	6.00
SynFormer	ROSHAMBO2	128.04	650.02	0.00	252.98	-1.79	8.28	0.00	7.00	0.00	10.00
	GeoDiff 2D	261.18	1116.36	31.06	278.13	-0.57	12.75	0.00	8.00	1.00	14.00
	GeoDiff 3D	224.15	690.99	6.48	193.57	-0.68	8.07	0.00	6.00	1.00	9.00
	GeoDiff Concat	73.09	1089.41	0.00	306.02	-3.83	11.91	0.00	6.00	0.00	18.00
	MoLFormer	219.08	823.37	31.06	227.01	-0.05	10.77	0.00	6.00	1.00	9.00

Table S8. Percentages of top-5,000 generated molecules by REINVENT and SynFormer in good drug-likeness range.


Model	Method	MW ([200, 500])	TPSA ([20, 130])	LogP ([-1, 6])	HBD ([0, 5])	HBA ([0, 10])	QED ([0.4, 1])	SA ([1, 5])
REINVENT	ROSHAMBO2	100.00%	99.82%	100.00%	100.00%	100.00%	100.00%	100.00%
	GeoDiff 2D	100.00%	99.98%	100.00%	100.00%	100.00%	99.98%	100.00%
	GeoDiff 3D	100.00%	94.38%	100.00%	100.00%	100.00%	99.98%	100.00%
	GeoDiff Concat	100.00%	97.56%	100.00%	100.00%	100.00%	100.00%	100.00%
	MoLFormer	100.00%	98.40%	100.00%	100.00%	100.00%	99.92%	100.00%
SynFormer	ROSHAMBO2	92.06%	94.18%	99.26%	99.82%	100.00%	94.24%	100.00%
	GeoDiff 2D	57.18%	84.92%	75.74%	98.18%	99.00%	49.18%	98.92%
	GeoDiff 3D	97.04%	98.06%	99.02%	99.92%	100.00%	86.42%	100.00%
	GeoDiff Concat	36.10%	84.26%	96.40%	99.82%	99.62%	92.84%	99.72%
	MoLFormer	94.48%	98.36%	95.78%	99.92%	100.00%	75.16%	100.00%

G. Top-5 Generated Molecules by REINVENT and SynFormer


Reference




ROSHAMBO2




Total Score=0.80
Predicted pIC50=7.98




Total Score=0.79
Predicted pIC50=8.03



Total Score=0.79
Predicted pIC50=9.00




Total Score=0.79
Predicted pIC50=7.18




Total Score=0.78
Predicted pIC50=7.18


GeoDiff 2D




Total Score=0.94
Predicted pIC50=7.20




Total Score=0.94
Predicted pIC50=7.21



Total Score=0.94
Predicted pIC50=7.80




Total Score=0.94
Predicted pIC50=7.38




Total Score=0.94
Predicted pIC50=7.32


GeoDiff 3D




Total Score=0.93
Predicted pIC50=8.78




Total Score=0.93
Predicted pIC50=12.75



Total Score=0.93
Predicted pIC50=12.62




Total Score=0.93
Predicted pIC50=12.85




Total Score=0.93
Predicted pIC50=8.99


GeoDiff Concat




Total Score=0.95
Predicted pIC50=8.28




Total Score=0.95
Predicted pIC50=8.42



Total Score=0.95
Predicted pIC50=7.24




Total Score=0.95
Predicted pIC50=8.18




Total Score=0.95
Predicted pIC50=8.24


MolFormer




Total Score=0.88
Predicted pIC50=11.52




Total Score=0.88
Predicted pIC50=10.59



Total Score=0.88
Predicted pIC50=11.01




Total Score=0.88
Predicted pIC50=10.66




Total Score=0.88
Predicted pIC50=11.08

Figure S11. The reference molecule and REINVENT-generated molecules in top-5 total score across methods.


Reference




ROSHAMBO2




Total Score=0.67
Predicted pIC50=8.44




Total Score=0.66
Predicted pIC50=8.18



Total Score=0.66
Predicted pIC50=6.96




Total Score=0.66
Predicted pIC50=7.38




Total Score=0.66
Predicted pIC50=7.81


GeoDiff 2D




Total Score=0.94
Predicted pIC50=9.13




Total Score=0.93
Predicted pIC50=10.12



Total Score=0.93
Predicted pIC50=7.95




Total Score=0.93
Predicted pIC50=9.82




Total Score=0.93
Predicted pIC50=10.15


GeoDiff 3D




Total Score=0.91
Predicted pIC50=9.97




Total Score=0.90
Predicted pIC50=7.19



Total Score=0.90
Predicted pIC50=9.56




Total Score=0.90
Predicted pIC50=8.26




Total Score=0.90
Predicted pIC50=9.49


GeoDiff Concat




Total Score=0.97
Predicted pIC50=8.36




Total Score=0.97
Predicted pIC50=9.14



Total Score=0.97
Predicted pIC50=8.83




Total Score=0.97
Predicted pIC50=8.59




Total Score=0.97
Predicted pIC50=8.31


MoLFormer




Total Score=0.94
Predicted pIC50=8.42




Total Score=0.93
Predicted pIC50=9.03



Total Score=0.93
Predicted pIC50=7.59



Total Score=0.93
Predicted pIC50=7.07



Total Score=0.92
Predicted pIC50=9.51

Figure S12. The reference molecule and SynFormer-generated molecules in top-5 total score across methods.