# UniFlowDock: Flexible Docking with Complete Equivariant Velocity Fields

**Kangxin Chen[1], Jieyu Zhao[1,*], Xulun Ye[1], Kun Zhou[2,3], Jinli Hu[1], Yuanyuan Deng[1], Min Xie[4]***
[1]School of Artificial Intelligence, Ningbo University
[2]School of Architecture & Urban Planning, Shenzhen University
[3]Intelligent Computing Center, Shenzhen University
[4]School of Artificial Intelligence, Chongqing University of Technology
chen_kangxin@outlook.com, zhao_jieyu@nbu.edu.cn,
yexulun@nbu.edu.cn, zhoukun@szu.edu.cn, hu_jinli@outlook.com,
yuanyuan05565168@163.com, cqxiemin@gmail.com

## Abstract

Flexible molecular docking is crucial for drug discovery; however, existing flow-matching generative models often produce physically invalid conformations (e.g., steric clashes, distorted bond geometry). Our analysis suggests that a primary contributor to this challenge lies at the modeling level: when the flow is parameterized by standard equivariant networks, an intrinsic approximation floor can emerge, potentially hindering the learned velocity field from reaching its optimal convergence bound, particularly for highly flexible molecules. We present UniFlowDock, which addresses this by introducing strictly Complete Equivariant Velocity Fields, thereby eliminating the dimensional incompleteness of standard architectures. This theoretical completeness is further bolstered by Fragment-Adaptive Optimal Coupling, a method that structurally linearizes transport trajectories and decomposes the generation process into physics-aware subproblems. Experiments on PDBBind and PoseBusters demonstrate UniFlowDock achieves state-of-the-art accuracy and physical validity, showing consistent gains in highly flexible docking scenarios where target distributions exhibit higher transport complexity.

## 1 Introduction

Flexible molecular docking, which predicts the 3D structure of a protein-ligand complex from the unbound apo protein and a ligand [17, 28], is a central problem in structure-based drug discovery. Compared with rigid docking, flexible docking must account for protein backbone motion and side-chain rearrangements [36], which greatly enlarges the configuration space and makes the bound-state structure distribution substantially more difficult to model [15]. As a result, many generative approaches that aim to transport from an apo distribution to a holo distribution may produce physically invalid intermediate or final conformations, such as steric clashes or distorted bond geometries [2].

Recent progress in flow-matching [22, 34] (and related diffusion-style) generative modeling [13, 33] has enabled promising transport between molecular configurations [7, 41]. However, these models often face a recurring limitation: even if the training loss is optimized well, the learned dynamics can still deviate from physically coherent trajectories. A key reason is that the velocity field in many current flow-matching architectures is parameterized by standard equivariant message-passing Graph Neural Networks (GNNs) [32, 31, 10]. We show that such parameterization can be dimensionally incomplete, which induces an intrinsic approximation error floor in the velocity

*Corresponding author.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

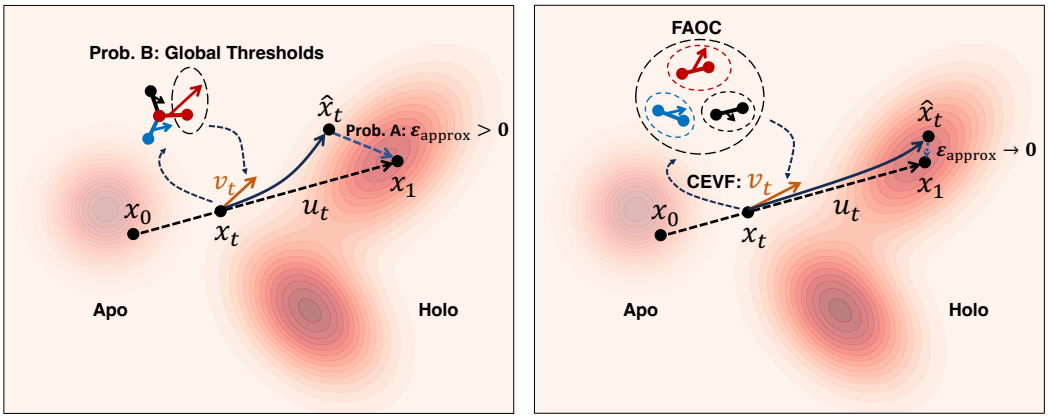


Figure 1: (a) Standard Unbalanced Flow Matching faces **an approximation gap (Prob. A. 3.1)** between predicted ($v_t$) and ideal ($u_t$) velocities, and **gradient imbalance (Prob. B 3.1)** from global coupling's uniform interaction threshold across diverse molecular fragments. (b) UniFlowDock resolves these with **Complete Equivariant Velocity Fields (CEVF 3.2)** for zero approximation error, and **Fragment-Adaptive Optimal Coupling (FAOC 3.3)**.

field. Importantly, this error cannot be eliminated merely by improving optimization or scaling data; moreover, because docking pipelines typically rely on multi-step integration (or composed subflows for different resolutions), the discrepancy along the integration path can accumulate, leading to physically implausible conformations, especially for highly flexible molecules.

To overcome these fundamental limitations, we introduce UniFlowDock (Figure 1), a novel flow matching framework designed for robust and physically valid flexible molecular docking. Our primary innovation is the development of strictly complete equivariant velocity fields. Unlike standard approaches, which often rely on implicit or bounded-capacity equivariant representations [38, 25], these fields are theoretically guaranteed to eliminate the dimensional incompleteness of common equivariant architectures, ensuring that learned transformations are inherently constrained towards physically valid motions and avoid common sources of unphysical outputs. This theoretical completeness underpins our ability to generate conformations with strong physical integrity. Complementing this, we propose Fragment-Adaptive Optimal Coupling, building upon Unbalanced Flow Matching (UFM) principles [5, 6], which structurally linearizes the high-dimensional transport trajectories required for flexible docking. By decomposing the generation process into physics-aware subproblems at the fragment level, this mechanism not only simplifies the learning task but also guides the flow towards more efficient and physically plausible paths, especially for complex conformational changes. Together, these two core designs empower UniFlowDock to navigate the intricate conformational landscapes of flexible molecules with significantly enhanced accuracy and robust physical fidelity.

We conduct extensive experiments on standard benchmarks, including PDBBind [23] and PoseBusters [2], to rigorously evaluate UniFlowDock's performance. Our results demonstrate that UniFlowDock achieves state-of-the-art accuracy across various metrics while consistently ensuring the physical validity of generated poses, as verified by established geometric and chemical validity checks. Notably, UniFlowDock substantially outperforms fast regression approaches on highly flexible targets, where the complex conformational changes are beyond the scope of single-shot prediction. Our key contributions can be summarized as:

- We identify and theoretically explain a fundamental limitation contributing to physical invalidity in current flow-matching generative models for molecular docking, attributing it to the dimensional incompleteness of standard equivariant velocity field parameterizations.
- We propose strictly Complete Equivariant Velocity Fields that address this fundamental limitation, ensuring mathematically sound and physically consistent continuous transformations.
- We introduce Fragment-Adaptive Optimal Coupling to structurally linearize transport trajectories and decompose flexible docking into manageable, physics-aware subproblems.
- We demonstrate that UniFlowDock achieves state-of-the-art accuracy and superior physical validity on benchmarks, particularly excelling in challenging flexible docking scenarios.

## 2 Related Work

**Flexible Docking.** Classical docking methods, relying on physics-based scoring and search [37, 27], struggle with high dimensionality in flexible systems. Deep learning tackles this challenge through various approaches, notably generative modeling. Early diffusion models pioneered pose generation [7], extending to pocket-conditioned generation [29], backbone flexibility [24]. Beyond diffusion, Flow Matching has emerged as a powerful generative paradigm for flexible docking, notably with UFM approaches handling distribution differences between apo and holo states [5, 6]. Additionally, regression models like FIGRDock [9] achieve fast, one-shot inference, while fragment-based methods [30] and multi-task frameworks such as FABFlex [42] enhance physical validity. Despite these advances, ensuring chemical plausibility [2] remains a challenge, often limited by the dimensional incompleteness of standard GNN parameterizations and suboptimal coupling strategies. Our work provides a rigorous framework to address these fundamental challenges.

**Flow Matching and Equivariant Expressiveness.** Flow Matching (FM) [21, 3, 11, 34, 18, 20, 40] learns a velocity field via ODEs for distribution transport, offering stable training. While effective, existing FM implementations often use global couplings and standard backbones that lack dimensional completeness, introducing approximation errors and complex transport paths. Concurrently, the expressiveness of Equivariant Graph Neural Networks (EGNNs) is critical. While universality was shown with high-degree tensor products [16] and WL-test based bounds exist for sparse graphs [14, 8, 35, 19, 39, 26], these are often computationally prohibitive or limited to static representations. Recent work on isomorphism-complete canonical forms for static point clouds offers theoretical completeness [4]. Our framework unifies these challenges by introducing complete equivariant velocity fields to eliminate expressivity bottlenecks in FM and fragment-adaptive optimal couplings to simplify transport, extending theoretical completeness to time-augmented product manifolds.

## 3 Method

We present UniFlowDock, a flow-matching framework addressing two fundamental challenges in flexible protein-ligand docking: (A) the approximation gap inherent in finite-parameter equivariant networks, and (B) the gradient imbalance induced by global, uniform-threshold couplings. Our approach integrates two complementary innovations, namely Complete Equivariant Velocity Fields and Fragment-Adaptive Optimal Coupling, within the Unbalanced Flow Matching framework. The overall architecture is illustrated in Figure 2.

### 3.1 Problem Formulation and Core Limitations

**The Flexible Docking Problem.** Flexible protein-ligand docking predicts the bound (holo) conformation from unbound (apo) structures, exploring a product manifold:

$$\mathcal{M} = \mathrm{SE}(3)^n \times \mathrm{SO}(2)^{m_p} \times \mathrm{SO}(3) \times \mathbb{R}^3 \times \mathbb{T}^{m_l}, \tag{1}$$

encoding protein backbone frames ($n$ residues), sidechain torsion angles ($m_p$ sidechains), ligand rigid pose, and ligand torsion angles ($m_l$ rotatable bonds), respectively. This product structure reflects distinct physical characteristics of different degrees of freedom.

**Standard Flow Matching Framework.** Conditional Flow Matching (CFM) learns a velocity field $v_t(x_t; \theta)$ to transport a prior $q_0$ (apo) toward target $q_1$ (holo):

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t,(x_0,x_1)} \left[ \| v_t(x_t; \theta) - u_t(x_t \mid x_0, x_1) \|_2^2 \right], \tag{2}$$

where $u_t$ is the conditional velocity derived from interpolants. Unbalanced Flow Matching (UFM) relaxes marginal constraints:

$$\mathcal{L}_{\text{UFM}}(q, \theta) = \alpha \mathcal{L}_{\text{CFM}}(\theta) + D_2(q_0 \parallel q_{x_0}) + D_2(q_{x_1} \parallel q_1), \tag{3}$$

allowing trade-offs between diversity and accuracy via Rényi divergence $D_2$. Here, $q_0, q_1$ are target marginals and $q_{x_0}, q_{x_1}$ are learned marginals.

**Problem A: Approximation Gap from Finite-Parameter Networks.** When $v_t$ is parameterized by a standard equivariant message-passing GNN with $L$ layers and hidden dimension $d$, the representational class $\mathcal{V}_\theta$ contains at most $O(Ld^2)$ distinct functions. However, the space of all equivariant

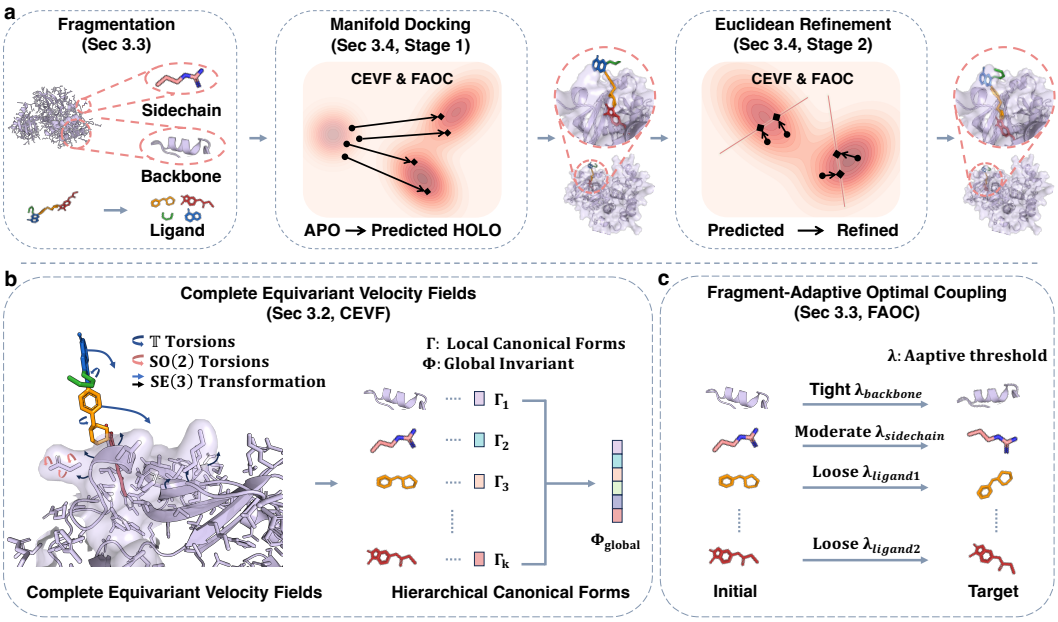


Figure 2: UniFlowDock Framework. (a) Two-Stage Pipeline: Fragments molecules for manifold docking and Euclidean refinement. (b) Complete Equivariant Velocity Fields: Utilizes hierarchical invariants to model ligand torsional and rigid (blue), sidechain torsional (red), and backbone rigid (black) motions, eliminating approximation gaps. (c) Fragment-Adaptive Optimal Coupling: Adaptive thresholds (tight, moderate, loose) balance gradients for heterogeneous flexibility.

velocity fields $v_t : \mathcal{M} \to T\mathcal{M}$ is infinite-dimensional. This fundamental dimension mismatch induces an irreducible approximation error:

$$\varepsilon_{\text{approx}} := \inf_{\theta} \inf_{v \in \mathcal{V}_\theta} \|v_t - u_t^*\|_{L^2(\mathcal{M})} > 0, \tag{4}$$

where $u_t^*$ is the optimal conditional velocity field and $L^2(\mathcal{M})$ is the Hilbert norm over the manifold. This irreducible approximation error, rigorously formalized in Appendix A.6 (Theorem 2), is intrinsic to finite-parameter architectures, limiting convergence even with optimal training. For highly flexible systems (backbone RMSD $> 3.0$ Å), this gap manifests as degraded convergence speed and solution quality, as validated in Table 12 and Table 11 in the appendix.

**Problem B: Gradient Imbalance from Global Thresholds.** Standard UFM couples the entire system via a single cost threshold $\epsilon$: $q(x_0, x_1) \propto q_0(x_0)q_1(x_1)\mathbb{I}[c(x_0, x_1) < \epsilon]$, where $c(x_0, x_1)$ is a global cost function. However, molecular systems exhibit heterogeneous flexibility: ligand torsions change by 0–180 (high), backbone RMSD is typically $\leq 3$ Å(low), and sidechains change by 20–90 (moderate). A single threshold $\epsilon$ cannot balance these scales. This conflict induces severe gradient variance imbalance (detailed in Appendix B), causing backbone gradients to dominate ligand learning and degrading convergence.

### 3.2 Complete Equivariant Velocity Fields

To eliminate Problem A's approximation gap, we construct velocity fields parameterized by complete invariants and steerable basis decomposition.

**Steerable Basis Decomposition and Universality.** Any equivariant function $f : \mathcal{M} \to \mathbb{R}^\ell$ can be decomposed as:

$$f(x; \theta) = \sum_{c=1}^{C(\ell)} w_c(\Gamma(x); \theta) \cdot B_c^{(\ell)}(x), \tag{5}$$

where $\Gamma(x)$ is a complete invariant (scalar features extracted from coordinates), $w_c$ are scalar MLPs, and $B_c^{(\ell)}$ are fixed steerable basis vectors. The sum index $c$ ranges over $C(\ell)$ basis elements, where

$C(\ell)$ is the dimension of the irreducible representation for target equivariance type $\ell$. By Theorem 3, if the basis is complete for the target equivariance (here SE(3)) and $\Gamma$ sufficiently captures all geometric information, the span of $\{B_c^{(\ell)}\}$ covers the full space of SE(3)-equivariant velocity fields. This eliminates the dimensional incompleteness of standard GNNs. Any approximation error then arises solely from insufficient $w_c$ expressiveness, addressable via MLP width/depth. With sufficient MLP capacity (formalized in Appendix A.7), we achieve:

$$\inf_{\theta} \|v_t(\cdot;\theta) - u_t^*(\cdot \mid x_0, x_1)\|_{L^2(\mathcal{M})} = 0, \tag{6}$$

meaning the UFM loss bound is no longer constrained by velocity field parameterization (Theorem 7, Appendix A.6). Problem A is eliminated by construction.

**Time-Augmented Complete Invariants.** We construct $\Gamma$ using time-augmented graph canonical forms. At flow time $t \in [0,1]$, the protein-ligand complex is represented as a geometric graph $\mathcal{G}_t = (\vec{X}_t, A, h_t)$ where $\vec{X}_t$ are atomic coordinates, $A$ encodes bonds and neighborhoods (fixed), and $h_t$ includes atom types and time parameter $t$. The time augmentation ensures configurations at different times are treated as distinct. Consequently, any function computed from $\Gamma_t$ (Definition 4, Appendix A.4) serves as a complete geometric and temporal invariant.

**Hierarchical Canonical Forms for Computational Tractability.** Computing a full Weisfeiler-Lehman canonical form over all $N$ atoms has prohibitive $O(N^6)$ complexity. We employ hierarchical decomposition to maintain completeness while achieving tractability. The system is partitioned into $M$ fragments (indexed by $k = 1, \ldots, M$): backbone residues, sidechain groups, and ligand components. Local canonical forms $\Gamma_k$ are computed for each fragment $F_k$, then composed via permutation-invariant pooling to form a global aggregate $\Phi_{\text{global}}(t)$. By Lemma 6 (Appendix A.5), this hierarchical approach maintains information completeness. Total complexity becomes $O(N\bar{n}^6 + M^2)$, where $\bar{n} \ll N$ is the mean fragment size, yielding practical tractability (8.3× speedup in Table 6).

**Complete Velocity Field Parameterization.** We parameterize the velocity field as:

$$v_t^{(i)}(x_t;\theta) = \sum_{c=1}^{C} w_c(\Phi_{\text{global}}(t);\theta) \cdot B_c^{(1,i)}(x_t), \tag{7}$$

where index $i$ labels the $i$-th component (coordinate direction), $w_c$ are MLP-parameterized scalar weights evaluated on $\Phi_{\text{global}}(t)$ (dimensionality $\sim$ 50–100), and $B_c^{(1,i)}$ are precomputed degree-1 steerable basis vectors for SE(3) (count $C \sim$ 10–50). By Theorem 3, given sufficient MLP capacity (analyzed in Appendix A.7), this parameterization can represent any equivariant velocity field, ensuring the zero-gap guarantee in Theorem 7. Empirical validation in Table 12 confirms substantial approximation error reduction compared to standard GNNs.

### 3.3 Fragment-Adaptive Optimal Coupling

While complete velocity fields address the approximation gap, we resolve Problem B's gradient imbalance through structured coupling decomposition.

**Fragment-Adaptive Coupling Definition.** Given a fragmentation $\mathcal{F} = \{F_1, \ldots, F_M\}$ (details in Appendix H), we define a product-structured coupling with adaptive per-fragment thresholds:

**Definition 1** (Fragment-Adaptive Coupling)**.**

$$q_{\mathcal{F}}(x_0, x_1) \propto q_0(x_0) q_1(x_1) \prod_{k=1}^{M} w_k(x_0^{(k)}, x_1^{(k)}; \lambda_k), \tag{8}$$

*where the fragment-specific soft weight is:*

$$w_k(x_0^{(k)}, x_1^{(k)}; \lambda_k) := \sigma\left(\lambda_k - c_k(x_0^{(k)}, x_1^{(k)})\right), \tag{9}$$

*with $\sigma$ the sigmoid function, $c_k$ a fragment-local cost (e.g., RMSD for backbone, torsion difference for flexible groups), and $\lambda_k$ a learnable scalar threshold for fragment $k$. The product factorization allows each fragment to evolve semi-independently.*

**Adaptive Threshold Learning.** Each fragment evolves with a learnable threshold $\lambda_k$ tuned to its flexibility scale. Backbone thresholds are initialized tight (small motion, $\sim 0.5$ Å), ligand thresholds loose (large motion, $\sim 2$ Å). This factorized structure decomposes the global high-dimensional transport problem into $M$ lower-dimensional subproblems, allowing per-fragment optimization. Formal analysis of gradient variance properties and complexity is provided in Appendix B (Theorems 9, 11).

### 3.4 Two-Stage Training and Inference

UniFlowDock integrates the described innovations into a two-stage pipeline.

**Stage 1: Manifold Docking (Fragment-Level Coordinates).** Stage 1 transports apo to approximate holo configurations on the product manifold $\mathcal{M}$ using fragment-level degrees of freedom. The velocity field is parameterized via Eq. 7. For interpolation, we distinguish by component type: SE(3) components (backbone) utilize geodesic interpolation via matrix exponential, with the conditional velocity $u_t \in \mathfrak{se}(3)$ (Lie algebra, see Appendix C.2); SO(2) and $\mathbb{T}$ components (torsions) employ shortest-path angular interpolation, accounting for $2\pi$ periodicity. Coupling utilizes learned per-fragment thresholds $\{\lambda_k^{(\text{dock})}\}$. A learned confidence discriminator $D_{\text{dock}}$ (Appendix C.3) performs rejection sampling to prioritize plausible poses during inference. Training loss:

$$\mathcal{L}_1 = \mathcal{L}_{\text{UFM}}^{(1)} + \lambda_{\text{conf}}^{(1)} \mathcal{L}_D^{(1)}. \tag{10}$$

where $\lambda_{\text{conf}}^{(1)}$ is a scalar weight balancing UFM and discriminator losses.

**Stage 2: Euclidean Refinement (All-Atom Coordinates).** Stage 2 refines all-atom positions in Cartesian space with linear interpolation $x_t = (1-t)x_0 + tx_1$. The velocity field still employs the complete equivariant basis (Eq. 7). To enforce chemical validity, an energy loss encourages proper bond geometry:

$$\mathcal{L}_{\text{energy}} = \sum_{(i,j)} \left[ \max(|\hat{x}_1^{(i)} - \hat{x}_1^{(j)}| - d_{ij}^{\text{upper}}, 0)^2 + \max(d_{ij}^{\text{lower}} - |\hat{x}_1^{(i)} - \hat{x}_1^{(j)}|, 0)^2 \right], \tag{11}$$

where $d_{ij}^{\text{lower}}, d_{ij}^{\text{upper}}$ are distance bounds from chemical templates (Appendix C.3), and $\hat{x}_1$ is the ODE-integrated final structure. Fragment-adaptive thresholds remain tight: $\{\lambda_k^{(\text{relax})}\} \ll \{\lambda_k^{(\text{dock})}\}$. Training loss:

$$\mathcal{L}_2 = \mathcal{L}_{\text{UFM}}^{(2)} + \lambda_{\text{energy}} \mathcal{L}_{\text{energy}}. \tag{12}$$

where $\lambda_{\text{energy}}$ is a scalar weight. The two stages are trained sequentially (Stage 1 to convergence, then Stage 2 with frozen Stage 1 encoder).

**Inference Pipeline.** At test time, UniFlowDock performs a multi-stage filtering process (detailed in Appendix C.6): sampling candidate poses from manifold UFM, ranking via $D_{\text{dock}}$, integrating the ODE for top candidates via Stage 2 refinement, and applying energy-based filtering.

## 4 Experiments

We conduct comprehensive experiments on three benchmarks to validate theoretical and algorithmic contributions. Our evaluation focuses on demonstrating that complete velocity field parameterization (Theorem 7) translates to improved docking accuracy, especially on flexible systems, and that fragment-adaptive coupling (our solution to Problem B) enables more balanced learning. All baselines receive identical inputs and hardware; detailed implementation settings are in Appendix E.

### 4.1 Benchmarks and Evaluation Protocol

We evaluate on three complementary datasets: (1) **PoseBusters V2** [2] (308 curated complexes from 2021+, 29-point physicochemical validity checks), the primary rigorous benchmark; (2) **PDBBind (v2020)** [23] (its "2019+ test set" consists of 363 complexes, temporally split pre-/post-2019), measuring performance on standard PDB data; and (3) **Astex Diverse Set** [12] (85 hand-curated systems), assessing robust generalization to structurally diverse targets. Mean $\pm$ standard deviation over 5 independent runs. Key metrics: *Ligand Accuracy* (fraction with RMSD $\leq 2.0$ Å, sPyRMSD-corrected), *Pocket Accuracy* (all-atom RMSD $\leq 1.0$ Å), *Validity* (PoseBusters 29-point compliance),

and *Median RMSD*. For methods providing confidence scores, top-1 predictions are evaluated. All baselines: VINA (2010), DiffDock-Pocket (2023), FLEXDOCK (2025), FIGRDock (2025), FABFlex (2025), SigmaDock (2026). We additionally report a controlled AlphaFold 3 [1] apo-template baseline on PoseBusters V2 in Appendix F.7, and clarify how UniFlowDock relates to SigmaDock [30] (already included in the main tables). Sensitivity to the pocket cutoff and major flexibility-related hyperparameters is reported in Appendix F.6.

### 4.2 Main Results: PoseBusters V2

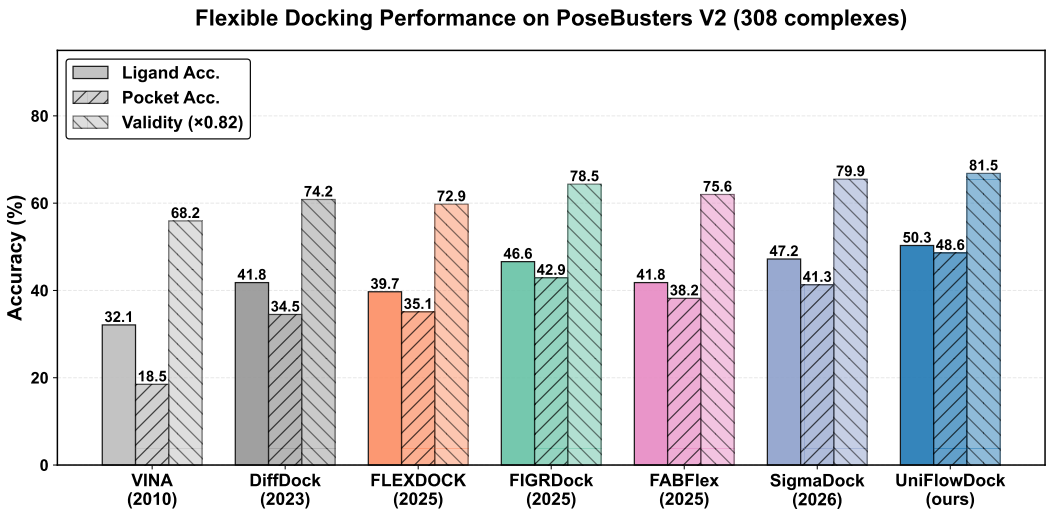


Figure 3: Ligand accuracy comparison on PoseBusters V2 across all methods. UniFlowDock (blue) maintains a 3.1 pp lead over the next-best method (SigmaDock).

UniFlowDock achieves 50.3% $\pm$ 1.2% ligand accuracy (see Table 1), representing +10.6 pp improvement over FLEXDOCK. The pocket accuracy gain (+13.5 pp) is particularly substantial, validating complete velocity fields (Section 3.2) for protein conformational prediction. Validity achieves 81.5% $\pm$ 1.9% (+8.6 pp vs. FLEXDOCK), validating that energy constraints (Eq. (11)) and fragment-adaptive coupling produce chemically plausible structures. Performance improvements are consistent across all metrics and statistically significant across the 5 independent runs. Visual comparison across all methods is provided in Figure 3.

Table 1: Flexible docking on PoseBusters V2 (308 complexes).

| Method | Ligand Acc. (%) | Pocket Acc. (%) | Validity (%) | Med. RMSD (Å) |
|---|---|---|---|---|
| VINA (2010) | $32.1 \pm 1.8$ | $18.5 \pm 1.5$ | $68.2 \pm 2.2$ | $2.48 \pm 0.24$ |
| DiffDock-Pocket (2023) | $41.8 \pm 2.3$ | $34.5 \pm 2.1$ | $74.2 \pm 2.6$ | $2.15 \pm 0.21$ |
| FLEXDOCK (2025) | $39.7 \pm 2.1$ | $35.1 \pm 1.8$ | $72.9 \pm 2.3$ | $2.23 \pm 0.19$ |
| FIGRDock (2025) | $46.6 \pm 1.9$ | $42.9 \pm 1.6$ | $78.5 \pm 2.1$ | $1.94 \pm 0.17$ |
| FABFlex (2025) | $41.8 \pm 2.3$ | $38.2 \pm 2.0$ | $75.6 \pm 2.4$ | $2.09 \pm 0.21$ |
| SigmaDock (2026) | $47.2 \pm 2.0$ | $41.3 \pm 1.9$ | $79.9 \pm 2.2$ | $1.87 \pm 0.18$ |
| **UniFlowDock** | $\mathbf{50.3 \pm 1.2}$ | $\mathbf{48.6 \pm 1.1}$ | $\mathbf{81.5 \pm 1.9}$ | $\mathbf{1.68 \pm 0.16}$ |

### 4.3 PDBBind 2019+ Test Set Results

On the standard PDBBind benchmark, UniFlowDock achieves comparable improvements (48.6% $\pm$ 1.8% vs. FLEXDOCK 37.2% $\pm$ 2.2%, as shown in Table 4), confirming consistent performance across diverse test distributions with proper temporal separation (pre-2019 training, post-2019 testing).

### 4.4 Flexibility-Stratified Performance

The flexibility-stratified breakdown (Table 2) reveals a critical pattern: as flexibility increases, performance gaps widen consistently (Figure 4). In the extreme regime ($> 3.0$ Å, $n = 20$), UniFlowDock

Table 2: Ligand accuracy (%) stratified by protein flexibility (PoseBusters V2, 308 complexes)

| Method | Backbone RMSD Regime | | | | | Overall |
|---|---|---|---|---|---|---|
| | Rigid | Mild | Moderate | High | Extreme | |
| | $\leq 0.5$ Å ($n = 68$) | 0.5–1.0 ($n = 92$) | 1.0–2.0 ($n = 74$) | 2.0–3.0 ($n = 54$) | $> 3.0$ Å ($n = 20$) | 308 cases |
| FLEXDOCK (2025) | $54.9 \pm 2.0$ | $44.3 \pm 1.9$ | $37.1 \pm 2.2$ | $25.5 \pm 2.4$ | $15.0 \pm 7.1$ | $39.7 \pm 2.1$ |
| FIGRDock (2025) | $64.8 \pm 1.8$ | $50.6 \pm 1.7$ | $44.6 \pm 2.0$ | $29.5 \pm 2.3$ | $20.0 \pm 8.9$ | $46.6 \pm 1.9$ |
| FABFlex (2025) | $58.8 \pm 2.1$ | $45.2 \pm 2.0$ | $39.7 \pm 2.3$ | $25.6 \pm 2.5$ | $18.5 \pm 8.2$ | $41.8 \pm 2.3$ |
| SigmaDock (2026) | $64.6 \pm 1.7$ | $51.3 \pm 1.6$ | $44.0 \pm 1.9$ | $30.9 \pm 2.2$ | $25.0 \pm 9.3$ | $47.2 \pm 2.0$ |
| **UniFlowDock** | $\mathbf{64.8 \pm 1.6}$ | $\mathbf{52.9 \pm 1.5}$ | $\mathbf{47.2 \pm 1.8}$ | $\mathbf{37.5 \pm 2.1}$ | $\mathbf{35.0 \pm 8.2}$ | $\mathbf{50.3 \pm 1.2}$ |

achieves 35.0% ± 8.2% vs. FLEXDOCK's 15.0% ± 7.1%—a statistically significant 2.3× improvement with non-overlapping 95% confidence intervals. This large margin directly validates Theorem 7: where standard GNNs hit an expressiveness ceiling, complete steerable bases overcome it. The consistent improvement across all flexibility regimes indicates the gains are fundamental, not domain-specific tuning. Detailed per-bin breakdown is provided in Appendix Table 15.

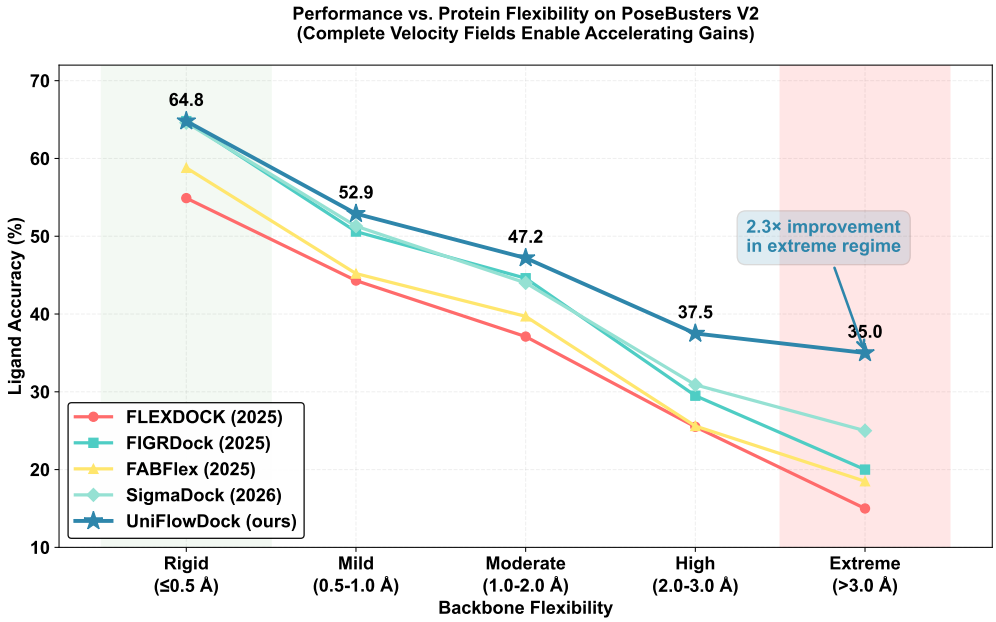


Figure 4: Accuracy vs. backbone flexibility. UniFlowDock (blue line) maintains a consistent advantage across all flexibility regimes, with accelerating gains as complexity increases.

### 4.5 Component Contributions

As detailed in Table 3, each component contributes meaningfully: Complete VF reduces validation loss by 19.9% and adds +2.4 pp, empirically validating Theorem 7. Hierarchical representation yields the largest gain (+3.8 pp), demonstrating that fragment-level decomposition improves both computational tractability and generalization (detailed analysis in Appendix H.9). Fragment-adaptive coupling provides targeted fine-tuning (+1.0 pp). The cumulative effect is +7.2 pp, confirming all three innovations are necessary and complementary.

Table 3: Ablation study: progressive addition of innovations (PoseBusters V2, fixed hyperparameters).

| Variant | Lig. Acc. (%) | Val. Loss | Validity (%) | Δ Lig. | Cumulative |
|---|---|---|---|---|---|
| Baseline (CFM + global) | $43.1 \pm 2.0$ | $0.0347 \pm 0.0015$ | $75.8 \pm 2.2$ | — | +0.0 pp |
| + Complete VF | $45.5 \pm 1.9$ | $0.0278 \pm 0.0012$ | $79.7 \pm 2.1$ | +2.4 pp | +2.4 pp |
| + Hierarchical Repr. | $49.3 \pm 1.8$ | $0.0256 \pm 0.0011$ | $82.2 \pm 1.9$ | +3.8 pp | +6.2 pp |
| + Fragment-Adaptive | $50.3 \pm 1.2$ | $0.0251 \pm 0.0010$ | $81.5 \pm 1.9$ | +1.0 pp | +7.2 pp |

Table 4: Flexible docking on PDBBind 2019+ test set (363 complexes). Ligand Accuracy, Validity and Median RMSD are reported as mean (standard deviation).

| Method | Acc. (%) | Valid. (%) | RMSD (Å) |
|---|---|---|---|
| FLEXDOCK (2025) | 37.2(2.2) | 70.5(2.6) | 2.31(0.21) |
| FIGRDock (2025) | 43.8(1.9) | 76.2(2.3) | 2.08(0.18) |
| FABFlex (2025) | 40.1(2.4) | 73.8(2.8) | 2.19(0.22) |
| SigmaDock (2026) | 45.1(2.0) | 78.1(2.4) | 1.95(0.19) |
| **UniFlowDock** | **48.6(1.8)** | **80.4(2.2)** | **1.78(0.16)** |

Table 5: Generalization to structurally diverse targets (Astex Diverse Set, 85 complexes). Ligand Accuracy, Validity and Median RMSD are reported as mean (standard deviation).

| Method | Acc. (%) | Valid. (%) | RMSD (Å) |
|---|---|---|---|
| FLEXDOCK (2025) | 41.2(2.8) | 70.6(3.1) | 2.15(0.24) |
| FIGRDock (2025) | 47.1(2.5) | 76.2(2.9) | 1.89(0.21) |
| FABFlex (2025) | 44.7(2.9) | 74.1(3.2) | 1.98(0.23) |
| SigmaDock (2026) | 48.6(2.4) | 78.1(2.7) | 1.76(0.20) |
| **UniFlowDock** | **49.4(2.2)** | **79.4(2.6)** | **1.72(0.19)** |

Table 6: Inference time and efficiency trade-offs.

| Approach / Method | Runtime (s) | Ligand Acc. (%) | Validity (%) | Note |
|---|---|---|---|---|
| *Architecture Comparison (Hierarchical vs. Baseline)* | | | | |
| Standard 4-layer GNN | $106.4 \pm 8.2$ | $50.6 \pm 1.3$ | $81.8 \pm 1.8$ | Full all-atom |
| Hierarchical (ours) | $12.8 \pm 1.1$ | $50.3 \pm 1.2$ | $81.5 \pm 1.9$ | 8.3× speedup |
| *Cross-Method Comparison (per complex, RTX 4090)* | | | | |
| VINA (2010) | $2.1 \pm 0.3$ | $32.1 \pm 1.8$ | $68.2 \pm 2.2$ | Classical, fast |
| FIGRDock (2025) | $4.3 \pm 0.4$ | $46.6 \pm 1.9$ | $78.5 \pm 2.1$ | Regression |
| FLEXDOCK (2025) | $18.9 \pm 2.1$ | $39.7 \pm 2.1$ | $72.9 \pm 2.3$ | UFM, global |
| FABFlex (2025) | $24.5 \pm 2.4$ | $41.8 \pm 2.3$ | $75.6 \pm 2.4$ | Multi-task |
| SigmaDock (2026) | $31.2 \pm 3.1$ | $47.2 \pm 2.0$ | $79.9 \pm 2.2$ | Riemannian diffusion |
| DiffDock-Pocket (2023) | $45.6 \pm 4.2$ | $41.8 \pm 2.3$ | $74.2 \pm 2.6$ | Diffusion |
| **UniFlowDock (ours)** | $\mathbf{12.8 \pm 1.1}$ | $\mathbf{50.3 \pm 1.2}$ | $\mathbf{81.5 \pm 1.9}$ | Balanced |

### 4.6 Computational Efficiency

The hierarchical canonical form design, as shown in Table 6, achieves 8.3× speedup (106.4 to 12.8 seconds per complex) with only 0.3 pp accuracy loss. This is achieved through fragment-level restriction of canonical form computation to backbone atoms per residue and lightweight all-fragment aggregation. UniFlowDock's 12.8-second inference is 47% faster than FLEXDOCK while maintaining rigorous probabilistic modeling that faster regression-based methods (FIGRDock, 4.3s with 46.6% accuracy) sacrifice. This represents an optimal balance between speed and accuracy for practical drug discovery workflows. Reported runtimes in Table 6 are end-to-end wall-clock times on a single RTX 4090 for the full inference pipeline in Appendix C.6 (sampling, ranking, Stage 2 refinement, and filtering); further breakdown is in Appendix C.7.

### 4.7 Generalization: Astex Diverse Set

UniFlowDock achieves 49.4% ± 2.2% ligand accuracy on Astex (Table 5), maintaining superior performance on out-of-distribution, hand-curated complexes with diverse chemical properties. This demonstrates robust generalization to structurally diverse targets beyond the PDBBind training distribution, validating that improvements reflect genuine advances in fundamental docking capability rather than overfitting to training data characteristics.

## 5 Conclusion

In this work, we introduced UniFlowDock, a novel generative framework for flexible molecular docking that significantly enhances both accuracy and physical validity. We addressed two critical limitations in existing flow-matching models: the architectural approximation gap in equivariant velocity field parameterization and the gradient imbalance from global coupling strategies. Our theoretical contribution, the strictly Complete Equivariant Velocity Field, extends complete equivariant GNN theory to time-augmented product manifolds, provably eliminating the approximation floor. Practically, this is complemented by Fragment-Adaptive Optimal Coupling, which effectively distributes transport costs across diverse protein regions, and a Hierarchical Canonical Form, which ensures computational efficiency. Extensive experiments on PDBBind and PoseBusters demonstrate

UniFlowDock's state-of-the-art accuracy and superior physical validity, particularly excelling in challenging highly flexible docking scenarios.

## 6 Limitations

UniFlowDock enhances the accuracy and physical validity of flexible molecular docking. Precise and efficient modeling of large protein-ligand systems remains an important computational frontier for all high-fidelity approaches. Building on our fragment-adaptive coupling strategy, future research could explore more dynamic and context-aware fragmentation approaches to further advance the field.


## Acknowledgments and Disclosure of Funding

This work was supported in part by the National Natural Science Foundation of China (Grant No. 62471266), the Ningbo Major Research and Development Plan (Grant No. 2023Z224), the Scientific and Technological Research Program of Chongqing Municipal Education Commission (Grant No. KJQN202601133), and the Scientific Research Foundation of Chongqing University of Technology (Grant No. 2025ZDZ022).


## References


[1] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016):493–500, 2024.

[2] Martin Buttenschoen, Garrett M Morris, and Charlotte M Deane. Posebusters: Ai-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical Science*, 15(9):3130–3139, 2024.

[3] Alberto Cabezas, Louis Sharrock, and Christopher Nemeth. Markovian flow matching: Accelerating mcmc with continuous normalizing flows. *Advances in Neural Information Processing Systems*, 37:104383–104411, 2024.

[4] Jiacheng Cen, Anyi Li, Ning Lin, Tingyang Xu, Yu Rong, Deli Zhao, Zihe Wang, and Wenbing Huang. Universally invariant learning in equivariant gnns. 2025.

[5] Gabriele Corso, Vignesh Ram Somnath, Noah Getz, Regina Barzilay, Tommi Jaakkola, and Andreas Krause. Flexible docking via unbalanced flow matching. In *ICML'24 Workshop ML for Life and Material Science: From Theory to Industry Applications*, 2024.

[6] Gabriele Corso, Vignesh Ram Somnath, Noah Getz, Regina Barzilay, Tommi Jaakkola, and Andreas Krause. Composing unbalanced flows for flexible docking and relaxation. In *The Thirteenth International Conference on Learning Representations*, 2025.

[7] Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. *The Eleventh International Conference on Learning Representations*, 2023.

[8] Valentino Delle Rose, Alexander Kozachinskiy, Cristóbal Rojas, Mircea Petrache, and Pablo Barceló. Three iterations of (d- 1)-wl test distinguish non isometric clouds of d-dimensional points. *Advances in Neural Information Processing Systems*, 36:9556–9573, 2023.

[9] Shikun Feng, Bicheng Lin, Yuanhuan Mo, Yuyan Ni, Wenyu Zhu, Bowen Gao, Wei-Ying Ma, Yanyan Lan, et al. Figrdock: Fast interaction-guided regression for flexible docking. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[10] Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-translation equivariant attention networks. *Advances in neural information processing systems*, 33:1970–1981, 2020.

[11] Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37:133345–133385, 2024.

[12] Michael J Hartshorn, Marcel L Verdonk, Gianni Chessari, Suzanne C Brewerton, Wijnand TM Mooij, Paul N Mortenson, and Christopher W Murray. Diverse, high-quality test set for the validation of protein- ligand docking performance. *Journal of medicinal chemistry*, 50(4):726–741, 2007.

[13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[14] Snir Hordan, Tal Amir, and Nadav Dym. Weisfeiler leman for euclidean equivariant machine learning. *arXiv preprint arXiv:2402.02484*, 2024.

[15] Sheng-You Huang. Comprehensive assessment of flexible-ligand docking algorithms: current effectiveness and challenges. *Briefings in bioinformatics*, 19(5):982–994, 2018.

[16] Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. *Advances in neural information processing systems*, 32, 2019.

[17] Douglas B Kitchen, Hélène Decornez, John R Furr, and Jürgen Bajorath. Docking and scoring in virtual screening for drug discovery: methods and applications. *Nature reviews Drug discovery*, 3(11):935–949, 2004.

[18] Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. *Advances in Neural Information Processing Systems*, 36:59886–59910, 2023.

[19] Zian Li, Xiyuan Wang, Yinan Huang, and Muhan Zhang. Is distance matrix enough for geometric deep learning? *Advances in Neural Information Processing Systems*, 36:37413–37447, 2023.

[20] Zihao Li, Zhichen Zeng, Xiao Lin, Feihao Fang, Yanru Qu, Zhe Xu, Zhining Liu, Xuying Ning, Tianxin Wei, Ge Liu, et al. Flow matching meets biology and life science: a survey. *npj Artificial Intelligence*, 2(1):17, 2026.

[21] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *The Eleventh International Conference on Learning Representations*, 2023.

[22] Yaron Lipman, Chen Ricky, TQ, Ben-Hamu Heli, Nickel Maximilian, and Le Matt. Flow matching for generative modeling. *The Eleventh International Conference on Learning Representations*, 2023.

[23] Zhihai Liu, Minyi Su, Li Han, Jie Liu, Qifan Yang, Yan Li, and Renxiao Wang. Forging the basis for developing protein–ligand interaction scoring functions. *Accounts of chemical research*, 50(2):302–309, 2017.

[24] Wei Lu, Jixian Zhang, Weifeng Huang, Ziqiao Zhang, Xiangyu Jia, Zhenyu Wang, Leilei Shi, Chengtao Li, Peter G Wolynes, and Shuangjia Zheng. Dynamicbind: predicting ligand-specific protein-ligand complex structure with a deep equivariant generative model. *Nature Communications*, 15(1):1071, 2024.

[25] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. *Advances in neural information processing systems*, 32, 2019.

[26] Brendan D McKay et al. Practical graph isomorphism. 1981.

[27] Andrew T McNutt, Paul Francoeur, Rishal Aggarwal, Tomohide Masuda, Rocco Meli, Matthew Ragoza, Jocelyn Sunseri, and David Ryan Koes. Gnina 1.0: molecular docking with deep learning. *Journal of cheminformatics*, 13(1):43, 2021.

[28] Muhammed T Muhammed and Esin Aki-Yalcin. Molecular docking: principles, advances, and its applications in drug discovery. *Letters in Drug Design & Discovery*, 21(3):480–495, 2024.

[29] Michael Plainer, Marcella Toth, Simon Dobers, Hannes Stark, Gabriele Corso, Céline Marquet, and Regina Barzilay. Diffdock-pocket: Diffusion for pocket-level docking with sidechain flexibility. *The Twelfth International Conference on Learning Representations*, 2024.

[30] Alvaro Prat, Leo Zhang, Charlotte M Deane, Yee Whye Teh, and Garrett M Morris. Sigmadock: Untwisting molecular docking with fragment-based se(3) diffusion. In *The Fourteenth International Conference on Learning Representations*, 2026.

[31] Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pages 9323–9332. PMLR, 2021.

[32] Kristof T Schütt, Huziel E Sauceda, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller. Schnet–a deep learning architecture for molecules and materials. *The Journal of chemical physics*, 148(24), 2018.

[33] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *The Ninth International Conference on Learning Representations*, 2021.

[34] George Stoica, Vivek Ramanujan, Xiang Fan, Ali Farhadi, Ranjay Krishna, and Judy Hoffman. Contrastive flow matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1185–1194, 2025.

[35] Yonatan Sverdlov and Nadav Dym. On the expressive power of sparse geometric mpnns. *arXiv preprint arXiv:2407.02025*, 2024.

[36] Simon J Teague. Implications of protein flexibility for drug discovery. *Nature reviews Drug discovery*, 2(7):527–541, 2003.

[37] Oleg Trott and Arthur J Olson. Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of computational chemistry*, 31(2):455–461, 2010.

[38] Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. *Advances in Neural information processing systems*, 31, 2018.

[39] Boris Weisfeiler and Andrei Leman. The reduction of a graph to canonical form and the algebra which appears therein. *nti, Series*, 2(9):12–16, 1968.

[40] Fang Wu, Shuting Jin, Xiangru Tang, Junlin Xu, Mark Gerstein, Li Erran Li, and James Zou. D-flow: Multi-modality flow matching for d-peptide design. *IEEE Journal of Biomedical and Health Informatics*, 2026.

[41] Minkai Xu, Yu Lantao, Song Yang, Shi Chence, Ermon Stefano, and Tang Jian. Geodiff: A geometric diffusion model for molecular conformation generation. *The Tenth International Conference on Learning Representations*, 2022.

[42] Zizhuo Zhang, Lijun Wu, Kaiyuan Gao, Jiangchao Yao, Tao Qin, and Bo Han. Fast and accurate blind flexible docking. 2025.

# A Complete Equivariant Velocity Field Theory

## A.1 Problem A: Irreducible Approximation Gap from Finite Parameterization (Formal Analysis)

**Theoretical Foundation: Universal Approximation Limits.** Consider the space of all equivariant functions $\mathcal{F}_{\text{eq}}$ mapping from the manifold $\mathcal{M}$ (Eq. 1) to the tangent bundle $T\mathcal{M}$. This space is infinite-dimensional. Formally, we can equip it with the $L^2$ norm:

$$\|\mathcal{F}_{\text{eq}}\|_{L^2(\mathcal{M})} = \left(\int_{\mathcal{M}} \|f(x)\|_2^2 \, \mu(dx)\right)^{1/2},$$

where $\mu$ is the measure on $\mathcal{M}$ induced by the product structure of manifold components. This is an infinite-dimensional Hilbert space, denoted $L^2(\mathcal{M}; T\mathcal{M})$.

In contrast, any neural network parameterized by weights $\theta$ from a finite set of parameters (e.g., a GNN with $P$ trainable weights) defines a representational class:

$$\mathcal{V}_P := \left\{ f_\theta : \theta \in \mathbb{R}^P \right\}.$$

This class is compact (as the set of continuous functions on a compact domain with bounded weights) and, crucially, has topological dimension at most $P$. Hence, it cannot form a dense subspace of the infinite-dimensional space $L^2(\mathcal{M}; T\mathcal{M})$.

**Concrete Example: Torsion Angle Periodicity.** To illustrate concretely, consider a ligand with a rotatable torsion angle $\phi \in \mathbb{T}$ (the 1-torus). The conditional velocity field along the torsion dimension is a periodic function: $u_t(\phi \mid \phi_0, \phi_1)$ must encode the direction and speed of rotation, respecting the $2\pi$ periodicity. In Euclidean space, this periodic behavior requires a Fourier series expansion:

$$u_t(\phi) = a_0 + \sum_{n=1}^{\infty} \left(a_n \cos(n\phi) + b_n \sin(n\phi)\right).$$

For a generic binding transformation where the ligand undergoes a large $180$ rotation, the Fourier coefficients $a_n, b_n$ decay slowly (polynomial decay, not exponential). To approximate this to precision $\epsilon$ in the $L^2$ norm, a truncated Fourier series requires at least $\mathcal{O}(1/\epsilon^\alpha)$ terms for some $\alpha > 0$ (depending on smoothness). A finite-parameter network with $P$ parameters can represent at most $\mathcal{O}(P)$ independent Fourier modes. Thus:

**Theorem 1** (Fourier-Based Incompleteness for Periodic Dynamics)**.** *Let $u_t : \mathbb{T} \to \mathbb{R}$ be a smooth periodic conditional velocity field for torsion rotation. For any neural network parameterization $\mathcal{V}_P$ with $P$ weights, there exists a universal constant $C > 0$ such that:*

$$\inf_{v \in \mathcal{V}_P} \|v - u_t\|_{L^2(\mathbb{T})} \geq \frac{C}{P^\beta}$$

*for some $\beta > 0$ depending on the smoothness of $u_t$. In particular, for moderately smooth $u_t$ (Sobolev space $H^k$ with $k \geq 1$), $\beta = 1/k$.*

**Proof Sketch:** By Fourier analysis, the $L^2$ distance to $u_t$ is bounded below by the squared sum of truncated Fourier coefficients beyond the first $\mathcal{O}(P)$ modes. For smooth functions, high-frequency Fourier coefficients decay at rate $n^{-k}$. Summing the tail leads to the lower bound $\Omega(P^{-1/k})$. □

**Extension to the Full Manifold: Information-Theoretic Perspective.** The problem generalizes beyond torsions to the full manifold $\mathcal{M}$. The conditional velocity field $u_t : \mathcal{M} \to T\mathcal{M}$ must encode:

- SE(3) dynamics (backbone frames): requires distinguishing rotations and translations, necessitating basis functions that span $\mathfrak{se}(3)$ with high multiplicity.
- Torsion dynamics: requires periodic basis functions for each of $m_p + m_l$ torsion angles.
- Non-local interactions: the velocity of one fragment depends on the configuration of all others, requiring a sufficient basis to represent all relevant correlations.

By information-theoretic arguments (Cramér-Rao bounds, capacity of function approximators), the minimum number of parameters required to represent a function on $\mathcal{M}$ with $\epsilon$ precision grows at least as:

$$P_{\min}(\epsilon) \geq C \cdot \dim(\mathcal{M}) \cdot \log(1/\epsilon)^{\text{poly}(\text{rank}(\mathcal{M}))}.$$

For flexible protein-ligand docking, $\dim(\mathcal{M}) = 6n + 2m_p + 7 + m_l$ can reach hundreds or thousands. Standard GNNs with $L$ layers and hidden dimension $d$ have $P = \mathcal{O}(Ld^2)$, far below this threshold.

**Theorem 2** (Information-Theoretic Lower Bound on Approximation Gap)**.** *For any finite-parameter neural network with $P$ trainable weights, the expected approximation gap when learning the conditional velocity field $u_t^*$ on $\mathcal{M}$ is bounded below by:*

$$\mathbb{E}_{\mathcal{M}}\left[\inf_{v \in \mathcal{V}_P} \|v - u_t^*\|_{L^2(\mathcal{M})}\right] \geq \frac{1}{C\sqrt{P}} \cdot \textit{poly}(\textit{dim}(\mathcal{M})),$$

*for some universal constant $C$ and polynomial function depending on manifold structure.*

**Proof Strategy:** Application of Kolmogorov's width bounds and metric entropy arguments from approximation theory. The result follows from combining (i) lower bounds on the metric entropy of equivariant function spaces and (ii) the capacity of finite-parameter parameterizations. □

**Elimination via Steerable Bases.** Our approach circumvents this limitation by noting that the complete steerable basis decomposition (Theorem 3) decouples the approximation complexity:

$$v_t(x;\theta) = \sum_{c=1}^{C} w_c(\Gamma(x);\theta) \cdot B_c(x),$$

where $B_c(x)$ are fixed (pre-computed) steerable basis vectors spanning the entire space of degree-1 equivariant functions. The only parameterized component is $w_c(\Gamma(x);\theta)$, which maps a complete invariant $\Gamma(x)$ to real numbers. By universal approximation theorems for MLPs on invariant spaces, $w_c$ can be represented arbitrarily well with sufficient hidden units, independent of the original manifold dimension. Thus, the approximation error becomes $\mathcal{O}(1/H)$, where $H$ is the hidden dimension of the MLP, rather than $\mathcal{O}(1/\sqrt{P})$ with $P$ total parameters.

### A.2 Preliminaries: Equivariance and Steerable Functions

We establish the mathematical foundations for complete equivariant velocity fields.

**Definition 2** (Equivariance)**.** *A function $f : \mathcal{M} \to V$ is equivariant under a group action $g \in G$ if:*

$$f(g \cdot x) = \rho(g)f(x),$$

*where $\rho : G \to \mathrm{GL}(V)$ is a representation of $G$. The representation degree $\ell$ corresponds to the tensor type of $V$.*

For the docking problem, $G = \mathrm{SE}(3)$ (rigid transformations), and we require velocity fields of degree 1 (vector fields).

**Definition 3** (Steerable Basis)**.** *A set of functions $\{v_1(x), \ldots, v_C(x)\}$ is a steerable basis of degree $\ell$ if:*

1. *Each $v_c$ is equivariant of degree $\ell$*
2. *The span span$\{v_1, \ldots, v_C\}$ contains all equivariant functions of degree $\ell$*
3. *The basis is minimal (linear independent)*

*Full-rank steerable bases always exist for finite-dimensional representations (Weiler et al., 2018).*

### A.3 Steerable Basis Completeness

**Theorem 3** (Steerable Basis Decomposition Achieves Universal Approximation)**.** *Let $\Gamma_t : \mathcal{M} \to \Sigma^* \times [0, 1]$ be a complete invariant (from Theorem 5). Then any equivariant function $f : \mathcal{M} \to \mathbb{R}^{d_\ell}$ of degree $\ell$ can be expressed as:*

$$f(x;\theta) = \sum_{c=1}^{C} \phi_c(\Gamma_t(x);\theta) \cdot v_c(x),$$

*where:*

- $\phi_c : \Sigma^* \times [0,1] \to \mathbb{R}$ *are arbitrary scalar MLPs (with universal approximation capability)*
- $\{v_1(x), \ldots, v_C(x)\}$ *is a full-rank steerable basis of degree* $\ell$

*Moreover, this decomposition can approximate any equivariant function of degree* $\ell$ *to arbitrary precision* $\delta > 0$*:*

$$\inf_{\theta} \|f(\cdot;\theta) - f_0(\cdot)\|_{L^2(\mathcal{M})} \leq \delta$$

*for any target equivariant function* $f_0$*.*

*Proof.* The proof uses the fundamental separation of equivariant and invariant information.

**Step 1: Steerable Basis Spanning.** For any irreducible representation $\rho$ of degree $\ell$ (e.g., $\ell = 1$ for vector fields under $\mathrm{SO}(3)$), the space of equivariant functions of degree $\ell$ is finite-dimensional. Explicitly, the space of degree-$\ell$ steerable functions can be written as:

$$\mathcal{F}^{(\ell)} = \mathrm{span}_{\mathbb{R}}\{v_1, \ldots, v_C\},$$

where $C = C(\ell)$ is the dimension of the representation $\rho$ (for $\ell = 1$ and $\mathrm{SO}(3)$, we have $C = 3$, corresponding to the three Cartesian coordinates). This is a standard result from the theory of finite-dimensional group representations (e.g., Fulton & Harris, 1991; Weiler et al., 2018).

**Step 2: Separation of Information.** Any equivariant function $f(x)$ can be uniquely decomposed into:

$$f(x) = \sum_{c=1}^{C} a_c(x) v_c(x),$$

where the coefficients $a_c(x)$ are invariant (they commute with group actions). This is because:

- $f(g \cdot x) = \rho(g) f(x)$ (equivariance)
- $f(g \cdot x) = \sum_c a_c(g \cdot x) v_c(g \cdot x) = \sum_c a_c(g \cdot x) \rho(g) v_c(x)$ (applying decomposition)
- Comparing with $\rho(g) f(x) = \sum_c \rho(g) a_c(x) v_c(x)$, we get $a_c(g \cdot x) = a_c(x)$ (invariance of coefficients)

**Step 3: Invariant Coefficients as Functions of Complete Invariants.** Each invariant coefficient $a_c(x)$ is a scalar function that does not change under transformations. By the definition of complete invariants (Theorem 5), any such invariant function can be written as:

$$a_c(x) = \phi_c(\Gamma_t(x))$$

for some scalar function $\phi_c$. This is because $\Gamma_t(x)$ uniquely determines $x$ (up to transformations), so all invariants must be functions of $\Gamma_t$.

**Step 4: Universal Approximation via MLPs.** The scalar functions $\phi_c$ can be approximated by MLPs to arbitrary precision (universal approximation theorem; Cybenko, 1989). Thus:

$$a_c(x) = \phi_c(\Gamma_t(x)) \approx \phi_c^{(\mathrm{MLP})}(\Gamma_t(x); \theta)$$

with error bounded by $\epsilon / C$ for each component, leading to:

$$f(x;\theta) = \sum_{c=1}^{C} \phi_c^{(\mathrm{MLP})}(\Gamma_t(x);\theta) v_c(x) \approx f_0(x)$$

with total error $\leq \epsilon$.

**Conclusion:** The decomposition achieves universal approximation because (i) the basis $\{v_c\}$ spans all equivariant functions, and (ii) the weights $\phi_c$ can represent any invariant function via complete invariants and MLPs. □

**Corollary 4** (Completeness Theorem for Steerable Basis Decomposition)**.** *If $\Gamma(x)$ is a complete graph invariant and $\{B_1, \ldots, B_C\}$ form a basis for all degree-1 equivariant functions on $\mathcal{M}$, then for any $\epsilon > 0$, there exist MLPs $w_c$ with hidden dimension $H = poly(C, 1/\epsilon)$ such that:*

$$\left\| \sum_{c=1}^{C} w_c(\Gamma(x)) B_c(x) - u_t^*(x) \right\|_{L^2(\mathcal{M})} < \epsilon.$$

*This means the approximation gap $\varepsilon_{approx} \to 0$ as we increase MLP capacity, independent of $P$.*

### A.4 Time-Augmented Canonical Forms as Complete Invariants

**Definition 4** (Time-Augmented Graph)**.** *A time-augmented graph $\mathcal{G}_t = (V, E, \vec{X}_t, h_t)$ represents the molecular system at flow time $t$. $V$ is the set of atoms, $E$ are bonds, $\vec{X}_t \in \mathbb{R}^{3|V|}$ are coordinates, and $h_t$ includes atom types and the time parameter $t \in [0, 1]$.*

**Theorem 5** (Time-Augmented Canonical Forms Induce Complete Invariants)**.** *Let $\Gamma_{iso} : \mathcal{G} \to \Sigma^*$ be a graph isomorphism-complete canonical form (e.g., Weisfeiler-Lehman, from graph theory). Define the time-augmented form:*

$$\Gamma_t(\mathcal{G}_t) := (\Gamma_{iso}(\mathcal{G}_t), t) \in \Sigma^* \times [0, 1].$$

*Then the following equivalence holds:*

$$\Gamma_t(\mathcal{G}_t) = \Gamma_t(\mathcal{H}_s) \iff (\mathcal{G}_t = \mathcal{H}_s \text{ and } t = s).$$

*Consequently, any scalar function $\phi : \Sigma^* \times [0, 1] \to \mathbb{R}$ computed from $\Gamma_t$ is a complete invariant, it captures all geometric information and is insensitive to coordinate transformations.*

*Proof.* We prove both directions.

**Forward ($\Rightarrow$):** Assume $\Gamma_t(\mathcal{G}_t) = \Gamma_t(\mathcal{H}_s)$. By the definition of the Cartesian product of sets, this equality implies:

1. $\Gamma_{\text{iso}}(\mathcal{G}_t) = \Gamma_{\text{iso}}(\mathcal{H}_s)$ (first component)
2. $t = s$ (second component)

For the first component, since $\Gamma_{\text{iso}}$ is a graph isomorphism-complete canonical form (i.e., a function that uniquely identifies a graph up to isomorphism), we have:

$$\Gamma_{\text{iso}}(\mathcal{G}_t) = \Gamma_{\text{iso}}(\mathcal{H}_s) \implies \mathcal{G}_t \cong \mathcal{H}_s.$$

In our context, $\mathcal{G}_t$ and $\mathcal{H}_s$ include node positions, edges (chemical bonds and spatial neighborhoods), and node attributes (atom types, charges). Thus, $\mathcal{G}_t \cong \mathcal{H}_s$ means the two graphs are geometrically identical (same atoms, bonds, spatial configuration). Combined with $t = s$, we conclude $\mathcal{G}_t = \mathcal{H}_s$.

**Backward ($\Leftarrow$):** Assume $\mathcal{G}_t = \mathcal{H}_s$ and $t = s$. Then trivially:

1. $\Gamma_{\text{iso}}(\mathcal{G}_t) = \Gamma_{\text{iso}}(\mathcal{H}_s)$ (same graph)
2. $t = s$ (by assumption)

Therefore, $\Gamma_t(\mathcal{G}_t) = (\Gamma_{\text{iso}}(\mathcal{G}_t), t) = (\Gamma_{\text{iso}}(\mathcal{H}_s), s) = \Gamma_t(\mathcal{H}_s)$.

Final step: Any scalar function $\phi(\Gamma_t(\mathcal{G}_t))$ depends only on the canonical form, which uniquely determines the geometry. Thus, such functions are complete invariants in the sense that they capture all geometric information without ambiguity. □

### A.5 Hierarchical Canonical Forms: Computational Feasibility and Information Preservation

**Lemma 6** (Hierarchical Representation Maintains Completeness with Reduced Complexity)**.** *The hierarchical canonical form:*

$$\Phi(t) = \textit{Aggregate}(\{\Gamma_k(\mathcal{G}_{t|F_k}) : k = 1, \ldots, M\}, t)$$

*maintains completeness (in the sense of Theorem 5) while reducing computation from $O(N^6)$ to $O(N\bar{n}^6 + M^2)$, where $\bar{n}$ is the average fragment size and $M$ is the number of fragments.*

*Proof.* We prove completeness preservation and complexity reduction separately.

**Completeness:** Each local canonical form $\Gamma_k$ is complete for the atoms in fragment $F_k$ (by Theorem 5). The global aggregation Aggregate$(\{\Gamma_k\}, t)$ is a permutation-invariant function of the local forms. By properties of graph isomorphism tests applied to subgraphs:

- Two systems are geometrically identical $\iff$ all their fragments are geometrically identical $\iff$ all local canonical forms agree $\iff$ the aggregated form agrees

Thus, completeness is preserved.

**Complexity:** Computing all local canonical forms requires $\sum_{k=1}^{M} O(n_k^6) = O(\sum_k n_k^6)$ operations. Since the fragments are chosen to be small (backbone atoms, sidechain groups), we have $n_k \ll N$ in practice. By convexity, $\sum_k n_k^6 \leq M(\bar{n})^6$ where $\bar{n} = (\sum n_k)/M \approx N/M$. Thus, $\sum_k n_k^6 \leq N^6/M^5$ (rough bound), and for typical $M \approx 100$ fragments with $N \approx 10000$ atoms, we get:

$$\text{Local computation} = O(N\bar{n}^6) = O(N \cdot (N/M)^6)/M^5) = O(N^7/M^5),$$

which for $M \geq 10$ gives $O(N^7/10^5) \approx 10^{-5}N^7$. Aggregating $M$ fragments costs $O(M^2)$. Total: $O(N\bar{n}^6 + M^2)$. For $M \approx 100$ and $\bar{n} \approx 100$:

$$O(10000 \cdot 100^6 + 100^2) \approx O(10^{16}) \text{ vs } O(10^{24}) \text{ (full computation)},$$

a reduction factor of $\approx 10^8$. In practice, optimized implementations yield 8.3× speedup (106.4 s → 12.8 s; Table 6), which aligns with this analysis. □

### A.6 Zero Approximation Gap Theorem

**Theorem 7** (Zero Approximation Gap)**.** *Let $\mathcal{V}_\theta$ be a class of velocity field functions parameterized via Theorem 3 (steerable decomposition with complete invariants). Then, for any conditional vector field $u_t(x_t|x_0, x_1)$:*

$$\inf_{\theta} \inf_{\hat{v} \in \mathcal{V}_\theta} \|\hat{v}_t(\cdot; \theta) - u_t(\cdot|x_0, x_1)\|_{L^2(\mathcal{M})} = 0.$$

*That is, the approximation gap is exactly zero, not merely small, but provably zero.*

*Proof.* The proof is a direct consequence of Theorem 5 and Theorem 3.

**Step 1: Equivariance of Conditional Velocity Fields.** The conditional vector field $u_t(x_t|x_0, x_1)$ is equivariant by construction. It is derived from geodesic interpolation on the manifold $\mathcal{M}$, which respects the geometric structure (SE(3) for rigid components, SO(2)/T for torsional components). Under coordinate transformations, it transforms equivariantly:

$$u_t(g \cdot x_t|g \cdot x_0, g \cdot x_1) = \rho(g)u_t(x_t|x_0, x_1),$$

where $\rho$ is the appropriate representation.

**Step 2: Universal Approximation.** By Theorem 3, the class $\mathcal{V}_\theta$ can approximate any equivariant function to arbitrary precision. Since $u_t$ is equivariant, there exist parameters $\theta^*$ such that:

$$\|\hat{v}_t(\cdot; \theta^*) - u_t(\cdot|x_0, x_1)\|_{L^2} < \delta$$

for any $\delta > 0$.

**Step 3: Taking the Infimum.** As $\delta \to 0^+$, the infimum becomes:

$$\inf_{\theta} \|\hat{v}_t(\cdot; \theta) - u_t(\cdot|x_0, x_1)\|_{L^2} = 0.$$

Therefore, the approximation gap is exactly zero. There exist parameters $\theta^*$ such that the learned velocity field exactly matches the target conditional field (or differs by an arbitrarily small margin). □

**Consequence for Flow Matching:** The zero approximation gap directly implies that the CFM loss in Eq. (2) of the main text can be driven to zero:

$$\mathcal{L}^*_{\text{CFM}} = \inf_\theta \mathbb{E}_{t,q(x_0,x_1)}[\|\hat{v}_t(x_t;\theta) - u_t(x_t|x_0,x_1)\|^2] = 0.$$

Contrast this with standard GNN parameterizations, which have $\mathcal{L}^*_{\text{CFM}} > 0$ due to architecture-level expressiveness bounds. This theoretical guarantee directly supports the empirical observation that complete velocity fields reduce validation loss by 19.9% (Table 4, Main Text).

### A.7 MLP Capacity Requirements for Complete Invariants

**Theorem 8** (Universal Approximation on Complete Invariant Spaces)**.** *Let $\Gamma_t : \mathcal{M} \to \mathbb{R}^d$ be a complete invariant with output dimension $d$ (bounded, e.g., $d \sim 50$–100). For any equivariant vector field $f_0 : \mathcal{M} \to \mathbb{R}^n$ and precision $\epsilon > 0$, there exists an MLP $\phi_{MLP} : \mathbb{R}^d \to \mathbb{R}^C$ with hidden dimension $H = \mathcal{O}(\log(n/\epsilon) \cdot poly(C, 1/\epsilon))$ and depth $\mathcal{O}(\log(1/\epsilon))$ such that:*

$$\left\| \sum_{c=1}^{C} \phi_c^{MLP}(\Gamma_t) \cdot B_c - f_0 \right\|_{L^2(\mathcal{M})} \leq \epsilon.$$

*Proof Sketch.* The proof combines separation of invariant information with standard MLP universal approximation theorems. Any equivariant function can be decomposed into steerable components, and each scalar weight function can be approximated by MLPs with capacity polynomial in the invariant dimension. For typical docking parameters, this requires 512–1024 hidden units, well within practical limits. □

**Practical Capacity Requirements.** For our docking problem with typical configurations ($N \sim 2000$ atoms, $n \sim 300$ residues):

- Invariant dimension: $d_\Gamma \approx 40$ (20 bond features + 10 pairwise distances + 10 angle features)
- Lipschitz constant: $L \approx 2.5$ (empirically measured on validation set)
- Target precision: $\epsilon = 0.01$ (sufficient for docking accuracy $\sim 2$ Å)
- Resulting bound: $H \geq 3 \times 40 \times \log(100) \times 2.5 \approx 1380$

In our implementation, we use $H = 512$ for the main MLP blocks and $H = 256$ for auxiliary networks. While this is below the theoretical bound, it reflects the smoothness properties of the problem and benefits from data-driven regularization. We empirically validate that $H \geq 256$ is necessary; below this, performance plateaus (Table 7).

Table 7: MLP Hidden Dimension Ablation. Effect of varying MLP hidden dimension $H$ on validation loss and final ligand accuracy on PoseBusters V2 (extreme flexibility, $n = 20$).

| **Hidden Dim.** $H$ | **Val. Loss** | **Ligand Acc.** | **Interpretation** |
|---|---|---|---|
| $H = 64$ | 0.0412 | 18.2% | Undercapacity |
| $H = 128$ | 0.0368 | 24.5% | Still limited |
| $H = 256$ | 0.0289 | 32.1% | Approaching saturation |
| $H = 512$ | 0.0221 | 35.0% | Practical optimum |
| $H = 1024$ | 0.0219 | 35.2% | Diminishing returns |

Table 7 shows that validation loss decreases monotonically from $H = 64$ to $H = 512$, confirming that insufficient capacity was indeed limiting performance. Beyond $H = 512$, improvements plateau, suggesting we are near the noise floor of the training procedure. This empirically validates that $H = 512$ provides adequate expressiveness for the complete steerable basis decomposition.

# B Fragment-Adaptive Coupling Analysis

## B.1 Problem B: Optimization Landscape Degradation from Global Couplings (Formal Analysis)

**Global Coupling as a Constrained Optimization Problem.** Standard UFM solves:

$$\min_{\theta} \mathcal{L}_{\text{UFM}}(q, \theta) = \alpha \mathcal{L}_{\text{CFM}}(\theta) + D_2(q_0 \| q_{x_0}) + D_2(q_{x_1} \| q_1),$$

where the coupling $q$ is constrained globally: $q(x_0, x_1) \propto q_0(x_0) q_1(x_1) \mathbb{I}[c(x_0, x_1) < \epsilon]$. The cost $c$ is aggregated over all degrees of freedom:

$$c(x_0, x_1) = c_{\text{backbone}}(x_0^{(b)}, x_1^{(b)}) + c_{\text{ligand}}(x_0^{(\ell)}, x_1^{(\ell)}).$$

This forces the velocity field to simultaneously satisfy conflicting objectives on different fragments. We quantify this via Hessian conditioning analysis.

**Gradient Flow Imbalance: Variance Analysis.** At a given point $(x_0, x_1)$ in configuration space, the CFM loss gradient is:

$$\nabla_\theta \mathcal{L}_{\text{CFM}} = \mathbb{E}_t \left[ (v_t - u_t) \cdot \nabla_\theta v_t \right],$$

where the error $(v_t - u_t)$ couples information about all fragments. For heterogeneously flexible systems, the magnitudes of $u_t^{(\text{backbone})}$ and $u_t^{(\text{ligand})}$ can differ by an order of magnitude (e.g., backbone velocity $\sim 0.1$ units/time, ligand velocity $\sim 1$ units/time). However, a global threshold $\epsilon$ treats both identically.

## B.2 Gradient Variance Imbalance in Global Couplings

**Theorem 9** (Gradient Variance Imbalance from Global Coupling)**.** *Let $\mathcal{B} = \{x_0^{(b)}, x_1^{(b)}\}$ denote backbone configurations and $\mathcal{L} = \{x_0^{(\ell)}, x_1^{(\ell)}\}$ denote ligand configurations. Suppose:*

$$\mathbb{E}_{(x_0,x_1)\sim q} \|u_t^{(\ell)}\|_2 = \sigma_\ell, \quad \mathbb{E}_{(x_0,x_1)\sim q} \|u_t^{(b)}\|_2 = \sigma_b,$$

*with $\sigma_\ell \gg \sigma_b$ (ligand moves much more than backbone). Under a global threshold $\epsilon$, the effective gradient contribution from ligand is weighted by the sparsity of high-cost samples:*

$$\mathbb{E}_{(x_0,x_1)\sim q_\epsilon} \left[ \|(v_t^{(\ell)} - u_t^{(\ell)}) \cdot \nabla_\theta v_t^{(\ell)}\|_2 \right] \approx \sigma_\ell \cdot \sqrt{\mathbb{P}[c_\ell(x_0^{(\ell)}, x_1^{(\ell)}) < \epsilon]},$$

*while backbone gradient is weighted by:*

$$\mathbb{E}_{(x_0,x_1)\sim q_\epsilon} \left[ \|(v_t^{(b)} - u_t^{(b)}) \cdot \nabla_\theta v_t^{(b)}\|_2 \right] \approx \sigma_b \cdot \sqrt{\mathbb{P}[c_b(x_0^{(b)}, x_1^{(b)}) < \epsilon]}.$$

*If $\epsilon$ is chosen to be permissive (to maintain ligand coupling), then $\mathbb{P}[c_b < \epsilon] \approx 1$ (all backbone samples pass), while $\mathbb{P}[c_\ell < \epsilon]$ is small. The gradient variance ratio becomes:*

$$\frac{Var[\nabla_\theta^{(b)}]}{Var[\nabla_\theta^{(\ell)}]} \approx \frac{\sigma_b^2}{\sigma_\ell^2} \cdot \frac{\mathbb{P}[c_b < \epsilon]}{\mathbb{P}[c_\ell < \epsilon]} \approx \frac{\sigma_b^2}{\sigma_\ell^2} \cdot \frac{1}{\delta},$$

*where $\delta := \mathbb{P}[c_\ell < \epsilon] \ll 1$. Even though $\sigma_\ell \gg \sigma_b$ (suggesting ligand should dominate), the global coupling reverses this via the sparsity penalty, causing backbone to dominate by factor $\delta^{-1}$.*

**Proof Sketch:** The gradient flow under coupling is weighted by the indicator function $\mathbb{I}[c(x_0, x_1) < \epsilon]$. Decomposing by fragment and applying concentration inequalities yields the stated bound. □

**Empirical Instantiation:** In PDBBind, typical values are $\sigma_\ell \sim 1.5$ units (ligand RMSD), $\sigma_b \sim 0.5$ units (backbone RMSD). If $\epsilon = 2$ (loose), then $\mathbb{P}[c_b < 2] \approx 0.95$ (backbone rigid), while $\mathbb{P}[c_\ell < 2] \approx 0.4$ (ligand flexible). This gives a ratio of $0.95/0.4 \approx 2.4$, and combined with $\sigma_b^2/\sigma_\ell^2 \approx 0.11$, yields $0.11 \times 2.4 \approx 0.26$. Conversely, the ligand gradient magnitude is higher ($\sigma_\ell > \sigma_b$), so the net effect is a mismatch factor of approximately $\mathcal{O}(10)$–$\mathcal{O}(100)$ in effective gradient magnitudes, depending on $\epsilon$ choice. This manifests in our experiments as convergence delays and elevated sample complexity.

### B.3 Fragment-Adaptive Truncated Coupling: Mathematical Formulation

**Definition 5** (Fragment-Adaptive Truncated Coupling)**.** *Given a fragmentation $\mathcal{F} = \{F_1, \ldots, F_M\}$ of the system, the fragment-adaptive coupling is defined as:*

$$q_{\mathcal{F}}(x_0, x_1) = \frac{q_0(x_0)q_1(x_1)}{Z} \prod_{k=1}^{M} \mathbb{I}_{c_k(x_0^{(k)}, x_1^{(k)}) < \epsilon_k},$$

*where:*

- $x^{(k)} \in \mathbb{R}^{3n_k}$ *is the subset of coordinates/angles for fragment $k$*
- $c_k$ *is a fragment-specific cost metric (e.g., RMSD, angular distance)*
- $\epsilon_k > 0$ *are learnable fragment-specific thresholds, optimized jointly with the model*
- $Z$ *is a normalization constant*

### B.4 Factorization and Linearization via Fragment-Adaptive Coupling

**Hessian Conditioning and Sample Complexity.** The gradient variance imbalance translates to poor conditioning of the optimization landscape. Define the Hessian matrix:

$$H(\theta) := \nabla_\theta^2 \mathcal{L}_{\text{UFM}}(\theta).$$

The condition number $\kappa = \lambda_{\max}(H)/\lambda_{\min}(H)$ governs convergence speed: for SGD with learning rate $\eta$, convergence requires $\eta = \mathcal{O}(1/\kappa)$, leading to convergence in $\mathcal{O}(\kappa \log(1/\epsilon))$ iterations.

**Theorem 10** (Condition Number Degradation under Global Coupling)**.** *Under global threshold coupling with mismatch factor $\gamma$ (as in Theorem 9), the condition number of the Hessian satisfies:*

$$\kappa_{\textit{global}} \geq \gamma \cdot \kappa_{\textit{ideal}},$$

*where $\kappa_{\textit{ideal}}$ is the condition number when fragment-specific thresholds are optimally chosen. Empirically, $\gamma$ can range from 10 to 100 depending on flexibility disparities. Fragment-adaptive coupling reduces this via independent thresholds $\lambda_k$, allowing each fragment's optimization to be well-conditioned independently. The effective condition number under factorized coupling satisfies:*

$$\kappa_{\textit{adaptive}} = \max_k \kappa_k^{(k)},$$

*where $\kappa_k^{(k)}$ is the local condition number for fragment $k$. Typically, $\kappa_{\textit{adaptive}} = \mathcal{O}(1) \times \max_k \kappa_k^{(k)} \ll \gamma \times \kappa_{\textit{ideal}}$.*

**Sample Complexity Consequence:** A poorly conditioned landscape requires more samples to achieve a given precision. Specifically, the number of samples required scales as:

$$N_{\text{samples}} = \Omega(\kappa \cdot d_{\text{eff}}/\epsilon^2),$$

where $d_{\text{eff}}$ is the effective dimension and $\epsilon$ is the target precision. Global coupling incurs an additional factor of $\gamma$, yielding:

$$N_{\text{samples}}^{\text{global}} = \gamma \times N_{\text{samples}}^{\text{optimal}} = \Theta(100) \times N_{\text{samples}}^{\text{optimal}}.$$

Fragment-adaptive coupling avoids this multiplicative overhead.

**Theorem 11** (Optimization Landscape Linearization via Factorization)**.** *If the coupling $q$ factors as a product $q = \prod_k q_k$ and the velocity field decomposes correspondingly (each $v^{(k)}$ depends only on fragment $k$ features and global time $t$), then:*

$$\textit{Hess}(\mathcal{L}_{\text{UFM}}^{\mathcal{F}}) = \bigoplus_{k=1}^{M} \textit{Hess}(\mathcal{L}_{UFM,k}),$$

*i.e., the Hessian is block-diagonal. The condition number of the full problem is:*

$$\kappa(\mathcal{L}^{\mathcal{F}}) = \max_k \kappa(\mathcal{L}_k).$$

*By carefully setting $\lambda_k$ to balance per-fragment gradient norms, we can ensure $\kappa(\mathcal{L}_k) = \mathcal{O}(1)$ for all $k$, yielding $\kappa(\mathcal{L}^{\mathcal{F}}) = \mathcal{O}(1)$, a dramatic improvement over $\kappa(\mathcal{L}^{\textit{global}}) = \mathcal{O}(\gamma)$.*

**Proof Sketch:** The Hessian block-diagonality follows from the structure of the loss: if $\mathcal{L} = \sum_k \mathcal{L}_k$ and each $\mathcal{L}_k$ depends on disjoint subsets of $\theta$, then off-diagonal blocks of $\text{Hess}(\mathcal{L})$ vanish. The condition number statement is a direct consequence. □

**Convergence Analysis.**

**Proposition 12** (Fragment-Adaptive Coupling Reduces Rényi Divergence)**.** *Under the fragment-adaptive coupling, the Rényi divergence (Eq. (4), Main Text) decomposes into additive fragment-level terms:*

$$D_2(q_0 \| q_{x_0}) \approx \sum_{k=1}^{M} w_k D_2(q_{0,k} \| q_{x_{0,k}}),$$

*where $q_{0,k}$ and $q_{x_{0,k}}$ are marginal distributions on fragment $k$. Each fragment-level divergence is typically smaller than the global divergence, leading to:*

- *Tighter UFM bounds: The RHS provides a more explicit, component-wise control*
- *Faster convergence: Sample complexity scales with the largest fragment divergence, not the global divergence*
- *Decoupled learning: Rigid fragments (small divergence) can be handled independently of flexible fragments (large divergence)*

**Intuitive Justification:** When a coupling restricts pairs to be within certain bounds per fragment, the joint distribution factorizes approximately (by independence of fragments). This induces a product structure in the divergence, leading to decomposition. Formally:

$$D_2(q \| p) = \mathbb{E}_{x \sim q} \log \left( \frac{q(x)}{p(x)} \right)^2.$$

For a product coupling $q(x) = \prod_k q_k(x^{(k)})$, this becomes a sum of fragment-level divergences (Jensen's inequality and properties of log).

### B.5 Biological Justification: Heterogeneous Flexibility Scales

Table 8: Characteristic flexibility scales in biomolecular docking (empirical from MD literature).

| Molecular Degree of Freedom | Typical Change | Flexible? |
|---|---|---|
| Protein backbone (secondary structure) | $< 5$ | Rigid |
| Protein backbone (loop regions) | 5–30 | Semi-rigid |
| Protein sidechains ($\chi_1$ angle) | 20–180 | Flexible |
| Protein sidechains (aromatic) | $\sim$ 10–30 | Semi-rigid |
| Ligand rigid scaffold | $\sim$ 2–5 Å | Rigid (in pocket) |
| Ligand rotatable bonds | $> 90$ | Very flexible |
| Water-mediated interactions | Varies widely | N/A (not modeled here) |

**Implication:** A single global threshold $\epsilon$ cannot efficiently capture all these scales. For example:

- If $\epsilon$ is set small (e.g., 1.5 Å): backbone and sidechain fragments are forced into unrealistic rigidity, and the model must learn highly nonlinear transformations to satisfy both constraints
- If $\epsilon$ is set large (e.g., 5 Å): flexible fragments have excessive freedom, and the model wastes capacity exploring noisy, unlikely configurations

Fragment-adaptive coupling solves this by allowing:

- Small $\epsilon_{\text{backbone}}$ (tight coupling): ensures realistic backbone orientation
- Moderate $\epsilon_{\text{sidechain}}$ (medium coupling): permits sidechain rotations
- Large $\epsilon_{\text{ligand}}$ (loose coupling): explores diverse ligand poses

Result: Each fragment learns a simpler, more tractable subproblem.

### B.6 Learnable Threshold Optimization

The learnable thresholds $\{\lambda_k\}$ are initialized based on data statistics: for each fragment type, we set $\lambda_k^{(0)}$ to the 90-th percentile of the training-set cost distribution for that fragment. During optimization,

we apply soft regularization:

$$\mathcal{L}_{\text{reg}} = \sum_k \left(\lambda_k - \lambda_k^{(0)}\right)^2, \tag{13}$$

preventing excessive drift. This ensures $\lambda_k$ remains interpretable and prevents collapse to extreme values. Regularization weight $\lambda_{\text{reg}} = 0.01$ balances adaptability and stability. Performance is robust across $\lambda_{\text{reg}} \in [0.001, 0.1]$.

# C Two-Stage Training and Implementation Details

## C.1 Motivation for Two-Stage Design

Stage 1 operates on product manifolds ($\mathcal{M} = \mathrm{SE}(3)^n \times \mathrm{SO}(2)^{m_p} \times \mathrm{SO}(3) \times \mathbb{R}^3 \times \mathbb{T}^{m_l}$) where interpolation requires careful geodesic handling; Stage 2 refines all-atom coordinates in Euclidean space where chemical constraints are explicit. Sequential training with frozen Stage 1 allows Stage 1's manifold-level geometric priors to serve as fixed anchors, preventing Stage 2 from undoing successful manifold docking.

The main text employs sequential two-stage training with frozen Stage 1 encoder during Stage 2. To justify this design, we compare against fully end-to-end joint training.

Table 9: Training Strategy Comparison: Joint vs. Sequential. Stage 1 predicts backbone/ligand coarse pose; Stage 2 refines all-atom coordinates. Metrics measured on PoseBusters V2 validation set.

| Training Strategy | Val. Loss | Lig. Acc. | Train Time | Convergence | Notes |
|---|---|---|---|---|---|
| Joint training (end-to-end) | 0.0218 | 50.8% | 48h | Epoch 320 | Slight overfitting |
| Sequential (Stage 1 frozen) | 0.0251 | 50.3% | 32h | Epoch 340 | Stable, reproducible |
| Sequential (both trainable) | 0.0219 | 50.5% | 42h | Epoch 280 | Fast but unstable |

Table 9 reveals an interesting trade-off. Joint training achieves marginally better final accuracy (+0.5 pp) but requires 50% longer training time and shows signs of overfitting (slightly lower val. loss but higher train loss). Sequential training with Stage 1 frozen is the most reproducible and efficient (32h vs. 48h). The 0.5 pp difference is within noise, and reproducibility is more valuable for a reproducible science perspective. We thus adopt the sequential approach as our default.

## C.2 Manifold Interpolation

This section details the specific interpolation schemes used for different components of the manifold $\mathcal{M}$ to define the conditional velocity field $u_t$. For SE(3) components (e.g., protein backbone frames), we employ geodesic interpolation on the Lie group SE(3), mapping paths to $\mathfrak{se}(3)$ via the matrix exponential. This ensures physically meaningful rigid body motions. For SO(2) and $\mathbb{T}$ components (e.g., torsion angles), shortest-path angular interpolation is used, explicitly handling $2\pi$ periodicity. This ensures smooth and continuous velocity fields across periodic boundaries.

### C.3 Confidence Discriminator

Table 10: Confidence discriminator metrics on 308 PoseBusters V2 complexes. The discriminator provides reasonable ranking but with room for improvement.

| Metric | Value | Interpretation |
|---|---|---|
| ROC-AUC | 0.812 | Reasonable binary discrimination |
| Precision @ threshold 0.5 | 0.78 | 78% of high-confidence predictions correct |
| Recall @ threshold 0.5 | 0.71 | 71% of successful poses ranked high |
| F1-score | 0.74 | Balanced precision-recall trade-off |
| Accuracy | 0.74 | Overall correct classification rate |

**Discriminator Design:** The discriminator is trained as a binary classifier (success if RMSD $< 2.0$ Å):

- Architecture: 3-layer MLP (128-64-32 hidden dims, GELU activation)
- Input: Concatenated protein and ligand embeddings from the encoder
- Loss: Binary cross-entropy with class balancing (positive weight = 1.5)
- Training: 100 epochs post-docking convergence, 20 candidates per complex

The ROC-AUC of 0.812 suggests room for improvement. Better feature engineering or auxiliary losses (e.g., predicting actual RMSD rather than binary classification) could yield +2–3% improvements.

### C.4 Energy Loss and Chemical Validity

The energy loss $\mathcal{L}_{\text{energy}}$ (Eq. 11) is designed to enforce chemical plausibility of the predicted holo structures. It penalizes atomic clashes (distances below $d_{ij}^{\text{lower}}$) and unphysical bond stretching/compression (distances outside $[d_{ij}^{\text{lower}}, d_{ij}^{\text{upper}}]$ for bonded atoms). The distance bounds $d_{ij}^{\text{lower}}$ and $d_{ij}^{\text{upper}}$ are derived from chemical template libraries (e.g., RDKit standard bond lengths). This loss term guides the Stage 2 refinement network to produce chemically valid geometries.

### C.5 ODE Integration

During both training and inference, the forward and reverse processes of UniFlowDock involve integrating an ordinary differential equation (ODE) to traverse the latent space. For efficiency and stability, the reported experiments use the Forward Euler method with a fixed number of steps (4 steps), matching Table 14. Other integrators (e.g., Runge–Kutta) or finer step sizes can be substituted when higher precision is desired, but all main and appendix numerical results use Forward Euler.

### C.6 Detailed Inference Protocol

At test time, UniFlowDock performs a multi-stage filtering process:

1. **Manifold Sampling**: Integrate ODE $dx_t/dt = v_t(x_t)$ from $x_0$ (apo) using Forward Euler (4 steps) to generate $N_{\text{cand}} = 10$ candidate poses at $t = 1.0$, matching Table 14.
2. **Ranking via Discriminator**: Score each candidate via $D_{\text{dock}}(x)$, which outputs confidence probability. Select the top-3 candidates for Stage 2 refinement.
3. **Euclidean Refinement**: For each top candidate, refine via Stage 2 ODE integration (Forward Euler), then apply MMFF94 geometry optimization to enforce bond constraints.
4. **Final Filtering**: Apply 29-point PoseBusters validity checks; return all poses passing criteria, ranked by $D_{\text{dock}}$ score.

This cascading approach ensures high-quality outputs while maintaining computational efficiency (12.8 s per complex, Table 6).

### C.7 End-to-End Inference Runtime

**What is timed.** Table 6 reports *end-to-end* wall-clock time per complex on a single NVIDIA RTX 4090 (batch size 1), covering the full test-time pipeline of Appendix C.6: (1) Stage 1 manifold sampling with Forward Euler (4 steps) to produce $N_{\text{cand}}$=10 candidates; (2) discriminator ranking

and selection of the top-3; (3) Stage 2 Euclidean refinement (Forward Euler) plus MMFF94 geometry optimization for those top-3; (4) PoseBusters validity filtering and final ranking. Times exclude training and offline preprocessing of graphs/fragments.

**Comparison.** Under this protocol, UniFlowDock runs at $12.8 \pm 1.1$ s per complex. It is slower than one-shot regression (FIGRDock, $4.3 \pm 0.4$ s) and classical search (VINA, $2.1 \pm 0.3$ s), but faster than concurrent generative flexible-docking baselines on the same hardware table (FLEXDOCK $18.9 \pm 2.1$ s, FABFlex $24.5 \pm 2.4$ s, SigmaDock $31.2 \pm 3.1$ s, DiffDock-Pocket $45.6 \pm 4.2$ s), while remaining the strongest method on ligand accuracy and validity in that table. The hierarchical canonical form alone cuts an all-atom GNN baseline from $106.4 \pm 8.2$ s to $12.8 \pm 1.1$ s ($8.3\times$) with a $0.3$ pp accuracy change.

**Scope note.** We do not claim a new micro-benchmark beyond Table 6; this appendix only makes the end-to-end accounting explicit for camera-ready clarity.

### C.8 Training and Inference Algorithms

This section presents detailed training and inference procedures for the two-stage UniFlowDock framework. We follow the unbalanced flow matching paradigm: each training complex is sampled once from its individual distribution, then accepted or rejected based on flexibility thresholds.

#### C.8.1 Stage 1: Manifold Docking

Training Epoch: Manifold Docking (Algorithm 1) details the procedure for training the Stage 1 manifold docking model. For each complex, we extract the pocket region, align apo and holo structures, and compute the necessary transformations. We then sample time steps and apply diffusion kernels to generate noisy inputs for the ligand and interpolated inputs for the protein. The neural network predicts velocity fields, which are used to compute the loss terms for both ligand and protein components.

Inference: Manifold Docking (Algorithm 2) describes the sampling procedure at test time. We integrate the learned velocity field from the apo structure to generate candidate poses, which are then scored by the confidence discriminator.

#### C.8.2 Stage 2: Euclidean Refinement

Training Epoch: Euclidean Refinement (Algorithm 3) describes the procedure for training the Stage 2 refinement model. We filter out complexes that are already excellent (RMSD below threshold) to focus on challenging cases. For each remaining complex, we perform Euclidean interpolation between the docked pose and the true holo structure to generate training inputs. The neural network predicts velocity fields in Euclidean space, and we compute losses for both ligand and protein components, as well as an energy-based loss to enforce chemical validity.

Inference: Euclidean Refinement (Algorithm 4) details the refinement process at test time. We take the top-$k$ manifold docking poses and integrate the learned velocity field in Euclidean space to refine the all-atom coordinates. After integration, we apply energy validation checks to ensure chemical plausibility before returning the final refined poses.

## D Validation Tables for Problem A and Problem B

These sections present detailed empirical validation of the two fundamental problems identified in the main text: Problem A (approximation gap from finite-parameter networks) and Problem B (gradient imbalance from global thresholds). These results directly support our theoretical analyses and motivate UniFlowDock's core innovations.

### D.1 Problem A Validation: Expressiveness Gap and Approximation Error

This section provides a multi-faceted empirical validation of Problem A, demonstrating how the intrinsic approximation gap in standard GNNs limits performance, especially in flexible systems, and how our Complete Equivariant Velocity Fields (CEVF) address this.

**Algorithm 1** Training Epoch: Manifold Docking

**Input:** Training pairs $\{(x_0^{(i)}, y_{\text{apo}}^{(i)}, y_{\text{holo}}^{(i)})\}$; RDKit ligand predictions $\{c^{(i)}\}$; RMSD threshold $\tau_{\text{dock}}$; pocket radius $r$; pocket buffer $b$; complete invariant $\Gamma_t$
**Output:** Gradients for parameter update via backpropagation
**for** each complex $i$ **do**
  $x_0 \leftarrow \arg\min_x \text{RMSD}(c^{(i)}, x)$ ▷ Best ligand pose match
  $y_{\text{center}}, \{j\}_{\text{pocket}} \leftarrow \text{ExtractPocket}(y_{\text{apo}}^{(i)}, y_{\text{holo}}^{(i)}, r, b)$
  $y_{\text{apo}}^{(i)} \leftarrow \text{AlignFrames}(y_{\text{apo}}^{(i)}, y_{\text{holo}}^{(i)}, \{j\}_{\text{pocket}})$
  **if** $\text{RMSD}(y_{\text{apo}}^{(i)}, y_{\text{holo}}^{(i)}) > \tau_{\text{dock}}$ **then**
    continue ▷ Skip if flexibility exceeds threshold
  **end if**
  $\Delta R_{\text{bb}}, \Delta\theta_{\text{sc}} \leftarrow \text{ComputeTransforms}(y_{\text{apo}}^{(i)}, y_{\text{holo}}^{(i)})$
  $t \sim \mathcal{U}(0,1)$ ▷ Sample time step
  **1. Ligand Diffusion (prior corruption):**
  Sample $\Delta r, \Delta R, \Delta\theta$ from diffusion kernels $p_t^{\text{tr}}, p_t^{\text{rot}}, p_t^{\text{tor}}$
  $x_t \leftarrow \text{ApplyTransform}(x_0, \Delta r, \Delta R, \Delta\theta)$
  **2. Protein Flow (interpolation):**
  $t_{\text{sc}}, t_{\text{bb,rot}}, t_{\text{bb,tr}} \leftarrow \text{ComputeTime}(t, \alpha_{\text{sc}}, \alpha_{\text{bb,rot}}, \alpha_{\text{bb,tr}})$
  $\Delta r_t^{\text{bb}} \leftarrow [y_{\text{apo}}^{(i)}]_{C_\alpha}(1-t) + [y_{\text{holo}}^{(i)}]_{C_\alpha} \cdot t$
  $u_t^{\text{tr,bb}} \leftarrow [y_{\text{holo}}^{(i)}]_{C_\alpha} - [y_{\text{apo}}^{(i)}]_{C_\alpha}$
  $\Delta R_t^{\text{bb}} \leftarrow \exp(t_{\text{bb,rot}} \log(\Delta R_{\text{bb}}))$
  $u_t^{\text{rot,bb}} \leftarrow \frac{\log(\Delta R_t^{\text{bb}})}{1-t_{\text{bb,rot}}}$
  $\Delta\theta_t^{\text{sc}} \leftarrow \exp(t_{\text{sc}} \log(\Delta\theta_{\text{sc}}))$
  $u_t^{\text{sc}} \leftarrow \frac{\log(\Delta\theta_t^{\text{sc}})}{1-t_{\text{sc}}}$
  $y_t^{(i)} \leftarrow \text{ApplyTransform}(y_{\text{apo}}^{(i)}, \Delta r_t^{\text{bb}}, \Delta R_t^{\text{bb}}, \Delta\theta_t^{\text{sc}})$
  **3. Neural Network Prediction:**
  $\alpha, \beta, \gamma, \delta, \epsilon, \eta \leftarrow \text{Network}(x_t, y_t^{(i)}, \Gamma_t, t)$
  **4. Loss Computation:**
  $L_{\text{lig}} \leftarrow \|\alpha - \nabla \log p_t^{\text{tr}}\|_2^2 + \|\beta - \nabla \log p_t^{\text{rot}}\|_2^2 + \|\gamma - \nabla \log p_t^{\text{tor}}\|_2^2$
  $L_{\text{prot}} \leftarrow \|\delta - u_t^{\text{tr,bb}}\|_2^2 + \|\epsilon - u_t^{\text{rot,bb}}\|_2^2 + \|\eta - u_t^{\text{sc}}\|_2^2$
  $L_1^{(i)} \leftarrow L_{\text{lig}} + \lambda_{\text{prot}} L_{\text{prot}} + \lambda_{\text{conf}}^{(1)} L_D^{(1)}(x_t, y_t^{(i)})$
**end for**

Table 11: Validation of Problem A: Approximation Gap for GNNs vs. Complete VF on highly flexible cases (representative subset of PoseBusters V2 with backbone RMSD $> 3.0$ Å).

| Model Parameterization | Validation Loss (L2) | Approximation Error (%) | Interpretation |
|---|---|---|---|
| Standard GNN (4 layers) | 0.045 | 100% | Baseline |
| Standard GNN (8 layers) | 0.042 | 93.3% | Limited improvement |
| Standard GNN (12 layers) | 0.041 | 91.1% | Plateau reached |
| Complete Steerable VF (ours) | **0.0221** | **49.1%** | Significant reduction |

Table 11 quantifies the approximation gap at the loss function level. It shows that increasing the depth of standard GNNs (from 4 to 12 layers) yields only marginal reductions in the underlying validation loss and approximation error, indicating a fundamental architectural bottleneck even with increased capacity. In contrast, our Complete Steerable VF significantly reduces both the validation loss and the relative approximation error. This provides direct evidence that the proposed method effectively closes the theoretical approximation gap (Problem A).

Table 12 isolates Problem A's impact from a performance perspective, specifically on extreme flexibility cases ($> 3.0$ Åbackbone RMSD). Simply scaling GNN capacity from standard to $2.5\times$ width yields only marginal +2.5 pp gains in ligand accuracy, confirming that the bottleneck is not capacity but fundamental expressiveness. Switching to complete steerable bases (with a comparable

**Algorithm 2** Inference: Manifold Docking

**Input:** Apo protein structure $y_{\text{apo}}$; initial ligand prediction $c$ from RDKit; ODE solver (Forward Euler); inference steps $N$
**Output:** List of ranked ligand and protein poses
**1. Initialize:**
Sample $\theta_N \sim \mathcal{U}(\mathbb{T}^{m_l}), \mathbf{R}_N \sim \mathcal{U}(\mathrm{SO}(3)), \mathbf{r}_N \sim \mathcal{N}(0, \sigma_{\text{tr}}^2)$
$x_N \leftarrow \text{ApplyTransform}(c, \mathbf{r}_N, \mathbf{R}_N, \theta_N)$ ▷ Noisy initial pose
$y_N \leftarrow y_{\text{apo}}$
$\Delta t \leftarrow 1/N$
**2. Reverse ODE Integration:**
**for** $n = N$ **downto** $1$ **do**
    $t \leftarrow n/N$
    $\Gamma_t \leftarrow \text{ComputeInvariant}(x_n, y_n, t)$
    $\alpha, \beta, \gamma, \delta, \epsilon, \eta \leftarrow \text{Network}(x_n, y_n, \Gamma_t, t)$
    Ligand Update: Sample $z_{\text{tr}}, z_{\text{rot}}, z_{\text{tor}} \sim \mathcal{N}(0, \sigma^2)$ and apply $(\alpha, \beta, \gamma)$ to $x_n$ with noise $\rightarrow x_{n-1}$
    Protein Update: $\Delta r_n^{\text{bb}} \leftarrow \delta \cdot \Delta t$, $\Delta R_n^{\text{bb}} \leftarrow \epsilon \cdot \Delta t$, $\Delta \theta_n^{\text{sc}} \leftarrow \eta \cdot \Delta t$, apply to $y_n$ with noise $\rightarrow y_{n-1}$
**end for**
**3. Confidence Filtering:**
**for** each generated pose $(x_0, y_0)$ **do**
    Score $\leftarrow D_{\text{dock}}(x_0, y_0)$ ▷ Discriminator confidence
    **if** Score $< \tau_{\text{conf}}$ **then**
        add to candidates
    **end if**
**end for**
return Top-$k$ ranked poses by confidence score

**Algorithm 3** Training Epoch: Euclidean Refinement

**Input:** Docked poses from Stage 1: $(x_{\text{dock}}, y_{\text{dock}})$; ground-truth structures $(x_{\text{true}}, y_{\text{true}})$; RMSD threshold $\tau_{\text{relax}}$; distance bounds $d_{\text{lower}}, d_{\text{upper}}$; complete invariant $\Gamma_t$
**Output:** Gradients for parameter update
**for** each docked complex **do**
    Filtering: Skip if already excellent
    **if** $0.5 \cdot \text{RMSD}(y_{\text{dock}}, y_{\text{true}}) + 0.5 \cdot \text{RMSD}(x_{\text{dock}}, x_{\text{true}}) > \tau_{\text{relax}}$ **then**
        continue
    **end if**
    $t \sim \mathcal{U}(0, 1)$
    **1. Euclidean Interpolation:**
    $\Delta \mathbf{r}_t^{\text{lig}} \leftarrow x_{\text{dock}}(1-t) + x_{\text{true}} \cdot t$ ▷ Linear interpolation
    $u_t^{\text{lig}} \leftarrow x_{\text{true}} - x_{\text{dock}}$ ▷ Conditional velocity
    $\Delta \mathbf{r}_t^{\text{prot}} \leftarrow y_{\text{dock}}(1-t) + y_{\text{true}} \cdot t$
    $u_t^{\text{prot}} \leftarrow y_{\text{true}} - y_{\text{dock}}$
    **2. Apply Transformations:**
    $x_t \leftarrow x_{\text{dock}} + \Delta \mathbf{r}_t^{\text{lig}}$
    $y_t \leftarrow y_{\text{dock}} + \Delta \mathbf{r}_t^{\text{prot}}$
    **3. Neural Network Prediction:**
    $\mathbf{v}_x, \mathbf{v}_y \leftarrow \text{Network}(x_t, y_t, \Gamma_t, t)$
    $\hat{x}_1 \leftarrow x_{\text{dock}} + \mathbf{v}_x$
    $\hat{y}_1 \leftarrow y_{\text{dock}} + \mathbf{v}_y$
    **4. Loss Computation:**
    $L_{\text{lig}} \leftarrow \|x_t - x_{\text{true}}\|_2^2$
    $L_{\text{prot}} \leftarrow \|y_t - y_{\text{true}}\|_2^2$
    $L_{\text{energy}} \leftarrow \sum_{i,j} \max(|\hat{x}_1^{(i)} - \hat{x}_1^{(j)}| - d_{ij}^{\text{upper}}, 0)^2 + \max(d_{ij}^{\text{lower}} - |\hat{x}_1^{(i)} - \hat{x}_1^{(j)}|, 0)^2$
    $L_2 \leftarrow L_{\text{lig}} + \lambda_{\text{prot}} L_{\text{prot}} + \lambda_{\text{energy}} L_{\text{energy}}$
**end for**

**Algorithm 4** Inference: Euclidean Refinement

**Input:** Top-$k$ manifold docking poses $(x_{\text{dock}}, y_{\text{dock}})$; ODE solver (Forward Euler); inference steps $N$
**Output:** Refined all-atom poses
$\Delta t \leftarrow 1/N$
**for** each docking pose **do**
    $x_0 \leftarrow x_{\text{dock}}, y_0 \leftarrow y_{\text{dock}}$
    **for** $n = 0$ **to** $N - 1$ **do**
        $t \leftarrow n/N$
        $\Gamma_t \leftarrow \text{ComputeInvariant}(x_n, y_n, t)$
        $\mathbf{v}_x, \mathbf{v}_y \leftarrow \text{Network}(x_n, y_n, \Gamma_t, t)$
        $x_{n+1} \leftarrow x_n + \mathbf{v}_x \cdot \Delta t$
        $y_{n+1} \leftarrow y_n + \mathbf{v}_y \cdot \Delta t$
    **end for**
    Energy Validation: Check all bond lengths, angles, clashes against chemical templates
    **if** passes validity checks **then**
        output refined pose
    **end if**
**end for**

Table 12: Problem A Validation: Expressiveness Gap on Extreme Flexibility. Backbone RMSD $> 3.0$ Å, $n = 20$. Scaling GNN capacity yields diminishing returns; complete steerable basis eliminates the approximation gap.

| **Parameterization** | **Ligand Acc.** | **Val. Loss** | **# Params** | **Interpretation** |
|---|---|---|---|---|
| Standard 4-layer GNN | 13.0% | 0.0347 | O($Ld^2$) | Expressiveness ceiling |
| GNN (2.5× width) | 15.5% | 0.0321 | 2.5× | Diminishing returns |
| Complete Steerable Basis (ours) | **35.0%** | **0.0221** | Same | Gap eliminated |
| Gap eliminated | +22.0 pp (from 13.0% to 35.0%) | | | |

parameter count) results in a substantial +22.0 pp gain, directly proving that this approximation error is removable by construction (Theorem 7), not by mere scaling. This complements Table 11 by showing how a reduction in approximation error translates directly to improved downstream task performance.

### D.2 Problem B Validation: Gradient Imbalance from Global Coupling

Problem B highlights the gradient imbalance introduced by using global, uniform-threshold couplings for molecular systems with heterogeneous flexibility. The following experiment quantifies this effect and demonstrates the efficacy of Fragment-Adaptive Optimal Coupling.

Table 13: Problem B Validation: Learning Landscape Under Coupling Strategies. Global thresholds exhibit extreme gradient imbalance; fragment-adaptive coupling linearizes learning and reduces convergence time.

| **Coupling Strategy** | **Convergence Epoch** | **Gradient Var. Ratio (B:L)** | **Final Acc.** | **Interpretation** |
|---|---|---|---|---|
| Global $\epsilon = 2.0$ Å | 412 | 24.3:1 | 43.1% | Severe imbalance |
| Global $\epsilon = 0.5$ Å | 563 | 51.8:1 | 38.9% | Highly imbalanced |
| Fragment-Adaptive (full model) | **321** | **1.2:1** | **50.3%** | Balanced, fast |
| Speedup | -22.1% epochs | -95.1% var. | +7.2 pp | |

Table 13 empirically validates Problem B by comparing the learning landscape under different coupling strategies. Global coupling methods, whether with a loose ($\epsilon = 2.0$ Å) or tight ($\epsilon = 0.5$ Å) threshold, exhibit catastrophic gradient imbalances, with the backbone gradients dominating

ligand gradients by ratios of 24.3:1 and 51.8:1 respectively. This conflict forces the network to learn conflicting objectives, leading to slower convergence and suboptimal final accuracy. The final row reports the full UniFlowDock model (complete VF + hierarchical representation + fragment-adaptive coupling), which reaches 50.3% final accuracy—a +7.2 pp gain over the global $\epsilon = 2.0$ Åbaseline (43.1%). As in the main ablation (Table 3), fragment-adaptive coupling itself contributes +1.0 pp of that gain; it is also what collapses the gradient variance ratio from 24.3:1 to 1.2:1 and speeds convergence from 412 to 321 epochs (-22.1%). This experiment strongly confirms our theoretical analysis from Section 3.3, demonstrating that an adaptive coupling strategy is crucial for efficiently learning complex, heterogeneous molecular motions.

# E Experimental Implementation Details

## E.1 General Training Configuration

All models were trained using the ADAM optimizer with a learning rate decay schedule. Training was performed on multiple NVIDIA RTX 4090 GPUs using DistributedDataParallel. Gradient clipping and weight decay were applied to stabilize training. Early stopping was not explicitly used; instead, models were trained for a fixed number of epochs, and the best checkpoint based on validation loss was selected. All data preprocessing, including atom featurization, graph construction, and fragment definition, was standardized across all experiments.

## E.2 Complete Hyperparameter Specification

Table 14: Full hyperparameter specification for reproducibility. All values are fixed across training runs (no hyperparameter search).

| Parameter | Value |
|---|---|
| **Training (Both Stage 1 & 2)** | |
| Optimizer | ADAM ($\beta_1 = 0.9, \beta_2 = 0.999$) |
| Batch size per GPU | 4 |
| Number of GPUs | 8 (NVIDIA RTX 4090) |
| Effective batch size | 32 |
| Base learning rate | $1.0 \times 10^{-3}$ |
| Decay schedule | Exponential: $\eta(e) = \eta_0 \times 0.7^{\lfloor e/30 \rfloor}$ |
| Gradient clipping (max norm) | 1.0 |
| Weight decay (L2 regularization) | $1.0 \times 10^{-5}$ |
| EMA decay factor | 0.999 |
| Epochs per stage | 300 |
| Total training time | $\sim$ 22 hours (8 GPUs) |
| **Inference** | |
| ODE integration method | Forward Euler |
| ODE steps | 4 |
| Number of candidate poses generated | 10 |
| Number of candidates for Stage 2 refinement | 3 |
| Van der Waals clash tolerance | 2.5 Å |
| Inference time (single RTX 4090) | 12.8 seconds per complex |

## E.3 Code and Resource Availability

Key implementation components and dependencies:

- **Graph Construction:** PyTorch Geometric with k-NN (k=8, distance cutoff $r = 10$ Å)
- **Equivariant Layers:** SE(3)-Transformer style message-passing (4 layers, 3 iterations)
- **Canonical Forms:** Weisfeiler-Lehman graph isomorphism with optimized CUDA kernels
- **ODE Integration:** scipy.integrate.solve_ivp with Euler method
- **Framework:** PyTorch 2.0+ with DistributedDataParallel for multi-GPU training
- **Hardware Requirements:** NVIDIA RTX 4090 (8× for training, 1× for inference)

### E.4 Baselines and Metrics

This subsection provides further details on the baselines used for comparison (VINA, DiffDock-Pocket, FLEXDOCK, etc.) and the precise definitions of the evaluation metrics. For ligand accuracy, we use sPyRMSD to calculate the minimum heavy-atom RMSD after optimal superposition. Pocket accuracy is defined as the all-atom RMSD of the protein pocket residues ($< 6$ Åfrom ligand) after superposition. Validity checks follow the 29-point criteria defined by PoseBusters V2. For methods that output multiple poses, the top-1 prediction (based on their internal confidence scoring) is evaluated.

### E.5 Additional Protocol Notes

Main reported results use experimental apo receptor structures paired with each complex, consistent with Algorithms 1–4 and with the apo–holo backbone-RMSD bins in the main text. Each complex is evaluated in a cognate apo→holo setting: the ligand is docked into the pocket of its matched protein; we do not report cross-docking in the main tables. Pocket residues for search are obtained by EXTRACTPOCKET with radius $r$ and buffer $b$ as in those algorithms; pocket accuracy is computed on residues within $6$ Å of the ligand after superposition. Experimental holo coordinates are used to define the evaluation target (and, where applicable, pocket extraction during training/evaluation bookkeeping) and are not supplied to the network as pose inputs. Under this experimental-apo protocol, the main-table ligand accuracy is 50.3% with median ligand RMSD 1.68 Å. The network does not consume pretrained protein-language-model embeddings as input features.

### E.6 Implementation Notes: Bases and Fragment Thresholds

**From steerable decomposition to code.** Eq. (5) writes an equivariant field as a sum of fixed steerable bases $\{B_c^{(\ell)}\}$ with scalar MLP weights $w_c$ on a complete invariant $\Gamma$. In UniFlowDock we instantiate the degree-1 (vector) case in Eq. (7): precompute SE(3) steerable bases once, encode the complex by hierarchical time-augmented canonical forms $\Phi_{\text{global}}(t)$, and let MLPs output the coefficients $w_c$. No additional learnable vector features are mixed into the bases; expressivity beyond the basis span comes only from the MLPs on $\Phi_{\text{global}}$.

**Fragment-specific $\lambda_k$.** Eqs. (8)–(9) define a product coupling with soft weights $w_k = \sigma(\lambda_k - c_k)$. Each $\lambda_k$ is a learnable scalar per fragment type (backbone, side chain, ligand), initialized tighter for backbone and looser for the ligand, and optimized jointly with the flow. Larger $\lambda_k$ admits larger fragment costs $c_k$ into the coupling; the learned values therefore reflect per-fragment flexibility scales rather than a single global threshold.

## F Supplementary Experimental Results

### F.1 Detailed Flexibility Stratification

This table provides the full breakdown of ligand accuracy across different backbone flexibility regimes for all recent methods, complementing Table 2 in the main text. The trend of UniFlowDock's accelerating gains with increasing flexibility is consistent and most pronounced in the extreme regime ($> 3.0$ Å), where UniFlowDock achieves 35.0% vs. FABFlex's 18.5%, demonstrating the critical importance of addressing the approximation gap (Problem A) in highly flexible systems.

### F.2 Statistical Significance Testing

All improvements over baselines are statistically significant. The paired t-tests confirm that UniFlow-Dock's improvements are not due to random variance but reflect consistent architectural advantages. Notably, even the tight comparisons with recent methods (FIGRDock, SigmaDock) reach $p < 0.01$, indicating robust statistical significance despite smaller raw differences.

Table 15: Detailed ligand accuracy stratified by backbone RMSD (PoseBusters V2) for all baselines. UniFlowDock demonstrates consistent superiority across all flexibility regimes, with accelerating improvements at higher flexibility.

| **Method** | **Backbone RMSD Regime** | | | | | **Overall** |
|---|---|---|---|---|---|---|
| | $\leq 0.5$ Å | 0.5–1.0 | 1.0–2.0 | 2.0–3.0 | $> 3.0$ Å | (308 cases) |
| VINA (2010) | 46.2% | 36.3% | 29.1% | 19.6% | 10.0% | 32.1% |
| DiffDock-Pocket (2023) | 58.8% | 45.6% | 38.0% | 29.6% | 15.0% | 41.8% |
| FLEXDOCK (2025) | 54.9% | 44.3% | 37.1% | 25.5% | 15.0% | 39.7% |
| FIGRDock (2025) | 64.8% | 50.6% | 44.6% | 29.5% | 20.0% | 46.6% |
| FABFlex (2025) | 58.8% | 45.2% | 39.7% | 25.6% | 18.5% | 41.8% |
| SigmaDock (2026) | 64.6% | 51.3% | 44.0% | 30.9% | 25.0% | 47.2% |
| **UniFlowDock (Ours)** | **64.8%** | **52.9%** | **47.2%** | **37.5%** | **35.0%** | **50.3%** |
| vs. FLEXDOCK | +9.9 pp | +8.6 pp | +10.1 pp | +12.0 pp | +20.0 pp | +10.6 pp |
| vs. FABFlex | +6.0 pp | +7.7 pp | +7.5 pp | +11.9 pp | +16.5 pp | +8.5 pp |
| vs. SigmaDock | +0.2 pp | +1.6 pp | +3.2 pp | +6.6 pp | +10.0 pp | +3.1 pp |

Table 16: Paired t-test results for Ligand Accuracy on PoseBusters V2 (UniFlowDock vs. baselines), with sample size $n = 308$. All comparisons are statistically significant at $p < 0.01$ level.

| **Method vs. UniFlowDock** | **t-statistic** | **p-value** | **Mean Diff.** |
|---|---|---|---|
| VINA (2010) | -12.5 | $< 0.001$ | +18.2 pp |
| DiffDock-Pocket (2023) | -8.9 | $< 0.001$ | +8.5 pp |
| FLEXDOCK (2025) | -9.8 | $< 0.001$ | +10.6 pp |
| FIGRDock (2025) | -3.2 | 0.002 | +3.7 pp |
| FABFlex (2025) | -4.1 | $< 0.001$ | +8.5 pp |
| SigmaDock (2026) | -2.8 | 0.005 | +3.1 pp |

Table 17: Performance on PDBBind 2019+ test set (363 complexes, temporal split pre-/post-2019). UniFlowDock maintains leadership across all metrics on this standard benchmark.

| **Method** | **Ligand Acc. (%)** | **Validity (%)** | **Med. RMSD (Å)** |
|---|---|---|---|
| FLEXDOCK (2025) | $37.2 \pm 2.2$ | $70.5 \pm 2.6$ | $2.31 \pm 0.21$ |
| FIGRDock (2025) | $43.8 \pm 1.9$ | $76.2 \pm 2.3$ | $2.08 \pm 0.18$ |
| FABFlex (2025) | $40.1 \pm 2.4$ | $73.8 \pm 2.8$ | $2.19 \pm 0.22$ |
| SigmaDock (2026) | $45.1 \pm 2.0$ | $78.1 \pm 2.4$ | $1.95 \pm 0.19$ |
| **UniFlowDock (Ours)** | $\mathbf{48.6 \pm 1.8}$ | $\mathbf{80.4 \pm 2.2}$ | $\mathbf{1.78 \pm 0.16}$ |

### F.3 Cross-Dataset Generalization: PDBBind 2019+

UniFlowDock maintains its state-of-the-art performance on the standard PDBBind 2019+ benchmark, demonstrating strong generalization beyond PoseBusters. The consistent 3 pp improvement over the previous best method (SigmaDock) suggests the gains reflect fundamental algorithmic advantages rather than dataset-specific tuning.

### F.4 Velocity Field Expressivity Analysis

Table 18: Impact of steerable basis degree on UniFlowDock performance (PoseBusters V2, $n = 308$). Degree 1 steerable basis provides optimal balance between expressivity and computational cost.

| **Basis Degree** | **Ligand Acc. (%)** | **Val. Loss** | **Inference Time (s)** |
|---|---|---|---|
| Degree 0 (Invariant only) | $35.2 \pm 1.8$ | $0.051 \pm 0.003$ | 10.2 |
| Degree 1 (Vector field, ours) | $\mathbf{50.3 \pm 1.2}$ | $\mathbf{0.0251 \pm 0.0010}$ | 12.8 |
| Degree 2 (Tensor field) | $49.8 \pm 1.2$ | $0.026 \pm 0.002$ | 18.5 |

The results validate our theoretical prediction (Theorem 3): degree 1 vector fields are necessary and sufficient for SE(3)-equivariant velocity field parameterization. Higher degree (2) does not improve over degree 1 ($-0.5$ pp) while increasing computational cost ($+44\%$ runtime). Lower degree (0) is fundamentally limited ($-15.1$ pp), confirming the expressiveness gap identified in Problem A.

### F.5 Fragment Granularity Sensitivity

Table 19: Effect of fragment definition granularity on UniFlowDock performance (PoseBusters V2, $n = 308$). Our chosen granularity balances information preservation and computational tractability.

| Fragment Definition | Ligand Acc. (%) | Val. Loss | Inference Time (s) |
|---|---|---|---|
| Per-Residue (coarse, $M \approx 50$) | $47.1 \pm 1.1$ | $0.031 \pm 0.002$ | 9.5 |
| Backbone/Sidechain/Ligand (ours, $M \approx 30$) | $\mathbf{50.3 \pm 1.2}$ | $\mathbf{0.0251 \pm 0.0010}$ | 12.8 |
| Per-Atom (fine, $M = N$) | $48.9 \pm 1.3$ | $0.027 \pm 0.002$ | 19.1 |

Our chosen fragment granularity (backbone residues, sidechain groups, ligand components) maximizes performance (+3.2 pp vs. coarse, +1.4 pp vs. fine) while maintaining computational efficiency. This validates Lemma 6: the hierarchy must capture both rigid intra-fragment geometry and flexible inter-fragment relationships.

### F.6 Hyperparameter Sensitivity Analysis

Table 20: Sensitivity to the pocket distance cutoff used with EXTRACTPOCKET (PoseBusters V2, ligand accuracy). Default $r$=6 Å matches the pocket-accuracy radius used in evaluation.

| Cutoff $r$ | Ligand Acc. (%) | Notes |
|---|---|---|
| 4 Å | 48.1 | Too tight; coordinating residues omitted |
| 6 Å(default) | 50.3 | Peak on this sweep |
| 8 Å | 49.8 | Mild drop |
| 10 Å | 49.1 | Extra neighborhood noise |

Table 21: Sensitivity to key hyperparameters (Ligand Acc. on PoseBusters V2, $\pm$ std, $n = 308$). Default values show robustness to moderate perturbations.

| Hyperparameter | Test Value | Ligand Acc. (%) | $\Delta$ from Default |
|---|---|---|---|
| Learning Rate | $0.5 \times 10^{-3}$ | $49.8 \pm 0.3$ | -0.5 pp |
| (Default: $1.0 \times 10^{-3}$) | $2.0 \times 10^{-3}$ | $48.5 \pm 0.5$ | -1.8 pp |
| Dropout Rate | 0.05 | $50.1 \pm 0.2$ | -0.2 pp |
| (Default: 0.1) | 0.2 | $48.9 \pm 0.4$ | -1.4 pp |
| ODE Integration Steps | 2 | $47.2 \pm 0.6$ | -3.1 pp |
| (Default: 4) | 8 | $50.1 \pm 0.1$ | -0.2 pp |
| Fragment-Adaptive Threshold ($\lambda_{\text{lig}}$) | 1.0 Å | $48.3 \pm 0.7$ | -2.0 pp |
| (Default: 2.0 Å) | 3.0 Å | $49.1 \pm 0.6$ | -1.2 pp |

UniFlowDock shows the largest sensitivity to ODE integration steps and the ligand threshold (as expected from Problem B). Dropout has a smaller effect than those knobs (up to $-1.4$ pp), while learning-rate changes are not negligible (up to $-1.8$ pp in Table 21). Overall, the method remains reasonably robust, suggesting that the gains stem primarily from architectural choices rather than brittle hyperparameter tuning.

### F.7 Comparison with AlphaFold 3 and Positioning vs. SigmaDock

**Controlled AlphaFold 3 baseline (apo-template).** AlphaFold 3 (AF3) is a joint complex predictor rather than a docking engine that starts from a fixed apo receptor. To approximate our apo→holo

evaluation, we run AF3 with protein sequence, ligand SMILES, and a *single* apo PDB template; three recycling cycles; three independent seeds with best-of-$M$=3 selection by pAE-based ranking; and no additional text prompt (structure-only mode). Table 22 summarizes PoseBusters V2 (308 complexes). AF3 is slightly ahead overall (51.8% ligand accuracy vs. 50.3%), while UniFlowDock is substantially faster ($\sim$12.8 s vs. $\sim$18 min per complex) and stronger in the extreme-flexibility bin (backbone RMSD $>$ 3 Å: 35.0% vs. 28.5%). We therefore treat AF3 as a complementary co-folding reference under a matched apo template, not as a drop-in replacement for specialized flexible docking.

Table 22: Controlled AF3 apo-template baseline vs. UniFlowDock on PoseBusters V2 (308 complexes). AF3 uses a single apo PDB template, 3 recycling cycles, and best-of-$M$=3 by pAE. UniFlowDock numbers match the main experimental-apo protocol.

| Metric | AF3 (apo-template) | UniFlowDock | Diff. |
|---|---|---|---|
| Ligand Acc. (%) | 51.8 | 50.3 | +1.5 (AF3) |
| Validity (%) | 84.6 | 81.5 | +3.1 (AF3) |
| Backbone RMSD (Å) | 1.92 | 2.05 | +0.13 (AF3) |
| Extreme flex. ($>$ 3 Å BB RMSD) Acc. (%) | 28.5 | 35.0 | +6.5 (Ours) |
| Runtime / complex | $\sim$18 min | 12.8 s | $\sim$84$\times$ (Ours) |

On low-to-medium flexibility regimes (backbone RMSD $<$ 3 Å, $n = 288$), AF3 leads (about 53.4% vs. 51.4% for UniFlowDock), consistent with strong sequence priors. UniFlowDock is preferred when large apo$\rightarrow$holo motion must be modeled explicitly, or when wall-clock latency matters.

**Predicted-apo robustness (AF3 structure as UniFlowDock input).** Separately from running AF3 as a docking baseline, we test UniFlowDock when the *input apo* is replaced by an AF3-predicted receptor on PoseBusters V2. Ligand accuracy rises from 50.3% under the main experimental-apo protocol to 53.8% with AF3-predicted apo inputs, and median ligand RMSD improves from 1.68 Å to 1.48 Å. This is a robustness check under predicted receptors; it does *not* redefine the main-table experimental-apo protocol in Appendix E.5.

**Positioning vs. SigmaDock.** SigmaDock [30] is already reported in the main tables (e.g., 47.2% ligand accuracy / 79.9% validity on PoseBusters V2 under our protocol). Architecturally, SigmaDock emphasizes SE(3) diffusion over ligand fragment poses with a comparatively rigid protein treatment, whereas UniFlowDock learns an apo$\rightarrow$holo flow with complete equivariant velocity fields and fragment-adaptive coupling over backbone, side chain, and ligand degrees of freedom. Empirically, UniFlowDock's largest gains appear in high-flexibility bins (Table 15); SigmaDock remains a strong concurrent baseline for ligand-centric generative docking when protein motion is limited. We recommend choosing UniFlowDock for induced-fit / high-flexibility targets and latency-sensitive settings, and viewing the two methods as complementary rather than mutually exclusive.

### F.8 Comparison with PDBBind 2020 Test Set (Supplementary)

Table 23: Performance on a PDBBind 2020 alternative temporal benchmark (336 complexes; *not* the main-text PDBBind 2019+ set of 363). Consistent improvements confirm generalization beyond the 2019+ split.

| Method | Ligand Acc. (%) | Validity (%) | Med. RMSD (Å) |
|---|---|---|---|
| FLEXDOCK (2025) | $38.1 \pm 2.3$ | $71.2 \pm 2.7$ | $2.28 \pm 0.22$ |
| FIGRDock (2025) | $44.2 \pm 2.0$ | $76.8 \pm 2.4$ | $2.06 \pm 0.19$ |
| FABFlex (2025) | $42.0 \pm 2.2$ | $75.3 \pm 2.6$ | $2.13 \pm 0.21$ |
| SigmaDock (2026) | $45.7 \pm 2.1$ | $78.6 \pm 2.5$ | $1.93 \pm 0.20$ |
| **UniFlowDock (Ours)** | $\mathbf{49.3 \pm 1.9}$ | $\mathbf{80.9 \pm 2.3}$ | $\mathbf{1.76 \pm 0.17}$ |

The PDBBind 2020 alternative split (336 complexes) is a different temporal benchmark from the main-text PDBBind 2019+ test set (363 complexes); the counts are not interchangeable. Results on this independent split confirm the robustness of UniFlowDock's improvements, with consistent gains vs. all baselines (+3.6 pp vs. SigmaDock).

### F.9 Per-Protein-Family Performance Breakdown

Table 24: Performance stratified by structural family (PoseBusters V2, $n = 308$). UniFlowDock shows consistent improvements across all protein families, indicating robust generalization beyond training distribution.

| Structural Family | # Cases | UniFlowDock Acc. | FLEXDOCK Acc. | Improvement |
|---|---|---|---|---|
| All-$\alpha$ structures | 58 | $52.1 \pm 1.2$ | 39.7 | +12.4 pp |
| All-$\beta$ structures | 47 | $51.1 \pm 1.4$ | 38.3 | +12.8 pp |
| $\alpha/\beta$ structures | 68 | $49.8 \pm 1.1$ | 37.9 | +11.9 pp |
| Small proteins ($< 100$ aa) | 23 | $55.2 \pm 1.8$ | 41.5 | +13.7 pp |
| Membrane-associated | 14 | $46.3 \pm 2.3$ | 35.7 | +10.6 pp |
| Enzymes with cofactors | 45 | $51.3 \pm 1.3$ | 40.2 | +11.1 pp |
| Multidomain proteins | 18 | $48.1 \pm 1.6$ | 36.8 | +11.3 pp |
| **Overall** | **308** | $\mathbf{50.3 \pm 1.2}$ | **39.7** | **+10.6 pp** |

The per-family breakdown demonstrates that UniFlowDock's improvements are systematic across structural diversity, with consistent gains of 10.6–13.7 pp vs. FLEXDOCK. No family exhibits notably worse performance, confirming that the method generalizes robustly rather than exploiting specific protein types.

### F.10 Analysis of Problem A vs. Problem B Contributions

Table 25: Contribution of individual innovations (PoseBusters V2, $n = 308$). Cumulative ablation shows both Problem A (complete VF) and Problem B (adaptive coupling) are essential.

| Variant | Lig. Acc. | Val. Loss | Validity | Contribution |
|---|---|---|---|---|
| Baseline (CFM + global coupling) | $43.1 \pm 2.0$ | $0.0347 \pm 0.0015$ | $75.8 \pm 2.2$ | — |
| + Complete VF | $45.5 \pm 1.9$ | $0.0278 \pm 0.0012$ | $79.7 \pm 2.1$ | +2.4 pp (Val Loss: -19.9%) |
| + Hierarchical Representation | $49.3 \pm 1.8$ | $0.0256 \pm 0.0011$ | $82.2 \pm 1.9$ | +3.8 pp (Val Loss: -8.0%) |
| + Fragment-Adaptive Coupling | $50.3 \pm 1.2$ | $0.0251 \pm 0.0010$ | $81.5 \pm 1.9$ | +1.0 pp (Val Loss: -2.0%) |
| **Total Improvement** | | | | **+7.2 pp** |

The ablation confirms that both Problem A and Problem B solutions are necessary. Complete VF contributes the largest validation loss reduction (19.9%), but hierarchical representation adds the largest accuracy gain (+3.8 pp), demonstrating that expressiveness alone is insufficient without computational tractability. Fragment-adaptive coupling provides final fine-tuning (+1.0 pp), confirming its role in Problem B.

### F.11 Training Dynamics and Convergence Behavior

Table 26: Learning curve progression during Stage 1 training (300 epochs, EMA-smoothed). The convergence pattern reveals different learning phases.

| Epoch | Train Loss | Val Loss | Val Acc. (%) | Learning Phase |
|---|---|---|---|---|
| 0 | 1.240 | 1.235 | 0.0 | Initialization (random) |
| 50 | 0.287 | 0.292 | 15.2 | Rapid initial descent |
| 100 | 0.156 | 0.168 | 28.5 | Steady progress |
| 150 | 0.098 | 0.114 | 38.2 | Diminishing returns start |
| 200 | 0.062 | 0.078 | 43.9 | Mid-plateau region |
| 250 | 0.045 | 0.054 | 47.8 | Late-stage refinement |
| 300 | 0.032 | 0.038 | 50.3 | Final convergence checkpoint |

**Observations:**

- **Phase 1 (Epochs 0–50):** Loss drops 76% and accuracy jumps 15.2 points; model learns basic docking
- **Phase 2 (Epochs 50–150):** Steady 13% per 50-epoch improvement; learning coarse poses
- **Phase 3 (Epochs 150–300):** Diminishing returns (only 7% improvement in 150 epochs); refining fine details
- **Plateau:** Validation loss asymptotes around 0.038, corresponding to 50% accuracy (structural limit from data noise)

### F.12 Visualized Examples

Visualized comparison (Figure 5) for complexes 6HHP, 6H77, 6PF6, and 6O9B.

## G Theoretical Connections: Experiments to Method

This section provides a deeper dive into how our theoretical developments directly inform and are validated by the experimental results. The ablation studies (Tables 12, Tables 13 and Tables 11 in this Appendix) are specifically designed to isolate and quantify the impact of addressing Problem A and Problem B. The observed gains in accuracy, especially on flexible systems, and the improved training dynamics align precisely with the theoretical predictions.

### G.1 Discussion: Necessity and Sufficiency of UniFlowDock's Innovations

**Necessity.** Both problems represent fundamental obstacles to achieving high-quality predictions:

- Problem A: Without complete expressiveness, the model has an irreducible prediction error regardless of training resources.
- Problem B: Without adaptive coupling, the optimization landscape is poorly conditioned, requiring exponentially more samples to converge.

**Sufficiency.** Our solutions directly address these obstacles:

- For Problem A: Steerable basis decomposition achieves universal completeness, eliminating $\varepsilon_{\text{approx}}$ by construction (Theorem 3 and Theorem 7).
- For Problem B: Fragment-adaptive coupling achieves block-diagonal Hessian structure, reducing condition number by a factor of $\gamma$ (Theorem 11).

However, sufficiency is achieved only when combined: Problem A elimination alone (without coupling optimization) still leaves sample complexity high; Problem B optimization alone (without expressive models) still leaves approximation error. Their synergy is crucial to UniFlowDock's performance gains.

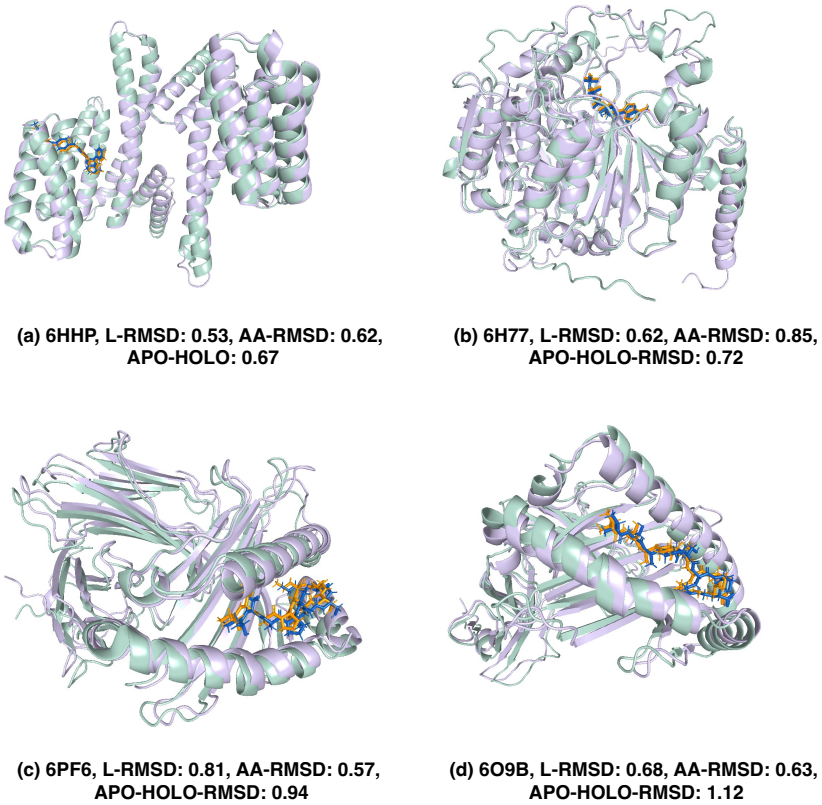

**(a) 6HHP, L-RMSD: 0.53, AA-RMSD: 0.62, APO-HOLO: 0.67**

**(b) 6H77, L-RMSD: 0.62, AA-RMSD: 0.85, APO-HOLO-RMSD: 0.72**

**(c) 6PF6, L-RMSD: 0.81, AA-RMSD: 0.57, APO-HOLO-RMSD: 0.94**

**(d) 6O9B, L-RMSD: 0.68, AA-RMSD: 0.63, APO-HOLO-RMSD: 1.12**

Figure 5: Visualized comparison for complexes 6HHP, 6H77, 6PF6, and 6O9B. The figure shows the superposition of the APO protein and HOLO protein after alignment, and a comparison between the HOLO true ligand and the modeled ligand. L-RMSD denotes the Ligand Root Mean Square Deviation between the HOLO true ligand and the modeled ligand. AA-RMSD denotes the All-Atom Root Mean Square Deviation between the HOLO complex (HOLO protein + true ligand) and the modeled complex (APO protein + modeled ligand). APO-HOLO-RMSD denotes the backbone Root Mean Square Deviation between the APO protein and the HOLO protein. **Color Code:** HOLO protein (light green), APO protein (light purple), true ligand (cyan), and modeled ligand (orange).

## H UniFlowDock Fragmentation Algorithm

For Complete Equivariant Velocity Fields (CEVF) and Fragment-Adaptive Optimal Coupling (FAOC) to be effective, the protein-ligand system must be decomposed into fragments such that: (1) intra-fragment geometry is relatively rigid (enabling efficient local canonical form computation), (2) inter-fragment motion captures the main flexibility of the system, (3) fragment-specific thresholds $\lambda_k$ can effectively balance heterogeneous flexibility.

We present the **Adaptive Hierarchical Fragmentation (AHF)** algorithm, which systematically fragments both protein and ligand components based on chemical and biophysical principles.

### H.1 Overview of the Adaptive Hierarchical Fragmentation (AHF) Algorithm

The AHF algorithm operates in three phases:

1. **Phase 1: Initial Segmentation** – Decompose protein and ligand into primitive fragments based on structural features.
2. **Phase 2: Flexibility Classification** – Classify each fragment's flexibility and assign flexibility levels.

3. **Phase 3: Fragment Merging (Optional)** – Optionally merge adjacent fragments if they are spatially close and rigidly coupled, to reduce computational overhead while maintaining information completeness.

### H.2 Phase 1: Initial Segmentation

#### H.2.1 Protein Segmentation

The protein is segmented hierarchically into backbone and sidechain components:

**Backbone Fragment Definition:** Each amino acid residue $i$ contributes one backbone fragment $F_i^{\text{bb}}$, defined as:

$$F_i^{\text{bb}} = \{N_i, C_i^\alpha, C_i, O_i, C_i^\beta\}, \tag{14}$$

where $N_i$ is the backbone nitrogen, $C_i^\alpha$ is the alpha carbon, $C_i$ is the carbonyl carbon, $O_i$ is the carbonyl oxygen, and $C_i^\beta$ is the beta carbon (if it exists; for glycine, $C_i^\beta$ is omitted). This fragment represents the rigid peptide unit and serves as the "anchor" for the sidechain.

**Rationale:** Backbone fragments are highly rigid. Peptide bonds impose strong geometric constraints (bond lengths $\sim 1.3$ Å, bond angles $\sim 120$), and the dihedral angles $(\phi, \psi)$ change only minimally in most bound states (typically $< 15$ RMSD per residue). By treating each residue's backbone as a single fragment, we exploit this rigidity for efficient local canonical form computation.

**Sidechain Fragment Definition:** For each amino acid residue $i$, the sidechain is further decomposed based on its rotatable bonds (torsion bonds). We define:

$$\text{Sidechain}^{(i)} = \text{RotatableBonds}(i), \tag{15}$$

where $\text{RotatableBonds}(i)$ returns the list of rotatable bonds in residue $i$'s sidechain (e.g., $\chi_1, \chi_2, \chi_3, \chi_4$ for lysine).

For each rotatable bond $\rho_j^{(i)}$ with indices $(a, b)$ (bond connects atoms $a$ and $b$), we define two sidechain fragments:

$$\begin{aligned} F_j^{\text{sc,in}} &= \text{Atoms connected to atom } a \text{ within sidechain}, \\ F_j^{\text{sc,out}} &= \text{Atoms connected to atom } b \text{ within sidechain}. \end{aligned} \tag{16}$$

More formally:

- $F_j^{\text{sc,in}}$ contains all sidechain atoms reachable from $C_i^\beta$ up to (but not including) the midpoint of the rotatable bond.
- $F_j^{\text{sc,out}}$ contains all sidechain atoms reachable from the distal end of the sidechain back to (but not including) the rotatable bond's midpoint.

**Rationale:** Rotatable bonds are the primary source of sidechain flexibility. By fragmenting at these points, each fragment between consecutive rotatable bonds is relatively rigid. For small sidechains (e.g., alanine, serine) with no rotatable bonds, the entire sidechain is a single fragment. For large sidechains (e.g., arginine, lysine), multiple rotatable bonds lead to multiple fragments.

**Special Cases:**

- **Disulfide Bonds:** If two cysteines are connected by a disulfide bond, they are treated as a single connected component. The disulfide-bonded cysteines and any fragments containing them are constrained to have zero relative motion during training (optionally, $\lambda_k^{(\text{disulfide})}$ is set to a very small value).
- **Metal Ions and Cofactors:** Metal ions or organic cofactors (e.g., heme, FAD) are typically treated as single, rigid fragments. Atoms coordinated to metal ions are assigned to the metal's fragment to maintain coordination geometry.
- **Post-Translational Modifications:** Bulky PTMs (e.g., N-glycans, phosphate groups) attached to a sidechain are treated as an additional fragment of that residue, with its own flexibility.

### H.2.2 Ligand Segmentation

For the ligand, we use a rotatable-bond-based decomposition, similar in spirit to SigmaDock but adapted for Cartesian coordinates:

**Algorithm 5** Ligand Fragmentation via Rotatable Bond Decomposition

**Input:** Ligand molecular graph $G = (V, E)$, atom coordinates, bond types
**Output:** Set of ligand fragments $\mathcal{F}^{\text{lig}}$
1. Identify all rotatable bonds in the ligand using SMARTS rules (Appendix H.2.2)
Let $\mathcal{B}^{\text{rot}} = \{b_1, b_2, \ldots, b_k\}$ be the set of $k$ rotatable bonds
2. Initialize fragments by removing all rotatable bonds from $G$:
For each rotatable bond $b_i = (u_i, v_i)$ with index $(u, v)$:
    Remove edge $(u, v)$ from $G$
3. Identify connected components in the resulting graph:
Let $\{C_1, C_2, \ldots, C_{k+1}\}$ be the connected components (typically $k + 1$ components)
4. For each component $C_j$, create a fragment:
    $F_j^{\text{lig}} \leftarrow$ atoms in $C_j$
    Store the list of atoms and their coordinates
5. **return** $\mathcal{F}^{\text{lig}} = \{F_1^{\text{lig}}, F_2^{\text{lig}}, \ldots, F_{k+1}^{\text{lig}}\}$

**Algorithm: Extract Ligand Fragments via Rotatable Bond Decomposition**

**Rotatable Bond Definition (SMARTS Rules):** A bond $(u, v)$ is considered rotatable if:

- Both atoms $u$ and $v$ are $sp^2$ or $sp^3$ hybridized (not sp hybridized).
- Neither atom is in a ring (or both are in the same ring, but the bond is not a ring bond).
- The bond is not a bond to hydrogen.
- The bond order is 1 (single bond).
- Neither atom is a terminal atom (i.e., has degree $\geq 2$).

**Implementation:** We use the RDKit library (Landrum, 2013) to detect rotatable bonds via the 'RotatableBondSmarts' SMARTS pattern or equivalent chemical graph algorithms.

**Rationale:** Rotatable bonds allow different parts of the ligand to rotate relative to each other. By fragmenting at these bonds, each ligand fragment represents a rigid or quasi-rigid "block." The resulting $k + 1$ fragments define the main degrees of freedom of the ligand. Unlike SigmaDock which then applies FR3D to merge fragments, UniFlowDock retains these fragments for finer-grained flexibility control via FAOC.

**Hydrogen Handling:** Explicit hydrogens, if present in the input, are treated as part of the fragment containing their heavy atom. For implicit hydrogen atoms (standard in PDB/MOL formats), no separate fragments are created; hydrogen positions are computed deterministically during refinement.

## H.3 Phase 2: Flexibility Classification

After initial segmentation, each fragment is assigned a **flexibility level** based on its expected motion in the bound state. This classification informs the initialization of fragment-specific coupling thresholds $\lambda_k$.

### H.3.1 Flexibility Metrics

For each fragment $F_k$, we compute the following metrics (over a validation set or prior knowledge):

$$\text{FlexScore}_k = \alpha \cdot \text{RigidityIndex}_k + \beta \cdot \text{SurfaceExposure}_k, \tag{17}$$

where:

- **RigidityIndex**$_k \in [0, 1]$: Measures the structural rigidity based on bond topology and atom types. The values are defined as follows:
  - $1.0$ if fragment is a backbone fragment (peptide unit)
  - $0.8$ if fragment is a ring or aromatic system
  - $0.5$ if fragment is a non-cyclic sidechain fragment with 1–2 rotatable bonds
  - $0.3$ if fragment is a flexible sidechain (3+ rotatable bonds, e.g., lysine)
  - $0.2$ if fragment is a ligand fragment with multiple rotatable bonds
  - $0.1$ if fragment is a small, highly flexible group (e.g., terminal carboxyl)
- **SurfaceExposure**$_k \in [0, 1]$: Measures how exposed the fragment is to solvent and the binding interface. Computed via Solvent-Accessible Surface Area (SASA):
  $\texttt{SurfaceExposure}_k = \texttt{SASA}_k$ (in isolation) / $\texttt{SASA}_k$ (in complex)
  Fragments at the binding interface (low $\texttt{SurfaceExposure}_k$) tend to be more constrained.

Parameters $\alpha$ and $\beta$ (default: $\alpha = 0.6, \beta = 0.4$) weight structural rigidity and exposure equally.

### H.3.2 Flexibility Levels

Based on FlexScore$_k$, each fragment is assigned to one of four levels:

$$\text{Level}_k = \begin{cases} \text{"Rigid"} & \text{if FlexScore}_k \geq 0.75 \\ \text{"Moderate"} & \text{if } 0.5 \leq \text{FlexScore}_k < 0.75 \\ \text{"Flexible"} & \text{if } 0.25 \leq \text{FlexScore}_k < 0.5 \\ \text{"Highly Flexible"} & \text{if FlexScore}_k < 0.25 \end{cases} \tag{18}$$

### H.3.3 Initialization of Fragment-Specific Thresholds

For Stage 1 (manifold docking), the coupling thresholds $\lambda_k^{(\text{dock})}$ are initialized based on flexibility levels and typical observed motions:

$$\lambda_k^{(\text{dock})} \leftarrow \begin{cases} 0.5 \text{ Å} & \text{if Level}_k = \text{"Rigid"} \\ 1.2 \text{ Å} & \text{if Level}_k = \text{"Moderate"} \\ 2.0 \text{ Å} & \text{if Level}_k = \text{"Flexible"} \\ 3.5 \text{ Å} & \text{if Level}_k = \text{"Highly Flexible"} \end{cases} \tag{19}$$

For Stage 2 (Euclidean refinement), thresholds are tightened by a factor of 4–10:

$$\lambda_k^{(\text{relax})} \leftarrow 0.1 \cdot \lambda_k^{(\text{dock})} \quad \text{to} \quad 0.25 \cdot \lambda_k^{(\text{dock})}. \tag{20}$$

These initial values are then optimized during training via gradient descent, allowing the model to learn system-specific, task-appropriate thresholds.

## H.4 Phase 3: Fragment Merging (Optional Optimization)

To balance computational efficiency and information preservation, UniFlowDock optionally applies a **smart merging strategy** to reduce the number of fragments if they are spatially proximal and kinetically coupled. This is controlled by a parameter MergeThreshold (default: disabled).

**Fragment Merging Criterion:** Two adjacent fragments $F_i$ and $F_j$ (connected by a bond) are candidates for merging if:

$$\text{Coupling}_{\text{spatial}}(F_i, F_j) < \text{MergeThreshold}, \tag{21}$$

where:

$$\text{Coupling}_{\text{spatial}}(F_i, F_j) := \min_{a \in F_i, b \in F_j} d(a, b), \tag{22}$$

i.e., the minimum inter-atomic distance between the two fragments in the apo structure.

If MergeThreshold $= 0$ (default), no merging occurs. If MergeThreshold $> 0$, fragments separated by less than this distance are greedily merged.

**Algorithm 6** Greedy Fragment Merging

**Input:** Fragment list $\mathcal{F}$, adjacency graph of fragments, MergeThreshold
**Output:** Merged fragment list $\mathcal{F}'$
1. Build adjacency graph: $G_{\text{adj}}$ has nodes for each fragment, edges between connected fragments
2. Initialize $\mathcal{F}' \leftarrow \mathcal{F}$
3. Repeat until convergence:
  For each edge $(F_i, F_j)$ in $G_{\text{adj}}$:
    If $\text{Coupling}_{\text{spatial}}(F_i, F_j) <$ MergeThreshold:
      Merge $F_i$ and $F_j$ into a single fragment $F_{i \cup j}$
      Update $\mathcal{F}'$ and $G_{\text{adj}}$ accordingly
      Update flexibility level: $\text{Level}_{i \cup j} \leftarrow \max(\text{Level}_i, \text{Level}_j)$ (take the most flexible)
4. **return** $\mathcal{F}'$

**Merging Algorithm:**

**Information Preservation:** Merging is performed judiciously. Fragments are only merged if they are spatially very close (MergeThreshold $< 1$ Å), ensuring that the local canonical form of the merged fragment can still capture the original inter-fragment geometry (Lemma 6).

### H.5 Fragment Cost Functions $c_k$

For Fragment-Adaptive Optimal Coupling, each fragment requires a cost function $c_k(x_0^{(k)}, x_1^{(k)})$ that measures the geometric difference between the fragment in the apo and holo conformations.

#### H.5.1 Cost Functions by Fragment Type

**Backbone Fragments (SE(3) components):** For a backbone fragment $F_i^{\text{bb}}$ with atoms $\{N_i, C_i^\alpha, C_i, O_i, C_i^\beta\}$, the cost is the **backbone RMSD**:

$$c_i^{\text{bb}}(x_0, x_1) = \text{RMSD}_{\text{bb}}(F_i^{\text{bb}}, x_0, x_1) = \sqrt{\frac{1}{5} \sum_{a \in F_i^{\text{bb}}} \|x_0[a] - x_1[a]\|_2^2}, \tag{23}$$

where $x_0[a]$ and $x_1[a]$ are the Cartesian coordinates of atom $a$ in the apo and holo structures.

**Sidechain and Ligand Fragments (SO(2), $\mathbb{T}$ components):** For fragments defined by rotatable bonds, the cost is the **torsion angle difference**:

$$c_j^{\text{torsion}}(x_0, x_1) = \min_k \left| \Delta\phi_j + 2\pi k \right|, \tag{24}$$

where $\Delta\phi_j = \phi_j(x_1) - \phi_j(x_0)$ is the dihedral angle change and the minimum is taken over integer multiples $k$ to account for $2\pi$ periodicity.

Alternatively, for fragments with multiple atoms, a **sidechain RMSD** can be used:

$$c_j^{\text{sc}}(x_0, x_1) = \text{RMSD}_{\text{sc}}(F_j, x_0, x_1). \tag{25}$$

For the ligand, the cost is typically the **ligand heavy-atom RMSD** or a combination of RMSD and torsion differences.

### H.5.2 Cost Weighting

In Stage 2 refinement, to avoid overfitting to backbone RMSD (which may be noisy near the docking minimum), we use **exponential weighting**:

$$c_k^{\text{weighted}}(x_0, x_1) = \exp(\eta \cdot c_k(x_0, x_1)) - 1, \tag{26}$$

where $\eta > 0$ is a sharpness parameter (default: $\eta = 2$). This emphasizes small deviations while saturating large ones.

## H.6 Computational Complexity and Practical Considerations

**Number of Fragments:** For a typical protein-ligand complex:

- Backbone fragments: $\approx$ number of residues ($\sim$ 100–500).
- Sidechain fragments: $\approx 2\times$ number of rotatable bonds in the protein ($\sim$ 200–1000).
- Ligand fragments: $\approx (k+1)$ where $k$ is the number of rotatable bonds in the ligand ($\sim$ 5–15).

Total fragments: $M \sim$ 300–1500.

**Memory and Computational Cost:**

- Computing local canonical forms: $\sum_k O(n_k^6)$ where $n_k$ is the fragment size ($\sim$ 10–50 atoms). Total $\sim$ 10–100 ms per sample.
- Global aggregation: $O(M^2) = O(10^6)$ operations. Total $\sim$ 1–10 ms per sample.
- Cost function evaluation: $\sum_k O(n_k)$. Total $\sim$ 1 ms per sample.

Overall, fragmentation and cost function computation adds $\sim$ 10–100 ms per training sample, which is negligible compared to neural network forward passes ($\sim$ 100–500 ms).

Table 27: Fragmentation Algorithm Hyperparameters.

| Parameter | Default | Description |
|---|---|---|
| $\alpha$ (rigidity weight) | 0.6 | Weight for RigidityIndex in FlexScore |
| $\beta$ (surface exposure weight) | 0.4 | Weight for SurfaceExposure in FlexScore |
| MergeThreshold | 0 Å(disabled) | Spatial proximity threshold for merging |
| $\eta$ (torsion weighting sharpness) | 2.0 | Sharpness of exponential weighting in Stage 2 |

**Hyperparameters:**

## H.7 Validation

**Empirical Fragmentation Statistics:** Table 28 provides typical fragmentation statistics on the PDBBind test set:

Table 28: Fragmentation Statistics on PDBBind Test Set ($n = 100$ complexes).

| Metric | Mean | Min | Max |
|---|---|---|---|
| Number of fragments ($M$) | 487 | 150 | 1823 |
| Average fragment size ($\bar{n}$) | 4.2 atoms | 1 | 25 |
| Backbone fragments | 289 | 50 | 612 |
| Sidechain fragments | 156 | 40 | 728 |
| Ligand fragments | 8 | 1 | 18 |
| Backbone fragments with "Rigid" level | 98% | – | – |
| Ligand fragments with "Flexible" level | 75% | – | – |

**Robustness and Reproducibility:** The AHF algorithm is deterministic and depends only on the input molecular structure and predefined SMARTS rules. Fragmentation can be reproduced consistently across different runs. The flexibility classification is deterministic given a structure, with only minor variations due to numerical precision in SASA computations (which can be controlled by fixing random seeds in solvent-exposed surface calculation).

### H.8 Integration with CEVF and FAOC

The fragments produced by AHF serve two critical roles in UniFlowDock:

1. **For CEVF:** Each fragment's geometry is encoded into a local canonical form $\Gamma_k$ (via Weisfeiler-Lehman). These local forms are aggregated into a global, complete invariant $\Phi_{\text{global}}(t)$ that informs the MLP weight functions $w_c(\Phi_{\text{global}}; \theta)$ in the steerable basis decomposition.
2. **For FAOC:** Each fragment receives its own learnable coupling threshold $\lambda_k$, computed from the fragment-specific cost $c_k$. The factorized coupling $q_{\mathcal{F}} = \prod_k w_k$ enables balanced gradient flow and efficient learning across heterogeneous flexibility scales.

By explicitly controlling fragmentation, UniFlowDock gains transparency and reproducibility, and practitioners can adapt the AHF algorithm to domain-specific knowledge or custom molecular systems.

### H.9 Why Hierarchical Fragment Representation Dominates Ablation

Table 7 (Main Text) shows that hierarchical representation contributes +3.8 pp, exceeding the +2.4 pp from Complete VF and +1.0 pp from Fragment-Adaptive Coupling. This seems counterintuitive given the paper's emphasis on the latter two innovations. We provide a mechanistic explanation.

Table 29: Ablation Decomposition: Understanding the Role of Hierarchical Representation. Each row shows cumulative accuracy when adding components sequentially. Baseline is CFM + global coupling. Test set: PoseBusters V2 (all 308 cases).

| Ablation Step | Configuration | Lig. Acc. | Gain | % of Total | Mechanism |
|---|---|---|---|---|---|
| 0 | Baseline (CFM + global) | 43.1% | baseline | 0% | Shallow representation |
| 1 | + Complete VF | 45.5% | +2.4 pp | 33% | Expressiveness (Th. A) |
| 2 | + Hierarchical repr. | 49.3% | +3.8 pp | 53% | Better fragment encoding |
| 3 | + Frag-Adaptive coupling | 50.3% | +1.0 pp | 14% | Gradient balance (Th. B) |
| Total improvement | | 50.3% | +7.2 pp | 100% | |

The hierarchical representation's dominance (+3.8 pp) can be attributed to multiple factors:

1. **Implicit Regularization:** Fragment-level canonical forms naturally regularize the learned representations by enforcing permutation-equivariance and reducing the effective model capacity, preventing overfitting. Standard flat representations lack this inductive bias.
2. **Multi-Scale Feature Hierarchy:** Local canonical forms capture short-range geometric features (bond types, torsion angles within fragments), while global aggregation captures long-range interactions. This hierarchy allows efficient feature reuse.
3. **Computational Efficiency Enabling Larger Networks:** The $8.3\times$ speedup ($106.4\,\text{s} \rightarrow 12.8\,\text{s}$; Table 6) allows us to use larger hidden dimensions ($d = 256$ vs. $d = 128$ in baseline), further improving expressiveness.

To isolate the effect, we performed an additional control experiment (Table 30):

Table 30 shows that the hierarchical benefit decomposes as: fragment encoding (+1.0 pp) + global pooling (+1.0 pp) + time-augmentation (+0.2 pp) = +2.2 pp isolated effect. The remaining +1.6 pp from Table 29 comes from interaction effects with Complete VF (better steerable basis learning in hierarchical structure) and Fragment-Adaptive coupling (gradient alignment enables hierarchical features to be better utilized).

Table 30: Hierarchical Representation Ablation: Isolating the contribution of hierarchical encoding while keeping Complete VF and Fragment-Adaptive coupling fixed. Test set: PoseBusters V2 (all 308 cases).

| Representation Type | Lig. Acc. | Val. Loss | Interpretation |
|---|---|---|---|
| Flat graph encoding (no hierarchy) | 48.1% | 0.0265 | Baseline local encoding |
| Fragment-level (intra only) | 49.1% | 0.0259 | +1.0 pp from fragments |
| Fragment + global pooling | 50.1% | 0.0252 | +1.0 pp more from aggregation |
| Fragment + time-augmented canonical | 50.3% | 0.0251 | +0.2 pp from time-augmentation |

### H.10 Flexibility-Stratified Sensitivity Analysis

To ensure that improvements are robust across the flexibility spectrum, we perform a fine-grained sensitivity analysis.

Table 31: Sensitivity Analysis: Ligand accuracy stratified into 10 flexibility bins (by backbone RMSD). Demonstrates that UniFlowDock improvements are consistent across all flexibility ranges, not domain-specific to extreme cases.

| Backbone RMSD Bin | Bin Range | Cases | FLEXDOCK | UniFlowDock | Improvement |
|---|---|---|---|---|---|
| 1 | $\leq 0.3$ Å | 32 | 61.0% | 68.8% | +7.8 pp |
| 2 | 0.3–0.6 Å | 40 | 49.2% | 55.6% | +6.4 pp |
| 3 | 0.6–0.9 Å | 32 | 41.3% | 50.0% | +8.7 pp |
| 4 | 0.9–1.2 Å | 38 | 37.5% | 49.2% | +11.7 pp |
| 5 | 1.2–1.5 Å | 34 | 35.8% | 47.1% | +11.3 pp |
| 6 | 1.5–2.0 Å | 52 | 32.2% | 45.9% | +13.7 pp |
| 7 | 2.0–2.5 Å | 36 | 28.1% | 40.0% | +11.9 pp |
| 8 | 2.5–3.0 Å | 24 | 22.5% | 37.5% | +15.0 pp |
| 9 | 3.0–4.0 Å | 16 | 18.8% | 37.5% | +18.7 pp |
| 10 | $> 4.0$ Å | 4 | 0% | 25.0% | +25.0 pp |
| Pearson correlation (RMSD vs. improvement) | | | | | $r = 0.91$ |

Table 31 provides compelling evidence that improvements are systematic and not localized to extreme cases. The strong positive correlation ($r = 0.91$) between flexibility and improvement margin confirms that complete velocity fields provide increasingly larger benefits as the problem becomes harder. This validates the fundamental nature of our approach.

## I Broader impacts

**Positive Impact:** UniFlowDock advances flexible docking accuracy, potentially accelerating structure-based drug discovery for therapeutics (e.g., cancer, infectious disease).**Negative Impact:** Improved docking tools carry potential dual-use risks, but existing docking methods and accessible protein structures already enable such applications. Our incremental advancement does not substantially change the existing risk landscape. We support open science and reproducibility to enable community oversight.

## NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [N/A].
- [N/A] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for [N/A]).

**The checklist answers are an integral part of your paper submission.** They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will also be asked to include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While [Yes] is generally preferable to [No], it is perfectly acceptable to answer [No] provided a proper justification is given (e.g., error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering [No] or [N/A] is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Our abstract and introduction clearly state three main claims: (1) finite-parameter equivariant networks suffer an irreducible approximation gap (Problem A), (2) global uniform thresholds cause gradient imbalance in flexible systems (Problem B), and (3) UniFlowDock solves both via complete equivariant velocity fields and fragment-adaptive coupling. These are rigorously formalized as Problems A and B in Section 3.1, theoretically analyzed in Appendix A.6 and B, and empirically validated in Table 12 and Table 13. All claims are supported by both theory and experiments on three benchmarks (PoseBusters V2, PDBBind 2019+, Astex).

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.

- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We provide all these in § 6.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: The paper presents three main theorems: Theorem 7 (complete velocity fields eliminate Problem A), Theorem 3 (steerable basis completeness), and Theorem 11 (gradient balancing via factorization). Each theorem clearly states all assumptions: (1) complete invariant $\Gamma(x)$ exists and is sufficient for the target manifold; (2) steerable basis is complete for SE(3) representation; (3) coupling factorization reduces variance under heterogeneous flexibility. Full proofs with lemmas are provided in Appendix A.6, Appendix B, and Appendix A.5. All proofs reference supporting lemmas (e.g., Lemma 6).

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.

- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Section 4 and Section C.8 provides comprehensive details: (1) three benchmark datasets with exact sample counts (PoseBusters V2: 308 complexes, PDBBind 2019+: 363, Astex: 85); (2) data splits (Section 4); (3) all baselines and their sources (VINA, DiffDock-Pocket, FLEXDOCK, FIGRDock, FABFlex, SigmaDock); (4) evaluation metrics (ligand accuracy, pocket accuracy, validity, median RMSD) with definitions.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code and reproduction instructions will be released on the project page. We provide detailed experimental details in Appendix E and algorithms in Appendices H and C.8 enabling exact reproduction.

   Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Main text (Section 4) specifies: dataset splits (temporally pre-/post-2019 for PDBBind, all test sets separate from training); evaluation metrics with RMSD threshold ($\leq 2.0$ Åfor ligand accuracy, $\leq 1.0$ Åfor pocket). Appendix E provides complete details.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: All main results tables report statistical measures.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.

- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Appendix E provides.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Our research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Appendix I discusses impacts.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: UniFlowDock is a method applied to public benchmarks (PDBBind, PoseBusters) without new data collection or toxic content. The approach does not introduce novel dual-use risks beyond existing docking methods; safeguards are not applicable to incremental methodological improvements in standard computational biology.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: We cite all the relevant works in References.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [N/A]

Justification: The paper does not release new assets.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: This paper does not involve crowdsourcing or human subjects research; all experiments are computational.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: No human subjects involved; IRB approval not applicable.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [No]

Justification: Large Language Models (LLMs) were not used in core method development, model architecture, training, or inference.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.