

AI-Driven Protein–Ligand Interaction Prediction: Emerging Computational Strategies for Drug Development

Chris Kucharik

Professor University of Wisconsin–Madison

Abstract

Protein–ligand interaction prediction has become a central component of modern computational drug discovery, supporting the identification, prioritisation and optimisation of therapeutic candidates. Conventional approaches, including molecular docking, molecular dynamics simulations and structure-based virtual screening, have provided valuable computational frameworks but remain constrained by scoring-function limitations, conformational complexity, computational cost and incomplete representation of biological environments. Recent advances in artificial intelligence have introduced new strategies for predicting protein–ligand binding affinity, interaction poses, molecular compatibility and pharmacological properties. This paper examines emerging AI-driven approaches for protein–ligand interaction prediction, including graph neural networks, transformer architectures, geometric deep learning, generative models, multimodal learning and hybrid physics-informed approaches. A conceptual computational framework is proposed that integrates protein structures, ligand representations, molecular dynamics, experimental data and AI-based prediction. The paper discusses challenges involving data quality, protein flexibility, ligand-induced conformational changes, dataset bias, out-of-distribution prediction, interpretability and experimental validation. Particular attention is given to the integration of AI predictions with molecular docking and molecular dynamics rather than treating machine-learning outputs as independent replacements for established computational chemistry methods. The analysis suggests that future drug-discovery platforms will increasingly rely on hybrid computational architectures combining data-driven learning with physical principles and experimental feedback. Such systems could accelerate virtual screening, lead optimisation and target discovery while reducing computational and experimental costs. However, reliable deployment requires rigorous benchmarking, uncertainty estimation, prospective validation and careful integration with medicinal chemistry expertise.

Keywords: Artificial Intelligence, Protein–Ligand Interaction, Drug Discovery, Molecular Docking, Deep Learning, Graph Neural Networks, Protein Structure, Binding Affinity, Virtual Screening, Computational Chemistry.

1. Introduction

Drug discovery is a complex and resource-intensive process.

A typical discovery pipeline includes:

Target Identification → Hit Discovery → Hit Validation → Lead Optimisation → Preclinical Evaluation → Clinical Development

Computational methods can accelerate several stages of this process.

Protein–ligand interaction prediction is particularly important because many therapeutic effects depend on the ability of a small molecule to bind to a specific biological target.

The fundamental computational question is:

How strongly and in what manner will a ligand interact with a protein?

AI is increasingly being used to address this question.

2. Protein–Ligand Interactions

Proteins contain binding sites capable of interacting with small molecules.

Interactions can involve:

- Hydrogen bonding.
- Electrostatic interactions.
- Hydrophobic interactions.
- π – π interactions.
- Van der Waals forces.
- Metal coordination.

The combined effects determine binding affinity and molecular recognition.

3. Importance in Drug Development

Accurate interaction prediction can support:

- Virtual screening.
- Hit identification.
- Lead optimisation.
- Drug repurposing.
- Target validation.
- Molecular design.

If computational systems can accurately prioritise promising molecules, experimental screening requirements can potentially be reduced.

4. Conventional Computational Approaches

Traditional protein–ligand modelling includes:


1. Molecular docking.
2. Molecular dynamics.
3. Molecular mechanics.
4. Quantum chemical calculations.
5. Pharmacophore modelling.
6. Similarity-based virtual screening.

These methods remain important and increasingly serve as complementary components of AI-driven workflows.

5. Molecular Docking

Molecular docking attempts to identify favourable ligand orientations within a protein binding site.

A simplified workflow is:



The challenge is that docking scores do not always accurately correspond to experimentally measured binding affinities.

6. Limitations of Docking

Docking can be affected by:

- Protein flexibility.
- Ligand flexibility.
- Solvent effects.
- Incomplete structural information.
- Simplified scoring functions.

AI can potentially improve ranking by learning complex relationships from experimental data.

7. Molecular Dynamics

Molecular dynamics simulations model the movement of atoms over time.

They can provide information about:

- Protein flexibility.
- Ligand stability.
- Binding-site dynamics.
- Conformational changes.

However, long simulations can be computationally expensive.

AI may help identify relevant conformational states or accelerate selected components of simulation workflows.

8. AI for Protein–Ligand Prediction

AI systems can learn relationships between:

Protein Representation

And

Ligand Representation

to predict:

- Binding affinity.
- Interaction probability.
- Binding poses.
- Contact maps.
- Molecular compatibility.

9. Protein Representation

Proteins can be represented using:

- Amino-acid sequences.
- 3D structures.
- Contact maps.
- Molecular graphs.
- Surface representations.
- Learned embeddings.

Different representations capture different aspects of molecular information.

10. Ligand Representation

Small molecules can be represented as:


- Molecular graphs.
- SMILES strings.
- Molecular fingerprints.
- 3D coordinates.
- Learned molecular embeddings.

Graph-based representations are particularly useful because atoms and chemical bonds naturally form graph structures.

11. Graph Neural Networks

Graph neural networks can learn molecular representations by passing information between connected atoms.

Conceptually:



GNNs can capture structural relationships that may be difficult to encode using conventional fingerprints.

12. Protein-Ligand Graphs

A joint molecular graph can represent:

- Protein residues.
- Ligand atoms.
- Intramolecular interactions.
- Protein-ligand contacts.

The model can then learn interaction patterns directly from molecular structure.

13. Transformer Architectures

Transformers have become important for biological sequence modelling.

They can process:

- Protein sequences.
- Ligand strings.
- Molecular descriptions.
- Scientific text.

Large protein language models can generate contextual representations of amino-acid sequences that can be incorporated into interaction-prediction models.

14. Protein Language Models

Protein language models learn statistical patterns from large collections of protein sequences.

These representations can encode information related to:

- Protein families.
- Functional regions.
- Evolutionary relationships.
- Structural properties.

Such embeddings can provide useful inputs for protein–ligand prediction.

15. Geometric Deep Learning

Protein–ligand interactions occur in three-dimensional space.


Therefore, models must account for:

- Distances.
- Angles.
- Spatial orientation.
- Molecular symmetry.

Geometric deep learning is designed to incorporate these spatial relationships.

16. 3D Molecular Learning

A three-dimensional model may represent:



This can provide a more physically meaningful representation than purely sequence-based models.

17. Binding Affinity Prediction

One major application is predicting binding affinity.

The objective is to estimate how strongly a ligand binds to a protein.

Experimental measures may include:


- K_d
- K_i
- IC_{50}
- EC_{50}

The precise relationship between these measurements must be interpreted carefully because they represent different experimental concepts.

18. Virtual Screening

AI can evaluate very large chemical libraries.

A conceptual workflow is:



This hierarchical strategy can reduce computational expense.

19. AI-Based Pose Prediction

Beyond affinity, AI models can predict possible ligand binding poses.

The objective is to identify:

- Binding orientation.
- Contact residues.
- Interaction geometry.

Accurate pose prediction is important for medicinal chemistry because structural information can guide molecular optimisation.

20. Protein Flexibility

Proteins are dynamic molecules.

Their structures can change due to:

- Ligand binding.
- Temperature.
- Environmental conditions.
- Intrinsic conformational dynamics.

A model assuming a completely rigid protein may therefore produce inaccurate predictions.

21. Ensemble-Based Prediction

One strategy is to use multiple protein conformations.

For example:

Protein Conformation A

Protein Conformation B

Protein Conformation C

↓

AI Interaction Model

↓

Integrated Prediction

This can better represent protein flexibility.

22. Induced Fit

Ligand binding can cause structural changes in the protein.

This phenomenon is known as induced fit.

AI models capable of considering both ligand and protein flexibility could potentially improve interaction prediction.

23. Solvent Effects

Water molecules can play important roles in molecular recognition.

They may:


- Mediate hydrogen bonds.
- Stabilise interactions.
- Influence binding energetics.

Accurate prediction therefore requires consideration of the molecular environment rather than relying solely on static protein–ligand geometry.

24. Multimodal AI

A promising direction is combining multiple information sources.

For example:



Such models may capture complementary biological and chemical information.

25. Physics-Informed AI

Purely data-driven models may produce chemically unrealistic predictions.


Physics-informed approaches incorporate principles such as:

- Molecular geometry.
- Energy constraints.
- Symmetry.
- Distance relationships.
- Physical interaction rules.

This can improve model plausibility and generalisation.

26. Hybrid AI–Physics Approaches

A powerful architecture may combine:




This approach uses the strengths of both machine learning and computational chemistry.

27. Generative AI for Molecular Design

Generative AI can create new molecular structures.

The workflow can be:



Experimental Testing

This connects interaction prediction with de novo drug design.

28. Lead Optimisation

AI can identify structural modifications that may improve:

- Binding affinity.
- Selectivity.
- Solubility.
- Stability.
- Pharmacokinetic properties.

Medicinal chemists can then evaluate these suggestions.

29. Selectivity Prediction

A drug candidate should ideally bind strongly to the desired target while avoiding undesirable proteins.

AI can compare interaction profiles across multiple targets.

Conceptually:

Target A: High Affinity

Target B: Low Affinity

Target C: Low Affinity

↓

Potentially Improved Selectivity

30. Off-Target Prediction

Off-target interactions can contribute to:

- Toxicity.
- Side effects.
- Unexpected pharmacological activity.

AI-based interaction prediction can potentially identify possible off-target interactions earlier in drug development.

31. Drug Repurposing

Existing drugs can be screened against new targets.

AI can integrate:

- Known molecular structures.
- Protein targets.
- Experimental interaction data.
- Biological pathways.

This may identify candidates for new therapeutic applications.

32. Data Quality

AI performance depends heavily on training data.

Protein–ligand datasets may contain:

- Experimental noise.
- Measurement differences.
- Duplicates.
- Inconsistent protocols.
- Structural errors.

Poor-quality data can produce misleading models.

33. Dataset Bias

A dataset may overrepresent:

- Particular protein families.
- Certain chemical scaffolds.
- Specific therapeutic areas.

A model may therefore perform well on familiar chemistry but poorly on novel targets or molecules.

34. Data Leakage

Data leakage is a major concern.

If highly similar molecules or protein targets occur in both training and test sets, reported performance may be artificially high.

Appropriate splitting strategies are therefore essential.

35. Out-of-Distribution Prediction

Drug discovery frequently involves molecules unlike those used during model training.

A trustworthy AI system should estimate whether a new prediction lies within its validated domain.

This is particularly important for novel chemical scaffolds.

36. Uncertainty Quantification

AI predictions should ideally include uncertainty.

For example:

Predicted affinity: high

Confidence: moderate

Chemical novelty: high

Such information helps researchers prioritise experimental validation.

37. Explainability

Researchers may want to understand why a model predicts a particular interaction.

Useful outputs could include:

- Important residues.
- Important ligand atoms.
- Predicted interaction types.
- Binding-site regions.

Interpretability can help medicinal chemists assess whether predictions are chemically plausible.

38. AI-Generated Interaction Maps

An AI system could potentially produce an interaction map showing:

- Hydrogen bonds.
- Hydrophobic contacts.
- Electrostatic interactions.
- Aromatic interactions.

Such representations can connect AI predictions with conventional structural chemistry.

39. Experimental Validation

Computational predictions should ultimately be tested experimentally.

Possible validation methods include:


- Surface plasmon resonance.
- Isothermal titration calorimetry.
- Biochemical binding assays.
- Enzyme inhibition assays.
- Structural methods.

The strongest computational models are those that demonstrate prospective predictive performance.

40. Active Learning

Active learning can connect AI models with experiments.


The process is:



This can reduce the number of experiments required to improve model performance.

41. Closed-Loop Drug Discovery

An advanced workflow could operate continuously:



This creates a closed-loop drug-discovery system.

42. Proposed Computational Framework

A comprehensive AI-driven protein–ligand prediction framework can contain eight stages:

Stage 1 — Target Preparation

Protein sequence and structural data.

Stage 2 — Ligand Representation

Molecular graph and 3D representation.

Stage 3 — AI Feature Learning

Protein and ligand embeddings.

Stage 4 — Interaction Prediction

Affinity and interaction probability.

Stage 5 — Pose Prediction

Binding geometry estimation.

Stage 6 — Physics-Based Refinement

Docking and molecular simulation.


Stage 7 — Candidate Ranking

Multi-objective optimisation.

Stage 8 — Experimental Validation

Laboratory testing and model updating.

43. Conceptual Architecture



44. Evaluation Metrics

Metric	Purpose
RMSE	Affinity prediction error
MAE	Average prediction error
Pearson correlation	Linear agreement
Spearman correlation	Ranking agreement
ROC-AUC	Classification performance
Precision	Candidate identification
Recall	Detection of relevant interactions

Metric	Purpose
Top-k enrichment	Virtual screening effectiveness
Pose RMSD	Structural pose accuracy
Calibration	Reliability of confidence estimates

45. Research Questions

1. Which AI architectures provide the most reliable protein–ligand interaction predictions?
2. How can protein flexibility be incorporated into deep-learning models?
3. Can geometric neural networks outperform conventional docking approaches for novel targets?
4. How can physics-based constraints improve AI predictions?
5. How should models be evaluated under realistic out-of-distribution conditions?
6. Can multimodal models integrate sequence, structure and experimental data effectively?
7. How can uncertainty estimates guide experimental prioritisation?
8. Can AI accurately predict off-target interactions?
9. How can generative models design molecules with target-specific interaction profiles?
10. What combination of AI and molecular simulation provides the best cost–accuracy trade-off?

46. Future Research Directions

Future research is likely to focus on:

- 3D foundation models for molecular interactions.
- Protein language model integration.
- Equivariant neural networks.
- Physics-informed molecular AI.
- AI-accelerated molecular dynamics.
- Generative molecular design.
- Multitarget interaction prediction.
- Uncertainty-aware virtual screening.
- Autonomous experimental platforms.
- AI–robotic drug-discovery laboratories.

47. Challenges

Despite rapid advances, important challenges remain.

Technical Challenges

- Protein flexibility.
- Limited high-quality data.

- Chemical-space diversity.
- Computational complexity.

Scientific Challenges

- Generalisation to novel targets.
- Experimental reproducibility.
- Biological complexity.

Practical Challenges

- Integration with existing drug-discovery workflows.
- Computational infrastructure.
- Regulatory requirements.

48. Role of Medicinal Chemists

AI should complement rather than eliminate expert judgement.

Medicinal chemists can evaluate:

- Chemical feasibility.
- Synthetic accessibility.
- Binding-site plausibility.
- Selectivity.
- Pharmacological relevance.

This human–AI collaboration can improve the practical value of computational predictions.

49. From Prediction to Decision Support

The ultimate objective is not simply:

Predict binding affinity.

It is:

Identify which molecules should be synthesised and tested next.

This requires combining:

Affinity + Selectivity + ADMET + Synthetic Accessibility + Uncertainty + Biological Context

into an integrated decision-making framework.

50. Conclusion

AI-driven protein–ligand interaction prediction is emerging as an important component of computational drug development. Advances in graph neural networks, transformers, geometric deep learning, protein language models, generative AI and multimodal learning are expanding the ability of computational systems to analyse molecular interactions.

However, AI should not be viewed as a complete replacement for molecular docking, molecular dynamics, physical modelling or laboratory experimentation. The most promising strategy is likely to involve hybrid AI–physics–experiment workflows in which machine

learning rapidly explores chemical space, physics-based methods refine molecular hypotheses and experimental testing establishes biological validity.

Future drug-discovery platforms are likely to become increasingly iterative and autonomous:

Design → Predict → Simulate → Test → Learn → Redesign

The success of these systems will depend not only on predictive accuracy but also on their ability to generalise to novel chemistry, quantify uncertainty, provide chemically meaningful explanations and demonstrate prospective experimental performance.

Ultimately, AI-driven protein–ligand interaction prediction has the potential to transform drug discovery from a largely sequential search process into a data-driven, adaptive and continuously learning molecular design ecosystem, accelerating the identification and optimisation of safer and more effective therapeutic candidates.

References

1. Jumper, J., Evans, R., Pritzel, A., et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596, 583–589.
2. Senior, A. W., Evans, R., Jumper, J., et al. (2020). Improved protein structure prediction using potentials from deep learning. *Nature*, 577, 706–710.
3. Stokes, J. M., Yang, K., Swanson, K., et al. (2020). A deep learning approach to antibiotic discovery. *Cell*, 180(4), 688–702.
4. Zhavoronkov, A., Ivanenkov, Y. A., Aliper, A., et al. (2019). Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, 37, 1038–1040.
5. Gawehn, E., Hiss, J. A., & Schneider, G. (2016). Deep learning in drug discovery. *Molecular Informatics*, 35(1), 3–14.
6. Vamathevan, J., Clark, D., Czodrowski, P., et al. (2019). Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18, 463–477.
7. Stärk, H., Ganea, O.-E., Pattanaik, L., et al. (2022). EquiBind: Geometric deep learning for drug binding structure prediction. *Proceedings of the 39th International Conference on Machine Learning*.
8. Corso, G., Stärk, H., Jing, B., Barzilay, R., & Jaakkola, T. (2023). DiffDock: Diffusion steps, twists, and turns for molecular docking. *International Conference on Learning Representations*.
9. McNutt, J., He, X., & Li, W. (2021). Deep learning methods for molecular property prediction. *Journal of Chemical Information and Modeling*.
10. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. (2017). Neural message passing for quantum chemistry. *Proceedings of the 34th International Conference on Machine Learning*, 1263–1272.

11. Schütt, K. T., Saucedo, H. E., Kindermans, P.-J., Tkatchenko, A., & Müller, K.-R. (2018). SchNet—A deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148, 241722.
12. Fout, A., Byrd, J., Shariat, B., & Ben-Hur, A. (2017). Protein interface prediction using graph convolutional networks. *Advances in Neural Information Processing Systems*, 30.
13. Jiménez, J., Škalič, M., Martínez-Rosell, G., & De Fabritiis, G. (2018). KDEEP: Protein–ligand absolute binding affinity prediction via 3D-convolutional neural networks. *Journal of Chemical Information and Modeling*, 58(2), 287–296.
14. Feinberg, E. N., Joshi, E., Pande, V. S., & Cheng, A. C. (2018). Improvement in ADMET prediction with multitask deep learning. *Journal of Medicinal Chemistry*, 61(24), 10883–10894.
15. Rifaioğlu, A. S., Atas, H., Martin, M. J., et al. (2019). Recent applications of deep learning and machine intelligence on in silico drug discovery: Methods, tools and databases. *Briefings in Bioinformatics*, 20(5), 1878–1900.