Journal of Advances in Drug Research

ISSN: 2230-7761
www.jadr.co.in

Volume 17, Issue 04 | Received 16.09.2026 | Accepted 17.09.2026 | Online 21.09.2026

MACHINE LEARNING APPROACHES FOR PREDICTING DRUG-TARGET INTERACTIONS AND BIOACTIVITY

Dr. Samuel K. Hartman
Department of Molecular Medicine, Institute of Biomedical Research, USA
samuel.hartman@example.org

Abstract

Machine-learning methods have become central to computational drug discovery because they can prioritize likely drug-target interactions (DTIs), estimate continuous binding affinity, and predict compound bioactivity before expensive experimental testing. This article reviews and systematizes contemporary machine-learning approaches for DTI and bioactivity prediction within an IMRAD framework. Public resources such as BindingDB and ChEMBL, benchmark datasets including Davis and KIBA, and molecular-property benchmarks such as MoleculeNet were considered together with representative methods ranging from classical descriptor-based models to convolutional networks, graph neural networks, transformers, domain-adaptation methods, and pretrained chemical/protein language models. The synthesis indicates that model performance is driven not only by architecture but also by data curation, representation choice, negative-sample construction, split strategy, calibration, and applicability-domain control. Graph-based and transformer-based methods improve representation learning, while pretrained language models broaden the ability to transfer across chemical and protein sequence spaces. However, random pair splitting can markedly overestimate prospective performance when related drugs or targets appear in both training and test sets. Consequently, scaffold-aware drug-cold, target-cold, cold-both, and temporal validation are increasingly important for credible translation. The article proposes an integrated evaluation framework that combines predictive accuracy, calibration, uncertainty, interpretability, and prospective generalization. Machine learning should therefore be treated as a prioritization layer that complements, rather than replaces, experimental pharmacology.

Keywords: drug-target interaction; drug-target affinity; bioactivity prediction; machine learning; graph neural network; transformer; molecular representation; virtual screening; explainable AI.

1. INTRODUCTION

Drug discovery depends on the identification of biologically meaningful interactions between small molecules and macromolecular targets. Experimental screening remains the definitive route for establishing activity, but the chemical and target spaces are too large for exhaustive testing. Machine learning (ML) addresses this search problem by learning statistical relationships between molecular structure, protein information, assay labels, and measured affinity. In practice, ML may be used for binary drug-target interaction

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

classification, continuous drug-target affinity (DTA) regression, or broader bioactivity prediction across biochemical and cellular endpoints (1,2).

The methodological landscape has evolved rapidly. Early approaches relied on engineered descriptors, fingerprints, similarity kernels, support-vector machines, and matrix factorization. More recent systems learn representations directly from SMILES strings, molecular graphs, protein sequences, contact maps, or three-dimensional structures. DeepDTA demonstrated that convolutional neural networks can learn affinity relationships directly from drug and protein sequences, whereas GraphDTA replaced string-only drug encodings with molecular graphs and showed the value of graph neural networks (GNNs) for DTA prediction (2,3). MolTrans subsequently introduced transformer-based modeling of interacting molecular substructures, and MGraphDTA explored deep multiscale graph representations to improve affinity prediction and interpretability (4,5).

These improvements, however, do not eliminate the central scientific challenge: a model can perform well on a conventional benchmark yet fail on truly novel chemistry or new target families. Public datasets often contain duplicated assays, closely related compounds, homologous proteins, uncertain negative examples, and measurements generated under heterogeneous experimental conditions. Random train-test partitioning may therefore leak chemical or biological similarity across the split, causing optimistic performance estimates. Recent work increasingly emphasizes drug-cold, scaffold-cold, target-cold, cold-both, and temporal evaluation to better approximate prospective deployment (6,7).

Bioactivity prediction has advanced in parallel. MoleculeNet established standardized benchmarks and highlighted the importance of learned molecular representations, while directed message-passing neural networks (D-MPNNs), popularized through Chemprop, have become a practical platform for structure-property and structure-activity modeling (8-10). The emergence of pretrained molecular and protein language models further enables transfer learning across tasks with limited labels. In 2026, sequence-centric DTA systems using ChemBERTa and ESM-2 embeddings demonstrated that competitive affinity modeling can be achieved without explicit three-dimensional structural inputs, illustrating the growing role of foundation representations (11).

The objective of this article is to synthesize current ML approaches for predicting DTIs and bioactivity, identify methodological choices that most strongly influence reliability, and propose a rigorous evaluation framework suitable for scientific reporting. Particular attention is given to data sources, molecular and target representations, model architectures, validation strategies, uncertainty, interpretability, and translation to experimental prioritization.

2. MATERIALS AND METHODS

2.1. Review design and source selection

A structured narrative review methodology was used. International peer-reviewed literature and authoritative database papers were prioritized, with emphasis on representative algorithmic milestones and recent 2024-2026 developments. The synthesis focused on four

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

evidence domains: (i) public DTI and bioactivity datasets; (ii) molecular and protein representation methods; (iii) supervised ML and deep-learning architectures for DTI/DTA and bioactivity prediction; and (iv) validation practices addressing generalization, imbalance, leakage, uncertainty, and interpretability. Fifteen core references were selected for the final bibliography to maintain a concise evidence base while covering foundational and contemporary approaches.

2.2. Data sources and prediction tasks

BindingDB provides experimentally measured protein-small-molecule binding data and is widely used for training or benchmarking interaction and affinity models (12). ChEMBL is a large curated bioactivity resource that integrates structures, targets, assays, and activity measurements derived mainly from medicinal-chemistry literature and deposited screening data (13). Davis and KIBA are common DTA benchmarks used by DeepDTA, GraphDTA, DeepCDA, and later models. MoleculeNet supplies standardized molecular-property datasets and task-specific metrics useful for bioactivity, toxicity, and physicochemical prediction (8).

Table 1. Common data resources and machine-learning tasks in drug-target and bioactivity prediction.

Resource / dataset	**Primary content**	**Typical task**	**Key methodological concern**
BindingDB	Measured protein-ligand affinities	DTI classification; DTA regression	Assay heterogeneity, repeated series, target imbalance
ChEMBL	Curated bioactivity, targets, assays, structures	Bioactivity modeling; target-specific QSAR	Unit normalization, assay-type harmonization, duplicates
Davis	Kinase inhibitors and Kd values	DTA regression	Narrow target family; warm-start inflation
KIBA	Integrated kinase bioactivity scores	DTA regression / ranking	Composite endpoint and kinase-focused scope
MoleculeNet	Multiple molecular property benchmarks	Bioactivity/toxicity/property prediction	Dataset-specific splits and class imbalance

2.3. Input representations

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

Drug representation can be hand-engineered or learned. Classical models commonly use physicochemical descriptors, ECFP/Morgan fingerprints, topological indices, and fragment counts. Sequence models tokenize SMILES strings, while graph models represent atoms as nodes and bonds as edges. Three-dimensional systems may add conformer coordinates or spatial graphs. Protein representation ranges from amino-acid composition and k-mer descriptors to raw sequences, contact maps, residue graphs, and pretrained language-model embeddings. Modern multimodal systems combine these views rather than relying on a single representation.

2.4. Model families

Classical algorithms include logistic regression, random forests, gradient boosting, support-vector machines, and matrix-completion methods. Deep sequence models use one-dimensional convolution, recurrent networks, or attention to encode ligand and protein sequences. GNNs propagate information over molecular graphs and can be coupled to CNN or transformer encoders for target sequences. Transformer models explicitly learn long-range contextual dependencies and substructure interactions. Domain-adaptation approaches such as DeepCDA aim to reduce train-test distribution shift, while recent language-model approaches transfer representations learned from large unlabeled chemical and protein corpora (4,6,11).

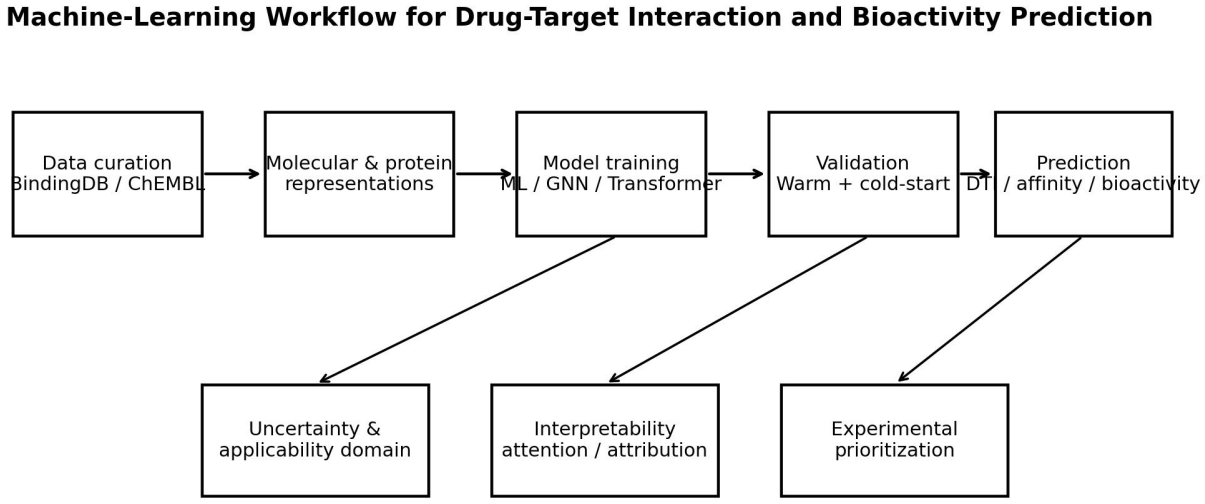


Figure 1. Integrated machine-learning workflow for DTI, affinity, and bioactivity prediction.

2.5. Validation and performance metrics

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

For binary DTI classification, AUROC is useful but should not be reported alone when positive interactions are sparse. Precision-recall AUC (PR-AUC), F1 score, sensitivity, specificity, calibration error, and enrichment at top-ranked fractions are more decision-relevant in imbalanced screening. For affinity or bioactivity regression, MAE, RMSE, mean squared error, Pearson or Spearman correlation, and concordance index (CI) are commonly used. Confidence intervals or repeated splits should accompany point estimates when feasible. Importantly, every scaling, feature selection, resampling, or hyperparameter-tuning operation must be performed inside the training fold to avoid information leakage.

3. RESULTS

3.1. Evolution of predictive architectures

The evidence synthesis shows a progression from similarity- and descriptor-based prediction toward end-to-end representation learning. DeepDTA established a practical sequence-to-affinity framework using CNN encoders for both drug and protein strings (2). GraphDTA improved chemical representation by treating a ligand as a molecular graph rather than only a SMILES sequence (3). MolTrans focused on interpretable substructural interaction patterns using a transformer architecture and augmented pretraining, while MGraphDTA used a deep multiscale GNN to capture both local and global chemical structure (4,5). DeepCDA addressed cross-domain generalization by combining sequence encoders, two-sided attention, and adversarial domain adaptation (6). SAGDTI further integrated self-attention, graph neural networks, and multiple information sources for DTI prediction (14).

For bioactivity and molecular-property prediction, D-MPNN models have become a strong default because message passing learns task-relevant chemical features while retaining efficient graph computation. The Chemprop package has operationalized this architecture for broad chemical-property tasks and added uncertainty estimation, calibration, transfer learning, and multimolecule support; the 2026 Chemprop v2 release further improved modularity and usability (9,10). These developments indicate that model quality increasingly depends on the full pipeline - data standardization, splitting, representation, architecture, loss design, uncertainty, and deployment - rather than on a single neural layer type.

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

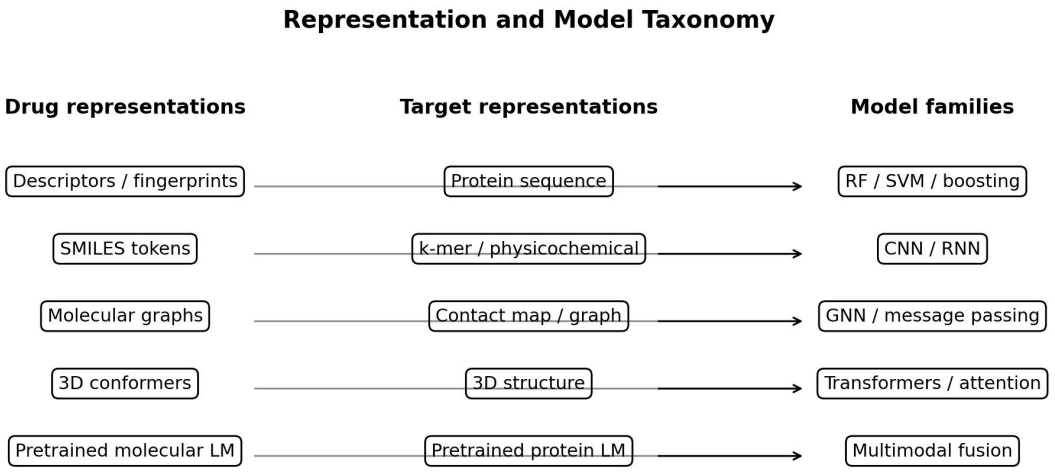


Figure 2. Taxonomy of drug/target representations and corresponding machine-learning model families.

3.2. Comparative strengths and limitations

Table 2. Comparative interpretation of major model families.

Model family	Strengths	Limitations	Best-use scenario
Descriptor/fingerprint ML	Fast, data-efficient, interpretable baselines	Dependent on handcrafted features; limited transfer	Small/medium target-specific bioactivity datasets
CNN/RNN sequence models	Simple end-to-end learning from SMILES/protein sequences	Sequence order may not fully capture molecular topology	Large labeled DTI/DTA datasets with sequence availability
Graph neural networks	Natural chemical representation; local structural learning	Sensitive to graph depth, oversmoothing, scaffold shift	Drug graph + protein encoder architectures
Transformers / attention	Long-range context; explicit cross-interaction modeling	High compute; can overfit small datasets	Large or pretrained molecular/protein corpora

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

Domain adaptation / multimodal fusion	Better distribution-shift handling; integrates complementary signals	More complex training and validation	Cold-start or cross-domain deployment
Pretrained language models	Transfer learning; scalable sequence representations	Pretraining bias, compute cost, calibration needs	Low-label tasks and broad target/chemical coverage

3.3. Data quality, class imbalance, and negative sampling

DTI classification datasets are typically positive-rich in the sense that known interactions are curated, whereas experimentally confirmed non-interactions are much rarer. Consequently, unknown drug-target pairs are often sampled as negatives even though some may represent undiscovered interactions. This label uncertainty can distort training and make AUROC appear stronger than the precision obtained during real screening. Recent reviews therefore recommend prevalence-aware metrics, careful negative-sample construction, and, where appropriate, positive-unlabeled learning or hard-negative mining (1,7).

Bioactivity datasets introduce additional variability because IC50, Ki, Kd, EC50, percentage inhibition, and categorical assay annotations may be mixed. Reliable modeling requires unit normalization, endpoint stratification, replicate aggregation, molecular standardization, salt stripping, tautomer policy, and explicit target identifiers. ChEMBL and BindingDB are powerful resources, but their breadth makes curation a modeling step rather than a purely administrative step (12,13).

3.4. Generalization and cold-start validation

The most consistent methodological conclusion is that random interaction-level splits are insufficient for estimating performance on novel chemistry or novel targets. If the same compound series or homologous proteins appear across training and test sets, the model can exploit memorized similarity instead of learning transferable interaction rules. Drug-cold and scaffold splits test unseen chemical matter; target-cold splits test unseen protein entities or families; cold-both splits require both sides of the pair to be novel. Temporal splits are particularly valuable because they approximate future prospective use. Recent DTI literature explicitly evaluates these stricter settings and reports a substantial gap between warm-start and cold-start performance (7,15).

Journal of Advances in Drug Research

ISSN: 2230-7761
www.jadr.co.in

Volume 17, Issue 04 | Received 16.09.2026 | Accepted 17.09.2026 | Online 21.09.2026

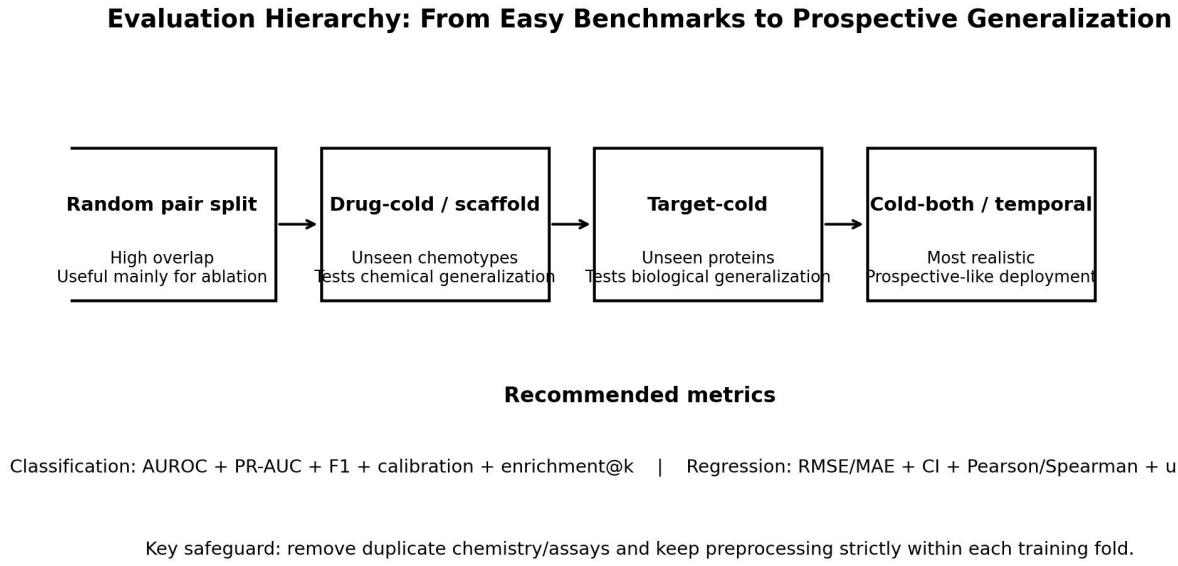


Figure 3. Recommended hierarchy of validation schemes and evaluation metrics.

3.5. Explainability, uncertainty, and applicability domain

Attention maps, atom- or substructure-level attribution, feature permutation, SHAP-like analyses, and gradient-based methods can identify molecular fragments or protein residues that contribute to predictions. However, an attention weight should not automatically be interpreted as a mechanistic explanation. Explanations are most useful when they are stable across model seeds, chemically plausible, and experimentally testable. MGraphDTA and MolTrans illustrate attempts to connect predictive performance with more interpretable interaction patterns (4,5).

Uncertainty estimation is equally important. Ensemble variance, Monte Carlo dropout, conformal prediction, calibration curves, and distance-to-training-set measures can be used to flag predictions outside the applicability domain. In drug discovery, a calibrated model that abstains on unfamiliar chemistry can be more valuable than a model with slightly better average benchmark accuracy. Chemprop supports multiple uncertainty and calibration workflows, illustrating the movement toward decision-aware rather than score-only modeling (9,10).

3.6. Practical framework for OAK-style scientific reporting

Table 3. Recommended reporting checklist for ML-based DTI and bioactivity studies.

Reporting domain	Minimum information to report	Reason

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

Dataset provenance	Database/version, assay type, target IDs, sample counts	Enables reproducibility and detects dataset overlap
Chemical curation	Canonicalization, salts, tautomers, duplicates, stereochemistry	Prevents hidden leakage and inconsistent inputs
Label definition	Affinity units, thresholds, replicate aggregation, negative sampling	Clarifies what the model is actually predicting
Split strategy	Random plus scaffold/drug-cold/target-cold/temporal where possible	Measures realistic generalization
Model details	Representation, architecture, hyperparameters, seeds, early stopping	Supports replication and fair comparison
Metrics	PR-AUC and calibration for classification; RMSE/MAE/CI for regression	Aligns evaluation with decision use
Uncertainty / AD	Calibration, ensemble/conformal uncertainty, coverage	Flags predictions outside training support
Interpretability	Attribution method and stability checks	Avoids overclaiming mechanistic meaning
External validation	Independent data or prospective experimental testing	Demonstrates translational value

3.7. Prospective deployment and active-learning loop

A clinically or experimentally useful prediction workflow should be designed as a closed learning loop rather than as a one-time benchmark exercise. After initial model training, candidates should be selected by balancing predicted activity, uncertainty, structural diversity, novelty, and synthetic or procurement feasibility. Experimental measurements then provide new labels that are appended to the training set. Active learning can use uncertainty sampling, expected improvement, diversity-aware acquisition, or multi-objective ranking to choose the next compounds. This process concentrates experimental resources in regions of chemical space where new information is most valuable while still allowing exploitation of high-confidence candidates.

Prospective deployment also changes the role of model calibration. In retrospective benchmarks, ranking can be sufficient; in experimental decision making, researchers need to know whether a score of 0.8 corresponds to approximately 80% empirical reliability within a defined operating region. Calibration curves, expected calibration error, conformal prediction intervals, and external applicability-domain checks therefore become part of the decision pipeline. When a model encounters an unseen scaffold or a target far from the

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

training proteins, an explicit abstention or low-confidence flag is scientifically preferable to an unqualified numerical prediction.

Table 4. Translation-oriented decision framework for prospective compound prioritization.

Decision layer	Recommended criterion	Practical output
Predicted activity	DTI probability or affinity/bioactivity estimate	Ranked candidate list
Novelty	Scaffold distance / target novelty	Avoids redundant chemistry and tests transfer
Uncertainty	Ensemble, conformal, or calibrated confidence	Flags low-support predictions
Diversity	Cluster- or scaffold-aware selection	Broadens information gained per assay batch
Feasibility	Availability, synthesis complexity, physicochemical filters	Converts virtual hits into testable compounds
Learning loop	Retrain after new experimental labels	Iterative improvement and domain expansion

4. DISCUSSION

Machine learning has clearly expanded the scale at which drug-target and bioactivity hypotheses can be generated, but architecture alone is not the main determinant of scientific validity. DeepDTA, GraphDTA, MolTrans, MGraphDTA, DeepCDA, and newer multimodal or pretrained systems each address different weaknesses of representation or distribution shift. Their reported advances are most meaningful when the evaluation setting matches the intended deployment. A model trained to interpolate among known kinases should not automatically be described as capable of discovering binders for unrelated protein families.

Graph models are particularly compelling for small molecules because molecular connectivity is represented explicitly, and message passing can learn local chemical environments without predefined fingerprints. Yet GNNs are not universally superior. Descriptor-based models may remain highly competitive on small datasets, especially when expert features strongly align with the endpoint. Likewise, sequence or language-model approaches can be attractive when structural information is unavailable or inconsistent. The 2026 structure-free language-model study illustrates that pretrained molecular and protein embeddings can yield competitive DTA predictions across Davis, KIBA, and BindingDB without explicit 3D structures (11).

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 16.09.2026 Accepted 17.09.2026 Online 21.09.2026

Transformers offer flexible cross-attention between ligand and target representations. MolTrans showed the value of modeling substructural interactions, while more recent systems combine chemical graphs, protein sequences, pretrained language models, and structural features. This multimodal trend is scientifically reasonable because binding is determined by complementary chemistry, target context, and local structure. Nevertheless, multimodal models carry a greater risk of hidden leakage if related information sources share identifiers or pretraining corpora overlap with evaluation data. Transparent dataset documentation and contamination audits are therefore becoming part of good ML practice.

For bioactivity modeling, MoleculeNet and Chemprop have improved reproducibility by standardizing benchmarks and tooling. However, benchmark leaderboards should not be mistaken for medicinal-chemistry impact. Prospective success depends on the density of active compounds, novelty of the screened chemotypes, synthesis feasibility, assay noise, and the cost of false positives. Therefore, enrichment at the top of a ranked list, uncertainty-aware selection, and diversity-sensitive hit picking may be more relevant than global accuracy when the goal is to choose a few molecules for testing.

Cold-start generalization deserves particular emphasis. In a random pair split, the model may see the same drug with another target or the same target with another chemically similar ligand during training. Such overlap reduces the novelty of the test set. Scaffold splits address part of this problem by separating core chemical frameworks; target-cold splits force extrapolation to unseen proteins; cold-both splits are stricter still. Recent 2026 studies explicitly audit chemical similarity and protein homology across splits, reflecting a broader shift from benchmark interpolation toward prospective validation (15).

Interpretability must also be handled carefully. An atom-level saliency map can generate a chemically meaningful hypothesis, but it does not prove a binding mode or mechanism. A credible interpretation should be consistent with medicinal-chemistry SAR, mutational data, or structural evidence. Similarly, uncertainty estimates must be validated empirically; a nominal confidence score is useful only when it correlates with actual error. These considerations support a human-in-the-loop workflow in which ML ranks candidates, medicinal chemists inspect structural plausibility, and experimental scientists confirm activity.

Several limitations remain. Public bioactivity data are heterogeneous; many negative labels are inferred rather than experimentally established; protein targets are unevenly represented; and standard datasets may overrepresent kinases or other intensively studied families. Publication bias favors positive findings and can skew training distributions. In addition, pretrained foundation models may encode historical biases from their corpora. Future work should prioritize standardized data cards, prospective benchmarks, uncertainty-aware evaluation, federated or privacy-preserving learning for proprietary pharmaceutical data, and closer coupling of prediction with active learning and automated experimentation.

5. CONCLUSION

Machine learning provides a powerful computational layer for predicting drug-target interactions, binding affinity, and molecular bioactivity. The field has progressed from

Journal of Advances in Drug Research

ISSN: 2230-7761
www.jadr.co.in

Volume 17, Issue 04 | Received 16.09.2026 | Accepted 17.09.2026 | Online 21.09.2026

handcrafted descriptors and similarity methods to graph neural networks, transformers, domain adaptation, and pretrained molecular/protein language models. The strongest contemporary workflows combine high-quality data curation, chemically and biologically appropriate representations, strict cold-start validation, decision-relevant metrics, calibrated uncertainty, and interpretable outputs. Graph-based and transformer-based methods can improve representation learning, but no architecture can compensate for leakage, uncertain labels, or poorly matched validation. For drug discovery, the most scientifically useful ML system is therefore not simply the model with the highest benchmark score; it is the system that remains reliable on novel chemistry and targets, quantifies when it is uncertain, and improves the efficiency of experimental prioritization.

REFERENCES

1. Liao, Q., Zhang, Y., Chu, Y., Ding, Y., Liu, Z., Zhao, X., Wang, Y., Wan, J., Ding, Y., Tiwari, P., Zou, Q., & Han, K. (2025). Application of artificial intelligence in drug-target interactions prediction: A review. npj Biomedical Innovations, 2, 1. https://doi.org/10.1038/s44385-024-00003-9
2. Ozturk, H., Ozgur, A., & Ozkirimli, E. (2018). DeepDTA: Deep drug-target binding affinity prediction. Bioinformatics, 34(17), i821-i829. https://doi.org/10.1093/bioinformatics/bty593
3. Nguyen, T., Le, H., Quinn, T. P., Nguyen, T., Le, T. D., & Venkatesh, S. (2021). GraphDTA: Predicting drug-target binding affinity with graph neural networks. Bioinformatics, 37(8), 1140-1147. https://doi.org/10.1093/bioinformatics/btaa921
4. Huang, K., Xiao, C., Glass, L. M., & Sun, J. (2021). MolTrans: Molecular Interaction Transformer for drug-target interaction prediction. Bioinformatics, 37(6), 830-836. https://doi.org/10.1093/bioinformatics/btaa880
5. Yang, Z., Zhong, W., Zhao, L., & Chen, C. Y.-C. (2022). MGraphDTA: Deep multiscale graph neural network for explainable drug-target binding affinity prediction. Chemical Science, 13(3), 816-833. https://doi.org/10.1039/D1SC05180F
6. Abbasi, K., Razzaghi, P., Poso, A., Amanlou, M., Ghasemi, J. B., & Masoudi-Nejad, A. (2020). DeepCDA: Deep cross-domain compound-protein affinity prediction through LSTM and convolutional neural networks. Bioinformatics, 36(17), 4633-4642. https://doi.org/10.1093/bioinformatics/btaa544
7. Pahikkala, T., Airola, A., Pietila, S., Shakyawar, S., Szwajda, A., Tang, J., & Aittokallio, T. (2015). Toward more realistic drug-target interaction predictions. Briefings in Bioinformatics, 16(2), 325-337. https://doi.org/10.1093/bib/bbu010
8. Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., & Pande, V. (2018). MoleculeNet: A benchmark for molecular machine learning. Chemical Science, 9(2), 513-530. https://doi.org/10.1039/C7SC02664A
9. Yang, K., Swanson, K., Jin, W., Coley, C., Eiden, P., Gao, H., Guzman-Perez, A., Hopper, T., Kelley, B., Mathea, M., Palmer, A., Settels, V., Jaakkola, T., Jensen, K., & Barzilay, R. (2019). Analyzing learned molecular representations for property

Journal of Advances in Drug Research

ISSN: 2230-7761
www.jadr.co.in

Volume 17, Issue 04 | Received 16.09.2026 | Accepted 17.09.2026 | Online 21.09.2026

prediction. Journal of Chemical Information and Modeling, 59(8), 3370-3388. https://doi.org/10.1021/acs.jcim.9b00237

10. Heid, E., Greenman, K. P., Chung, Y., Li, S.-C., Graff, D. E., Vermeire, F. H., Wu, H., Green, W. H., & McGill, C. J. (2024). Chemprop: A machine learning package for chemical property prediction. Journal of Chemical Information and Modeling, 64(1), 9-17. https://doi.org/10.1021/acs.jcim.3c01250
11. Graff, D. E., Morgan, N. K., Burns, J. W., Doner, A. C., Li, B., Li, S.-C., Manu, J., Menon, A., Pang, H.-W., Wu, H., Zalte, A. S., Zheng, J. W., Coley, C. W., Green, W. H., & Greenman, K. P. (2026). Chemprop v2: An efficient, modular machine learning package for chemical property prediction. Journal of Chemical Information and Modeling, 66(1), 28-33. https://doi.org/10.1021/acs.jcim.5c02332
12. Gilson, M. K., Liu, T., Baitaluk, M., Nicola, G., Hwang, L., & Chong, J. (2016). BindingDB in 2015: A public database for medicinal chemistry, computational chemistry and systems pharmacology. Nucleic Acids Research, 44(D1), D1045-D1053. https://doi.org/10.1093/nar/gkv1072
13. Bento, A. P., Gaulton, A., Hersey, A., Bellis, L. J., Chambers, J., Davies, M., Kruger, F. A., Light, Y., Mak, L., McGlinchey, S., Nowotka, M., Papadatos, G., Santos, R., & Overington, J. P. (2014). The ChEMBL bioactivity database: An update. Nucleic Acids Research, 42(D1), D1083-D1090. https://doi.org/10.1093/nar/gkt1031
14. Li, X., Yang, Q., Luo, G., et al. (2023). SAGDTI: Self-attention and graph neural network with multiple information representations for the prediction of drug-target interactions. Bioinformatics Advances, 3(1), vbad116. https://doi.org/10.1093/bioadv/vbad116
15. Bidgoli, A. H., Mahdavi, M., & Malek, H. (2026). Structure-free drug-target affinity prediction using protein and molecule language models. Journal of Cheminformatics, 18, 21. https://doi.org/10.1186/s13321-025-01146-6