InteraFold-Chem: Integrating pretrained ligand representations for protein – ligand pose prediction

MindRank Foundation Model Team

Abstract

Accurate ligand placement remains challenging even when an all-atom model recovers the overall structure of a protein–ligand complex. We present InteraFold-Chem, which integrates pretrained ligand representations into an AlphaFold 3–inspired backbone through a gated residual interface. Two independent branches extend a shared checkpoint: continued structural training (InteraFold-Chem) and policy-gradient reinforcement learning that updates only the diffusion module (InteraFold-Chem-RL). The RL branch uses a protein structure as a soft template, allowing protein coordinates to change. On an internal validation set of 445 samples, InteraFold-Chem achieves 54.4% Success@2Å and a median protein-aligned ligand RMSD of 1.426 Å, compared with 48.5% and 2.316 Å for Protenix-v2. InteraFold-Chem-RL achieves 52.1% and 1.745 Å, with lower complex lDDT than InteraFold-Chem (0.779 versus 0.924). Reward feedback is used only during RL training. These descriptive results come from a non-public validation cohort under a uniform low-step configuration, with differing protein conditioning and no matched-compute or model-specific recommended-setting comparison. They do not establish external generalization or isolate the benefit of the RL objective.

Introduction

Predicting how a small molecule binds to a protein is a central problem in structure-based drug discovery. Classical docking combines a search over ligand configurations with an approximate scoring function, as exemplified by AutoDock4 and AutoDock Vina (Morris et al., 2009; Trott & Olson, 2010; Eberhardt et al., 2021). Learned scoring functions, including those used in GNINA, provide another way to evaluate candidate poses (McNutt et al., 2021). These approaches depend on the receptor structure and the effectiveness of conformational search and scoring. A plausible ligand conformation alone is insufficient: its position, orientation and contacts must also agree with the binding site.

Recent structure models have expanded the scope of joint molecular prediction. AlphaFold 2, AlphaFold-Multimer and RoseTTAFold established protein and protein-complex modeling frameworks that informed later all-atom systems (Jumper et al., 2021; Evans et al., 2021; Baek et al., 2021). AlphaFold 3 extends structure prediction to complexes containing proteins and small molecules, among other molecular components (Abramson et al., 2024). Boltz-1, Chai-1 and Boltz-2 further broaden the available approaches to biomolecular modeling (Wohlwend et al., 2024; Chai Discovery, 2024; Passaro et al., 2025). However, accurate global structure does not guarantee accurate ligand placement: a model can perform well on a complex-level metric while missing the binding pose.

Specialized docking models offer complementary approaches to ligand placement. EquiBind, TANKBind and DiffDock use learned geometric information to predict binding poses (Stärk et al., 2022; Lu et al., 2022; Corso et al., 2023), while Uni-Mol Docking V2 draws on molecular pretraining (Alcaide et al., 2024). DynamicBind and NeuralPLexer also account for changes in the protein environment (Lu et al., 2024; Qiao et al., 2024). Comparisons between these methods must consider their differing assumptions about the protein structures available as input.

Pretrained molecular representations provide another way to improve ligand modeling. Geometric learning methods such as SchNet and directional message passing capture relationships between atoms, while molecular pretraining learns transferable representations from three-dimensional molecular data (Schütt et al., 2017; Gasteiger et al., 2020; Zhou et al., 2023; Ji et al., 2024). We investigate whether these representations can be incorporated into an all-atom complex model through a small ligand-specific interface. This design preserves the structure-generation pathway and provides a means to study the contribution of chemical pretraining.

Training can also be adapted to favor more accurate ligand poses. Diffusion models learn to recover structures from noisy inputs, and prior work has explored reward-driven adaptation and distillation (Ho et al., 2020; Song & Ermon, 2019; Song et al., 2021; Karras et al., 2022; Black et al., 2023; Salimans & Ho, 2022). Reinforcement Learning (RL) adapts a decision-making policy using reward feedback, with the aim of increasing expected return (Sutton & Barto, 2018). Here, the diffusion denoiser defines a pose-generation policy under protein conditioning and is adapted with policy gradients. Only diffusion-module parameters are updated in the RL branch. Reference-based pose rewards are used during training; inference uses the adapted model without reward evaluation or a reference ligand.

We introduce InteraFold-Chem, a model for predicting small-molecule binding poses within protein–ligand complexes. It combines an AlphaFold 3–inspired backbone, pretrained ligand representations and a ligand-specific feature interface, with independent continued-training and policy-gradient RL branches (Fig. 1). We evaluate pose accuracy and complex lDDT on a non-public internal validation set of 445 samples and examine three selected systems in detail (Table 1; Fig. 3). Data are organized through PLINDER and PoseBusters (Durairaj et al., 2024; Buttenschoen et al., 2024). The continued-training branch generates complete complexes without a supplied protein template, whereas the RL branch uses a reference protein structure as a soft template rather than a fixed receptor.

Results

Pretrained ligand representations enter through a gated interface

InteraFold-Chem retains the main structure-generation pathway of an AlphaFold 3 – inspired model (Fig. 1a). A pretrained molecular model supplies ligand features with its parameters held fixed. These features are combined with the backbone's native representations through a learnable gate (Fig. 1b). Direct updates are restricted to ligand tokens, while the shared network allows information to pass between ligand and protein representations.

Preserving the backbone's starting behavior requires a zero effective fusion residual at initialization, not merely zero sigmoid logits (Methods). InteraFold-Chem continues training the backbone, including the Pairformer, and the ligand-fusion interface on protein–ligand data; InteraFold-Chem-RL uses independent policy-gradient adaptation (Fig. 1c). In RL, only the diffusion module is trainable; the Pairformer, ligand-fusion interface and all other non-diffusion modules are frozen. The pretrained molecular model remains fixed in both branches. The interface uses pretrained atom features only.

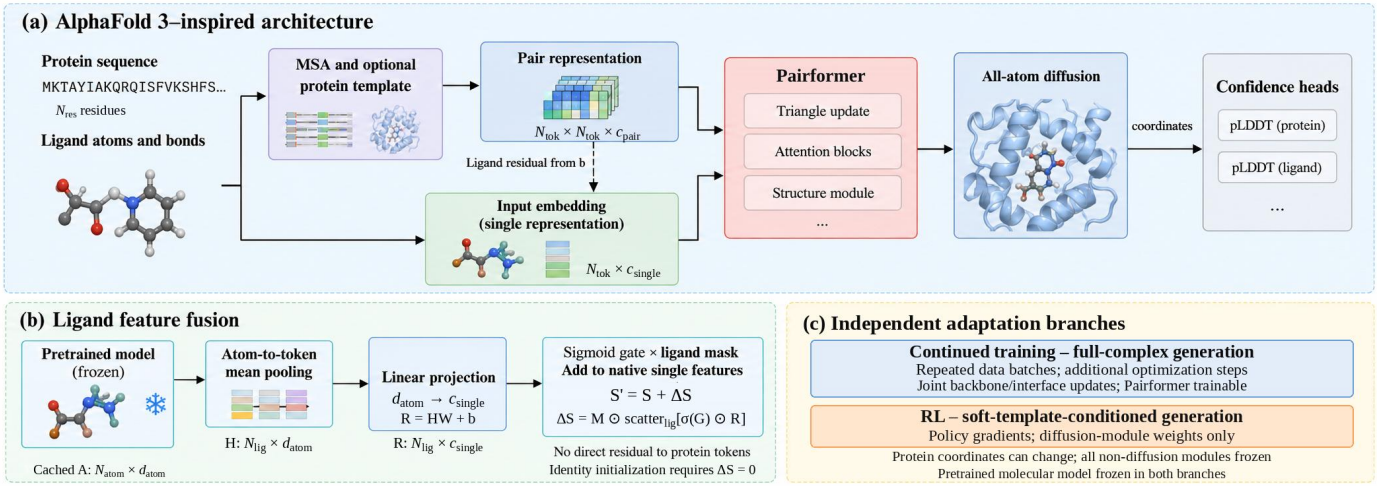

Fig. 1 | Architecture and adaptation of InteraFold-Chem. (a) An AlphaFold 3–inspired model predicts complex coordinates and confidence from protein and ligand inputs. (b) Pretrained ligand features enter through a gated residual interface; dimensions and the zero-residual initialization condition are defined in Methods. (c) Independent branches use continued structural training or policy-gradient RL. Continued training updates the backbone and ligand-fusion interface; RL updates only the diffusion module. RL uses a protein soft template, allowing protein geometry to change. All non-diffusion modules are frozen in RL; the pretrained molecular model remains fixed in both branches.

Internal comparison under a uniform low-step configuration

Under the uniform low-step configuration on the internal validation set of 445 samples, the reported Success@2Å values are 54.4% for InteraFold-Chem, 52.1% for InteraFold-Chem-RL and 48.5% for the baseline labeled Protenix-v2 (Zhang et al., 2026; Table 1; Fig. 2a). Chai-1 succeeds on 205 of 445 samples (46.1%), and Boltz-2 succeeds on 102 of 445 samples (22.9%). These descriptive point estimates place InteraFold-Chem highest in this restricted comparison.

InteraFold-Chem also has the lowest median protein-aligned ligand RMSD and the highest complex lDDT point estimates in this comparison (Table 1; Fig. 2b,c). Its median ligand RMSD is 1.426 Å, compared with 2.316 Å for Protenix-v2, 2.500 Å for Chai-1 and 7.753 Å for Boltz-2. The corresponding complex lDDT values are 0.924, 0.910, 0.852 and 0.600.

Policy-gradient RL adapts the diffusion module during training

InteraFold-Chem-RL applies policy-gradient reinforcement learning to protein-conditioned pose generation. Only diffusion-module weights are updated; the upstream backbone, ligand-fusion interface and other modules remain frozen. Training rewards favor ligand poses with lower RMSD after fitting the predicted protein to the reference protein. Policy-gradient updates are combined with inherited structural supervision. The reference protein supplies a soft template, not a coordinate constraint, so the generated protein can change conformation. Inference uses the adapted model without training-time exploration, reward evaluation or a reference ligand.

Table 1 | Internal validation under a uniform low-step configuration on a non-public cohort (n = 445).

Method	Success@2Å	Median ligand RMSD (Å)	Complex lDDT
InteraFold-Chem	**54.4%**	**1.426**	**0.924**
InteraFold-Chem-RL	52.1%	1.745	0.779
Protenix-v2	48.5%	2.316	0.910
Chai-1	46.1%	2.500	0.852
Boltz-2	22.9%	7.753	0.600

Success@2Å is the percentage of samples with ligand RMSD below 2 Å after protein-backbone fitting, without a ligand-only refit (Methods); percentages are rounded to one decimal place. Chai-1 and Boltz-2 have 205 and 102 successes out of 445. The non-public internal validation cohort includes the systems in Fig. 3. Nominal settings are 4 recycling cycles and 20 diffusion steps, not matched compute or model-specific recommended settings. Values are descriptive; uncertainty intervals, significance tests and recommended-configuration results are not reported. InteraFold-Chem-RL uses a reference-protein soft template and allows protein motion; Protenix-v2 and InteraFold-Chem use no supplied protein template. Chai-1 and Boltz-2 conditioning details are not fully documented here. Differences in conditioning, protein-loss weighting and trainable modules prevent attribution to the RL objective alone; these results do not establish external generalization.

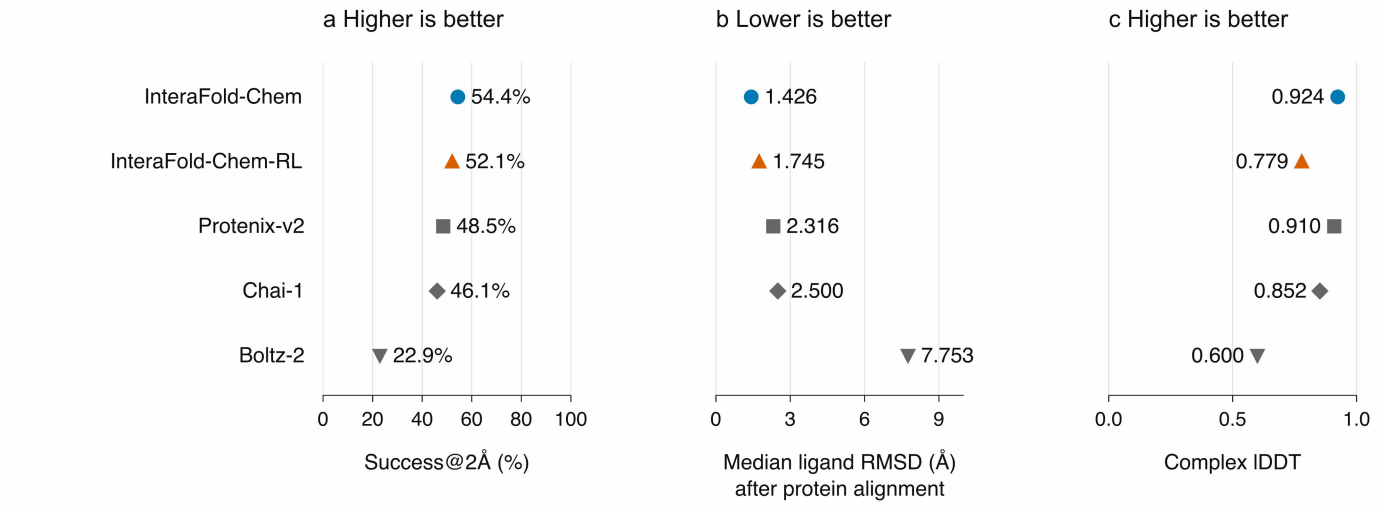

Fig. 2 | Internal validation under a uniform low-step configuration (n = 445). (a) Success@2Å (%). (b) Median protein-aligned ligand RMSD (Å), with no ligand-only refit. (c) Complex lDDT. Values reproduce Table 1; its notes specify the non-public cohort, differing inputs and configuration limitations. Higher values are better in a and c; lower values are better in b.

Three selected case studies illustrate differences in ligand placement

We examined 2FWP, 8GFD and 7UAW, three systems evaluated under the uniform low-step configuration in the internal validation set. Figure 3 presents the reference structures and model predictions in separate views, ordered from top to bottom as 2FWP, 8GFD and 7UAW. Table 2 reports the corresponding per-system metrics.

For 2FWP, InteraFold-Chem and InteraFold-Chem-RL achieve ligand RMSDs of 0.47 Å and 0.49 Å, respectively. Protenix-v2 also remains below the 2 Å threshold at 0.67 Å, while Chai-1 and Boltz-2 have RMSDs of 2.53 Å and 3.43 Å (Table 2). For 8GFD, InteraFold-Chem and InteraFold-Chem-RL achieve 1.10 Å and 1.57 Å, respectively; the values for Protenix-v2, Chai-1 and Boltz-2 are 12.21 Å, 15.24 Å and 8.37 Å. For 7UAW, both InteraFold-Chem variants reach 0.24 Å. Protenix-v2 also falls below the 2 Å threshold at 1.64 Å, whereas Chai-1 and Boltz-2 have RMSDs of 5.84 Å and 6.59 Å, respectively (Table 2).

In these three selected systems, both InteraFold-Chem variants produce ligand poses below the 2 Å RMSD threshold and have lower ligand RMSDs than the reported baselines. InteraFold-Chem-RL closely follows InteraFold-Chem but does not improve on it in these examples. The structural views in Figure 3 and measurements in Table 2 illustrate pose recovery in individual cases, not robustness across unselected systems.

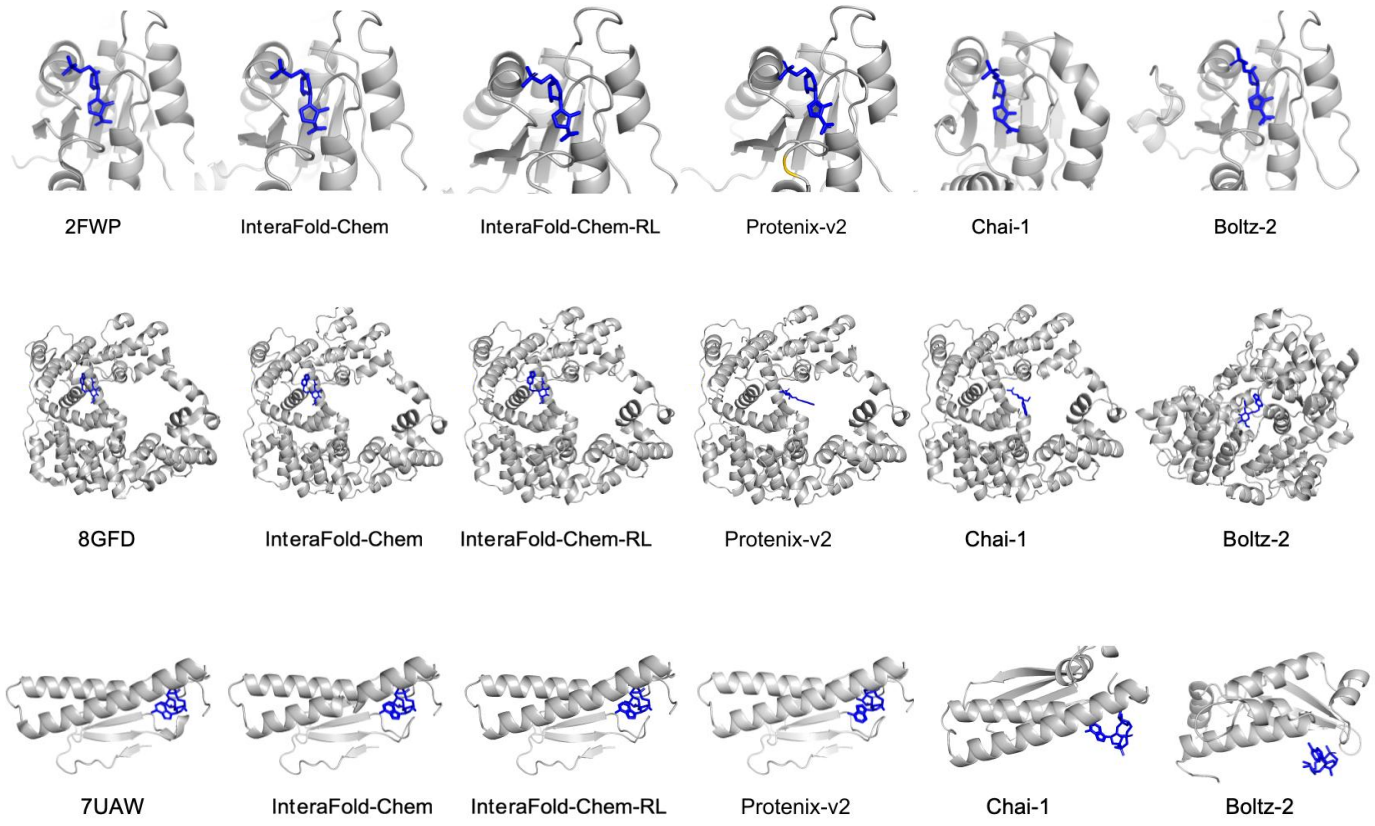

Fig. 3 | Structural comparisons for three selected systems under the uniform low-step configuration. Rows show 2FWP, 8GFD and 7UAW; columns show the reference, InteraFold-Chem, InteraFold-Chem-RL, Protenix-v2, Chai-1 and Boltz-2. Proteins are shown as gray cartoons and ligands as blue sticks. Panels are separate views; Table 2 provides quantitative protein-aligned ligand RMSDs.

Table 2 | Per-case pose accuracy under the uniform low-step configuration for the systems in Fig. 3.

System	Model	Ligand RMSD (Å)	Complex lDDT
2FWP	InteraFold-Chem	0.47	0.976
2FWP	InteraFold-Chem-RL	0.49	0.852
2FWP	Protenix-v2	0.67	0.951
2FWP	Chai-1	2.53	0.927
2FWP	Boltz-2	3.43	0.703
8GFD	InteraFold-Chem	1.10	0.953
8GFD	InteraFold-Chem-RL	1.57	0.861
8GFD	Protenix-v2	12.21	0.937
8GFD	Chai-1	15.24	0.874
8GFD	Boltz-2	8.37	0.467
7UAW	InteraFold-Chem	0.24	0.974
7UAW	InteraFold-Chem-RL	0.24	0.988
7UAW	Protenix-v2	1.64	0.891
7UAW	Chai-1	5.84	0.528
7UAW	Boltz-2	6.59	0.445

Systems and models follow the order in Fig. 3. Ligand RMSD is calculated after fitting the predicted protein to the reference and applying the same transform to the predicted ligand, without a ligand-only refit. Complex lDDT measures local distance agreement across the complex. Protenix-v2 refers to the reported baseline; InteraFold-Chem-RL is the policy-gradient branch with diffusion-only weight updates. Values are taken from the per-system summaries.

Discussion

InteraFold-Chem incorporates pretrained ligand representations into an AlphaFold 3 – inspired model of protein – ligand complexes and supports two forms of adaptation. A pretrained molecular model supplies chemical features with its parameters held fixed, while the structure model learns to use these features. Although the interface directly updates only ligand representations, information can propagate through the complex and affect protein geometry.

Within the uniform low-step configuration, the reported point estimates favor InteraFold-Chem across all three aggregate metrics (Table 1). InteraFold-Chem-RL achieves higher pose success and lower median ligand RMSD than Protenix-v2, but lower complex lDDT; it does not improve on InteraFold-Chem. Pose accuracy and local distance agreement across the complex should therefore be assessed separately.

The evaluation is reported as a non-public, study-specific internal comparison rather than a standard public benchmark evaluation. The full sample manifest is not released, and this report does not establish separation from adaptation or pretraining data; the results therefore do not establish generalization to unseen targets or ligands.

Values are descriptive point estimates without uncertainty intervals or significance tests. Equal nominal recycling and diffusion-step counts do not ensure matched compute or suitable settings for each model. Recommended-configuration, step-sensitivity and runtime measurements are not reported, so neither an efficiency advantage nor a causal explanation for Boltz-2's low score can be inferred.

The RL branch uses policy gradients together with inherited structural supervision and changes only diffusion-module weights. Its protein soft template conditions generation without fixing protein coordinates: frozen upstream weights do not imply a frozen receptor. Nevertheless, this conditioning provides structural information unavailable to the no-template branches. Differences in conditioning, protein-loss weighting and updated modules therefore prevent attribution of the observed changes to the policy-gradient objective alone. Matched-condition ablations are needed to isolate that contribution.

Accurate ligand placement alone does not establish chemical validity or reliable binding prediction. Pose quality should therefore be assessed along three complementary dimensions: protein-aligned pose accuracy, chemical and physical validity, and recovery of reference protein–ligand interactions. Cohort-wide PoseBusters validity and interaction-recovery statistics are not reported here. The three selected structural examples do not close this evidence gap, and the reported Success@2Å values should not be interpreted as the fraction of chemically valid or interaction-faithful predictions.

An expanded chemical-validity assessment should use PoseBusters to report intraligand steric clash rate, bond-length validity, bond-angle validity, chirality validity, internal-energy and geometry checks, and protein–ligand clash rate (Buttenschoen et al., 2024). These checks should include tetrahedral and double-bond stereochemistry, aromatic-ring and double-bond planarity, and a specified internal-energy plausibility test; internal energy is not a binding-affinity estimate. Per-check pass or failure rates should accompany an overall PoseBusters pass rate, defined by passing all applicable checks in a predeclared full chemical-and-physical validity panel, independently of the RMSD threshold. A separate joint endpoint, PB-valid Success@2Å, should report the fraction of all cohort samples whose selected pose both has protein-aligned ligand RMSD below 2 Å and passes the validity panel. Clash rates should denote the fraction of poses with at least one violation, rather than an unlabeled fraction of atom pairs.

Protein–ligand interaction fingerprints (PLIFs), derived from contacts identified with ProLIF or PLIP, would complement these tests by measuring recovery of reference interactions (Bouysset & Fiorucci, 2021; Adasme et al., 2021). Evaluation should include hydrogen-bond precision and recall, hydrophobic contact recovery, salt-bridge recovery, and π–π interaction recovery. Precision measures the fraction of predicted contacts also present in the reference; recall measures the fraction of reference contacts recovered. Contact matching should specify interaction type, corresponding protein residue, donor/acceptor roles where relevant, and ligand atom or group correspondence. Per-type precision and recall, together with an aggregate PLIF similarity, would distinguish missed reference contacts from additional non-reference contacts. Cases with no reference contacts of a given type require an explicit applicability and aggregation rule.

These analyses should use the same prespecified candidate-selection rule and full 445-sample cohort for all models, with missing or unreadable predictions retained as failures for overall pass and joint-success rates. Software versions, thresholds, molecular preparation, hydrogenation/protonation, and per-metric applicability and denominators should be documented. Because the models can change protein geometry, clash and interaction checks should first assess each predicted ligand with its own predicted receptor and compare the resulting contacts with the reference complex; any evaluation against a common reference receptor should be reported separately after protein alignment. Any refinement or minimization should likewise be evaluated separately from the original predictions. Until these measurements are available, the present findings support pose recovery under the reported configuration, not improved chemical validity, interaction recovery, or binding affinity.

Affinity prediction and ranking are prospective extensions of InteraFold-Chem. PDBbind, BindingDB and CASF provide resources for such studies. These extensions would require separate validation and careful treatment of assay differences and screening biases (Wang et al., 2004; Liu et al., 2007; Su et al., 2019; Mysinger et al., 2012; Chen et al., 2019). These datasets and objectives provide no additional evidence for the pose results reported here.

This study describes a practical way to incorporate pretrained ligand representations into an all-atom structure model. Independent evaluation and matched-condition component studies, including uncertainty estimates and chemical-validity checks, are needed to determine whether the observed pose improvements extend beyond the present internal validation setting.

Methods

Architecture and implementation provenance

The structure model follows an AlphaFold 3–inspired architecture, with protein and ligand embeddings, multiple-sequence-alignment information, optional protein conditioning, a Pairformer trunk, an all-atom diffusion denoiser and confidence heads (Abramson et al., 2024). A ligand-specific interface introduces pretrained ligand representations into this backbone. In the RL branch, the supplied reference protein structure is used as a soft template: it conditions structure generation but does not clamp protein coordinates. Both protein and ligand coordinates can change during denoising. Freezing upstream modules refers to their parameters, not to the predicted protein geometry.

Data sources and non-public internal validation

The study focuses on protein–small-molecule interfaces from Protein Data Bank structures, with data organized not only through PLINDER and PoseBusters but also other protein-compound structures from the Protein Data Bank database (Berman et al., 2000; Durairaj et al., 2024; Buttenschoen et al., 2024). Training emphasizes ligand-containing interfaces and applies size filtering or cropping where needed. Evaluation uses a filtered, internally curated validation cohort of 445 protein–small-molecule samples. The cohort is reported as a study-specific internal validation set, rather than as the complete benchmark or a standard split of either public resource. The same cohort is

used for the aggregate comparisons in Table 1 and Fig. 2; the three selected systems in Fig. 3 and Table 2 are members of this cohort, not an additional test set.

The complete sample manifest and processed validation data are not publicly released. The reported cohort is therefore identified as non-public internal validation, rather than a publicly reproducible benchmark. This report does not establish cohort disjointness from adaptation or model-pretraining data, and no claim of unseen-target or unseen-ligand generalization is made. Reference ligand coordinates are used for evaluation metrics; they are not provided as inference inputs.

Evaluation protocols and metrics

Ligand RMSD is calculated by first fitting the predicted protein to the reference protein, then comparing the ligands in that protein-aligned frame. A rigid-body least-squares fit is obtained from corresponding protein backbone atoms (Kabsch, 1976). The resulting rotation R and translation t are applied unchanged to the predicted ligand. RMSD is then calculated between corresponding transformed predicted and reference ligand atoms, without any additional ligand-only fitting. For predicted ligand coordinates x and reference ligand coordinates y, the metric is:

$$RMSD_{lig} = \sqrt{\frac{1}{N_{atom}} \sum_{i=1}^{N_{atom}} \|Rx_i + t - y_i\|^2}$$

Here N_{atom} is the number of corresponding ligand atoms included in the RMSD calculation, and R and t are determined only from the protein fit. Success@2Å is the percentage of samples with this ligand RMSD below 2 Å. We report its cohort median and complex lDDT, which measures local distance agreement across the complex without superposition (Mariani et al., 2013).

Inference configurations, sensitivity and computational cost

The reported internal comparison uses nominal inference settings of 4 recycling cycles and 20 diffusion steps for the 445-sample cohort (Table 1). These counts define the uniform low-step configuration; they do not establish equal compute or model-specific recommended operating points. Protein-conditioning differences are summarized in Table 1. Results at recommended configurations and controlled within-model step-count sensitivity measurements are not reported. A step-sensitivity experiment would need to hold the checkpoint, seeds, candidate budget, recycling, inputs, sampler family, other noise-schedule parameters and ranking fixed.

Wall-clock timing, GPU-hours and peak-memory measurements are not reported, so no inference-efficiency or sampling-speed advantage is established. A compute-matched comparison would require a common measured budget and hardware, with all generated candidates included. Model-only timing should be distinguished from end-to-end processing, including MSA/template preparation, feature construction, generation, ranking and any relaxation.

Ligand feature fusion and structural objective

A pretrained molecular model supplies ligand atom features with its parameters held fixed. In Fig. 1, N_{res}, N_{lig}, N_{tok} and N_{atom} denote protein residues, ligand tokens, all complex tokens and ligand atoms, respectively. $N_{tok} = N_{res} + N_{lig}$

only for one-token-per-residue protein–ligand inputs without other components. Single and pair representations span the full N_{tok} axis; c_{single} and c_{pair} denote their feature widths.

With pretrained feature width d_{atom}, atom features A ($N_{atom} \times d_{atom}$) are mean-pooled to H ($N_{lig} \times d_{atom}$) and linearly projected to R = HW + b ($N_{lig} \times c_{single}$); W has shape $d_{atom} \times c_{single}$, and b has width c_{single}. Pooling preserves feature width; no concatenation is implied by this schematic. The schematic masked gated residual can be expressed as:

$$S' = S + \Delta S, \quad \Delta S = M \odot \text{scatter}_{lig}[\sigma(G) \odot (HW + b)].$$

Here S has shape $N_{tok} \times c_{single}$, M is an $N_{tok} \times 1$ binary ligand mask broadcast over channels, G denotes broadcast-compatible gate logits, and $scatter_{lig}$ inserts ligand rows into the complex-token axis with zeros elsewhere; ⊙ denotes elementwise multiplication. A ligand mask restricts direct updates to ligand tokens, while the shared structure model allows information exchange across the complex. The interface uses atom features only; pair features are not used.

Identity at initialization requires ΔS = 0 for every valid input. Since σ(0) = 0.5, zero gate logits alone do not close the residual branch. For the affine residual above, setting both W and b to zero (or using a bias-free zero projection) is sufficient, irrespective of the gate logits.

Before any optimizer update, initialization should be tested against the unmodified backbone with identical backbone weights, inputs, evaluation mode and stochastic draws. The fused features must equal the native features; denoiser outputs and final coordinates should match within stated dtype-specific tolerances. A nonzero-residual mask test should also confirm that non-ligand features are unchanged at the fusion site.

Training retains the inherited structural objective, which combines coordinate, geometry, local-distance and confidence terms. The diffusion formulation follows established denoising methods (Ho et al., 2020; Song & Ermon, 2019; Song et al., 2021; Karras et al., 2022). The RL branch additionally reduces the weight assigned to protein-atom errors.

Continued training on protein–ligand data

InteraFold-Chem continues optimization on protein–ligand training data for additional steps under the inherited structural objective. Training repeatedly uses batches from these data to extend the optimization trajectory. The structure backbone, including the Pairformer, and ligand-fusion module are updated jointly, while the pretrained molecular model remains fixed.

Continued training for InteraFold-Chem and policy-gradient RL for InteraFold-Chem-RL are independent branches initialized from the same model checkpoint (Table 3). The reported InteraFold-Chem results use an early checkpoint from the continued-training trajectory and therefore do not represent the final planned training state.

Policy-gradient reinforcement learning for ligand pose generation

InteraFold-Chem-RL uses policy-gradient reinforcement learning to adapt the diffusion denoiser with feedback from training-reference structures. The noisy coordinates, denoising timestep and conditioning features define the policy

state, and a denoising continuation is treated as an action. Policy-gradient updates use pose rewards to optimize the expected return of this denoising policy (Sutton & Barto, 2018; Black et al., 2023). Protein conditioning is supplied as a soft template, so it does not require the predicted protein to retain the template coordinates.

During RL training, candidate poses are generated through diffusion and receive higher rewards when their ligand RMSD to the training reference is lower. The reward uses the same protein-fit-then-ligand-RMSD convention described above, with no ligand-only refit. Reward feedback drives policy-gradient updates to the diffusion policy; differentiating the RMSD reward itself is not required. The RL update is combined with the inherited structural objective, with reduced weighting of protein-atom errors. Reference ligand coordinates provide training-time reward feedback and structural supervision, not an inference input.

Only the diffusion-module weights are updated in this branch. All non-diffusion modules, including the upstream representation backbone, Pairformer, ligand-fusion interface and confidence heads, remain frozen; the pretrained molecular model is also fixed. Inference retains soft-template protein conditioning and uses the adapted diffusion module without the training-time exploration procedure, reward evaluation or a reference ligand. The generated protein and ligand coordinates remain free to change under this conditioning.

Table 3 | Overview of the independent continued-training and policy-gradient RL branches. RL updates only the diffusion module; all non-diffusion modules are frozen. The pretrained molecular model remains fixed in both branches.

Item	Continued training	RL
Initialization	Shared model checkpoint	Shared model checkpoint
Training objective	Inherited structural supervision; continued optimization on training data	Inherited structural supervision and policy-gradient RL
Model updates	Structure backbone (including Pairformer) and ligand-fusion module	Diffusion module only; all non-diffusion modules frozen
Inference	Full-complex generation	Soft-template-conditioned generation; protein coordinates can change

Data and code availability

PLINDER, PoseBusters and the Protein Data Bank are public source resources cited above. The curated 445-sample validation set, its full sample manifest and processed evaluation files are not publicly released. Aggregate results and three selected case studies are reported in Tables 1–2 and Fig. 3. The MindRank Foundation Model Team plans to announce the usage release of InteraFold-Chem through a future website according to company news.

References

1. Morris, G. M. et al. AutoDock4 and AutoDockTools4: automated docking with selective receptor flexibility. J. Comput. Chem. 30, 2785–2791 (2009). https://doi.org/10.1002/jcc.21256

2. Trott, O. & Olson, A. J. AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. J. Comput. Chem. 31, 455–461 (2010). https://doi.org/10.1002/jcc.21334

3. Eberhardt, J., Santos-Martins, D., Tillack, A. F. & Forli, S. AutoDock Vina 1.2.0: new docking methods, expanded force field, and Python bindings. J. Chem. Inf. Model. 61, 3891–3898 (2021). https://doi.org/10.1021/acs.jcim.1c00203

4. McNutt, A. T. et al. GNINA 1.0: molecular docking with deep learning. J. Cheminform. 13, 43 (2021). https://doi.org/10.1186/s13321-021-00522-2

5. Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. Nature 596, 583–589 (2021). https://doi.org/10.1038/s41586-021-03819-2

6. Evans, R. et al. Protein complex prediction with AlphaFold-Multimer. Preprint at bioRxiv (2021). https://doi.org/10.1101/2021.10.04.463034

7. Baek, M. et al. Accurate prediction of protein structures and interactions using a three-track neural network. Science 373, 871–876 (2021). https://doi.org/10.1126/science.abj8754

8. Abramson, J. et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature 630, 493–500 (2024). https://doi.org/10.1038/s41586-024-07487-w

9. Wohlwend, J. et al. Boltz-1: democratizing biomolecular interaction modeling. Preprint at bioRxiv (2024). https://doi.org/10.1101/2024.11.19.624167

10. Chai Discovery. Chai-1: decoding the molecular interactions of life. Preprint at bioRxiv (2024). https://doi.org/10.1101/2024.10.10.615955

11. Passaro, S. et al. Boltz-2: towards accurate and efficient binding affinity prediction. Preprint at bioRxiv (2025). https://doi.org/10.1101/2025.06.14.659707

12. Stärk, H., Ganea, O., Pattanaik, L., Barzilay, R. & Jaakkola, T. EquiBind: geometric deep learning for drug binding structure prediction. Proc. Mach. Learn. Res. 162, 20503–20521 (2022). https://proceedings.mlr.press/v162/stark22b.html

13. Lu, W. et al. TANKBind: trigonometry-aware neural networks for drug-protein binding structure prediction. Adv. Neural Inf. Process. Syst. 35, 7236–7249 (2022). https://openreview.net/forum?id=MSBDFwGYwwt

14. Corso, G., Stärk, H., Jing, B., Barzilay, R. & Jaakkola, T. DiffDock: diffusion steps, twists, and turns for molecular docking. International Conference on Learning Representations (2023). https://arxiv.org/abs/2210.01776

15. Alcaide, E. et al. Uni-Mol Docking V2: towards realistic and accurate binding pose prediction. Preprint at arXiv (2024). https://arxiv.org/abs/2405.11769

16. Lu, W. et al. DynamicBind: predicting ligand-specific protein–ligand complex structure with a deep equivariant generative model. Nat. Commun. 15, 1071 (2024). https://doi.org/10.1038/s41467-024-45461-2

17. Qiao, Z., Nie, W., Vahdat, A., Miller, T. F. III & Anandkumar, A. State-specific protein–ligand complex structure prediction with a multiscale deep generative model. Nat. Mach. Intell. 6, 195–208 (2024). https://doi.org/10.1038/s42256-024-00792-z

18. Schütt, K. T. et al. SchNet: a continuous-filter convolutional neural network for modeling quantum interactions. Adv. Neural Inf. Process. Syst. 30, 992–1002 (2017). https://arxiv.org/abs/1706.08566

19. Gasteiger, J., Groß, J. & Günnemann, S. Directional message passing for molecular graphs. International Conference on Learning Representations (2020). https://arxiv.org/abs/2003.03123

20. Zhou, G. et al. Uni-Mol: a universal 3D molecular representation learning framework. International Conference on Learning Representations (2023). https://openreview.net/forum?id=6K2RM6wVqKu

21. Ji, X. et al. Uni-Mol2: exploring molecular pretraining model at scale. Preprint at arXiv (2024). https://arxiv.org/abs/2406.14969

22. Ho, J., Jain, A. & Abbeel, P. Denoising diffusion probabilistic models. Adv. Neural Inf. Process. Syst. 33, 6840–6851 (2020). https://arxiv.org/abs/2006.11239

23. Song, Y. & Ermon, S. Generative modeling by estimating gradients of the data distribution. Adv. Neural Inf. Process. Syst. 32 (2019). https://arxiv.org/abs/1907.05600

24. Song, Y. et al. Score-based generative modeling through stochastic differential equations. International Conference on Learning Representations (2021). https://arxiv.org/abs/2011.13456

25. Karras, T., Aittala, M., Aila, T. & Laine, S. Elucidating the design space of diffusion-based generative models. Adv. Neural Inf. Process. Syst. 35 (2022). https://arxiv.org/abs/2206.00364

26. Black, K., Janner, M., Du, Y., Kostrikov, I. & Levine, S. Training diffusion models with reinforcement learning. Preprint at arXiv (2023). https://arxiv.org/abs/2305.13301

27. Salimans, T. & Ho, J. Progressive distillation for fast sampling of diffusion models. International Conference on Learning Representations (2022). https://arxiv.org/abs/2202.00512

28. Durairaj, J. et al. PLINDER: the protein–ligand interactions dataset and evaluation resource. Preprint at bioRxiv (2024). https://doi.org/10.1101/2024.07.17.603955

29. Buttenschoen, M., Morris, G. M. & Deane, C. M. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. Chem. Sci. 15, 3130–3139 (2024). https://doi.org/10.1039/D3SC04185A

30. Sutton, R. S. & Barto, A. G. Reinforcement Learning: An Introduction. 2nd edn. MIT Press (2018). https://mitpress.mit.edu/9780262039246/reinforcement-learning/

31. Corso, G. et al. Deep confident steps to new pockets: strategies for docking generalization. Preprint at arXiv (2024). https://arxiv.org/abs/2402.18396

32. Francoeur, P. G., Masuda, T. & Koes, D. R. 3D convolutional neural networks and a CrossDocked dataset for structure-based drug design. Preprint at ChemRxiv, version 2 (2020). https://doi.org/10.26434/chemrxiv.11833323.v2

33. Bouysset, C. & Fiorucci, S. ProLIF: a library to encode molecular interactions as fingerprints. J. Cheminform. 13, 72 (2021). https://doi.org/10.1186/s13321-021-00548-6

34. Adasme, M. F. et al. PLIP 2021: expanding the scope of the protein–ligand interaction profiler to DNA and RNA. Nucleic Acids Res. 49, W530–W534 (2021). https://doi.org/10.1093/nar/gkab294

35. Wang, R., Fang, X., Lu, Y. & Wang, S. The PDBbind database: collection of binding affinities for protein–ligand complexes with known three-dimensional structures. J. Med. Chem. 47, 2977–2980 (2004). https://doi.org/10.1021/jm030580l

36. Liu, T., Lin, Y., Wen, X., Jorissen, R. N. & Gilson, M. K. BindingDB: a web-accessible database of experimentally determined protein–ligand binding affinities. Nucleic Acids Res. 35, D198–D201 (2007). https://doi.org/10.1093/nar/gkl999

37. Su, M. et al. Comparative assessment of scoring functions: the CASF-2016 update. J. Chem. Inf. Model. 59, 895–913 (2019). https://doi.org/10.1021/acs.jcim.8b00545

38. Mysinger, M. M., Carchia, M., Irwin, J. J. & Shoichet, B. K. Directory of useful decoys, enhanced (DUD-E): better ligands and decoys for better benchmarking. J. Med. Chem. 55, 6582–6594 (2012). https://doi.org/10.1021/jm300687e

39. Chen, L. et al. Hidden bias in the DUD-E dataset leads to misleading performance of deep learning in structure-based virtual screening. PLoS ONE 14, e0220113 (2019). https://doi.org/10.1371/journal.pone.0220113

40. ByteDance AML AI4Science Team et al. Protenix—advancing structure prediction through a comprehensive AlphaFold3 reproduction. Preprint at bioRxiv (2025). https://doi.org/10.1101/2025.01.08.631967

41. Zhang, Y. et al. Protenix-v2: Broadening the Reach of Structure Prediction and Biomolecular Design. Preprint at bioRxiv (2026). https://doi.org/10.64898/2026.04.10.717613

42. Ahdritz, G. et al. OpenFold: retraining AlphaFold2 yields new insights into its learning mechanisms and capacity for generalization. Nat. Methods 21, 1514–1524 (2024). https://doi.org/10.1038/s41592-024-02272-z

43. Mirdita, M. et al. ColabFold: making protein folding accessible to all. Nat. Methods 19, 679–682 (2022). https://doi.org/10.1038/s41592-022-01488-1

44. Berman, H. M. et al. The Protein Data Bank. Nucleic Acids Res. 28, 235–242 (2000). https://doi.org/10.1093/nar/28.1.235

45. Kabsch, W. A solution for the best rotation to relate two sets of vectors. Acta Crystallogr. A 32, 922–923 (1976). https://doi.org/10.1107/S0567739476001873

46. Mariani, V., Biasini, M., Barbato, A. & Schwede, T. lDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. Bioinformatics 29, 2722–2728 (2013). https://doi.org/10.1093/bioinformatics/btt473

47. Meli, R. & Biggin, P. C. spyrmsd: symmetry-corrected RMSD calculations in Python. J. Cheminform. 12, 49 (2020). https://doi.org/10.1186/s13321-020-00455-2

48. Boltz developers. Prediction documentation, docs/prediction.md, revision 296ec7e. Official Boltz repository (accessed 17 September 2026). https://github.com/jwohlwend/boltz/blob/296ec7e/docs/prediction.md

49. Chai Discovery. Chai-1 inference implementation, chai_lab/chai1.py, revision 66c38d1. Official chai-lab repository (accessed 17 September 2026). https://github.com/chaidiscovery/chai-lab/blob/66c38d1/chai_lab/chai1.py

50. Boltz developers. Boltz-2 model implementation, src/boltz/model/models/boltz2.py, revision b1ebfc4. Official Boltz repository (accessed 17 September 2026). https://github.com/jwohlwend/boltz/blob/b1ebfc4/src/boltz/model/models/boltz2.py

51. ByteDance AML AI4Science Team. Supported models, docs/supported_models.md, section "Protenix-v2 (Enhanced Capacity Model)". Official Protenix repository, main branch (accessed 17 September 2026). https://github.com/bytedance/Protenix/blob/main/docs/supported_models.md