

---

# PIGEON: Pocket-Inferred Geometric Ensembles for Flexible Docking

---

Anonymous Authors<sup>1</sup>

## Abstract

Molecular docking is a central tool for structure-based drug discovery, but realistic docking remains challenging because receptor conformations are flexible, a receptor-ligand pair may admit multiple stable holo states, and the true binding-pocket residues are often unknown in prospective applications. We introduce **PIGEON** (Pocket-Inferred Geometric Ensembles for Flexible Docking), a manifold flow-matching framework that performs ensemble flexible docking without ground-truth pocket residues. PIGEON represents receptor-ligand complexes with a time-conditioned heterogeneous graph and predicts residual motions for ligand translation, ligand rotation, ligand torsions, receptor residue frames, and receptor side-chain torsions. To guide ligand translation in the absence of pocket labels, PIGEON constructs a receptor-intrinsic buriedness field, a ligand-centered spherical harmonic pocket field, and a compact bank of pocket-aware candidate directions, then uses a learnable staged selector to choose between pocket-ingress and local-fitting motions. Empirically, PIGEON recovers accurate ligand poses while maintaining physically plausible full-complex geometry, providing a practical route to generating coupled ligand-receptor ensembles under realistic pocket-unknown conditions.

## 1 Introduction

Molecular docking is a central tool in structure-based drug discovery, where predicting how a small molecule binds to a protein target supports virtual screening, lead optimization, and mechanistic interpretation (Morris et al., 2009; Trott & Olson, 2010; Friesner et al., 2004). In many therapeutic targets, rigid-receptor assumptions break down: enzyme

<sup>1</sup>Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

active sites can require side-chain or loop rearrangements to admit catalytic or inhibitory binding modes (Sherman et al., 2006), GPCRs and kinases stabilize distinct ligand-specific conformations that drive different pharmacological outcomes (Wootten et al., 2018; Vijayan et al., 2015), and viral proteases such as HIV-1 protease undergo substantial flap and pocket rearrangements upon inhibitor binding (Heaslet et al., 2007). These applications expose three coupled challenges that remain difficult for current docking models: receptors are flexible and may change conformation upon ligand binding, a single receptor-ligand pair may admit multiple stable holo conformations, and the true binding-pocket residues are usually unknown in real applications.

Recent generative docking models have made substantial progress on flexible protein-ligand docking. Building on early geometric one-shot predictors (Stärk et al., 2022; Lu et al., 2022) and diffusion-based pose generation (Corso et al., 2023), methods such as DiffDock-Pocket (Plainer et al., 2023), Re-Dock (Huang et al., 2024), DynamicBind (Lu et al., 2024), FlexDock (Corso et al., 2025), and FlowDock (Morehead & Cheng, 2025) use diffusion or flow matching to predict ligand poses jointly with receptor adaptation. While these models capture important aspects of ligand-induced receptor motion, most are evaluated as single-endpoint pose generation rather than ensemble prediction.

Classical ensemble docking instead models receptor flexibility by docking ligands against multiple pre-generated conformations from crystal structures, molecular dynamics, or enhanced sampling (Lin et al., 2002; Amaro et al., 2008; 2018; Bottegoni et al., 2011). These pipelines are typically two-stage: receptor conformations are generated upstream, and docking, scoring, and ranking are performed separately against each. As a result, ensemble docking workflows focus on selecting from precomputed receptor states rather than directly generating ligand and receptor holo conformations as a coupled conditional distribution.

A further limitation is the treatment of pocket information. Site-specific docking methods generally require an explicit pocket center or residue set (Plainer et al., 2023; Huang et al., 2024), supplied by externally predicted pockets (Le Guilloux et al., 2009; Krivák & Hoksza, 2018), which is reasonable for retrospective redocking but biases the model

toward a fixed region in prospective use, where the holo pocket residues are unknown. Pocket-free methods such as DiffDock (Corso et al., 2023) and FlowDock (Morehead & Cheng, 2025) avoid explicit pocket inputs but do not explicitly model receptor-intrinsic pocket geometry to guide ligand motion, leaving open the challenge of inferring pocket-like directions while jointly predicting ligand and receptor conformational changes.

To address these gaps, we introduce **PIGEON: Pocket-Inferred Geometric Ensemble for Flexible Docking**, a manifold flow-matching framework for ensemble flexible docking without ground-truth pocket residues. Given an apo-like receptor-ligand state, PIGEON predicts residual motion toward holo ensemble states across ligand translation, rotation, torsions, receptor residue frames, and side-chain torsions. The key idea is to replace unconstrained ligand translation regression with a pocket-inferred geometric decision process: receptor-intrinsic buriedness and a ligand-centered spherical harmonic pocket field summarize local receptor geometry, from which we derive a compact bank of pocket-aware candidate translation directions. A learnable staged selector then chooses among these candidates using atom-aware probe features measuring contact quality, clash avoidance, void avoidance, contact uniformity, and contact count. The selected direction is combined with a learned positive magnitude, while separate heads predict ligand rotation, torsions, and receptor flexibility.

Our main contributions are:

1. We formulate ensemble flexible docking as manifold residual prediction over multiple apo-holo state pairs, allowing the model to learn multiple valid holo conformations for the same receptor-ligand system.
2. We introduce a pocket-inferred geometric translation module that derives ligand motion cues from receptor-intrinsic buriedness, ligand-centered spherical harmonic fields, pocket-potential gradients, and long-range buriedness directions without using true pocket residue indices.
3. We design a learnable staged direction selector that preserves the interpretability of geometric docking rules while learning how to balance pocket-ingress and local-fitting behavior from data.

## 2 Problem Formulation

### 2.1 Flexible Docking

We consider local flexible docking for protein-ligand complexes, where both the ligand pose and the receptor conformation may change upon binding. Let the receptor contain  $N$  residues and let the ligand contain  $A$  atoms. For recep-

tor residue  $i$ , we represent its local backbone frame by a rotation and frame origin  $(R_i^{\text{rec}}, \mathbf{p}_i^{\text{rec}}) \in SO(3) \times \mathbb{R}^3$ , and its side-chain torsions by  $\chi_i^{\text{rec}} \in \mathbb{T}^{K_i}$ , where  $\mathbb{T}$  denotes the angular torus. For the ligand, we represent the rigid pose by a centroid  $\mathbf{c}^{\text{lig}} \in \mathbb{R}^3$  and rotation  $R^{\text{lig}} \in SO(3)$ , and its internal conformation by rotatable-bond torsions  $\chi^{\text{lig}} \in \mathbb{T}^{K_L}$ . Since ligand local coordinates are centered at  $\mathbf{c}^{\text{lig}}$ , this centroid serves as the ligand rigid translation coordinate. The complete flexible docking state lies on the product manifold

$$\mathcal{M} = \left( \prod_{i=1}^N (SO(3) \times \mathbb{R}^3 \times \mathbb{T}^{K_i}) \right) \times \mathbb{R}^3 \times SO(3) \times \mathbb{T}^{K_L}.$$

Each receptor-ligand pair may admit multiple apo and holo states, so the model learns a distribution of flexible docking transitions rather than a deterministic apo-to-holo mapping, without access to ground-truth pocket residue indices.

### 2.2 Manifold Flow Matching

We formulate flexible docking as flow matching on  $\mathcal{M}$  (Lipman et al., 2022; 2024). For each apo-holo state pair  $(x_0, x_1)$  and scalar schedule  $\alpha(t)$ , the geodesic interpolant is

$$x_t = \text{Exp}_{x_0}(\alpha(t) \text{Log}_{x_0}(x_1)),$$

applied component-wise: linear interpolation in  $\mathbb{R}^3$  for frame origins and the ligand centroid,  $SO(3)$  geodesics for frame and ligand rotations, and wrapped angular interpolation for receptor and ligand torsions. For the ligand rigid pose, this gives

$$\begin{aligned} \mathbf{c}_t^{\text{lig}} &= (1 - \alpha(t))\mathbf{c}_0^{\text{lig}} + \alpha(t)\mathbf{c}_1^{\text{lig}}, \\ R_t^{\text{lig}} &= R_0^{\text{lig}} \exp\left(\alpha(t) \log\left((R_0^{\text{lig}})^{-1} R_1^{\text{lig}}\right)\right). \end{aligned}$$

We use the polynomial schedule  $\alpha(t) = t^2$ .

The standard flow-matching objective trains a vector field  $v_\theta(x_t, t) \in T_{x_t}\mathcal{M}$  to match the target velocity  $u_t^* = \frac{d}{dt}x_t$ :

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{x_0, x_1, t} \left[ \|v_\theta(x_t, t) - u_t^*\|_{T_{x_t}\mathcal{M}}^2 \right].$$

For the geodesic interpolant, the path velocity is proportional to the endpoint residual  $r_t^* = \text{Log}_{x_t}(x_1)$  along each component, so we equivalently predict residuals from  $x_t$  to  $x_1$ :

$$\mathcal{L}_{\text{res}} = \mathbb{E}_{x_0, x_1, t} \left[ \|\widehat{r}_\theta(x_t, t) - \text{Log}_{x_t}(x_1)\|_{T_{x_t}\mathcal{M}}^2 \right].$$

Concretely, the ligand translation residual is  $\Delta \mathbf{c}_{t \rightarrow 1}^{\text{lig},*} = \mathbf{c}_1^{\text{lig}} - \mathbf{c}_t^{\text{lig}}$ , and the ligand rotation residual is  $\omega_{t \rightarrow 1}^{\text{lig},*} = \log\left((R_t^{\text{lig}})^{-1} R_1^{\text{lig}}\right)$ . This residual formulation is equivalent to velocity matching up to a known time-dependent scaling factor.

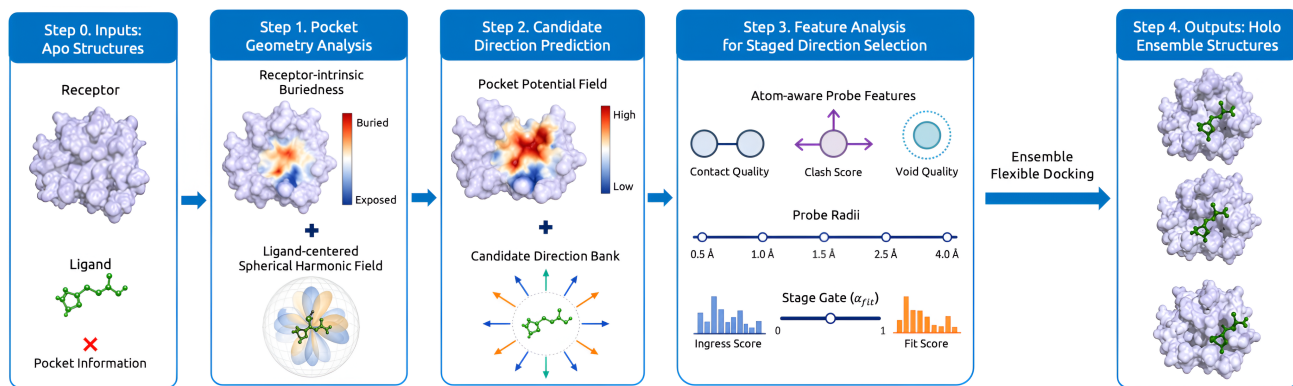


Figure 1. **Overview of PIGEON.** PIGEON performs pocket-unknown ensemble flexible docking by inferring receptor pocket geometry, constructing pocket-aware candidate translation directions, selecting a direction with a learnable staged selector, and predicting ligand and receptor residual motions.

### 3 PIGEON: Pocket-Inferred Geometric Ensemble Flexible Docking

#### 3.1 Heterogeneous Complex Representation

We represent each interpolated state  $x_t$  as a heterogeneous graph spanning three structural resolutions: ligand atoms, receptor atoms, and receptor residues. Ligand-atom nodes encode atom type, chemical features, current coordinates, and time; receptor-atom nodes encode atom type, van der Waals radius, current coordinates, and time; receptor-residue nodes encode residue features, current frame, side-chain torsions, and time. The graph contains ligand bond edges, receptor atom-radius edges, receptor residue-radius edges, receptor atom-to-residue edges, and ligand-receptor cross edges. A time-conditioned message-passing network with FiLM-style time modulation and edge-type-specific projections produces final embeddings  $\mathbf{h}_j^{\text{lig}}$ ,  $\mathbf{h}_a^{\text{rec-atom}}$ , and  $\mathbf{h}_i^{\text{rec-res}}$ , used by all downstream output heads (Perez et al., 2018). Architecture details are provided in the Appendix.

#### 3.2 Receptor-Intrinsic Buriedness

Atoms lining a cavity are more geometrically enclosed than solvent-exposed atoms. For receptor atom  $a$  with position  $\mathbf{q}_a$  and van der Waals radius  $r_a^{\text{rec}}$ , we cast  $K_{\text{probe}}$  probe directions  $\{\mathbf{u}_k\}_{k=1}^{K_{\text{probe}}}$  and mark a direction occluded if another receptor atom lies within the probe radius and subtends it. With  $\mathbf{d}_{ab} = \mathbf{q}_b - \mathbf{q}_a$ ,

$$o_{ak} = \mathbb{1} \left[ \exists b : 0 < \|\mathbf{d}_{ab}\| < r_{\text{probe}} \wedge \frac{\mathbf{u}_k^\top \mathbf{d}_{ab}}{\|\mathbf{d}_{ab}\|} > \cos \left( \arcsin \frac{r_a^{\text{rec}} + r_b^{\text{rec}}}{\|\mathbf{d}_{ab}\| + \epsilon} \right) \right]. \quad (1)$$

The buriedness of atom  $a$  is

$$b_a = \frac{1}{K_{\text{probe}}} \sum_{k=1}^{K_{\text{probe}}} o_{ak}.$$

Buriedness depends only on receptor geometry, so it requires no pocket labels at inference.

#### 3.3 Ligand-Centered Spherical Harmonic Pocket Field

To capture pocket geometry local to the current ligand pose, we construct a spherical harmonic field around  $\mathbf{c}_t^{\text{lig}}$ . For each receptor atom  $a$ ,

$$\mathbf{r}_{a,t} = \mathbf{q}_{a,t} - \mathbf{c}_t^{\text{lig}}, \quad r_{a,t} = \|\mathbf{r}_{a,t}\|, \quad \hat{\mathbf{r}}_{a,t} = \mathbf{r}_{a,t} / (r_{a,t} + \epsilon).$$

We use Gaussian radial shells with centers  $\{R_s\}_{s=1}^M$  and widths  $\{\sigma_s\}_{s=1}^M$ :

$$\phi_s(r_{a,t}) = \exp \left( -\frac{(r_{a,t} - R_s)^2}{2\sigma_s^2} \right),$$

with  $M = 6$  shells:

$$R = (1.5, 2.5, 4.0, 8.0, 14.0, 20.0) \text{ \AA},$$

$$\sigma = (0.75, 1.0, 2.0, 3.0, 4.0, 5.0) \text{ \AA}.$$

Per shell, we compute buriedness-weighted real spherical harmonic coefficients up to order  $L = 2$ :

$$\mathbf{f}_\ell^{(s)}(\mathbf{c}_t^{\text{lig}}) = \sum_a b_a \phi_s(r_{a,t}) \mathbf{Y}_\ell(\hat{\mathbf{r}}_{a,t}), \quad \ell = 0, 1, 2,$$

with  $L = 0$  summarizing buried mass at each radius,  $L = 1$  giving a dipole direction toward asymmetric buried mass, and  $L = 2$  capturing anisotropic pocket shape. The full representation is

$$\mathbf{F}_{\text{SH}}(\mathbf{c}_t^{\text{lig}}) = \bigoplus_{s=1}^M \bigoplus_{\ell=0}^L \mathbf{f}_\ell^{(s)}(\mathbf{c}_t^{\text{lig}}).$$

#### 3.4 Pocket-Aware Translation Candidate Bank

Rather than regressing a free 3D direction, we define a compact bank of receptor-derived candidates. The first group comes from the six shellwise  $L = 1$  dipoles. Letting  $\mathbf{d}_s^{\text{SH}}$  denote the 3D vector associated with the  $L = 1$  coefficients in shell  $s$ ,

$$\mathbf{v}_s^{\text{SH}} = \mathbf{d}_s^{\text{SH}} / (\|\mathbf{d}_s^{\text{SH}}\| + \epsilon).$$

The second group comes from multi-scale pocket-potential gradients. For a preferred pocket distance  $r_{\text{opt}}$ ,

$$U_{r_{\text{opt}}}(\mathbf{c}_t^{\text{lig}}) = \sum_a b_a \exp\left(-\frac{(\|\mathbf{q}_{a,t} - \mathbf{c}_t^{\text{lig}}\| - r_{\text{opt}})^2}{2\sigma_{\text{pot}}^2}\right),$$

$$\mathbf{v}_{r_{\text{opt}}}^{\text{pot}} = \nabla_{\mathbf{c}_t^{\text{lig}}} U_{r_{\text{opt}}}(\mathbf{c}_t^{\text{lig}}) / (\|\nabla_{\mathbf{c}_t^{\text{lig}}} U_{r_{\text{opt}}}(\mathbf{c}_t^{\text{lig}})\| + \epsilon).$$

We use five potential scales, giving five gradient candidates. We add two long-range buriedness directions: one toward the global buriedness-weighted receptor centroid,

$$\bar{\mathbf{q}}_{\text{buried}} = \frac{\sum_a b_a \mathbf{q}_{a,t}}{\sum_a b_a + \epsilon},$$

$$\mathbf{v}_{\text{centroid}} = \frac{\bar{\mathbf{q}}_{\text{buried}} - \mathbf{c}_t^{\text{lig}}}{\|\bar{\mathbf{q}}_{\text{buried}} - \mathbf{c}_t^{\text{lig}}\| + \epsilon},$$

and one toward the nearest highly buried receptor atom,

$$a^* = \arg \min_{a: b_a > \tau_b} \|\mathbf{q}_{a,t} - \mathbf{c}_t^{\text{lig}}\|,$$

$$\mathbf{v}_{\text{nearest}} = \frac{\mathbf{q}_{a^*,t} - \mathbf{c}_t^{\text{lig}}}{\|\mathbf{q}_{a^*,t} - \mathbf{c}_t^{\text{lig}}\| + \epsilon}.$$

The 13-direction candidate bank is

$$\mathcal{V}_t = \{\mathbf{v}_1, \dots, \mathbf{v}_{13}\}, \quad \|\mathbf{v}_k\| = 1.$$

### 3.5 Learnable Direction Selector

Different docking states call for different motions: when the ligand is outside the pocket, an *ingress* move toward buried receptor regions is preferred; when the ligand is already near a pocket, a *fit* move that improves local packing is preferred. We use a learnable staged selector that interpolates between an ingress and a fit scoring rule, learning only the coefficients controlling feature importance, probe-distance importance, and the ingress-fit transition.

**Probe geometry.** For each candidate  $\mathbf{v}_k$  and probe distance  $\rho_s$ , the hypothetical translated centroid and ligand atoms are

$$\mathbf{p}_{k,s} = \mathbf{c}_t^{\text{lig}} + \rho_s \mathbf{v}_k, \quad \tilde{\mathbf{x}}_{j,k,s} = \mathbf{x}_{j,t} + \rho_s \mathbf{v}_k,$$

with probe schedule  $\mathcal{R} = \{0.5, 1.0, 1.5, 2.5, 4.0\}$  Å. Letting  $r_{\text{lig}} = \max_j \|\tilde{\mathbf{x}}_{j,k,s} - \mathbf{p}_{k,s}\|$ , the local receptor atom set is

$$\mathcal{A}_{k,s} = \{a : \|\mathbf{q}_{a,t} - \mathbf{p}_{k,s}\| < r_{\text{probe}} + r_{\text{lig}} + \gamma_{\text{out}} + 1\}.$$

For  $a \in \mathcal{A}_{k,s}$  and ligand atom  $j$ , the surface gap is

$$g_{aj}^{(k,s)} = \|\mathbf{q}_{a,t} - \tilde{\mathbf{x}}_{j,k,s}\| - (r_a^{\text{rec}} + r_j^{\text{lig}}),$$

with closest-gap reductions

$$g_{a,\min}^{(k,s)} = \min_j g_{aj}^{(k,s)}, \quad h_{j,\min}^{(k,s)} = \min_{a \in \mathcal{A}_{k,s}} g_{aj}^{(k,s)}.$$

**Probe features.** We define five probe features. With contact width  $\sigma_{\text{contact}}$  and preferred contact gap  $\gamma_{\text{pref}}$ ,

$$C_a^{(k,s)} = \exp\left(-\frac{(g_{a,\min}^{(k,s)} - \gamma_{\text{pref}})^2}{2\sigma_{\text{contact}}^2}\right),$$

$$F_{\text{contact}}(k,s) = \log\left(1 + \sum_{a \in \mathcal{A}_{k,s}} b_a C_a^{(k,s)}\right),$$

$$P_{\text{clash}}(k,s) = \sum_{a \in \mathcal{A}_{k,s}} b_a \text{ReLU}(-g_{a,\min}^{(k,s)}),$$

$$F_{\text{clash}}(k,s) = -\log(1 + P_{\text{clash}}(k,s)),$$

$$P_{\text{void}}(k,s) = \frac{1}{A} \sum_{j=1}^A \text{ReLU}(h_{j,\min}^{(k,s)} - \gamma_{\text{void}}),$$

$$F_{\text{void}}(k,s) = -\log(1 + P_{\text{void}}(k,s)).$$

With contact-shell set

$$\mathcal{C}_{k,s} = \left\{a \in \mathcal{A}_{k,s} : \gamma_{\text{in}} < g_{a,\min}^{(k,s)} < \gamma_{\text{out}} \text{ and } b_a > \tau_b\right\},$$

the contact-count feature is  $F_{\text{count}}(k,s) = \log(1 + |\mathcal{C}_{k,s}|)$ . For uniformity, define

$$\mathbf{u}_{a,k,s} = \frac{\mathbf{q}_{a,t} - \mathbf{p}_{k,s}}{\|\mathbf{q}_{a,t} - \mathbf{p}_{k,s}\| + \epsilon},$$

and, for  $|\mathcal{C}_{k,s}| \geq 3$ ,

$$F_{\text{uniform}}(k,s) = 1 - \left\| \frac{1}{|\mathcal{C}_{k,s}|} \sum_{a \in \mathcal{C}_{k,s}} \mathbf{u}_{a,k,s} \right\|,$$

with  $F_{\text{uniform}} = 0$  when  $|\mathcal{C}_{k,s}| < 3$ . Together,  $F_{\text{contact}}$  rewards good buried contacts,  $F_{\text{clash}}$  penalizes overlap,  $F_{\text{void}}$  penalizes empty space,  $F_{\text{count}}$  counts plausible buried contacts, and  $F_{\text{uniform}}$  rewards balanced contact directions.

We collect features in fixed order:

$$\mathbf{F}(k,s) = [F_{\text{contact}}, F_{\text{clash}}, F_{\text{void}}, F_{\text{uniform}}, F_{\text{count}}](k,s),$$

and within-sample normalize across the  $K = 13$  candidates:

$$\mu_{s,f} = \frac{1}{K} \sum_{k=1}^K F_f(k,s),$$

$$\sigma_{s,f} = \sqrt{\frac{1}{K} \sum_{k=1}^K (F_f(k,s) - \mu_{s,f})^2 + \epsilon},$$

$$Z(k,s,f) = \frac{F_f(k,s) - \mu_{s,f}}{\sigma_{s,f}}.$$

**Stage gate.** We summarize the current pose with a stage-feature vector  $\mathbf{z}_{\text{stage}} = [z_d, z_q, z_c, z_u, z_n]$ , where  $z_d$  is a smoothed pocket-distance signal and  $z_q, z_c, z_u, z_n$  are bounded versions of  $F_{\text{contact}}, P_{\text{clash}}, F_{\text{uniform}}, F_{\text{count}}$  evaluated at the untranslated ligand position (full definitions in the Appendix B.9). The learnable fit gate is

$$\alpha_{\text{fit}} = \sigma \left( s_{\text{gate}} (\mathbf{w}_{\text{gate}}^{\top} \mathbf{z}_{\text{stage}} + b_{\text{gate}}) \right),$$

where  $\mathbf{w}_{\text{gate}}, b_{\text{gate}}, s_{\text{gate}}$  are initialized from the predefined rule and learned.

**Ingress and fit scores.** Both rules use  $Z(k, s, f)$  with separate learnable feature logits  $\boldsymbol{\eta}^{\text{ing}}, \boldsymbol{\eta}^{\text{fit}}$  and probe-distance logits  $\boldsymbol{\zeta}^{\text{ing}}, \boldsymbol{\zeta}^{\text{fit}}$ :

$$\begin{aligned} \boldsymbol{\lambda}^{\text{ing}} &= \text{softmax}(\boldsymbol{\eta}^{\text{ing}}), & \boldsymbol{\lambda}^{\text{fit}} &= \text{softmax}(\boldsymbol{\eta}^{\text{fit}}), \\ \boldsymbol{\rho}^{\text{ing}} &= \text{softmax}(\boldsymbol{\zeta}^{\text{ing}}), & \boldsymbol{\rho}^{\text{fit}} &= \text{softmax}(\boldsymbol{\zeta}^{\text{fit}}). \end{aligned}$$

$$s_k^{\text{ing}} = \sum_s \rho_s^{\text{ing}} \sum_f \lambda_f^{\text{ing}} Z(k, s, f),$$

$$s_k^{\text{fit}} = \sum_s \rho_s^{\text{fit}} \sum_f \lambda_f^{\text{fit}} Z(k, s, f).$$

The stage-weighted selector logit and selected direction are

$$\ell_k = (1 - \alpha_{\text{fit}}) s_k^{\text{ing}} + \alpha_{\text{fit}} s_k^{\text{fit}}, \quad k^* = \arg \max_k \ell_k, \quad \hat{\mathbf{u}} = \mathbf{v}_{k^*}.$$

### 3.6 Other Prediction Heads

A magnitude head consumes the pooled graph context, the time embedding, and invariant pocket-field summaries, and outputs a positive scalar through a Softplus readout; the final translation prediction is the selected unit direction times this magnitude. Additional heads predict the ligand  $SO(3)$  log residual, wrapped ligand torsion residuals, per-residue receptor-frame residuals, and per-residue side-chain torsion residuals, all consuming the embeddings defined above. Architectural details are in the Appendix.

### 3.7 Training Objective

For direction selection, the normalized ground-truth direction is

$$\mathbf{u}^* = \Delta \mathbf{c}_{t \rightarrow 1}^{\text{lig},*} / (\|\Delta \mathbf{c}_{t \rightarrow 1}^{\text{lig},*}\| + \epsilon),$$

with oracle cosines  $q_k^* = \mathbf{v}_k^{\top} \mathbf{u}^*$  and soft oracle distribution

$$p_k^* = \frac{\exp(q_k^*/T_{\text{oracle}})}{\sum_{j=1}^K \exp(q_j^*/T_{\text{oracle}})}.$$

The selector is trained with the listwise rank loss

$$\mathcal{L}_{\text{rank}} = - \sum_{k=1}^K p_k^* \log \frac{\exp(\ell_k)}{\sum_{j=1}^K \exp(\ell_j)},$$

*Table 1.* Ensemble flexible docking performance. We compare PIGEON against representative generative docking baselines using ligand RMSD, all-atom RMSD, PoseBusters validity, joint PoseBusters-valid and ligand-RMSD success rate, and runtime. Higher is better for percentage metrics marked with  $\uparrow$ , and lower is better for median RMSD and runtime.

| Method                 | Ligand RMSD        |                       | All-Atom RMSD      | % PB valid and L-RMSD < 2 Å $\uparrow$ | Runtime (s) |
|------------------------|--------------------|-----------------------|--------------------|--|-------------|
|                        | % < 2 Å $\uparrow$ | Median Å $\downarrow$ | % < 1 Å $\uparrow$ |  |             |
| TankBind               | 2.39               | 7.14                  | 30.39              | 1.05                                   | 1.77        |
| FABFlex                | 13.95              | 3.78                  | 40.05              | 0.40                                   | 0.98        |
| FlowDock               | 26.28              | 3.17                  | 7.46               | 14.83                                  | 29.4        |
| DiffDock (rigid)       | 7.01               | 3.76                  | N.A.               | 0.99                                   | <b>0.31</b> |
| FlexDock               | <b>51.61</b>       | <b>1.95</b>           | 6.07               | <b>43.78</b>                           | 11.68       |
| <b>PIGEON (1-step)</b> | 33.89              | 2.48                  | <b>49.21</b>       | 12.99                                  | 1.49        |
| <b>PIGEON (4-step)</b> | 36.36              | 2.40                  | <b>44.18</b>       | 16.52                                  | 4.78        |

*Table 2. Ensemble docking coverage and assignment success.* We evaluate whether each method can recover multiple holo ensemble conformations for the same receptor-ligand pair. GT coverage@1Å all-atom measures the fraction of ground-truth holo states covered by at least one prediction within 1Å all-atom RMSD. Hungarian success@1Å all-atom measures assignment-based ensemble recovery after optimal matching between predictions and ground-truth holo states. Higher is better for both metrics.

| Method                 | Mean GT coverage@1Å | Mean Hungarian success@1Å |
|------------------------|---------------------|---------------------------|
|                        | all-atom $\uparrow$ | all-atom $\uparrow$       |
| TankBind               | 32.69               | 30.59                     |
| FABFlex                | 39.08               | 40.09                     |
| FlowDock               | 17.71               | 16.92                     |
| DiffDock (rigid)       | N.A.                | N.A.                      |
| FlexDock               | 11.57               | 7.33                      |
| <b>PIGEON (1-step)</b> | 36.56               | <b>49.40</b>              |
| <b>PIGEON (4-step)</b> | <b>39.49</b>        | 44.35                     |

regularized toward the rule-based initialization  $\theta_0$ :

$$\mathcal{L}_{\text{reg}} = \|\theta - \theta_0\|_2^2.$$

Magnitude is supervised in log space:

$$\mathcal{L}_{\text{mag}} = \left( \log(\hat{m} + \epsilon) - \log(\|\Delta \mathbf{c}_{t \rightarrow 1}^{\text{lig},*}\| + \epsilon) \right)^2.$$

Non-translation outputs use fixed-scale normalized residual losses; for prediction  $\hat{\mathbf{y}}$ , target  $\mathbf{y}^*$ , scale  $s$ ,

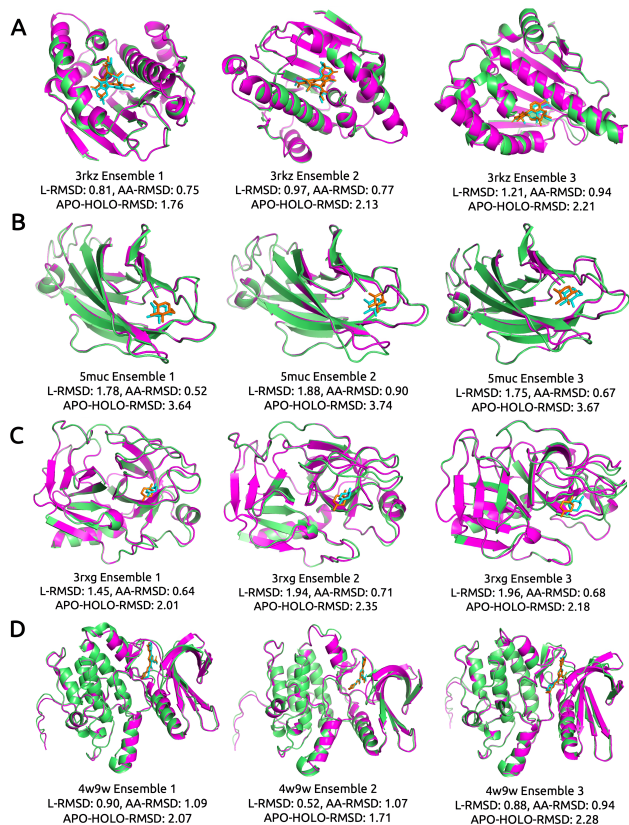
$$\mathcal{L}_{\text{res}}(\hat{\mathbf{y}}, \mathbf{y}^*; s) = \text{SmoothL1}(\hat{\mathbf{y}}/s, \mathbf{y}^*/s).$$

The full objective is

$$\mathcal{L} = \lambda_{\text{rank}} \mathcal{L}_{\text{rank}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_{\text{mag}} \mathcal{L}_{\text{mag}} + \lambda_{\text{frame}} \mathcal{L}_{\text{frame}} + \lambda_{\chi} \mathcal{L}_{\chi} + \lambda_{\text{rot}} \mathcal{L}_{\text{rot}} + \lambda_{\text{tor}} \mathcal{L}_{\text{tor}}.$$

## 4 Results

We evaluate PIGEON on an ensemble flexible docking benchmark built from apo structures from FlexDock (Corso et al., 2025) and their corresponding holo molecular-dynamics ensembles from MISATO (Siebenmorgen et al., 2024). The apo structures define the local docking inputs,



**Figure 2. Visualized ensemble flexible docking predictions.** We show representative ensemble predictions for four receptor-ligand complexes from the test set: (A) 3rkz, (B) 5muc, (C) 3rxg, and (D) 4w9w. For each complex, three holo ensemble targets are shown to illustrate that PIGEON can generate multiple flexible docking conformations for the same receptor-ligand pair. L-RMSD and AA-RMSD denote ligand RMSD and all-atom complex RMSD, respectively, after aligning the predicted and ground-truth complexes. APO-HOLO-RMSD denotes the ligand RMSD between the ligand pose in the ground-truth apo structure and the ligand pose in the corresponding ground-truth holo structure. Color code: holo protein in lime, predicted protein in magenta, true ligand in cyan, and predicted ligand in orange.

with frame-specific ligand augmentations around the binding region, while representative MISATO frames provide multiple holo targets for each receptor-ligand pair. This benchmark therefore tests whether a model can recover both ligand poses and the corresponding receptor conformations across holo ensemble states. Dataset construction details are provided in Appendix A.

#### 4.1 Docking Accuracy, Ensemble Recovery, and Efficiency

PIGEON achieves strong performance under the pocket-unknown ensemble flexible docking setting (Table 1). Among methods that do not require explicit pocket input, PIGEON performs favorably on both ligand-level and full-complex metrics. Compared with TankBind (Lu et al.,

2022), FABFlex (Zhang et al., 2025), and FlowDock (Morehead & Cheng, 2025), PIGEON improves ligand placement while also achieving substantially better all-atom complex recovery. This is important because ligand RMSD measures only pose accuracy, whereas all-atom RMSD also reflects whether the predicted receptor conformation is consistent with the holo complex. The improvement on both metrics suggests that PIGEON’s pocket-inferred geometric representation and coupled residual prediction heads help recover coordinated ligand-receptor holo geometry, rather than only placing the ligand near a plausible local region.

FlexDock remains the strongest baseline for ligand RMSD success, which is expected because it receives explicit pocket information. In contrast, PIGEON must infer pocket-like directions from receptor geometry and the current ligand state. The remaining ligand-RMSD gap to FlexDock should therefore be interpreted in light of this harder input setting. However, ligand RMSD alone does not measure whether the model recovers the holo receptor conformation associated with each ensemble state. In our benchmark, all-atom RMSD is evaluated against frame-specific MIS-ATO holo complexes, so a method can place the ligand accurately while still failing to reproduce the corresponding receptor geometry. Under this stricter full-complex criterion, PIGEON outperforms FlexDock, showing that it is better suited to recovering frame-specific coupled ligand-receptor holo conformations rather than only producing accurate pocket-localized ligand poses.

This distinction becomes clearer in the ensemble recovery metrics (Table 2). FlexDock performs well as a pocket-conditioned single-pose docking model, but its coverage remains limited because accurate docking to a supplied pocket does not necessarily produce diverse predictions matching different holo receptor-ligand conformations. Some pocket-free baselines obtain higher coverage than FlexDock, likely because their predictions vary more across runs or inputs; however, this variability is not always accompanied by accurate ligand placement or full-complex recovery. PIGEON achieves the strongest ensemble recovery because it combines pocket-inferred geometric conditioning with supervision over multiple apo-holo transitions per receptor-ligand pair, allowing it to cover distinct holo ensemble members while maintaining accurate coupled ligand-receptor geometry.

The Hungarian assignment metric provides a stricter test of ensemble generation than coverage alone (Kuhn, 1955). A model can obtain reasonable coverage by producing many predictions near a subset of easy holo states, but high Hungarian success requires one-to-one matching between predicted and ground-truth ensemble conformations. PIGEON performs strongly on this metric, showing that it does not merely generate repeated variants of a single bound state.

Table 3. Input-to-holo structural gap on the test set. Q1 and Q3 denote the 25th and 75th percentiles, respectively, and P90 denotes the 90th percentile. All values are in Å.

| Metric  | Mean  | Median | Q1    | Q3    | P90   |
|---|-------|--------|-------|-------|-------|
| <i>Ligand apo-holo gap</i>                      |       |        |       |       |       |
| Ligand RMSD, $x_0 \rightarrow x_1$              | 3.318 | 2.813  | 1.937 | 3.992 | 5.139 |
| Ligand centroid displacement                    | 2.673 | 2.230  | 1.539 | 3.328 | 3.980 |
| <i>Whole-receptor apo-holo gap</i>              |       |        |       |       |       |
| Receptor CA RMSD                                | 0.720 | 0.633  | 0.502 | 0.810 | 1.079 |
| Receptor mapped-atom RMSD                       | 0.890 | 0.803  | 0.651 | 1.002 | 1.283 |
| <i>Local pocket and active-site flexibility</i> |       |        |       |       |       |
| Pocket CA RMSD                                  | 1.112 | 0.974  | 0.763 | 1.280 | 1.713 |
| Pocket mapped-atom RMSD                         | 1.366 | 1.226  | 1.010 | 1.533 | 1.982 |
| Active-site CA RMSD                             | 1.098 | 0.962  | 0.771 | 1.245 | 1.643 |
| Active-site mapped-atom RMSD                    | 1.363 | 1.226  | 1.030 | 1.510 | 1.920 |
| <i>Aggregate complex gap</i>                    |       |        |       |       |       |
| Complex all-atom RMSD, $x_0 \rightarrow x_1$    | 1.013 | 0.893  | 0.727 | 1.125 | 1.503 |

Instead, its predictions are distributed across multiple holo endpoints, consistent with the goal of modeling ensemble flexible docking as a conditional distribution over coupled ligand-receptor conformations.

PIGEON also maintains competitive physical validity. Its PoseBusters-valid rate is comparable to other methods that do not use explicit pocket input, indicating that the structured translation and flexible receptor prediction do not come at the cost of severe geometric invalidity (Buttenschoen et al., 2024). The joint PoseBusters-valid and ligand-RMSD success metric remains more challenging, because it requires both chemically valid ligand geometry and accurate ligand placement. FlexDock performs strongly on this metric, likely benefiting from explicit pocket conditioning and stronger pose-level refinement. This suggests a clear direction for improving PIGEON: adding a lightweight local relaxation module, stronger stereochemical regularization, or validity-aware losses could improve chemically strict ligand success while preserving the model’s ensemble recovery advantages.

Finally, PIGEON is efficient at inference time. Even one-step generation already gives strong docking and ensemble recovery performance, while four-step generation improves ligand placement and coverage-related metrics with only a modest runtime increase (Table 1). This efficiency comes from the factorized translation design: PIGEON selects from a compact bank of pocket-aware directions and predicts a scalar magnitude, rather than relying on expensive unconstrained pose sampling. Thus, PIGEON provides a practical tradeoff between speed, ligand accuracy, full-complex recovery, and ensemble diversity under the realistic setting where the true pocket is not provided.

Table 4. Effect of frame-specific ligand augmentation. Statistics are computed over 7,300 frame pairs after receptor alignment.  $x_0$  denotes the unaugmented apo ligand pose,  $x_0^{\text{aug}}$  denotes the augmented ligand pose used as model input, and  $\Delta$  denotes the per-frame change in each metric induced by augmentation. Q1 and Q3 are the 25th and 75th percentiles, and P90 is the 90th percentile. All values are in Å.

| Metric                | Input                              | Mean  | Median | Q1    | Q3    | P90   |
|-----------------------|------------------------------------|-------|--------|-------|-------|-------|
| Ligand RMSD           | $x_0 \rightarrow x_1$              | 2.524 | 1.727  | 1.065 | 2.994 | 5.012 |
|                       | $x_0^{\text{aug}} \rightarrow x_1$ | 3.318 | 2.813  | 1.937 | 3.992 | 5.139 |
|                       | $\Delta$                           | 0.794 | 0.511  | 0.000 | 1.304 | 2.294 |
| Ligand centroid       | $x_0 \rightarrow x_1$              | 1.804 | 1.107  | 0.640 | 2.044 | 3.642 |
|                       | $x_0^{\text{aug}} \rightarrow x_1$ | 2.673 | 2.230  | 1.539 | 3.328 | 3.980 |
|                       | $\Delta$                           | 0.869 | 0.631  | 0.000 | 1.430 | 2.486 |
| Complex all-atom RMSD | $x_0 \rightarrow x_1$              | 0.972 | 0.851  | 0.691 | 1.081 | 1.437 |
|                       | $x_0^{\text{aug}} \rightarrow x_1$ | 1.013 | 0.893  | 0.727 | 1.125 | 1.503 |
|                       | $\Delta$                           | 0.041 | 0.011  | 0.000 | 0.037 | 0.119 |

## 4.2 Qualitative Analysis

The qualitative examples further illustrate PIGEON’s ability to recover multiple holo-like conformations for the same receptor-ligand pair (Figure 2). Across 3rkz, 5muc, 3rxg, and 4w9w, the three predicted ensemble members are not identical copies of a single docked pose; instead, they follow different holo targets with distinct ligand placements and receptor conformations, which is consistent with the ensemble recovery metrics. The visualized examples also show that the task is not a trivial reconstruction of the input structure. The apo-holo ligand displacements are often substantial, and the test-set statistics confirm that the initial ligand and receptor states have a non-negligible gap to the corresponding holo targets (Table 3). Moreover, Table 4 shows that the frame-specific ligand augmentation in the dataset construction increases the ligand input-to-holo displacement while keeping the overall complex perturbation controlled. Thus, successful predictions require PIGEON to move the ligand away from its augmented apo-like initialization, adjust its orientation and internal torsions, and predict receptor conformational changes that remain consistent with the corresponding MISATO holo frame. In the examples, PIGEON frequently reduces the apo-holo ligand gap while maintaining low all-atom complex error, indicating that the model learns meaningful residual motions for both ligand and receptor components rather than merely preserving the starting pose.

## 5 Discussion

Here, we have presented PIGEON, a pocket-inferred manifold flow-matching framework for ensemble flexible docking under the realistic setting where the true binding pocket is not provided. By integrating receptor-intrinsic buriedness fields, ligand-centered spherical harmonic pocket representations, pocket-aware translation candidates, and coupled

residual prediction heads for ligand and receptor degrees of freedom, PIGEON generates holo-like receptor-ligand ensembles without relying on ground-truth pocket residues. Across our ensemble docking benchmark, PIGEON demonstrates strong full-complex recovery, competitive ligand placement among pocket-free methods, efficient few-step generation, and improved recovery of multiple holo conformations for the same receptor-ligand pair. These results highlight the value of treating flexible docking as coupled ensemble generation rather than as single-pose ligand placement, providing a practical route toward modeling receptor adaptation and conformational heterogeneity in structure-based drug discovery.

Looking forward, several directions could further strengthen pocket-unknown ensemble flexible docking. While PIGEON already recovers diverse and physically plausible holo complexes, incorporating lightweight local relaxation, stronger stereochemical or torsional regularization, and validity-aware training objectives may further improve chemically strict ligand recovery. In addition, uncertainty-aware sampling and learned ranking could help prioritize the most relevant generated ensemble states for downstream virtual screening and prospective design. Extending the framework to broader blind-docking settings, larger receptor ensembles, and experimental validation would further test its generality. Nevertheless, PIGEON lays a solid foundation for pocket-unknown ensemble flexible docking, showing that receptor-inferred geometric cues can support efficient generation of diverse coupled ligand-receptor holo conformations without requiring explicit pocket inputs.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning, specifically generative modeling for structure-based drug discovery. PIGEON predicts protein-ligand binding geometries and does not generate novel chemical entities, synthesis routes, or activity predictions. The societal consequences of advancing flexible docking methods are largely those associated with computational drug discovery as a whole, none of which we feel must be specifically highlighted here.

## References

- Amaro, R. E., Baron, R., and McCammon, J. A. An improved relaxed complex scheme for receptor flexibility in computer-aided drug design. *Journal of Computer-Aided Molecular Design*, 22(9):693–705, 2008. doi: 10.1007/s10822-007-9159-2.
- Amaro, R. E., Baudry, J., Chodera, J., Demir, Ö., McCammon, J. A., Miao, Y., and Smith, J. C. Ensemble docking in drug discovery. *Biophysical Journal*, 114(10):2271–

2278, 2018. doi: 10.1016/j.bpj.2018.02.038.

- Bottegoni, G., Rocchia, W., Rueda, M., Abagyan, R., and Cavalli, A. Systematic exploitation of multiple receptor conformations for virtual ligand screening. *PLOS ONE*, 6(5):e18845, 2011. doi: 10.1371/journal.pone.0018845.
- Buttenschoen, M., Morris, G. M., and Deane, C. M. Posebusters: Ai-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical Science*, 15(9):3130–3139, 2024.
- Corso, G., Stärk, H., Jing, B., Barzilay, R., and Jaakkola, T. DiffDock: Diffusion steps, twists, and turns for molecular docking. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. URL <https://arxiv.org/abs/2210.01776>.
- Corso, G., Somnath, V. R., Getz, N., Barzilay, R., Jaakkola, T., and Krause, A. Composing unbalanced flows for flexible docking and relaxation. In *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025. URL <https://openreview.net/forum?id=gHLWTzKiZV>.
- Friesner, R. A., Banks, J. L., Murphy, R. B., Halgren, T. A., Klicic, J. J., Mainz, D. T., Repasky, M. P., Knoll, E. H., Shelley, M., Perry, J. K., Shaw, D. E., Francis, P., and Shenkin, P. S. Glide: A new approach for rapid, accurate docking and scoring. 1. method and assessment of docking accuracy. *Journal of Medicinal Chemistry*, 47(7):1739–1749, 2004. doi: 10.1021/jm0306430.
- Heaslet, H., Rosenfeld, R., Giffin, M., Lin, Y.-C., Tam, K., Torbett, B. E., Elder, J. H., McRee, D. E., and Stout, C. D. Conformational flexibility in the flap domains of ligand-free HIV protease. *Acta Crystallographica Section D: Biological Crystallography*, 63(8):866–875, 2007. doi: 10.1107/S0907444907029125.
- Huang, Y., Zhang, O., Wu, L., Tan, C., Lin, H., Gao, Z., Li, S., and Li, S. Z. Re-Dock: Towards flexible and realistic molecular docking with diffusion bridge. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pp. 20474–20489. PMLR, 2024.
- Kabsch, W. A solution for the best rotation to relate two sets of vectors. *Foundations of Crystallography*, 32(5): 922–923, 1976.
- Krivák, R. and Hoksza, D. P2Rank: machine learning based tool for rapid and accurate prediction of ligand binding sites from protein structure. *Journal of Cheminformatics*, 10(1):39, 2018. doi: 10.1186/s13321-018-0285-8.
- Kuhn, H. W. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

- 440 Le Guilloux, V., Schmidtke, P., and Tuffery, P. Fpocket:  
441 An open source platform for ligand pocket detection.  
442 *BMC Bioinformatics*, 10(1):168, 2009. doi: 10.1186/  
443 1471-2105-10-168.
- 444 Lin, J.-H., Perryman, A. L., Schames, J. R., and McCam-  
445 mon, J. A. Computational drug design accommodating  
446 receptor flexibility: the relaxed complex scheme. *Journal*  
447 *of the American Chemical Society*, 124(20):5632–5633,  
448 2002. doi: 10.1021/ja0260162.
- 450 Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and  
451 Le, M. Flow matching for generative modeling. *arXiv*  
452 *preprint arXiv:2210.02747*, 2022.
- 453 Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M.,  
454 Karrer, B., Chen, R. T., Lopez-Paz, D., Ben-Hamu, H.,  
455 and Gat, I. Flow matching guide and code. *arXiv preprint*  
456 *arXiv:2412.06264*, 2024.
- 458 Lu, W., Wu, Q., Zhang, J., Rao, J., Li, C., and Zheng, S.  
459 Tankbind: Trigonometry-aware neural networks for drug-  
460 protein binding structure prediction. *Advances in neural*  
461 *information processing systems*, 35:7236–7249, 2022.
- 463 Lu, W., Zhang, J., Huang, W., Zhang, Z., Jia, X., Wang,  
464 Z., Shi, L., Li, C., Wolynes, P. G., and Zheng, S. Dy-  
465 namicBind: predicting ligand-specific protein-ligand  
466 complex structure with a deep equivariant generative  
467 model. *Nature Communications*, 15(1):1071, 2024. doi:  
468 10.1038/s41467-024-45461-2.
- 469 Morehead, A. and Cheng, J. FlowDock: Geometric flow  
470 matching for generative protein–ligand docking and affin-  
471 ity prediction. *Bioinformatics*, 41(Supplement\_1):i198–  
472 i206, 2025. doi: 10.1093/bioinformatics/btaf187.
- 474 Morris, G. M., Huey, R., Lindstrom, W., Sanner, M. F.,  
475 Belew, R. K., Goodsell, D. S., and Olson, A. J.  
476 AutoDock4 and AutoDockTools4: Automated docking  
477 with selective receptor flexibility. *Journal of Com-  
478 putational Chemistry*, 30(16):2785–2791, 2009. doi:  
479 10.1002/jcc.21256.
- 480 Perez, E., Strub, F., De Vries, H., Dumoulin, V., and  
481 Courville, A. Film: Visual reasoning with a general con-  
482 ditioning layer. In *Proceedings of the AAAI conference*  
483 *on artificial intelligence*, volume 32, 2018.
- 485 Plainer, M., Toth, M., Dobers, S., Stärk, H., Corso, G., Mar-  
486 quet, C., and Barzilay, R. DiffDock-Pocket: Diffusion  
487 for pocket-level docking with sidechain flexibility. In  
488 *NeurIPS 2023 Workshop on Machine Learning in Struc-  
489 tural Biology (MLSB)*, 2023.
- 490 Sherman, W., Day, T., Jacobson, M. P., Friesner, R. A., and  
491 Farid, R. Novel procedure for modeling ligand/receptor  
492 induced fit effects. *Journal of Medicinal Chemistry*, 49  
493 (2):534–553, 2006. doi: 10.1021/jm050540c.
- 494 Siebenmorgen, T., Menezes, F., Benassou, S., Merdivan, E.,  
Didi, K., Mourão, A. S. D., Kitel, R., Liò, P., Kesselheim,  
S., Piraud, M., et al. Misato: machine learning dataset  
of protein–ligand complexes for structure-based drug dis-  
covery. *Nature computational science*, 4(5):367–378,  
2024.
- Stärk, H., Ganea, O.-E., Pattanaik, L., Barzilay, R., and  
Jaakkola, T. EquiBind: Geometric deep learning for drug  
binding structure prediction. In *Proceedings of the 39th*  
*International Conference on Machine Learning (ICML)*,  
volume 162 of *Proceedings of Machine Learning Re-  
search*, pp. 20503–20521. PMLR, 2022.
- Trott, O. and Olson, A. J. AutoDock Vina: Improving the  
speed and accuracy of docking with a new scoring func-  
tion, efficient optimization, and multithreading. *Journal*  
*of Computational Chemistry*, 31(2):455–461, 2010. doi:  
10.1002/jcc.21334.
- Vijayan, R. S. K., He, P., Modi, V., Duong-Ly, K. C., Ma,  
H., Peterson, J. R., Dunbrack, Jr., R. L., and Levy, R. M.  
Conformational analysis of the DFG-out kinase motif  
and biochemical profiling of structurally validated type  
II inhibitors. *Journal of Medicinal Chemistry*, 58(1):  
466–479, 2015. doi: 10.1021/jm501603h.
- Wootten, D., Christopoulos, A., Marti-Solano, M., Babu,  
M. M., and Sexton, P. M. Mechanisms of signalling and  
biased agonism in G protein-coupled receptors. *Nature*  
*Reviews Molecular Cell Biology*, 19(10):638–653, 2018.  
doi: 10.1038/s41580-018-0049-3.
- Zhang, Z., Wu, L., Gao, K., Yao, J., Qin, T., and Han, B.  
Fast and accurate blind flexible docking. *arXiv preprint*  
*arXiv:2502.14934*, 2025.

# Appendix

## A Dataset Preparation

We construct an ensemble flexible docking dataset by combining apo receptor-ligand complexes from FlexDock (Corso et al., 2025) with holo molecular-dynamics ensemble structures from MISATO (Siebenmorgen et al., 2024). The FlexDock data provide apo-like receptor structures and initial ligand information, while MISATO provides trajectory coordinates for the corresponding holo complexes. For each receptor-ligand pair, we treat the 100 MISATO trajectory frames as a set of possible holo ensemble states. Since training on all 100 frames for every complex is computationally expensive and contains redundant conformations, we select 5 representative holo frames for each PDB entry and use these frames as the ensemble targets during training.

**Representative holo-frame selection.** For each PDB entry, we first load the 100 MISATO trajectory frames and use the receptor and ligand atom mappings to extract the matched receptor and ligand atoms. To select representative ensemble states, we compute pairwise distances between trajectory frames using both receptor-pocket motion and ligand-pose variation. Specifically, for two frames, we align the receptor using a stable receptor subset and compute a weighted distance that combines receptor RMSD in the local pocket region and ligand RMSD after receptor alignment. The combined distance is defined as

$$D(i, j) = w_{\text{rec}} \text{RMSD}_{\text{rec}}(i, j) + w_{\text{lig}} \text{RMSD}_{\text{lig}}(i, j),$$

where  $w_{\text{rec}} = 0.4$  and  $w_{\text{lig}} = 0.6$  in our preprocessing. We then run  $k$ -medoids clustering with  $k = 5$  on this distance matrix and select the five medoid frames as representative holo ensemble states. This procedure preserves diverse ligand placements and receptor conformational changes while avoiding the redundancy of using all trajectory frames.

**Apo-gauge alignment of holo frames.** For each selected holo frame, we align the MISATO holo coordinates into the apo coordinate system. This step is necessary because the model should learn the conformational transition from an apo-like input state to a holo target state in a consistent global frame. We identify a stable receptor core using mapped receptor atoms, excluding chain termini and optionally refining the core by receptor RMSF across sampled trajectory frames. A Kabsch alignment is then computed between the selected holo frame and the apo receptor using this stable core (Kabsch, 1976). The same rigid transformation is applied to both receptor and ligand holo coordinates, producing apo-gauge holo receptor coordinates  $P_1$  and ligand coordinates  $L_1$ . We also save the corresponding apo receptor and ligand coordinates  $P_0$  and  $L_0$ , together with alignment metrics and the selected core-anchor metadata.

**Active-region and pocket labels for supervision and analysis.** Although PIGEON does not receive true pocket residue indices as model input, we still compute pocket and active-region labels during preprocessing for supervision, filtering, and analysis. For each selected holo frame, we define pocket residues based on the minimum heavy-atom distance between receptor residues and the holo ligand. A residue is labeled as pocket if any of its heavy atoms lies within a fixed cutoff of the ligand heavy atoms. We then define an active region by expanding the pocket through local residue-neighborhood rules, including one-hop sequence or structural shells and distance-based residue-center expansion. The resulting masks identify residues that directly contact the ligand and residues in the nearby flexible binding region. These labels are saved as per-frame masks and are used only as training targets or diagnostics, not as inference-time pocket inputs.

**Receptor flexible-state construction.** For every selected holo frame, we convert receptor coordinates into residue-level flexible degrees of freedom. Each residue is represented by a local backbone frame computed from its  $N$ ,  $C_\alpha$ , and  $C$  atoms, together with side-chain torsions computed from standard amino-acid chi definitions. We first construct target receptor latents from the aligned holo frame. We then perform frame-based receptor alignment and side-chain conformational matching to obtain receptor coordinates that preserve the target holo conformation while maintaining chemically valid local geometry. In particular, side-chain torsions are matched by rotating downstream atoms around chi axes, with clash-aware acceptance criteria. A subsequent backbone cleanup stage repairs local peptide-continuity issues and reduces artifacts introduced by frame-wise reconstruction. The final receptor target state contains post-cleanup receptor atom coordinates, residue-frame rotations and translations, frame-valid masks, receptor chi values, and receptor chi masks.

**Ligand topology, pose, and torsion matching.** We also construct a consistent ligand representation for each receptor-ligand pair. From the mapped ligand structure, we identify heavy atoms, atomic numbers, rotatable bonds, torsion-defining atom quadruplets, anchor atoms, and downstream masks for torsion rotations. The apo ligand coordinates are stored as the reference ligand state, while each selected holo frame provides a frame-specific target ligand conformation. For each target frame, we fit a rigid ligand pose and torsion vector so that the reconstructed ligand coordinates match the aligned holo ligand coordinates. The saved ligand state therefore contains the target ligand atom coordinates, ligand pose rotation and translation, ligand torsion values, torsion masks, and reconstruction diagnostics. This representation is consistent with the decoupled ligand-pose parameterization used by the model.

**Frame-specific ligand augmentation.** To better represent the ensemble local docking setting, we augment the initial ligand state separately for each selected holo frame. Rather than using only one fixed ligand initialization for all holo states of a receptor-ligand pair, we generate frame-specific apo-like ligand starts by applying controlled rigid perturbations and selecting poses that remain near the binding region while avoiding severe receptor-ligand clashes. Candidate augmented poses are evaluated using center-distance criteria, ligand RMSD improvement criteria, and clash constraints against the local receptor environment. The selected augmented pose is saved as the frame-specific ligand initial state, including augmented atom coordinates, pose rotation and translation, and torsion values. Thus, each receptor-ligand pair can contribute multiple apo-like initial states and multiple holo target states, matching our ensemble flexible docking formulation.

**Final dataset package.** After preprocessing, each training item corresponds to a PDB entry and one selected holo ensemble frame. The per-frame directory contains the aligned apo and holo coordinates, receptor active-region masks, receptor target latents, cleaned receptor target coordinates, ligand target pose and torsions, and optional frame-specific ligand augmentation. Static per-complex files contain ligand topology, receptor topology, receptor local templates, receptor chi topology, atomic numbers, and apo receptor latents. During training, the dataloader samples these processed apo-holo state pairs and constructs online interpolated states  $x_t$  between the frame-specific initial state and the selected holo target. This dataset construction enables PIGEON to learn ensemble flexible docking transitions while preserving the realistic constraint that the true pocket residue indices are not provided as model input.

## B Model Implementation

This section provides implementation details for PIGEON that are omitted from the main Method section for clarity. We describe the heterogeneous graph encoder, time conditioning, pocket-field construction, learnable selector implementation, output heads, online interpolation, and training protocol.

### B.1 Heterogeneous Graph Construction

For every interpolated state  $x_t$ , we reconstruct the current receptor atom coordinates and ligand atom coordinates before graph construction. The graph contains three node types: ligand atoms, receptor atoms, and receptor residues. Ligand atoms represent the chemically detailed ligand geometry. Receptor atoms provide atom-level steric and contact information. Receptor residues provide coarse flexible receptor state information, including residue identity, local frame, and side-chain torsions.

The graph uses five edge types. Ligand bond edges connect covalently bonded ligand atoms and carry ligand bond features. Receptor atom-radius edges connect receptor atoms within a fixed spatial cutoff and encode local atom-level receptor geometry. Receptor residue-radius edges connect nearby residues and encode residue-level geometry. Receptor atom-to-residue edges connect receptor atoms to their parent residues. Ligand-receptor cross edges connect ligand atoms to nearby receptor atoms and provide the main local interaction channel between ligand and receptor.

For a generic spatial edge from node  $u$  to node  $v$ , with positions  $\mathbf{p}_u$  and  $\mathbf{p}_v$ , we compute the relative displacement, distance, and normalized direction as  $\mathbf{d}_{uv} = \mathbf{p}_v - \mathbf{p}_u$ ,  $r_{uv} = \|\mathbf{d}_{uv}\|$ , and  $\hat{\mathbf{d}}_{uv} = \mathbf{d}_{uv}/(r_{uv} + \epsilon)$ . Distances are further encoded with radial basis functions  $\phi_\ell(r_{uv}) = \exp(-(r_{uv} - \mu_\ell)^2/(2\sigma_{\text{rbf}}^2))$ . The final edge feature concatenates the scalar distance, unit direction, radial basis encoding, and any edge-type-specific attributes.

## B.2 Node and Edge Feature Projections

Each node type has its own input projection into the shared hidden dimension. Ligand atom features include atom identity, ligand atom descriptors, atomic number, van der Waals radius, current atom coordinate information, and time features. Receptor atom features include atom identity, atomic number, van der Waals radius, current atom coordinates, and time features. Receptor residue features include residue identity features, residue-level embeddings, current receptor frame origin, current receptor frame rotation, encoded side-chain torsion values, and time features.

For torsion features, angular values are encoded with sine and cosine channels rather than raw angles. This avoids discontinuities at the  $-\pi$  and  $\pi$  boundary. For receptor frames, rotations are represented in flattened matrix form in the input feature vector, while output frame residuals are predicted as 6D local rigid-frame residuals.

Each edge type also has a separate projection into the shared edge hidden dimension. This allows ligand bonds, receptor atom contacts, receptor residue contacts, atom-residue edges, and ligand-receptor cross edges to retain distinct geometric and chemical meanings while being processed by the same heterogeneous interaction backbone.

## B.3 Time Conditioning with FiLM Modulation

PIGEON uses explicit time conditioning throughout the encoder. The scalar interpolation time  $t$  is first embedded with a sinusoidal embedding and projected to the model hidden dimension. Let  $\mathbf{e}_t$  denote the resulting time embedding. For each node type, the initial projected node state is modulated by a FiLM transformation:

$$\text{FiLM}(\mathbf{h}, \mathbf{e}_t) = \text{LN}(\mathbf{h}) \odot (1 + \gamma(\mathbf{e}_t)) + \beta(\mathbf{e}_t).$$

Here  $\gamma(\cdot)$  and  $\beta(\cdot)$  are learned linear maps specific to the node type. We also apply analogous FiLM-style time modulation after each heterogeneous interaction block. This ensures that the same molecular geometry can be interpreted differently depending on whether the state is close to the apo start, close to the holo target, or in the middle of the interpolation path.

## B.4 Heterogeneous Interaction Blocks

The encoder consists of multiple heterogeneous interaction blocks. Each block updates ligand atom states, receptor atom states, and receptor residue states using the corresponding projected edge features and edge indices. Within each block, messages are computed separately for different edge types and aggregated into the appropriate target node type. Residual connections and layer normalization are applied to stabilize training.

After the final interaction block, we denote the output embeddings by  $\mathbf{h}_j^{\text{lig}}$  for ligand atom  $j$ ,  $\mathbf{h}_a^{\text{rec-atom}}$  for receptor atom  $a$ , and  $\mathbf{h}_i^{\text{rec-res}}$  for receptor residue  $i$ . These embeddings are used by all prediction heads. We also compute pooled graph-level summaries by mean pooling ligand atom embeddings, receptor atom embeddings, and receptor residue embeddings over each complex in the batch. These pooled summaries are normalized and concatenated with the time embedding to form the global complex context.

## B.5 Local Ligand-Receptor Context

In addition to global pooling, PIGEON builds a local ligand-receptor context from ligand-receptor cross edges. This context is designed to summarize receptor atoms near the current ligand placement. A ligand-derived query attends over nearby receptor atom embeddings, producing a local pocket-context vector. This vector is combined with pooled ligand and receptor summaries and the explicit time embedding to form the sample-level feature used by several output heads. This local context is important because translation, rotation, and receptor flexibility are primarily determined by the ligand’s immediate receptor environment rather than by the entire protein equally.

## B.6 Pocket Field Tensor and Invariant Summaries

The graph builder computes a ligand-centered pocket field from receptor geometry before running the translation module. For receptor atom  $i$  with current position  $\mathbf{q}_{i,t}$ , let  $b_i^{\text{local}}$  denote the local receptor-intrinsic buriedness score computed by probe occlusion. We use the normalized pocket score

$$s_i^{\text{pocket}} = \text{clip}\left(\frac{b_i^{\text{local}}}{b_{\text{ref}}}, 0, 1\right), \quad (2)$$

where  $b_{\text{ref}}$  is a fixed reference buriedness value.

At the current ligand centroid  $\mathbf{c}_t^{\text{lig}}$ , define

$$\mathbf{d}_{i,t} = \mathbf{q}_{i,t} - \mathbf{c}_t^{\text{lig}}, \quad r_{i,t} = \|\mathbf{d}_{i,t}\|, \quad \hat{\mathbf{d}}_{i,t} = \frac{\mathbf{d}_{i,t}}{r_{i,t} + \epsilon}. \quad (3)$$

We use six radial shells with centers

$$R = (1.5, 2.5, 4.0, 8.0, 14.0, 20.0) \quad (4)$$

and widths

$$\sigma = (0.75, 1.0, 2.0, 3.0, 4.0, 5.0). \quad (5)$$

The soft shell membership is

$$\phi_s(r_{i,t}) = \exp\left(-\frac{(r_{i,t} - R_s)^2}{2\sigma_s^2}\right). \quad (6)$$

For outer shells, the atom weight is

$$w_{is}^{\text{std}} = s_i^{\text{pocket}} \phi_s(r_{i,t}), \quad (7)$$

while for the two inner shells we use a gradient-like radial upweighting

$$w_{is}^{\text{grad}} = \frac{s_i^{\text{pocket}} \phi_s(r_{i,t})}{r_{i,t} + \epsilon_r}. \quad (8)$$

The implementation uses  $\epsilon_r = 0.5 \text{ \AA}$ . The effective shell weight is

$$w_{is} = w_{is}^{\text{grad}} \mathbf{1}[s \in \{1, 2\}] + w_{is}^{\text{std}} \mathbf{1}[s \in \{3, 4, 5, 6\}]. \quad (9)$$

For inner shells, weights are normalized by the total shell mass, while outer shells use raw weights:

$$\tilde{w}_{is} = \frac{w_{is}}{\sum_i w_{is} + \epsilon} \mathbf{1}[s \in \{1, 2\}] + w_{is} \mathbf{1}[s \in \{3, 4, 5, 6\}]. \quad (10)$$

For the real  $l = 2$  components, we use

$$Y_2(\hat{\mathbf{d}}) = [xy, yz, 3z^2 - 1, xz, x^2 - y^2], \quad (11)$$

where  $\hat{\mathbf{d}} = (x, y, z)$ . The shellwise spherical harmonic coefficients are

$$f_0^{(s)} = \sum_i \tilde{w}_{is}, \quad (12)$$

$$\mathbf{f}_1^{(s)} = \sum_i \tilde{w}_{is} \hat{\mathbf{d}}_{i,t}, \quad (13)$$

$$\mathbf{f}_2^{(s)} = \sum_i \tilde{w}_{is} Y_2(\hat{\mathbf{d}}_{i,t}). \quad (14)$$

The shellwise dipole norm, dipole direction, and quadrupole power are

$$n_1^{(s)} = \|\mathbf{f}_1^{(s)}\|, \quad \hat{\mathbf{f}}_1^{(s)} = \frac{\mathbf{f}_1^{(s)}}{n_1^{(s)} + \epsilon}, \quad P_2^{(s)} = \|\mathbf{f}_2^{(s)}\|_2^2. \quad (15)$$

The full shellwise spherical harmonic tensor is

$$\mathbf{F}_{\text{SH}}(\mathbf{c}_t^{\text{lig}}) = \bigoplus_{s=1}^6 [f_0^{(s)}, \mathbf{f}_1^{(s)}, \mathbf{f}_2^{(s)}] \in \mathbb{R}^{54}. \quad (16)$$

The six shellwise dipole directions  $\{\hat{\mathbf{f}}_1^{(s)}\}_{s=1}^6$  form the first six translation candidates.

For the multi-scale pocket-potential branch, we use five preferred distances

$$r_{\text{opt}} = (2.0, 3.5, 5.0, 8.0, 12.0) \quad (17)$$

with widths

$$\sigma_{\text{pot}} = (1.0, 1.5, 2.0, 3.0, 5.0). \quad (18)$$

For potential scale  $a$ , define

$$\psi_a(r_{i,t}) = \exp\left(-\frac{(r_{i,t} - r_{\text{opt},a})^2}{2\sigma_{\text{pot},a}^2}\right). \quad (19)$$

The raw pocket-potential gradient is

$$\mathbf{g}_a = \sum_i s_i^{\text{pocket}} \psi_a(r_{i,t}) \frac{r_{i,t} - r_{\text{opt},a}}{\sigma_{\text{pot},a}^2} \hat{\mathbf{d}}_{i,t}. \quad (20)$$

We store its norm, unit direction, and log-norm feature:

$$n_a^{\text{pot}} = \|\mathbf{g}_a\|, \quad \hat{\mathbf{g}}_a = \frac{\mathbf{g}_a}{n_a^{\text{pot}} + \epsilon}, \quad \eta_a = \log(1 + n_a^{\text{pot}}). \quad (21)$$

The five normalized vectors  $\{\hat{\mathbf{g}}_a\}_{a=1}^5$  form the next five translation candidates.

The pocket field also includes four distance-to-surface features. Let  $\tau_b$  be a buriedness threshold. The nearest-buried distance is

$$d_{\text{buried}} = \min_{i: s_i^{\text{pocket}} > \tau_b} r_{i,t}. \quad (22)$$

If no receptor atom passes the threshold, we use the centroid-distance fallback. For a finite radius  $R_c$ , define the local pocket-score-weighted centroid

$$\bar{\mathbf{q}}_{R_c} = \frac{\sum_{i: r_{i,t} < R_c} s_i^{\text{pocket}} \mathbf{q}_{i,t}}{\sum_{i: r_{i,t} < R_c} s_i^{\text{pocket}} + \epsilon}, \quad d_{\text{centroid}} = \|\mathbf{c}_t^{\text{lig}} - \bar{\mathbf{q}}_{R_c}\|. \quad (23)$$

For a near radius  $R_n$  and Gaussian weight

$$\varphi_i = \exp\left(-\frac{1}{2} \left(\frac{r_{i,t}}{\sigma_n}\right)^2\right), \quad (24)$$

the local buriedness mean and variance are

$$\mu_b = \frac{\sum_{i: r_{i,t} < R_n} \varphi_i s_i^{\text{pocket}}}{\sum_{i: r_{i,t} < R_n} \varphi_i + \epsilon}, \quad (25)$$

$$\text{Var}_{\text{local}} = \frac{\sum_{i: r_{i,t} < R_n} \varphi_i (s_i^{\text{pocket}} - \mu_b)^2}{\sum_{i: r_{i,t} < R_n} \varphi_i + \epsilon}. \quad (26)$$

For local anisotropy, within radius  $R_a$  we compute

$$\boldsymbol{\mu}_a = \frac{\sum_{i: r_{i,t} < R_a} s_i^{\text{pocket}} \mathbf{q}_{i,t}}{\sum_{i: r_{i,t} < R_a} s_i^{\text{pocket}} + \epsilon}, \quad (27)$$

$$C_a = \frac{\sum_{i: r_{i,t} < R_a} s_i^{\text{pocket}} (\mathbf{q}_{i,t} - \boldsymbol{\mu}_a)(\mathbf{q}_{i,t} - \boldsymbol{\mu}_a)^\top}{\sum_{i: r_{i,t} < R_a} s_i^{\text{pocket}} + \epsilon}. \quad (28)$$

If  $\lambda_{\min}$  and  $\lambda_{\max}$  are the smallest and largest eigenvalues of  $C_a$ , then

$$A_{\text{local}} = 1 - \frac{\lambda_{\min}}{\lambda_{\max} + \epsilon}. \quad (29)$$

The distance-to-surface feature vector is

$$\mathbf{d}_{\text{surf}} = [d_{\text{buried}}, d_{\text{centroid}}, \text{Var}_{\text{local}}, A_{\text{local}}]. \quad (30)$$

The invariant pocket summary used by the magnitude head is

$$\mathbf{z}_{\text{pocket inv}} = [f_0^{(1)}, \dots, f_0^{(6)}, P_2^{(1)}, \dots, P_2^{(6)}, \eta_1, \dots, \eta_5, \mathbf{d}_{\text{surf}}]. \quad (31)$$

Thus, the translation module uses both directional pocket cues, through  $\hat{\mathbf{f}}_1^{(s)}$  and  $\hat{\mathbf{g}}_a$ , and invariant pocket summaries, through  $\mathbf{z}_{\text{pocket inv}}$ .

## B.7 Candidate Direction Bank Implementation

The candidate direction bank is constructed deterministically from the pocket geometry at the current ligand centroid. Six directions are obtained from the  $L = 1$  spherical harmonic dipoles, one per radial shell. Five directions are obtained from the multi-scale pocket-potential gradients. Two directions are obtained from long-range buriedness cues: the direction to the buriedness-weighted receptor centroid and the direction to the nearest highly buried receptor atom. All candidate directions are normalized to unit length with an  $\epsilon$  stabilizer. If a direction has near-zero norm, the implementation uses numerical stabilization to avoid invalid values.

The final candidate bank has 13 directions. This bank is recomputed for every interpolated state  $x_t$ , so the candidate directions are ligand-position dependent. As the ligand moves along the interpolation path, the spherical harmonic field and pocket-potential gradients change accordingly, allowing the candidate bank to adapt to the current local docking geometry.

## B.8 Learnable Staged Selector Implementation

The staged direction selector is implemented as a small set of learnable coefficients rather than a free neural network over candidate directions. For each candidate and probe radius, the graph builder computes five atom-aware features: contact quality, clash score, void quality, uniformity, and contact count. These features are z-scored across the 13 candidates at each probe radius. The selector then learns feature weights and probe-radius weights separately for the ingress and fit regimes.

The selector contains four softmax-normalized coefficient vectors: ingress feature weights, fit feature weights, ingress probe-radius weights, and fit probe-radius weights. The selector also contains a learnable stage gate that maps the current-state stage-feature vector to a scalar  $\alpha_{\text{fit}} \in [0, 1]$ . The final candidate logit is the convex interpolation between ingress and fit scores. All selector coefficients are initialized from the hand-designed geometric rule, and the training objective includes a regularization term that penalizes drift away from the initialization. This design lets the model learn from data while preserving the interpretability of the original rule.

## B.9 Stage-Feature Definitions

The learnable staged selector uses a five-dimensional current-state feature vector

$$\mathbf{z}_{\text{stage}} = [z_d, z_q, z_c, z_u, z_n]. \quad (32)$$

These features indicate whether the current ligand state is in an ingress-like regime, where the ligand should move toward a pocket, or a local-fit regime, where the ligand should refine contacts near the current pose.

We first compute a local pocket centroid around the current ligand centroid. For receptor atom  $i$ , define

$$\omega_i^{\text{stage}} = s_i^{\text{pocket}} \exp\left(-\frac{1}{2} \left(\frac{\|\mathbf{q}_{i,t} - \mathbf{c}_t^{\text{lig}}\|}{\sigma_{\text{stage}}}\right)^2\right) \mathbf{1}[\|\mathbf{q}_{i,t} - \mathbf{c}_t^{\text{lig}}\| < R_{\text{stage}}]. \quad (33)$$

The stage pocket centroid is

$$\bar{\mathbf{q}}_{\text{stage}} = \frac{\sum_i \omega_i^{\text{stage}} \mathbf{q}_{i,t}}{\sum_i \omega_i^{\text{stage}} + \epsilon}. \quad (34)$$

If the total weight is numerically zero, we set  $\bar{\mathbf{q}}_{\text{stage}} = \mathbf{c}_t^{\text{lig}}$ . Let the current ligand radius be

$$r_{\text{lig}} = \max_j \|\mathbf{x}_{j,t} - \mathbf{c}_t^{\text{lig}}\|. \quad (35)$$

The normalized pocket-distance feature is

$$z_d = \text{clip}\left(\frac{\|\mathbf{c}_t^{\text{lig}} - \bar{\mathbf{q}}_{\text{stage}}\|}{r_{\text{lig}} + 1}, 0, 2\right). \quad (36)$$

The remaining four stage features are computed by evaluating the same atom-aware probe feature function at the current untranslated ligand pose, i.e., at destination centroid  $\mathbf{p} = \mathbf{c}_t^{\text{lig}}$ . Let

$$[Q_{@c_t}, C_{@c_t}, V_{@c_t}, U_{@c_t}, N_{@c_t}] \quad (37)$$

denote contact quality, clash score, void quality, uniformity, and contact count at the current pose. We also retain the raw clash penalty  $P_{\text{clash}, @c_t}$  before the negative log transformation. The bounded stage signals are

$$z_q = \text{clip}\left(\frac{Q_{@c_t} - 2.5}{1.5}, 0, 1\right), \quad (38)$$

$$z_c = \text{clip}\left(\frac{P_{\text{clash}, @c_t}}{1.5}, 0, 1\right), \quad (39)$$

$$z_u = \text{clip}(U_{@c_t}, 0, 1), \quad (40)$$

$$z_n = \text{clip}\left(\frac{N_{@c_t} - 2.5}{2.0}, 0, 1\right). \quad (41)$$

Thus,  $z_q$  measures current contact quality,  $z_c$  measures current clash,  $z_u$  measures current contact uniformity, and  $z_n$  measures current contact count.

## B.10 Translation Magnitude Head

After the selector chooses a unit direction  $\hat{\mathbf{u}} = \mathbf{v}_{k^*}$ , PIGEON predicts a nonnegative translation magnitude from graph-level and pocket-invariant features. Let  $\mathbf{h}^{\text{graph}}$  denote the pooled graph context and let  $\mathbf{h}^{\text{time}}$  denote the time embedding. We first project the invariant pocket summary from Eq. 31:

$$\mathbf{h}^{\text{pocket inv}} = \phi_{\text{inv}}(\mathbf{z}_{\text{pocket inv}}). \quad (42)$$

The magnitude head input is

$$\mathbf{h}^{\text{mag}} = [\mathbf{h}^{\text{graph}}, \mathbf{h}^{\text{time}}, \mathbf{h}^{\text{pocket inv}}]. \quad (43)$$

The predicted magnitude is

$$\hat{m} = \text{softplus}(f_{\text{mag}}(\mathbf{h}^{\text{mag}})). \quad (44)$$

The final ligand translation residual is

$$\widehat{\Delta}\mathbf{c}_{t \rightarrow 1}^{\text{lig}} = \hat{m}\mathbf{v}_{k^*}. \quad (45)$$

This factorization separates the direction decision from the distance prediction: the selected direction comes from receptor-derived geometric candidates, while the scalar magnitude is learned from pooled complex context and invariant pocket-field summaries.

## B.11 Ligand Rotation and Torsion Heads

The ligand rotation head predicts a 3D  $SO(3)$  log residual. It uses pooled ligand atom embeddings, the sample-level complex context, and explicit time features. The output is interpreted as the remaining rotation vector from  $R_t^{\text{lig}}$  to  $R_1^{\text{lig}}$ .

The ligand torsion head predicts wrapped residuals for ligand rotatable bonds. For each rotatable bond, the model builds a torsion feature from the embeddings of the atoms defining the torsion and from torsion-local geometric information. A shared torsion MLP predicts one residual per valid torsion. The output is bounded with a smooth angular nonlinearity so that the predicted torsion residual remains in a physically meaningful angular range. Invalid or padded torsions are masked out of the loss.

## B.12 Receptor Frame and Side-Chain Heads

The receptor frame head predicts one 6D residual for each receptor residue. The output represents the remaining local rigid-frame motion from the current receptor frame to the target holo receptor frame. The receptor chi head predicts side-chain torsion residuals for each residue. The number of chi angles varies by residue type, so the receptor chi loss uses a residue-specific chi mask. Both receptor heads condition on receptor residue embeddings and sample-level context, allowing receptor flexibility to be predicted in a ligand-aware and time-aware manner.

## B.13 Output Dictionary and Ligand Pose Assembly

For compatibility with the loss and evaluation code, the predicted ligand pose is assembled as a 6D vector by concatenating the translation and rotation predictions:

$$\hat{\mathbf{y}}_{\text{pose}}^{\text{lig}} = [\widehat{\Delta \mathbf{c}}_{t \rightarrow 1}^{\text{lig}}; \widehat{\boldsymbol{\omega}}_{t \rightarrow 1}^{\text{lig}}].$$

The translation portion is produced by the direction selector and magnitude head, while the rotation portion is produced by the ligand rotation head. The model output dictionary also stores diagnostic quantities such as candidate logits, selector probabilities, selected candidate index, oracle candidate cosine values, predicted translation magnitude, and selector-stage gate value. These diagnostics are used to monitor whether the model is learning the intended geometric behavior.

## B.14 Loss Implementation Details

The selector is trained with a soft listwise ranking objective. For each sample, the oracle score for candidate  $k$  is the cosine similarity between candidate direction  $\mathbf{v}_k$  and the normalized ground-truth translation direction. A low-temperature softmax over oracle cosines produces the target distribution. The selector logits are trained with cross entropy against this soft oracle distribution. This supervision gives high probability to all candidates close to the true direction, rather than treating direction selection as a brittle one-hot classification problem.

The magnitude head is trained in log space against the ground-truth residual translation norm. This reduces sensitivity to the absolute scale of the translation and encourages relative magnitude accuracy. Receptor frame residuals, receptor chi residuals, ligand rotation residuals, and ligand torsion residuals are trained with fixed-scale normalized Smooth L1 losses. Fixed scales are used instead of target-dependent normalization so that the loss remains stable when the target residual norm is small.

The complete loss is a weighted sum of selector ranking loss, selector regularization, translation log-magnitude loss, receptor frame loss, receptor chi loss, ligand rotation loss, and ligand torsion loss. The final weights are selected based on validation performance and diagnostic stability.

**Selector regularization.** To keep the selector close to the intended geometric rule family, we regularize the learnable feature logits, radius logits, and stage-gate parameters toward their initialized values. Let  $\boldsymbol{\theta}_{\text{feat}}^{\text{fit}}$  and  $\boldsymbol{\theta}_{\text{feat}}^{\text{ing}}$  denote the fit and ingress feature logits, and let  $\boldsymbol{\theta}_{\text{rad}}^{\text{fit}}$  and  $\boldsymbol{\theta}_{\text{rad}}^{\text{ing}}$  denote the fit and ingress probe-radius logits. The selector regularization is

$$\begin{aligned} \mathcal{L}_{\text{selector reg}} = & \|\boldsymbol{\theta}_{\text{feat}}^{\text{fit}} - \boldsymbol{\theta}_{\text{feat},0}^{\text{fit}}\|_2^2 + \|\boldsymbol{\theta}_{\text{feat}}^{\text{ing}} - \boldsymbol{\theta}_{\text{feat},0}^{\text{ing}}\|_2^2 \\ & + \|\boldsymbol{\theta}_{\text{rad}}^{\text{fit}} - \boldsymbol{\theta}_{\text{rad},0}^{\text{fit}}\|_2^2 + \|\boldsymbol{\theta}_{\text{rad}}^{\text{ing}} - \boldsymbol{\theta}_{\text{rad},0}^{\text{ing}}\|_2^2 + \mathcal{L}_{\text{gate reg}}, \end{aligned} \quad (46)$$

where

$$\mathcal{L}_{\text{gate reg}} = \|\mathbf{w}_{\text{gate}} - \mathbf{w}_{\text{gate},0}\|_2^2 + (b_{\text{gate}} - b_{\text{gate},0})^2 + (\log s_{\text{gate}} - \log s_{\text{gate},0})^2. \quad (47)$$

This regularizer allows the selector to adapt to data while preserving the inductive bias of the hand-designed ingress and fit rules.

## B.15 Training and Inference Details

We train PIGEON using the selector configuration described in the main text. All experiments use random seed 42 and are trained on four NVIDIA RTX A6000 GPUs with 48GB memory each. Training uses PyTorch Lightning distributed data parallelism with `ddp_find_unused_parameters_true`, bf16 mixed precision, batch size 1, gradient accumulation over 8 steps, and gradient clipping with global norm threshold 1.0. The model is trained for 20 epochs, with validation performed once per epoch. We use AdamW with learning rate  $1.0 \times 10^{-4}$ , weight decay 0.01,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and

935  $\epsilon = 1.0 \times 10^{-8}$ . The dataset is split by PDB ID into train, validation, and test partitions with fractions 0.70/0.15/0.15  
936 using split seed 42.

937 The model uses hidden dimension 256, edge hidden dimension 256, time embedding dimension 128, and 6 heterogeneous  
938 interaction blocks with dropout 0.0. The receptor-residue, receptor-atom, and ligand-atom input feature dimensions are  
939 1307, 6, and 523, respectively. The edge feature dimensions are 20 for receptor atom-radius edges, receptor residue-radius  
940 edges, receptor atom-to-residue edges, and ligand-receptor cross edges, and 23 for ligand bond edges. The graph builder  
941 uses a receptor atom-radius cutoff of 6.0 Å, receptor residue-radius cutoff of 12.0 Å, and ligand-receptor cross cutoff of  
942 15.0 Å. The maximum numbers of receptor atom neighbors, receptor residue neighbors, and cross-edge neighbors are 64,  
943 32, and 128, respectively. Distances are encoded with 16 radial basis functions, and dense attention masks and dense node  
944 features are included.  
945

946 For the pocket geometry module, the spherical harmonic field uses 6 radial shells with radii (1.5, 2.5, 4.0, 8.0, 14.0, 20.0) Å  
947 and widths (0.75, 1.0, 2.0, 3.0, 4.0, 5.0) Å. We use real spherical harmonics up to  $L = 2$ , giving a 54-dimensional shell field.  
948 Buriedness is computed with 48 probe directions, a buriedness probe radius of 6.0 Å, local and global probe radii of 5.0 Å  
949 and 10.0 Å, an occlusion radius scale of 0.8, neighbor exclusion margin 0.5 Å, and buriedness reference value 0.3. The  
950 inner two shells are used for gradient-like SH weighting with  $\epsilon = 0.5$ . The pocket-potential branch uses five preferred radii  
951 (2.0, 3.5, 5.0, 8.0, 12.0) Å with Gaussian widths (1.0, 1.5, 2.0, 3.0, 5.0) Å. The distance-to-surface branch uses buriedness  
952 threshold 0.25, centroid radius 20.0 Å, near radius 6.0 Å, near sigma 2.0 Å, and anisotropy radius 6.0 Å.  
953

954 The translation module uses pocket-guided translation with a 13-direction candidate bank. The selector probe radii are  
955 (0.5, 1.0, 1.5, 2.5, 4.0) Å. The selector uses contact-probe radius 10.0 Å, buriedness threshold 0.15, preferred contact gap  
956 0.5 Å, contact sigma 0.75 Å, contact-shell inner and outer gaps 0.0 Å and 1.5 Å, void-gap threshold 2.5 Å, stage radius  
957 12.0 Å, and stage sigma 6.0 Å. Candidate-type bias is disabled. The translation magnitude head uses translation hidden  
958 dimension 256, pocket SH field dimension 74, pocket distance feature dimension 4, pocket invariant projection dimension  
959 256, and translation time dropout 0.25.

960 The interpolant uses residual targets with polynomial schedule  $\alpha(t) = t^2$  and time stabilizer  $10^{-4}$ . Intermediate states are  
961 sampled with uniform-alpha time sampling. Clearance projection is disabled in the current setting. Although declashing and  
962 cleanup hyperparameters remain specified in the configuration for compatibility, they are not used as a clearance-projection  
963 step during training because `use_clearance_projection` is set to false.  
964

965 The training objective uses Smooth L1 residual losses and the structured selector losses described in the main text. The  
966 receptor frame, receptor chi, and ligand torsion losses each have weight 1.0. The ligand pose loss weight is set to 0.0,  
967 while ligand rotation diagnostics are still computed with a rotation-specific weight parameter. Direct Cartesian translation  
968 loss, translation direction cosine loss, geometry-magnitude loss, steer-direction loss, geometry-direction loss, and residual  
969 regularization are disabled. Translation is trained through the selector rank loss and log-magnitude loss: the selector rank loss  
970 has weight 1.0 with oracle temperature 0.05, the selector coefficient regularization has weight 100, and the log-magnitude  
971 loss has weight 1.0 with a 300-step warmup and  $\epsilon = 10^{-4}$ . Fixed normalization scales are 5.0 for receptor frame residuals,  
972 0.05 for receptor chi residuals, 0.25 for ligand rotation residuals, and 0.25 for ligand torsion residuals.

973 At inference time, we use the same configuration, and model predictions as the residual source. Unless otherwise stated,  
974 inference is performed on the test split with 4 integration steps and terminal time  $t_{\text{end}} = 0.85$ . We run PoseBusters v2  
975 during inference to compute physical-validity diagnostics. The ODE inference command uses one GPU and writes predicted  
976 structures and evaluation outputs to the specified checkpoint output directory.  
977

## 978 C Algorithms

979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989

**Algorithm 1** Pocket-Aware Translation Direction Selection

990 **Algorithm 1** Pocket-Aware Translation Direction Selection  
 991 1: **Input:** interpolated complex state  $x_t$  at time  $t$   
 992 2: receptor atom positions  $\{\mathbf{q}_{a,t}\}$  with vdW radii  $\{r_a^{\text{rec}}\}$   
 993 3: ligand atom positions  $\{\mathbf{x}_{j,t}\}$  with vdW radii  $\{r_j^{\text{lig}}\}$ , centroid  $\mathbf{c}_t^{\text{lig}}$   
 994 4: learnable selector parameters  $\theta_{\text{sel}} = (\boldsymbol{\eta}^{\text{ing}}, \boldsymbol{\eta}^{\text{fit}}, \boldsymbol{\zeta}^{\text{ing}}, \boldsymbol{\zeta}^{\text{fit}}, \mathbf{w}_{\text{gate}}, b_{\text{gate}}, s_{\text{gate}})$   
 995 5: **Hyperparameters:** probe directions  $K_{\text{probe}}$ , shells  $M = 6$ , harmonic order  $L = 2$ , probe schedule  $\mathcal{R} =$   
 996  $\{0.5, 1.0, 1.5, 2.5, 4.0\}$  Å, potential scales, buriedness threshold  $\tau_b$ , contact gaps  $(\gamma_{\text{in}}, \gamma_{\text{out}}, \gamma_{\text{pref}}, \gamma_{\text{void}})$   
 997  $\triangleright$  Step 1: receptor-intrinsic buriedness, independent of ligand pose  
 998 6:  
 999 7: **for** each receptor atom  $a$  **do**  
 1000 8: | Cast  $K_{\text{probe}}$  probe directions  $\{\mathbf{u}_k\}$  around  $\mathbf{q}_{a,t}$   
 1001 9: | Compute occlusion indicators  $\{o_{ak}\}$   $\triangleright$  Eq. 1  
 1002 10: |  $b_a \leftarrow \frac{1}{K_{\text{probe}}} \sum_k o_{ak}$   
 1003 11: **end for**  
 1004 12:  $\triangleright$  Step 2: ligand-centered spherical harmonic pocket field  
 1005 13: **for** each shell  $s = 1, \dots, M$  **do**  
 1006 14: | **for** each harmonic degree  $\ell = 0, 1, 2$  **do**  
 1007 15: | |  $\mathbf{f}_\ell^{(s)}(\mathbf{c}_t^{\text{lig}}) \leftarrow \sum_a b_a \phi_s(r_{a,t}) \mathbf{Y}_\ell(\hat{\mathbf{r}}_{a,t})$   
 1008 16: | **end for**  
 1009 17: **end for**  
 1010 18: Construct invariant pocket summaries  $\mathbf{F}_{\text{pocket}}$  from shell masses, SH powers, potential strengths, and distance-to-surface  
 1011 features  
 1012 19:  $\triangleright$  Step 3: build 13-direction candidate bank  $\mathcal{V}_t$   
 1013 20:  $\mathcal{V}_t \leftarrow \emptyset$   
 1014 21: **for**  $s = 1, \dots, 6$  **do**  $\triangleright$  shellwise dipole directions  
 1015 22: |  $\mathbf{v}_s^{\text{SH}} \leftarrow \mathbf{d}_s^{\text{SH}} / (\|\mathbf{d}_s^{\text{SH}}\| + \epsilon)$   
 1016 23: |  $\mathcal{V}_t \leftarrow \mathcal{V}_t \cup \{\mathbf{v}_s^{\text{SH}}\}$   
 1017 24: **end for**  
 1018 25: **for** each potential scale  $r_{\text{opt}}$  **do**  $\triangleright$  five gradient candidates  
 1019 26: |  $\mathbf{v}_{r_{\text{opt}}}^{\text{pot}} \leftarrow \nabla_{\mathbf{c}_t^{\text{lig}}} U_{r_{\text{opt}}} / (\|\nabla_{\mathbf{c}_t^{\text{lig}}} U_{r_{\text{opt}}}\| + \epsilon)$   
 1020 27: |  $\mathcal{V}_t \leftarrow \mathcal{V}_t \cup \{\mathbf{v}_{r_{\text{opt}}}^{\text{pot}}\}$   
 1021 28: **end for**  
 1022 29: Add long-range directions  $\mathbf{v}_{\text{centroid}}$  and  $\mathbf{v}_{\text{nearest}}$  to  $\mathcal{V}_t$   
 1023 30:  $\triangleright$  Step 4: probe each candidate at every probe distance  
 1024 31: **for** each candidate  $\mathbf{v}_k \in \mathcal{V}_t$  **do**  
 1025 32: | **for** each probe distance  $\rho_s \in \mathcal{R}$  **do**  
 1026 33: | |  $\mathbf{p}_{k,s} \leftarrow \mathbf{c}_t^{\text{lig}} + \rho_s \mathbf{v}_k$   
 1027 34: | |  $\tilde{\mathbf{x}}_{j,k,s} \leftarrow \mathbf{x}_{j,t} + \rho_s \mathbf{v}_k$  for all ligand atoms  $j$   
 1028 35: | | Build local receptor set  $\mathcal{A}_{k,s}$  and compute surface gaps  $g_{aj}^{(k,s)}$   
 1029 36: | | Compute features  $F_{\text{contact}}, F_{\text{clash}}, F_{\text{void}}, F_{\text{uniform}}, F_{\text{count}}$   
 1030 37: | **end for**  
 1031 38: **end for**  
 1032 39:  $\triangleright$  Step 5: within-sample normalization across candidates  
 1033 40: **for** each probe distance  $s$  and feature  $f$  **do**  
 1034 41: |  $Z(k, s, f) \leftarrow (F_f(k, s) - \mu_{s,f}) / \sigma_{s,f}$  for all candidates  $k$   
 1035 42: **end for**  
 1036 43:  $\triangleright$  Step 6: stage gate from current untranslated ligand pose  
 1037 44: Compute  $\mathbf{z}_{\text{stage}} = [z_d, z_q, z_c, z_u, z_n]$   $\triangleright$  App. B.9  
 1038 45:  $\alpha_{\text{fit}} \leftarrow \sigma(s_{\text{gate}}(\mathbf{w}_{\text{gate}} \mathbf{z}_{\text{stage}} + b_{\text{gate}}))$   
 1039 46:  $\triangleright$  Step 7: compute ingress/fit scores and select direction  
 1040 47:  $\boldsymbol{\lambda}^{\text{ing}}, \boldsymbol{\lambda}^{\text{fit}} \leftarrow \text{softmax}(\boldsymbol{\eta}^{\text{ing}}), \text{softmax}(\boldsymbol{\eta}^{\text{fit}})$   
 1041 48:  $\boldsymbol{\rho}^{\text{ing}}, \boldsymbol{\rho}^{\text{fit}} \leftarrow \text{softmax}(\boldsymbol{\zeta}^{\text{ing}}), \text{softmax}(\boldsymbol{\zeta}^{\text{fit}})$   
 1042 49: **for** each candidate  $k$  **do**  
 1043 50: |  $s_k^{\text{ing}} \leftarrow \sum_s \rho_s^{\text{ing}} \sum_f \lambda_f^{\text{ing}} Z(k, s, f)$   
 1044 51: |  $s_k^{\text{fit}} \leftarrow \sum_s \rho_s^{\text{fit}} \sum_f \lambda_f^{\text{fit}} Z(k, s, f)$   
 52: |  $\ell_k \leftarrow (1 - \alpha_{\text{fit}}) s_k^{\text{ing}} + \alpha_{\text{fit}} s_k^{\text{fit}}$  19  
 53: **end for**  
 54:  $k^* \leftarrow \arg \max_k \ell_k; \hat{\mathbf{u}} \leftarrow \mathbf{v}_{k^*}$   
 55: **return** selected direction  $\hat{\mathbf{u}}$ , candidate logits  $\{\ell_k\}$ , candidate bank  $\mathcal{V}_t$ , pocket summaries  $\mathbf{F}_{\text{pocket}}$

---

**Algorithm 2 PIGEON Training Step on a Sampled Apo-Holo Pair**

---

1045  
1046  
1047  
1048  
1049  
1050 1: **Input:** dataset  $\mathcal{D}$  of apo-holo state pairs  $(x_0, x_1)$   
1051 2: heterogeneous message-passing network  $f_\theta$  with output heads  
1052 3: rule-based selector initialization  $\theta_0$   
1053 4: **Hyperparameters:** schedule  $\alpha(t) = t^2$ , oracle temperature  $T_{\text{oracle}}$ , loss weights  
1054  $\{\lambda_{\text{rank}}, \lambda_{\text{reg}}, \lambda_{\text{mag}}, \lambda_{\text{frame}}, \lambda_\chi, \lambda_{\text{rot}}, \lambda_{\text{tor}}\}$ , residual scales  $\{s_{\text{frame}}, s_\chi, s_{\text{rot}}, s_{\text{tor}}\}$   
1055 5:  $\triangleright$  Sample apo-holo pair and interpolation time  
1056 6: Sample  $(x_0, x_1) \sim \mathcal{D}$ ;  $\alpha \sim \mathcal{U}(0, 1)$ ;  $t \leftarrow \sqrt{\alpha}$   
1057 7:  $\triangleright$  Geodesic interpolation on the product manifold  
1058 8:  $x_t \leftarrow \text{Exp}_{x_0}(\alpha \text{Log}_{x_0}(x_1))$   $\triangleright$  component-wise on  $\mathcal{M}$   
1059 9:  $\triangleright$  Encode complex with time-conditioned heterogeneous GNN  
1060 10:  $\{\mathbf{h}_j^{\text{lig}}, \mathbf{h}_a^{\text{rec-atom}}, \mathbf{h}_i^{\text{rec-res}}\} \leftarrow f_\theta(x_t, t)$   
1061 11: Let  $\mathbf{h}$  denote pooled ligand, receptor-atom, receptor-residue, local-context, and time-conditioned features  
1062 12:  $\triangleright$  Predict ligand translation via pocket-aware selector  
1063 13:  $\hat{\mathbf{u}}, \{\ell_k\}_{k=1}^K, \{\mathbf{v}_k\}_{k=1}^K, \mathbf{F}_{\text{pocket}} \leftarrow \text{Algorithm 1}(x_t, t, \theta_{\text{sel}})$   
1064 14:  $\hat{m} \leftarrow \text{Softplus}(\text{MagnitudeHead}(\mathbf{h}, t, \mathbf{F}_{\text{pocket}}))$   
1065 15:  $\hat{\Delta}\mathbf{c}_{t \rightarrow 1}^{\text{lig}} \leftarrow \hat{m} \hat{\mathbf{u}}$   
1066 16:  $\triangleright$  Predict remaining residuals from output heads  
1067 17:  $\hat{\omega}^{\text{lig}} \leftarrow \text{LigandRotationHead}(\mathbf{h})$   
1068 18:  $\hat{\chi}^{\text{lig}} \leftarrow \text{LigandTorsionHead}(\mathbf{h})$   
1069 19:  $\hat{\Delta}^{\text{frame}} \leftarrow \text{FrameHead}(\mathbf{h})$   
1070 20:  $\hat{\chi}^{\text{rec}} \leftarrow \text{ChiHead}(\mathbf{h})$   
1071 21:  $\triangleright$  Compute ground-truth residuals from  $x_t$  to  $x_1$   
1072 22:  $\Delta\mathbf{c}_{t \rightarrow 1}^{\text{lig},*} \leftarrow \mathbf{c}_1^{\text{lig}} - \mathbf{c}_t^{\text{lig}}$   
1073 23:  $\mathbf{u}^* \leftarrow \Delta\mathbf{c}_{t \rightarrow 1}^{\text{lig},*} / (\|\Delta\mathbf{c}_{t \rightarrow 1}^{\text{lig},*}\| + \epsilon)$   
1074 24: Compute  $\omega_{t \rightarrow 1}^{\text{lig},*}$ , ligand torsion residuals, receptor frame residuals, and receptor side-chain residuals  
1075 25:  $\triangleright$  Selector loss: listwise rank against soft oracle  
1076 26: **for** each candidate  $k$  **do**  
1077 27: |  $q_k^* \leftarrow \mathbf{v}_k^\top \mathbf{u}^*$   
1078 28: **end for**  
1079 29: **for** each candidate  $k$  **do**  
1080 30: |  $p_k^* \leftarrow \exp(q_k^*/T_{\text{oracle}}) / \sum_j \exp(q_j^*/T_{\text{oracle}})$   
1081 31: **end for**  
1082 32:  $\mathcal{L}_{\text{rank}} \leftarrow -\sum_k p_k^* \log(\exp(\ell_k) / \sum_j \exp(\ell_j))$   
1083 33:  $\mathcal{L}_{\text{reg}} \leftarrow \|\theta_{\text{sel}} - \theta_0\|_2^2$   
1084 34:  $\triangleright$  Magnitude loss in log space  
1085 35:  $\mathcal{L}_{\text{mag}} \leftarrow (\log(\hat{m} + \epsilon) - \log(\|\Delta\mathbf{c}_{t \rightarrow 1}^{\text{lig},*}\| + \epsilon))^2$   
1086 36:  $\triangleright$  Fixed-scale residual losses for non-translation outputs  
1087 37:  $\mathcal{L}_{\text{rot}} \leftarrow \text{SmoothL1}(\hat{\omega}^{\text{lig}}/s_{\text{rot}}, \omega_{t \rightarrow 1}^{\text{lig},*}/s_{\text{rot}})$   
1088 38:  $\mathcal{L}_{\text{tor}}, \mathcal{L}_{\text{frame}}, \mathcal{L}_\chi \leftarrow$  analogous SmoothL1 residual losses with fixed scales  
1089 39:  $\triangleright$  Backpropagate combined objective  
1090 40:  $\mathcal{L} \leftarrow \lambda_{\text{rank}}\mathcal{L}_{\text{rank}} + \lambda_{\text{reg}}\mathcal{L}_{\text{reg}} + \lambda_{\text{mag}}\mathcal{L}_{\text{mag}} + \lambda_{\text{frame}}\mathcal{L}_{\text{frame}} + \lambda_\chi\mathcal{L}_\chi + \lambda_{\text{rot}}\mathcal{L}_{\text{rot}} + \lambda_{\text{tor}}\mathcal{L}_{\text{tor}}$   
1091 41: Take an optimizer step on  $\nabla_\theta \mathcal{L}$   
1092 42: **return** updated parameters  $\theta$

---

---

**Algorithm 3 PIGEON Inference: Apo-to-Holo Sampling via Manifold Flow Integration**

---

```

1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111 1: Input: initial apo-like state  $x_0 \in \mathcal{M}$ , trained network  $f_\theta$ 
1112 2: Hyperparameters: number of integration steps  $T$ , terminal time  $t_{\text{end}}$ , schedule  $\alpha(t) = t^2$ , ensemble size  $S$ 
1113 3:  $\mathcal{E} \leftarrow \emptyset$   $\triangleright$  predicted holo ensemble
1114 4: for  $s = 1, \dots, S$  do
1115 5:   Sample or choose apo-like initialization  $x^{(0)}$  from  $x_0$ 
1116 6:   Set  $\alpha_n \leftarrow \frac{n}{T}\alpha(t_{\text{end}})$  and  $t_n \leftarrow \sqrt{\alpha_n}$  for  $n = 0, \dots, T$ 
1117 7:    $\triangleright$  Iteratively integrate the manifold residual field
1118 8:   for  $n = 0, \dots, T - 1$  do
1119 9:      $\{\mathbf{h}_j^{\text{lig}}, \mathbf{h}_a^{\text{rec-atom}}, \mathbf{h}_i^{\text{rec-res}}\} \leftarrow f_\theta(x^{(n)}, t_n)$ 
1120 10:    Let  $\mathbf{h}$  denote pooled ligand, receptor-atom, receptor-residue, local-context, and time-conditioned features
1121 11:     $\triangleright$  Predict ligand translation residual
1122 12:     $\hat{\mathbf{u}}, \cdot, \cdot, \mathbf{F}_{\text{pocket}} \leftarrow \text{Algorithm 1}(x^{(n)}, t_n, \theta_{\text{sel}})$ 
1123 13:     $\hat{m} \leftarrow \text{Softplus}(\text{MagnitudeHead}(\mathbf{h}, t_n, \mathbf{F}_{\text{pocket}}))$ 
1124 14:     $\hat{\Delta}\mathbf{c}^{\text{lig}} \leftarrow \hat{m} \hat{\mathbf{u}}$ 
1125 15:     $\triangleright$  Predict remaining endpoint residuals
1126 16:     $\hat{\omega}^{\text{lig}} \leftarrow \text{LigandRotationHead}(\mathbf{h})$ 
1127 17:     $\hat{\chi}^{\text{lig}} \leftarrow \text{LigandTorsionHead}(\mathbf{h})$ 
1128 18:     $\hat{\Delta}^{\text{frame}} \leftarrow \text{FrameHead}(\mathbf{h})$ 
1129 19:     $\hat{\chi}^{\text{rec}} \leftarrow \text{ChiHead}(\mathbf{h})$ 
1130 20:     $\triangleright$  Convert endpoint residuals to this integration step
1131 21:     $\gamma_n \leftarrow (\alpha_{n+1} - \alpha_n) / (1 - \alpha_n + \epsilon)$ 
1132 22:     $\triangleright$  Update each manifold component
1133 23:     $\mathbf{c}_{\text{lig}}^{(n+1)} \leftarrow \mathbf{c}_{\text{lig}}^{(n)} + \gamma_n \hat{\Delta}\mathbf{c}^{\text{lig}}$ 
1134 24:     $R_{\text{lig}}^{(n+1)} \leftarrow R_{\text{lig}}^{(n)} \exp(\gamma_n \hat{\omega}^{\text{lig}})$ 
1135 25:     $\chi_{\text{lig}}^{(n+1)} \leftarrow \text{wrap}(\chi_{\text{lig}}^{(n)} + \gamma_n \hat{\chi}^{\text{lig}})$ 
1136 26:    Update receptor residue frames and side-chain torsions analogously using  $\hat{\Delta}^{\text{frame}}$  and  $\hat{\chi}^{\text{rec}}$ 
1137 27:    Assemble  $x^{(n+1)}$  from updated components
1138 28:   end for
1139 29:    $\mathcal{E} \leftarrow \mathcal{E} \cup \{x^{(T)}\}$ 
1140 30: end for
1141 31: return predicted holo ensemble  $\mathcal{E}$ 

```

---