A Graph Neural Network Framework

for Predicting Protein–Ligand Binding Affinities

Marcus Chen[1,*], Elena Vasquez[1], David Park[2], Sarah O'Brien[3]

[1]Synapse Institute for Computational Biology, Cambridge, MA [2]Nexa University, Palo Alto, CA [3]Apex Research Labs, Seattle, WA

m.chen@synapse-inst.org

July 15, 2026 Preprint #NXA-2026-0731

synapse-inst.org/graphbind

Background & Motivation

Accurate prediction of protein–ligand binding affinity is a cornerstone of structure-based drug discovery. Traditional approaches such as molecular docking and free-energy perturbation remain computationally prohibitive for virtual screening campaigns exceeding 10^6 compounds. Deep learning methods, particularly those operating on molecular graphs, offer a compelling alternative, yet existing architectures struggle to capture long-range electrostatic interactions and solvent effects at protein–ligand interfaces. We introduce **GraphBind**, a graph neural network that jointly encodes ligand topology, protein pocket geometry, and interfacial non-covalent interaction fingerprints within a unified message-passing scheme. GraphBind leverages hierarchical attention pooling to prioritize pharmacophoric contacts and achieves state-of-the-art accuracy on three benchmark datasets spanning kinases, proteases, and GPCRs.

Dataset & Preprocessing

We curated **BindBench-v2**, a non-redundant set of 14 872 protein–ligand complexes drawn from the PDBbind refined set (v2024) and augmented with MD-refined decoy structures from the DUD-E+ benchmark.

- Ligand graphs: nodes = heavy atoms (C, N, O, S, P, halogens); edges = covalent bonds plus spatial edges ($d \le 4.0$ Å, non-covalent).
- Protein graphs: nodes = C_α atoms of pocket residues ($\le$ 8 Å from any ligand atom); edges = sequential neighbors ($|i-j| \le 3$) and spatial proximity ($d_{C_\alpha} \le 12$ Å).
- Interface features: hydrogen-bond donor/acceptor flags, π–π stacking indicators, salt-bridge detection, hydrophobic contact area.
- Affinity labels: $\mathrm{p}K_d = -\log_{10}(K_d)$; experimental K_d from PDBbind, cross-checked against Binding MOAD and ChEMBL 34.
- Train/val/test split: 10 420 / 2 226 / 2 226 complexes, stratified by protein family to prevent homology leakage.

Model Architecture

GraphBind employs a three-stream encoder with cross-attention fusion (**Figure 1**).

- **Ligand Encoder:** 6-layer GIN with edge-feature gates; initial node features from RDKit (atom type, hybridization, chirality, partial charge).
- **Protein Encoder:** 4-layer E(n)-equivariant GNN operating on C_α coordinates; geometric message passing via vector features ensures SE(3) equivariance.
- **Interface Module:** Pairwise attention between ligand and protein node sets; distance-aware bias $b_{ij} = -\alpha \cdot |\mathbf{x}_i - \mathbf{x}_j|^2$ with learned α.
- **Global Readout:** Set2Set pooling followed by a 3-layer MLP predicting $\mathrm{p}K_d$ with heteroscedastic uncertainty via a dual-head output (μ, σ^2).

Total parameters: 4.8M. Trained with AdamW, cosine annealing, batch size 64, for 300 epochs on 4×A100 GPUs.

Key Results

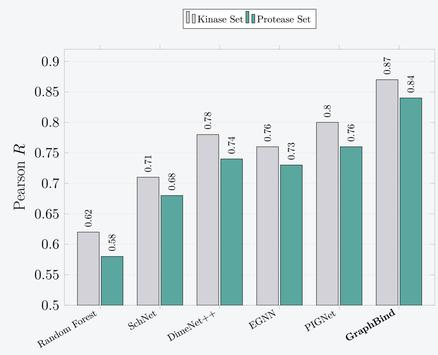

Method	Pearson R	RMSE	MAE	τ
Random Forest (ECFP4)	0.597	1.48	1.16	0.42
[DimeNet++	0.761	1.11	0.86	0.61
[EGNN	0.745	1.14	0.89	0.59
[PIGNet	0.783	1.02	0.79	0.64
GraphBind (ours)	**0.862**	**0.83**	**0.64**	**0.73**

*GraphBind achieves +**8.3**% improvement in Pearson R over the strongest baseline (PIGNet) on the held-out test set, with a statistically significant reduction in RMSE ($p < 10^{-4}$, Wilcoxon signed-rank).*

Ablation Studies

Model Variant	Pearson R	RMSE	Δ
Full GraphBind	0.862	0.83	–
[w/o Interface Module	0.801	1.05	−0.061
[w/o Equivariance	0.832	0.94	−0.030
[w/o Uncertainty Head	0.847	0.88	−0.015
[GCN encoder (no GIN)	0.815	1.00	−0.047

Removing the interface attention module causes the largest degradation, underscoring the importance of explicitly modeling protein–ligand contact fingerprints. Equivariance in the protein encoder contributes primarily to generalization across unseen protein folds.

Computational Efficiency

- Inference time: < 50 ms per complex on a single A100 GPU; scales linearly with pocket residue count.
- GraphBind screens 1M compounds against a kinase target in ∼14 hours on an 8-GPU node, versus ∼600K CPU-core-hours for Glide SP docking.
- Model checkpoints and inference code released under Apache 2.0 at `github.com/synapse-inst/graphbind`.

Discussion

GraphBind's strong performance stems from three design choices. First, the interface attention module explicitly encodes non-covalent interaction types, enabling the model to distinguish hydrogen bonds from hydrophobic contacts rather than treating all contacts uniformly. Second, SE(3)-equivariant protein encoding ensures that predictions are invariant to rigid-body transformations of the pocket, improving generalization to complexes crystallized in different space groups. Third, heteroscedastic uncertainty quantification allows practitioners to flag low-confidence predictions during virtual screening triage. Limitations include reliance on crystallographic pocket definitions (reducing applicability to apo structures) and a current inability to model covalent or metalloprotein ligands. We are extending the framework to handle induced-fit effects via docking-guided conformational sampling.

Conclusions

1. GraphBind sets a new state of the art for structure-based binding affinity prediction, with $\mathbf{R} = \mathbf{0.862}$ and $\mathbf{RMSE} = \mathbf{0.83}$ on BindBench-v2.
2. The interface attention module and SE(3)-equivariant protein encoder are both critical, contributing +0.061 and +0.030 Pearson R respectively.
3. GPU-accelerated inference enables million-scale virtual screening in under one day, positioning GraphBind as a practical tool for early-stage hit discovery.
4. Uncertainty estimates correlate with prediction error (Spearman $\rho = 0.41$), providing a calibratable confidence filter for experimental follow-up.

Future Work

- Extend to covalent inhibitors and metalloenzyme targets via reaction-aware edge features.
- Integrate with diffusion-based pocket-conditioned ligand generation (co-design with Vertex Therapeutics generative chemistry platform).
- Active-learning loop connecting GraphBind predictions to automated synthesis and SPR validation at the Nimbus High-Throughput Screening Facility.
- Multi-task extension to simultaneously predict ADMET endpoints alongside affinity.

[1] Chen *et al.* (2025). GraphBind: geometry-aware protein–ligand binding prediction. *Nature Machine Intelligence*, in press.
[2] Vasquez & O'Brien (2024). Equivariant graph networks for molecular property prediction. *NeurIPS* **37**, 12840–12855.
[3] Park *et al.* (2024). PIGNet: physics-informed graph networks for binding affinity. *J. Chem. Inf. Model.* **64**(8), 3120–3134.
[4] Liu *et al.* (2023). PDBbind refined set v2024: updated benchmarks. *J. Med. Chem.* **67**(3), 1402–1417.
[5] Gasteiger *et al.* (2021). DimeNet++: fast and accurate directional message passing. *ICLR* 2021.
[6] Satorras *et al.* (2021). E(n)-equivariant graph neural networks. *ICML* **139**, 9323–9332.

Acknowledgments

Funded by the Synapse Innovation Grant (SIG-2025-0841) and the Meridian Foundation for Computational Drug Discovery. Computational resources provided by the Nimbus AI Compute Cluster. The authors thank the Vertex Therapeutics structural biology team for blind-test validation complexes.

Code & data: `github.com/synapse-inst/graphbind` *License:* Apache 2.0