DOI: 10.34749/3061-1466.2026.9

LigForge: Physics-Informed Diffusion for Structure-Based Drug Design

Tobias Schopper
tobias.schopper@innophore.com

Michael Hetmann **Lena Parigger** **Chiara Gasbarri** **Christian Gruber**

Georg Steinkellner

Abstract

Structure-based drug design aims to generate novel small molecules that bind favorably to a defined protein target. While recent generative models have demonstrated impressive results, they typically face challenges such as limited specificity, insufficient diversity, poor physical or chemical validity, weak binding affinity, and difficulties generalizing beyond the binders and targets seen during training. Here we present LigForge, a physics-informed pipeline that decouples key competences to overcome these challenges.

1 Methods

LigForge operates in three stages that decouple the modelling of the target, binder and their interactions, with the aim of generalizing beyond known complexes.

1.1 Stage 1: Interaction Field Prediction

A 3D U-Net is trained on the entire Protein Data Bank (PDB) to predict spatial probability maps for eleven types of protein–ligand interactions within a binding cavity. These interaction types include hydrogen bonding, halogen bonding, π–π stacking, edge-to-face π interactions, metal coordination, and others. Training is performed on pairs of voxelized potential fields and annotated interaction heatmaps derived from experimentally resolved crystal structures, with a 3 Å restricted receptive field to ensure focus on local interaction features rather than memorization of specific pocket geometries. Training on potential fields, agnostic of sidechains and sequence, forces the model to understand fundamental interaction patterns, enabling it to generalize across protein families, for example from non-human to human targets.

1.2 Stage 2: Guided Diffusion with Catalophore Potentials

An atomistic diffusion model, trained on the GEOM dataset for generating chemically and physically viable molecules and their conformers via iterative denoising, forms the generative core. Crucial for generalisation, this model has never been trained on proteins or protein–ligand complexes. We extend the standard diffusion sampling procedure with energy-based guidance through potential fields extracted from the target protein, termed “Catalophores,” computed using CavitomixPro (Innophore GmbH). To avoid bottlenecks during inference, we implemented rapid, differentiable, and highly accurate descriptor calculation for localized quantum charges, orbital assignment, acceptor/donor classification and entropy related quantities, enabling efficient physical steering during inference. To simultaneously improve drug-likeness without sacrificing binding potential or diversity, we introduce

The Third Austrian Symposium on AI and Vision (AIROV26).

a suite of structural entropy penalties also applied during inference. These reduce the enthalpy–entropy tradeoff and improve synthesizability and specificity by penalizing excessive rotatable bonds, structural complexity indicators such as bridgehead atoms, spiro centers, and saturated ring systems, while optimizing Sp^3 ratio and total polar surface area.

1.3 Stage 3: Candidate Refinement Pipeline

Post-generation, molecules are filtered and reranked by multiple quantities of interest, including synthesizability (assessed via RAScore), solubility, and structural entropy metrics. From a typical generation run of over 20,000 molecules, candidates are first scored and ranked using GNINA, UNet interaction densities as well as simple entropic penalty terms. The top-ranked ligands are then subjected to molecular dynamics (MD) simulations to estimate average unbinding times, continuing until approximately 100 candidates with acceptable unbinding times are obtained. Finally, highly accurate quantum mechanics (QM) computations are performed on the top MD hits to more precisely estimate binding free energies. This cascaded filtering strategy balances throughput with accuracy: the full pipeline from pocket input to QM-validated candidates can be completed within a single day, with throughput exceeding 20,000 molecules per day on a single RTX 4090, exact numbers depending on the average molecule size and the geometric complexity of the target.

2 Results

Systematic tests on six targets—encompassing buried hydrophobic virus cavities, superficial binding sites, metal-binding sites, and enzymatic active sites—demonstrated that LigForge consistently generates binders with superior computed affinities relative to known inhibitors. Across all tested cavity types, GNINA-predicted $\mathrm{p}K_i$ values for generated molecules exceeded those of reference inhibitors by more than one log unit (corresponding to approximately one order of magnitude improvement in predicted binding affinity), with predicted VINA binding free energies (ΔG) improving by approximately -4 kcal/mol. In many cases, the model re-discovered nontrivial functional groups found in known inhibitors—we consider this a noteworthy preliminary result given that these motifs emerge solely due to physical guidance.

Counterintuitively, both scaffold diversity and internal diversity of generated ligands increased with increasing guidance strength: among 1,000 generated molecules, approximately 900 unique scaffolds were observed, with an average pairwise Tanimoto similarity of only 0.1. This suggests that incorporating physics into the diffusion process improves the generalization abilities of the model, unlocking regions in chemical space not covered in the training data.

2.1 Multi-Target Design

For applications requiring pathway inhibition or mitigation of off-target binding, the Catalophore potential fields can be computed and aligned for multiple targets and antitargets simultaneously using consensus cavities calculated in CavitomixPro, enabling structure- and sequence-agnostic comparison and characterization of protein binding sites.

3 Conclusion

LigForge provides a modular, physics-based generative framework for structure-based drug design that integrates learned pharmacophore prediction with guided diffusion sampling. By decoupling interaction field estimation from molecular generation, the pipeline can be readily adapted to new targets without retraining the generative model. The Catalophore guidance mechanism offers a differentiable route to incorporate physical knowledge beyond pharmacophores into diffusion-based design, and the structural entropy penalties provide a practical lever for generating binders that complement target pockets without excessive rotatable bonds. Together, these components establish LigForge as a powerful and efficient platform for exploration and optimization of drug-like chemical space conditioned on 3D protein structures.