CORAL: Learning Amyloid Fibril Ligand Docking with Cooperative Binding Rewards

Yasheng Sun[1*†] **Bohan Li**[2*] **Youqi Tao**[3†] **Jürgen Schmidhuber**[1]

[1]Center of Excellence for Generative AI, KAUST
[2]Department of Biomedical Engineering, National University of Singapore
[3]Stern Laboratory, Brigham and Women's Hospital, Harvard University

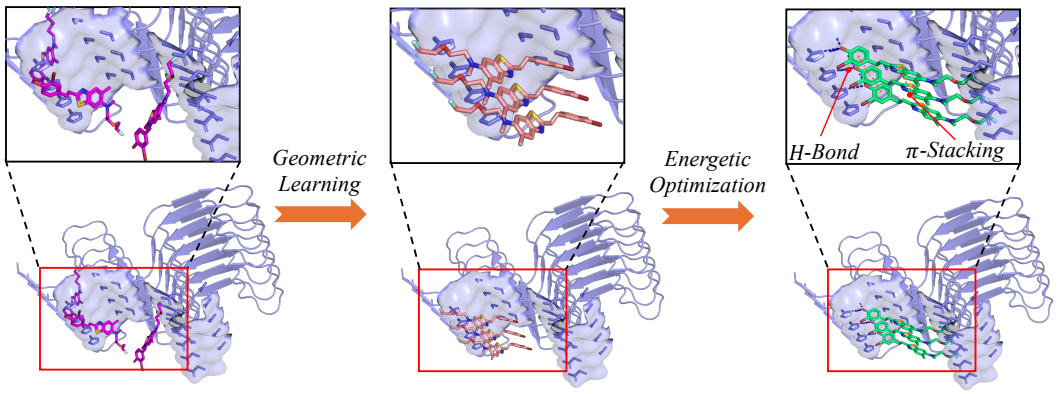

Figure 1: **Overview of CORAL.** A two-stage training recipe progressively shapes the model's ligand-pose distribution on amyloid fibrils. The *geometric learning* stage (left to middle) supervises the model on curated stacking poses, teaching it to align ligands along the cross-β groove into periodic stacking arrangements. The *energetic optimisation* stage (middle to right) further refines the pose distribution via online reinforcement learning under a composite reward that combines protein–ligand binding affinity with cooperative inter-ligand stacking energy, leading to tighter contacts such as the H-bond and π–π stacking interactions highlighted in the right inset.

Abstract

A hallmark of neurodegenerative diseases such as Alzheimer's and Parkinson's is the aberrant aggregation of proteins into amyloid fibrils, and small molecules that selectively bind to these fibrils hold promise as diagnostics, imaging probes, and therapeutics. Predicting how such ligands bind to fibril targets, however, presents two fundamental challenges. First, resolved co-crystal structures of amyloid–ligand complexes are exceptionally scarce; even with recent advances in cryo-EM only a handful have been structurally characterized, making supervised training of docking models impractical for this target class. Second, amyloid fibrils present a binding mode fundamentally different from globular proteins: ligands intercalate into longitudinal cross-β grooves and stack cooperatively along the fibril axis, a geometry that existing docking models are not designed to capture. To address these challenges, we present **CORAL** (**CO**ope**R**ative **A**myloid **L**igand docking), a reinforcement learning framework that trains a generative docking model to produce ligand pose

*Equal contribution.
†Corresponding authors.

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

distributions tailored to the cross-β groove geometry. Our reward explicitly incorporates cooperative ligand–ligand stacking energy alongside protein–ligand docking affinity, directly capturing the distinctive binding geometry of amyloid fibrils. We further introduce a curated evaluation set of amyloid–ligand complexes constructed from model-generated poses validated by domain experts. Experiments on both experimentally resolved structures and this evaluation set demonstrate improved pose quality and binding affinity correlation over existing docking baselines. Our project page is available at sunyasheng.github.io/projects/CORAL.

1 Introduction

> *"Of his bones are coral made; nothing of him that doth fade, but doth suffer a sea-change into something rich and strange."* — Shakespeare, *The Tempest* (1611)

Predicting how a small molecule binds to a target protein, known as the molecular docking problem, is a cornerstone of computational drug discovery, and recent deep learning approaches have achieved remarkable progress on globular protein targets Corso et al. [2023], Qiao et al. [2024], Abramson et al. [2024]. However, amyloid fibrils (the insoluble protein aggregates implicated in Alzheimer's, Parkinson's, and related neurodegenerative diseases) represent an increasingly important yet largely neglected class of docking targets. Developing small molecules that selectively bind amyloid fibrils is critical for both diagnostics (PET tracers, imaging probes) and therapeutics Chisholm and Hunter [2024, 2025], yet predicting how a given compound docks to a fibril remains an open challenge.

Modelling the binding modes of amyloid fibril–ligand complexes, however, faces two fundamental challenges. *First,* unlike globular proteins, amyloid fibrils cannot be crystallised, and cryo-EM, the only viable route to atomic-resolution fibril structures, often lacks sufficient resolution to unambiguously identify bound small-molecule ligands Zielinski et al. [2021]; co-crystal structures of amyloid–ligand complexes therefore remain exceptionally rare. Existing computational approaches Smith et al. [2024, 2025] rely on predefined potential energy functions to enumerate and score ligand poses; such physics-based methods are inherently limited by the brittleness of hand-crafted scoring functions and the rigid-body docking assumption. *Second,* unlike globular proteins, where a single ligand occupies a well-defined pocket, amyloid fibrils exhibit collective binding: multiple ligands intercalate cooperatively into cross-β grooves, with inter-ligand π–π stacking often dominating over direct protein contacts Merz et al. [2023]. The prevalence of fibril polymorphism, where the same protein can adopt multiple distinct fibril folds, further compounds this challenge, as the binding groove geometry varies across polymorphs. While recent work Guo et al. [2025] has applied learning-based methods to model amyloid fibril structures, how to model cooperative multi-ligand binding within these structures remains unexplored.

To address these challenges, we propose **CORAL** (**CO**ope**R**ative **A**myloid **L**igand docking), a cooperative multi-ligand docking framework that *mimics collaborative amyloid–ligand binding behaviours without demanding real-world co-crystal structures*. Our key insight is to design a mechanism that encourages the model to explore how ligands interact with the fibril groove, so that generated structures naturally satisfy the cooperative binding mode. Because co-crystal structures are scarce, we first construct a synthetic training set: diverse small-molecule conformers are generated with a drug generative model Schneuing et al. [2025], each conformer is replicated along the fibril axis to approximate cooperative stacking, and rule-based filters are applied to curate plausible binding poses. A structure generative model built on an AlphaFold 3-style Pairformer–Diffusion architecture Abramson et al. [2024] is supervised on this curated set, providing the model with a basic understanding of groove-intercalation geometry. To further align the generated poses with the energetics of cooperative binding, we refine the model via online reinforcement learning on the diffusion process Zheng et al. [2026], guided by a composite reward that jointly considers protein–ligand binding affinity and inter-ligand energetic complementarity along the periodic fibril axis. On the modelling side, amyloid fibrils are ribbon-like assemblies of repeated β-strand layers with a characteristic inter-chain Cα–Cα spacing of $\sim$4.8 Å. To faithfully capture this periodicity, we encode the inter-chain distance constraints directly into the Pairformer pair representation as an explicit structural prior for the fibril architecture.

Our main contributions can be summarised as follows: **1)** We present **CORAL**, to our knowledge the first deep learning-based docking framework for amyloid fibril–ligand complexes. **2)** We design a

two-stage training recipe that enables the model to learn fibril-specific binding modes: a warm-up stage on synthetically curated stacking poses, followed by online reinforcement-learning fine-tuning with a cooperative binding reward. **3)** We evaluate on two complementary datasets: a public set of fibril–ligand co-crystal structures and an expert-curated set whose reference poses are generated by computational agents and validated by domain experts. **CORAL** produces higher-quality poses than existing docking baselines on both.

2 Related Work

Amyloid Fibril Structures and Their Ligands. Small-molecule binding to tau, α-synuclein, and amyloid-β fibrils is central to diagnostics and therapeutics for neurodegenerative diseases Chisholm and Hunter [2024, 2025], Li and Liu [2023]. Cryo-EM structures reveal that aromatic ligands intercalate into longitudinal fibril grooves and stack cooperatively, with ligand–ligand π–π interactions often dominating over protein contacts Merz et al. [2023], Kunach et al. [2024], Smith et al. [2024], Chisholm and Hunter [2024, 2025], Liu et al. [2024]. SymDOCK Smith et al. [2024] extends classical docking to predict symmetric ligand stacks, reproducing experimentally determined tracer geometries, while a cooperative-binding model Smith et al. [2025] provides a thermodynamic framework for fibril–ligand affinities. RibbonFold Guo et al. [2025] adapts AlphaFold2 to generate polymorph-diverse amyloid structures but models only a single protofilament, and benchmarks show that general-purpose co-folding models still struggle with fibril targets Wojciechowska et al. [2024]. However, learning-based modelling of fibril–ligand binding modes remains an open problem.

Protein-Ligand Docking. Classical docking methods Jones et al. [1997], Friesner et al. [2004], Trott and Olson [2010], Eberhardt et al. [2021] score conformational samples with hand-crafted energy functions, yielding physically interpretable predictions Buttenschoen et al. [2024] but degrading in non-standard contexts Koes et al. [2013], McNutt et al. [2021]. Learning-based methods replace these hand-crafted components with data-driven models: regression approaches Stark et al. [2022], Lu et al. [2022], Zhang et al. [2023], Zhou et al. [2023], Pei et al. [2023] predict poses in a single forward pass, while diffusion-based docking Corso et al. [2023, 2024], Lu et al. [2024], Qiao et al. [2024], Morehead and Cheng [2025] and co-folding methods Jumper et al. [2021], Abramson et al. [2024], Krishna et al. [2024], Boitreaud et al. [2024], Wohlwend et al. [2024] further improve structural accuracy Morehead et al. [2025], building on a long line of neural protein structure prediction Pollastri et al. [2002], Hochreiter et al. [2007], Golkov et al. [2016]. All data-driven approaches share the same fundamental limitation, however: trained predominantly on globular proteins, they fail to generalise to non-standard binding geometries Morehead et al. [2025], Buttenschoen et al. [2024], leaving amyloid fibril targets, with their distinctive cross-β groove binding mode, largely out of reach.

Reinforcement Learning for Structure Generation. Conditional diffusion models have been applied to controllable image, motion, and video generation Sun et al. [2023, 2024, 2025, 2026]. For molecular structures, early work adopted offline preference optimisation: AliDiff Gu et al. [2024], DecompDPO Cheng et al. [2024], and AbDPO Zhou et al. [2024] extended DPO Rafailov et al. [2023] to molecule, structure-based drug, and antibody design respectively. The recent extension of GRPO Shao et al. [2024], DeepSeek-AI [2025], Schmidhuber [2015, 2018, 1992] to continuous diffusion and flow matching Xue et al. [2025], Liu et al. [2025] has opened the door to *online* RL fine-tuning of structural generative models, with concurrent work applying it to tasks such as protein inverse folding Wang et al. [2025b] and showing that online RL improves task-specific performance beyond what distribution matching alone can achieve.

3 Methodology

3.1 Overview

Problem Formulation. We formulate amyloid fibril–ligand docking as a co-folding problem Morehead et al. [2025]. An amyloid fibril is a periodic assembly of identical protein monomers (subunits) stacked along the fibril axis; we represent a fibril complex as K_p copies of the monomer amino acid sequence $\mathbf{A}$ and K_l copies of the same ligand molecular graph $\mathbf{G}$, where K_p and K_l may differ. Given $(\mathbf{A}^{1:K_p}, \mathbf{G}^{1:K_l})$, we jointly sample the 3D heavy-atom coordinates of all protein monomers

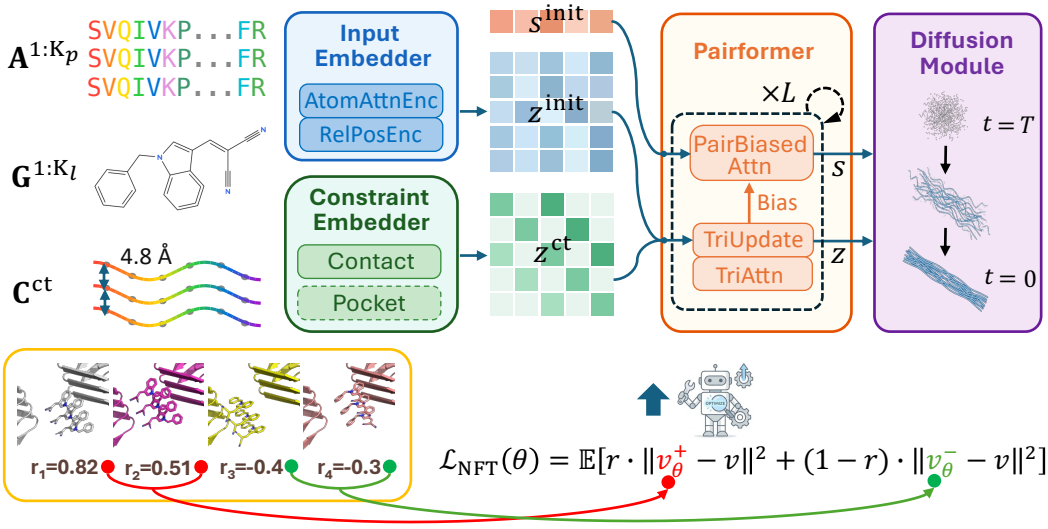

Figure 2: Overview of **CORAL**. A Pairformer–Diffusion structure model with fibril-periodicity priors generates 3D fibril–ligand complexes from protein sequences and ligand graphs. Generated poses are scored by a reward that combines protein–ligand binding affinity with cooperative inter-ligand stacking energy, and the model is refined via online reinforcement learning.

$\mathbf{x}_{\text{prot}} \in \mathbb{R}^{K_p N_p \times 3}$ and all ligands $\mathbf{x}_{\text{lig}} \in \mathbb{R}^{K_l N_l \times 3}$ from a generative model:

$$(\mathbf{x}_{\text{prot}},\, \mathbf{x}_{\text{lig}}) \sim p_\theta(\mathbf{x}_{\text{prot}}, \mathbf{x}_{\text{lig}} \mid \mathbf{A}^{1:K_p}, \mathbf{G}^{1:K_l}), \tag{1}$$

where N_p and N_l denote the number of heavy atoms per monomer and per ligand, respectively.

Preliminaries. Recent structure prediction models Abramson et al. [2024], Wohlwend et al. [2024] adopt flow matching as their generative backbone due to its stable training and efficient sampling. *Flow matching* Lipman et al. [2023], Liu et al. [2023] learns a velocity field $v_\theta(\mathbf{x}_t, t)$ that transports samples along a straight-line interpolant $\mathbf{x}_t = (1-t)\,\mathbf{x}_0 + t\,\mathbf{x}_1$ between data $\mathbf{x}_0$ and noise $\mathbf{x}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, trained via:

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t,\mathbf{x}_0,\mathbf{x}_1}\left[\|v_\theta(\mathbf{x}_t, t) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2\right]. \tag{2}$$

DiffusionNFT Zheng et al. [2026] aligns a flow-matching model with downstream objectives by operating on the forward-process velocity field, avoiding sampler coupling and trajectory storage. Given a scalar reward $r(\mathbf{x}_0)$, we obtain a weight $\tilde{r}(\mathbf{x}_0) \in [0, 1]$ via group-relative scaling. Two implicit velocity targets are constructed from the current policy v_θ and the sampling policy v_{old}:

$$v_\theta^+ = (1-\beta)\, v_{\text{old}} + \beta\, v_\theta, \tag{3}$$

$$v_\theta^- = (1+\beta)\, v_{\text{old}} - \beta\, v_\theta. \tag{4}$$

The training objective is:

$$\mathcal{L}_{\text{NFT}}(\theta) = \mathbb{E}\left[\tilde{r} \cdot \|v_\theta^+ - v\|^2 + (1-\tilde{r}) \cdot \|v_\theta^- - v\|^2\right], \tag{5}$$

where v is the target velocity derived from the forward process. High-reward samples are reinforced through v_θ^+, while low-reward samples are suppressed through v_θ^-.

3.2 Model Design

Our framework adopts a typical Pairformer–Diffusion architecture Abramson et al. [2024]. As shown in Figure 2, the input embedder produces initial single ($\mathbf{s}^{\text{init}}$) and pair ($\mathbf{z}^{\text{init}}$) representations, the Pairformer trunk refines them through triangle updates and pair-biased attention, and the diffusion module decodes atomic coordinates conditioned on the refined representations. We augment this pipeline with a *ConstraintEmbedder* that injects fibril-specific geometric priors into $\mathbf{z}^{\text{init}}$.

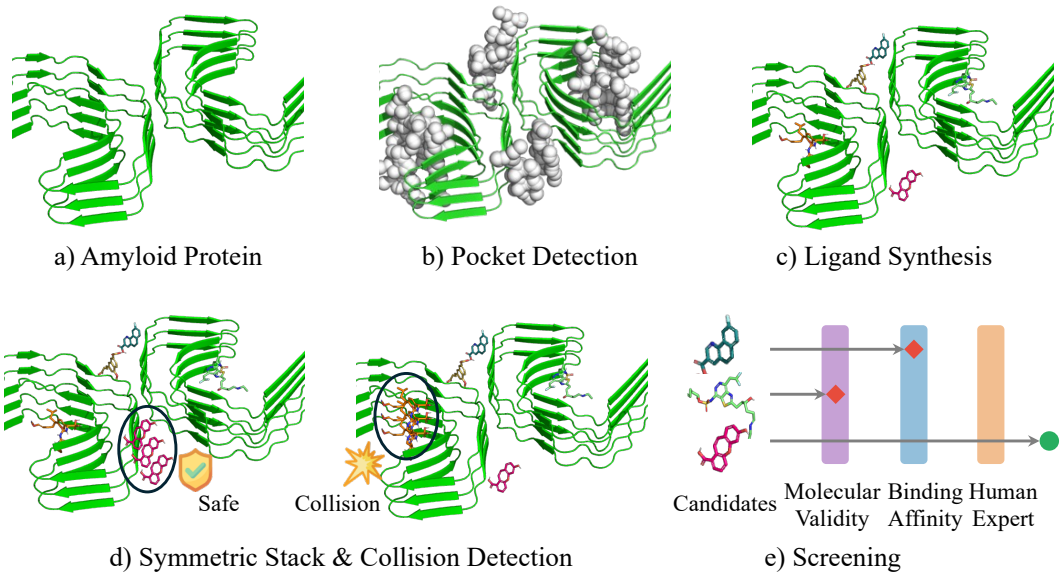

Figure 3: Training data curation pipeline. Starting from an amyloid fibril structure (a), we detect candidate binding sites on the fibril surface (b) and generate diverse small-molecule conformers at each site (c). Each conformer is replicated along the fibril axis to form a symmetric stacking arrangement, retaining only collision-free configurations (d). Surviving candidates pass sequential LLM, RDKit, binding-affinity, and expert checks (e).

Constraint-Aware Input Embedding. A target fibril complex comprises K_p peptide chains that share an identical amino acid sequence and K_l ligand copies that share an identical molecular graph, which we denote by $\mathbf{A}^{1:K_p}$ and $\mathbf{G}^{1:K_l}$, respectively. These inputs are tokenised and embedded into initial per-token single representations $\mathbf{s}_i^{\text{init}} \in \mathbb{R}^{c_s}$ and per-token-pair pair representations $\mathbf{z}_{ij}^{\text{init}} \in \mathbb{R}^{c_z}$, where each token i corresponds to either a residue (for protein chains) or a heavy atom (for ligands). Here $\mathbf{s}_i$ encodes the identity of token i, while $\mathbf{z}_{ij}$ encodes the pairwise relationship between tokens i and j, including spatial proximity, bonding, and relative position.

To inject fibril-specific geometry into the pair representation, we introduce a *ConstraintEmbedder* that supports two types of constraints. The *contact constraint* $\mathbf{C}^{\text{ct}}$ captures the translational periodicity of the fibril, i.e., the rigid-body shift of $\sim$4.8 Å between neighbouring chains along the fibril axis. It is realised as the C_α–C_α distance $\mathbf{C}_{ij}^{\text{ct}} = \|\mathbf{r}_i - \mathbf{r}_j\|_2$ between residue tokens i and j, where $\mathbf{r}_i, \mathbf{r}_j \in \mathbb{R}^3$ are reference coordinates taken from the target fibril structure. The distance is then added to the pair representation through a linear projection:

$$\mathbf{z}_{ij}^{\text{init}} \leftarrow \mathbf{z}_{ij}^{\text{init}} + \mathbf{W}_{\text{ct}}\, \mathbf{C}_{ij}^{\text{ct}}. \tag{6}$$

The *pocket constraint* $\mathbf{C}^{\text{pkt}}$ marks the target binding region on the fibril surface, giving the ligand a spatial prior over where to bind; it is injected in the same manner via a separate projection $\mathbf{W}_{\text{pkt}}$. All projections in the ConstraintEmbedder are zero-initialised, so that the added branch does not perturb the pair representation at initialization.

Pairformer Trunk and Diffusion Module. The Pairformer trunk serves as the primary representation-learning engine of the model: it progressively refines $(\mathbf{s}^{\text{init}}, \mathbf{z}^{\text{init}})$ through L blocks that enforce triangle-inequality consistency on the pair track and propagate this geometry into the single track via pair-biased attention Abramson et al. [2024]. Because the constraints enter at $\mathbf{z}^{\text{init}}$, they are carried through every Pairformer block and are ultimately encoded into the refined representations $(\mathbf{s}^{(L)}, \mathbf{z}^{(L)})$. The diffusion module then acts as a generative decoder: conditioned on $(\mathbf{s}^{(L)}, \mathbf{z}^{(L)})$, it parameterises the velocity field $v_\theta(\mathbf{x}_t, t; \mathbf{s}^{(L)}, \mathbf{z}^{(L)})$ introduced in Eq. (2), and at inference integrates it from t=1 to t=0 to denoise random coordinates into a plausible fibril–ligand complex $(\mathbf{x}_{\text{prot}}, \mathbf{x}_{\text{lig}})$.

3.3 Training

Data Curation. To address the scarcity of fibril–ligand co-crystal structures, we curate synthetic stacking poses for supervised warm-up (Figure 3). We retain two-protofilament fibrils with at least three consecutive chains per protofilament to capture both intra- and inter-protofilament binding interfaces. For each fibril, fpocket Le Guilloux et al. [2009] identifies candidate pockets along the cross-β grooves, which are inspected before DrugFlow Schneuing et al. [2025] generates diverse ligand conformers at each site. Each conformer is replicated along the fibril axis using the inter-chain translation to mimic the symmetric stacking observed in cryo-EM structures; arrangements with steric clashes or chemically unreasonable conformations are discarded. Surviving poses pass through sequential filters: coarse LLM-based pose-plausibility screening, RDKit checks of chemical validity, docking-based affinity screening, and final expert structural review. The retained poses form the *AmyDock-Synthetic* warm-up training set.

Warm-up Stage. This stage warms up the model to the symmetric binding mode at a geometric level: how ligands intercalate into cross-β grooves, align along the fibril axis, and form periodically repeating stacking arrangements. We supervise the model on the final curated *AmyDock-Synthetic* training set, using a typical structure-prediction training objective Abramson et al. [2024] that combines three groups of losses: a flow-matching loss $\mathcal{L}_{\text{FM}}$ that trains the diffusion module to denoise atomic coordinates towards the ground truth; a distogram loss $\mathcal{L}_{\text{dist}}$ that supervises the trunk's predicted distribution over pairwise residue distances; and three auxiliary confidence losses ($\mathcal{L}_{\text{pLDDT}}$, $\mathcal{L}_{\text{PDE}}$, $\mathcal{L}_{\text{PAE}}$) that supervise per-atom and per-pair error predictions. Following the classifier-free guidance strategy Ho and Salimans [2021], the pocket constraint $\mathbf{C}^{\text{pkt}}$ is dropped with 50% probability during training. This way, the network learns to generate both with and without the pocket prior, and we can combine the two at inference to modulate the strength of the prior.

Reinforcement Learning with Cooperative Binding Reward. The warm-up stage captures the overall shape of symmetric stacking, but generated poses still deviate from realistic binding energetics. To close this gap, we fine-tune the warmed-up model v_{ref} with reinforcement learning Zheng et al. [2026] under a composite reward:

$$r(\mathbf{x}) \;=\; -R_{\text{bind}}(\mathbf{x}) + \lambda\, R_{\text{sym}}(\mathbf{x}), \tag{7}$$

with two complementary terms. The *binding score* combines the AutoDock Vina protein–ligand docking score ΔG_{PL} Trott and Olson [2010], Eberhardt et al. [2021] with the ligand–ligand stacking energy E_{stack}, computed between the centre ligand and its symmetry-generated neighbours using a Mie (8–4) potential with short-distance regularisation:

$$R_{\text{bind}} \;=\; \Delta G_{\text{PL}} + R_{\text{stack}} \;\equiv\; \Delta G_{\text{eff}}, \qquad R_{\text{stack}} = f \cdot E_{\text{stack}}, \tag{8}$$

where $f > 0$ weights the stacking contribution. Lower values of this combined score yield higher rewards through the minus sign in Eq. (7). The *symmetry consistency* $R_{\text{sym}} = -\,\text{RMSD}(\mathbf{T}\,\mathbf{x}_k,\, \mathbf{x}_{k+1})$ penalises deviations from the inter-chain rigid-body transformation $\mathbf{T}$ (a $\sim$4.8 Å translation along the fibril axis) that should map ligand $\mathbf{x}_k$ at chain k onto $\mathbf{x}_{k+1}$ at chain $k{+}1$. We minimise the RL objective $\mathcal{L}_{\text{NFT}}(\theta)$ in Eq. (5) with this reward, adding a KL regulariser $\beta_{\text{KL}}\|v_\theta - v_{\text{ref}}\|^2$ towards v_{ref} to stabilise training.

4 Experiments

4.1 Experimental Setup

Implementation. We implement the Pairformer–Diffusion backbone using Protenix AI for Science Team, ByteDance Research [2025], an open-source PyTorch reproduction of AlphaFold 3 Abramson et al. [2024]. Training follows the two-stage recipe described in §3.3.[3] In the warm-up stage, we supervise the model on the final curated *AmyDock-Synthetic* fibril–ligand stacking poses. In the RL stage, the amyloid fibril–ligand complexes used as roll-out inputs are drawn from the amyloid ligand database of Chisholm and Hunter [2025] (covering α-synuclein, amyloid-β, and tau fibril targets), filtered to entries with a two-protofilament architecture and excluding evaluation ligands.

[3] See Appendix A.2 for data curation, Appendix B.2 for expanded reward equations, and Appendix A.3 for implementation details.

Table 1: Docking accuracy on AmyDock-CryoEM and AmyDock-Synthetic, and binding-affinity rank correlation on αSyn-Bind (from Chu et al. [2015]). RMSD values are in Å; Hit@5 Å denotes the fraction of pairs with ligand RMSD ≤ 5 Å; Spearman ρ is computed between predicted binding scores and experimental $\log(K_d/\mathrm{M})$. "–" denotes not available (NA).

	AmyDock-CryoEM			**AmyDock-Synthetic**			**αSyn-Bind**	
Method	Mean ↓	Median ↓	Hit@5Å ↑	Mean ↓	Median ↓	Hit@5Å ↑	ρ ↑	p ↓
PatchDock	23.6	24.8	0.0%	26.9	25.7	0.0%	–	–
DiffDock	17.67	13.68	0.0%	19.19	19.68	3.5%	0.191	0.197
DynamicBind	26.63	29.02	0.0%	18.90	17.93	1.8%	0.102	0.544
FABind	26.50	24.90	0.0%	30.34	26.87	0.0%	0.157	0.292
NeuralPLexer	17.43	19.00	5.6%	19.78	16.56	0.0%	−0.291	0.047
FlowDock	19.95	21.52	0.0%	20.69	19.88	0.0%	0.148	0.322
Boltz-2	28.27	28.67	0.0%	37.13	33.87	0.0%	0.223	0.132
Chai-1	19.85	19.92	0.0%	19.40	16.52	0.0%	−0.079	0.597
CORAL (ours)	**8.70**	**4.89**	**50.0%**	**7.35**	**3.84**	**68.4%**	**0.419**	**0.004**

Datasets. We evaluate primarily on two complementary docking datasets. *AmyDock-CryoEM* draws ground-truth poses directly from PDB-deposited cryo-EM models; after filtering out solvent, buffer, and other non-specific small molecules, the set comprises 18 fibril–ligand complexes spanning tau, α-synuclein, and related fibril-forming proteins. *AmyDock-Synthetic* generates reference poses through an agent-based sampling pipeline; after plausibility screening, binding-affinity scoring, and expert validation, the held-out evaluation set comprises 57 fibril–ligand pairs that cover a broader range of ligand chemotypes than the cryo-EM structures alone. Additionally, to probe whether the predicted poses carry signal of the underlying binding energetics, we use a set of 47 small molecules with measured dissociation constants (K_d) against α-synuclein fibrils published by Chu et al. [2015]. This set has no ground-truth poses and is used solely for binding-affinity rank correlation; we refer to it as *αSyn-Bind*.

Baselines. We compare against eight docking methods spanning geometry-based docking, regression, diffusion, flow matching, and co-folding paradigms: PatchDock Schneidman-Duhovny et al. [2005], DiffDock Corso et al. [2023], DynamicBind Lu et al. [2024], FABind Pei et al. [2023], NeuralPLexer Qiao et al. [2024], FlowDock Morehead and Cheng [2025], Boltz-2 Passaro et al. [2025], and Chai-1 Boitreaud et al. [2024]. For a fair comparison, each learning-based baseline is also fine-tuned on the *AmyDock-Synthetic* training set, as is our approach.

4.2 Evaluation

Quantitative Results. On *AmyDock-CryoEM* and *AmyDock-Synthetic*, we measure docking accuracy with two metrics: **Mean / Median RMSD**, the heavy-atom RMSD between the best predicted pose and the ground truth (averaged or median across all pairs); and **Hit rate (≤ 5 Å)**, the fraction of pairs with RMSD ≤ 5 Å. On *αSyn-Bind*, we measure binding-affinity correlation with **Spearman** ρ and its p-value between predicted binding scores and experimental $\log(K_d/\mathrm{M})$. For baselines that only support single-ligand docking, we compare the predicted pose against the nearest ground-truth ligand copy.

As shown in Table 1, **CORAL** reaches a hit rate of $\mathbf{68.4}\%$ at ≤ 5 Å on *AmyDock-Synthetic* (median RMSD $\mathbf{3.84}$ Å) and $\mathbf{50.0}\%$ on *AmyDock-CryoEM*, outperforming baselines whose single-ligand pocket formulations do not explicitly model the cooperative, periodic binding geometry of fibrils. On *αSyn-Bind*, **CORAL**'s predicted binding scores show a rank correlation with experimental K_d values (Spearman ρ=$\mathbf{0.419}$, p=$\mathbf{0.004}$), suggesting potential for ranking ligands by binding affinity.

Qualitative Results. Figure 4 visualizes predictions on two representative *AmyDock-CryoEM* targets. Regression-based dockers such as FABind Pei et al. [2023] collapse to a single averaged pose that drifts from the binding region, while the generative baselines dock each ligand independently and, lacking awareness of the fibril's periodic geometry, tend to plaster the ligand vertically against the fibril surface or miss the intra-chain groove entirely—both failure modes yield large RMSD even

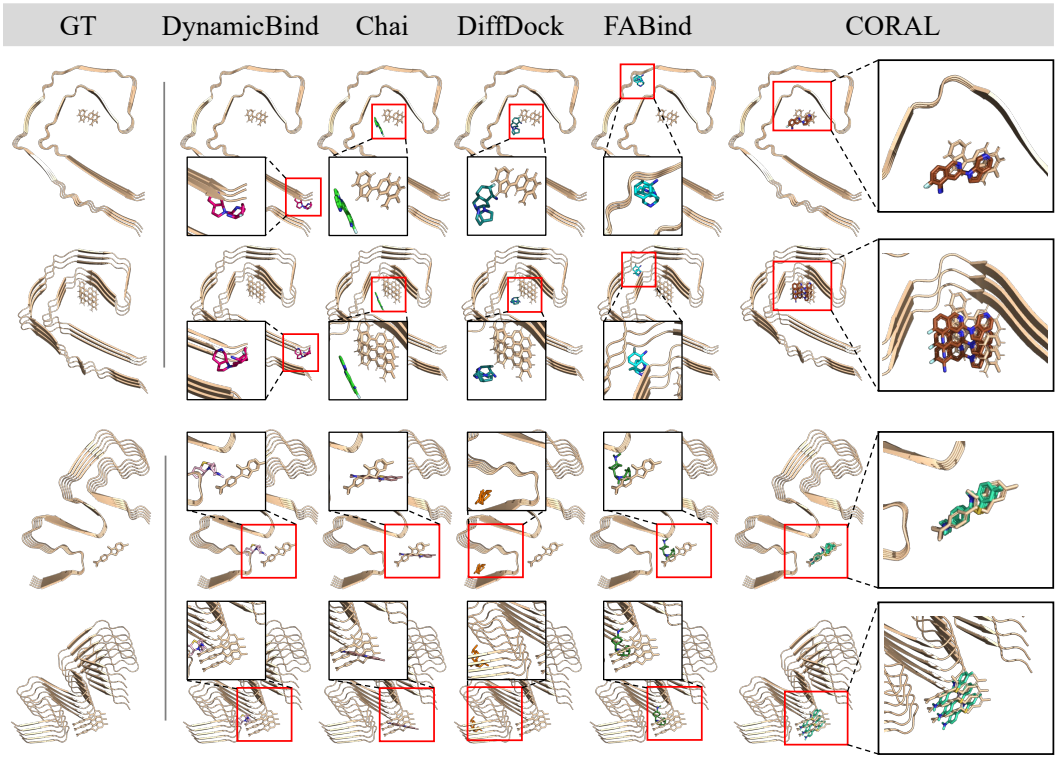

Figure 4: Qualitative comparison of predicted fibril–ligand poses on two representative *AmyDock-CryoEM* targets. Each target spans two rows—an axial view along the fibril axis (top) and an oblique top-down view (bottom)—while the columns compare the ground truth against each docking method. The ground-truth ligand is overlaid on every prediction panel; each red inset marks the ligand position that is enlarged in the adjacent black box.

when the per-ligand chemistry looks plausible. **CORAL** instead reproduces the periodic stacking of ligand copies across neighbouring peptide chains. The visual comparison corroborates the quantitative gap in Table 1 and indicates that learning the inter-ligand cooperative structure (not just single-pose protein–ligand fit) is the key driver of **CORAL**'s improvement.

4.3 Further Analysis

Ablation Study. We ablate two key design choices of **CORAL**—the RL reward and the two-stage training recipe—on *AmyDock-Synthetic* (Table 2). Each variant modifies a single factor while keeping all other components identical to the full model. Removing R_{stack} collapses R_{bind} to plain protein–ligand affinity, eliminating the cooperative inter-ligand signal and degrading ΔG_{eff}; removing R_{sym} has a milder effect on RMSD because the warm-up stage has already established much of the stacking geometry, so R_{sym} mainly serves as a fine-tuning anchor during RL; removing the KL term destabilises training and the model diverges from the warmed-up shape prior. Finally, RL without warm-up converges to a worse local optimum—without a shape prior to start from, the policy explores from scratch and lands on noisier stacking geometry—while warm-up without RL produces plausible stacking but under-performs energetically; only the combined recipe achieves both accurate geometry and favourable binding energetics.

Pose Refinement during RL. During RL, the mean reward stabilises and R_{bind} generally becomes more favourable.[4] Figure 5 illustrates these training-stage changes at the pose level for four representative training complexes, comparing their first and last visits. For each complex, the pose at its first

[4] Appendix B.3 reports training dynamics and the effects of stacking weight on surface burial and interface complementarity.

Table 2: Ablation study on AmyDock-Synthetic. Each variant modifies one factor of the full model; "$-X$" removes component X. ΔG_{eff} (kcal/mol) is the combined binding score defined in Eq. (8); we report both mean and median.

Variant	**Mean RMSD ↓**	**Median RMSD ↓**	**Mean ΔG_{eff} ↓**	**Median ΔG_{eff} ↓**
CORAL (full)	**7.35**	**3.84**	**-2.24**	**-2.44**
RL reward components				
$-$ R_{sym} (no symmetry term)	7.92	4.31	-2.07	-2.21
$-$ R_{stack} (no cooperative term)	8.21	4.46	-0.78	-0.95
$-$ KL (no stability regulariser)	12.07	8.13	-1.32	-1.45
Training recipe				
Warm-up only (no RL)	8.67	5.18	-1.41	-1.58
RL only (no warm-up)	13.45	9.27	-0.72	-0.88

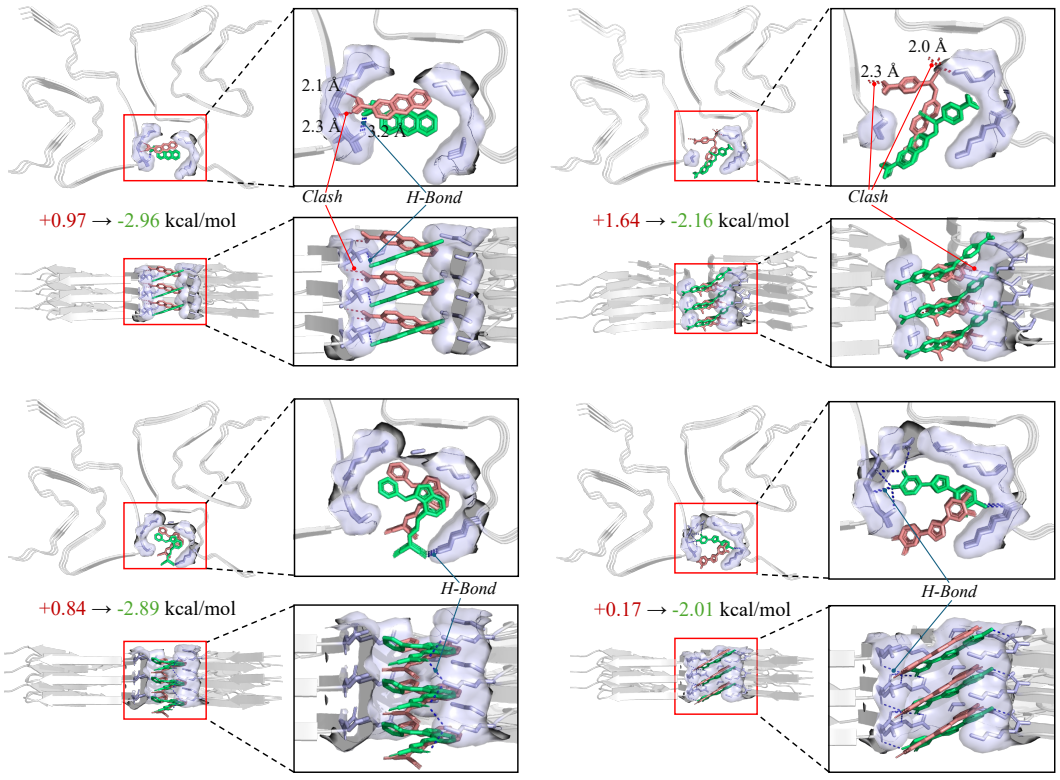

Figure 5: **Pose refinement during reinforcement learning.** Four representative training complexes. For each complex, the pose at its first visit during RL training (brown) is shown alongside the pose at its last visit (green), with zoomed insets in both top-down (along the fibril axis) and side views; red labels and arrows mark steric clashes (with closest distances in Å), and blue dashed lines mark hydrogen bonds. Numbers report the corresponding shift in R_{bind} from first to last visit, moving from positive (unfavourable) to clearly negative (favourable) values, with ∼2–4 kcal/mol of improvement per complex.

visit during training (brown) fits awkwardly into the cross-β groove and exhibits steric clashes with the protein surface. By the last visit (green), the ligand has adapted to the local pocket geometry, and the resulting shape complementarity yields a richer network of favourable protein–ligand interactions: clashes are eliminated, well-defined hydrogen bonds form, and van der Waals packing improves. R_{bind} correspondingly shifts from positive (unfavourable) to clearly negative values.

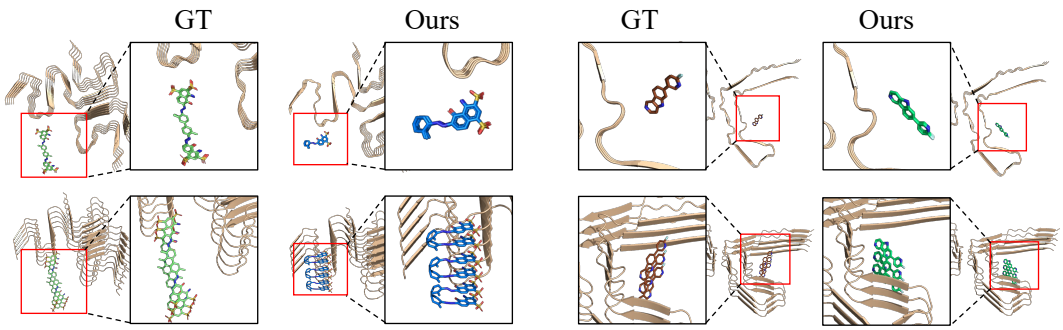

Figure 6: Representative failure cases on *AmyDock-CryoEM*. Two cases (columns 1–2 and 3–4) are shown, each pair comparing the ground-truth pose (GT) against **CORAL**'s prediction (Ours) in two views (top row: axial along the fibril axis; bottom row: oblique). Red insets are enlarged in the adjacent panels. The left case shows a prediction with distorted ligand geometry, while the right case shows an orientation mismatch relative to the ground-truth groove site.

Case Study. We highlight a few cases where **CORAL** underperforms (Figure 6). *First*, the model folds the ligand back on itself so that one segment stacks with neighbouring copies while the other bends toward the protein surface (Figure 6, left), letting a single pose gain both stacking and protein-contact energy at once. *Second*, the model places the ligand at the correct binding site but rotates it sideways from the canonical axial orientation to enlarge protein contact (Figure 6, right).

5 Conclusion

We present **CORAL**, the first deep learning-based docking framework tailored to amyloid fibril–ligand complexes. Our framework tackles cooperative multi-ligand docking on amyloid fibrils—a binding mode that single-pocket docking models are not designed to capture: it injects fibril-periodicity priors into the structure model and refines pose distributions through online reinforcement learning under a stacking-aware reward, sidestepping the scarcity of fibril–ligand co-crystal supervision. Across the two docking benchmarks, our framework substantially outperforms existing docking methods. Results on α*Syn-Bind* also suggest potential for ranking ligands by binding affinity.

Future Work. We plan to explore more accurate numerical simulations, such as molecular dynamics, to develop reward signals that better reflect realistic fibril–ligand binding modes. We also aim to extend our evaluation to other amyloid fibrils, including those formed by transthyretin and immunoglobulin light chains, and to integrate *de novo* ligand generation, our docking framework, and downstream filtering into a single agent-orchestrated closed-loop pipeline Lu et al. [2026], Boiko et al. [2023] for automated discovery of fibril-binding candidates with minimal human intervention.

Limitations. Our evaluation covers periodic ligand-stacking geometries in tau, α-synuclein, and amyloid-β fibrils. The synthetic data pipeline may introduce biases during generation and screening, with its LLM pre-screen relying on prompt-based judgments. Equipping this screening agent with established structural and chemical analysis tools could improve reliability. The training reward approximates intermolecular interactions through a docking score and a pairwise Mie (8–4) stacking potential. The geometric distortions and orientation mismatches in Figure 6 motivate a more complete treatment of conformational and intermolecular energetics. Our affinity results suggest potential for ligand ranking; broader experimental binding datasets and more accurate numerical models of fibril–ligand interactions could support a more comprehensive assessment and refinement of affinity prediction.

Broader Impact. The framework aims to accelerate early-stage drug discovery for Alzheimer's, Parkinson's, and related neurodegenerative diseases. Predictions of cooperative fibril–ligand binding provide structural hypotheses that can help prioritise candidate imaging probes and therapeutic ligands for experimental evaluation.

Acknowledgments and Disclosure of Funding

The research reported in this publication was supported by a gift from Google, administered through the King Abdullah University of Science and Technology (KAUST) – Center of Excellence for Generative AI under award number 5940. We also thank KAUST for providing GPU computing resources on the Ibex cluster. The authors declare no competing interests.

References

J. Abramson, J. Adler, J. Dunger, R. Evans, T. Green, A. Pritzel, O. Ronneberger, L. Willmore, A. J. Ballard, J. Bambrick, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016):493–500, 2024.

AI for Science Team, ByteDance Research. Protenix: advancing structure prediction through a comprehensive AlphaFold3 reproduction. *bioRxiv*, 2025. doi: 10.1101/2025.01.08.631967.

G. W. Bemis and M. A. Murcko. The properties of known drugs. 1. molecular frameworks. *Journal of Medicinal Chemistry*, 39(15):2887–2893, 1996. doi: 10.1021/jm9602928.

D. A. Boiko, R. MacKnight, B. Kline, and G. Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023. doi: 10.1038/s41586-023-06792-0.

J. Boitreaud, J. Dent, M. McPartlon, J. Meier, V. R. Reis, A. Rogozhnikov, and K. Wu. Chai-1: decoding the molecular interactions of life. *bioRxiv*, 2024.

M. Buttenschoen, G. M. Morris, and C. M. Deane. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical Science*, 15(9): 3130–3139, 2024.

X. Cheng, X. Zhou, Y. Yang, Y. Bao, and Q. Gu. Decomposed direct preference optimization for structure-based drug design. *Transactions on Machine Learning Research*, 2024.

T. S. Chisholm and C. A. Hunter. A closer look at amyloid ligands, and what they tell us about protein aggregates. *Chemical Society Reviews*, 53(3):1354–1374, 2024. doi: 10.1039/D3CS00518F.

T. S. Chisholm and C. A. Hunter. Ligands for protein fibrils of amyloid-β, α-synuclein, and tau. *Chemical Reviews*, 125(11):5282–5348, 2025. doi: 10.1021/acs.chemrev.4c00838.

W. Chu, D. Zhou, V. Gaba, J. Liu, S. Li, X. Peng, J. Xu, D. Dhavale, D. P. Bagchi, A. d'Avignon, et al. Design, synthesis, and characterization of 3-(benzylidene)indolin-2-one derivatives as ligands for α-synuclein fibrils. *Journal of Medicinal Chemistry*, 58(15):6002–6017, 2015. doi: 10.1021/acs.jmedchem.5b00571.

G. Corso, H. Stärk, B. Jing, R. Barzilay, and T. Jaakkola. DiffDock: diffusion steps, twists, and turns for molecular docking. In *International Conference on Learning Representations*, 2023.

G. Corso, A. Deng, B. Fry, N. Polizzi, R. Barzilay, and T. Jaakkola. Deep confident steps to new pockets: strategies for docking generalization. In *International Conference on Learning Representations*, 2024.

J. Cremer, R. Irwin, A. Tibo, J. P. Janet, S. Olsson, and D.-A. Clevert. FLOWR: Flow matching for structure-aware de novo, interaction- and fragment-based ligand generation. *arXiv preprint arXiv:2504.10564*, 2025.

DeepSeek-AI. DeepSeek-R1: incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

J. Durairaj, Y. Adeshina, Z. Cao, X. Zhang, V. Oleinikovas, T. Duignan, Z. McClure, X. Robin, G. Studer, D. Kovtun, E. Rossi, G. Zhou, S. P. Veccham, C. Isert, Y. Peng, P. Sundareson, M. Akdel, G. Corso, H. Stärk, G. Tauriello, Z. W. Carpenter, M. M. Bronstein, E. Kucukbenli, T. Schwede, and L. Naef. PLINDER: The protein-ligand interactions dataset and evaluation resource. *bioRxiv*, 2024. doi: 10.1101/2024.07.17.603955.

J. Eberhardt, D. Santos-Martins, A. F. Tillack, and S. Forli. AutoDock Vina 1.2.0: new docking methods, expanded force field, and Python bindings. *Journal of Chemical Information and Modeling*, 61(8):3891–3898, 2021.

J. Filipovič, O. Vávra, J. Plhák, D. Bednář, S. M. Marques, J. Brezovský, L. Matyska, and J. Damborský. CaverDock: A novel method for the fast analysis of ligand transport. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 17(5):1625–1638, 2020. doi: 10.1109/TCBB.2019.2907492.

R. A. Friesner, J. L. Banks, R. B. Murphy, T. A. Halgren, J. J. Klicic, D. T. Mainz, M. P. Repasky, E. H. Knoll, M. Shelley, J. K. Perry, et al. 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.

V. Golkov, M. J. Skwark, A. Golkov, A. Dosovitskiy, T. Brox, J. Meiler, and D. Cremers. Protein contact prediction from amino acid co-evolution using convolutional networks for graph-valued images. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, 2016.

S. Gu, M. Xu, A. Powers, W. Nie, T. Geffner, K. Kreis, J. Leskovec, A. Vahdat, and S. Ermon. AliDiff: aligning target-aware molecule diffusion models with exact energy optimization. In *Advances in Neural Information Processing Systems*, 2024.

L. Guo, Q. Yu, D. Wang, X. Wu, P. G. Wolynes, and M. Chen. Generating the polymorph landscapes of amyloid fibrils using AI: RibbonFold. *Proceedings of the National Academy of Sciences*, 2025. doi: 10.1073/pnas.2501321122.

J. Ho and T. Salimans. Classifier-free diffusion guidance. In *NeurIPS Workshop on Deep Generative Models and Downstream Applications*, 2021.

S. Hochreiter, M. Heusel, and K. Obermayer. Fast model-based protein homology detection without alignment. *Bioinformatics*, 23(14):1728–1736, 2007.

L. Hu, M. L. Benson, R. D. Smith, M. G. Lerner, and H. A. Carlson. Binding MOAD (mother of all databases). *Proteins: Structure, Function, and Bioinformatics*, 60(3):333–340, 2005. doi: 10.1002/prot.20512.

G. Jones, P. Willett, R. C. Glen, A. R. Leach, and R. Taylor. Development and validation of a genetic algorithm for flexible docking. *Journal of Molecular Biology*, 267(3):727–748, 1997.

J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021.

W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, 1976. doi: 10.1107/S0567739476001873.

D. R. Koes, M. P. Baumgartner, and C. J. Camacho. Lessons learned in empirical scoring with SMINA from the CSAR 2011 benchmarking exercise. *Journal of Chemical Information and Modeling*, 53(8):1893–1904, 2013.

R. Krishna, J. Wang, W. Ahern, P. Sturmfels, P. Venkatesh, I. Kalvet, G. R. Lee, F. S. Morey-Burrows, I. Anishchenko, I. R. Humphreys, R. McHugh, D. Vafeados, X. Li, G. A. Sutherland, A. Hitchcock, C. N. Hunter, A. Kang, E. Brackenbrough, A. K. Bera, M. Baek, F. DiMaio, and D. Baker. Generalized biomolecular modeling and design with RoseTTAFold All-Atom. *Science*, 384(6693): eadl2528, 2024. doi: 10.1126/science.adl2528.

P. Kunach, J. Vaquer-Alicea, M. S. Smith, J. Monistrol, R. Hopewell, B. K. Shoichet, P. Rosa-Neto, M. I. Diamond, S. H. Shahmoradian, et al. Cryo-EM structure of Alzheimer's disease tau filaments with PET ligand MK-6240. *Nature Communications*, 15(1):8497, 2024. doi: 10.1038/s41467-024-52265-x.

M. C. Lawrence and P. M. Colman. Shape complementarity at protein/protein interfaces. *Journal of Molecular Biology*, 234(4):946–950, 1993.

V. Le Guilloux, P. Schmidtke, and P. Tuffery. fpocket: an open source platform for ligand pocket detection. *BMC Bioinformatics*, 10:168, 2009.

D. Li and C. Liu. Molecular rules governing the structural polymorphism of amyloid fibrils in neurodegenerative diseases. *Structure*, 31(11):1335–1347, 2023. doi: 10.1016/j.str.2023.08.006.

Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, and M. Nickel. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

J. Liu, G. Liu, J. Liang, Y. Li, J. Liu, X. Wang, P. Wan, D. Zhang, and W. Ouyang. Flow-GRPO: training flow matching models via online RL. In *Advances in Neural Information Processing Systems*, 2025.

K. Liu, Y. Tao, Q. Zhao, W. Xia, X. Li, S. Zhang, Y. Yao, H. Xiang, C. Han, L. Tan, B. Sun, D. Li, A. Li, and C. Liu. Binding adaptability of chemical ligands to polymorphic α-synuclein amyloid fibrils. *Proceedings of the National Academy of Sciences*, 121(35):e2321633121, 2024. doi: 10.1073/pnas.2321633121.

X. Liu, C. Gong, and Q. Liu. Flow straight and fast: learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations*, 2023.

C. Lu, C. Lu, R. T. Lange, Y. Yamada, S. Hu, J. Foerster, D. Ha, and J. Clune. Towards end-to-end automation of AI research. *Nature*, 651(8107):914–919, 2026. doi: 10.1038/s41586-026-10265-5.

W. Lu, Q. Wu, J. Zhang, J. Rao, C. Li, and S. Zheng. TANKBind: trigonometry-aware neural networks for drug-protein binding structure prediction. In *Advances in Neural Information Processing Systems*, 2022.

W. Lu, J. Zhang, W. Huang, Z. Zhang, X. Jia, Z. Wang, L. Shi, C. Li, P. G. Wolynes, and S. Zheng. DynamicBind: predicting ligand-specific protein-ligand complex structure with a deep equivariant generative model. *Nature Communications*, 15(1):1071, 2024.

A. T. McNutt, P. Francoeur, R. Aggarwal, T. Masuda, R. Meli, M. Ragoza, J. Sunseri, and D. R. Koes. GNINA 1.0: molecular docking with deep learning. *Journal of Cheminformatics*, 13(1):43, 2021.

G. E. Merz, M. J. Chalkley, S. K. Tan, E. Tse, J. Lee, S. B. Prusiner, N. A. Paras, W. F. DeGrado, D. R. Southworth, et al. Stacked binding of a PET ligand to Alzheimer's tau paired helical filaments. *Nature Communications*, 14(1):3048, 2023. doi: 10.1038/s41467-023-38537-y.

A. Morehead and J. Cheng. FlowDock: geometric flow matching for generative protein–ligand docking and affinity prediction. *Bioinformatics*, 41(Supplement_1):i198–i206, 2025.

A. Morehead, N. Giri, J. Liu, P. Neupane, and J. Cheng. Assessing the potential of deep learning for protein–ligand docking. *Nature Machine Intelligence*, 2025.

S. Passaro, G. Corso, J. Wohlwend, M. Reveiz, S. Thaler, V. R. Somnath, N. Getz, T. Portnoi, J. Roy, H. Stark, D. Kwabi-Addo, D. Beaini, T. Jaakkola, and R. Barzilay. Boltz-2: towards accurate and efficient binding affinity prediction. *bioRxiv*, 2025. doi: 10.1101/2025.06.14.659707.

Q. Pei, K. Gao, L. Wu, J. Zhu, Y. Xia, S. Xie, T. Qin, K. He, T.-Y. Liu, and R. Yan. FABind: fast and accurate protein-ligand binding. In *Advances in Neural Information Processing Systems*, 2023.

G. Pollastri, D. Przybylski, B. Rost, and P. Baldi. Improving the prediction of protein secondary structure in three and eight classes using recurrent neural networks and profiles. *Proteins: Structure, Function, and Bioinformatics*, 47(2):228–235, 2002.

Z. Qiao, W. Nie, A. Vahdat, T. F. Miller III, and A. Anandkumar. State-specific protein–ligand complex structure prediction with a multi-scale deep generative model. *Nature Machine Intelligence*, 6: 195–208, 2024.

R. Rafailov, A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, and C. Finn. Direct preference optimization: your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, 2023.

J. Schmidhuber. Learning complex, extended sequences using the principle of history compression. *Neural Computation*, 4(2):234–242, 1992.

J. Schmidhuber. On learning to think: Algorithmic information theory for novel combinations of reinforcement learning controllers and recurrent neural world models. *arXiv preprint arXiv:1511.09249*, 2015.

J. Schmidhuber. One big net for everything. *arXiv preprint arXiv:1802.08864*, 2018.

D. Schneidman-Duhovny, Y. Inbar, R. Nussinov, and H. J. Wolfson. PatchDock and SymmDock: Servers for rigid and symmetric docking. *Nucleic Acids Research*, 33(Web Server):W363–W367, 2005. doi: 10.1093/nar/gki481.

A. Schneuing, I. Igashov, A. W. Dobbelstein, T. Castiglione, M. M. Bronstein, and B. Correia. Multi-domain distribution learning for de novo drug design. *arXiv preprint arXiv:2508.17815*, 2025.

Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. Li, Y. Wu, et al. DeepSeekMath: pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

A. Shrake and J. A. Rupley. Environment and exposure to solvent of protein atoms. lysozyme and insulin. *Journal of Molecular Biology*, 79(2):351–371, 1973.

M. S. Smith, I. S. Knight, R. C. Kormos, J. G. Pepe, P. Kunach, M. I. Diamond, S. H. Shahmoradian, J. J. Irwin, W. F. DeGrado, and B. K. Shoichet. Docking for molecules that bind in a symmetric stack with SymDOCK. *Journal of Chemical Information and Modeling*, 64(2):425–434, 2024. doi: 10.1021/acs.jcim.3c01749.

M. S. Smith, W. F. DeGrado, M. Grabe, and B. K. Shoichet. A cooperative model for symmetric ligand binding to protein fibrils. *Biochemistry*, 64(15):3382–3392, 2025. doi: 10.1021/acs.biochem.5c00068.

H. Stark, O.-E. Ganea, L. Pattanaik, R. Barzilay, and T. Jaakkola. EquiBind: geometric deep learning for drug binding structure prediction. In *International Conference on Machine Learning*, 2022.

Y. Sun, Y. Yang, H. Peng, Y. Shen, Y. Yang, H. Hu, L. Qiu, and H. Koike. ImageBrush: Learning visual in-context instructions for exemplar-based image manipulation. In *Advances in Neural Information Processing Systems*, volume 36, pages 48723–48743, 2023. doi: 10.52202/075280-2115.

Y. Sun, W. Chu, H. Zhou, K. Wang, and H. Koike. AVI-Talking: Learning audio-visual instructions for expressive 3D talking face generation. *IEEE Access*, 12:57288–57301, 2024. doi: 10.1109/ACCESS.2024.3390182.

Y. Sun, B. Li, M. Zhuge, D.-P. Fan, S. Khan, F. S. Khan, and H. Koike. Connecting dreams with visual brainstorming instruction. *Visual Intelligence*, 3(1), 2025. doi: 10.1007/s44267-025-00081-2. Article 12.

Y. Sun, Z. Xu, H. Zhou, J. Guan, Q. Yang, K. Wang, B. Liang, Y. Li, H. Feng, J. Wang, Z. Liu, and H. Koike. Cosh-DiT: Co-speech gesture video synthesis via hybrid audio-visual diffusion transformers. *International Journal of Computer Vision*, 134(3), 2026. doi: 10.1007/s11263-026-02752-z. Article 85.

O. Trott and A. J. Olson. AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of Computational Chemistry*, 31(2):455–461, 2010.

Y. Wang, K. Sun, J. Li, X. Guan, O. Zhang, D. Bagni, and T. Head-Gordon. A workflow to create a high-quality protein–ligand binding dataset for training, validation, and prediction tasks. *Digital Discovery*, 2025a. doi: 10.1039/D4DD00357H.

Z. Wang, J. Fan, R. Guo, T. Nguyen, H. Ji, and G. Liu. ProteinZero: self-improving protein generation via online reinforcement learning. *arXiv preprint arXiv:2506.07459*, 2025b.

J. Wohlwend, G. Corso, S. Passaro, M. Reveiz, K. Leidal, W. Swiderski, T. Portnoi, I. Chinn, J. Silterra, T. Jaakkola, and R. Barzilay. Boltz-1: democratizing biomolecular interaction modeling. *bioRxiv*, 2024. doi: 10.1101/2024.11.19.624167.

A. W. Wojciechowska, J. W. Wojciechowski, and M. Kotulska. Non-standard proteins in the lens of AlphaFold 3 — a case study of amyloids. *bioRxiv*, 2024. doi: 10.1101/2024.07.09.602655.

Z. Xue, J. Wu, Y. Gao, F. Kong, L. Zhu, M. Chen, Z. Liu, W. Liu, Q. Guo, W. Huang, et al. DanceGRPO: unleashing GRPO on visual generation. *arXiv preprint arXiv:2505.07818*, 2025.

Y. Zhang, H. Cai, C. Shi, B. Zhong, and J. Tang. E3Bind: an end-to-end equivariant network for protein-ligand docking. In *International Conference on Learning Representations*, 2023.

K. Zheng, H. Chen, H. Ye, H. Wang, Q. Zhang, K. Jiang, H. Su, S. Ermon, J. Zhu, and M.-Y. Liu. DiffusionNFT: online diffusion reinforcement with forward process. In *International Conference on Learning Representations*, 2026.

G. Zhou, Z. Gao, Q. Ding, H. Zheng, H. Xu, Z. Wei, L. Zhang, and G. Ke. Uni-Mol: a universal 3D molecular representation learning framework. In *International Conference on Learning Representations*, 2023.

X. Zhou, D. Xue, R. Chen, Z. Zheng, L. Wang, and Q. Gu. Antigen-specific antibody design via direct energy-based preference optimization. In *Advances in Neural Information Processing Systems*, 2024.

M. Zielinski, C. Röder, and G. F. Schröder. Challenges in sample preparation and structure determination of amyloids by cryo-EM. *Journal of Biological Chemistry*, 297(1):100938, 2021. doi: 10.1016/j.jbc.2021.100938.

Appendix

This appendix provides additional details supporting the main paper. §A describes the architecture, supervised warm-up, the synthetic data-curation pipeline used to construct *AmyDock-Synthetic*, and the *AmyDock* benchmark. §B presents the full reinforcement-learning algorithm, the cooperative binding reward design, and the pairwise stacking potential. §C provides additional qualitative visualisations of **CORAL**'s predicted fibril–ligand poses.

A Warm-up and Data Curation

This section provides additional details on the warm-up training stage and the synthetic data-curation pipeline used to construct *AmyDock-Synthetic*. Because amyloid fibril–ligand co-crystal structures are scarce, we curate a synthetic dataset of symmetric stacking poses via an agent-driven pipeline (§A.2), and warm up the model on this dataset to first establish the symmetric binding mode at a geometric level (§A.1) before reinforcement-learning fine-tuning toward cooperative binding energetics (Appendix B).

A.1 Initialisation and Warm-up Training

Initialisation and Generic Docking Pre-Training. We train the Protenix Pairformer–Diffusion backbone AI for Science Team, ByteDance Research [2025] from random initialisation on a diverse collection of protein–ligand co-crystals using the supervised objective defined below. This generic docking pre-training combines PLINDER Durairaj et al. [2024], Binding MOAD Hu et al. [2005], HiQBind Wang et al. [2025a], and SPINDR Cremer et al. [2025] to expose the model to a broad range of small-molecule geometries and chemotypes. Amyloid-specific warm-up then teaches periodic stacking geometry, followed by RL refinement of cooperative binding energetics. The added ConstraintEmbedder projections are zero-initialised as described below.

Warm-up Objective. For the warm-up stage, we adopt the standard AlphaFold-3-style structure-prediction objective Abramson et al. [2024], AI for Science Team, ByteDance Research [2025], which combines three groups of losses: (i) a *diffusion* loss $\mathcal{L}_{\text{diff}}$ that supervises atomic coordinates (instantiating $\mathcal{L}_{\text{FM}}$ in Eq. (2)), (ii) a *distogram* loss $\mathcal{L}_{\text{dist}}$ that supervises pairwise residue distances, and (iii) a *confidence* loss $\mathcal{L}_{\text{conf}}$ that trains the model to predict its own per-atom and per-pair errors. The total warm-up objective is

$$\mathcal{L}_{\text{SFT}} \;=\; \alpha_{\text{diff}}\,\mathcal{L}_{\text{diff}} \;+\; \alpha_{\text{dist}}\,\mathcal{L}_{\text{dist}} \;+\; \alpha_{\text{conf}}\,\mathcal{L}_{\text{conf}}, \tag{9}$$

with each component defined as

$$\mathcal{L}_{\text{diff}} = \mathcal{L}_{\text{MSE}} + \alpha_{\text{bond}}\,\mathcal{L}_{\text{bond}} + w_{\text{lddt}}\,\mathcal{L}_{\text{smooth-LDDT}}, \tag{10}$$

$$\mathcal{L}_{\text{conf}} = \mathcal{L}_{\text{pLDDT}} + \mathcal{L}_{\text{PDE}} + \alpha_{\text{pae}}\,\mathcal{L}_{\text{PAE}}. \tag{11}$$

Here $\mathcal{L}_{\text{MSE}}$ is the weighted per-atom coordinate MSE (with element-specific weights $w_{\text{protein}}{=}1$, $w_{\text{DNA}}{=}w_{\text{RNA}}{=}5$, $w_{\text{ligand}}{=}10$), $\mathcal{L}_{\text{bond}}$ is a bond-length consistency loss, $\mathcal{L}_{\text{smooth-LDDT}}$ is the smooth-LDDT loss, $\mathcal{L}_{\text{dist}}$ is the distogram cross-entropy over binned C_β–C_β distances, and $\mathcal{L}_{\text{pLDDT}}, \mathcal{L}_{\text{PDE}}, \mathcal{L}_{\text{PAE}}$ are the three confidence-head classification losses. The global weights are empirically set as $\alpha_{\text{diff}}{=}4.0$, $\alpha_{\text{dist}}{=}0.03$, $\alpha_{\text{conf}}{=}10^{-4}$, $w_{\text{lddt}}{=}1.0$, while α_{bond} and α_{pae} act as switches that are set to 0 during early training and turned on ($=1$) in the final fine-tuning stage.

Geometry Constraint. Beyond the loss design, we inject two fibril-specific priors into the pair representation $\mathbf{z}^{\text{init}}$ via the ConstraintEmbedder (§3.2). The *contact constraint* $\mathbf{C}^{\text{ct}}$ encodes the translational periodicity of cross-β amyloid fibrils, which are stacks of identical β-strand layers separated by a characteristic inter-strand C_α–C_α spacing of $\sim$4.8 Å. For each pair of residue tokens (i, j) we compute the pairwise distance $\mathbf{C}^{\text{ct}}_{ij} = \|\mathbf{r}_i - \mathbf{r}_j\|_2$ from the reference fibril structure and inject it through a learned linear projection $\mathbf{W}_{\text{ct}}$ (Eq. 6); the constraint then propagates to every Pairformer block, biasing attention toward the periodic ladder of inter-chain contacts characteristic of cross-β assemblies. The *pocket constraint* $\mathbf{C}^{\text{pkt}}$ marks the candidate binding region on the fibril surface (the residues bordering a cross-β groove), providing a soft spatial prior over where ligands should bind. It is injected with the same mechanism through a separate projection $\mathbf{W}_{\text{pkt}}$. Both projections are

zero-initialised so that the constraint branch does not perturb the pair representation at the start of training. Following classifier-free guidance Ho and Salimans [2021], $\mathbf{C}^{\text{pkt}}$ is dropped with probability 0.5 during warm-up; the model thus learns both with and without the pocket prior, and the prior strength can be modulated at inference by interpolating between the conditional and unconditional velocity fields.

A.2 Data Curation

Resolved co-crystal structures of amyloid fibril–ligand complexes are exceptionally scarce: even with recent cryo-EM advances, only a handful have been structurally characterised, making supervised training of fibril-specific docking models infeasible from experimental data alone. To overcome this bottleneck, we synthesise plausible cooperative stacking poses through an agent-driven pipeline (Figure 7) and validate them via geometric, LLM-based, and expert checks, yielding the *AmyDock-Synthetic* training set used in the warm-up stage. The remainder of this subsection describes each stage in detail.

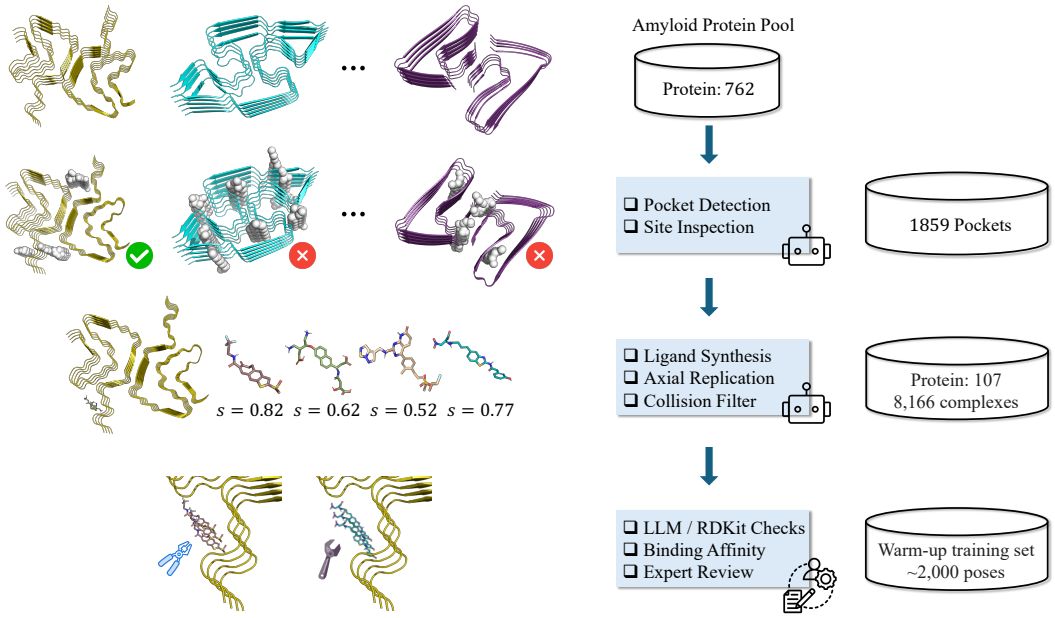

Figure 7: Data curation for the *AmyDock-Synthetic* warm-up training set. Pocket detection and inspection precede ligand generation, axial replication, and collision filtering. LLM/RDKit checks, docking-based affinity screening, and expert review yield approximately 2,000 training poses; $s \in [0, 1]$ denotes the LLM pose-plausibility screening score.

Curation Workflow. The workflow comprises three sequential stages: (1) pocket detection and site inspection; (2) ligand generation, axial replication, and collision filtering; and (3) LLM-based plausibility screening, RDKit validation, docking-based affinity screening, and expert structural review. We detail these operations and their intermediate outputs below.

Pocket detection and site inspection. We identify candidate binding sites on the fibril surface using two complementary tools, Fpocket Le Guilloux et al. [2009] and CaverDock Filipovič et al. [2020], and merge their outputs followed by manual expert inspection to discard non-physical or duplicate sites. This yields **1,859 candidate pockets** across all fibril targets.

Conformer generation. For each pocket, we generate **100 candidate ligands** using a fine-tuned variant of DrugFlow Schneuing et al. [2025]; we fine-tune DrugFlow on training targets that contain β-sheet structures, biasing the generative distribution toward chemotypes more compatible with cross-β groove binding. This stage takes approximately **one week on 4 NVIDIA A100 GPUs**.

Symmetric stacking and collision filtering. Each conformer is replicated along the fibril axis and tested for inter-ligand and ligand–protein steric clashes (atom-pair distance < 2.0 Å). After this filter, **8,166 candidate stacking arrangements** remain.

Pose screening and expert review. Surviving arrangements undergo coarse LLM-based screening for binding-mode plausibility (prompts in §A.2) and chemical-validity checks with RDKit's sanitisation routines. Candidates that pass these checks undergo docking-based affinity screening followed by expert structural review; experts additionally adjust ambiguous cases. The resulting *AmyDock-Synthetic* warm-up training set contains approximately **2,000 poses**.

LLM-Based Pose Screening. GPT-5 provides a qualitative plausibility screen for the stacking poses produced by DrugFlow and axial replication. Given a PDB excerpt of a fibril–ligand complex, it assesses whether the ligand occupies a groove, aligns with the fibril axis, and makes plausible contacts. This screen removes obvious binding-mode mismatches before the remaining chemical, affinity, and expert assessments. The model returns a binary verdict and a screening score in $[0, 1]$ using the following prompts.

System Prompt

```
You are an expert structural biologist and computational chemist specializing
in amyloid fibril structures and small molecule binding.  Your task is to
evaluate whether a small molecule is properly docked in the grooves/channels
of a fibril protein structure.  Fibril proteins have characteristic cross-β
sheet structures with grooves or channels running along the fibril axis where
small molecules can bind.

Key characteristics of proper fibril groove docking:
(1) Location:  the molecule should be positioned in the groove/channel between
protein chains, not on the surface or outside the fibril.
(2) Orientation:  the molecule should align along the fibril axis, not
perpendicular to it.
(3) Interactions:  favorable contacts with groove walls, including Van der
Waals contacts (3.0-5.0 Å), hydrogen bonds (2.5-3.5 Å), π-π stacking (3.3-4.5
Å) where applicable, and hydrophobic contacts.
(4) Shape complementarity:  the molecule should fit well into the groove
geometry.
(5) No clashes:  no severe steric clashes (<2.0 Å).
```

User Prompt

```
I have a PDB structure file containing a fibril protein with a docked small
molecule.  Please evaluate whether the small molecule is properly docked in
the fibril groove/channel.

PDB File Content:  {pdb_content}

Evaluation Criteria:
(1) Location:  is the molecule positioned in a groove/channel between protein
chains (not on surface)?
(2) Orientation:  is the molecule aligned along the fibril axis?
(3) Interactions:  does the molecule form favorable interactions with groove
residues (VdW, H-bond, π-π, hydrophobic)?
(4) Shape complementarity:  does the molecule fit well into the groove
geometry?
(5) No clashes:  no severe steric clashes (<2.0 Å).

Output Format:
Answer:  [Yes or No]
Score:  [number between 0.0 and 1.0]
```

Candidates are retained when the LLM returns a positive verdict and a screening score of at least 0.5; the score is not used for ranking. It expresses qualitative pose plausibility rather than a physical energy, binding affinity, or calibrated probability. Axial replication, collision filtering, RDKit sanitisation, affinity screening, and expert review provide the complementary geometric and chemical assessment.

AmyDock Training and Evaluation Sets. The approximately 2,000 curated *AmyDock-Synthetic* poses produced by this pipeline are used for warm-up training. The *AmyDock-Synthetic* evaluation benchmark comprises a separate held-out set of 57 fibril–ligand pairs. Each evaluation pose undergoes stricter validation than the warm-up training poses: a more stringent binding-affinity criterion, an energy-minimisation check requiring a small heavy-atom RMSD change between the initial and relaxed poses, and manual expert structural review. Both the synthetic evaluation set and the experimentally-resolved *AmyDock-CryoEM* set consist of two-protofilament fibrils with at least three consecutive chains per protofilament, covering both intra- and inter-protofilament binding interfaces.

Ligand Scaffold Overlap. To assess chemical overlap between the training and evaluation sets beyond exact ligand identity, we compare Bemis–Murcko scaffolds Bemis and Murcko [1996], which represent a molecule's ring systems and connecting linkers. Amyloid binders often share extended, planar aromatic motifs, so fingerprint similarity can remain high across this ligand class. Scaffold comparison therefore examines whether ligands share the same core molecular framework. We count an overlap when an evaluation ligand has the same scaffold as any training ligand. Every evaluation ligand in *AmyDock-CryoEM* and *AmyDock-Synthetic* has a Bemis–Murcko scaffold absent from the training set, yielding zero scaffold overlap for both benchmarks. This separation concerns core scaffolds; local chemical motifs may still be shared.

A.3 Implementation Details

Table 3 summarises the backbone dimensions, sampling settings, and optimisation parameters. The warm-up stage uses a learning rate of 1×10^{-4}, while the RL stage uses a lower learning rate of 3×10^{-6} for training stability. Both stages are trained on 8 NVIDIA A100 80G GPUs; the warm-up stage takes approximately 3 days and the RL stage approximately 7 days. The weight EMA with decay 0.999 is distinct from the sampling-policy EMA in §B.1. Additional RL hyperparameters are given in Table 4.

Table 3: Model architecture and optimisation settings. Stage-specific values are labelled explicitly.

Hyperparameter	**Value**
Representations	
Single (token) width	384
Pair width	128
Atom / atom-pair width	128 / 16
Pairformer trunk	
Blocks	48
Attention heads	16
Dropout	0.25
Recycling iterations	10
Diffusion module	
Token width	768
Atom encoder / decoder blocks	3 / 3
Atom encoder / decoder heads	4 / 4
Token transformer blocks	24
Token transformer heads	16
Sampling steps (training roll-outs)	20
Sampling steps (inference)	200
Heads and embedders	
Distogram bins	64
ConstraintEmbedder	Contact and pocket projections, zero-initialised
MSA module / template embedder	Disabled / disabled

Continued on next page

Table 3: Model architecture and optimisation settings (continued).

Hyperparameter	Value
Optimisation	
Optimiser	Adam ($\beta_1 = 0.9$, $\beta_2 = 0.95$)
Weight decay	10^{-8}
Learning rate (warm-up / RL)	1×10^{-4} / 3×10^{-6}
Gradient clipping (warm-up / RL)	10 / 1
Weight EMA decay	0.999
Crop size (warm-up)	384 tokens
Roll-outs per complex G	24

B Reinforcement Learning with Cooperative Reward

This section details the reinforcement-learning stage of **CORAL** used in §3.3. We first present the training objective based on DiffusionNFT Zheng et al. [2026] (§B.1), then describe how the composite reward combines protein–ligand docking, ligand–ligand stacking modelled by a Mie (8–4) potential, and symmetry consistency (§B.2).

B.1 Training Objective

DiffusionNFT Zheng et al. [2026] is an online reinforcement-learning algorithm for diffusion (and flow-matching) models that operates directly on the forward process and avoids explicit likelihood estimation.

Forward Process and Sampling. Given a clean sample $\mathbf{x}_0 \sim p_{\text{data}}$ and a noise sample $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the *forward process* interpolates between data and noise according to a schedule (α_t, σ_t):

$$\mathbf{x}_t \;=\; \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}, \qquad t \in [0, 1], \tag{12}$$

with target velocity

$$\mathbf{v}(\mathbf{x}_t, t) \;=\; \dot{\alpha}_t \mathbf{x}_0 + \dot{\sigma}_t \boldsymbol{\epsilon}. \tag{13}$$

Standard flow matching trains a velocity field v_θ by $\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}\,\|v_\theta(\mathbf{x}_t, t) - \mathbf{v}\|_2^2$. At inference, samples are generated by integrating the learned velocity field along the reverse direction: starting from $\mathbf{x}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and solving $\mathrm{d}\mathbf{x}_t/\mathrm{d}t = v_\theta(\mathbf{x}_t, t)$ from $t = 1$ to $t = 0$ to obtain a clean sample $\mathbf{x}_0$.

Implicit Positive / Negative Policies. Given a frozen sampling policy v_{old} and the current policy v_θ, DiffusionNFT defines two *implicit* velocity fields through reparameterisation:

$$v_\theta^+(\mathbf{x}_t, t) := (1 - \beta)\, v_{\text{old}}(\mathbf{x}_t, t) + \beta\, v_\theta(\mathbf{x}_t, t), \tag{14}$$

$$v_\theta^-(\mathbf{x}_t, t) := (1 + \beta)\, v_{\text{old}}(\mathbf{x}_t, t) - \beta\, v_\theta(\mathbf{x}_t, t), \tag{15}$$

where $\beta \in (0, 1)$ is a mixing coefficient. Here v_θ^+ is the *positive* policy (interpolating from v_{old} towards v_θ) and v_θ^- is the *negative* policy (extrapolating away from v_θ); both are fully determined by v_θ, with no auxiliary guidance networks.

Reward-Weighted Loss. Let $\tilde{r}(\mathbf{x}_0) \in [0, 1]$ denote the *normalised reward weight* of a sampled clean trajectory endpoint $\mathbf{x}_0$ (we describe how we obtain $\tilde{r}$ from the raw composite reward r in the next paragraph). The NFT loss minimised at each optimisation step is

$$\mathcal{L}_{\text{NFT}}(\theta) \;=\; \mathbb{E}\big[\, \tilde{r}(\mathbf{x}_0)\, \|v_\theta^+(\mathbf{x}_t, t) - \mathbf{v}\|_2^2 \;+\; (1 - \tilde{r}(\mathbf{x}_0))\, \|v_\theta^-(\mathbf{x}_t, t) - \mathbf{v}\|_2^2 \,\big], \tag{16}$$

with the expectation taken over $(\mathbf{x}_0, \boldsymbol{\epsilon}, t)$. High-reward samples are regressed by the *positive* branch v_θ^+ and low-reward samples by the *negative* branch v_θ^-, so the policy is simultaneously pushed towards v_θ^+ and away from v_θ^-. We refer to Zheng et al. [2026] for the gradient analysis showing this corresponds to a soft importance-weighted policy improvement without explicit likelihood ratios.

Reward Normalisation. Our composite reward $r(\mathbf{x}) = -R_{\text{bind}}(\mathbf{x}) + \lambda\, R_{\text{sym}}(\mathbf{x})$ (Eq. (7)) assigns higher values to more favourable poses. Group-relative normalisation maps this reward to the weight $\tilde{r} \in [0, 1]$ used in the NFT loss, comparing G on-policy rollouts $\{\mathbf{x}^{(i)}\}_{i=1}^{G}$ sampled from the same complex:

$$\tilde{r}^{(i)} \;=\; \frac{r(\mathbf{x}^{(i)}) - \min_j r(\mathbf{x}^{(j)})}{\max_j r(\mathbf{x}^{(j)}) - \min_j r(\mathbf{x}^{(j)}) + \varepsilon}. \tag{17}$$

This increasing map preserves the reward ordering, assigning larger weights to higher-reward rollouts and hence greater emphasis to the positive branch v_θ^+. The small constant $\varepsilon = 10^{-6}$ prevents division by zero when all rewards are equal.

Sampling Policy Update. To keep the on-policy assumption approximately valid while preventing v_{old} from lagging too far behind v_θ, we update v_{old} by a soft exponential moving average (EMA) after every N_{upd} gradient steps:

$$\theta_{\text{old}} \;\leftarrow\; \eta\, \theta_{\text{old}} + (1 - \eta)\, \theta, \tag{18}$$

with $\eta \in [0, 1]$ the EMA coefficient.

KL Regularisation. To preserve the stacking geometry learned during warm-up, we regularise the RL policy towards the frozen warm-up model v_{ref}:

$$\mathcal{L}_{\text{KL}}(\theta) \;=\; \mathbb{E}_{\mathbf{x}_t, t}\big[\, \|v_\theta(\mathbf{x}_t, t) - v_{\text{ref}}(\mathbf{x}_t, t)\|_2^2 \,\big], \tag{19}$$

and optimise the combined objective $\mathcal{L}_{\text{NFT}}(\theta) + \beta_{\text{KL}}\, \mathcal{L}_{\text{KL}}(\theta)$. This velocity-space penalty acts as a proximal regulariser rather than an explicit KL divergence.

Hyperparameters. The mixing coefficient β, EMA coefficient η, and group size G follow the defaults of DiffusionNFT Zheng et al. [2026]. We update v_{old} every N_{upd}=8 epochs (rather than every step) to give the behaviour policy a more stable optimisation window on our relatively small cooperative-binding training set, and increase the KL weight from the DiffusionNFT default of 10^{-3} to β_{KL}=10^{-2} to constrain policy drift from the warmed-up reference. The symmetry-reward weight λ was tuned on a held-out subset of fibrils.

Table 4: DiffusionNFT hyperparameters used in the RL stage.

Symbol	Value	Description
β	0.1	Mixing coefficient in $v_\theta^{\pm}$ (Eq. 14–15)
λ	0.01	Weight of R_{sym} in the composite reward
β_{KL}	0.01	KL regularisation weight
η	0.5	EMA coefficient for v_{old} (Eq. 18)
N_{upd}	8	Gradient steps between v_{old} updates
G	24	Rollouts per complex for reward normalisation (Eq. 17)

B.2 Reward Design

The composite reward $r(\mathbf{x}) = -R_{\text{bind}}(\mathbf{x}) + \lambda\, R_{\text{sym}}(\mathbf{x})$ (§3.3) favours energetically compatible ligand stacks with consistent geometry across neighbouring binding sites. The binding term R_{bind} combines protein–ligand docking and ligand–ligand packing, while R_{sym} penalises deviations from the fibril's repeating geometry. With $\lambda > 0$, maximising r favours lower binding scores and smaller symmetry deviations.

Binding Score. Small-molecule ligands that bind amyloid fibrils occupy periodic binding sites along the cross-β groove, where adjacent ligands are separated by the inter-strand distance ($\sim$4.7–4.9 Å). This periodic arrangement gives rise to cooperative ligand–ligand interactions that are not captured by standard protein–ligand docking scores Smith et al. [2024, 2025]. We therefore augment the AutoDock Vina protein–ligand affinity with a stacking-energy correction:

$$R_{\text{bind}} \;=\; \Delta G_{\text{PL}} + f \cdot E_{\text{stack}}, \tag{20}$$

where ΔG_{PL} is the empirical Vina docking score (kcal/mol) Trott and Olson [2010], Eberhardt et al. [2021], and E_{stack} sums Mie (8–4) pair interactions with short-distance regularisation between the centre ligand and its neighbouring copies. We use ΔG_{eff} as shorthand for this combined binding score. It provides an approximate energetic objective for pose optimisation and ligand ranking, rather than a calibrated estimate of solution-phase binding free energy.

Ligand–Ligand Stacking Potential. We model inter-ligand packing using a Mie (8–4) potential with short-distance regularisation. For an atom pair (i, j) separated by distance r, let $d_{ij} = R_i + R_j$ denote its preferred contact distance, with R_i and R_j the assigned atomic contact radii. The regularised pair potential is

$$E_{\mathrm{Mie}}(r; d_{ij}) = \varepsilon \left[\left(\frac{d_{ij}}{\bar{r}} \right)^8 - 2 \left(\frac{d_{ij}}{\bar{r}} \right)^4 \right], \qquad \bar{r} = \max(r, s\, d_{ij}), \tag{21}$$

where $\varepsilon > 0$ sets the well depth and $s \in (0, 1)$ specifies a short-distance floor that keeps the energy finite for overlapping atoms. The underlying Mie (8–4) potential is a generalised Lennard–Jones potential: the eighth-power term provides short-range repulsion and the fourth-power term provides attraction, with a minimum at $r = d_{ij}$ and energy $-\varepsilon$. The distance floor regularises the short-range interaction while preserving this preferred contact distance.

Let A contain the heavy atoms of the centre ligand and B_k those of neighbouring copy k, generated by the fibril's inter-chain symmetry operation. The stacking contribution is

$$E_{\mathrm{stack}} = \sum_{k \in \mathcal{N}} \sum_{i \in A} \sum_{j \in B_k} \mathbf{1}\left[r_{ij}^{(k)} < r_c \right] E_{\mathrm{Mie}}\left(r_{ij}^{(k)}; d_{ij} \right), \tag{22}$$

where $\mathcal{N}$ is the set of neighbouring ligand copies, $r_{ij}^{(k)}$ is the corresponding atom-pair distance, and r_c is the pairwise cutoff. This stacking term contains no separate Coulomb contribution.

Stacking Weight. The empirical weight $f = 0.6$ balances ligand–ligand packing against the protein–ligand docking score. Its effective strength depends on the well depth ε and the atom-pair summation: for fixed geometry and contact radii, the weighted stacking term scales with $f\varepsilon$. This term models local packing interactions; explicit solvation and configurational entropy are outside its scope.

Symmetry Consistency. The second reward term encourages ligand poses at neighbouring sites to follow the protein's repeating geometry. Let $\mathbf{T} = (\mathbf{R}, \mathbf{t})$ map chain k onto its nearest neighbour $k{+}1$ within the same protofilament, corresponding to a translation of approximately $4.8\,\text{Å}$ along the fibril axis. We estimate $\mathbf{T}$ by Kabsch alignment Kabsch [1976] of the corresponding C_α atoms. For ligand poses $\mathbf{x}_k$ and $\mathbf{x}_{k+1}$ at these adjacent sites, symmetry consistency is measured by:

$$R_{\mathrm{sym}} \;=\; -\,\mathrm{RMSD}(\mathbf{T}\,\mathbf{x}_k,\; \mathbf{x}_{k+1})\,. \tag{23}$$

B.3 Training Dynamics and Stacking-Weight Sensitivity

Training Dynamics. Figure 8 shows how RL improves the cooperative-binding objective during training. The mean per-batch reward rises from a negative initial value to a positive plateau within approximately 25 epochs and remains stable thereafter (left). Improvement is also observed across individual complexes: comparing each complex's first and last training visits, R_{bind} decreases for 25 of 30 complexes (83%), with the largest improvement reaching approximately 9.5 kcal/mol (right). The remaining 5 complexes have increases smaller than 1.5 kcal/mol. Thus, the overall reward gain is accompanied by more favourable binding scores for most training complexes.

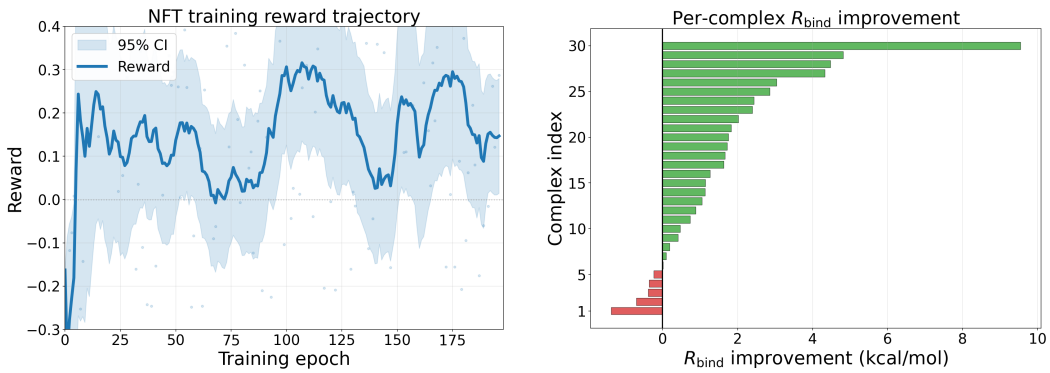

Figure 8: **RL training dynamics. Left:** mean per-batch reward over 200 epochs (shading: 95% CI). **Right:** per-complex binding-score improvement (kcal/mol; positive is better), sorted by magnitude.

Table 5: **Stacking-weight sensitivity and interface complementarity.** Mean geometric measures at three stacking weights f. Surface areas are reported per ligand; single and stacked refer to the corresponding interface calculations.

Weight	**Ligand area**	**Buried area** (Å^2)		ϕ_{bury}		S_c	
f	(Å^2)	Protein	Neighbours	Single	Stacked	Single	Stacked
0.2	688	340	230	0.507	0.836	0.559	0.662
0.5	686	347	236	0.514	0.853	0.554	0.657
0.8	689	340	246	0.507	0.855	0.540	0.644

Stacking-Weight Sensitivity. We examine how the stacking weight f affects surface burial and interface complementarity (Table 5). Solvent-accessible surface areas are computed with the Shrake–Rupley algorithm Shrake and Rupley [1973] using a 1.4 Å probe. For each ligand, we separate burial by the protein from the additional burial by neighbouring ligand copies; the buried fraction ϕ_{bury} normalises the buried area by the isolated ligand's surface area. The single-ligand reference evaluates each pose without neighbouring copies, whereas the stacked calculation includes them. The Lawrence–Colman shape complementarity S_c Lawrence and Colman [1993] compares the fibril groove with either the single ligand or the exposed surface of the ligand column.

Across the three weights, including neighbouring copies increases the buried fraction from approximately 0.51 to 0.84–0.86, while the ligand column has higher groove complementarity ($S_c = 0.64$–0.66) than the single ligand (0.54–0.56). Increasing the weight raises the area buried against neighbouring copies from 230 to 246 Å^2, accompanied by a modest decrease in stacked S_c from 0.662 to 0.644. These trends reflect a trade-off between inter-ligand packing and complementarity to the fibril groove.

B.4 Evaluation on General Protein–Ligand Docking

To assess docking beyond amyloid fibrils, we evaluate two training and input configurations on the 130-target PoseBusters subset of Morehead et al. [2025] (PDB depositions after 2021-09-30): the single-sequence model described above (**CORAL-single-seq**) and **CORAL-MSA**, trained on mixed PDB and amyloid data and evaluated with multiple-sequence alignments (MSAs). Blind inference uses no reference receptor or pocket constraints and selects the highest-confidence of five samples (200 diffusion steps, 10 recycling cycles). Relaxation and interaction scoring follow the PoseBench protocol, with all 130 targets included in the averages. Without relaxation, CORAL-MSA achieves 48.7% RMSD $\leq$ 2 Å and 35.1% RMSD $\leq$ 2 Å $\wedge$ PB-Valid, compared with 8.5% and 6.2% for CORAL-single-seq (Figure 9). PLIF-WM scores are similar (52.3 for CORAL-single-seq and 53.7 for CORAL-MSA). OpenMM relaxation does not further improve CORAL-MSA's pose success rates on this benchmark.

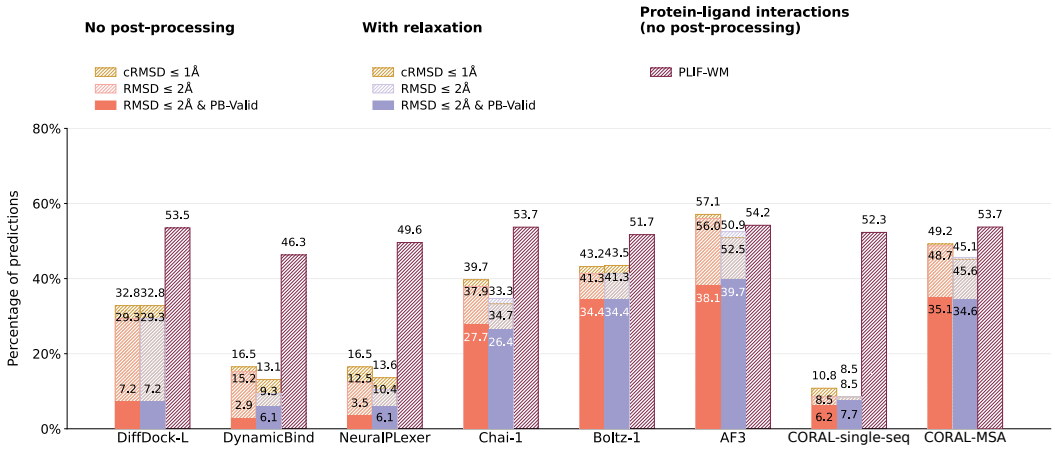

Figure 9: **General protein–ligand docking on PoseBusters.** Each method's three bars show pose accuracy before relaxation, after relaxation, and protein–ligand interaction similarity (PLIF-WM), respectively. Baseline results are from Figure 4 of Morehead et al. [2025].

C Additional Qualitative Results

We provide additional qualitative visualisations of **CORAL**'s predicted fibril–ligand poses on *AmyDock-CryoEM* and *AmyDock-Synthetic* targets, complementing the comparisons in §4.2.

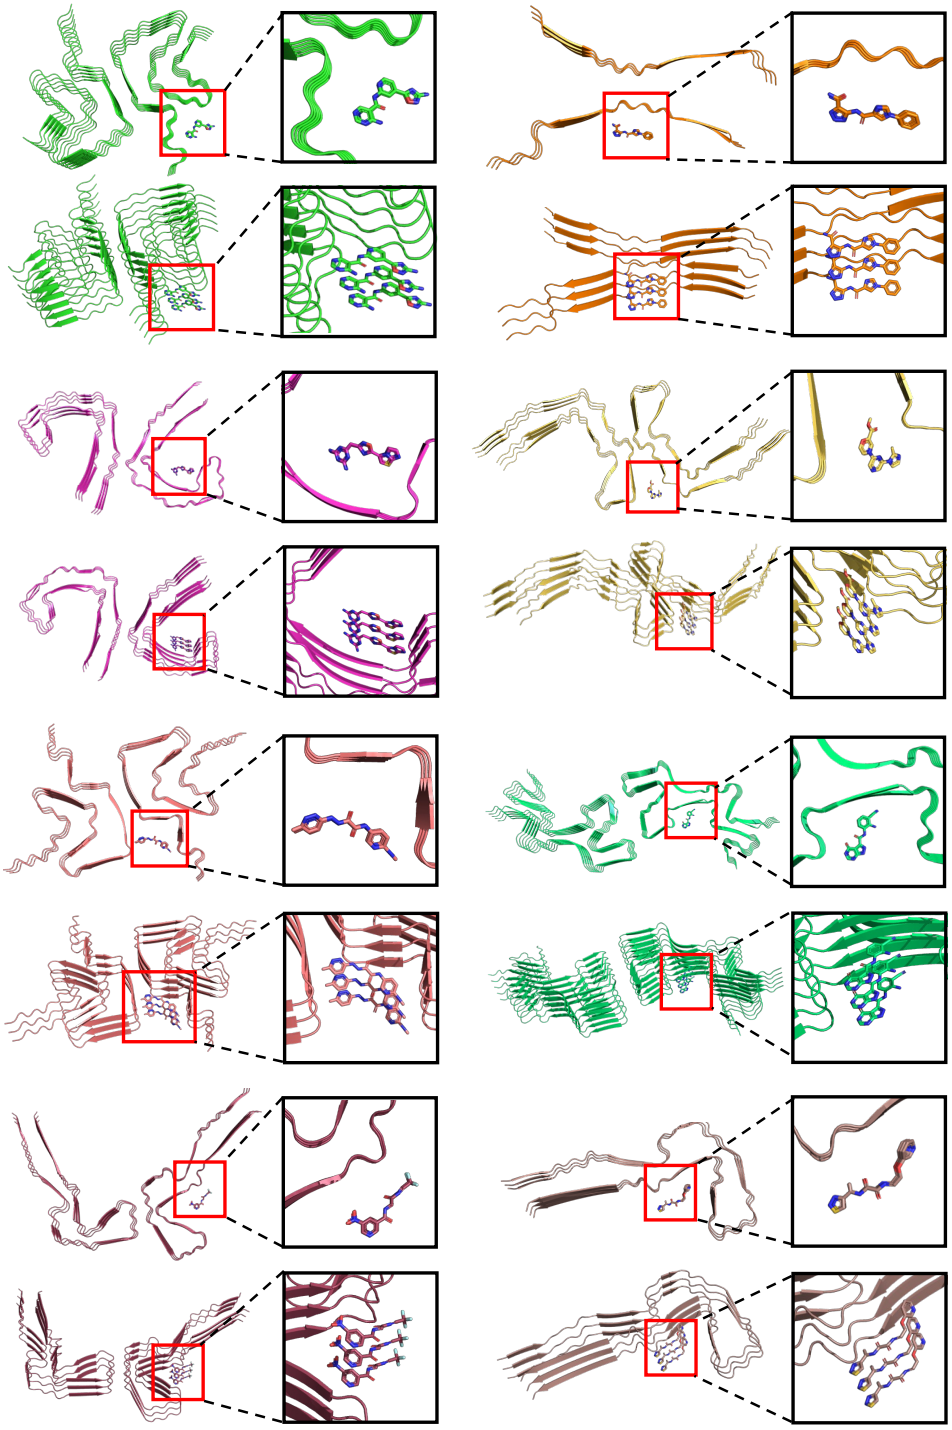

Figure 10: Additional qualitative visualisations of **CORAL**'s predicted fibril–ligand poses across *AmyDock* targets. Each case pairs a top-down view along the fibril axis with an oblique view, highlighting the local geometry of cooperative stacking and protein–ligand contacts.

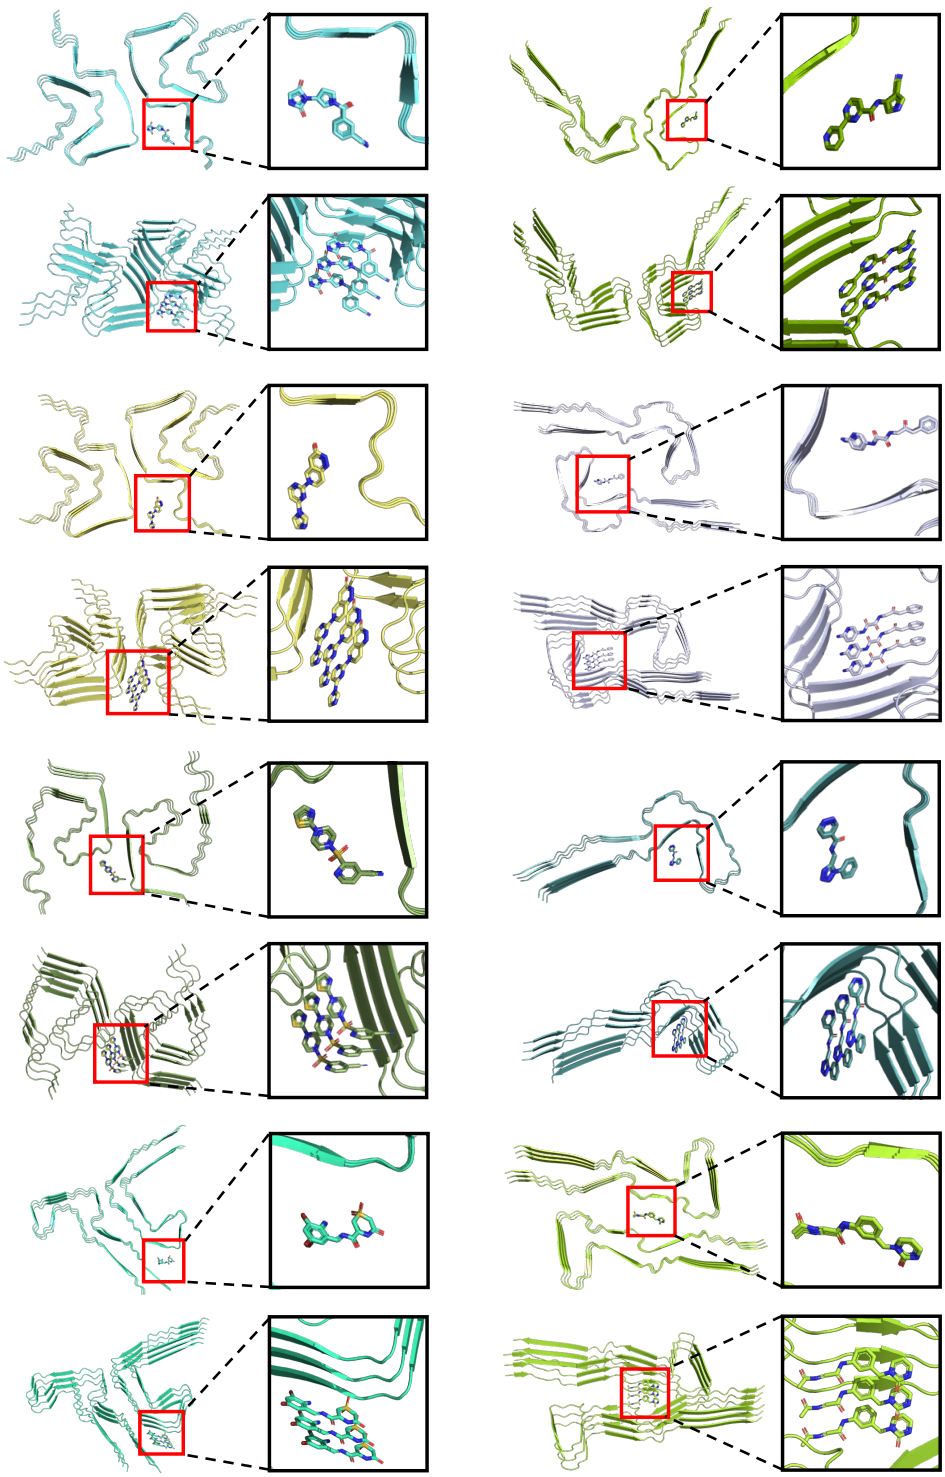

Figure 11: Additional qualitative visualisations of **CORAL**'s predicted fibril–ligand poses across *AmyDock* targets. Each case pairs a top-down view along the fibril axis with an oblique view, highlighting the local geometry of cooperative stacking and protein–ligand contacts.

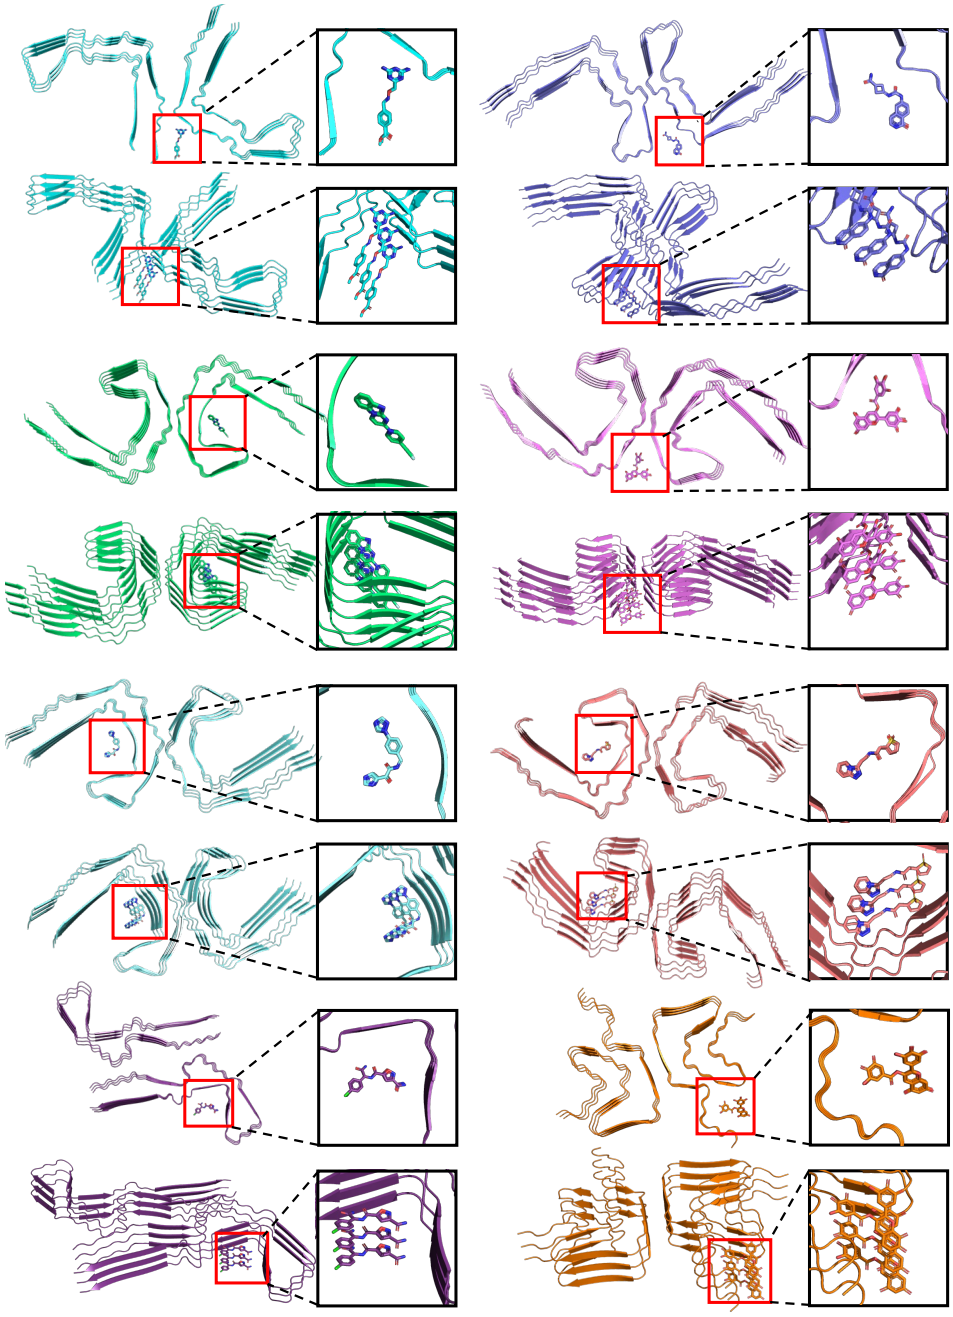

Figure 12: Additional qualitative visualisations of **CORAL**'s predicted fibril–ligand poses across *AmyDock* targets. Each case pairs a top-down view along the fibril axis with an oblique view, highlighting the local geometry of cooperative stacking and protein–ligand contacts.

NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: The abstract and Section 1 state three contributions: a deep-learning docking framework for amyloid fibril–ligand complexes, a two-stage training recipe that combines warm-up on curated stacking poses with reinforcement learning under a cooperative binding reward, and an evaluation on *AmyDock-CryoEM* and *AmyDock-Synthetic*. Sections 3.2 and 3.3 describe the method, and Table 1, Figure 4 and §4.3 support the results. The Limitations paragraph in §5 states the scope of these claims.

   Guidelines:

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

2. **Limitations**

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

   Answer: [Yes]

   Justification: The Limitations paragraph in §5 states that the evaluation covers periodic ligand-stacking geometries in tau, α-synuclein and amyloid-β fibrils, that the synthetic data pipeline and its prompt-based LLM pre-screen may introduce biases, that the training reward approximates intermolecular interactions with a docking score and a pairwise Mie (8–4) stacking potential, and that the affinity results need broader experimental binding data. §4.3 and Figure 6 show the failure modes.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

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

   Answer: [N/A]

   Justification: The paper presents an empirical method and contains no theoretical results or proofs. The training objectives and the reward are defined with their assumptions in §3.3 and Appendix B.

   Guidelines:

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

4. **Experimental result reproducibility**

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

   Answer: [Yes]

   Justification: §3.2 and Table 3 describe the architecture. §3.3 and Appendices A–B describe the data curation, the warm-up objective, the RL objective, the reward and the RL hyperparameters (Table 4). §4.1 and §4.3 define the datasets, baselines, metrics and ablations. We also release the inference code and the trained checkpoint through the project page linked in the abstract.

   Guidelines:

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

(a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
(b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
(c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

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

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

   Answer: [No]

   Justification: We release the inference code and the trained **CORAL** checkpoint, with installation and usage instructions, through the project page linked in the abstract. The training code and the curated *AmyDock* datasets are not part of this release. §3.3, §4.1 and Appendices A–B describe the data curation, the training recipe and the hyperparameters.

   Guidelines:

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

6. **Experimental setting/details**

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

   Answer: [Yes]

   Justification: §4.1 specifies the implementation, the evaluation datasets, the baselines and their fine-tuning on the *AmyDock-Synthetic* training set. Appendix A.3, Table 3 and Table 4 give the learning rates, the optimisation settings and the RL hyperparameters. Appendix A.2 separates the approximately 2,000 warm-up training poses from the 57 held-out evaluation pairs.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

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

   Answer: [No]

   Justification: We do not report error bars over repeated training runs, because each full warm-up and RL run takes approximately 10 days on 8 NVIDIA A100 GPUs (Appendix A.3). Table 1 reports the p-value of the Spearman correlation on α*Syn-Bind*, and Figure 8 shows the 95% confidence interval of the per-batch RL reward.

   Guidelines:

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

8. **Experiments compute resources**

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

   Answer: [Yes]

   Justification: Appendix A.3 reports that both training stages use 8 NVIDIA A100 80G GPUs, with approximately 3 days for warm-up and 7 days for RL. Appendix A.2 reports that conformer generation for data curation takes approximately one week on 4 NVIDIA A100 GPUs.

   Guidelines:

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

9. **Code of ethics**

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

Answer: [Yes]

Justification: The research conforms to the NeurIPS Code of Ethics. It does not involve human subjects or personal data, and it uses published protein structures and published ligand data. The released assets are inference code and model weights for docking-pose prediction.

Guidelines:

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

10. **Broader impacts**

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

Answer: [Yes]

Justification: The Broader Impact paragraph in §5 describes the intended benefit for early-stage drug discovery for neurodegenerative diseases and presents predicted poses as structural hypotheses for experimental evaluation. The Limitations paragraph describes where the predictions can be unreliable. The model predicts binding geometries for given fibril–ligand pairs and does not generate new compounds, so we see limited potential for direct misuse.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

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

Answer: [N/A]

Justification: The released model predicts docking geometries for user-supplied protein–ligand inputs. It does not generate new compounds, personal data or other content with a high risk for misuse, and we release no scraped data.

Guidelines:

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

12. **Licenses for existing assets**

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

    Answer: [Yes]

    Justification: We cite the original publications of all existing assets: the Protenix backbone AI for Science Team, ByteDance Research [2025] (Apache 2.0); the docking baselines PatchDock Schneidman-Duhovny et al. [2005], DiffDock Corso et al. [2023], DynamicBind Lu et al. [2024], FABind Pei et al. [2023], NeuralPLexer Qiao et al. [2024], FlowDock Morehead and Cheng [2025], Boltz-2 Passaro et al. [2025] and Chai-1 Boitreaud et al. [2024]; the pre-training datasets PLINDER Durairaj et al. [2024], Binding MOAD Hu et al. [2005], HiQBind Wang et al. [2025a] and SPINDR Cremer et al. [2025]; the tools fpocket Le Guilloux et al. [2009], CaverDock Filipovič et al. [2020], DrugFlow Schneuing et al. [2025] and AutoDock Vina Trott and Olson [2010], Eberhardt et al. [2021]; the ligand database of Chisholm and Hunter [2025], the αSyn-Bind affinity data Chu et al. [2015] and the PoseBusters benchmark Buttenschoen et al. [2024], Morehead et al. [2025]. Protein structures come from the Protein Data Bank, which is freely available under CC0. We use these assets for research under their respective licences and terms of use; most are released under permissive open-source licences such as MIT or Apache 2.0.

    Guidelines:

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

13. **New assets**

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

    Answer: [Yes]

    Justification: We release the **CORAL** inference code and the trained checkpoint under the Apache 2.0 licence through the project page linked in the abstract, with documentation of installation, inputs, outputs and examples. §3.3 and Appendix A.2 document how the *AmyDock* sets were built, filtered and reviewed, and how they are used for training and evaluation.

Guidelines:

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

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

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

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or research with human subjects. Biology researchers reviewed candidate fibril–ligand poses as domain experts during data curation (§3.3), not as paid crowdworkers or study participants.

    Guidelines:

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

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

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

    Answer: [N/A]

    Justification: The research does not involve human subjects.

    Guidelines:

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

16. **Declaration of LLM usage**

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

    Answer: [Yes]

Justification: GPT-5 provides a coarse pose-plausibility screen in the data-curation pipeline (§3.3). Appendix A.2 gives the full prompts, the acceptance rule (a positive verdict and a screening score of at least 0.5) and the limited role of the score.

Guidelines:

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