Journal of Advances in Drug Research

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

Volume 17, Issue 04 | Received 13.09.2026 | Accepted 15.09.2026 | Online 20.09.2026

ARTIFICIAL INTELLIGENCE–ASSISTED DRUG DISCOVERY: RECENT ADVANCES IN LEAD IDENTIFICATION AND OPTIMIZATION

Dr. Adrian Keller
Department of Medicinal Chemistry and Computational Drug Design, European Institute of Pharmaceutical Sciences, Berlin, Germany
adrian.keller@example.com

ABSTRACT

Artificial intelligence (AI) has moved from a largely exploratory technology in pharmaceutical research to a practical set of methods for molecular representation, virtual screening, de novo design, structure prediction, property forecasting and multi-parameter lead optimization. The present study provides a structured narrative review of recent advances in AI-assisted small-molecule drug discovery, with particular emphasis on the transition from hit identification to experimentally actionable lead series. Fifteen peer-reviewed international publications from 2018–2026 were selected to represent major methodological and translational milestones. The reviewed evidence shows that graph neural networks, self-supervised molecular representations, chemical language models, reinforcement learning, deep quantitative structure–activity relationship models and structure-prediction systems can improve prioritization across large chemical spaces and support iterative medicinal chemistry decisions. Prospective examples include the rapid generation of DDR1 kinase inhibitors and deep-learning-guided identification of the structurally distinct antibiotic halicin. AlphaFold2 and AlphaFold3 further expanded the structural information available for target assessment and biomolecular interaction modeling. However, benchmark performance does not automatically translate into clinical value. Data quality, dataset shift, uncertainty calibration, interpretability, synthetic feasibility, experimental validation and the choice of biologically meaningful endpoints remain decisive constraints. The most defensible near-term model is therefore not autonomous drug discovery, but a closed-loop human–AI workflow in which algorithms prioritize experiments, new assay data update the models and medicinal chemists retain control over multi-objective trade-offs. This framework can shorten design–make–test–analyze cycles while preserving scientific accountability and translational relevance.

Keywords: artificial intelligence; drug discovery; machine learning; deep learning; virtual screening; lead optimization; ADMET; generative models.

1. INTRODUCTION

Drug discovery is an iterative decision problem in which enormous chemical and biological search spaces must be reduced to a small number of molecules that combine potency, selectivity, developability and acceptable safety. Conventional computational chemistry has long supported this process, but the growing volume of chemical structures, omics data, phenotypic measurements, imaging data and protein information has created conditions in which modern machine-learning methods can extract relationships that are

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

difficult to encode manually. AI is therefore increasingly used as a decision-support layer across target identification, hit finding, molecular design and preclinical candidate selection (1,2).

The term AI in drug discovery covers a heterogeneous family of approaches rather than a single technology. Supervised learning can map molecular representations to measured activities or properties; graph neural networks (GNNs) learn directly from atomic connectivity; self-supervised and contrastive learning exploit large unlabeled molecular collections; natural-language-inspired models treat SMILES strings or protein sequences as symbolic languages; generative models propose new structures; reinforcement learning searches for molecules that satisfy reward functions; and structure-prediction systems use deep learning to infer three-dimensional biomolecular organization. Their common purpose is to improve prioritization under uncertainty rather than to replace experimental pharmacology or medicinal chemistry (1–3).

A major historical barrier has been representation. Traditional descriptors and fingerprints compress chemical structures into engineered variables, whereas neural networks can learn task-specific representations from graphs, strings and three-dimensional geometry. MoleculeNet helped formalize comparative evaluation by collecting diverse molecular machine-learning tasks and demonstrating the importance of learned representations while also showing persistent weaknesses in data-scarce and imbalanced settings (3). Subsequent GNN and geometry-aware approaches have further improved the capacity to encode local chemistry and spatial context (9,10).

Prospective demonstrations have strengthened the case for AI-assisted hit finding. Zhavoronkov et al. reported a generative reinforcement-learning platform that produced potent DDR1 kinase inhibitors within a compressed design cycle, with biochemical, cellular and pharmacokinetic validation (4). Stokes et al. trained a deep neural network for antibacterial activity and identified halicin, a compound structurally distinct from established antibiotic classes, followed by experimental validation in bacterial models (5). These studies are important because they move beyond retrospective benchmark accuracy and test whether an algorithm can influence real experimental choices.

At the same time, critical analyses have emphasized that high predictive accuracy on curated datasets does not guarantee useful medicines. Drug discovery data are sparse, biased by historical project choices, affected by assay noise and frequently non-stationary when chemical series or biological contexts change. The clinically relevant objectives—human efficacy, safety, exposure and therapeutic index—are often weakly represented in early-stage datasets. For this reason, AI performance must be judged by prospective decision quality, uncertainty calibration and experimentally verified outcomes rather than by leaderboard metrics alone (6,7).

The structural biology landscape has also changed. AlphaFold2 achieved near-experimental accuracy for many protein structures and dramatically expanded the availability of structural hypotheses for target assessment and structure-based design (8). AlphaFold3 extended this concept to complexes that can include proteins, nucleic acids, small molecules, ions and modified residues, reporting improved performance across several biomolecular interaction categories (14). These systems do not eliminate the need for crystallography, cryo-electron microscopy, biophysical assays or molecular simulation,

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

but they alter the starting point for many discovery programs by reducing the number of targets for which structural modeling begins from little or no information.

The objective of this article is to synthesize recent international evidence on AI-assisted drug discovery with a focus on lead identification and optimization. The analysis emphasizes practical integration: how AI methods can help select targets and hits, generate and rank molecules, optimize multiple properties and interact with experimental design. Particular attention is given to the limits that determine whether algorithmic advances translate into robust pharmaceutical decisions.

2. MATERIALS AND METHODS

A structured narrative review was performed to map major methodological and translational advances in AI-assisted drug discovery. Literature screening was completed in September 2026 using PubMed and publisher journal platforms. Search combinations included “artificial intelligence drug discovery”, “machine learning lead optimization”, “graph neural network molecular property prediction”, “generative molecular design”, “deep learning virtual screening”, “AI ADMET”, “AlphaFold drug discovery”, and “deep QSAR”.

Eligibility criteria were: (i) peer-reviewed international publication; (ii) direct relevance to target assessment, molecular representation, hit identification, de novo design, property prediction, structure modeling or lead optimization; (iii) methodological influence, prospective experimental validation, or high relevance to translational limitations; and (iv) publication in English. Fifteen publications from 2018–2026 were selected to provide methodological breadth rather than to support a formal meta-analysis. Foundational benchmark work was retained when it remained essential for interpreting newer models.

Evidence was extracted into five analytical categories: molecular/target representation, virtual screening and hit discovery, generative molecular design, lead optimization and ADMET prediction, and structural AI. For prospective case studies, preference was given to reports containing experimental follow-up. Methodological claims were interpreted conservatively: retrospective benchmark improvements were not treated as equivalent to prospective medicinal-chemistry performance, and structure predictions were not treated as substitutes for experimental validation. The synthesis was qualitative because the included publications used heterogeneous endpoints, datasets and validation designs that do not support a single pooled effect size.

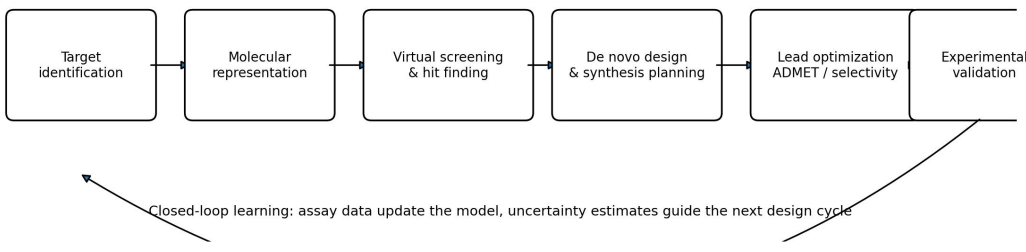
Table 1. Principal AI approaches used across the drug-discovery workflow

Approach	Typical input	Primary role	Main advantage	Key limitation	Refs.
Supervised ML / deep QSAR	Descriptors, fingerprints, graphs	Activity, affinity, ADMET and property prediction	Fast ranking of analogues; scalable prioritization	Strong dependence on assay quality and applicability domain	(1,13)

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

Approach	Typical input	Primary role	Main advantage	Key limitation	Refs.
Graph neural networks	Molecular graphs / 3D geometry	Property prediction, DTI, virtual screening	Learns local chemical environments and transferable representations	Data scarcity, oversmoothing, weak extrapolation to novel chemistry	(9,10)
Chemical language models	SMILES / reaction strings	Representation learning, molecule generation, reaction prediction	Can exploit very large unlabeled corpora	Validity and biological relevance require external constraints	(2,11)
Generative + reinforcement learning	Latent vectors, strings, graphs	De novo design and multi-objective optimization	Direct search for molecules satisfying several objectives	Reward hacking; synthetic infeasibility; uncertain novelty	(4,11)
Explainable AI	Model attributions / saliency / counterfactuals	Model audit and medicinal-chemistry interpretation	Supports hypothesis generation and error detection	Explanations can be unstable or method-dependent	(6)
Structure-prediction AI	Protein / complex sequences and structures	Target assessment and structure-enabled design	Expands structural coverage; informs binding hypotheses	Static predictions and confidence limitations; experimental confirmation needed	(8,14)

AI-assisted discovery is most useful as an iterative decision-support loop



Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

Figure 1. Closed-loop architecture for AI-assisted hit identification and lead optimization.

3. RESULTS

3.1. Molecular representation and predictive learning

The reviewed literature indicates that representation learning is a central technical advance because the usefulness of any downstream model is bounded by the information encoded in its input. MoleculeNet showed that learned molecular representations can outperform fixed representations on many benchmark tasks, but also documented clear limitations in small, imbalanced or physics-sensitive datasets (3). Modern GNNs address this problem by treating atoms as nodes and bonds as edges, allowing neighborhood information to be propagated through molecular graphs. Contrastive learning can reduce dependence on labeled data by training encoders to recognize related molecular views before task-specific fine-tuning (9).

Three-dimensional representation adds another layer of information. Geometry-enhanced learning combines graph topology with distances, angles or spatial relations, enabling models to capture conformational cues that are absent from two-dimensional connectivity alone. Fang et al. demonstrated that geometry-aware pretraining improved molecular property prediction across benchmark tasks (10). For medicinal chemistry, the practical implication is not that a single architecture is universally superior, but that representation should be matched to the causal structure of the endpoint. Solubility and permeability may often be modeled from ligand-level information, whereas binding affinity and selectivity can depend strongly on three-dimensional protein–ligand context.

Deep QSAR represents the convergence of classical structure–activity modeling with neural representation learning. Tropsha et al. argue that the field has expanded from descriptor-based regression toward deep generative design, synthetic planning and learning-enhanced structure-based screening (13). Yet the fundamental principles of QSAR remain: chemical curation, rigorous separation of training and test chemistry, applicability-domain analysis and prospective validation are still required. More complex models do not remove the risk of information leakage or activity cliffs.

3.2. AI-enabled virtual screening and prospective hit identification

Virtual screening is an attractive application for AI because computational ranking can be applied to libraries that are far larger than the number of compounds that can be tested experimentally. Neural models can learn activity patterns from historical assays and prioritize molecules for purchase or synthesis. Their greatest value appears when they complement, rather than simply reproduce, similarity search or docking. A useful model should enrich experimentally active chemistry while maintaining structural diversity and uncertainty awareness.

The DDR1 study by Zhavoronkov et al. is a frequently cited prospective example. The authors used generative tensorial reinforcement learning (GENTRL) to generate compounds optimized for novelty, synthetic feasibility and predicted activity; several molecules were synthesized, four were active in biochemical assays, two were validated in cell-based assays, and one lead showed favorable pharmacokinetics in mice (4). The significance of this study lies less in the exact development timeline than in the demonstration of a complete computational-to-experimental loop.

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

Stokes et al. used a deep neural network trained to predict antibacterial activity and screened large chemical collections in silico. The model identified halicin from a repurposing library, and subsequent experiments showed broad antibacterial activity, including activity in murine infection models. The same workflow was used to search more than 100 million structures in ZINC15, from which a small number of diverse predictions were experimentally tested and additional antibacterial hits were found (5). This example demonstrates how AI can search for non-obvious chemical similarity patterns and thereby increase the probability of discovering scaffolds outside familiar medicinal-chemistry neighborhoods.

Nevertheless, prospective hit finding remains difficult. Biased negative data, inactive compounds tested under different assay conditions, missing counter-screens and inconsistent measurement protocols can lead to overly optimistic models. A model that ranks millions of molecules can also generate a large number of false positives even when its average discrimination metric appears strong. Therefore, enrichment, precision at experimentally affordable cutoffs, diversity, novelty, calibration and hit confirmation should be reported together.

Table 2. Selected landmark evidence for AI-assisted drug discovery

Study / year	AI task	System or data	Main evidence	Validation level	Ref.
MoleculeNet benchmark (2018)	Benchmarking molecular ML	Multiple public molecular datasets	Learned representations performed strongly on many tasks, while data scarcity and imbalance remained important constraints	Retrospective benchmark	(3)
GENTRL / DDR1 (2019)	De novo lead generation	DDR1 kinase program	Generated and synthesized novel inhibitors; biochemical, cellular and mouse PK follow-up reported	Prospective experimental	(4)
Halicin (2020)	Antibiotic hit identification	Drug Repurposing Hub + ZINC15	Deep learning prioritized structurally unusual antibacterial chemistry followed by in vitro and in vivo validation	Prospective experimental	(5)

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

Study / year	AI task	System or data	Main evidence	Validation level	Ref.
AlphaFold2 (2021)	Protein structure prediction	CASP14 and proteome-scale modeling	Near-experimental accuracy for many targets; major expansion of structural hypotheses	Prospective blind benchmark / structural	(8)
AlphaFold3 (2024)	Biomolecular complex prediction	Protein, nucleic acid, small-molecule and ion complexes	Unified model reported improved accuracy for several interaction classes, including protein–ligand complexes	Structural benchmark	(14)
Lead-optimization review (2026)	Translational synthesis	Recent AI lead-optimization literature	Emphasizes trustworthy AI, uncertainty, data-centric practice and scientific decision quality	Critical review	(15)

3.3. Generative AI and de novo molecular design

Generative AI changes the role of computation from ranking existing molecules to proposing new ones. Chemical language models, variational architectures, graph generators, diffusion-based methods and reinforcement-learning systems can sample regions of chemical space conditioned on predicted objectives. The practical attraction is clear: medicinal chemists often face a multi-dimensional design problem in which potency must improve without sacrificing selectivity, solubility, permeability, metabolic stability, safety or synthetic accessibility (2,11,15).

Reinforcement learning formalizes this as reward optimization. Mokaya et al. examined curriculum and deep reinforcement-learning strategies for SMILES-based design and highlighted both the capability and the limits of these approaches (11). A central problem is that the reward function can become a proxy that the model exploits in chemically unrealistic ways. If synthetic accessibility, assay uncertainty or off-target effects are not encoded, the algorithm can optimize the measurable surrogate rather than the medicinal goal.

For this reason, successful generative design increasingly uses constraints and filters: valid valence and chemistry rules, similarity or novelty bounds, synthetic planning, uncertainty penalties, target-engagement models and experimentally calibrated property predictors. Human review remains important because medicinal chemistry includes qualitative considerations—reactive motifs, unexplained structure–activity relationships, route robustness, formulation risks and project-specific liabilities—that may not be represented in training data.

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

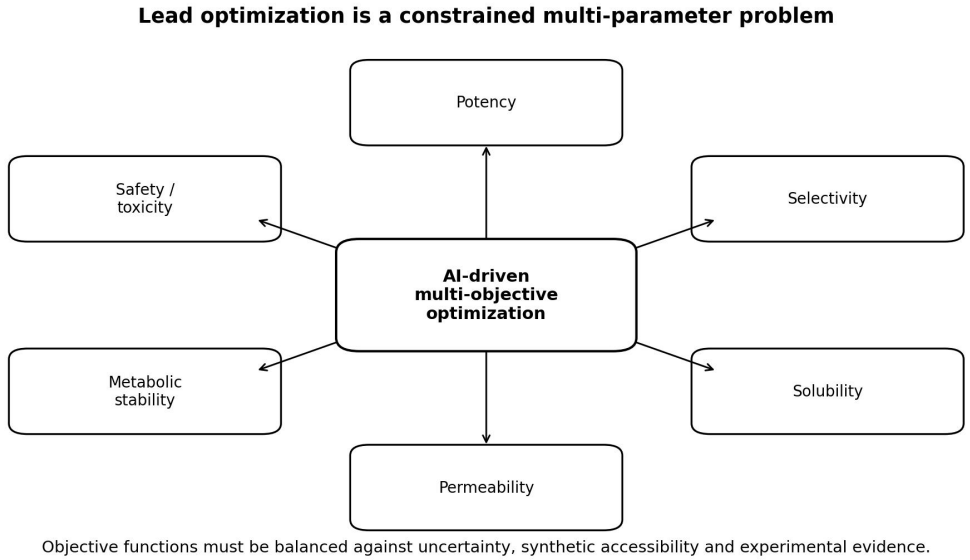


Figure 2. Multi-objective nature of AI-assisted lead optimization.

3.4. Lead optimization, ADMET prediction and scientific decision support

Lead optimization is where AI has the potential to deliver high practical value because each design cycle requires simultaneous reasoning over multiple properties. A candidate with excellent biochemical potency may fail because of poor solubility, insufficient cellular exposure, rapid clearance, CYP inhibition, hERG liability, reactive metabolites or weak selectivity. AI models can estimate many of these endpoints before synthesis, enabling compounds with obviously unfavorable profiles to be deprioritized and allowing chemists to compare trade-offs earlier (1,12,15).

The strongest conceptual shift is from single-endpoint prediction to multi-objective decision support. Instead of maximizing a single potency score, modern workflows can combine predicted activity with developability and synthetic constraints. Pareto optimization, reinforcement learning and Bayesian approaches can identify compounds that represent different compromises rather than forcing all objectives into one weighted score. This is useful because the relative importance of endpoints changes across programs and development stages.

Uncertainty is essential. A predicted logD or clearance value should not be interpreted in the same way when a molecule lies close to the training distribution and when it represents a novel scaffold. Confidence intervals, ensembles, conformal methods and out-of-distribution detection can help distinguish reliable predictions from extrapolation. The 2026 review by Mohamed and El-Sagheir places particular emphasis on trustworthy AI, uncertainty quantification, data-centric modeling and the distinction between benchmark optimization and real scientific decision-making (15).

Explainability can also improve model use, though it should not be overinterpreted. Feature attribution, atom-level importance maps, matched molecular-pair analysis and counterfactual examples can show why a model changes its prediction after a structural modification. Jiménez-Luna et al. noted that explainability is particularly relevant in drug

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

discovery because domain experts must understand model behavior well enough to identify spurious correlations and generate testable hypotheses (6). Explanations are best treated as model diagnostics rather than as proof of molecular mechanism.

Table 3. Practical validation framework for AI-assisted lead optimization

Objective	AI strategy	Recommended validation	Common failure mode
Potency / affinity	QSAR, GNN, protein–ligand models	Random split plus scaffold/time split; prospective assay	High accuracy may collapse on new scaffolds
Selectivity	Multi-task models; target panels	External off-target panel; negative controls	Missing negative data and assay inconsistency
ADME	Property-specific regression/classification	Independent experimental PK/ADME set; uncertainty calibration	Protocol dependence and interspecies translation
Toxicity	Multi-task safety models	Orthogonal assays and mechanistic follow-up	Rare events, class imbalance, label ambiguity
Synthetic accessibility	Reaction models; route planners	Chemist review and real synthesis attempts	A numerically easy synthesis may be operationally poor
Overall lead quality	Multi-objective / Pareto / RL	Closed-loop design–make–test–analyze cycles	Reward function may omit clinically important liabilities

3.5. Structural AI and the expansion of druggable structural space

Protein structure is a key source of mechanistic information for target validation and structure-based drug design. AlphaFold2 represented a major advance by achieving high accuracy in CASP14 and making large-scale structural prediction feasible for proteins lacking experimentally solved structures (8). For drug discovery, these predictions can support hypothesis generation, pocket assessment, mutagenesis planning and prioritization of constructs for structural experiments.

AlphaFold3 broadened the modeling scope to biomolecular complexes containing proteins, nucleic acids, small molecules, ions and modified residues. The reported benchmarks showed substantial improvements for several classes of interactions, including protein–ligand systems, relative to specialized baselines (14). This makes structural AI increasingly relevant to early target assessment and to programs in which experimentally determined complex structures are unavailable.

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 | Received 13.09.2026 | Accepted 15.09.2026 | Online 20.09.2026

However, a predicted complex is not identical to an experimentally verified binding mode or binding free energy. Drug targets are dynamic, protonation states vary, water networks can be decisive, induced fit may occur and alternative conformations can control selectivity. Structural AI should therefore be integrated with docking, molecular simulation, biophysical measurements and experimental structural biology. Confidence scores and model limitations should be included in the decision record, particularly when a predicted pose drives expensive chemistry.

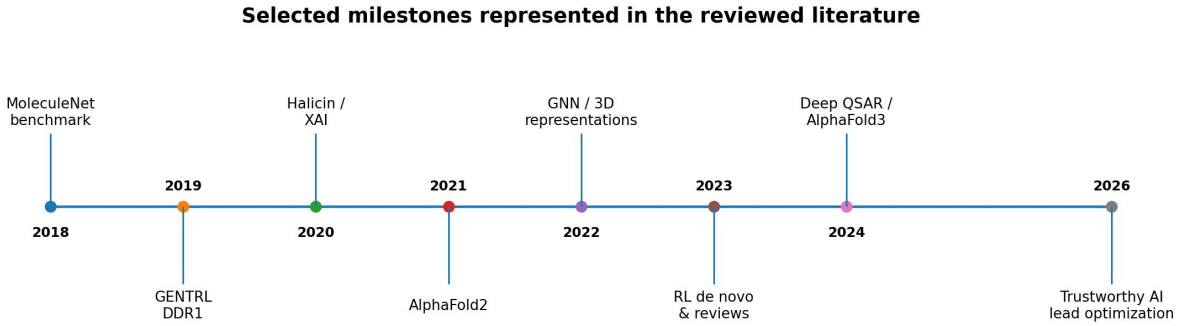


Figure 3. Selected milestones represented in the 2018–2026 review corpus.

4. DISCUSSION

The reviewed evidence supports a balanced conclusion: AI is already useful in drug discovery, but its value is task-dependent and strongly conditioned by data and experimental integration. The most mature applications are those in which the prediction target is measurable, reasonably abundant and closely related to the decision being made—for example, ranking analogues for a biochemical assay, predicting selected physicochemical properties, or learning structure–activity relationships within a defined chemical domain. The least reliable applications are those that attempt to infer distant clinical outcomes from weak early-stage proxies.

A recurring challenge is dataset shift. Medicinal chemistry projects intentionally move away from known molecules in order to discover better chemistry. This creates a paradox: the compounds of greatest interest are often the ones least similar to the data on which the model was trained. Random train–test splits can obscure this problem because near-neighbor analogues may appear in both sets. Scaffold splits, chronological splits and truly prospective evaluation give a more realistic estimate of whether a model can support future synthesis decisions (3,7,13).

Data quality is equally important. Biological labels are not abstract numbers; they arise from assay protocols, biological systems, concentrations, detection technologies and experimental error. Merging measurements without respecting context can create apparent scale while degrading scientific meaning. For lead optimization, a smaller internally consistent dataset may outperform a much larger heterogeneous dataset. This is why data-centric AI, assay metadata and curation remain central despite rapid algorithmic innovation (7,15).

Journal of Advances in Drug Research
ISSN: 2230-7761
www.jadr.co.in
Volume 17, Issue 04 Received 13.09.2026 Accepted 15.09.2026 Online 20.09.2026

The prospective examples in DDR1 and antibiotic discovery show what stronger evidence looks like: model output influenced compound selection, molecules were synthesized or acquired, biological activity was measured, and follow-up experiments tested the original hypothesis (4,5). Future AI papers would benefit from reporting the full funnel—number of virtual candidates, selection rules, number synthesized, assay success rate, chemical novelty, negative results and uncertainty—because these quantities determine operational usefulness.

Generative AI deserves particular caution. It can produce very large numbers of chemically valid structures, but novelty alone is not valuable. A useful generated molecule must be synthesizable, patentable where relevant, active in the correct biological system, selective, stable and compatible with the therapeutic context. Optimization against imperfect predictive models can amplify their errors. Closed-loop experimentation is therefore not an optional final step; it is the mechanism that keeps the generative system grounded in reality (2,11,15).

Explainability and governance are also becoming practical requirements. Medicinal chemists need to understand when model suggestions are supported by prior chemistry and when they are extrapolations. Reproducible pipelines, versioned datasets, documented training procedures and audit trails are important for scientific quality even before formal regulatory use. Explainable AI may help identify spurious correlations, but explanation methods themselves require validation because different techniques can provide different narratives for the same prediction (6).

Structural AI changes the availability of hypotheses rather than abolishing experimental structural biology. AlphaFold2 and AlphaFold3 can accelerate the generation of plausible models, especially for targets lacking structures, yet dynamic binding events, allostery and thermodynamic quantities still require complementary approaches (8,14). In practice, the best workflow is hybrid: structure prediction proposes a starting model; docking or physics-based simulation explores poses and dynamics; medicinal chemistry evaluates chemical plausibility; and experiments confirm binding and activity.

Taken together, the field is shifting from an algorithm-centered view to a decision-centered view. The relevant question is no longer whether a neural network can outperform another model on a benchmark, but whether the overall workflow produces better compounds with fewer unnecessary experiments while preserving the ability to detect failure. This perspective is consistent with the recent emphasis on trustworthy AI and scientific decision quality in lead optimization (15).

4.1. Limitations of the present review

This article is a structured narrative review rather than a systematic review or meta-analysis. The selected 15 references were chosen to represent methodological and translational milestones, and therefore do not capture every AI architecture, proprietary industrial result or disease-specific application. Publication bias is also likely because successful prospective examples are more visible than negative programs. In addition, many industrial AI systems are not described with enough experimental detail to permit independent comparison. Consequently, the synthesis should be interpreted as a framework for understanding current capabilities and constraints, not as a quantitative estimate of the effect of AI on pharmaceutical productivity.

Journal of Advances in Drug Research

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

Volume 17, Issue 04 | Received 13.09.2026 | Accepted 15.09.2026 | Online 20.09.2026

5. CONCLUSION

AI-assisted drug discovery has progressed from proof-of-concept modeling toward integrated use in target assessment, virtual screening, molecular generation, property prediction and lead optimization. The reviewed evidence shows clear technical gains in molecular representation, prospective hit prioritization and structural modeling. Landmark studies such as GENTRL-based DDR1 design, deep-learning-guided discovery of halicin, AlphaFold2 and AlphaFold3 illustrate the breadth of these advances (4,5,8,14).

The central translational lesson is that AI should be embedded in an iterative design–make–test–analyze cycle. Reliable application requires curated data, realistic validation splits, explicit uncertainty, experimentally meaningful endpoints and prospective confirmation. Generative and multi-objective systems are most useful when they operate within medicinal-chemistry and synthetic constraints rather than as unconstrained molecule generators. As of 2026, the strongest model for pharmaceutical use is therefore human-guided, trustworthy and experimentally closed-loop AI. Continued progress is likely to depend as much on data quality, assay design and decision science as on larger or more complex neural architectures (13,15).

DECLARATIONS

Funding: No external funding was specified for this manuscript.

Conflict of interest: The author should declare any relevant financial or non-financial conflicts before submission.

Data availability: No new experimental dataset was generated; all scientific claims are based on the cited literature.

Ethics statement: Not applicable; this study is a literature review and did not involve human participants or experimental animals.

REFERENCES

1. Vamathevan, J., Clark, D., Czodrowski, P., et al. Applications of machine learning in drug discovery and development. Nature Reviews Drug Discovery. 2019;18:463–477. https://doi.org/10.1038/s41573-019-0024-5
2. Schneider, P., Walters, W. P., Plowright, A. T., et al. Rethinking drug design in the artificial intelligence era. Nature Reviews Drug Discovery. 2020;19:353–364. https://doi.org/10.1038/s41573-019-0050-3
3. Wu, Z., Ramsundar, B., Feinberg, E. N., et al. MoleculeNet: a benchmark for molecular machine learning. Chemical Science. 2018;9:513–530. https://doi.org/10.1039/C7SC02664A
4. Zhavoronkov, A., Ivanenkov, Y. A., Aliper, A., et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. Nature Biotechnology. 2019;37:1038–1040. https://doi.org/10.1038/s41587-019-0224-x
5. Stokes, J. M., Yang, K., Swanson, K., et al. A deep learning approach to antibiotic discovery. Cell. 2020;180(4):688–702.e13. https://doi.org/10.1016/j.cell.2020.01.021
6. Jiménez-Luna, J., Grisoni, F., & Schneider, G. Drug discovery with explainable artificial intelligence. Nature Machine Intelligence. 2020;2:573–584. https://doi.org/10.1038/s42256-020-00236-4

Journal of Advances in Drug Research

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

Volume 17, Issue 04 | Received 13.09.2026 | Accepted 15.09.2026 | Online 20.09.2026

7. Bender, A., & Cortés-Ciriano, I. Artificial intelligence in drug discovery: what is realistic, what are illusions? Part 1: Ways to make an impact, and why we are not there yet. Drug Discovery Today. 2021;26(2):511–524. https://doi.org/10.1016/j.drudis.2020.12.009
8. Jumper, J., Evans, R., Pritzel, A., et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596:583–589. https://doi.org/10.1038/s41586-021-03819-2
9. Wang, Y., Wang, J., Cao, Z., et al. Molecular contrastive learning of representations via graph neural networks. Nature Machine Intelligence. 2022;4:279–287. https://doi.org/10.1038/s42256-022-00447-x
10. Fang, X., Liu, L., Lei, J., et al. Geometry-enhanced molecular representation learning for property prediction. Nature Machine Intelligence. 2022;4:127–134. https://doi.org/10.1038/s42256-021-00438-4
11. Mokaya, M., Imrie, F., van Hoorn, W. P., et al. Testing the limits of SMILES-based de novo molecular generation with curriculum and deep reinforcement learning. Nature Machine Intelligence. 2023;5:386–394. https://doi.org/10.1038/s42256-023-00636-2
12. Li, H., Zou, L., Kowah, J. A. H., et al. A compact review of progress and prospects of deep learning in drug discovery. Journal of Molecular Modeling. 2023;29(4):117. https://doi.org/10.1007/s00894-023-05492-w
13. Tropsha, A., Isayev, O., Varnek, A., Schneider, G., & Cherkasov, A. Integrating QSAR modelling and deep learning in drug discovery: the emergence of deep QSAR. Nature Reviews Drug Discovery. 2024;23:141–155. https://doi.org/10.1038/s41573-023-00832-0
14. Abramson, J., Adler, J., Dunger, J., et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature. 2024;630:493–500. https://doi.org/10.1038/s41586-024-07487-w
15. Mohamed, F. A. A. A., & El-Sagheir, A. M. K. Artificial intelligence-assisted lead optimization in drug discovery: bridging computational advances and translational challenges. Medicinal Chemistry Research. 2026. Online ahead of print. https://doi.org/10.1007/s00044-026-03616-4