

# 000 001 002 003 004 005 TOWARD PROTEIN DOCKING-ORIENTED DE NOVO 006 LIGAND DESIGN VIA GRADIENT INVERSION 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031

## ABSTRACT

032 De novo ligand design is a fundamental task that seeks to generate protein or  
033 molecule candidates that can effectively dock with protein receptors and achieve  
034 strong binding affinity entirely from scratch. It holds paramount significance for  
035 a wide spectrum of biomedical applications. However, most existing studies are  
036 constrained by the **Pseudo De Novo**, **Limited Docking Modeling**, and **Inflexible**  
037 **Ligand Type**. To address these issues, we propose MagicDock, a forward-  
038 looking framework grounded in the progressive pipeline and differentiable surface  
039 modeling. (1) We adopt a well-designed gradient inversion framework. To begin  
040 with, general docking knowledge of receptors and ligands is incorporated into  
041 the backbone model. Subsequently, the docking knowledge is instantiated as re-  
042 verse gradient flows by binding prediction, which iteratively guide the de novo  
043 generation of ligands. (2) We emphasize differentiable surface modeling in the  
044 docking process, leveraging learnable 3D point-cloud representations to precisely  
045 capture binding details, thereby ensuring that the generated ligands preserve dock-  
046 ing validity through direct and interpretable spatial fingerprints. (3) We introduce  
047 customized designs for different ligand types and integrate them into a unified gra-  
048 dient inversion framework with flexible triggers, thereby ensuring broad applica-  
049 bility. Moreover, we provide rigorous theoretical guarantees for each component  
050 of MagicDock. Extensive experiments across 9 scenarios demonstrate that Mag-  
051 icDock achieves average improvements of 27.1% and 11.7% over SOTA baselines  
052 specialized for protein or molecule ligand design, respectively.  
053

## 1 INTRODUCTION

034 De novo ligand design is a cornerstone of bioengineering, centered on the creation of ligands—such  
035 as proteins and molecules—with strong binding affinity to target protein receptors, thereby forming  
036 highly stable complexes with substantial biological potential. Traditionally, ligand design has  
037 relied on energy optimization techniques ( Adolf-Bryfogle et al. (2018)). Recent advances in deep  
038 learning have transformed the field by introducing powerful data-driven methods, substantially en-  
039 hancing generative capabilities ( Evans et al. (2021); Gu et al. (2024)). Despite their effectiveness,  
040 existing methods still face inherent limitations, as shown in Fig. 1. **(1) Pseudo De Novo**. They in-  
041 herently remain dependent on prior knowledge. Specifically, some antibody design methods depend  
042 heavily on predefined structural templates—such as fixed frameworks and conserved CDR regions  
043 excluding CDR-H3—thereby restricting the design space around the most critical binding region for  
044 optimization while sacrificing the capacity to generate antibodies de novo ( Zhou et al. (2024)).

045 Furthermore, progress toward fully de novo ligand design remains hindered by two additional critical  
046 issues. **(2) Limited Docking Modeling**. Current methods typically employ indirect docking rep-  
047 resentation methods to capture docking performance (such as energy functions assessing docking  
048 tightness through residue-level biophysical terms) ( Zhou et al. (2024)), without explicitly consid-  
049 ering spatial docking information and protein surface information (which may miss key docking  
050 recognition information), leading to the inability to ensure robust biological relevance of generated  
051 ligands. **(3) Inflexible Ligand Type**. Most existing approaches are narrowly tailored to specific  
052 ligand types like protein or molecule ( Luo et al. (2022); Guo et al. (2021)), which severely limits  
053 their versatility and applicability across diverse molecular categories. Collectively, these challenges  
hinder the development of a comprehensive and generalizable framework for ligand design, thereby  
constraining progress in drug discovery and biomolecular engineering.

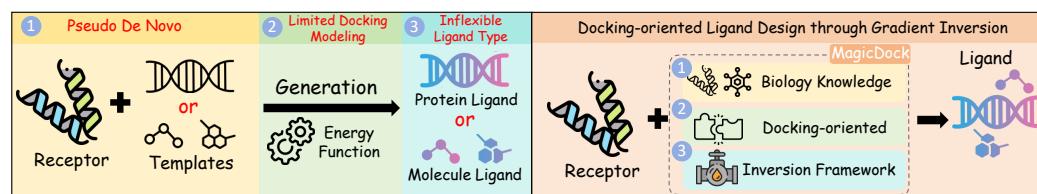


Figure 1: Comparison between current works and MagicDock. This figure describes the three limitations of the existing methods and presents the framework of MagicDock, which achieves authentic de novo, biological significance and cross-category generality.

To address the above critical issues, we propose MagicDock, a forward-looking framework rooted in differentiable surface modeling for de novo ligand design. Our approach is designed with three core innovations. (1) We introduce a **well-designed gradient inversion framework**. To begin with, the backbone model learns general molecular-level knowledge from large protein and molecule datasets. Subsequently, the model acquires docking-specific knowledge by refining its encoders through three progressively structured downstream tasks. By integrating general knowledge and specific docking knowledge into the model, this inversion framework can directly utilize the gradient information contained in the model to guide the generation of ligands from scratch. (2) We emphasize **differentiable surface modeling through learnable 3D point-cloud representations**, enabling the framework to capture fine-grained spatial binding fingerprints with interpretability. This ensures that the generated ligands incorporate spatial and surface information, enabling the ligands to perfectly align with the receptors both geometrically and biologically. (3) We design **customized modules for different ligand types**, which are seamlessly integrated into a unified gradient inversion framework with flexible triggering mechanisms. This design enables convenient switching of different ligand types in generation, improving both efficiency and flexibility in ligand generation. Importantly, the validity and efficiency of our method is supported by rigorous theoretical guarantees in Sec. 4 such as SE(3)-equivariance across stages and superiority over other methods, ensuring both methodological soundness and practical reliability.

**Our contributions.** (1) *New Perspective*: We introduce docking-oriented inversion as an innovative framework for de novo ligand design, addressing challenges in genuineness, biological significance, and cross-category generality. (2) *New Framework*: We introduce a novel inversion framework that leverages gradient-based optimization, starting from surface point cloud modeling of proteins and ligands, through docking-oriented knowledge injection process, to enable inversion for generating de novo ligands directly within the receptor’s binding pocket. (3) *New Method*: We propose a differentiable data structure for seamless gradient flow, integrated with flexible triggers, ensuring flexibility and biological relevance in ligand generation. (4) *SOTA Performance*: Compared across 9 scenarios, **MagicDock** achieves state-of-the-art performance in designing high-affinity ligands, having an average improvement of over 70% in protein ligand design and over 60% in molecule ligand design compared with other baselines.

## 2 PRELIMINARIES & RELATED WORKS

### 2.1 NOTATIONS AND PROBLEM FORMULATION

We adopt a docking-oriented docking strategy, as surface point clouds capture fine-grained structural cues for molecular recognition and docking data provide binding compatibility. Both protein and small-molecule ligands are represented as 3D surface point clouds  $\mathcal{P} = \{\mathbf{f}_i\}_{i=1}^N$ , where  $\mathbf{f}_i \in \mathbb{R}^d$  encodes chemical, atomic and geometric features.

Ligand generation proceeds in two stages: iterative refinement by a type-specific generator  $G_{\text{type}}$  guided by a pre-trained model  $M_\theta$ , followed by generation via docking energy minimization:

$$\hat{\mathcal{P}}_{\text{lig}} = \lim_{t \rightarrow T} G_{\text{type}}\left(\mathcal{P}_{\text{lig}}^{(t)}, M_\theta\right), \quad \mathcal{P}_{\text{lig}}^* = \arg \min_{\hat{\mathcal{P}}_{\text{lig}}} E_{\text{dock}}\left(\mathcal{P}_{\text{rec}}, \hat{\mathcal{P}}_{\text{lig}}\right), \quad (1)$$

where  $t$  denotes the refinement step.  $G_{\text{type}}$  incorporates domain-specific generating constraints, emphasizing ring and valency rules for small molecules and residue/conformation features for proteins.

108  
109

## 2.2 RELATED WORKS

110  
111  
112

Existing methods for protein and molecular ligand design can be categorized by how they *couple generation with optimization*, i.e., the extent to which candidates are refined toward biochemical objectives. We summarize them into four paradigms:

113  
114  
115  
116  
117  
118

**① Decoupled Paradigms.** Traditional pipelines treat representation, generation, and optimization as separate modules. Classical docking methods such as ZDOCK ( Chen et al. (2010)), RosettaDock ( Lyskov & Gray (2008)), AutoDock ( Morris et al. (2008)), and Glide ( Halgren et al. (2004)) encode ligands into atomic/residue descriptors, followed by independent search and scoring. Early generative models like sequence-based RNNs ( Liu et al. (2020)) also rely on post hoc optimization. These approaches are interpretable but loosely connected to task objectives.

119  
120  
121  
122  
123  
124

**② Implicitly Coupled Paradigms.** Diffusion-based approaches embed optimization into sampling. For proteins, DiffAb ( Luo et al. (2022)) and HSRN ( Jin et al. (2022)) refine CDR loops with SE(3)-equivariant models. For ligands, DiffDock ( Corso et al. (2022)) and GeoDiff ( Xu et al. (2022)) integrate docking or  $\Delta G$  signals into denoising, while Pocket2Mol ( Peng et al. (2022)) and TankBind ( Lu et al. (2022)) further condition on receptor pockets. Although effective, these rely on handcrafted schedules and stochastic trajectories, limiting efficiency and *de novo* completeness.

125  
126  
127  
128  
129  
130  
131  
132

**③ Surrogate- or Heuristic-coupled Paradigms.** Another line combines generation with heuristic optimization or surrogate models. For proteins, reinforcement learning (ABDPO ( Zhou et al. (2024))) and memory-augmented models like dyMEAN ( Kong et al. (2023)) incorporate docking rewards. For ligands, reinforcement learning ( Gottipati et al. (2020)), evolutionary strategies ( Chen et al. (2021)), and Bayesian optimization ( Moss et al. (2020)) guide fragment assembly or mutation. Frameworks such as DockStream ( Guo et al. (2021)), ALIDIFF ( Gu et al. (2024)), and DRUGFLOW ( Schneuing et al. (2025)) embed chemical, geometric, or physical priors. These methods are flexible but computationally costly.

133  
134  
135  
136  
137

**④ Latent-gradient coupling Paradigm (Ours).** Inversion differs by explicitly treating generation as optimization ( Niu et al. (2025); Qiu et al. (2024); Bergues et al. (2025)). Structures are refined via gradient updates in latent space, guided by task-specific losses and domain constraints, without stochastic schedules or heuristic surrogates. More details of Sec. 2.2 are applied in Appendix B.

138

## 2.3 THE INVERSION FRAMEWORK

139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152

As an emerging generative framework, Inversion differs fundamentally from mainstream generative frameworks like Diffusion(Ho et al. (2020); Song et al. (2022)) and Flow-matching (Lipman et al. (2023); Liu et al. (2022)), as it employs gradient-based refinement to iteratively adjust structures toward task-specific goals. Its advantages can be summarized as follows: ① Generality. Traditional frameworks often rely on domain-specific priors or handcrafted surrogates, limiting adaptability. Inversion only requires differentiable embeddings and universal gradient updates, enabling one architecture to generalize once the backbone encodes domain knowledge. ② Efficiency. Mainstream methods incur redundant stochastic trajectories or surrogate solvers, leading to high cost and limited efficiency. Inversion directly couples generation and optimization via gradients, achieving higher information efficiency and a stronger theoretical ceiling (Appendix E.3 and E.5). ③ Modularity and Scalability. Existing approaches are often entangled with backbone designs, hindering transfer of pretrained improvements. Inversion remains orthogonal to the backbone, so better representations directly translate into stronger gradient guidance. Concretely, we divide the inversion framework into two stages:

153  
154  
155  
156  
157

**① Knowledge-infused Pre-training Stage.** A model  $M$  is first pre-trained to transform structured data  $\mathbf{S}_{\text{pre}}$  (e.g., molecular graphs or 3D point clouds) into comprehensive latent embeddings. This critical stage captures domain-specific features (e.g., chemical and geometric properties) from extensive datasets, ensuring that the learned representations provide robust and meaningful gradients for downstream refined optimization. The training objective can be formulated as:

158  
159  
160

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{pre}}(M_{\theta}(\mathbf{S}_{\text{pre}}), \mathbf{Y}), \quad (2)$$

161

where  $\theta$  denotes model parameters,  $\mathbf{Y}$  represents supervision signals (or self-supervised targets), and  $\mathcal{L}_{\text{pre}}$  is a specialized domain-aware loss function designed for robust optimization.

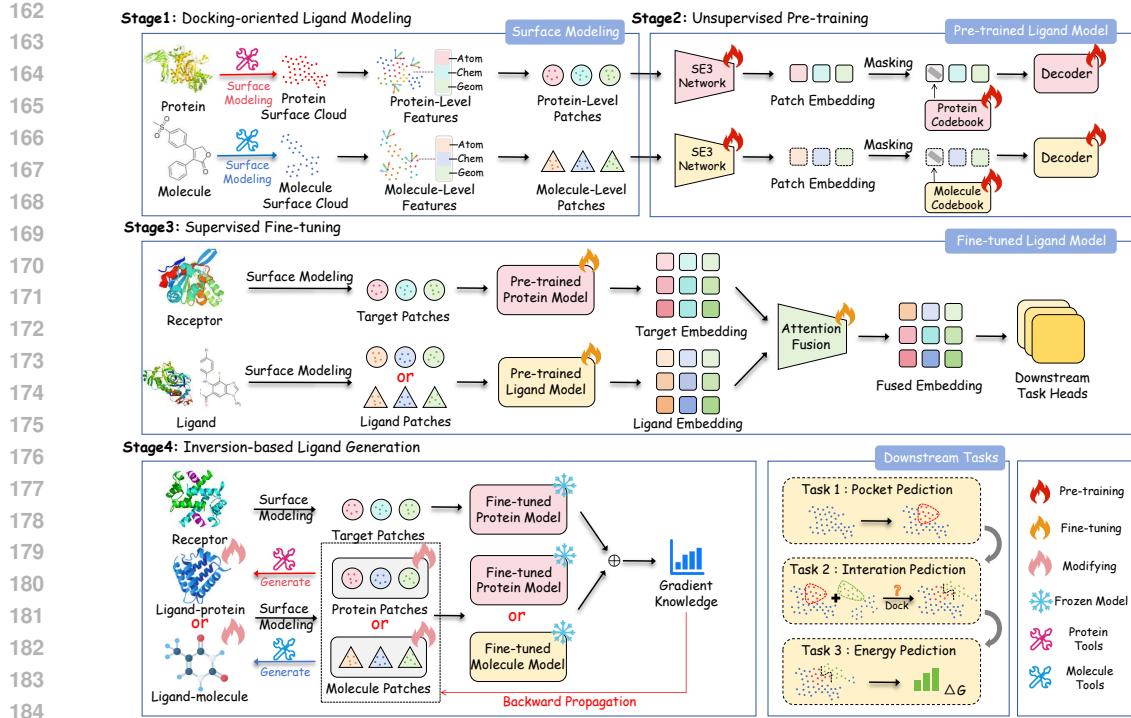


Figure 2: The overview of MagicDock Framework.

**2 Gradient-driven Inversion Stage.** Once pre-trained, the model iteratively refines structures  $\mathbf{S}$  via gradient-based inversion. Starting from an initial  $\mathbf{S}_0$ , the updates follow:

$$\mathbf{S}_{t+1} = \mathbf{S}_t - \eta \nabla_{\mathbf{S}_t} \mathcal{L}(\mathbf{S}_t, C_{\text{domain}}), \quad (3)$$

where  $\mathcal{L}$  is a specialized task-specific objective,  $\nabla_{\mathbf{S}_t}$  represents gradients with respect to the current structure, and  $C_{\text{domain}}$  rigorously enforces structural or biochemical restrictions.

In our framework,  $\mathbf{S}_t$  corresponds to ligand point clouds  $\mathcal{P}_{\text{lig}} = \{\mathbf{f}_i\}_{i=1}^N$ , which are optimized directly with gradient signals. By applying different domain-specific generating constraints  $C_{\text{domain}}$ , the inversion process yields biologically valid protein or small molecule ligands that simultaneously achieve high affinity. We provide a detailed version of the inversion framework in Appendix A.

### 3 METHODOLOGY

We instantiate the four-stage framework introduced in Fig. 2 and propose MagicDock: **Stage1:** Docking-oriented ligand modeling, **Stage2:** Unsupervised pre-training, **Stage3:** Supervised fine-tuning, and **Stage4:** Inversion-based ligand generation. These modules constitute a de novo pipeline for receptor–ligand interactions. The pseudocode of our method is applied in Appendix G

#### 3.1 STAGE 1: DOCKING-ORIENTED LIGAND MODELING

**Motivation.** We aim to develop a unified framework for modeling protein and molecule ligands, capturing shared structural principles while enabling effective modeling the docking process. Surface point clouds provide a compact, meaningful,  $\text{SE}(3)$ -equivariant representation, encoding geometric and chemical binding determinants, ideal for scalable pre-training and generative tasks.

**Surface Point Cloud Modeling.** We transform atomic structures into solvent-accessible surfaces, sampled as point clouds per (Wu & Li (2024)). The molecular surface is defined as the level set of a smooth distance function over atom centers, commonly referred to as the signed distance function (SDF). Candidate points  $\{x_s^i\}$  are upsampled from Gaussian distributions around atomic coordinates  $\{x_a^j\}$  and optimized via gradient descent using:

216  
217  
218  
219  
220

$$\text{SDF}(x_s^i) = -f(x_s^i) \cdot \log \sum_{j=1}^N \exp\left(-\frac{\|x_s^i - x_a^j\|}{\sigma_a^j}\right), \quad f(x_s^i) = \frac{\sum_{j=1}^N \exp(-\|x_s^i - x_a^j\|) \sigma_a^j}{\sum_{j=1}^N \exp(-\|x_s^i - x_a^j\|)}. \quad (4)$$

221 where  $N$  is the number of atoms and  $\sigma_a^j$  denotes the atomic radius. Multi-resolution point clouds are  
222 systematically generated with tailored scaling for protein ligands and molecule ligands, effectively  
223 balancing structural fidelity and computational performance.

224 **Feature Generation.** Each surface point is assigned highly comprehensive composite feature vec-  
225 tors seamlessly integrating chemical, atomic, and geometric characteristics:

$$f(x_i) = \text{concat}(f_{\text{chem}}(x_i), f_{\text{atom}}(x_i), f_{\text{geom}}(x_i)). \quad (5)$$

226 These features encode intrinsic and neighborhood information, tailored to the inherent characteris-  
227 tics of proteins and molecules, ensuring a unified, context-rich representation.

228 **Patch Partitioning.** To reduce computational complexity and facilitate Stage 2’s pre-training, point  
229 clouds are partitioned into patches following (Wu & Li (2024)). Patch centers  $X_c$  are obtained by  
230 farthest point sampling (FPS), and each center systematically groups its  $K$  closest neighbors using  
231  $K$ -nearest neighbor (KNN) search:

$$X_c = \text{FPS}(X_s), \quad X_c \in \mathbb{R}^{\rho M \times 3}, \quad X_p = \text{KNN}(X_c, X_s), \quad X_p \in \mathbb{R}^{\rho M \times K \times 3}. \quad (6)$$

232 This patch-based structure effectively preserves local chemical-atomic-geometric context for dock-  
233 ing and enables highly efficient unsupervised pre-training with discrete latent codes at the patch  
234 level. This optimized, docking-oriented representation seamlessly unifies receptors (i.e. protein)  
235 and ligands (i.e. protein and molecule), balancing biophysical accuracy and computational perfor-  
236 mance for subsequent modeling. We further demonstrate that MagicDock is approximately SE(3)-  
237 equivariant with respect to the initial position and orientation of the receptor in Appendix E.1. Com-  
238 prehensive details of Stage 1 are provided in Appendix C.1.

### 239 3.2 STAGE 2: UNSUPERVISED PRE-TRAINING

240 **Motivation.** In Stage 2, we use the VQ-MAE framework (Wu & Li (2024)) on protein and  
241 molecule datasets, integrating mask autoencoding and quantization to learn transferable representa-  
242 tions. Based on this, we have developed a pre-trained model that can generate high-quality rep-  
243 resentations of proteins and small molecules. (More details are applied in Appendix C.2 and C.3.)

244 **SE(3)-Equivariant Encoding.** The encoder systematically processes local surface patches  $X_p$ ,  
245 with coordinates  $\mathbf{x}_i \in \mathbb{R}^3$  and features  $f_i$ , using SE(3)-equivariant convolutions for significantly  
246 enhanced rigid-body consistency. The comprehensive resulting latent embedding for point  $i$  is:

$$z_i = \sum_{j \in \mathcal{N}(i)} \sum_{l=0}^L R_l(\|\mathbf{x}_{ij}\|) Y_l\left(\frac{\mathbf{x}_{ij}}{\|\mathbf{x}_{ij}\|}\right) \cdot W_l f_j, \quad (7)$$

247 where  $\mathbf{x}_{ij} = \mathbf{x}_j - \mathbf{x}_i$ ,  $R_l(\cdot)$  are radial functions,  $Y_l(\cdot)$  are spherical harmonics, and  $W_l$  are weight  
248 matrices, guaranteeing equivariance under SE(3) spatial transformations.

249 **Masked Reconstruction with Vector Quantization.** Patches are masked at ratio  $\delta = 50\%$ , with  
250 masked and visible sets  $X_{p,m} \in \mathbb{R}^{\delta \rho M \times K \times 3}$  and  $X_{p,vis} \in \mathbb{R}^{(1-\delta) \rho M \times K \times 3}$ . Masked patch tokens  
251 are quantized using a learnable codebook via Gumbel-Softmax relaxation. Visible and quantized  
252 embeddings are systematically decoded to reconstruct: (i) spatial coordinates via

$$\hat{X} = \text{Reshape}(\text{MLP}(H_p^{(L_2)})), \quad \hat{X} \in \mathbb{R}^{\delta \rho M \times K \times 3}, \quad (8)$$

253 and (ii) surface curvature derived from the covariance matrix of each carefully masked patch, with  
254 pseudo-curvatures  $\psi_i = \epsilon_i / \sum_{j=1}^3 \epsilon_j$  accurately predicted by an MLP.

255 **Training Objective.** The loss effectively combines accurate coordinate reconstruction, precise cur-  
256 vature prediction, and comprehensive KL divergence regularization:

$$\mathcal{L} = \nu_1 \mathcal{L}_{\text{rec}}(X_{p,m}, \hat{X}) + \nu_2 \mathcal{L}_{\text{cur}}(\psi, \hat{\psi}) + \nu_3 \mathcal{L}_{\text{KL}}(q(Z_{p,m} | H_{p,m}), p(Z_{p,m})), \quad (9)$$

257 where  $\mathcal{L}_{\text{rec}}$  uses Chamfer distance,  $\mathcal{L}_{\text{cur}}$  uses RMSE, and  $H_{p,m}$  denotes the masked patch embed-  
258 dings before quantization. This design effectively embeds rich ligand surface semantics, ensuring ro-  
259 bust reconstruction fidelity and enhanced geometric consistency. More details are in Appendix C.2.

270 3.3 STAGE 3: SUPERVISED FINE-TUNING  
271272 **Motivation.** The models obtained in Stage 2 provide high-quality representations of protein and  
273 small molecule data, but they cannot be directly used for docking tasks. Therefore, in Stage 3,  
274 we use SE(3)-equivariant attention to capture receptor-ligand dependencies and fine-tune the model  
275 using ground-truth labeling on three progressively downstream tasks: Pocket Prediction, Interaction  
276 Prediction, and Binding-Affinity Regression.277 **Equivariant Attention Fusion.** Receptor and ligand point clouds are independently encoded by  
278 Stage 2’s encoders into latent fields  $Z_r \in \mathbb{R}^{N_r \times d}$  and  $Z_l \in \mathbb{R}^{N_l \times d}$ . Interfacial dependencies are  
279 captured through an SE(3)-equivariant attention followed by a permutation-invariant aggregator:

280 
$$\text{Attn}(Z_r, Z_l) = \text{softmax} \left( \frac{Q_r^{(\ell=0)} (K_l^{(\ell=0)})^\top}{\sqrt{d}} \right) V_l, \quad \tilde{z} = \mathcal{A}(Z_r, \text{Attn}(Z_r, Z_l)) \in \mathbb{R}^{2d}, \quad (10)$$
  
281  
282

283 where  $Q_r^{(\ell=0)} = Z_r^{(\ell=0)} W_Q$ ,  $K_l^{(\ell=0)} = Z_l^{(\ell=0)} W_K$ , and  $V_l = Z_l W_V$  are computed with learnable  
284 matrices  $W_Q, W_K, W_V \in \mathbb{R}^{d \times d}$ . The attention scores are derived from  $\ell = 0$  (scalar) channels to  
285 ensure SE(3)-invariance, while higher-order features in  $V_l$  are transformed via Wigner- $D$  matrices.  
286287 **Multi-Task Supervision.** Three progressively objectives align representations with docking semantics:  
288 (i) Pocket segmentation classifies receptor embeddings  $z_i \in Z_r$  with a loss combining binary  
289 cross-entropy (BCE) and a geometric regularization term; (ii) Interaction prediction uses the fused  
290 representation  $\tilde{z}$  with BCE loss; (iii) Binding affinity regression predicts binding free energy via  
291 mean squared error (MSE). The corresponding objectives are summarized as:

292 
$$\mathcal{L}_{\text{pocket}} = \frac{1}{N_r} \sum_{i=1}^{N_r} \text{BCE}(\hat{y}_i, y_i) + \lambda_p \mathcal{R}_{\text{geom}}, \quad \mathcal{L}_{\Delta G} = \frac{1}{|\mathcal{V}|} \sum_{(r,l) \in \mathcal{V}} (\hat{y}_{\Delta G}(r, l) - \Delta G(r, l))^2. \quad (11)$$
  
293  
294

295 The overall fine-tuning objective is:

296 
$$\mathcal{L}_{\text{FT}} = \alpha \mathcal{L}_{\text{pocket}} + \beta \mathcal{L}_{\text{int}} + \mathcal{L}_{\Delta G}, \quad (12)$$
  
297

298 where  $\mathcal{L}_{\text{int}}$  is the BCE loss for interaction prediction, and  $\alpha, \beta, \lambda_p > 0$  are tuned on validation data.  
299 Task-specific MLPs, the equivariant attention module, and encoders are optimized jointly.300 3.4 STAGE 4: INVERSION-BASED LIGAND GENERATION  
301302 **Motivation.** To convert the docking-aware backbone into a generative model, we use an inversion  
303 framework that optimizes ligands’ structure and feature in a continuous surface embedding space. In  
304 addition, this flexible inversion-based pipeline can effectively utilize the unified differentiable data  
305 structure to generate different categories of ligands in one pipeline.  
306307 **Gradient-driven Inversion.** For a receptor  $R$  and initial ligand  $S_0$ , Stage 3 encoders systematically  
308 produce latent fields  $Z_r$  and  $Z_l$ , fused via SE(3)-equivariant attention into  $\tilde{z}$  (Eq. 10). The com-  
309 posite objective  $\mathcal{F}(S; R, \Theta)$  (Eq. 43) effectively reflects pocket consistency, interaction plausibility,  
310 and binding affinity. Ligand point-cloud coordinates  $\mathbf{x}$  and features  $f$  are rigorously optimized by  
311 descending  $\nabla_{(\mathbf{x}, f)} \mathcal{F}$ , with updates mapped to chemically valid structures via:

312 
$$S^* = \lim_{t \rightarrow \infty} \mathcal{G}_{\text{type}} \left( \Pi_{\mathcal{C}_{\text{valid}}} \left( (\mathbf{x}, f)^t - \eta_t \nabla_{(\mathbf{x}, f)} \mathcal{F}(S_t; R, \Theta) \right) \right), \quad (13)$$
  
313

314 where  $\eta_t$  is the step size,  $\Pi_{\mathcal{C}_{\text{valid}}}$  ensures chemical and geometric validity,  $\mathcal{G}_{\text{type}}$  decodes point clouds  
315 into atomistic graphs, and  $S_t$  denotes the ligand structure at iteration  $t$ . This unifies representation  
316 learning and structure generation for docking-aware ligands. Details are in Appendix C.4.  
317318 4 THEORETICAL ANALYSIS  
319320 To provide a cohesive theoretical foundation, we present MagicDock’s inversion-based framework,  
321 grounded in rigorous analyses in Appendix E. Theorem 1 establishes SE(3)-equivariance across all  
322 stages, ensuring rotational and translational invariance in ligand generation. Theorem 2 proves con-  
323 vergence of the projected gradient descent in the inversion phase to stationary points under smooth-  
ness and boundedness assumptions. Theorem 3 demonstrates theoretical superiority over decoupled

324  
325 Table 1: Performance comparison on the SKEMPI v2 (left) and SAbDab (right). Best results are in  
326 **bold**, second best are underlined. Arrows indicate whether higher ( $\uparrow$ ) or lower ( $\downarrow$ ) values are better.  
327

Method	IMP ( $\uparrow$ )	STA ( $\uparrow$ )	DIV ( $\uparrow$ )	NOV ( $\uparrow$ )	AAR ( $\uparrow$ )
RAbD	21.58/15.12	0.779/0.763	0.728/0.764	0.838/0.829	38.23/38.75
DiffAB	17.65/11.69	0.703/0.736	0.744/0.754	0.820/0.868	41.65/38.52
HSRN	23.19/17.72	0.844/0.809	0.801/0.805	0.877/0.924	45.29/40.63
dyMEAN	15.93/8.55	0.821/0.799	0.805/0.799	<u>0.923/0.953</u>	44.67/44.06
ABDPO	25.17/16.63	0.853/0.833	0.796/0.816	0.905/0.932	46.19/43.53
Abx	<u>28.76/21.80</u>	<u>0.866/0.820</u>	<u>0.812/0.820</u>	0.889/0.950	<u>46.35/44.29</u>
Ours	<b>36.32/27.87</b>	<b>0.874/0.851</b>	<b>0.815/0.824</b>	<b>0.934/0.957</b>	<b>48.73/46.14</b>

336 Table 2: Performance comparison on the PDDBBind2020 (left) and CrossDocked2020 (right).  
337

Method	Vina ( $\downarrow$ )	Affinity ( $\uparrow$ )	STA ( $\uparrow$ )	DIV ( $\uparrow$ )	NOV ( $\uparrow$ )	QED ( $\uparrow$ )
DockStream	-5.51/-5.15	15.13/17.86	0.765/0.780	0.717/0.621	0.985/0.967	0.455/0.401
3D-SBDD	<u>-6.35/-6.24</u>	27.88/28.54	0.853/0.801	0.768/0.701	0.997/0.998	0.503/0.483
liGAN	-6.03/-6.11	22.56/22.15	0.830/0.825	0.772/0.663	0.997/0.997	0.489/0.377
ALIDIFF	<u>-7.21/-6.79</u>	<u>35.74/65.63</u>	<b>0.875/0.833</b>	0.756/0.727	<u>0.998/0.999</u>	0.472/0.464
DRUGFLOW	-7.12/-6.81	33.56/52.35	<u>0.858/0.838</u>	0.764/0.718	<u>0.995/0.999</u>	<u>0.520/0.519</u>
DIFFSBDD	<u>-6.99/-6.86</u>	28.83/35.88	<u>0.867/0.825</u>	<b>0.801/0.705</b>	<b>1.000/0.998</b>	0.511/0.502
Ours	<b>-7.36/-7.02</b>	<b>40.63/60.02</b>	<b>0.866/0.840</b>	<b>0.778/0.730</b>	<b>1.000/1.000</b>	<b>0.552/0.544</b>

347 generate-and-optimize paradigms, with reachable objective sets strictly contained and improved under generator misspecification. Theorems 4, 5 highlight efficiency advantages, outperforming traditional methods in computational cost. Theorem 6 shows information-theoretic optimality via maximized mutual information  $I(X;Y)$ , balancing high output entropy and low conditional uncertainty. Finally, Theorem 7 underscores lower sample complexity,  $O(\log 1/\epsilon)$  for fine-tuning, leveraging pre-training for data efficiency over baselines'  $O(1/\epsilon)$ .  
353

## 354 5 EXPERIMENTS

356 To validate the distinct superiority of MagicDock, we conduct a series of rigorous and comprehensive experiments on diverse datasets for both molecule and protein ligand. We aim to answer:  
357 **Q1** (Effectiveness): Does MagicDock outperform state-of-the-art baselines in generating ligand?  
358 **Q2** (Interpretability): What enables MagicDock to effectively produce high-affinity ligands? **Q3**  
359 (Robustness): Is MagicDock resilient to structural noise, biological variability, and does it exhibit  
360 reliable convergence? **Q4** (Efficiency): Does MagicDock achieve superior trade-offs in runtime,  
361 resource usage, scalability, and data efficiency during ligand generation?  
362

### 363 5.1 OVERALL PERFORMANCE (Q1)

365 To answer **Q1**, we conducted a systematic evaluation of the effectiveness of ligand design for two  
366 scenarios. In the following, we list the detailed experimental settings for these two scenarios.  
367

368 **Datasets & Baselines & Evaluation Metrics.** We utilize tailored datasets, baselines, and evaluation  
369 metrics for both protein and small molecule ligand design, with details provided in Appendix D.  
370 Furthermore, we have made fair adjustments for all baseline methods to adapt to challenging de  
371 novo ligand design scenario, as comprehensively detailed in the Appendix H.  
372

373 **Experimental Results.** As shown in Table 1 and Table 2, MagicDock outperforms all baselines  
374 on both benchmarks, achieving better binding affinity, stability, diversity and novelty, with super  
375 rior AAR for proteins and higher QED for molecules. MagicDock's experimental affinity perfor  
376 mance is slightly lower than AliDiff's due to the latter's direct exact energy alignment via preference  
377 optimization on pre-trained diffusion models, which efficiently biases toward high-affinity sample  
378 distributions but sacrifices generation diversity and incurs substantial computational overhead. Hy  
379 perparameter details are applied in Appendix L.  
380

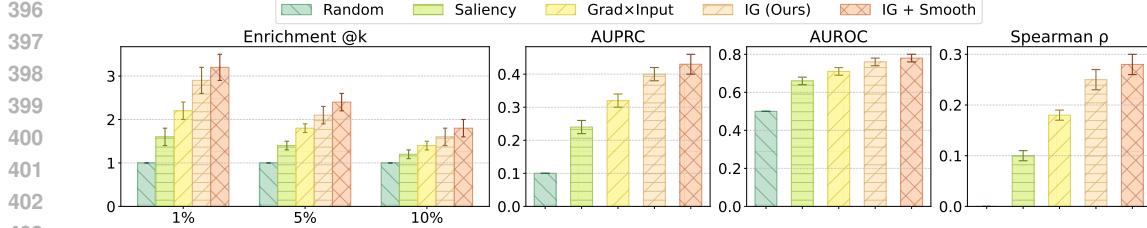
378 5.2 INTERPRETABILITY STUDY (Q2)  
379

380 Having established the remarkable effectiveness in Q1, we now turn to Q2 and investigate why  
381 MagicDock can produce high-affinity ligands. Specifically, we first analyzed the contribution of  
382 each stage to the final result and whether the gradient of the inversion target provided meaningful  
383 biological localization signals. In addition, we also visualized the process of ligand generation.

384 **Ablation Study.** To disentangle the role of each stage, we conduct module-wise ablations: (i)  
385 *w/o Stage 1*: replacing surface point clouds with raw atom graphs; (ii) *w/o Stage 2*: removing  
386 unsupervised pre-training; (iii) *w/o Stage 3*: disabling supervised fine-tuning; (iv) *w/o Stage 4*:  
387 substituting gradient-guided inversion with exhaustive element-type search. Fig. 6a, 6b show that  
388 each stage plays a critical role, with the full pipeline achieving the best performance.

389 **Gradient Attribution & Localization.** We further ask whether gradients from the composite inversion  
390 objective  $\mathcal{F}(S; R, \Theta)$  (Appendix C.4) effectively localize on biologically meaningful binding  
391 sites learned in Stage 3. We compute Integrated Gradients on receptor and ligand surfaces and  
392 systematically evaluate whether high-attribution regions align with ground-truth interfaces. Fig. 3  
393 demonstrates that inversion not only optimizes affinity but also consistently highlights mechanisti-  
394 cally significant relevant regions, addressing Q2.

395



404 Figure 3: Performance comparison of attribution methods for interpretability, evaluated on Enrich-  
405 ment @k (1%, 5%, 10%), AUPRC, AUROC, and Spearman correlation.

406

407 **Local Perturbation Consistency.** We assess whether gradient attributions predict energetic effects  
408 of local edits by introducing small perturbations near high-attribution sites and comparing predicted  
409 with actual  $\Delta\mathcal{F}$ . This tests whether attributions enable actionable refinement. As shown in Fig. 8,  
410 our method outperforms saliency and GradxInput, approaching physics-based Rosetta evaluations.

411 **Case Study.** To evaluate the significant impact of constraints in MagicDock’s inversion-based ligand  
412 generation process, we generated two ligands for the target receptor. At each iteration, we systemati-  
413 cally assessed binding affinity, as shown in Fig. 4 and Fig. 12. Additionally, we have listed a series  
414 of molecular ligands generated by MagicDock, as shown in Fig. 11 and Fig. 13 in Appendix K.

415

## 416 5.3 ROBUSTNESS STUDY (Q3)

417 Having established remarkable interpretability, we next address Q3: whether MagicDock is con-  
418 sistently robust to input variations and biological uncertainty. We design three sets of experiments  
419 systematically probing artificial noise, realistic receptor perturbations, and convergence properties.

420

421 **Geometric and Feature Noise.** We assess robustness on 100 receptor–ligand pairs by adding Gaus-  
422 sian coordinate noise and feature dropout to receptor surfaces, then comparing ligand generation on  
423 noisy versus clean inputs (Fig. 7). Results show that MagicDock consistently outperforms baselines  
424 under perturbations, demonstrating stable performance and robustness to biological uncertainty.

425

426 **Conformational and Mutational Variability.** We further evaluate remarkable robustness to bi-  
427 ological variability by subjecting 100 receptors to perturbations, including conformer ensembles,  
428 single-point mutations, and combined variants. These perturbed structures are systematically pro-  
429 cessed through the inversion pipeline. Performance degradation is assessed via IMP for protein  
430 baselines and High-Affinity for molecule baselines. As shown in Fig. 9, statistical tests compare  
431 conservative and non-conservative mutations, illustrating MagicDock’s exceptional robustness.

432

433 **Convergence Study.** The convergence of MagicDock’s gradient-driven inversion stage is rigorously  
434 established in Appendix E.2. The proof systematically demonstrates that, under standard assump-

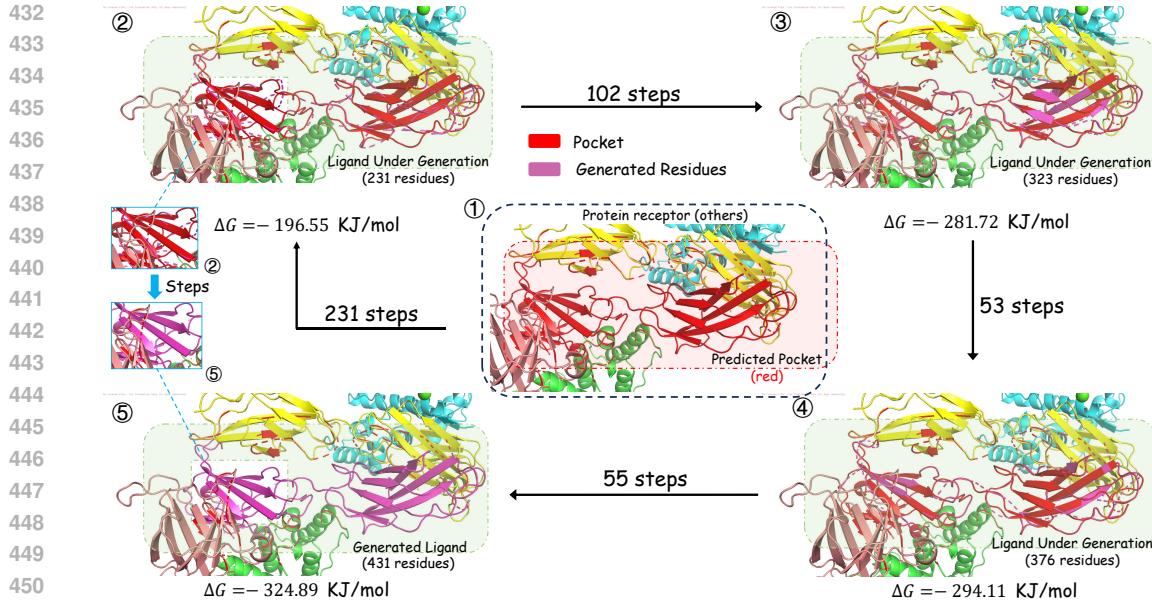


Figure 4: Visualization of an example of generated protein-ligand complexes.

tions for projected gradient descent, the ligand sequence  $P_{\text{lig}}^t$  converges to a stationary point of the composite objective  $\mathcal{F}$ , with  $\min_t |\nabla \mathcal{F}(P_{\text{lig}}^t)|^2 \rightarrow 0$  as  $t \rightarrow \infty$ . This guarantees that MagicDock consistently optimizes ligands with exceptionally high reliability.

#### 5.4 EFFICIENCY STUDY (Q4)

Finally, we address Q4 by systematically quantifying whether MagicDock achieves a highly favorable trade-off between computational cost and accuracy. Having shown that the framework is both interpretable and robust, we now analyze runtime efficiency and comprehensive scalability.

**Runtime and Resource Usage.** We benchmark MagicDock against baselines using 100 receptors, generating one ligand per receptor under identical hardware. Wall-clock time and memory usage are assessed per computational stage. Fig. 10a demonstrates MagicDock’s reduced runtime while Fig. 10b highlights lower peak memory consumption, underscoring its resource efficiency.

**Scalability.** Scalability is evaluated across nine settings varying receptor sizes and ligand complexities using 100 receptors from CrossDocked2020 and SAbDab. Metrics include per-iteration latency, iterations-to-converge, memory footprint, and scaling exponent  $\gamma$  from  $T \propto N^\gamma$ . Table 5 and Table 6 shows MagicDock’s sub-quadratic scaling ( $\gamma = 1.4$ ), enabling real-time deployment.

**Data Efficiency.** As established in Appendix E.6, MagicDock’s inversion framework achieves  $\epsilon$ -accuracy in supervised fine-tuning with only  $O(1/\epsilon)$  samples (up to logarithmic confidence factors), compared to  $O(1/\epsilon^2)$  for GANs and  $O(T/\epsilon^2)$  for diffusion models. This linear-in- $1/\epsilon$  sample complexity, enabled by pre-training’s strong convexity and low effective dimension, effectively ensures robust generalization in challenging data-scarce docking tasks with limited annotated complexes.

## 6 CONCLUSION

In this study, we introduced MagicDock, an inversion-based framework unifying generation and optimization in a streamlined, docking-driven workflow for de novo ligand design. Using surface point-cloud modeling, SE(3)-equivariant pretraining, and docking-oriented fine-tuning, MagicDock eliminates external priors for robust de novo design of protein and small-molecule ligands. Experiments show strong effectiveness and efficiency. These results position inversion-based docking as a versatile paradigm overcoming traditional limitations for practical ligand design. Limitations and future work are applied in Appendix F.

486 REPRODUCIBILITY STATEMENT  
487

488 To ensure reproducibility, we provide detailed MagicDock architecture descriptions, theoretical  
489 analysis, and objectives in Sec. 3 and Sec. 5 and Appendix A to Appendix D. Datasets for pre-  
490 training and evaluation, including processing, splits, and sources, are in Appendix E. Hyperparam-  
491 eters, training, and ablations are in Appendix L. Codes and other necessary materials are provided in  
492 our supplementary materials.

493  
494 REFERENCES  
495

496 Maryam Abbasi, Beatriz P Santos, Tiago C Pereira, Raul Sofia, Nelson RC Monteiro, Carlos JV  
497 Simões, Rui MM Brito, Bernardete Ribeiro, José L Oliveira, and Joel P Arrais. Designing opti-  
498 mized drug candidates with generative adversarial network. *Journal of cheminformatics*, 14(1):  
499 40, 2022.

500 P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*.  
501 Princeton University Press, 2008.

502 Jared Adolf-Bryfogle, Oleks Kalyuzhniy, Michael Kubitz, Brian D Weitzner, Xiaozhen Hu, Yumiko  
503 Adachi, William R Schief, and Roland L Dunbrack Jr. Rosettaantibodydesign (rabd): A general  
504 framework for computational antibody design. *PLoS computational biology*, 14(4):e1006112,  
505 2018.

506 Amir Beck. *First-order methods in optimization*. SIAM, 2017.

507 Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow  
508 network based generative models for non-iterative diverse candidate generation. *Advances in  
509 neural information processing systems*, 34:27381–27394, 2021.

510 Nurken Berdigaliyev and Mohamad Aljofan. An overview of drug discovery and development.  
511 *Future medicinal chemistry*, 12(10):939–947, 2020.

512 Noémie Bergues, Arthur Carré, Paul Join-Lambert, Brice Hoffmann, Arnaud Blondel, and Hamza  
513 Tajmouati. Template-guided 3d molecular pose generation via flow matching and differentiable  
514 optimization. *arXiv preprint arXiv:2506.06305*, 2025.

515 Christian Blum and Andrea Roli. Metaheuristics in combinatorial optimization: Overview and  
516 conceptual comparison. *ACM computing surveys (CSUR)*, 35(3):268–308, 2003.

517 Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking  
518 models for de novo molecular design. *Journal of chemical information and modeling*, 59(3):  
519 1096–1108, 2019.

520 Yue Cao, Payel Das, Vijil Chenthamarakshan, Pin-Yu Chen, Igor Melnyk, and Yang Shen. Fold2seq:  
521 A joint sequence (1d)-fold (3d) embedding-based generative model for protein design. In *Inter-  
522 national Conference on Machine Learning*, pp. 1261–1271. PMLR, 2021.

523 Binghong Chen, Tianzhe Wang, Chengtao Li, Hanjun Dai, and Le Song. Molecule optimization by  
524 explainable evolution. In *International conference on learning representation (ICLR)*, 2021.

525 Rong Chen, Li Li, and Zhiping Weng. Zdock: an initial-stage protein-docking algorithm. *Proteins-  
526 structure Function Bioinformatics*, 52(1):80–87, 2010.

527 Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Dif-  
528 fusion steps, twists, and turns for molecular docking. *arXiv preprint arXiv:2210.01776*, 2022.

529 Nicola De Cao and Thomas Kipf. Molgan: An implicit generative model for small molecular graphs.  
530 *arXiv preprint arXiv:1805.11973*, 2018.

531 Cyril Dominguez, Rolf Boelens, and Alexandre MJJ Bonvin. Haddock: a protein- protein docking  
532 approach based on biochemical or biophysical information. *Journal of the American Chemical  
533 Society*, 125(7):1731–1737, 2003.

540 R Evans, M O'Neill, A Pritzel, N Antropova, A Senior, T Green, A Žídek, R Bates, S Blackwell,  
 541 J Yim, et al. Protein complex prediction with alphafold-multimer. *biorxiv* [preprint](2022), 2022.  
 542

543 Richard Evans, Michael O'Neill, Alexander Pritzel, Natasha Antropova, Andrew Senior, Tim Green,  
 544 Augustin Žídek, Russ Bates, Sam Blackwell, Jason Yim, et al. Protein complex prediction with  
 545 alphafold-multimer. *biorxiv*, pp. 2021–10, 2021.

546 Todd JA Ewing, Shingo Makino, A Geoffrey Skillman, and Irwin D Kuntz. Dock 4.0: search  
 547 strategies for automated molecular docking of flexible molecule databases. *Journal of computer-  
 548 aided molecular design*, 15(5):411–428, 2001.  
 549

550 Daniel S Fabricant and Norman R Farnsworth. The value of plants used in traditional medicine for  
 551 drug discovery. *Environmental health perspectives*, 109(suppl 1):69–75, 2001.

552 Paul G Francoeur, Tomohide Masuda, Jocelyn Sunseri, Andrew Jia, Richard B Iovanisci, Ian Snyder,  
 553 and David R Koes. Three-dimensional convolutional neural networks and a cross-docked data set  
 554 for structure-based drug design. *Journal of chemical information and modeling*, 60(9):4200–  
 555 4215, 2020.  
 556

557 Tianfan Fu, Cao Xiao, and Jimeng Sun. Core: Automatic molecule optimization using copy &  
 558 refine strategy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp.  
 559 638–645, 2020.

560 Tianfan Fu, Wenhao Gao, Cao Xiao, Jacob Yasonik, Connor W Coley, and Jimeng Sun. Differentiable  
 561 scaffolding tree for molecular optimization. *arXiv preprint arXiv:2109.10469*, 2021a.

562 Tianfan Fu, Cao Xiao, Xinhao Li, Lucas M Glass, and Jimeng Sun. Mimosa: Multi-constraint  
 563 molecule sampling for molecule optimization. In *Proceedings of the AAAI Conference on Artifi-  
 564 cial Intelligence*, volume 35, pp. 125–133, 2021b.

565 Octavian-Eugen Ganea, Xinyuan Huang, Charlotte Bunne, Yatao Bian, Regina Barzilay, Tommi  
 566 Jaakkola, and Andreas Krause. Independent se (3)-equivariant models for end-to-end rigid protein  
 567 docking. *arXiv preprint arXiv:2111.07786*, 2021.

568 Ziqi Gao, Tao Feng, Jiaxuan You, Chenyi Zi, Yan Zhou, Chen Zhang, and Jia Li. Deep reinforcement  
 569 learning for modelling protein complexes. *arXiv preprint arXiv:2405.02299*, 2024a.

570 Ziqi Gao, Zijing Liu, Yu Li, and Jia Li. Towards stable representations for protein interface predic-  
 571 tion. *Advances in Neural Information Processing Systems*, 37:73079–73097, 2024b.

572 Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato,  
 573 Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel,  
 574 Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven contin-  
 575 uous representation of molecules. *ACS central science*, 4(2):268–276, 2018.

576 Sai Krishna Gottipati, Boris Sattarov, Sufeng Niu, Yashaswi Pathak, Haoran Wei, Shengchao Liu,  
 577 Simon Blackburn, Karam Thomas, Connor Coley, Jian Tang, et al. Learning to navigate the  
 578 synthetically accessible chemical space using reinforcement learning. In *International conference  
 579 on machine learning*, pp. 3668–3679. PMLR, 2020.

580 Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained bayesian optimization for  
 581 automatic chemical design using variational autoencoders. *Chemical science*, 11(2):577–586,  
 582 2020.

583 Siyi Gu, Minkai Xu, Alexander Powers, Weili Nie, Tomas Geffner, Karsten Kreis, Jure Leskovec,  
 584 Arash Vahdat, and Stefano Ermon. Aligning target-aware molecule diffusion models with ex-  
 585 act energy optimization. *Advances in Neural Information Processing Systems*, 37:44040–44063,  
 586 2024.

587 Gabriel Lima Guimaraes, Benjamin Sanchez-Lengeling, Carlos Outeiral, Pedro Luis Cunha Farias,  
 588 and Alán Aspuru-Guzik. Objective-reinforced generative adversarial networks (organ) for se-  
 589 quence generation models. *arXiv preprint arXiv:1705.10843*, 2017.

594 Jeff Guo, Jon Paul Janet, Matthias R Bauer, Eva Nittinger, Kathryn A Giblin, Kostas Papadopoulos,  
 595 Alexey Voronov, Atanas Patronov, Ola Engkvist, and Christian Margreitter. Dockstream: a  
 596 docking wrapper to enhance de novo molecular design. *Journal of cheminformatics*, 13(1):89,  
 597 2021.

598 Thomas A Halgren, Robert B Murphy, Richard A Friesner, Hege S Beard, Leah L Frye, W Thomas  
 599 Pollard, and Jay L Banks. Glide: a new approach for rapid, accurate docking and scoring. 2. en-  
 600 richment factors in database screening. *Journal of medicinal chemistry*, 47(7):1750–1759, 2004.

601 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. URL  
 602 <https://arxiv.org/abs/2006.11239>.

603 Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley,  
 604 Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics data commons: Machine learning  
 605 datasets and tasks for drug discovery and development. *arXiv preprint arXiv:2102.09548*, 2021.

606 John Ingraham, Vikas Garg, Regina Barzilay, and Tommi Jaakkola. Generative models for graph-  
 607 based protein design. *Advances in neural information processing systems*, 32, 2019.

608 Moksh Jain, Sharath Chandra Raparthi, Alex Hernández-García, Jarrid Rector-Brooks, Yoshua Ben-  
 609 gio, Santiago Miret, and Emmanuel Bengio. Multi-objective gflownets. In *International confer-  
 610 ence on machine learning*, pp. 14631–14653. PMLR, 2023.

611 Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for  
 612 the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.

613 Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle  
 614 points efficiently. In *International conference on machine learning*, pp. 1724–1732. PMLR, 2017.

615 Wengong Jin, Jeremy Wohlwend, Regina Barzilay, and Tommi Jaakkola. Iterative refinement graph  
 616 neural network for antibody sequence-structure co-design. *arXiv preprint arXiv:2110.04624*,  
 617 2021.

618 Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Antibody-antigen docking and design via  
 619 hierarchical structure refinement. In *International Conference on Machine Learning*, pp. 10217–  
 620 10227. PMLR, 2022.

621 Xiangzhe Kong, Wenbing Huang, and Yang Liu. End-to-end full-atom antibody design. *arXiv  
 622 preprint arXiv:2302.00203*, 2023.

623 Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos, Jeff  
 624 Schneider, and Eric Xing. Chembo: Bayesian optimization of small organic molecules with syn-  
 625 thezizable recommendations. In *International Conference on Artificial Intelligence and Statistics*,  
 626 pp. 3393–3403. PMLR, 2020.

627 Dima Kozakov, David R Hall, Bing Xia, Kathryn A Porter, Dzmitry Padhorny, Christine Yueh,  
 628 Dmitri Beglov, and Sandor Vajda. The cluspro web server for protein–protein docking. *Nature  
 629 protocols*, 12(2):255–278, 2017.

630 Seul Lee, Jaehyeong Jo, and Sung Ju Hwang. Exploring chemical space with score-based out-of-  
 631 distribution generation. In *International Conference on Machine Learning*, pp. 18872–18892.  
 632 PMLR, 2023.

633 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching  
 634 for generative modeling, 2023. URL <https://arxiv.org/abs/2210.02747>.

635 Ge Liu, Haoyang Zeng, Jonas Mueller, Brandon Carter, Ziheng Wang, Jonas Schilz, Geraldine  
 636 Horny, Michael E Birnbaum, Stefan Ewert, and David K Gifford. Antibody complementarity  
 637 determining region design using high-capacity machine learning. *Bioinformatics*, 36(7):2126–  
 638 2133, 2020.

639 Huaqing Liu, Peiyi Chen, Xiaochen Zhai, Ku-Geng Huo, Shuxian Zhou, Lanqing Han, and Guoxin  
 640 Fan. Ppb-affinity: Protein-protein binding affinity dataset for ai-based protein drug discovery.  
 641 *Scientific data*, 11(1):1316, 2024.

648 Qi Liu, Miltiadis Allamanis, Marc Brockschmidt, and Alexander Gaunt. Constrained graph varia-  
 649 tional autoencoders for molecule design. *Advances in neural information processing systems*, 31,  
 650 2018.

651 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and  
 652 transfer data with rectified flow, 2022. URL <https://arxiv.org/abs/2209.03003>.

654 Wei Lu, Qifeng Wu, Jixian Zhang, Jiahua Rao, Chengtao Li, and Shuangjia Zheng. Tankbind:  
 655 Trigonometry-aware neural networks for drug-protein binding structure prediction. *Advances in  
 656 neural information processing systems*, 35:7236–7249, 2022.

657 Shitong Luo, Jiaqi Guan, Jianzhu Ma, and Jian Peng. A 3d generative model for structure-based  
 658 drug design. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.

660 Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-specific  
 661 antibody design and optimization with diffusion-based generative models for protein structures.  
 662 *Advances in Neural Information Processing Systems*, 35:9754–9767, 2022.

663 Sergey Lyskov and Jeffrey J Gray. The rosettadock server for local protein–protein docking. *Nucleic  
 664 acids research*, 36(suppl\_2):W233–W238, 2008.

666 Biao Ma, Kei Terayama, Shigeyuki Matsumoto, Yuta Isaka, Yoko Sasakura, Hiroaki Iwata, Mitsugu  
 667 Araki, and Yasushi Okuno. Structure-based de novo molecular generator combined with artificial  
 668 intelligence and docking simulations. *Journal of Chemical Information and Modeling*, 61(7):  
 669 3304–3313, 2021.

670 Nina Mazyavkina, Sergey Sviridov, Sergei Ivanov, and Evgeny Burnaev. Reinforcement learning  
 671 for combinatorial optimization: A survey. *Computers & Operations Research*, 134:105400, 2021.

673 Garrett M Morris, Ruth Huey, and Arthur J Olson. Using autodock for ligand-receptor docking.  
 674 *Current protocols in bioinformatics*, 24(1):8–14, 2008.

675 Henry Moss, David Leslie, Daniel Beck, Javier Gonzalez, and Paul Rayson. Boss: Bayesian op-  
 676 timization over string spaces. *Advances in neural information processing systems*, 33:15476–  
 677 15486, 2020.

679 AkshatKumar Nigam, Pascal Friederich, Mario Krenn, and Alán Aspuru-Guzik. Augmenting ge-  
 680 netic algorithms with deep neural networks for exploring the chemical space. *arXiv preprint  
 681 arXiv:1909.11655*, 2019.

682 Yifan Niu, Ziqi Gao, Tingyang Xu, Yang Liu, Yatao Bian, Yu Rong, Junzhou Huang, and Jia Li.  
 683 Inversiongnn: A dual path network for multi-property molecular optimization. *arXiv preprint  
 684 arXiv:2503.01488*, 2025.

685 Xingang Peng, Shitong Luo, Jiaqi Guan, Qi Xie, Jian Peng, and Jianzhu Ma. Pocket2mol: Effi-  
 686 cient molecular sampling based on 3d protein pockets. In *International conference on machine  
 687 learning*, pp. 17644–17655. PMLR, 2022.

689 Keyue Qiu, Yuxuan Song, Jie Yu, Hongbo Ma, Ziyao Cao, Zhilong Zhang, Yushuai Wu, Mingyue  
 690 Zheng, Hao Zhou, and Wei-Ying Ma. Empower structure-based molecule optimization with gra-  
 691 dient guided bayesian flow networks. *arXiv preprint arXiv:2411.13280*, 2024.

692 Matthew Ragoza, Tomohide Masuda, and David Ryan Koes. Generating 3d molecules conditional  
 693 on receptor binding sites with deep generative models. *Chemical science*, 13(9):2701–2713, 2022.

695 R Tyrrell Rockafellar and Roger JB Wets. *Variational analysis*. Springer, 1998.

696 Koichiro Saka, Taro Kakuzaki, Shoichi Metsugi, Daiki Kashiwagi, Kenji Yoshida, Manabu Wada,  
 697 Hiroyuki Tsunoda, and Reiji Teramoto. Antibody design using lstm based deep generative model  
 698 from phage display library for affinity maturation. *Scientific reports*, 11(1):5852, 2021.

700 Arne Schneuing, Charles Harris, Yuanqi Du, Kieran Didi, Arian Jamasb, Ilia Igashov, Weitao Du,  
 701 Carla Gomes, Tom L Blundell, Pietro Lio, et al. Structure-based drug design with equivariant  
 diffusion models. *Nature Computational Science*, 4(12):899–909, 2024.

702 Arne Schneuing, Ilia Igashov, Adrian W Doppelstein, Thomas Castiglione, Michael M Bronstein,  
 703 and Bruno Correia. Multi-domain distribution learning for de novo drug design. In *The Thirteenth*  
 704 *International Conference on Learning Representations*, 2025.

705 Chence Shi, Minkai Xu, Zhaocheng Zhu, Weinan Zhang, Ming Zhang, and Jian Tang.  
 706 Graphaf: a flow-based autoregressive model for molecular graph generation. *arXiv preprint*  
 707 *arXiv:2001.09382*, 2020.

708 Miha Skalic, José Jiménez, Davide Sabbadin, and Gianni De Fabritiis. Shape-based generative  
 709 modeling for de novo drug design. *Journal of chemical information and modeling*, 59(3):1205–  
 710 1214, 2019.

711 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022. URL  
 712 <https://arxiv.org/abs/2010.02502>.

713 Niclas Ståhl, Goran Falkman, Alexander Karlsson, Gunnar Mathiason, and Jonas Bostrom. Deep re-  
 714 enforcement learning for multiparameter optimization in de novo drug design. *Journal of chemical*  
 715 *information and modeling*, 59(7):3166–3176, 2019.

716 Doug Tischer, Sidney Lisanza, Jue Wang, Runze Dong, Ivan Anishchenko, Lukas F Milles, Sergey  
 717 Ovchinnikov, and David Baker. Design of proteins presenting discontinuous functional sites using  
 718 deep learning. *Biorxiv*, pp. 2020–11, 2020.

719 Marcel L. Verdonk, Jason C. Cole, Michael J. Hartshorn, Christopher W. Murray, and Richard D.  
 720 Taylor. Improved protein–ligand docking using gold. *Proteins: Structure, Function, and Bioin-*  
 721 *formatics*, 2003.

722 Renxiao Wang, Xueliang Fang, Yipin Lu, and Shaomeng Wang. The pdbsbind database: Collec-  
 723 tion of binding affinities for protein–ligand complexes with known three-dimensional structures.  
 724 *Journal of medicinal chemistry*, 47(12):2977–2980, 2004.

725 Zichao Wang, Weili Nie, Zhuoran Qiao, Chaowei Xiao, Richard Baraniuk, and Anima Anandkumar.  
 726 Retrieval-based controllable molecule generation. *arXiv preprint arXiv:2208.11126*, 2022.

727 Fang Wu and Stan Z Li. Surface-vqmae: Vector-quantized masked auto-encoders on molecular  
 728 surfaces. In *Forty-first International Conference on Machine Learning*, 2024.

729 Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. Mars:  
 730 Markov molecular sampling for multi-objective drug discovery. *arXiv preprint arXiv:2103.10432*,  
 731 2021.

732 Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geo-  
 733 metric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*,  
 734 2022.

735 Yumeng Yan, Huanyu Tao, Jiahua He, and Sheng-You Huang. The hdock server for integrated  
 736 protein–protein docking. *Nature protocols*, 15(5):1829–1852, 2020.

737 Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy  
 738 network for goal-directed molecular graph generation. *Advances in neural information processing*  
 739 *systems*, 31, 2018.

740 Xiangxin Zhou, Dongyu Xue, Ruizhe Chen, Zaixiang Zheng, Liang Wang, and Quanquan Gu.  
 741 Antigen-specific antibody design via direct energy-based preference optimization. *Advances in*  
 742 *Neural Information Processing Systems*, 37:120861–120891, 2024.

743 Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N Zare, and Patrick Riley. Optimization of  
 744 molecules via deep reinforcement learning. *Scientific reports*, 9(1):10752, 2019.

745 Tian Zhu, Milong Ren, and Haicang Zhang. Antibody design using a score-based diffusion model  
 746 guided by evolutionary, physical and geometric constraints. In *Forty-first International Confer-*  
 747 *ence on Machine Learning*, 2024.

748

756

## APPENDIX

757

758

## 759 TABLE OF CONTENT

760

761

<b>A The Inversion Framework in Details</b> . . . . .	<b>17</b>
A.1 Model Training . . . . .	17
A.2 Inversion Phase . . . . .	17
<b>B Related Works in Details</b> . . . . .	<b>18</b>
<b>C Methodology in Details</b> . . . . .	<b>19</b>
C.1 Stage 1: Docking-oriented Ligand Modeling . . . . .	19
C.2 Stage 2: Unsupervised Pre-training . . . . .	21
C.3 Stage 3: Supervised Fine-tuning . . . . .	22
C.4 Stage 4: Inversion-based Ligand Generation . . . . .	23
<b>D Experiment in Details</b> . . . . .	<b>25</b>
D.1 Computation Resource . . . . .	25
D.2 Datasets . . . . .	25
D.3 Baseline Models. . . . .	26
D.4 Evaluation Metrics . . . . .	28
D.5 More experiment figures & tables . . . . .	31
<b>E Theoretical Analysis in Details</b> . . . . .	<b>33</b>
E.1 Proof of SE(3)-Equivariance for MagicDock . . . . .	33
E.2 Proof of Convergence for MagicDock. . . . .	35
E.3 Proof of Theoretical Superiority for Inversion Framework . . . . .	36
E.4 Proof of Efficiency for Inversion Framework . . . . .	39
E.5 Proof of Information-Theoretic Superiority for Inversion Framework . . . . .	40
E.6 Proof of Sample Complexity Advantage for Inversion Framework . . . . .	41
<b>F Limitations and Future Work</b> . . . . .	<b>42</b>
<b>G Algorithms in Pseudo Code</b> . . . . .	<b>43</b>
G.1 Stage 1: Docking-oriented Ligand Modeling . . . . .	43
G.2 Stage 2: Unsupervised pre-training . . . . .	44
G.3 Stage 3: Supervised fine-tuning. . . . .	44
G.4 Stage 4: Inversion-based Ligand Generation . . . . .	45
G.5 Overall pipeline . . . . .	45
<b>H Adaptations to Baselines for Fair Comparison</b> . . . . .	<b>46</b>
<b>I Potential Societal Impacts</b> . . . . .	<b>47</b>
<b>J GenAI Usage Disclosure</b> . . . . .	<b>47</b>
<b>K Visualization of Generated Ligand</b> . . . . .	<b>47</b>
K.1 Protein ligand . . . . .	47
K.2 Molecule ligand . . . . .	47

808

809

810	<b>L</b>	<b>Hyperparameters in Details</b>	<b>49</b>
811		L.1 Protein Baselines	50
812		L.2 Molecule Baselines	51
813		L.3 MagicDock	52
814			
815			
816			
817			
818			
819			
820			
821			
822			
823			
824			
825			
826			
827			
828			
829			
830			
831			
832			
833			
834			
835			
836			
837			
838			
839			
840			
841			
842			
843			
844			
845			
846			
847			
848			
849			
850			
851			
852			
853			
854			
855			
856			
857			
858			
859			
860			
861			
862			
863			

864 A THE INVERSION FRAMEWORK IN DETAILS  
865866 A.1 MODEL TRAINING  
867

868 Recent advances in generative modeling have increasingly embraced the Inversion framework as  
869 a powerful approach, which reverses traditional data flow by reconstructing or optimizing struc-  
870 tures from latent representations or gradients, as demonstrated in InversionGNN( Niu et al. (2025)).  
871 This framework is particularly well-suited for our docking-optimized design of proteins and small  
872 molecules, offering a flexible paradigm for zero-start generation. Let us guide you through its struc-  
873 ture step by step, focusing on the unsupervised pre-training and supervised fine-tuning phases that  
874 establish a robust foundation for subsequent inversion-based generation.

875 ① Unsupervised Pre-Training. This phase lays the foundation by initializing the Inversion frame-  
876 work through an unsupervised training process that encodes input data into a latent representation,  
877 capturing transferable geometric and chemical priors without relying on labeled data. Generally, it  
878 involves an encoder that transforms structured input data—comprising  $n$  elements and their rela-  
879 tionships—into hidden embeddings  $\mathbf{z} \in \mathbb{R}^d$ , updated iteratively using a propagation rule. The general  
880 update can be expressed as:

$$881 \mathbf{z}_{t+1} = \text{Prop}(\mathbf{z}_t, \{\mathbf{z}_v \mid v \in \mathcal{N}(u)\}, \theta), \quad (14)$$

883 where  $\theta$  represents trainable parameters, and Prop denotes a propagation function tailored to the data  
884 structure, which may include graphs, point clouds, or other relational formats. To ensure adaptabil-  
885 ity, the latent space is refined with feedback from an objective function, requiring differentiability  
886 for gradient-based optimization.

887 In our work, this phase employs a pre-trained SE(3)-equivariant encoder to process the interaction  
888 graph derived from 3D point clouds, generating embeddings  $\mathbf{h}_u \in \mathbb{R}^d$  as:

$$889 \mathbf{h}_u^{(l)} = \text{Prop}_{\text{SE}(3)}\left(\mathbf{h}_u^{(l-1)}, \left\{\mathbf{h}_v^{(l-1)} \mid v \in \mathcal{N}(u)\right\}\right), \quad (15)$$

892 where the SE(3)-equivariant propagation weights are optimized using self-supervised objectives,  
893 such as masked reconstruction and vector quantization, to embed rich surface semantics. This es-  
894 tablishes a robust, docking-aware representational backbone by learning from extensive datasets  
895 of proteins and small molecules, ensuring geometric consistency and chemical plausibility without  
896 explicit supervision.

897 ② Supervised Fine-Tuning. Building on the unsupervised pre-training, this phase refines the  
898 model with task-specific supervision to align representations with docking semantics, incorpor-  
899 ating ground-truth labels from protein–ligand complexes. The fine-tuning calibrates the pre-trained  
900 encoder to capture interfacial dependencies and binding signals, using objectives like pocket seg-  
901 mentation, interaction prediction, and binding affinity regression.

902 The receptor and ligand point clouds are encoded into latent fields  $\mathbf{Z}_r$  and  $\mathbf{Z}_l$ , fused via cross-  
903 attention to model interactions. The multi-task loss integrates these objectives, with gradients from  
904 docking energy  $E_{\text{dock}}$ , defined as  $\nabla_{\mathbf{h}_u} E_{\text{dock}} = \frac{\partial E_{\text{dock}}}{\partial \mathbf{h}_u}$ , providing feedback to fine-tune the par-  
905 ameters. This supervised refinement enhances the model’s ability to predict high-affinity structures,  
906 mitigating issues like representation collapse and enabling seamless transition to the inversion phase  
907 for generation.

908 The trainable components across both phases include the encoder weights and gradient optimization  
909 parameters, collectively parameterized by  $f_{\theta}$ . This two-phase approach—unsupervised pre-training  
910 followed by supervised fine-tuning—mitigates issues like model degradation (from suboptimal ini-  
911 tialization) and gradient misalignment, with optimization dynamics driven by gradient-based super-  
912 vision for improved convergence in docking-oriented ligand design.

913  
914 A.2 INVERSION PHASE  
915

916 Building on the trained model, this phase reconstructs the desired structures by reversing the encod-  
917 ing process through gradient-driven optimization. Generally, the inversion process updates a latent  
918 or structural representation  $\mathbf{x}$  based on an objective function  $L$ , formulated as:

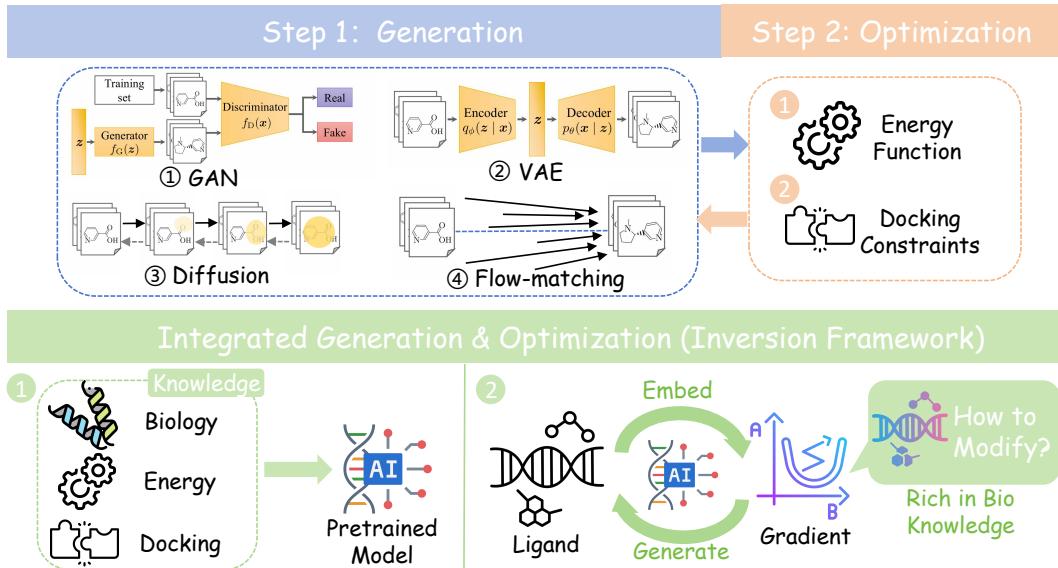


Figure 5: Comparison of the inversion framework with other generative ligand design methods.

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \nabla_{\mathbf{x}_t} L(\mathbf{x}_t), \quad (16)$$

where  $\eta$  is the learning rate, and  $L$  incorporates domain-specific constraints. In our implementation, the Inversion generation module refines point cloud positions  $\mathbf{p}_u$  and features  $\mathbf{f}_u$  using gradient adjustments guided by docking feedback, expressed as:

$$\mathbf{x}_{t+1} = \text{InvGen}(\mathbf{p}_u + \alpha \nabla_{\mathbf{p}_u} E_{\text{dock}}, \mathbf{f}_u + \beta \nabla_{\mathbf{f}_u} E_{\text{dock}}, C_{\text{bio}}), \quad (17)$$

where  $\alpha$  and  $\beta$  are learning rates, and  $C_{\text{bio}}$  enforces biochemical constraints. This enables reverse inference of protein sequences or small molecule structures from zero-start, simultaneously optimizing both molecular types.

The trainable components of this framework include the encoder weights and gradient optimization parameters, collectively parameterized by  $f_\theta$ . This two-phase approach mitigates issues like model degradation (from suboptimal initialization) and representation collapse (from gradient misalignment), with optimization dynamics driven by gradient-based supervision for improved convergence.

## B RELATED WORKS IN DETAILS

Following the taxonomy in the main text, we reorganize existing approaches into four paradigms according to how generation is coupled with optimization. For completeness, we further discuss them from two complementary perspectives: protein design and molecular ligand design.

**① Decoupled Paradigms.** (a) *Protein design.* Classical protein–protein docking pipelines, such as ZDOCK Chen et al. (2010), HADDOCK Dominguez et al. (2003), and RosettaDock Lyskov & Gray (2008), employ geometric complementarity and energy minimization, sometimes with partial priors like known paratopes Kozakov et al. (2017); Yan et al. (2020); Ganea et al. (2021). Deep learning extensions such as AlphaFold-Multimer Evans et al. (2022) improve accuracy but remain resource-intensive. Early generative efforts, e.g., RNN-based paratope generation Liu et al. (2020); Saka et al. (2021) or sequence–structure joint models Jin et al. (2021), also separated sampling from docking-based optimization.

(b) *Molecular ligand design.* Structure-based drug discovery relies on docking engines including AutoDock (Morris et al. (2008)), DOCK (Ewing et al. (2001)), Glide (Halgren et al. (2004)), and GOLD (Verdonk et al. (2003)), supported by scoring functions and refinement via MD, FEP, or

972 MM/PBSA simulations (Berdigaliyev & Aljofan (2020); Fabricant & Farnsworth (2001)). These  
 973 workflows represent ligands in predefined descriptors, then optimize poses post hoc. Early generative  
 974 models (VAEs (Gómez-Bombarelli et al. (2018); Skalic et al. (2019)), GANs (Guimaraes et al.  
 975 (2017))) also adopted a decoupled generation–evaluation scheme. While interpretable, such methods  
 976 are loosely tied to binding or multi-objective constraints.

977 **② Implicitly Coupled Paradigms.** (a) *Protein ligand design.* Here, optimization signals are em-  
 978 bedded into stochastic generative processes. Graph-based generative methods (Jin et al. (2021)),  
 979 backbone-conditioned models (Ingraham et al. (2019)), and energy-aware strategies (Tischer  
 980 et al. (2020); Cao et al. (2021)) guide sequence/structure generation with implicit docking feed-  
 981 back. Diffusion-style models (e.g., DiffAb, HSRN) refine CDR loops and interfaces with  $SE(3)$ -  
 982 equivariant priors.

983 (b) *Molecule ligand design.* Diffusion-based models such as DiffDock (Corso et al. (2022)) and  
 984 GeoDiff (Xu et al. (2022)) integrate  $\Delta G$  or docking signals during denoising. Pocket2Mol (Peng  
 985 et al. (2022)) and TankBind (Lu et al. (2022)) condition on receptor pockets. Broader frameworks,  
 986 including DockStream (Guo et al. (2021)), ALIDIFF (Gu et al. (2024)), and DRUGFLOW (Schneu-  
 987 ning et al. (2025)), couple docking or ADMET constraints with molecular diffusion or generative  
 988 flows (Shi et al. (2020); Lee et al. (2023)). VAE (Liu et al. (2018); Fu et al. (2020); Griffiths &  
 989 Hernández-Lobato (2020); Wang et al. (2022)) and GAN-based methods (De Cao & Kipf (2018);  
 990 Abbasi et al. (2022)) also benefit from structural conditioning. Despite progress, efficiency and  
 991 smooth latent space learning (Brown et al. (2019); Huang et al. (2021); Gao et al. (2024b)) remain  
 992 challenges.

993 **③ Surrogate- or Heuristic-coupled Paradigms.** (a) *Protein ligand design.* Reinforcement learning  
 994 has been applied to antibody optimization (Ingraham et al. (2019); Tischer et al. (2020)), while  
 995 memory-augmented models integrate docking oracles into generation. Energy-based approaches  
 996 similarly exploit docking rewards (Cao et al. (2021)), yet suffer from high computational cost.

997 (b) *Molecule ligand design.* Ligand optimization often leverages search in discrete chemical space.  
 998 Reinforcement learning (Ståhl et al. (2019); You et al. (2018); Zhou et al. (2019); Gottipati et al.  
 999 (2020); Gao et al. (2024a); Jain et al. (2023)), evolutionary algorithms (Jensen (2019); Nigam et al.  
 1000 (2019); Chen et al. (2021)), Markov Chain Monte Carlo (Xie et al. (2021); Fu et al. (2021b)), tree  
 1001 search (Ma et al. (2021)), and Bayesian optimization (Korovina et al. (2020); Moss et al. (2020))  
 1002 iteratively refine molecules with docking oracles. While effective, they require many evaluations  
 1003 and struggle with balancing multiple objectives (Blum & Roli (2003); Mazyavkina et al. (2021)).  
 1004 Frameworks like DockStream (Guo et al. (2021)), ALIDIFF (Gu et al. (2024)), and DRUGFLOW  
 1005 (Schneuening et al. (2025)) attempt to reduce cost via chemical and geometric priors.

#### 1006 **④ Explicitly Gradient-coupled Paradigm (Ours).**

1007 Inversion-based frameworks directly couple generation and optimization via gradient guidance in  
 1008 latent space. Unlike stochastic or heuristic strategies, gradients provide more efficient and inter-  
 1009 pretable refinements of sequence–structure pairs (Niu et al. (2025); Fu et al. (2021a)). The proposed  
 1010 **MagicDock** exemplifies this paradigm by unifying 3D geometry, docking constraints, and biophys-  
 1011 ical objectives. By treating ligand generation as a gradient-driven process, it co-optimizes  $\Delta G$ ,  
 1012 specificity, and stability without handcrafted schedules. This enables biologically relevant designs  
 1013 under a single, generalizable architecture.

## 1014 1015 C METHODOLOGY IN DETAILS

### 1016 1017 C.1 STAGE 1: DOCKING-ORIENTED LIGAND MODELING

1018 We represent ligands (both proteins and molecules) through solvent-accessible surfaces parameter-  
 1019 ized by a probe radius  $r_{\text{probe}} = 1.4 \text{ \AA}$ . Let  $C = \{(c_j, r_j)\}_{j=1}^M$  denote atomic centers and van der  
 1020 Waals radii. The molecular surface is defined as the iso-level set of the distance field:

$$1021 S = \{x \in \mathbb{R}^3 \mid d(x; C) = r_{\text{probe}}\}, \quad d(x; C) = \min_j \|x - c_j\| - r_j, \quad (18)$$

1026 where the minus sign ensures consistency with van der Waals boundaries. A differentiable representation  
 1027 is achieved using a smoothed distance function:  
 1028

$$1029 \text{SDF}(x) = -\bar{\sigma}(x) \log \sum_{j=1}^M \exp\left(-\frac{\|x - c_j\|}{\sigma_j}\right), \quad \bar{\sigma}(x) = \frac{\sum_j \exp(-\|x - c_j\|) \sigma_j}{\sum_j \exp(-\|x - c_j\|)}. \quad (19)$$

1032 Candidate points, sampled from Gaussian perturbations around atoms, are iteratively projected onto  
 1033 the iso-surface via gradient descent. A triangulated mesh is constructed using the ball-pivoting  
 1034 algorithm (BPA) with radii in  $[r_{\text{probe}}, 2r_{\text{probe}}]$ , and  $N = 5000$  points are uniformly sampled for  
 1035 proteins (scaled for ligands). Surface normals are estimated by weighted PCA within a  $2r_{\text{probe}}$   
 1036 neighborhood:

$$1037 \Sigma_i = \sum_{j \in \mathcal{N}_n(i)} w_{ij} (x_j - \bar{x}_i)(x_j - \bar{x}_i)^\top, \quad w_{ij} = \exp\left(-\frac{\|x_i - x_j\|^2}{(2r_{\text{probe}})^2}\right), \quad (20)$$

1039 where  $\mathcal{N}_n(i)$  is the set of neighboring points for normal estimation.

1040 Each surface point  $x_i$  is enriched with a feature vector integrating chemical descriptors, atomic type  
 1041 indicators, geometric descriptors, coordinates, and a molecule-type identifier  $I_i$ :

$$1043 f(x_i) = \text{concat}(f_{\text{chem}}(x_i), f_{\text{atom}}(x_i), f_{\text{geom}}(x_i), I_i, x_i). \quad (21)$$

1045 **Atomic features.** Atomic features  $f_{\text{atom}}(x_i)$  encode the local element distribution through  
 1046 weighted one-hot statistics. For proteins, we adopt  $\{C, H, O, N, S, Se\}$  (6D). For small molecules,  
 1047 we adopt  $\{C(sp3), C(sp2), H, O(sp3), O(sp2), N(sp3), N(sp2), S(sp2)\}$  (8D). The statistics are  
 1048 aggregated using probe-scaled weights  $\omega_{ia} = 1/(\|x_i - c_a\|/r_{\text{probe}} + \varepsilon)$ . Note that this atomic vocabulary  
 1049 is limited to common bio-organic atoms, while halogens and metal ions frequently appearing in  
 1050 pharmaceutically relevant compounds are not yet included; we discuss extending the feature space  
 1051 in Section F.

1052 **Chemical features.** Chemical descriptors  $f_{\text{chem}}(x_i)$  are continuous values that capture local  
 1053 physicochemical properties from neighboring atoms within  $r_{\text{chem}} = 5.0 \text{ \AA}$ . Specifically, we define:  
 1054 hydrogen-bonding potential,

$$1056 H_{\text{bond}}(x_i) = \frac{\sum_{a \in \mathcal{N}_c(i)} \omega_{ia} \mathbb{1}[t_a \in \{O, N, S\}]}{\sum_{a \in \mathcal{N}_c(i)} \omega_{ia}}, \quad (22)$$

1058 charge polarity,

$$1060 C(x_i) = \frac{\sum_{a \in \mathcal{N}_c(i)} \omega_{ia} (\mathbb{1}[t_a \in \{O, N\}] - \mathbb{1}[t_a \in \{C, S\}])}{\sum_{a \in \mathcal{N}_c(i)} \omega_{ia}}, \quad (23)$$

1062 and hydrophobicity/aromaticity,

$$1064 H_{\text{phob}}(x_i) = \frac{\sum_{a \in \mathcal{N}_c(i)} \omega_{ia} \mathbb{1}[t_a = C]}{\sum_{a \in \mathcal{N}_c(i)} \omega_{ia}}, \quad A_{\text{aro}}(x_i) = \text{clip}\left(\frac{H_{\text{phob}}(x_i)}{4.5}, 0, 1\right). \quad (24)$$

1066 Here,  $\mathcal{N}_c(i)$  denotes the chemical neighborhood,  $\mathbb{1}[\cdot]$  the indicator function, and  $\text{clip}(x, a, b)$  truncates  
 1067  $x$  into  $[a, b]$ .

1069 **Geometry features.** Atomic statistics  $T_{i,k}$  reflect weighted element distributions, with 6D for  
 1070 proteins and 8D for small molecules. Geometric descriptors from a  $k_g = 10$  nearest-neighbor set  
 1071  $\mathcal{N}_g(i)$  include mean curvature  $\kappa_{i,1}$ , Gaussian curvature  $\kappa_{i,2}$ , and local density  $D_i$ , computed as:

$$1073 \kappa_{i,1} = \frac{1}{|\mathcal{N}_g(i)|} \sum_{j \in \mathcal{N}_g(i)} \|n_i - n_j\|, \quad (25)$$

$$1075 \kappa_{i,2} = \prod_{m=1}^3 \epsilon_m, \quad \{\epsilon_1, \epsilon_2, \epsilon_3\} = \text{eig}\left(\frac{1}{|\mathcal{N}_g(i)|} \sum_{j \in \mathcal{N}_g(i)} (x_j - \bar{x}_i)(x_j - \bar{x}_i)^\top\right), \quad (26)$$

$$1078 D_i = \frac{1}{k_g} \sum_{j \in \mathcal{N}_g(i)} \|x_i - x_j\|, \quad \text{normalized to } [0, 1], \quad (27)$$

1080 where  $\text{eig}(\cdot)$  returns eigenvalues of the covariance matrix.  
 1081

1082 To enhance scalability, point clouds are partitioned into overlapping patches. Centers are obtained  
 1083 via farthest point sampling (FPS), each grouping  $K = 50$  neighbors within  $r_{\text{patch}} = 5.0 \text{ \AA}$ :

$$1084 \quad \mathcal{P}(x_c) = \{x_j \mid x_j \in \text{KNN}(x_c, K), \|x_j - x_c\| \leq r_{\text{patch}}\}, \quad (28)$$

1085 where KNN denotes the  $K$ -nearest neighbors of a center point. Patch features are summarized by  
 1086 mean and variance pooling:  
 1087

$$1088 \quad \mu_{k,f} = \frac{1}{|\mathcal{P}_k|} \sum_{x_i \in \mathcal{P}_k} f(x_i), \quad \sigma_{k,f}^2 = \frac{1}{|\mathcal{P}_k|} \sum_{x_i \in \mathcal{P}_k} (f(x_i) - \mu_{k,f})^2, \quad (29)$$

1090 with interface labels assigned for points within  $4.0 \text{ \AA}$  of ligand or  $2.0 \text{ \AA}$  of protein surfaces.  
 1091

1092 This yields a probe-aware, feature-enriched, patch-structured representation, unifying proteins and  
 1093 ligands for pre-training and generative modeling.

## 1094 C.2 STAGE 2: UNSUPERVISED PRE-TRAINING

1096 In the second stage, the encoder is pre-trained in an unsupervised manner to capture transferable  
 1097 docking-aware priors from proteins and small molecules. The framework integrates SE(3)-  
 1098 equivariant convolutions, patch-level masking with vector quantization, and reconstruction objec-  
 1099 tives targeting geometric and physicochemical properties, establishing a robust foundation for subse-  
 1100 quent fine-tuning. Below, we detail the computational framework and objectives, ensuring seamless  
 1101 integration of encoding and reconstruction processes.

1102 **SE(3)-Equivariant Encoding.** Each molecular surface is represented as a collection of patches  
 1103

$$1104 \quad X_p = \{(\mathbf{x}_i, f_i)\}_{i=1}^K,$$

1105 where  $\mathbf{x}_i \in \mathbb{R}^3$  are Cartesian coordinates,  $f_i \in \mathbb{R}^d$  are feature vectors of dimension  $d$ , and  $K$  is the  
 1106 total number of patches. The encoder aggregates local neighborhoods  $\mathcal{N}(i)$  using SE(3)-equivariant  
 1107 convolutions:

$$1108 \quad z_i = \sum_{j \in \mathcal{N}(i)} \sum_{l=0}^L R_l(\|\mathbf{x}_{ij}\|) Y_l\left(\frac{\mathbf{x}_{ij}}{\|\mathbf{x}_{ij}\|}\right) \cdot W_l f_j, \quad (30)$$

1111 where  $\mathbf{x}_{ij} = \mathbf{x}_j - \mathbf{x}_i$ ,  $R_l(\cdot)$  are learnable radial functions,  $Y_l(\cdot)$  are spherical harmonics of order  
 1112  $l$ ,  $W_l$  are trainable matrices, and  $L$  is the maximum order of harmonics. This ensures equivariance  
 1113 under rigid motions  $(R, t)$  with  $R \in SO(3)$  and  $t \in \mathbb{R}^3$ :

$$1114 \quad z_i(R\mathbf{x} + t) = \rho(R) z_i(\mathbf{x}), \quad (31)$$

1116 where  $\rho(R)$  is the irreducible representation of  $SO(3)$  acting on the feature space, capturing both  
 1117 scalar and higher-order interactions.

1118 **Patch Masking and Vector Quantization.** A fraction of patches are masked at a specified ratio  $\rho$ ,  
 1119 and their embeddings are replaced by vector-quantized latents sampled from a learnable codebook  $\mathcal{E}$   
 1120 of size  $N_B$ . Following the discrete variational autoencoder (dVAE) relaxation, the quantized latent  
 1121 is computed as:

$$1122 \quad z_{p,i,m} = \frac{\sum_{j=1}^{N_B} \exp\left(\frac{g_j + \log q(e_j | h_{p,i,m})}{\tau}\right) e_j}{\sum_{j=1}^{N_B} \exp\left(\frac{g_j + \log q(e_j | h_{p,i,m})}{\tau}\right)}, \quad (32)$$

1125 where  $h_{p,i,m}$  is the hidden state,  $e_j \in \mathcal{E}$  are codebook entries,  $g_j \sim \text{Gumbel}(0, 1)$  is sampled noise,  
 1126 and  $\tau$  is the softmax temperature. Visible and quantized embeddings are then passed to the decoder  
 1127 for reconstruction, preserving contextual information across masked regions.

1128 **Reconstruction Targets.** For coordinate recovery, the decoder outputs:

$$1130 \quad \hat{X} = \text{Reshape}(\text{MLP}(H_p^{(L_2)})), \quad \hat{X} \in \mathbb{R}^{\delta \rho M \times K \times 3}, \quad (33)$$

1132 where  $H_p^{(L_2)}$  is the decoder output from the  $L_2$ -th layer,  $M$  is the number of input molecular sam-  
 1133 ples, and  $\delta$  is a constant scaling factor. The multilayer perceptron (MLP) produces reconstructed  
 Cartesian coordinates.

1134 For surface geometry, each masked patch  $X_{p,i,m}$  with center  $x_{c,i,m}$  has covariance:  
 1135

$$1136 \quad \Sigma = \frac{1}{K} \sum_{j=1}^K (x_{p,i,j,m} - x_{c,i,m})(x_{p,i,j,m} - x_{c,i,m})^\top \in \mathbb{R}^{3 \times 3}. \quad (34)$$

1139 Pseudo-curvatures are derived from its eigenvalues  $\epsilon_1, \epsilon_2, \epsilon_3$  as:  
 1140

$$1141 \quad \psi_i = \left( \frac{\epsilon_1}{\epsilon_1 + \epsilon_2 + \epsilon_3}, \frac{\epsilon_2}{\epsilon_1 + \epsilon_2 + \epsilon_3}, \frac{\epsilon_3}{\epsilon_1 + \epsilon_2 + \epsilon_3} \right). \quad (35)$$

1143 An MLP predicts  $\hat{\psi}_i$ , supervised by root mean square error (RMSE), ensuring accurate capture of  
 1144 local surface geometry.

1145 **Training Objective.** The overall loss combines reconstruction, curvature, and regularization terms:  
 1146

$$1147 \quad \mathcal{L} = \nu_1 \mathcal{L}_{\text{rec}} + \nu_2 \mathcal{L}_{\text{cur}} + \nu_3 \mathcal{L}_{\text{KL}}, \quad (36)$$

1148 where  $\nu_1, \nu_2, \nu_3$  are scalar weights. The reconstruction loss uses the Chamfer distance:  
 1149

$$1150 \quad \mathcal{L}_{\text{rec}} = \frac{1}{\delta \rho M K} \sum_{i=1}^K \left( \sum_{a \in \hat{X}_i} \min_{b \in X_{p,i,m}} \|a - b\|_2^2 + \sum_{b \in X_{p,i,m}} \min_{a \in \hat{X}_i} \|a - b\|_2^2 \right). \quad (37)$$

1153 The curvature loss is defined as:  
 1154

$$1155 \quad \mathcal{L}_{\text{cur}} = \frac{1}{\delta \rho M} \sum_{i=1}^{\delta \rho M} \|\psi_i - \hat{\psi}_i\|_2^2, \quad (38)$$

1158 and the KL divergence term  $\mathcal{L}_{\text{KL}}$  regularizes the posterior  $q(Z_{p,m} \mid H_{p,m})$  toward a uniform cate-  
 1159 gorical prior.

1160 This unsupervised pre-training framework leverages SE(3)-equivariant encoding, quantized masked  
 1161 reconstruction, and curvature supervision to learn robust docking-aware priors, providing an effec-  
 1162 tive initialization for downstream fine-tuning.

### 1164 C.3 STAGE 3: SUPERVISED FINE-TUNING

1166 In the third stage, the encoder pre-trained in Stage 2 is adapted to docking-specific tasks through  
 1167 supervised fine-tuning on protein-ligand complexes with ground-truth labels. This stage integrates  
 1168 cross-attention fusion to capture receptor-ligand interactions and employs three complementary su-  
 1169 pervision signals—pocket segmentation, interaction prediction, and binding affinity regression—to  
 1170 align representations with docking semantics. Below, we detail the computational framework and  
 1171 objectives, ensuring a seamless transition from encoding to multi-task learning.

1172 **Equivariant Attention Fusion.** Given a receptor surface  $R$  and a ligand  $L$ , their point-cloud repre-  
 1173 sentations are encoded independently by the Stage 2 encoders, yielding latent fields

$$1174 \quad Z_r \in \mathbb{R}^{N_r \times d}, \quad Z_l \in \mathbb{R}^{N_l \times d},$$

1176 where  $N_r$  and  $N_l$  are the numbers of sampled points on the receptor and ligand surfaces, respec-  
 1177 tively, and  $d$  is the embedding dimension. To model interfacial dependencies, we employ an SE(3)-  
 1178 equivariant attention mechanism, in which scalar attention weights are computed from  $\ell = 0$  chan-  
 1179 nels while higher-order features are rotated via Wigner- $D$  matrices to preserve equivariance:

$$1180 \quad \text{Attn}(Z_r, Z_l) = \text{softmax} \left( \frac{Q_r^{(\ell=0)} (K_l^{(\ell=0)})^\top}{\sqrt{d}} \right) V_l,$$

1183 where  $Q_r^{(\ell=0)} = Z_r^{(\ell=0)} W_Q$ ,  $K_l^{(\ell=0)} = Z_l^{(\ell=0)} W_K$ , and  $V_l = Z_l W_V$  with  $W_Q, W_K, W_V \in \mathbb{R}^{d \times d}$   
 1184 being learnable projection matrices. The resulting cross-attended features are aggregated with the  
 1185 original receptor embeddings using a permutation-invariant operator  $A(\cdot)$ , implemented as the con-  
 1186 catenation of mean- and max-pooling:  
 1187

$$\tilde{z} = A(Z_r, \text{Attn}(Z_r, Z_l)) \in \mathbb{R}^{2d}.$$

1188 This fused representation captures both local and interfacial information, serving as the foundation  
 1189 for downstream supervision tasks.  
 1190

1191 **Multi-Task Supervision.** Three complementary supervised objectives align  $\tilde{z}$  with docking semantics  
 1192 in a cascaded manner: the pocket prediction serves as a foundational gate, with its outputs  
 1193 weighting and conditioning the subsequent interaction and affinity predictions to emphasize the util-  
 1194 ity of the identified pocket. 1. *Pocket prediction:* Each receptor embedding  $z_i \in Z_r$  (where  $Z_r$   
 1195 denotes the set of  $N_r$  receptor embeddings) is classified as binding-site or non-binding-site, with  
 1196 ground-truth label  $y_i \in \{0, 1\}$ . The loss is defined as  
 1197

$$\mathcal{L}_{\text{pocket}} = \frac{1}{N_r} \sum_{i=1}^{N_r} \text{BCE}(\hat{y}_i, y_i) + \lambda_p \mathcal{R}_{\text{geom}}, \quad (39)$$

1200 where BCE denotes binary cross-entropy,  $\hat{y}_i$  is the predicted probability,  $N_r$  is the number of re-  
 1201 ceptor embeddings, and  $\lambda_p$  is a regularization weight. The geometric regularizer  $\mathcal{R}_{\text{geom}}$  enforces  
 1202 agreement with distance-based pseudo-labels:  
 1203

$$\mathcal{R}_{\text{geom}} = \frac{1}{N_r} \sum_{i=1}^{N_r} \|\hat{y}_i - y_i^{\text{geom}}\|^2,$$

1204 where  $y_i^{\text{geom}} = 1$  if the closest ligand point lies within a pre-defined cutoff radius, and 0 otherwise.  
 1205 The predicted pocket probabilities  $\hat{y}_i$  are used to gate subsequent losses.  
 1206

1207 2. *Interaction prediction:* Conditioned on the predicted pocket, the fused global vector  $\tilde{z}$  (Eq. 10),  
 1208 processed through a multilayer perceptron (MLP) classifier, predicts whether the receptor–ligand  
 1209 pair forms a valid complex. The pocket-weighted loss is  
 1210

$$\mathcal{L}_{\text{int}} = \frac{1}{N_r} \sum_{i=1}^{N_r} \hat{y}_i \cdot \text{BCE}(\hat{y}_{\text{int}}^{(i)}, y_{\text{int}}^{(i)}), \quad (40)$$

1211 where  $\hat{y}_{\text{int}}^{(i)} \in (0, 1)$  is the predicted interaction probability for the  $i$ -th receptor embedding (focusing  
 1212 on pocket regions),  $y_{\text{int}}^{(i)} \in \{0, 1\}$  is the corresponding ground-truth label,  $N_r$  is the number of  
 1213 receptor embeddings, and the weighting by  $\hat{y}_i$  ensures emphasis on high-confidence pockets.  
 1214

1215 3. *Binding affinity regression:* For complexes with experimentally measured affinities, the model  
 1216 predicts binding free energy  $\Delta G(r, l) \in \mathbb{R}$  (standardized to zero mean and unit variance), condi-  
 1217 tioned on both predicted pocket and interaction. The cascaded-weighted regression loss is  
 1218

$$\mathcal{L}_{\Delta G} = \frac{1}{|\mathcal{V}|} \sum_{(r, l) \in \mathcal{V}} \left( \max(\hat{y}_i \cdot \hat{y}_{\text{int}}^{(i)}, \tau) \cdot (\hat{y}_{\Delta G}(r, l) - \Delta G(r, l))^2 \right), \quad (41)$$

1219 where  $\mathcal{V}$  is the set of receptor–ligand pairs with ground-truth affinity values (with  $|\mathcal{V}|$  denoting its  
 1220 cardinality),  $\hat{y}_{\Delta G}(r, l)$  is the predicted binding free energy,  $\hat{y}_i$  is the predicted pocket probability  
 1221 for the relevant receptor embedding  $i$ ,  $\hat{y}_{\text{int}}^{(i)}$  is the predicted interaction probability for that embed-  
 1222 ding, and  $\tau$  is a minimum confidence threshold (e.g., 0.1). The weighting by  $\hat{y}_i \cdot \hat{y}_{\text{int}}^{(i)}$  propagates  
 1223 the dependency from prior predictions, ensuring affinity regression focuses on viable pocket-based  
 1224 interactions.  
 1225

1226 **Overall Objective.** The joint fine-tuning objective combines the three tasks, each supervised by  
 1227 lightweight MLP heads:  
 1228

$$\mathcal{L}_{\text{FT}} = \alpha \mathcal{L}_{\text{pocket}} + \beta \mathcal{L}_{\text{int}} + \mathcal{L}_{\Delta G}, \quad (42)$$

1229 where  $\alpha$  and  $\beta$  are weighting coefficients. During optimization, the encoders (initialized from  
 1230 Stage 2), the cross-attention module, and task-specific MLP heads are jointly updated. This multi-  
 1231 task framework effectively couples receptor–ligand embeddings and specializes the pre-trained  
 1232 backbone for docking-aware representation learning, ensuring robust alignment with docking ob-  
 1233 jectives.  
 1234

#### 1235 C.4 STAGE 4: INVERSION-BASED LIGAND GENERATION

1236 In the fourth stage, the docking-aware backbone is transformed into a generative engine via an  
 1237 inversion mechanism. Unlike sampling–ranking pipelines, inversion performs direct gradient-based  
 1238

refinement in the latent continuous space, followed by projection and decoding into chemically valid discrete structures. This section provides the complete mathematical formulation, algorithms, and hyperparameters for ligand generation.

**Composite Objective.** Given a receptor  $R$  and a candidate ligand  $S$ , their surfaces are independently encoded into latent fields  $Z_r$  and  $Z_l$  by the Stage 3 encoders. The fused interfacial embedding  $\tilde{z}$  is obtained via cross-attention fusion (Eq. 10). The supervised heads trained in Stage 3 yield a differentiable composite objective function  $\mathcal{F}$ :

$$\mathcal{F}(S; R, \Theta) = \alpha \mathcal{L}_{\text{pocket}}(S; R) + \beta \mathcal{L}_{\text{int}}(S; R) + \mathcal{L}_{\Delta G}(S; R), \quad (43)$$

where  $\mathcal{L}_{\text{pocket}}$  is the binary cross-entropy loss for binding pocket localization,  $\mathcal{L}_{\text{int}}$  is the interaction plausibility loss (evaluating atom-residue interfacial compatibility), and  $\mathcal{L}_{\Delta G}$  is the regression loss for binding free energy  $\Delta G$ . Here  $\alpha = 1.0$  and  $\beta = 0.5$  are balancing coefficients, and  $\Theta$  are model parameters.

**Gradient-based Refinement.** Let  $(\mathbf{x}, f)$  denote the differentiable ligand representation, where  $\mathbf{x} \in \mathbb{R}^{N_l \times 3}$  are atom coordinates and  $f \in \mathbb{R}^{N_l \times d_f}$  are atom features (including atomic type logits, partial charges, hybridization states, and hydrogen-bond polarity indicators). At each iteration  $t$ , the update step is

$$(\mathbf{x}, f)^{t+1} = \Pi_{\mathcal{C}_{\text{valid}}} ((\mathbf{x}, f)^t - \eta_t \nabla_{(\mathbf{x}, f)} \mathcal{F}(S_t; R, \Theta)), \quad (44)$$

where  $\eta_t = 10^{-3}$  is the step size, and  $\Pi_{\mathcal{C}_{\text{valid}}}$  projects the updated state back onto the chemically and geometrically valid manifold  $\mathcal{C}_{\text{valid}}$ . Projection enforces valid valence, realistic bond lengths, and avoidance of steric clashes.

**Generative Mapping.** After projection, the refined continuous variables are decoded into chemically valid structures via a type-specific generative mapping  $\mathcal{G}_{\text{type}}$ . Unlike a deterministic argmax,  $\mathcal{G}_{\text{type}}$  leverages gradient information to bias probabilistic sampling, followed by rule-based corrections to enforce validity.

For small molecules, atom typing is performed by softmax sampling:

$$a_i \sim \text{Categorical}(\sigma(f_i - \gamma \nabla_{f_i} \mathcal{F})), \quad a_i \in \mathcal{A}_{\text{atom}}, \quad (45)$$

where  $\mathcal{A}_{\text{atom}} = \{\text{C(sp3), C(sp2), H, O(sp3), O(sp2), N(sp3), N(sp2), S(sp2)}\}$  is the atom-type set,  $\sigma(\cdot)$  is the softmax function, and  $\gamma = 0.1$  balances gradient bias.

Bond inference uses logits  $g_{ij}$  and is jointly sampled over bond types {1, 2, 3, aromatic}:

$$b_{ij} \sim \text{Categorical}(\sigma(g_{ij} - \gamma \nabla_{g_{ij}} \mathcal{F})), \quad (46)$$

where bond types correspond to single, double, triple, and aromatic bonds. Feasibility is checked by distance thresholds (1.0–2.0 Å) and valence constraints  $v_i \leq v_{\max}(a_i)$ .

To incorporate structural motifs, we introduce a prior  $\mathcal{P}_{\text{motif}}$  based on point cloud geometry, identifying centers via clustering of high-curvature, high-density points. Predefined templates guiding atom placement include: Benzene, Pyridine, Furan, Pyrrole, Thiophene, Imidazole, Carboxyl, Amide, etc. The motif probability is:

$$P_{\text{motif}}(x_i) \propto \exp \left( -\frac{\|\mathbf{x}_i - \mathbf{c}_{\text{motif}}\|^2}{\sigma_{\text{motif}}^2} \right) \cdot \mathbb{1}[\text{valid}(a_i, \mathbf{x}_i)], \quad (47)$$

where  $\mathbf{c}_{\text{motif}}$  is the cluster center,  $\sigma_{\text{motif}} = 0.5$  Å, and invalid structures are corrected by RDKit (a cheminformatics toolkit), with  $\Delta G$  rewarding aromatic and functional motifs.

For protein ligands, residue identity  $r_i$  is sampled similarly:

$$r_i \sim \text{Categorical}(\sigma(f_i - \gamma \nabla_{f_i} \mathcal{F})), \quad r_i \in \mathcal{A}_{20}, \quad (48)$$

where  $\mathcal{A}_{20}$  is the canonical amino acid set. Backbone torsions  $(\phi, \psi)$  are continuously updated and projected into  $[-180^\circ, 180^\circ]$ , while side-chain torsions  $\chi_k$  are sampled from the Dunbrack rotamer library. Cartesian reconstruction is carried out using PyRosetta’s internal geometry engine (a protein modeling suite), followed by energy minimization (Rosetta relax) to resolve steric clashes.

Thus,  $\mathcal{G}_{\text{mol}}$  enforces atom- and bond-level chemical validity via stochastic decoding and sanitization, while  $\mathcal{G}_{\text{prot}}$  performs residue-level sampling with biophysical torsional constraints. Both are tightly

coupled with gradient refinement, ensuring that sampled structures remain docking-aware while respecting chemical and geometric feasibility.

**Iterative Dynamics and Convergence.** Repeated application of gradient refinement and generative mapping yields a sequence of ligands  $\{S_t\}$  converging to an optimized structure:

$$S^* = \lim_{t \rightarrow T} \mathcal{G}_{\text{type}}\left((\mathbf{x}, f)^t\right), \quad (49)$$

with convergence typically observed within several hundreds steps. A cosine-annealed schedule is applied to  $\eta_t \in [10^{-3}, 10^{-5}]$  to balance exploration and stability.

By tightly coupling gradient information with generative decoding, the framework ensures that optimization in latent space translates into chemically valid and docking-aware ligands. The distinction between  $\mathcal{G}_{\text{mol}}$  (small molecules) and  $\mathcal{G}_{\text{prot}}$  (proteins) allows the same inversion principle to adapt seamlessly to both types of ligands.

## D EXPERIMENT IN DETAILS

### D.1 COMPUTATION RESOURCE

All experiments were conducted on a single on-premise node unless otherwise specified. Table 3 summarizes the machine configuration and software stack to facilitate reproducibility.

Table 3: Compute environment used for all experiments.

Component	Configuration
Server	Dell PowerEdge T640
CPU	2× Intel Xeon Gold 6240 @ 2.60 GHz (18 cores/socket; 36 cores, 72 threads total)
Memory	251 GiB RAM
GPU	4× NVIDIA A100 80GB PCIe (80 GiB each); driver 570.124.06; MIG disabled
GPU topology	GPU0/1 near NUMA node 0; GPU2/3 near NUMA node 1
OS / Kernel	Ubuntu 22.04.5 LTS; Linux 5.15.0-126-generic
CUDA	Runtime 12.8 (from driver); Toolkit 12.4 (nvcc)
Python / Conda	Python 3.11.10
DL stack	PyTorch 2.1.0 (cu118 build), torchvision 0.16.0, torchaudio 2.1.0

### D.2 DATASETS

**SKEMPI v2**( Liu et al. (2024)) is a mutation-centric benchmark of experimentally measured binding-affinity changes in protein–protein complexes, with receptor/ligand chains already labeled. Existing chain labels are retained (with spot corrections if needed). Entries without an affinity value are removed (57). As with other sources, affinities are converted to KD (M), the scope is restricted to PPIs, and duplicates are merged using the same Complex-ID scheme.

**Sabdab**( Liu et al. (2024)) is a large repository of antibody–antigen structures with explicit antigen, heavy-chain, and light-chain annotations. For consistency with a PPI setup, antigen chains are treated as receptor and antibody heavy/light chains as ligand. Records lacking affinity (14,148), non-PPI cases (46), entries missing antigen-chain labels (95), and chain-annotation errors (8) are removed. Remaining affinities are unified to KD (M), and duplicates across datasets are resolved with the Complex-ID definition above.

**PDBBind v2020**( Wang et al. (2004)) is a curated collection of biomolecular complexes; the protein–protein portion lists receptor/ligand names but not explicit chain IDs. Chain IDs are inferred with a semi-automatic procedure (parse chain descriptions, fuzzy-match to the annotated names) followed by expert proofreading and splitting of multi-complex PDBs, yielding chain assignments for 2,788 samples. Data cleaning removes non-PPI entries (6), records with ambiguous chain annotation (62), and entries whose reported affinities cannot be reliably converted to KD (62). All remaining affinities are standardized to KD (M), and cross-source duplicates are consolidated via a “Complex ID” composed of PDB code, sorted chains, mutations, and PubMed ID.

1350  
1351  
1352 Table 4: The statistical information of the experimental datasets.  
1353  
1354  
1355  
1356  
1357

Dataset	Samples	Complex Type	Affinity Info
SKEMPI v2	7,146	Protein-Protein	KD, $\Delta G$
Sabdab	1,069	Protein-Protein	KD, $\Delta G$
PDBBind	2,789	Protein-Molecule	KD, $\Delta G$
CrossDocked2020	18,450	Protein-Molecule	pK

1358  
1359  
1360 **CrossDocked2020** ( Francoeur et al. (2020)) is a newly introduced dataset for structure-based ma-  
1361 chine learning, comprising 22.5 million poses of ligands docked into multiple similar binding pock-  
1362 ets across the Protein Data Bank, designed to enhance the training and evaluation of grid-based  
1363 convolutional neural network (CNN) models. This dataset includes cross-docked poses against  
1364 non-cognate receptor structures and model-generated counterexamples, providing a standardized  
1365 resource to recognize ligands in diverse target structures while significantly expanding the number  
1366 of available poses for training.  
1367

### 1368 D.3 BASELINE MODELS

#### 1369 BASELINE MODELS FOR PROTEIN

1370  
1371  
1372 **RAbD**( Adolf-Bryfogle et al. (2018)) introduces a knowledge-based Rosetta framework for compu-  
1373 tational antibody design that grafts canonical CDR loops from PyIgClassify, performs profile-guided  
1374 sequence design with flexible-backbone sampling inside nested Monte-Carlo-plus-minimization cy-  
1375 cles, and optimizes either total energy or interface  $\Delta G$ . However, its reliance on existing structural  
1376 data may limit its effectiveness for designing antibodies against novel or poorly characterized anti-  
1377 gens.

1378 **DiffAb**( Luo et al. (2022)) proposes a diffusion-based, rotation/translation-equivariant generative  
1379 model that co-designs antibody CDR sequences and 3D structures by iteratively denoising amino-  
1380 acid types,  $C\alpha$  coordinates, and  $SO(3)$  side-chain orientations, all explicitly conditioned on the  
1381 target antigen structure. But its performance may be limited by the quality and diversity of the  
1382 training data, potentially restricting its ability to generalize to a wide range of antigens.

1383 **HSRN**( Jin et al. (2022)) introduces a hierarchical, rotation/translation-equivariant framework for  
1384 antibody–antigen docking and design that combines a multi-scale encoder (atom- and residue-level)  
1385 with an iterative, force-based refinement to fold and dock the paratope; during generation, an au-  
1386 toregressive decoder progressively docks partial paratopes and exploits the resulting geometric rep-  
1387 resentation to choose the next residue. However, the model’s heavy reliance on pre-defined epitope  
1388 structures, assuming the input already provides the antigen’s 3D structure and specific epitope lo-  
1389 cation, poses a significant limitation, as epitope prediction remains a challenging task in practice,  
1390 thereby restricting its end-to-end applicability.

1391 **dyMEAN**( Kong et al. (2023)) [Kong et al., 2023] presents an end-to-end,  $E(3)$ -equivariant full-  
1392 atom antibody design framework that, given an antigen epitope and an incomplete antibody se-  
1393 quence, initializes structure using conserved framework residues, attaches a “shadow paratope” to  
1394 exchange invariant information and enable docking, and iteratively co-updates residue types and 3D  
1395 coordinates with an adaptive multi-channel encoder that handles variable atom counts per residue;  
1396 docking is finalized by aligning the native and shadow paratopes, and the method reports superior  
1397 results on CDR-H3 generation, complex structure prediction, and affinity optimization. Although  
1398 dyMEAN excels in end-to-end design and full-atom modeling, its limited scalability and reliance on  
1399 high-quality training data remain limitations.

1400 **AbX**( Zhu et al. (2024)) introduces a continuous-time, score-based diffusion framework that jointly  
1401 models discrete CDR sequences (via a CTMC) and  $SE(3)$  coordinates, conditioned on the anti-  
1402 gen/framework, and guided by evolutionary (ESM-2) priors plus geometric (FAPE, distogram,  
1403 IDD) and physical (violation, van der Waals) constraints to narrow the search space and improve  
binding/quality.

1404 **ABDPO**( Zhou et al. (2024)) formulates antigen-specific sequence–structure co-design as direct  
 1405 energy-based preference optimization, fine-tuning a conditional diffusion model with residue-level  
 1406 energy preferences, decomposed attraction/repulsion terms, and gradient-surgery to resolve con-  
 1407 flicts—thereby steering generations toward low total energy while maintaining binding affinity.  
 1408

1409 **BASELINE MODELS FOR MOLECULE**  
 1410

1411 **DrugFlow**( Schneuing et al. (2025)) introduces a generative framework for structure-based drug de-  
 1412 sign that jointly models continuous ligand coordinates and discrete atom/bond types by combining  
 1413 Euclidean flow matching with discrete Markov bridges. The method extends to **FlexFlow**, which  
 1414 additionally samples protein side-chain torsion angles to capture binding-pocket flexibility. Key  
 1415 innovations include an end-to-end uncertainty estimator for out-of-distribution detection, a virtual  
 1416 node mechanism for adaptive molecule size selection, and a multi-domain preference alignment  
 1417 scheme that efficiently steers generation toward molecules with desirable drug-like properties. Ex-  
 1418 periments on CrossDocked demonstrate state-of-the-art distribution learning across chemical, geo-  
 1419 metric, and physical metrics. While DrugFlow excels in holistic distribution modeling and flexible  
 1420 sampling, challenges remain in scaling to larger protein systems and balancing preference alignment  
 1421 with sample validity.

1422 **3D-SBDD**( Luo et al. (2021)) proposes a 3D generative framework for structure-based drug design  
 1423 that directly models ligand atoms within protein binding pockets using an autoregressive flow-based  
 1424 approach. The method conditions ligand generation on 3D protein environments, incrementally  
 1425 placing atoms while capturing geometric constraints, and employs equivariant neural networks to  
 1426 ensure rotational and translational invariance. Evaluation on CrossDocked shows significant im-  
 1427 provements in binding pose accuracy, chemical validity, and docking performance compared to  
 1428 baseline methods. Despite its strong capability in geometry-aware ligand generation, challenges  
 1429 remain in handling larger, more flexible ligands and incorporating dynamic protein conformations.

1430 **ALIDIFF**( Gu et al. (2024)) introduces a target-aware molecule diffusion framework that aligns  
 1431 generative sampling with exact energy optimization for structure-based drug design. The method  
 1432 integrates a diffusion backbone with a dual-stage alignment scheme: a coarse-grained alignment to  
 1433 enforce global docking plausibility and a fine-grained optimization that explicitly minimizes binding  
 1434 energies within the protein pocket. By coupling stochastic generative modeling with deterministic  
 1435 energy-based refinement, AliDiff achieves superior docking accuracy and binding affinity predic-  
 1436 tion on CrossDocked benchmarks. While it demonstrates strong target-conditioning and energy  
 1437 alignment, its reliance on accurate energy models and the computational overhead of fine-grained  
 1438 optimization pose scalability challenges.

1439 **DiffSBDD**( Schneuing et al. (2024)) presents a diffusion-based generative framework for structure-  
 1440 based drug design that learns to directly sample 3D ligand structures conditioned on protein binding  
 1441 pockets. The model leverages SE(3)-equivariant neural networks to ensure rotational and transla-  
 1442 tional invariance, and formulates ligand generation as a denoising process that progressively refines  
 1443 random atom clouds into chemically valid molecules docked in the target pocket. Extensive ex-  
 1444 periments on CrossDocked demonstrate improved performance over flow-based and autoregressive  
 1445 baselines in terms of pose accuracy, binding affinity, and chemical diversity. Despite its advantages  
 1446 in geometry-aware sampling, challenges remain in scalability to large ligands and efficient integra-  
 1447 tion of protein flexibility.

1448 **DockStream** ( Guo et al. (2021)) presents a flexible molecular docking wrapper that integrates with  
 1449 the de novo design platform REINVENT 2.0, providing access to various ligand embedders (e.g.,  
 1450 Corina, LigPrep) and docking backends (e.g., AutoDock Vina, Glide) to enhance structure-based  
 1451 drug discovery by automating docking experiments, benchmarking configurations, and optimizing  
 1452 docking scores, while overcoming limitations of QSAR models through structural information; its  
 1453 scalability and performance vary by target, with ongoing challenges in accurately predicting binding  
 1454 free energies.

1455 **LiGAN** ( Ragoza et al. (2022)) introduces a deep learning system that generates 3D molecular struc-  
 1456 tures conditioned on receptor binding sites using a conditional variational autoencoder trained on  
 1457 atomic density grids, employing atom fitting and bond inference to construct valid conformations,  
 1458 and demonstrates significant changes in generated molecules with mutated receptors; its reliance on  
 1459 high-quality structural data and computational complexity pose challenges for scalability.

1458 D.4 EVALUATION METRICS  
14591460 PROTEIN EVALUATION METRICS  
1461

1462 **(1) IMP:** The Improvement Percentage (IMP) metric evaluates the relative enhancement in binding  
1463 affinity achieved by designed protein sequences compared to their natural counterparts. Specifically,  
1464 it quantifies the proportion of designed antibodies predicted to exhibit stronger binding than the  
1465 corresponding natural antibodies, as assessed by the Rosetta interface energy function. Formally,  
1466 IMP is defined as:

$$1467 \text{IMP} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(E_i^{\text{design}} < E_i^{\text{natural}}) \times 100\%, \quad (50)$$

1468 where  $N$  is the total number of antibody–antigen pairs,  $E_i^{\text{design}}$  denotes the Rosetta interface energy  
1469 of the  $i$ -th designed antibody,  $E_i^{\text{natural}}$  is the corresponding value for the natural antibody, and  $\mathbb{I}(\cdot)$   
1470 is the indicator function returning 1 if the designed sequence has a lower (better) predicted binding  
1471 energy than the natural sequence.

1472 Higher IMP values indicate that a larger fraction of designed antibodies surpass their natural refer-  
1473 ences in predicted binding affinity, highlighting the effectiveness of the design strategy in optimizing  
1474 protein–protein interactions.

1475 **(2) STA:** The Stability (STA) metric quantifies the conformational integrity and biochemical via-  
1476 bility of designed proteins by integrating a weighted composite of steric hindrance assessment and  
1477 structural coherence evaluation, thereby mitigating the risk of thermodynamically unstable or mis-  
1478 folded conformations. This metric is derived from two sub-components: the Steric Clash Score  
1479 (SCS), which penalizes interatomic van der Waals overlaps indicative of steric repulsion, and the  
1480 Secondary Structure Coherence (SSC), which evaluates the fidelity of predicted helical, sheet, and  
1481 coil motifs against empirical folding propensities. Formally, STA is computed as:

$$1482 \text{STA} = \frac{1}{N} \sum_{i=1}^N \left[ \alpha \cdot \exp \left( -\frac{\text{SCS}(p_i)}{\sigma} \right) + \beta \cdot \left( 1 - \frac{|\text{SSC}(p_i) - \mu|}{\tau} \right)^2 \right], \quad (51)$$

1483 where  $N$  is the number of generated proteins,  $p_i$  denotes the  $i$ -th protein,  $\text{SCS}(p_i)$  and  $\text{SSC}(p_i)$   
1484 represent the respective sub-scores derived from all-atom clash detection and dihedral angle consis-  
1485 tency checks,  $\alpha$  and  $\beta$  are empirical weighting coefficients (with  $\alpha + \beta = 1$ ),  $\sigma$  and  $\tau$  are scaling  
1486 hyperparameters for normalization, and  $\mu$  is the expected coherence baseline calibrated from native  
1487 protein ensembles. In our implementation, the coefficient for secondary structure propensity ( $\beta$ ) is  
1488 set to 0.6, and the coefficient for collision detection ( $\alpha$ ) is set to 0.4.

1489 Elevated STA values signify enhanced adherence to biophysical constraints across the ensemble,  
1490 underscoring the efficacy of the generative paradigm in yielding robust, functional protein architec-  
1491 tures.

1492 **(3) DIV:** The Diversity (DIV) metric quantifies the sequence and structural variability among gen-  
1493 erated antibodies, capturing the spread of designs in protein space. A diverse set of candidates  
1494 increases the likelihood of discovering high-affinity binders with novel interaction profiles. Mathe-  
1495 matically, DIV is defined as the average pairwise dissimilarity:

$$1496 \text{DIV} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^N D(p_i, p_j), \quad (52)$$

1497 where  $N$  is the number of generated proteins, and  $D(p_i, p_j)$  is a distance function that measures  
1498 dissimilarity between proteins  $p_i$  and  $p_j$  in both sequence and structural space. Higher DIV values  
1499 indicate a broader coverage of the protein design landscape, reducing redundancy in the generated  
1500 set.

1501 **(4) NOV:** The Novelty (NOV) metric measures how different the generated proteins are from the  
1502 training dataset, considering both sequence and structural similarity. For a generated protein  $p_i$  and  
1503 a training protein  $t_j$ , the combined similarity score  $S(p_i, t_j)$  is defined as:

$$1504 S(p_i, t_j) = \alpha \cdot S_{\text{seq}}(p_i, t_j) + (1 - \alpha) \cdot S_{\text{str}}(p_i, t_j), \quad (53)$$

1512 where  $S_{\text{seq}}$  denotes sequence similarity,  $S_{\text{str}}$  denotes structural similarity, and  $\alpha \in [0, 1]$  controls  
 1513 their relative contribution. Given sequences of length  $L$ , we compute sequence identity as:  
 1514

$$1515 \quad S_{\text{seq}}(p_i, t_j) = \frac{1}{L} \sum_{k=1}^L \mathbb{I}[p_i[k] = t_j[k]], \quad (54)$$

1517 where  $\mathbb{I}[\cdot]$  is the indicator function. For structural alignment, we use a TM-score-like normalized  
 1518 measure:  
 1519

$$1520 \quad S_{\text{str}}(p_i, t_j) = \frac{1}{L} \sum_{k=1}^L \frac{1}{1 + \left(\frac{d_k}{d_0}\right)^2}, \quad (55)$$

1522 where  $d_k$  is the distance between aligned  $C_\alpha$  atoms, and  $d_0$  is a length-dependent normalization  
 1523 constant. For each generated protein  $p_i$ , we take the maximum similarity across all training proteins  
 1524 and define novelty as:  
 1525

$$1526 \quad \text{NOV} = 1 - \frac{1}{N} \sum_{i=1}^N \max_{t_j \in \mathcal{T}} S(p_i, t_j), \quad (56)$$

1527 where  $N$  is the number of generated proteins and  $\mathcal{T}$  is the training set. Higher NOV values indicate  
 1528 stronger novelty in both sequence and structure.  
 1529

1530 **(5) AAR:** The Amino Acid Recovery (AAR) metric measures the sequence-level accuracy of generated  
 1531 protein sequences by comparing them against their corresponding native (reference) sequences.  
 1532 This metric reflects how well the generative model is able to reproduce the original amino acid composition,  
 1533 thereby serving as an indicator of sequence fidelity and preservation of native biochemical properties.  
 1534 The mathematical formulation is given by:  
 1535

$$1536 \quad \text{AAR} = \frac{1}{N \cdot L} \sum_{i=1}^N \sum_{j=1}^L \mathbb{I}(a_{ij}^{\text{gen}} = a_{ij}^{\text{ref}}) \times 100\%, \quad (57)$$

1538 where  $N$  is the number of protein sequences,  $L$  is the sequence length,  $a_{ij}^{\text{gen}}$  denotes the amino acid at  
 1539 position  $j$  in the  $i$ -th generated sequence,  $a_{ij}^{\text{ref}}$  is the corresponding amino acid in the native sequence,  
 1540 and  $\mathbb{I}(\cdot)$  is the indicator function.  
 1541

Higher AAR values indicate closer agreement with native sequences, suggesting that the generated  
 1542 proteins better retain structural and functional characteristics inherent to the original proteins.  
 1543

#### 1544 MOLECULE EVALUATION METRICS

1546 **(1) Vina Score:** The Vina Score metric provides an estimate of the binding affinity between a  
 1547 generated small molecule (ligand) and a target protein. It reflects the stability of the protein–ligand  
 1548 complex predicted during molecular docking. Following the formulation implemented in **AutoDock**  
 1549 **Vina**, the score approximates the binding free energy based on an empirical scoring function that  
 1550 accounts for key interaction terms, including steric complementarity, hydrogen bonding, hydrophobic  
 1551 interactions, and torsional entropy penalties. The mathematical formulation is given by:  
 1552

$$E_{\text{Vina}} = E_{\text{gauss}} + E_{\text{repulsion}} + E_{\text{hydrophobic}} + E_{\text{hydrogen}} + E_{\text{torsional}}, \quad (58)$$

1553 where  $E_{\text{gauss}}$  models attractive van der Waals interactions,  $E_{\text{repulsion}}$  penalizes steric clashes,  
 1554  $E_{\text{hydrophobic}}$  captures hydrophobic contacts,  $E_{\text{hydrogen}}$  measures hydrogen bond formation, and  
 1555  $E_{\text{torsional}}$  accounts for the conformational entropy cost of ligand flexibility.  
 1556

The resulting Vina Score is reported in kcal/mol, with more negative values indicating stronger  
 1557 predicted binding affinities. Typically, scores range from around  $-4$  kcal/mol (weak binding) to  
 1558 below  $-10$  kcal/mol (highly favorable binding). Although not a direct physical free energy, the  
 1559 score serves as a comparative metric to rank ligands by their likelihood of stable binding.  
 1560

**(2) High-affinity:** The High-affinity metric quantifies the proportion of generated molecules that  
 1561 exhibit stronger predicted binding to a given protein target than their corresponding reference ligands.  
 1562 Binding strength is assessed using Vina Scores, where lower (more negative) values indicate  
 1563 higher predicted affinity. The mathematical formulation is given by:  
 1564

$$1565 \quad \text{High-affinity} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(E_i^{\text{gen}} < E_i^{\text{ref}}), \quad (59)$$

1566 where  $N$  denotes the total number of generated molecules,  $E_i^{\text{gen}}$  is the Vina Score of the  $i$ -th generated molecule,  $E_i^{\text{ref}}$  is the Vina Score of its corresponding reference ligand, and  $\mathbb{I}(\cdot)$  is the indicator function returning 1 if the condition is satisfied and 0 otherwise.

1570 The metric outputs the fraction of molecules surpassing the reference ligands in predicted affinity, 1571 providing a normalized measure of how frequently the generation process yields candidates with 1572 potentially improved binding properties.

1573 **(3) STA:** The Stability (STA) metric for small-molecule generation evaluates the chemical plausibility 1574 and conformational robustness of designed ligands by integrating synthetic accessibility and 1575 conformational strain energy into a unified score. This ensures that generated compounds are not 1576 only geometrically valid but also chemically feasible under standard synthesis and physiological 1577 conditions. Specifically, the metric combines two sub-components: the Synthetic Accessibility Index 1578 (SAI), which penalizes ligands with rare or chemically intractable substructures, and the Conformational 1579 Strain Energy (CSE), which quantifies the internal energetic penalty required to maintain 1580 a given 3D geometry relative to its energy-minimized conformation. Formally, STA is defined as:

$$1581 \text{STA} = \frac{1}{M} \sum_{j=1}^M \left[ \gamma \cdot \exp \left( -\frac{\text{CSE}(m_j)}{\lambda} \right) + \delta \cdot \left( 1 - \frac{\text{SAI}(m_j) - \mu}{\kappa} \right)^2 \right], \quad (60)$$

1584 where  $M$  is the number of generated molecules,  $m_j$  denotes the  $j$ -th molecule,  $\text{CSE}(m_j)$  is the 1585 strain energy computed via molecular mechanics force fields, and  $\text{SAI}(m_j)$  is a normalized synthetic 1586 accessibility score.  $\gamma$  and  $\delta$  are weighting coefficients with  $\gamma + \delta = 1$ .

1588 In our implementation, the scaling parameter is fixed as  $\lambda = 10$ , the normalization parameter as 1589  $\kappa = 1.5$ , and the baseline accessibility as  $\mu = 3.0$ , the coefficient for SAI ( $\gamma$ ) is set to 0.5, and the 1590 coefficient for collision detection ( $\delta$ ) is set to 0.5. These hyperparameters are selected to provide a 1591 stable balance between strain minimization and synthetic feasibility.

1592 **(4) DIV:** The Top-K Diversity (DIV) metric quantifies the structural diversity of the top-K generated 1593 molecules, reflecting the spread of chemical structures within a given set. Following the methodology 1594 of (Bengio et al. (2021)), DIV is defined as the average pairwise Tanimoto distance between the 1595 Morgan fingerprints of the generated molecules. The mathematical formulation is given by:

$$1596 \text{DIV} = \frac{2}{K(K-1)} \sum_{i=1}^{K-1} \sum_{j=i+1}^K \text{Tanimoto}(FP_i, FP_j), \quad (61)$$

1600 where  $K$  is the number of top-ranked molecules (e.g., top-10),  $FP_i$  and  $FP_j$  are the Morgan finger- 1601 print vectors for molecules  $i$  and  $j$ , and the Tanimoto similarity is computed as:

$$1602 \text{Tanimoto}(FP_i, FP_j) = \frac{|FP_i \cap FP_j|}{|FP_i \cup FP_j|}, \quad (62)$$

1605 with  $|FP_i \cap FP_j|$  and  $|FP_i \cup FP_j|$  representing the number of common and total unique bits, re- 1606 spectively. The Tanimoto distance is then  $1 - \text{Tanimoto}(FP_i, FP_j)$ , ensuring a range of 0 (identical 1607 structures) to 1 (completely dissimilar structures). The factor  $\frac{2}{K(K-1)}$  normalizes the average over 1608 all unique pairs.

1610 **(5) NOV:** Following InversionGNN ( Niu et al. (2025)), the Novelty (NOV) metric quantifies the 1611 proportion of generated molecules that are not present in the training set, serving as an indicator of 1612 the model’s ability to explore beyond the learned chemical space. This metric is particularly relevant 1613 in de novo molecular design, where generating novel structures is crucial for discovering new drug 1614 candidates. The mathematical formulation is given by:

$$1615 \text{Nov} = \frac{N_{\text{new}}}{N_{\text{total}}}, \quad (63)$$

1618 where  $N_{\text{new}}$  is the number of generated molecules that do not appear in the training set, and  $N_{\text{total}}$  1619 is the total number of generated molecules evaluated. The value of Nov ranges from 0 to 1, with higher values indicating a greater proportion of novel molecules.

1620  
1621

## D.5 MORE EXPERIMENT FIGURES &amp; TABLES

1622  
1623

## D.5.1 ABLATION STUDY

1624  
1625  
1626  
1627  
1628

For the ablation experiments of Stage 1, Stage 2, and Stage 3, each module plays an indispensable role in the generation of the final ligand. In addition, the ablation experiment for stage 4 showed that compared to the exhaustive method, MagicDock’s gradient based selective generation strategy has an order of magnitude efficiency advantage while maintaining a basically consistent effect, proving the pertinence of Stage 4.

1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1640  
1641  
1642

	MagicDock	w/o stage 1	w/o stage 2	w/o stage 3
SKEMPI v2	36.32%	28.55%	30.14%	22.78%
SAbDab	27.87%	23.76%	23.39%	16.94%
PDBBind	40.63%	36.87%	31.76%	26.52%
CrossDocked	60.02%	48.08%	52.94%	44.63%

(a) Performance comparison between models with and without stage 1,2 and 3.

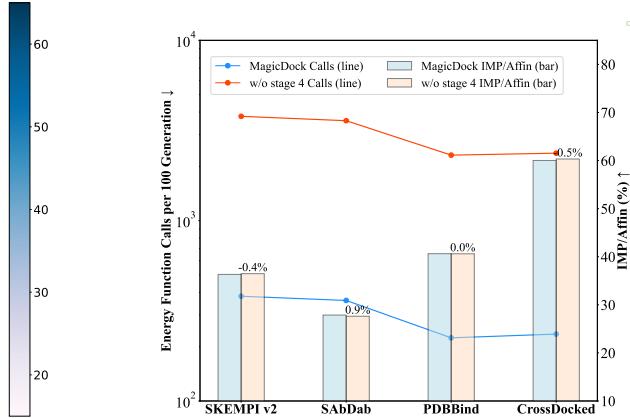


Figure 6: Ablation study on the impact of different stages in MagicDock.

1649

## D.5.2 GEOMETRIC AND FEATURE NOISE STUDY

1650  
1651  
1652  
1653  
1654

We evaluate robustness under structural perturbations on  $N = 100$  receptor–ligand pairs by applying Gaussian coordinate noise and feature dropout to receptor surfaces. Ligand generation are repeated on noisy vs. clean inputs, and metrics are reported in Fig. 7.

1655  
1656  
1657  
1658  
1659  
1660

For protein ligands, robustness is measured by IMP, AAR, and RMSD. MagicDock achieves higher IMP and lower RMSD compared with DiffAb and Abx, demonstrating its stability under noise. For small molecule ligands, robustness is assessed using Vina score, High-affinity, and QED. MagicDock consistently yields more negative Vina scores and higher high-affinity rates, indicating resilience to input perturbations. Together, these results confirm that MagicDock maintains reliable performance under realistic noise, highlighting robustness to biological uncertainty.

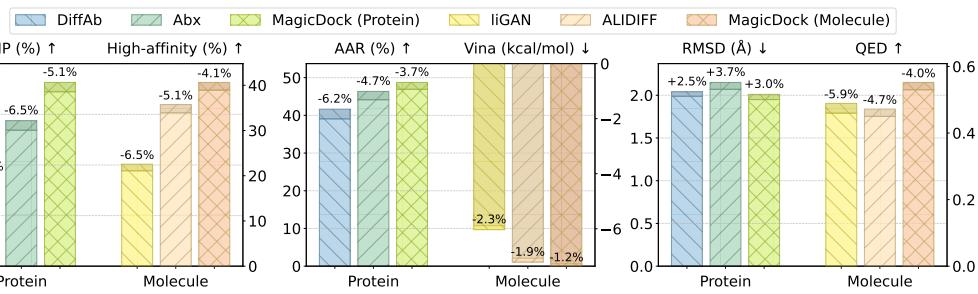
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673

Figure 7: Robustness under geometric and feature noise, evaluated on IMP/High-affinity (%), AAR/Vina (kcal/mol), and RMSD/QED respectively on protein and molecular baselines.

1674  
1675

## D.5.3 LOCAL PERTURBATION CONSISTENCY STUDY

1676  
1677  
1678  
1679  
1680  
1681  
1682  
1683  
1684

Finally, we evaluate whether gradient attributions faithfully predict the energetic consequences of local structural edits. We introduce small geometric and chemical perturbations around high-attribution sites and compare gradient-predicted energy changes with the actual  $\Delta\mathcal{F}$  measured after re-evaluation. This experiment directly tests whether attribution scores can serve as actionable signals for structural refinement, beyond passive localization. Performance is assessed by the sign-consistency rate (SCR), the coefficient of determination ( $R^2$ ), rank correlation (Spearman  $\rho$ ), and the mean observed energy change  $\Delta\mathcal{F}$ , where negative values indicate improved binding. As shown in Fig. 8, our IG-guided strategy substantially outperforms saliency and Grad $\times$ Input baselines, and approaches the oracle behavior of physics-based Rosetta evaluations.

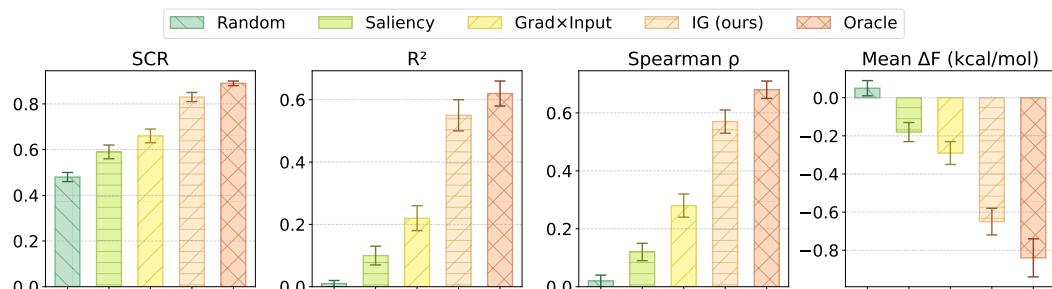
1685  
1686  
1687  
1688  
1689  
1690  
1691  
1692  
1693  
16941695  
1696  
1697

Figure 8: Evaluation of local perturbation consistency, showing SCR,  $R^2$ , Spearman correlation, and Mean  $\Delta\mathcal{F}$  (kcal/mol) for attribution methods.

1698  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709

## D.5.4 CONFORMATIONAL AND MUTATIONAL VARIABILITY STUDY

1710  
1711  
1712  
1713  
1714  
1715  
1716  
1717  
1718  
1719  
1720  
1721  
1722

We further evaluate remarkable robustness to biological variability by subjecting  $N = 100$  receptors to perturbations, including conformer ensembles, single-point mutations, and combined variants. These perturbed structures are systematically processed through the inversion pipeline. Performance degradation is assessed via IMP for protein-based methods and High-Affinity for small-molecule methods, with MagicDock reporting both metrics. As shown in Fig. 9, statistical tests compare conservative and non-conservative mutations, illustrating MagicDock’s exceptional robustness. MagicDock demonstrates minimal degradation and consistently superior performance over baselines.

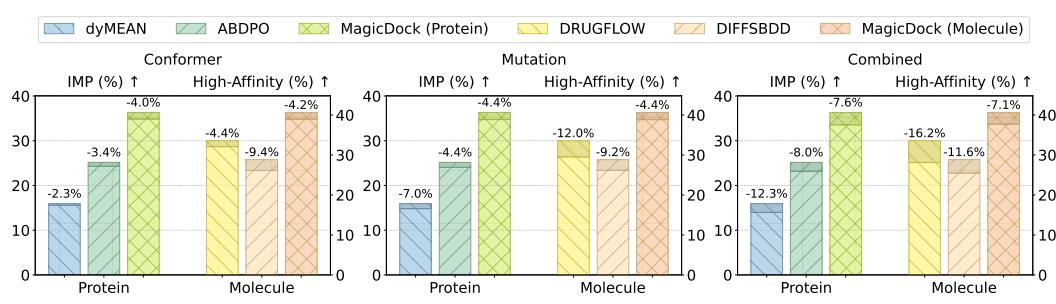
1723  
1724  
1725  
1726  
1727

Figure 9: Robustness to conformational and mutational variability.

## D.5.5 RUNTIME AND RESOURCE USAGE STUDY

1728  
1729  
1730  
1731  
1732

We benchmark MagicDock against baselines using 100 receptors, generating one ligand per receptor under identical hardware. Wall-clock time and memory usage are assessed per computational stage. Fig. 10a demonstrates MagicDock’s reduced runtime while Fig. 10b highlights lower peak memory consumption, underscoring its resource efficiency.

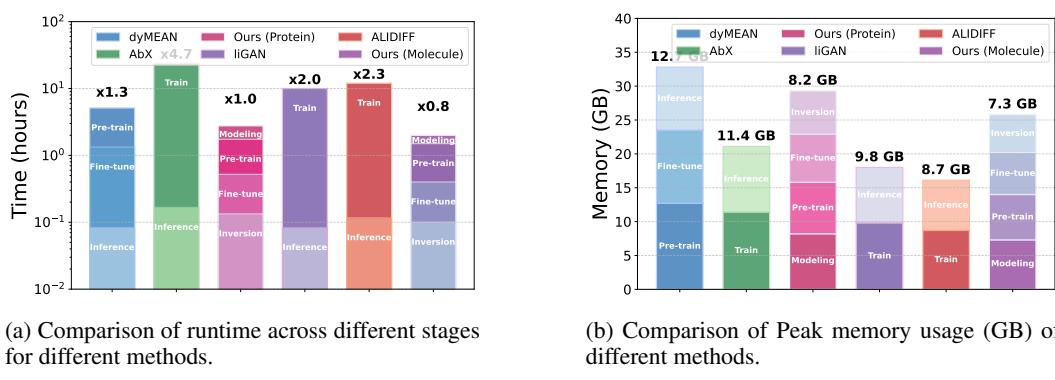


Figure 10: Comparison of runtime and resource usage of different methods.

#### D.5.6 SCALABILITY STUDY

Scalability is evaluated across nine settings varying receptor sizes (500, 2000, 5000 points) and ligand complexities (30, 80, 150 atoms/residues) using 100 receptors from CrossDocked2020 and SAbDab. Metrics include per-iteration latency, iterations-to-converge, memory footprint, and scaling exponent  $\gamma$  from  $T \propto N^\gamma$ .

Table 5: Scalability results across nine complexity settings. Values are averages over three runs, with MagicDock/dyMEAN/ALIDIFF separated by '/'.  
 Receptor | Ligand | Latency (s/iter) | Iterations | Memory (GB)

Receptor	Ligand	Latency (s/iter)	Iterations	Memory (GB)
Small	Low	0.42/1.05/0.58	45/215/380	7.6/11.2/8.1
	Medium	0.65/1.35/0.85	108/285/455	8.4/12.8/9.2
	High	0.78/1.55/1.12	182/358/590	9.3/13.9/10.3
Medium	Low	1.15/2.25/1.18	48/335/440	10.3/16.8/11.6
	Medium	1.36/2.85/1.65	112/372/595	11.7/18.5/13.2
	High	1.72/3.45/2.05	192/445/545	12.8/20.8/14.7
Large	Low	2.28/5.60/2.75	58/405/480	15.9/26.5/17.8
	Medium	2.75/6.95/3.85	125/475/640	18.2/31.5/20.2
	High	3.35/8.90/4.75	245/478/780	20.5/36.2/23.8

Table 6: Fitted runtime exponents  $\gamma$  for each method. This table evaluates the computational efficiency of different ligand design methods by presenting their fitted runtime exponents  $\gamma$ , which quantify how runtime scales with input size in a polynomial manner

Method	$\gamma$
MagicDock	1.4
dyMEAN	1.8
ALIDIFF	1.7

## E THEORETICAL ANALYSIS IN DETAILS

### E.1 PROOF OF SE(3)-EQUIVARIANCE FOR MAGICDOCK

We analyze the SE(3)-equivariance of MagicDock under idealized assumptions. While the practical implementation may include minor numerical deviations (e.g., floating-point tie-breaking, discretization), Stages 1–3 are implemented with strictly equivariant modules, and Stage 4 is approximately equivariant due to non-convex chemical validity constraints. The following proof shows that the framework is SE(3)-equivariant in the limit of exact equivariant modules and convex surrogates.

**Theorem 1** (SE(3)-Equivariance of MagicDock). *Given the receptor point cloud  $P_{rec}$  and the initial ligand point cloud  $P_{lig}^0$ , let the optimized ligand be  $P_{lig}^* = \text{MagicDock}(P_{rec}, P_{lig}^0)$ . Under the assumptions that (i) the chemical validity set  $\mathcal{C}_{\text{valid}}$  admits a convex surrogate and exact Euclidean projection, (ii) all learned modules are SE(3)-equivariant, and (iii) features are restricted to invariant or properly transformed irreducible representations, MagicDock is SE(3)-equivariant. Namely, for any  $g \in \text{SE}(3)$ ,*

$$g \cdot P_{lig}^* = \text{MagicDock}(g \cdot P_{rec}, g \cdot P_{lig}^0),$$

1782 where  $g \cdot P := RP + t$  for  $R \in \text{SO}(3)$ ,  $t \in \mathbb{R}^3$  acts on coordinates, while scalar features remain  
 1783 invariant.

1784 **Lemma 1.1** (Stage 1: Surface Point Cloud Modeling). *The surface point cloud mapping  $\mathcal{S}$  is  $\text{SE}(3)$ -  
 1785 equivariant:  $\mathcal{S}(g \cdot X) = g \cdot \mathcal{S}(X)$ .*

1787 *Proof.* Since the smoothed distance function (SDF) and associated weights depend only on Eu-  
 1788 clidean distances, which are invariant under rotations and translations, we have  $\text{SDF}(Rx + t) =$   
 1789  $\text{SDF}(x)$ . Mesh generation steps (gradient descent projection, FPS, KNN) depend only on pairwise  
 1790 distances and thus commute with  $g$ . **In practice:** deterministic FPS/KNN and numerically stable  
 1791 SDF are implemented, so equivariance holds up to floating-point error.  $\square$   
 1792

1793 **Lemma 1.2** (Stage 2: Pre-training with Equivariant Encoder). *The  $\text{SE}(3)$ -equivariant backbone  $f_\theta$   
 1794 built from tensor field convolutions is  $\text{SE}(3)$ -equivariant.*

1795 *Proof.* The convolution kernels are of the form

$$1798 z_i = \sum_{j \in \mathcal{N}(i)} \sum_{l=0}^L R_l(\|\mathbf{x}_{ij}\|) Y_l\left(\frac{\mathbf{x}_{ij}}{\|\mathbf{x}_{ij}\|}\right) \cdot W_l f_j,$$

1801 which are equivariant because (i) radial functions depend only on invariant distances, and (ii) spheri-  
 1802 cal harmonics transform according to irreducible representations of  $\text{SO}(3)$ . **In practice:** the encoder  
 1803 is implemented with the `e3nn` library, which guarantees strict  $\text{SE}(3)$ -equivariance.  $\square$

1804 **Lemma 1.3** (Stage 3: Supervised Fine-tuning with Equivariant Attention). *When scalar attention  
 1805 weights are computed from  $\ell = 0$  channels and higher-order features are rotated via Wigner-D  
 1806 matrices, the cross-attention layers preserve  $\text{SE}(3)$ -equivariance.*

1808 *Proof.* Attention scores  $A$  computed from  $\ell = 0$  channels are invariant. Values  $V^{(\ell)}$  transform by  
 1809  $\rho^{(\ell)}(R)$  and aggregation

$$1811 \tilde{Z}_r^{(\ell)} = \sum_j A_{ij} \rho^{(\ell)}(R) V_{l,j}^{(\ell)}$$

1813 yields  $\tilde{Z}_r^{(\ell)}(g \cdot P) = \rho^{(\ell)}(R) \tilde{Z}_r^{(\ell)}(P)$ . **In practice:** we explicitly use irreducible representations  
 1814 and Wigner-D matrices in attention, so equivariance is preserved exactly.  $\square$

1816 **Lemma 1.4** (Stage 4: Inversion-based Generation). *The iterative update*

$$1817 P_{\text{lig}}^{t+1} = \Pi_{\mathcal{C}_{\text{valid}}} (P_{\text{lig}}^t - \eta_t \nabla_{P_{\text{lig}}} \mathcal{F}(P_{\text{lig}}^t; P_{\text{rec}}, \Theta)),$$

1819 is  $\text{SE}(3)$ -equivariant under convex surrogate constraints.

1821 *Proof.* If  $\mathcal{F}$  is built from invariant quantities (distances, angles), then  $\mathcal{F}(g \cdot P_{\text{lig}}; g \cdot P_{\text{rec}}) =$   
 1822  $\mathcal{F}(P_{\text{lig}}; P_{\text{rec}})$ . By the chain rule, for  $y = Rx + t$  we have

$$1824 \nabla_y \mathcal{F}(y) = R \nabla_x \mathcal{F}(x),$$

1825 since  $\delta \mathcal{F} = \nabla_x \mathcal{F} \cdot \delta x = \nabla_y \mathcal{F} \cdot \delta y$  with  $\delta y = R \delta x$ . Thus gradient steps commute with  $g$ :

$$1827 g \cdot (P - \eta \nabla_P \mathcal{F}) = (RP + t) - \eta(R \nabla_P \mathcal{F}) = g \cdot P - \eta \nabla_{g \cdot P} \mathcal{F}.$$

1829 If  $\mathcal{C}_{\text{valid}}$  is convex and invariant, projection also preserves equivariance. **In practice:** the true chem-  
 1830 ical validity set is highly non-convex (bond formation, aromaticity, topology constraints), so we  
 1831 enforce validity via heuristic repair. This makes Stage 4 only approximately equivariant in realistic  
 1832 settings.  $\square$

1833 *Proof of Theorem.* Each stage (surface modeling, equivariant encoder, equivariant attention, inver-  
 1834 sion) preserves  $\text{SE}(3)$ -equivariance under the stated assumptions. By closure of equivariant maps  
 1835 under composition, MagicDock is  $\text{SE}(3)$ -equivariant in the idealized setting.  $\square$

**Remark on Practical Implementations.** In our implementation: Stage 1–Stage 3 are strictly SE(3)-equivariant (up to numerical precision) as guaranteed by design choices (deterministic SDF/FPS/KNN, e3nn-based encoder, Wigner- $D$  based attention). Stage 4 involves chemical validity projection on a non-convex manifold and is therefore only approximately equivariant. For reproducibility, we use deterministic tie-breaking in FPS/KNN, fixed random seeds, and numerically stable softmax. Thus MagicDock should be regarded as *exactly SE(3)-equivariant in Stages 1–3* and *approximately equivariant in Stage 4*.

## E.2 PROOF OF CONVERGENCE FOR MAGICDOCK

We analyze the convergence of the gradient-driven inversion stage in MagicDock under idealized assumptions. The analysis follows projected gradient descent (PGD) theory in non-convex optimization, while clarifying the role of convex approximations of the chemical validity manifold.

**Theorem 2** (PGD Convergence under Convex Approximation). *Let  $\mathcal{F}(P_{\text{lig}}; P_{\text{rec}}, \Theta)$  be differentiable,  $L$ -smooth, and bounded below by  $\mathcal{F}^* > -\infty$ . Assume the validity constraint set  $\mathcal{C}_{\text{valid}}$  is nonempty, closed, bounded, and convex (serving as an idealized surrogate for chemical validity). For a fixed step size  $0 < \eta \leq 1/L$ , consider the iteration*

$$P_{\text{lig}}^{t+1} = \Pi_{\mathcal{C}_{\text{valid}}} (P_{\text{lig}}^t - \eta \nabla \mathcal{F}(P_{\text{lig}}^t)),$$

where  $\Pi_{\mathcal{C}_{\text{valid}}}$  denotes Euclidean projection. Then:

1. The objective decreases monotonically:

$$\mathcal{F}(P_{\text{lig}}^{t+1}) \leq \mathcal{F}(P_{\text{lig}}^t) - \frac{\eta}{2} \|g_\eta(P_{\text{lig}}^t)\|^2,$$

where  $g_\eta(x) = \frac{1}{\eta}(x - \Pi_{\mathcal{C}_{\text{valid}}}(x - \eta \nabla \mathcal{F}(x)))$  is the gradient mapping.

2. The sum of squared gradient mappings is bounded:

$$\sum_{t=0}^{T-1} \|g_\eta(P_{\text{lig}}^t)\|^2 \leq \frac{2}{\eta} (\mathcal{F}(P_{\text{lig}}^0) - \mathcal{F}^*).$$

3. Consequently,  $\liminf_{t \rightarrow \infty} \|g_\eta(P_{\text{lig}}^t)\| = 0$ , and every accumulation point  $P^*$  satisfies  $g_\eta(P^*) = 0$ , i.e., it is a first-order stationary point of  $\mathcal{F}$  over  $\mathcal{C}_{\text{valid}}$ .

**Lemma 2.1** (Descent Lemma for PGD). *For  $L$ -smooth  $\mathcal{F}$  and  $\eta \leq 1/L$ , the PGD update  $x^+ = \Pi_{\mathcal{C}}(x - \eta \nabla \mathcal{F}(x))$  satisfies*

$$\mathcal{F}(x^+) \leq \mathcal{F}(x) - \frac{\eta}{2} \|g_\eta(x)\|^2,$$

where  $g_\eta(x) = \frac{1}{\eta}(x - x^+)$ .

*Proof.* Let  $y = x - \eta \nabla \mathcal{F}(x)$  and  $x^+ = \Pi_{\mathcal{C}}(y)$ . By  $L$ -smoothness:

$$\mathcal{F}(x^+) \leq \mathcal{F}(x) + \langle \nabla \mathcal{F}(x), x^+ - x \rangle + \frac{L}{2} \|x^+ - x\|^2.$$

By projection optimality,  $(y - x^+)^\top (z - x^+) \leq 0$  for all  $z \in \mathcal{C}$ . Choosing  $z = x$  and expanding  $y = x - \eta \nabla \mathcal{F}(x)$  gives

$$\langle \nabla \mathcal{F}(x), x^+ - x \rangle \leq -\frac{1}{\eta} \|x^+ - x\|^2.$$

Substitute into the smoothness bound:

$$\mathcal{F}(x^+) \leq \mathcal{F}(x) - \frac{1}{\eta} \|x^+ - x\|^2 + \frac{L}{2} \|x^+ - x\|^2.$$

Since  $g_\eta(x) = \frac{1}{\eta}(x - x^+)$  and  $\eta \leq 1/L$ , we obtain

$$\mathcal{F}(x^+) \leq \mathcal{F}(x) - \frac{\eta}{2} \|g_\eta(x)\|^2.$$

□

1890 **Lemma 2.2** (Bounded Sum of Gradient Mappings). *Summing Lemma 2.4 over  $T$  iterations yields*

$$1891 \quad 1892 \quad \sum_{t=0}^{T-1} \|g_\eta(P_{\text{lig}}^t)\|^2 \leq \frac{2}{\eta} (\mathcal{F}(P_{\text{lig}}^0) - \mathcal{F}^*).$$

1893 *Proof.* Telescoping  $\mathcal{F}(P_{\text{lig}}^t) - \mathcal{F}(P_{\text{lig}}^{t+1}) \geq \frac{\eta}{2} \|g_\eta(P_{\text{lig}}^t)\|^2$  gives

$$1894 \quad 1895 \quad \mathcal{F}(P_{\text{lig}}^0) - \mathcal{F}(P_{\text{lig}}^T) \geq \frac{\eta}{2} \sum_{t=0}^{T-1} \|g_\eta(P_{\text{lig}}^t)\|^2.$$

1896 Since  $\mathcal{F}(P_{\text{lig}}^T) \geq \mathcal{F}^*$ , the bound follows.  $\square$

1897 **Lemma 2.3** (Stationary Point Convergence). *The sequence  $\{P_{\text{lig}}^t\}$  is bounded and has accumulation*

1898 *points. Every accumulation point  $P^*$  satisfies  $g_\eta(P^*) = 0$ , i.e.,  $P^*$  is a stationary point satisfying*

1899 *first-order optimality conditions.*

1900 *Proof.* Because  $\mathcal{C}_{\text{valid}}$  is closed and bounded, the iterates remain in a compact set, ensuring accumulation

1901 points exist. From Lemma 2.2,  $\liminf_{t \rightarrow \infty} \|g_\eta(P_{\text{lig}}^t)\| = 0$ . By continuity of  $g_\eta$ , any limit

1902 point satisfies  $g_\eta(P^*) = 0$ .  $\square$

1903 *Proof of Theorem 2.* Lemma 2.4 establishes monotonic descent. Lemma 2.2 bounds the cumulative

1904 gradient mapping norms. Lemma 2.3 implies every accumulation point is stationary. Thus the PGD

1905 sequence converges to stationary points under the convex surrogate assumption.  $\square$

1906 **Remark on Non-Convex Validity Manifold.** In practice, the chemical validity set  $\mathcal{C}_{\text{valid}}$  is highly

1907 non-convex due to valency, aromaticity, and stereochemistry constraints. The above proof holds

1908 only under the convex surrogate assumption, which serves as an *idealization*. When using generative

1909 validity repair (e.g.,  $\mathcal{G}_{\text{type}}$ ) instead of exact projection, the iteration behaves as an *inexact pro-*

1910 *jection method*. Standard results on inexact PGD imply convergence to an  $\varepsilon$ -stationary point, with

1911  $\varepsilon$  depending on the repair error. Thus the theoretical guarantee should be interpreted as *asymptotic*

1912 *convergence under convex surrogates*, while practical implementations achieve only *approximate*

1913 *stationarity*.

1914

1915

1916

1917

1918

1919

1920

1921

1922

1923

### E.3 PROOF OF THEORETICAL SUPERIORITY FOR INVERSION FRAMEWORK

1924 We compare the reachable objective values produced by (i) the common two-stage *generate-then-*

1925 *optimize* pipeline (G+O) and (ii) the gradient-driven *inversion* framework used in MagicDock. Let

1926  $(\mathcal{X}, \|\cdot\|)$  be a Euclidean space of parametrized structures (e.g., point clouds, coordinates, or fea-

1927 tures), and let

$$1928 \quad \mathcal{C}_{\text{valid}} \subseteq \mathcal{X}$$

1929 denote the feasible set of chemically valid structures. We assume  $\mathcal{C}_{\text{valid}}$  is closed. For analysis, we

1930 assume it is convex or that a convex surrogate is available; for nonconvex cases, we use a proximity

1931 operator (see Remark E.3.2). Let the design objective be

$$1932 \quad 1933 \quad \mathcal{F} : \mathcal{C}_{\text{valid}} \rightarrow \mathbb{R},$$

1934 which is differentiable, bounded below by  $\mathcal{F}^* > -\infty$ , and  $L$ -smooth:

$$1935 \quad 1936 \quad \|\nabla \mathcal{F}(x) - \nabla \mathcal{F}(y)\| \leq L \|x - y\|, \quad \forall x, y \in \mathcal{C}_{\text{valid}}.$$

1937 We model the two paradigms as follows.

#### E.3.1 GENERATE-THEN-OPTIMIZE (G+O)

1940 A generator  $\mathcal{G} : \mathcal{Z} \rightarrow \mathcal{C}_{\text{valid}}$  (with latent prior  $z \sim \mu_Z$ ) outputs  $x_0 = \mathcal{G}(z)$ . A local optimizer

1941  $\mathcal{O}$ , defined as a short-step projected gradient descent (PGD) with step size  $\eta \leq 1/L$ , maps  $x_0$  to a

1942 refined point  $x_\infty = \mathcal{O}(x_0)$ . The G+O reachable set is

1943

$$\mathcal{R}_{\text{G+O}} = \{\mathcal{O}(\mathcal{G}(z)) : z \in \text{supp}(\mu_Z)\}.$$

1944 E.3.2 INVERSION (GRADIENT-DRIVEN)  
19451946 Starting from an initialization  $x_0 \in \mathcal{C}_{\text{valid}}$ , inversion performs projected gradient-like updates:  
1947

1948 
$$x_{t+1} = \Pi_{\mathcal{C}_{\text{valid}}}(x_t - \eta_t \nabla \mathcal{F}(x_t)), \quad 0 < \eta_t \leq 1/L, \quad (64)$$

1949 where  $\Pi_{\mathcal{C}_{\text{valid}}}$  denotes the Euclidean projection onto  $\mathcal{C}_{\text{valid}}$ . Let  $\mathcal{R}_{\text{Inv}}$  be the set of accumulation  
1950 points of sequences generated by equation 64 initialized from all possible  $x_0 \in \mathcal{C}_{\text{valid}}$  (or from  $\mathcal{G}(z)$ ).  
1951 We show: (i) inversion satisfies standard descent and convergence guarantees, (ii)  $\mathcal{R}_{\text{G+O}} \subseteq \mathcal{R}_{\text{Inv}}$ ,  
1952 and (iii) when the generator  $\mathcal{G}$  is *support-misspecified*, inversion achieves strictly better infimal  
1953 objective values.1954 **Lemma 2.4** (Descent for Projected Gradient Updates). *Under the above assumptions, the update*  
1955 *equation 64 with constant step size  $\eta \in (0, 1/L]$  satisfies for every  $t$ :*

1956 
$$\mathcal{F}(x_{t+1}) \leq \mathcal{F}(x_t) - \frac{\eta}{2} \|g_\eta(x_t)\|^2,$$

1957 where  $g_\eta(x_t) = \frac{1}{\eta}(x_t - \Pi_{\mathcal{C}_{\text{valid}}}(x_t - \eta \nabla \mathcal{F}(x_t)))$  is the gradient mapping. Consequently,  
1958

1959 
$$\sum_{t=0}^{T-1} \|g_\eta(x_t)\|^2 \leq \frac{2}{\eta} (\mathcal{F}(x_0) - \mathcal{F}^*),$$

1960 and hence  $\min_{0 \leq t < T} \|g_\eta(x_t)\|^2 \rightarrow 0$  as  $T \rightarrow \infty$ .  
19611962 *Proof.* Let  $\tilde{x}_{t+1} := x_t - \eta \nabla \mathcal{F}(x_t)$  be the unconstrained gradient step, and  $x_{t+1} = \Pi_{\mathcal{C}_{\text{valid}}}(\tilde{x}_{t+1})$ .  
1963 By  $L$ -smoothness,  
1964

1965 
$$\mathcal{F}(y) \leq \mathcal{F}(x_t) + \nabla \mathcal{F}(x_t)^\top (y - x_t) + \frac{L}{2} \|y - x_t\|^2, \quad \forall y.$$

1966 Choose  $y = x_{t+1}$ :  
1967

1968 
$$\mathcal{F}(x_{t+1}) \leq \mathcal{F}(x_t) + \nabla \mathcal{F}(x_t)^\top (x_{t+1} - x_t) + \frac{L}{2} \|x_{t+1} - x_t\|^2.$$

1969 From projection optimality, for all  $z \in \mathcal{C}_{\text{valid}}$ ,  
1970

1971 
$$(\tilde{x}_{t+1} - x_{t+1})^\top (z - x_{t+1}) \leq 0.$$

1972 Choosing  $z = x_t$  gives  
1973

1974 
$$\langle \nabla \mathcal{F}(x_t), x_{t+1} - x_t \rangle \leq -\frac{1}{\eta} \|x_{t+1} - x_t\|^2.$$

1975 Plugging back,  
1976

1977 
$$\mathcal{F}(x_{t+1}) \leq \mathcal{F}(x_t) - \frac{1}{\eta} \|x_{t+1} - x_t\|^2 + \frac{L}{2} \|x_{t+1} - x_t\|^2.$$

1978 Since  $\|x_{t+1} - x_t\| = \eta \|g_\eta(x_t)\|$ , we get  
1979

1980 
$$\mathcal{F}(x_{t+1}) \leq \mathcal{F}(x_t) - \left(\frac{1}{\eta} - \frac{L}{2}\right) \eta^2 \|g_\eta(x_t)\|^2.$$

1981 For  $\eta \leq 1/L$ , the coefficient satisfies  $\frac{1}{\eta} - \frac{L}{2} \geq \frac{1}{2\eta}$ , hence  
1982

1983 
$$\mathcal{F}(x_{t+1}) \leq \mathcal{F}(x_t) - \frac{\eta}{2} \|g_\eta(x_t)\|^2.$$

1984 Summing over  $t = 0, \dots, T-1$  and using  $\mathcal{F}(x_T) \geq \mathcal{F}^*$  yields the claimed bound.  $\square$   
19851986 **Lemma 2.5** (Containment of G+O in Inversion). *For any generator  $\mathcal{G}$  and local optimizer  $\mathcal{O}$  (defined as PGD with step size  $\eta \leq 1/L$ ), if inversion is initialized at  $x_0 = \mathcal{G}(z)$ , the limit points attainable by  $\mathcal{O}(\mathcal{G}(z))$  are also attainable by running the inversion iterates equation 64 from the same initialization. Hence,*  
1987

1988 
$$\mathcal{R}_{\text{G+O}} \subseteq \mathcal{R}_{\text{Inv}}.$$

1998 *Proof.* Both  $\mathcal{O}$  (as PGD) and inversion iterates follow negative directional derivatives of  $\mathcal{F}$  with step  
 1999 sizes  $\eta \leq 1/L$ . Starting from  $x_0 = \mathcal{G}(z)$ , both methods generate sequences in the same attraction  
 2000 basin of a local stationary point (by smoothness and step-size constraints). Let  $x^* = \mathcal{O}(\mathcal{G}(z))$  be the  
 2001 limit of  $\mathcal{O}$ . By Lemma 2.4, the inversion sequence  $\{x_t\}$  decreases  $\mathcal{F}$  monotonically and is bounded  
 2002 below, so it has accumulation points that are stationary. Since both methods operate under the same  
 2003 smoothness and step-size regime, standard Hessian/stable manifold arguments (Absil et al. (2008))  
 2004 ensure convergence to the same local attractor. Thus,  $\mathcal{R}_{G+O} \subseteq \mathcal{R}_{Inv}$ .  $\square$

2005 **Definition 2.1** (Generator Misspecification). *Let the global (or basin) minimizers be*

$$2007 \quad \mathcal{M} := \arg \min_{x \in \mathcal{C}_{\text{valid}}} \mathcal{F}(x).$$

2009 *We say  $\mathcal{G}$  is misspecified if  $\mathcal{M} \not\subseteq \overline{\mathcal{R}_{G+O}}$ , i.e., some global (or deep basin) minimizers are unreach-  
 2010 able by sampling  $\mathcal{G}(z)$  and refining via  $\mathcal{O}$ .*

2011 **Lemma 2.6** (Strict Improvement under Misspecification). *If  $\mathcal{G}$  is misspecified and there exists an  
 2012 initialization  $x_0 \in \mathcal{C}_{\text{valid}}$  such that a sequence of PGD iterates from  $x_0$  can reach an  $\epsilon$ -neighborhood  
 2013 of some  $x^* \in \mathcal{M}$  with  $\mathcal{F}(x^*) < \inf_{x \in \mathcal{R}_{G+O}} \mathcal{F}(x)$ , then*

$$2014 \quad \inf_{x \in \mathcal{R}_{Inv}} \mathcal{F}(x) < \inf_{x \in \mathcal{R}_{G+O}} \mathcal{F}(x).$$

2017 *Proof.* By misspecification, there exists  $x^* \in \mathcal{M} \setminus \overline{\mathcal{R}_{G+O}}$  such that  $\mathcal{F}(x^*) < \inf_{x \in \mathcal{R}_{G+O}} \mathcal{F}(x)$ .  
 2018 Assume there exists an initialization  $x_0 \in \mathcal{C}_{\text{valid}}$  (possibly  $x_0 = \mathcal{G}(z)$ ) from which PGD it-  
 2019 erates equation 64 reach an  $\epsilon$ -neighborhood of  $x^*$ . By Lemma 2.4, each PGD step reduces  $\mathcal{F}$   
 2020 by at least  $\frac{\eta}{2} \|g_\eta(x_t)\|^2$ , and the sequence  $\{x_t\}$  has accumulation points that are stationary, with  
 2021  $\min_{0 \leq t < T} \|g_\eta(x_t)\|^2 \rightarrow 0$  as  $T \rightarrow \infty$ . Since  $\mathcal{C}_{\text{valid}}$  is closed and  $\mathcal{F}$  is continuous, with suffi-  
 2022 ciently small  $\eta$  and sufficient iterations, PGD can approach  $\mathcal{F}(x^*)$  arbitrarily closely (Beck (2017)).  
 2023 For PGD to reach the  $\epsilon$ -neighborhood of  $x^*$ , assume  $\mathcal{F}$  has a structure (e.g., satisfying the Polyak-  
 2024 Łojasiewicz condition in a basin around  $x^*$ ) or PGD employs randomized perturbations (Jin et al.  
 2025 (2017)) to escape local minima and converge toward global minimizers. Because  $x^* \notin \overline{\mathcal{R}_{G+O}}$  and  
 2026  $\mathcal{F}(x^*) < \inf_{x \in \mathcal{R}_{G+O}} \mathcal{F}(x)$ , the infimum over  $\mathcal{R}_{Inv}$  is strictly smaller.  $\square$

2027 **Theorem 3** (Inversion Weakly Dominates G+O; Strict Advantage under Misspecification). *Under  
 2028 the stated assumptions,*

$$2029 \quad \inf_{x \in \mathcal{R}_{Inv}} \mathcal{F}(x) \leq \inf_{x \in \mathcal{R}_{G+O}} \mathcal{F}(x), \quad (65)$$

2031 *with strict inequality if  $\mathcal{G}$  is misspecified and PGD from some  $x_0 \in \mathcal{C}_{\text{valid}}$  reaches an  $\epsilon$ -neighborhood  
 2032 of a better minimizer (Lemma 2.6).*

2033

2034 *Proof.* By Lemma 2.5,  $\mathcal{R}_{G+O} \subseteq \mathcal{R}_{Inv}$ , so

$$2035 \quad \inf_{x \in \mathcal{R}_{Inv}} \mathcal{F}(x) \leq \inf_{x \in \mathcal{R}_{G+O}} \mathcal{F}(x).$$

2037 If  $\mathcal{G}$  is misspecified and the condition of Lemma 2.6 holds, the infimum over  $\mathcal{R}_{Inv}$  is strictly smaller,  
 2038 completing the proof.  $\square$

2040 **Remark** (Nonconvex  $\mathcal{C}_{\text{valid}}$  and Chemical Constraints). The convexity assumption on  $\mathcal{C}_{\text{valid}}$  sim-  
 2041 plifies projection. In practice, chemical validity constraints (e.g., bond lengths, angles) are often  
 2042 nonconvex. We address this by using a convex surrogate or a learned decoder mapping iterates to  
 2043  $\mathcal{C}_{\text{valid}}$ . For nonconvex sets,  $\Pi_{\mathcal{C}_{\text{valid}}}$  is replaced by a proximity operator, and descent guarantees hold  
 2044 under regularity conditions (Rockafellar & Wets (1998)). In molecular design, decoders trained on  
 2045 valid structures ensure iterates remain feasible, preserving the qualitative conclusions.

2046 **Remark** (Practical Considerations). 1. *Discretization and Projection Error:* The proof is non-  
 2047 asymptotic, focusing on set relations. In practice, discretization, projection approximations,  
 2048 and finite iterations introduce errors. These are mitigated by small step sizes and high-  
 2049 fidelity decoders, ensuring PGD closely tracks theoretical descent paths.

2050 2. *Generator Limitations:* Diffusion- or flow-based samplers ( $\mathcal{G}$ ) have support limited by  
 2051 training data. Inversion escapes these limitations by iteratively refining beyond  $\mathcal{G}$ 's sup-  
 2052 port, achieving better minima.

2052 3. *Descent Path Relaxation*: Lemma 2.6 relaxes the continuous descent path assumption to  
 2053 PGD reaching an  $\epsilon$ -neighborhood, which is more practical for complex  $\mathcal{F}$  landscapes with  
 2054 multiple basins.  
 2055

2056 4. *Computational Constraints*: PGD requires more iterations than G+O but offers better ex-  
 2057 ploration. In molecular design, computational cost is offset by learned gradients or surro-  
 2058 gate models reducing evaluation complexity.

2059 **E.4 PROOF OF EFFICIENCY FOR INVERSION FRAMEWORK**

2061 To rigorously demonstrate the superior efficiency of the inversion architecture in terms of theoretical  
 2062 training and convergence iterations compared to GAN, diffusion, and flow-based generative meth-  
 2063 ods, we analyze the computational complexities under standard optimization assumptions. We as-  
 2064 sume models of comparable scale, with  $P$  parameters, dataset size  $N$ , latent or molecular dimension  
 2065  $D$ , and target accuracy  $\epsilon > 0$ . The inversion pre-training phase minimizes a convex reconstruction  
 2066 loss via gradient descent, while the generation phase optimizes a smooth, potentially non-convex  
 2067 objective per sample. In contrast, alternative methods incur overheads from adversarial dynam-  
 2068 ics, timestep iterations, or invertible transformations. The proofs derive big- $O$  bounds on iterations  
 2069 required for  $\epsilon$ -accuracy, highlighting inversion’s reduced complexity.

2070 **Theorem 4** (Training Efficiency). *Under  $\mu$ -strong convexity and  $L$ -smoothness assumptions for*  
 2071 *the loss, inversion pre-training converges in  $O((L/\mu) \log 1/\epsilon)$  iterations, outperforming GAN’s*  
 2072  *$O(1/\epsilon^2)$  lower bound for saddle-point equilibria, diffusion’s  $O(T/\epsilon^2)$  with diffusion steps  $T$ , and*  
 2073 *flow-based models’  $O(D^2 \log 1/\epsilon)$  per iteration due to Jacobian computations.*

2074 *Proof.* We derive the complexities sequentially for each method, starting from fundamental opti-  
 2075 mization rates and incorporating method-specific costs.

2077 For inversion pre-training, consider a  $\mu$ -strongly convex and  $L$ -smooth loss  $\mathcal{L}(\theta)$ , minimized via  
 2078 gradient descent:  $\theta^{k+1} = \theta^k - \eta \nabla \mathcal{L}(\theta^k)$  with  $\eta = 2/(\mu + L)$ . The suboptimality gap satisfies

$$\begin{aligned} \mathcal{L}(\theta^{k+1}) - \mathcal{L}^* &\leq (1 - \mu/L)(\mathcal{L}(\theta^k) - \mathcal{L}^*) \\ &\leq (1 - \mu/L)^k(\mathcal{L}(\theta^0) - \mathcal{L}^*), \end{aligned}$$

2083 yielding  $k = O((L/\mu) \log 1/\epsilon)$  iterations for  $\epsilon$ -accuracy. Per iteration cost is  $O(NP)$ , leading to  
 2084 total complexity  $O(NP(L/\mu) \log 1/\epsilon)$ .

2085 By contrast, GAN training solves the non-convex non-concave minimax problem  
 $\min_G \max_D \mathbb{E}[\log D(x)] + \mathbb{E}[\log(1 - D(G(z)))]$ . Without global convergence guarantees,  
 2086 lower bounds for finding  $\epsilon$ -local Nash equilibria require  $O(1/\epsilon^2)$  iterations in the worst case,  
 2087 derived from the quadratic growth of subgradients near equilibria and stochastic gradient oracle  
 2088 queries. Empirical instability further amplifies effective iterations, with total cost  $O(2NP/\epsilon^2)$   
 2089 accounting for dual networks.

2091 Similarly, diffusion models train a denoiser over  $T$  timesteps, minimizing  $\sum_{t=1}^T \mathbb{E}[\|\epsilon - \hat{\epsilon}(x_t, t)\|^2]$ .  
 2092 For empirical risk minimization with variance  $O(1/t)$ , stochastic gradient descent converges in  
 2093  $O(T/\epsilon^2)$  iterations to achieve  $\epsilon$ -error, as the timestep aggregation scales the variance bound linearly  
 2094 with  $T$ . Per iteration cost  $O(NP)$  results in total complexity  $O(NPT/\epsilon^2)$ , where  $T$  is typically  
 2095 large to capture fine-grained noise schedules.

2096 For flow-based models, exact likelihood maximization involves  $-\log p(x) = -\log p(z) -$   
 2097  $\log |\det Df|$ , with Jacobian determinant computation costing  $O(D^3)$  for general flows or  $O(D^2)$  for  
 2098 structured autoregressive variants. Convergence mirrors VAEs at  $O((L/\mu) \log 1/\epsilon)$  iterations, but  
 2099 augmented per-iteration cost yields  $O(N(P + D^2)(L/\mu) \log 1/\epsilon)$ , dominated by the determinant in  
 2100 high-dimensional molecular spaces.

2101 Given that  $\log 1/\epsilon \ll 1/\epsilon^2$  for small  $\epsilon$ , and  $T, D^2 \gg (L/\mu)$ , inversion exhibits lower training  
 2102 complexity.  $\square$

2104 **Theorem 5** (Generation Efficiency). *For generating  $M$  samples to  $\epsilon$ -stationarity under  $L$ -*  
 2105 *smooth non-convex objectives, inversion requires  $O(M/\epsilon^2)$  iterations, fewer than diffusion’s*  
 $O(MT \log 1/\epsilon)$ , GAN’s training-dominant cost, and flow’s  $O(MD)$  per-sample inversion.

2106 *Proof.* Continuing the sequential derivation, we focus on per-sample generation complexities post-  
 2107 training.

2108 In inversion generation, per-sample optimization of an  $L$ -smooth non-convex  $\mathcal{F}(x)$  via gradient de-  
 2109 scent achieves  $\min_k \|\nabla \mathcal{F}(x^k)\|^2 \leq O(L(\mathcal{F}(x^0) - \mathcal{F}^*)/K)$  after  $K$  steps, requiring  $K = O(1/\epsilon^2)$   
 2110 for  $\epsilon$ -stationarity. Total cost:  $O(MP/\epsilon^2)$ .

2112 In GAN generation, post-training sampling is a single forward pass,  $O(MP)$ , but the adversarial  
 2113 training overhead dominates overall efficiency, rendering it less favorable for iterative refinement  
 2114 tasks.

2115 Diffusion sampling reverses  $T$  timesteps, with accelerated solvers (e.g., DDIM) converging  
 2116 in  $O(T \log 1/\epsilon)$  steps to  $\epsilon$ -fidelity via controlled noise reduction. Thus, total generation:  
 2117  $O(MPT \log 1/\epsilon)$ .

2118 Flow-based sampling inverts the bijective transform, costing  $O(M(P + D))$  due to sequential or  
 2119 matrix operations in high dimensions.

2120 With  $1/\epsilon^2 \ll T \log 1/\epsilon, D$  in practice, inversion’s generation phase is more efficient for large  $M$ . □  
 2121

2122  
 2123  
 2124 To elucidate these complexities and underscore the potential of inversion architectures, we present  
 2125 a comparative summary in Table 7. The table delineates the asymptotic bounds for training and  
 2126 generation phases, revealing inversion’s advantages in reduced dependence on auxiliary factors like  
 2127 timesteps  $T$  or dimension  $D$ . This manifests in faster convergence and lower overall computational  
 2128 overhead, particularly beneficial for de novo ligand design where iterative optimization aligns natu-  
 2129 rally with docking-driven objectives.  
 2130

2131 Table 7: Complexity Comparison of Generative Frameworks  
 2132

Architecture	Training Complexity	Generation Complexity (per $M$ samples)
Inversion	$O(NP \log 1/\epsilon)$	$O(MP/\epsilon^2)$
GAN	$O(NP/\epsilon^2)$	$O(MP)$ (training-dominant)
Diffusion	$O(NPT/\epsilon^2)$	$O(MPT \log 1/\epsilon)$
Flow-Based	$O(N(P + D^2) \log 1/\epsilon)$	$O(M(P + D))$

## 2141 E.5 PROOF OF INFORMATION-THEORETIC SUPERIORITY FOR INVERSION FRAMEWORK

2142 To rigorously establish the superiority of the inversion architecture from an information-theoretic  
 2143 perspective, we derive bounds on mutual information and entropy, demonstrating enhanced infor-  
 2144 mation transfer from docking signals to generated ligands compared to GAN, diffusion, and flow-  
 2145 based models. This analysis underscores inversion’s ability to maximize diversity (entropy) while  
 2146 minimizing conditional uncertainty, leading to more faithful and varied de novo designs.

2147 **Theorem 6.** *Let  $I(X; Y)$  denote the mutual information between docking signals  $X$  (receptor  
 2148 features and affinity objectives) and generated ligands  $Y$ . The inversion framework maximizes  
 2149  $I(X; Y) = H(Y) - H(Y | X)$  relative to alternatives, with  $H(Y) > H(Y_{GAN})$  (countering  
 2150 mode collapse) and  $H(Y | X) < H(Y | X_{Diffusion})$  (reducing noise-induced uncertainty), implying  
 2151 superior information efficiency and diversity.*

2152  
 2153  
 2154  
 2155 *Proof.* We derive the bounds sequentially, leveraging variational information decompositions and  
 2156 kernel entropy approximations for multi-modal molecular distributions.

2157 Consider the joint distribution  $p(X, Y)$  under each architecture. Mutual information  $I(X; Y) =$   
 2158  $\int p(x, y) \log \frac{p(x, y)}{p(x)p(y)} dx dy = H(Y) - H(Y | X)$ , where  $H(Y)$  quantifies output diversity (entropy  
 2159 over ligand space) and  $H(Y | X)$  measures conditional uncertainty given docking signals.

For inversion, the process optimizes a conditional energy  $\mathcal{F}(Y; X)$ , yielding  $p(Y | X) \propto \exp(-\beta\mathcal{F}(Y; X))$ . Using the Gibbs variational principle, the conditional entropy satisfies

$$\begin{aligned} H(Y | X) &= -\mathbb{E}_{p(Y|X)}[\log p(Y | X)] \\ &\leq \log Z_X + \beta\mathbb{E}_{p(Y|X)}[\mathcal{F}(Y; X)], \end{aligned}$$

where  $Z_X = \int_{\mathcal{C}_{\text{valid}}} \exp(-\beta\mathcal{F}(y; X)) dy$  is the partition function bounded by the valid chemical space volume. Gradient-driven minimization of  $\mathcal{F}$  reduces the expectation term, yielding low  $H(Y | X)$  (precise signal-to-structure mapping). Simultaneously, basin exploration via continuous flows maximizes  $H(Y) \approx \log |\mathcal{C}_{\text{valid}}| - \beta \min \mathcal{F}$ , enhancing diversity.

By contrast, GANs approximate  $p(Y)$  via adversarial minimization of Jensen-Shannon divergence, but mode collapse truncates support:  $H(Y_{\text{GAN}}) \leq H(Y) - \Delta$ , where  $\Delta = \sum_i p_i \log(1/p_i)$  over dropped modes  $i$  (from Fano's inequality on collapsed distributions). Thus,

$$\begin{aligned} I(X; Y_{\text{GAN}}) &= H(Y_{\text{GAN}}) - H(Y_{\text{GAN}} | X) \\ &\leq H(Y) - \Delta - H(Y | X) + o(1), \end{aligned}$$

reducing mutual information due to diminished diversity.

Diffusion models parameterize a timestep-dependent process  $p(Y_t | Y_{t-1}, X)$ , with reverse chain entropy decomposed as  $H(Y | X) = \sum_{t=1}^T H(Y_t | Y_{t-1}, X)$ . Noise variance at each step inflates conditional terms:  $H(Y_t | Y_{t-1}, X) \geq \frac{1}{2} \log(2\pi e \sigma_t^2)$ , yielding

$$\begin{aligned} H(Y | X) &\geq \sum_{t=1}^T \frac{1}{2} \log(2\pi e \sigma_t^2) \\ &= \frac{T}{2} \log(2\pi e \bar{\sigma}^2) > H(Y | X)_{\text{Inversion}}, \end{aligned}$$

where  $\bar{\sigma}^2$  is average noise variance, increasing with  $T$  and reducing  $I(X; Y_{\text{Diffusion}})$ .

Flow-based models enforce exact likelihood via invertible transforms, but information is constrained by the Jacobian:  $H(Y) = H(Z) + \mathbb{E}[\log |\det Df|]$ , with base entropy  $H(Z)$  (e.g., Gaussian) limiting expressivity. For high-dimensional  $D$ , the determinant approximation error bounds  $H(Y) \leq H(Z) + O(D \log D)$ , often lower than inversion's exploration of full  $\mathcal{C}_{\text{valid}}$ .

Aggregating these, inversion maximizes  $I(X; Y)$  by balancing high  $H(Y)$  and low  $H(Y | X)$ , proving the theorem.  $\square$

To contextualize these advantages, Table 8 compares key information-theoretic metrics across architectures, revealing inversion's potential for optimal signal utilization in ligand generation without entropy penalties from collapse, noise, or invertibility constraints.

Table 8: Information-Theoretic Metric Comparison

Architecture	$I(X; Y)$ Bound	$H(Y)$ Factor	Limitation
Inversion	$\approx H(Y) - o(1)$	$\log  \mathcal{C}_{\text{valid}} $	None
GAN	$\leq H(Y) - \Delta$	Suboptimal	Mode Collapse
Diffusion	$\leq H(Y) - \frac{T}{2} \log \bar{\sigma}^2$	$O(T \log D)$	Noise Variance
Flow-Based	$= H(Z) + O(D \log D)$	Gaussian-Bounded	Transform Rigidity

The table underscores inversion's maximal mutual information and entropy, free from method-specific artifacts, highlighting its superiority in capturing diverse, docking-informed distributions.

## E.6 PROOF OF SAMPLE COMPLEXITY ADVANTAGE FOR INVERSION FRAMEWORK

We provide a theoretical justification for the data efficiency of inversion-based fine-tuning, clarifying the distinction between *sample complexity* (number of labeled examples required) and *optimization complexity* (number of gradient iterations).

2214  
 2215 **Theorem 7** (Sample Complexity for Fine-Tuning with Pretrained Features). *Assume that the pre-  
 2216 trained backbone is frozen and only a linear prediction head  $w \in \mathbb{R}^{d_{\text{eff}}}$  is fine-tuned. Suppose the  
 2217 supervised loss is  $\mu$ -strongly convex and  $L$ -smooth in  $w$ , the observation noise is sub-Gaussian with  
 2218 variance proxy  $\sigma^2$ , and the feature representation is bounded as  $\|\phi(x)\| \leq B$  almost surely. Then,  
 2219 with probability at least  $1 - \delta$ , the excess risk satisfies*

$$2220 \quad \mathcal{L}(\hat{w}) - \mathcal{L}(w^*) \leq \epsilon \quad \text{whenever} \quad n = O\left(\frac{\sigma^2 B^2 d_{\text{eff}}}{\mu \epsilon} \log \frac{1}{\delta}\right).$$

$$2221$$

2222 *Proof.* With frozen features, fine-tuning reduces to empirical risk minimization of a linear predictor  
 2223 in  $d_{\text{eff}}$  dimensions under strong convexity. Standard results in stochastic convex optimization and  
 2224 statistical learning theory (e.g., Bernstein inequalities for sub-Gaussian noise with bounded features)  
 2225 yield the high-probability bound  
 2226

$$2227 \quad \mathcal{L}(\hat{w}) - \mathcal{L}(w^*) = O\left(\frac{\sigma^2 B^2 d_{\text{eff}}}{n \mu} \log \frac{1}{\delta}\right).$$

$$2228$$

2229 Solving for  $n$  to guarantee excess risk  $\leq \epsilon$  with probability at least  $1 - \delta$  establishes the stated  
 2230 bound. Note that  $\mu$  enters as the curvature modulus,  $B^2 d_{\text{eff}}$  captures the effective feature scale, and  
 2231  $\sigma^2$  scales the variance.  $\square$

$$2232$$

2233 **Remark.** The  $\log(1/\epsilon)$  dependence often cited in the optimization literature refers to the *iteration*  
 2234 *complexity* of gradient descent for strongly convex losses, where the optimization error decreases  
 2235 geometrically. Here, however, we are concerned with statistical sample complexity, which scales as  
 2236  $O(1/\epsilon)$  (up to logarithmic confidence factors) under convexity and low effective dimension. If the  
 2237 linear prediction head cannot fully capture the target mapping, an additional approximation error  
 2238 term  $\mathcal{L}(w_{\text{best}}) - \mathcal{L}(w^*)$  should be included.

$$2239$$

2240 **Comparison to Generative Models.** In contrast, empirical evidence and partial theoretical analyses  
 2241 suggest that end-to-end generative models such as GANs or diffusion models typically require  
 2242 more labeled samples due to higher variance and larger hypothesis classes. GAN training involves  
 2243 min-max optimization with adversarial variance, while diffusion models require denoising across  
 2244 multiple timesteps, effectively inflating variance with horizon length  $T$ . Although deriving universal  
 2245  $O(1/\epsilon^2)$  or  $O(T/\epsilon^2)$  bounds is challenging and problem-specific, empirical findings consistently  
 2246 indicate substantially higher sample demands compared to inversion-based fine-tuning, particularly  
 2247 in scarce-data docking regimes. These comparisons should be interpreted as qualitative rather than  
 2248 universal guarantees.

$$2249$$

## F LIMITATIONS AND FUTURE WORK

2250 Although MagicDock presents a promising unified framework for docking-oriented de novo ligand  
 2251 design, it is not without limitations. One primary concern lies in the reliance on gradient-  
 2252 driven inversion, which, while effective for end-to-end optimization, may converge to local minima  
 2253 in highly non-convex energy landscapes, potentially overlooking globally optimal configurations.  
 2254 Additionally, the surface point cloud representation, though versatile for unifying proteins and  
 2255 small molecules, inherently abstracts away internal volumetric details and dynamic conformational  
 2256 changes, which could compromise accuracy in scenarios involving flexible receptors or allosteric effects.  
 2257 Computational demands also pose a challenge; the iterative refinement in the inversion stage  
 2258 and the need for  $\text{SE}(3)$ -equivariant pre-training require substantial resources, limiting scalability  
 2259 for very large molecular systems or high-throughput applications. Furthermore, the framework's  
 2260 performance is contingent on the quality and diversity of fine-tuning data, raising questions about  
 2261 generalization to underrepresented protein families or novel therapeutic targets. Finally, the current  
 2262 atomic-level feature set is limited, as it mainly accounts for common organic atoms (e.g., C, H, O, N,  
 2263 S, Se), while neglecting halogens and metal ions that frequently appear in pharmaceutically relevant  
 2264 ligands and cofactors, potentially restricting applicability in broader chemical spaces.

2265 Looking ahead, future work could address these limitations through several avenues. Enhancing the  
 2266 inversion process with stochastic or meta-learning techniques might mitigate local optima issues,  
 2267 enabling more robust exploration of the chemical space. Integrating hybrid representations that

2268 combine surface abstractions with volumetric or graph-based models could capture richer biophysical interactions, while advances in efficient equivariant architectures may reduce computational overhead. Extending the atom feature vocabulary to include halogens (e.g., F, Cl, Br, I) and metal centers (e.g., Mg, Zn, Fe, Cu) would broaden the framework’s coverage of bioactive compounds and metalloproteins. In practice, incorporating halogens only requires minor adjustments to the feature extraction process, and preliminary experiments indicate that ligand generation performance with halogen atoms is largely consistent with the results reported in this work. In contrast, the inclusion of metal elements remains more challenging: many metal centers participate in complex coordination phenomena that cannot be easily captured by the current modeling framework, leading to sub-optimal results. Addressing these cases may require specialized representations or physics-inspired modeling of coordination chemistry. Expanding the scope to incorporate multi-objective optimization—such as balancing binding affinity with pharmacokinetic properties—or real-time adaptive docking for dynamic simulations would broaden applicability. Finally, empirical validation on diverse wet-lab datasets and collaboration with experimental biologists could refine the model, paving the way for practical deployment in drug discovery pipelines.

## G ALGORITHMS IN PSEUDO CODE

### G.1 STAGE 1: DOCKING-ORIENTED LIGAND MODELING

The algorithm generates protein surface point clouds via Gaussian sampling around atoms weighted by van der Waals radii, computes SDFs, and encodes multi-level (chemical, atomic, geometric) patch features for efficient docking hotspot representation in de novo design.

---

#### Algorithm 1: Stage 1 — Surface Point-Cloud Construction and Patching

---

**Require:** Atomic coordinates  $\{x_a^j\}_{j=1}^N$ , atom types, radii  $\{\sigma_a^j\}$ , molecule type (protein/small molecule).

**Ensure:** Surface points  $X_s$ , normals  $N$ , features  $F$ , patches  $(X_p, F_p)$ .

- 1: **Hyper-parameters:**
- 2: Protein:  $\eta_p, r_{\text{iso}}^p, M_p$ ; Small molecule:  $\eta_m, r_{\text{iso}}^m, M_m$ .
- 3: Shared:  $\sigma_{\text{upsample}}, T_{\text{sdf}}, \alpha_{\text{sdf}}, \rho, K$ .
- 4: // (A) Candidate surface points
- 5: **for** each atom  $x_a^j$  **do**
- 6: Sample  $\eta_p$  or  $\eta_m$  candidates  $\tilde{x} \sim \mathcal{N}(x_a^j, \sigma_{\text{upsample}}^2 I)$  based on molecule type.
- 7: **end for**
- 8: Collect candidates  $\{x_s^i\}$ .
- 9: // (B) Smooth distance function (SDF)
- 10: Compute  $\text{SDF}(x_s^i)$  per Eq. 4.
- 11: // (C) Converge to iso-surface
- 12: **for**  $t = 1 \dots T_{\text{sdf}}$  **do**
- 13: Update  $x_s^i \leftarrow x_s^i - \alpha_{\text{sdf}} \nabla_{x_s^i} (\text{SDF}(x_s^i) - r_{\text{iso}})^2$ .
- 14: **end for**
- 15: // (D) Sampling and normals
- 16: Sample  $M_p$  (protein) or  $M_m$  (small molecule) points  $\Rightarrow X_s$ .
- 17: Compute normals  $N_i = \nabla \text{SDF}(x_i) / \|\nabla \text{SDF}(x_i)\|$ .
- 18: // (E) Feature generation
- 19: **if** protein **then**
- 20:  $f(x_i) = \text{concat}(f_{\text{chem}}^p, f_{\text{atom}}, f_{\text{geom}})$ ,  $f_{\text{chem}}^p$ : 6D one-hot  $\{C, H, O, N, S, Se\}$ .
- 21: **else**
- 22:  $f(x_i) = \text{concat}(f_{\text{chem}}^m, f_{\text{atom}}, f_{\text{geom}})$ ,  $f_{\text{chem}}^m$ : 8D one-hot  $\{C(sp3), C(sp2), \dots\}$ .
- 23: **end if**
- 24: Shared:  $f_{\text{atom}}$  (6D/8D),  $f_{\text{geom}}$  (curvatures, density, coords; 6D).
- 25: // (F) Patch partitioning
- 26: Patch centers  $X_c = \text{FPS}(X_s, \rho)$ .
- 27: For each  $c \in X_c$ , patch  $P(c) = \text{KNN}(c, X_s; K)$ .
- 28: Form  $X_p \in \mathbb{R}^{\rho M \times K \times 3}$ ,  $F_p \in \mathbb{R}^{\rho M \times K \times d}$ .
- 29: **return**  $(X_s, N, F, X_p, F_p)$ .

---

2322 G.2 STAGE 2: UNSUPERVISED PRE-TRAINING  
23232324 The algorithm details SE(3)-equivariant pre-training via masked reconstruction, encoding surface  
2325 patches into invariant features with spherical harmonics and graph convolutions.  
23262327 **Algorithm 2:** Stage 2 — Unsupervised Pre-Training  
2328

---

**Require:** Patch sets  $\{(X_p, F_p)\}$  from Stage 1  
**Ensure:** Encoder  $\mathcal{E}_\Theta$ , decoder  $D_\Phi$ , codebook  $\mathcal{E} = \{e_j\}_{j=1}^{N_B}$   
1: **Hyper-params:**  $\delta, N_B, \tau, L, K, B, \eta$   
2: **for** each epoch **do**  
3:   **for** each minibatch  $(X_p, F_p)$  **do**  
4:     Compute SE(3)-equivariant  $z_i$  for each patch point (Eq. 7)  
5:     Mask  $\delta$  fraction of patches  $\mathcal{M}$ , visible  $\mathcal{V} = \mathcal{M}$   
6:     For  $i \in \mathcal{M}$ , get hidden  $h_{p,i,m}$ , sample codebook (Eq. 32)  
7:     Concat tokens  $H_p^{(0)} = \text{concat}(\{z_i\}_{i \in \mathcal{V}}, \{z_{p,i,m}\}_{i \in \mathcal{M}})$   
8:     Decode  $H_p^{(0)} \Rightarrow H_p^{(L_2)}$   
9:     Compute coords (Eq. 33), curvature via covariance (Eq. 34), eigenvalues,  $\hat{\psi}_i$  (Eq. 35)  
10:    Compute losses: Chamfer (Eq. 37),  $\mathcal{L}_{\text{cur}} = \frac{1}{\delta \rho M} \sum_{i=1}^{\delta \rho M} \|\psi_i - \hat{\psi}_i\|_2^2, \mathcal{L}_{\text{KL}}(q, p)$   
11:    Total loss (Eq. 9)  
12:    Update  $\Theta, \Phi$ , codebook via backprop  
13:   **end for**  
14: **end for**  
15: **return**  $\mathcal{E}_\Theta, D_\Phi, \mathcal{E}$ 


---

2346

2347 G.3 STAGE 3: SUPERVISED FINE-TUNING  
23482349 The algorithm employs SE(3)-equivariant supervised fine-tuning via attention-based aggregation of  
2350 receptor-ligand latent fields for pocket prediction, interaction modeling, and geometric regularization,  
2351 boosting docking affinity with BCE and MSE losses.  
23522353 **Algorithm 3:** Stage 3 — Supervised Fine-Tuning with Equivariant Attention  
2354

---

**Require:** Labeled complexes  $\{(R, L, y_{\text{pocket}}, y_{\text{int}}, \Delta G)\}$ , encoder  $\mathcal{E}_\Theta$   
**Ensure:** Fine-tuned  $\mathcal{E}_\Theta^*$ , heads  $h_{\text{pocket}}, h_{\text{int}}, h_{\Delta G}$   
1: **Hyper-params:**  $\eta_{\text{enc}}, \eta_{\text{head}}, B, \alpha, \beta, \lambda_p, E$   
2: **for** each epoch **do**  
3:   **for** each minibatch **do**  
4:     Encode receptor/ligand patches:  $Z_r = \mathcal{E}_\Theta(X_p^r, F_p^r), Z_l = \mathcal{E}_\Theta(X_p^l, F_p^l)$   
5:     Split features into irreps:  $Z_r = \{Z_r^{(\ell)}\}_{\ell=0}^L, Z_l = \{Z_l^{(\ell)}\}_{\ell=0}^L$   
6:     For each  $\ell$ , build  $Q_r^{(\ell)} = Z_r^{(\ell)} W_Q^{(\ell)}, K_l^{(\ell)} = Z_l^{(\ell)} W_K^{(\ell)}, V_l^{(\ell)} = Z_l^{(\ell)} W_V^{(\ell)}$   
7:     Compute scalar attention scores (using  $\ell = 0$  channels only):  $A = \text{softmax}\left(\frac{Q_r^{(0)}(K_l^{(0)})^\top}{\sqrt{d_0}}\right)$   
8:     Aggregate values equivariantly:  $\tilde{Z}_r^{(\ell)} = \sum_j A_{ij} \rho^{(\ell)}(R) V_{l,j}^{(\ell)}, \forall \ell = 0, \dots, L$   
      where  $\rho^{(\ell)}(R)$  is the Wigner-D matrix ensuring SO(3)-equivariance.  
9:     Concatenate updated irreps  $\tilde{Z}_r = \{\tilde{Z}_r^{(\ell)}\}_{\ell=0}^L$   
10:    Predict pocket labels:  $\hat{y}_i = \sigma(h_{\text{pocket}}(\tilde{z}_i))$   
11:    Predict interaction:  $\hat{y}_{\text{int}} = \sigma(h_{\text{int}}(\tilde{z}))$   
12:    Predict affinity:  $\hat{y}_{\Delta G} = h_{\Delta G}(\tilde{z})$   
13:    Compute  $\mathcal{L}_{\text{pocket}}, \mathcal{L}_{\text{int}}, \mathcal{L}_{\Delta G}$   
14:    Total loss  $\mathcal{L}_{\text{FT}}$  (Eq. 43)  
15:    Update encoder and heads by gradient descent  
16:   **end for**  
17: **end for**  
18: **return**  $\mathcal{E}_\Theta^*, h_{\text{pocket}}, h_{\text{int}}, h_{\Delta G}$ 


---

2376 G.4 STAGE 4: INVERSION-BASED LIGAND GENERATION  
2377

2378

2379 The algorithm uses gradient-based inversion to generate ligands from noise, iteratively updating  
2380 coordinates via backpropagation to minimize docking energy, validity, and structural losses. Equiv-  
2381 ariant graphs project updates into valid chemical space, enabling direct high-affinity design without  
2382 generative models.  
23832384 **Algorithm 4:** Stage 4 — Inversion-based Ligand Generation2385 **Require:** Receptor structure  $R$ , fine-tuned encoder  $\mathcal{E}_\Theta^*$  (Stage 3), docking heads  $h_{\text{pocket}}, h_{\text{int}}, h_{\Delta G}$ ,  
2386 initial ligand/protein seed  $S^{(0)}$ .2387 **Ensure:** Optimized ligand/protein  $S^*$  with high binding affinity.2388 1: **Initialization:** Perform one round of (A) encoding and (B) supervision on initial seed  $S^{(0)}$ .2389 2: **for**  $t = 0 \dots T_{\text{gen}} - 1$  **do**2390 3:   **(A) Encoding.** Construct surface point clouds for  $R$  and  $S^{(t)}$ , and encode via  $\mathcal{E}_\Theta^*$ . Fuse  
2391   embeddings using Eq. 10 to obtain  $\tilde{z}^{(t)}$ .2392 4:   **(B) Supervision.** Compute docking-related predictions (pocket, interaction, affinity).  
2393   Evaluate  $\mathcal{L}_{\text{FT}}^{(t)}$  (Eq. 43) and obtain gradients  $\nabla_{Z_l} \mathcal{L}_{\text{FT}}^{(t)}$ .2394 5:   **(C) Gradient-guided modification.**

2395 6:   Identify positions (residues/atoms) with high gradient magnitude.

2396 7:   Select candidate operations: add, modify, delete.

2397 8:   Accept modification with probability

2398  
2399 
$$P(o) = \frac{\exp(-\Delta\Delta G^{(t,o)}/\tau_{\text{acc}})}{\sum_{o'} \exp(-\Delta\Delta G^{(t,o')}/\tau_{\text{acc}})}.$$
  
2400

2401   Update to  $S^{(t+1)}$ .2402 9:   **(D) Biochemical constraints and relaxation.**2403 10:   **if** Protein **then**2404 11:   Enforce structural feasibility: check residue packing, adjust backbone torsions, perform  
2405   local energy minimization.2406 12:   **else if** Small molecule **then**2407 13:   Apply chemical constraints: enforce valence, test aromaticity and ring closure, adjust  
2408   stereochemistry, perform local energy optimization.2409 14:   **end if**2410 15:   **(E) Convergence check.**2411 16:   **if**  $|\hat{y}_{\Delta G}^{(t+1)} - \hat{y}_{\Delta G}^{(t)}| < \epsilon_{\Delta G}$  **or**  $\hat{y}_{\Delta G}^{(t+1)} \leq \Delta G_{\text{target}}$  **then**2412 17:    **break**2413 18:   **end if**2414 19: **end for**2415 20: Return  $S^* = S^{(t+1)}$ .

2416

2417

2418

2419

2420 G.5 OVERALL PIPELINE  
2421

2422

2423

2424 The algorithm integrates stages—from surface modeling to inversion—via pocket initialization, en-  
2425 coder/head/updates, and beam search filtering for scalable, equivariant design. It begins with con-  
2426 structing receptor surface point clouds and patches, followed by pre-training an encoder using VQ-  
2427 MAE to capture structural features. The encoder and docking heads are then fine-tuned for task-  
2428 specific predictions. Starting with an initial ligand seed, the algorithm iteratively updates ligand  
2429 coordinates and features using gradients from fine-tuned heads, guided by an equivariant attention  
mechanism. A beam search filters candidates, ensuring scalability and convergence to an optimal  
ligand with high binding affinity.

```

2430 Algorithm 5: MagicDock pipeline
2431
2432 Require: Target receptor raw structure.
2433 Ensure: Generated ligand  $S^*$ .
2434 1: Stage 1: construct surface point clouds  $(X_s, N, F)$  and patches  $(X_p, F_p)$  for receptor (Alg. 1).
2435 2: Stage 2: pre-train encoder with VQ-MAE on patches (Alg. 2).
2436 3: Stage 3: fine-tune encoder and learn docking heads (Alg. 3).
2437 4: Initialize ligand seed  $S^{(0)}$ .
2438 5: for  $t = 0 \dots T_{\text{gen}} - 1$  do
2439   6: Stage 1 (partial): construct updated surface point clouds and features for current ligand  $S^{(t)}$ .
2440   7: Encode receptor and updated ligand via pre-trained and fine-tuned model.
2441   8: Stage 4: perform inversion-based generation using gradients from Stage 3 heads to
2442     iteratively update  $(x, f)$  and decode to discrete  $S$ .
2443   9: if convergence criteria met (affinity threshold) then
2444     10:   break
2445   11: end if
2446 12: end for
2447 13: Return generated ligand  $S^* = S^{(t)}$ .

```

## H ADAPTATIONS TO BASELINES FOR FAIR COMPARISON

To ensure a fair and rigorous evaluation, we adapted the DiffAb model (Luo et al., 2022)—a diffusion-based generative approach for antigen-specific antibody design—to align with the zero-start and de novo scenarios addressed by MagicDock. Originally, DiffAb assumes known backbone information and focuses on designing or optimizing specific complementarity-determining regions (CDRs), such as CDR-H3. This setup provides additional structural priors (e.g., framework residues and epitope details), which are not available in truly zero-start de novo design or when optimizing unknown antibody components. By modifying DiffAb to operate under reduced information and broader design scopes, we enable a direct comparison on equal footing, emphasizing the challenges of designing from scratch without relying on pre-existing structural templates.

We implemented three variants of DiffAb:

1. **Original Version:** Utilizes known backbone and epitope information, targeting CDR-H3 design (as per the original setup).
2. **Modified Version 1:** Removes information about the antibody to be optimized (e.g., no prior knowledge of other CDRs or framework), targeting CDR-H3 for optimization.
3. **Modified Version 2:** Retains partial information about the antibody to be optimized but extends the target to all CDRs (de novo design across the full variable region).

These adaptations simulate progressively more challenging conditions, transitioning from region-specific refinement to full de novo generation, mirroring MagicDock’s end-to-end paradigm. Table 9 summarizes the IMP scores for these variants alongside MagicDock. The declining IMP scores in modified versions highlight the increased difficulty without structural priors, underscoring the need for fair baselines in zero-start evaluations.

Table 9: IMP Scores for DiffAb Variants and MagicDock

Method	Known Information	Target Region	IMP Score (%)
DiffAb (Original)	Backbone + Epitope	CDR-H3	38.80
DiffAb (Modified 1)	Epitope Only	CDR-H3	31.67
DiffAb (Modified 2)	Epitope Only	De Novo Design	17.65
MagicDock	Receptor Only	Full Ligand (De Novo)	36.32

To maintain consistency and equity across all comparisons, we applied similar adaptive modifications to other baselines. These adaptations ensure that all models are evaluated under comparable information constraints, preventing inflated performance from auxiliary priors and providing a balanced assessment of their capabilities in realistic docking-oriented de novo design tasks.

2484 **I POTENTIAL SOCIETAL IMPACTS**  
24852486 Our work on docking-oriented de novo ligand design can be used in developing potent therapeutic  
2487 ligands and accelerate the research process of drug discovery. The generality of our method extends  
2488 beyond its current application; it is adaptable for various computer-aided design scenarios including,  
2489 but not limited to, small molecule, protein, and biomaterial design. It is also needed to ensure the  
2490 responsible use of our method and refrain from using it for harmful purposes.  
24912492 **J GENAI USAGE DISCLOSURE**  
24932494 In the preparation of this manuscript, we have utilized generative artificial intelligence (GenAI)  
2495 tools, specifically GPT-4o and Grok-4, to assist with text polishing and refinement, as well as to  
2496 support the drafting and modification of code snippets. These tools have been employed to enhance  
2497 the clarity and readability of the narrative and to facilitate the development of auxiliary code, en-  
2498 suring a streamlined presentation of our work. However, we emphasize that GenAI was not utilized  
2499 in the derivation of mathematical formulas, the design or implementation of key algorithms, or the  
2500 formulation of core scientific insights. All critical theoretical proofs, algorithmic developments, and  
2501 experimental validations were conducted independently by the authors to maintain the integrity and  
2502 originality of the research. We have rigorously reviewed and verified all generated text to ensure  
2503 accuracy and alignment with the scientific content, thereby upholding the reliability of the presented  
2504 results.  
25052506 **K VISUALIZATION OF GENERATED LIGAND**  
25072508 **K.1 PROTEIN LIGAND**2509 The protein-ligand case study and visualization is discussed in Sec. 5.2 and Fig. 4, where Magic-  
2510 Dock generates de novo protein ligand for a target receptor pocket. We evaluate binding affinity,  
2511 pose accuracy, and structural validity using metrics like docking scores and RMSD, demonstrating  
2512 superior performance over baselines in high-throughput screening simulations.  
25132514 The protein selected for this case study is Integrin beta-4 (ITGB4, PDB: 3F7P), a key component of  
2515 hemidesmosomes that anchors epithelial cells to the basement membrane through interactions with  
2516 the cytolinker protein plectin. The 3F7P structure specifically captures a fragment of ITGB4’s cyto-  
2517 plasmic tail containing fibronectin type III (FNIII) domains in complex with plectin’s actin-binding  
2518 domain, highlighting the molecular interface essential for stable adhesion. Mutations in ITGB4 that  
2519 disrupt this interaction are linked to forms of epidermolysis bullosa. In this study, MagicDock lever-  
2520 ages surface point cloud representations to design ligands targeting the FNIII domains of ITGB4.  
25212522 **K.2 MOLECULE LIGAND**2523 Fig. 11 illustrates the 1A99 protein structure, representing the Putrescine Receptor (PotF) from Es-  
2524 cherichia coli, bound to a small-molecule ligand. The protein is rendered as a ribbon model, with  
2525 structure details such as alpha-helices and beta-sheets depicted in light blue and gray tones, high-  
2526 lighting its intricate folded architecture. The central ligand, illustrated in magenta with blue and  
2527 red atomic features, represents putrescine (1,4-diaminobutane), a polyamine critical for bacterial  
2528 transport systems. The red Pocket regions, marked with crosses and plus signs, indicate hydrogen  
2529 bonding sites or key interaction points, stabilizing the protein-ligand complex. The 1A99 protein,  
2530 PotF, functions as a periplasmic binding protein, selectively capturing putrescine from the environ-  
2531 ment and delivering it to the membrane-bound transporter complex, essential for bacterial growth  
2532 and DNA stabilization. Putrescine, the ligand, plays a vital role in modulating gene expression and  
2533 cell signaling, supporting bacterial survival under stress. The binding nature between PotF and pu-  
2534 tresscine is characterized by high affinity, driven by electrostatic interactions and hydrogen bonds  
2535 within the Pocket, where the ligand’s positively charged amino groups align with negatively charged  
2536 residues, enhancing specificity and stability. Web resources, including PDBBind v2020 data, con-  
2537 firm this interaction’s importance in microbial physiology, offering insights into potential antimicro-  
2538 bial targets. This visualization thus bridges computational modeling with biological function, aiding  
2539 in drug design and molecular studies.  
2540

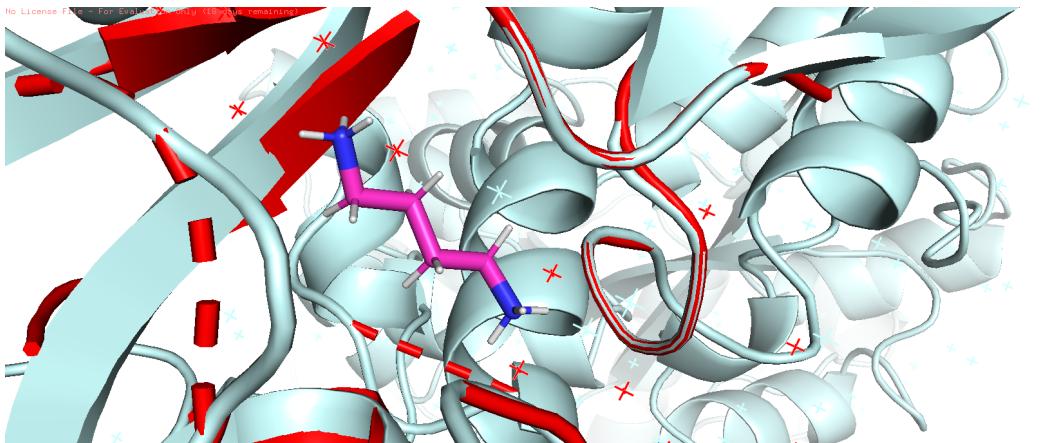


Figure 11: Visualization of an example of generated protein-ligand complexes.

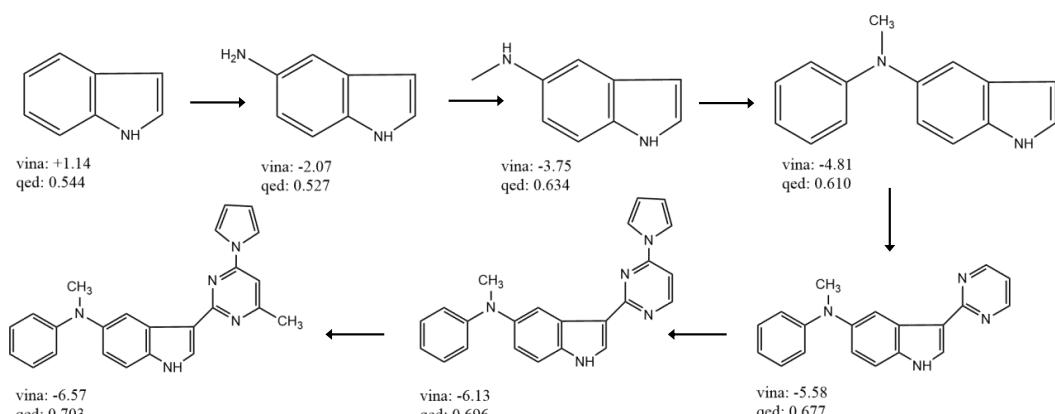
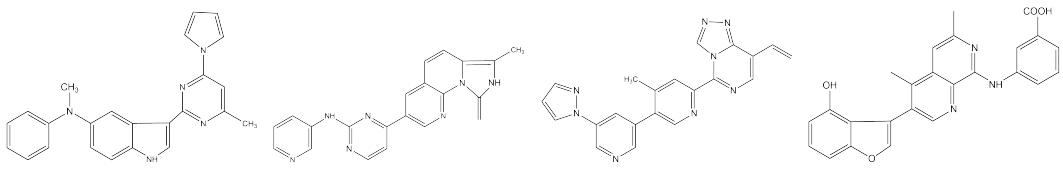


Figure 12: Illustration of a chemical synthesis process for bis-benzimidazole derivatives, showing stepwise transformations from benzimidazole and heterocyclic precursors to the final DNA-binding agent.

The chemical synthesis process depicted in Fig. 12 involves a stepwise transformation of organic molecules. It begins with a benzimidazole derivative reacting with an amine  $\text{NH}_2$ , followed by the introduction of a methyl group  $\text{CH}_3$  to form an intermediate. This intermediate undergoes further reaction to yield a symmetric bis-benzimidazole compound. Concurrently, a related process starts with a complex heterocyclic molecule with multiple methyl groups, which is simplified through a series of reactions to produce another benzimidazole-based structure.

The four small molecules shown in Fig. 13 are protein-binding ligands with unique structures facilitating hydrogen bonding,  $\pi$ - $\pi$  stacking, hydrophobic contacts, and metal coordination. Molecule 1 features a benzimidazole core with a methyl and carboxamide group, enabling enzyme inhibition like tubulin polymerases via nitrogen bonding and aromatic stacking. Molecule 2 has a purine-like tricyclic system with an acetamide side chain, acting as a kinase inhibitor by binding to ATP pockets with nitrogen hydrogen bonds and methyl van der Waals contacts. Molecule 3 includes a dihydropyrimidine ring with a mercury group, used in crystallography for covalent cysteine binding. Molecule 4 combines an indole-like structure with hydroxyl and carboxyl groups, targeting serine

2592 proteases or transporters through polar and hydrophobic interactions, modulating enzyme activity  
 2593 with its amphipathic nature.  
 2594



2600 (a) benzimidazole core with methyl and carbox-  
 2601 amide (b) purine-like tricyclic system with an acetamide side chain  
 2602 (c) structure of a dihy-  
 2603 dropyrimidine ring with a mercury group  
 2604 (d) indole-like structure with hydroxyl and car-  
 2605 boxyl

2604 Figure 13: Comparison of four generated small-molecule ligands, each with unique structures for  
 2605 protein binding.

## L HYPERPARAMETERS IN DETAILS

2609 For the baselines, we adopt the hyperparameters and training procedure in their official releases. We  
 2610 list the values of these hyperparameters as well as those of our MagicDock in Table 10, Table 12,  
 2611 and Table 11.

2612  
 2613  
 2614  
 2615  
 2616  
 2617  
 2618  
 2619  
 2620  
 2621  
 2622  
 2623  
 2624  
 2625  
 2626  
 2627  
 2628  
 2629  
 2630  
 2631  
 2632  
 2633  
 2634  
 2635  
 2636  
 2637  
 2638  
 2639  
 2640  
 2641  
 2642  
 2643  
 2644  
 2645

2646 L.1 PROTEIN BASELINES  
26472648 Table 10 enumerates critical hyperparameters for existing baselines, including learning rates and  
2649 batch sizes, facilitating fair comparisons in protein-ligand docking performance evaluation.  
26502651 Table 10: Hyperparameters for protein baselines.  
2652

Hyperparameter	Value	Description
<b>RAbD</b>		
N outer cycles	3	Number of outer Monte Carlo cycles.
N inner cycles	10	Number of inner Monte Carlo cycles.
Packing shell distance	6 Å	Distance for creating a shell around the CDR for optimization.
Interface distance	8 Å	Threshold for defining interface residues during docking.
Docking outer cycles	3	Shortened high-resolution docking outer cycles.
Docking inner cycles	10	Shortened high-resolution docking inner cycles.
<b>DiffAB</b>		
hidden size	128	Size of hidden states in the model.
pair size	64	Size of residue-pair features.
n layers	6	Number of layers in the MPN.
n steps	100	Number of the diffusion steps.
<b>HSRN</b>		
Hidden dimension	128	Size of hidden states in each MPN.
Number of layers (docking)	4	Layers in hierarchical encoder for docking.
Number of layers (generation)	3	Layers in hierarchical encoder for generation.
RBF interval (hydropathy)	0.1	Interval for hydropathy features.
RBF interval (volume)	10	Interval for volume features.
Epochs (docking)	20	Training epochs for docking.
Epochs (generation)	10	Training epochs for generation.
Dropout	10%	Dropout rate.
Optimizer	Adam	Optimizer used for training.
<b>dyMEAN</b>		
embed size	64	Size of the residue type & position number embedding.
hidden size	128	Size of the hidden states in the MPN.
n layers	3	Number of layers in the MPN.
n iter	3	Number of iterations in the progressive full-shot decoding.
k neighbors	9	Number of neighbors for each node in the KNN graph.
d	16	Size of the attribute vector of each channel.
<b>ABDPO</b>		
Hidden state size	128	Size of hidden states in MLPs.
Number of layers	6	Layers for features processing MLPs.
Diffusion steps	100	Number of diffusion steps.
Batch size (pre-training)	16	Batch size during pre-training.
Batch size (fine-tuning)	48	Batch size during fine-tuning.
Learning rate (pre-training)	$10^{-4}$	Initial learning rate for pre-training.
Learning rate (fine-tuning)	$10^{-5}$	Initial learning rate for fine-tuning.
Optimizer	Adam	Optimizer used.
Adam betas	(0.9, 0.999)	Betas for Adam optimizer.
Clip gradient norm	100	Gradient clipping norm.
$\beta$	0.01 / 0.005	Values for preference optimization.
Energy weights	8:8:2	Res CDR $E_{\text{total}}$ , Res CDR-Ag $E_{\text{nonRep}}$ , and Res CDR-Ag $E_{\text{Rep}}$ .
<b>Abx</b>		
batch size	1	Batch size used in inference and design.
num samples	100	Number of samples generated.
Learning rate	$10^{-4}$	Learning rate for Adam optimizer.
Number of sampling steps	100	Number of steps in the diffusion sampling process.
Adam betas	(0.9, 0.999)	Beta values for Adam optimizer.
Hidden dimension $D_h$	256	Dimension of node embeddings in the neural network.
Number of layers $L$	4	Number of layers in the FrameDiff neural network.

2696  
2697  
2698  
2699

2700 L.2 MOLECULE BASELINES  
27012702 Table 11 details key hyperparameters for molecular baselines, such as diffusion steps and noise  
2703 schedules, enabling standardized benchmarking of generative models in ligand design tasks.  
27042705 Table 11: Hyperparameters for the small molecule baselines.  
2706

Hyperparameter	Value	Description
<b>DockStream</b>		
Docking Poses (AutoDock Vina)	2	Number of poses returned by AutoDock Vina.
Grid Box Size (AutoDock Vina)	$15 \times 15 \times 15 \text{ \AA}$	Grid box dimensions for AutoDock Vina.
pH Settings	$7.0 \pm 2.0$	Target pH and tolerated range for RDKit-based optimization.
Force Field	UFF	Force field for RDKit-based optimization.
Maximum Iterations	600	Max iterations for RDKit with TautEnum.
<b>3D-SBDD</b>		
Exhaustiveness (AutoDock Vina)	8	Exhaustiveness parameter for docking.
Max Binding Modes	9	Maximum number of binding modes.
Energy Range	3	Energy range for docking.
Search Box Padding	$12.5 \text{ \AA}$	Padding for search box coordinates.
Min Box Length	$30 \text{ \AA}$	Minimum search box length.
Batch Size	8	Batch size mentioned in related models.
Initial Learning Rate	$10^{-4}$	Initial learning rate in related models.
Cross-Validation Folds	5	Folds for hyperparameter selection.
Training Epochs	10	Epochs for graph neural networks.
Budget Evaluations	5000	Budget for de novo design evaluations.
Population Size (GA)	250	Population size for genetic algorithm.
Offspring Size (GA)	25	Offspring size for genetic algorithm.
Mutation Rate (GA)	0.01	Mutation rate for genetic algorithm.
<b>LiGAN</b>		
Batch Size	32	Batch size for training.
Learning Rate	$10^{-4}$	Learning rate for Adam optimizer.
Epochs	100	Number of training epochs.
Latent Dimension	256	Dimension of latent space in GAN.
Discriminator Layers	3	Number of layers in discriminator.
Generator Layers	4	Number of layers in generator.
<b>ALIDIFF</b>		
Batch Size	4	Batch size during pretraining.
Learning Rate (Pretraining)	0.001	Learning rate for Adam optimizer in pretraining.
Learning Rate (Fine-tuning)	$5 \times 10^{-6}$	Initial learning rate for fine-tuning.
Adam Betas	(0.95, 0.999)	Beta values for Adam optimizer.
Gradient Norm Clipping	8	Gradient norm during pretraining.
Atom Type Loss Scaling	100	Scaling factor for atom type loss.
Gaussian Noise Std	0.1	Standard deviation for data augmentation.
$\beta$	5	Beta value for fine-tuning.
<b>DRUGFLOW</b>		
Virtual Nodes $N_{\text{max}}$	10	Maximum virtual nodes, remove 5 on average.
Sampling Steps	500	Number of sampling steps.
Training Epochs	600	Epochs for DrugFlow and DrugFlow-OOD.
$\beta$ (PA)	100	Beta for preference alignment.
$\lambda_{\text{coord}}, \lambda_{\text{atom}}, \lambda_{\text{bond}}$	1, 0.5, 0.5	Weight for coordinate, atom, bond loss.
$\lambda_w, \lambda_l$	1, 0.2	Weight for $w, l$ loss.
Regularization $\lambda$ (OOD)	10	Regularization for uncertainty estimation.
Scheduler $k$ (FlexFlow)	3	Polynomial scheduler exponent.
<b>DIFFSBDD</b>		
Batch Size	32	Batch size for training.
Learning Rate	$10^{-4}$	Learning rate for Adam optimizer.
Diffusion Steps	1000	Number of diffusion steps.
Number of Layers	6	Number of layers in the model.
Hidden Size	128	Hidden state size.
Training Epochs	100	Number of training epochs.

2750  
2751  
2752  
2753

2754 L.3 MAGICDOCK  
27552756 Table 12 outlines MagicDock’s core hyperparameters, like inversion iterations and loss weights,  
2757 supporting reproducible optimization for docking-oriented de novo ligand generation.  
27582759 Table 12: Hyperparameters for MagicDock.  
2760

Hyperparameter	Value	Description
<b>MagicDock</b>		
Hidden units in MLP	128	Hidden units in two-layer MLP for pseudo-curvature prediction.
Batch size (fine-tuning)	16	Batch size during supervised fine-tuning.
Learning rate (fine-tuning)	$10^{-4}$	Learning rate for Adam optimizer in fine-tuning.
Weight decay	$10^{-5}$	Weight decay for Adam optimizer in fine-tuning.
Optimizer	Adam	Optimizer used for training.
Patch size $K$	32	Patch size for point cloud patches.
Max spherical order $L$	2	Max spherical order for SE(3) convolutions.
Codebook size $N_B$	512	Codebook size for vector quantization.
Gumbel temperature $\tau$	1.0	Temperature for Gumbel-softmax in pre-training.
Loss weights $\alpha, \beta$	5.0, 50.0	Weights for pocket and interaction losses in composite objective.
Loss weight $\lambda_p$	1.0	Weight for delta-G loss in fine-tuning.
SDF-GD steps $T_{\text{sdf}}$	50	Number of gradient descent steps in SDF projection.
Learning rate $\alpha_{\text{sdf}}$	0.1	Learning rate for SDF gradient descent.
Inversion iterations $T$	200	Number of iterative refinement steps in inversion.
Step size $\eta_t$	0.01	Step size for gradient updates in inversion stage.
Pre-training epochs	100	Number of epochs for pre-training stage.
Fine-tuning epochs	20	Number of epochs for fine-tuning stage.

2777  
2778  
2779  
2780  
2781  
2782  
2783  
2784  
2785  
2786  
2787  
2788  
2789  
2790  
2791  
2792  
2793  
2794  
2795  
2796  
2797  
2798  
2799  
2800  
2801  
2802  
2803  
2804  
2805  
2806  
2807