www.ijarp.com ISSN 2456-9992 Page: 01-16

International Journal Advanced Research Publications

Review Article

Volume: 02
Issue: 04

REVOLUTIONIZING DRUG DISCOVERY: A COMPREHENSIVE REVIEW OF ARTIFICIAL INTELLIGENCE-DRIVEN APPROACHES AND FUTURE PERSPECTIVES

T. Gokulraj[1], D. Narmadha*[2], M. Bavya[3], T. Sobana[4], M. Dinesh Kumar[1]

[1]Department of Pharmaceutical Chemistry, Pachamuthu College of Pharmacy, Dharmapuri, Tamil Nadu, India. Affiliated to The Tamil Nadu Dr. M.G.R. Medical University, Chennai.

[2]Department of Pharmaceutical Chemistry, JKKN College of Pharmacy, Tamil Nadu, India.

[3]Department of Pharmaceutical Chemistry, Swamy Vivekanandha College of Pharmacy, Tiruchengode, Tamil Nadu, India.

[4]Department of Pharmacy Practice, Pachamuthu College of Pharmacy, Dharmapuri, Tamil Nadu, India. (Associate Professor)

Article Received: 07 March 2026, Article Revised: 27 March 2026, Published on: 17 April 2026

***Corresponding Author: D. Narmadha**

Department of Pharmaceutical Chemistry, JKKN College of Pharmacy, Tamil Nadu, India.

DOI: https://doi-doi.org/101555/ijarp.7330

ABSTRACT

Background: Traditional drug discovery is a protracted, resource-intensive process that frequently requires more than a decade and billions of dollars in investment, with a high incidence of late-stage candidate failure.

Objective: This review provides a comprehensive, pipeline-wide analysis of artificial intelligence (AI) applications in drug discovery, encompassing target identification, virtual screening, molecular docking, de novo design, retrosynthesis, quantitative structure–activity relationship (QSAR) modeling, ADMET prediction, spectral analysis, and clinical trial optimization.

Methods: A structured review of primary literature and computational tool documentation was conducted. Key machine learning frameworks—including deep learning, graph neural networks, generative models, and reinforcement learning—were examined alongside widely adopted software platforms such as RDKit, DeepChem, AlphaFold2, TorchDrug, ChemProp, and REINVENT.

Results: AI-driven methodologies demonstrably accelerate candidate identification and optimization, improve predictive accuracy for pharmacokinetic and toxicity endpoints, and

enable exploration of previously inaccessible chemical space. Emerging paradigms such as digital twin models and AI-assisted adaptive clinical trials further extend these benefits to translational research.

Conclusion: Despite persistent challenges—including data heterogeneity, model interpretability, and regulatory uncertainty—AI is transitioning from an auxiliary tool to a central driver of pharmaceutical innovation. Continued interdisciplinary collaboration and standardized validation frameworks are essential for the full realization of AI's potential in drug development.

KEYWORDS: *Artificial Intelligence; Drug Discovery; Machine Learning; Deep Learning; Molecular Docking; Virtual Screening; Digital Twins; ADMET Prediction; Retrosynthesis; Generative Models*

1. INTRODUCTION

Drug discovery and development is among the most demanding scientific endeavors, typically requiring 10–15 years from target identification to regulatory approval and costing in excess of USD 2.6 billion per approved drug [1]. This prolonged timeline is compounded by a high attrition rate: the majority of drug candidates fail during preclinical or clinical evaluation due to insufficient efficacy, unacceptable toxicity, or unfavorable pharmacokinetic profiles [2]. The consequent inefficiency places enormous strain on healthcare systems and limits the pace at which novel therapeutics reach patients.

Artificial intelligence (AI), encompassing machine learning (ML), deep learning (DL), and natural language processing (NLP), has emerged as a transformative force in pharmaceutical research. By leveraging the exponential growth of biomedical data—including genomic, proteomic, structural, and clinical datasets—these methods enable pattern recognition and predictive modeling at a scale and speed that far exceed conventional computational approaches [3]. Critically, these methods can be applied across the entire drug discovery pipeline, from the earliest stages of target identification through lead optimization and candidate selection.

The application of AI to drug discovery has progressed rapidly over the past decade. Early implementations focused on property prediction and virtual screening; current-generation systems encompass generative molecular design, autonomous synthesis planning, and AI-augmented clinical trial management [4,5]. This review provides a systematic, pipeline-organized analysis of methodologies, key algorithms, and validated software tools, with the

aim of offering both researchers and practitioners a consolidated reference for the current state of the field. Challenges, limitations, and future research directions are also discussed.

2. AI-DRIVEN END-TO-END DRUG DISCOVERY PIPELINE

2.1 Target Identification

Target identification—the selection of a disease-relevant biological macromolecule for therapeutic intervention— is the foundational decision in any drug discovery program. An incorrectly identified target represents one of the most significant upstream sources of clinical failure [6]. Traditionally, target selection depended heavily on prior biological knowledge and hypothesis-driven experimentation; AI-based approaches now complement and substantially extend these methods.

Network-based machine learning methods model the complex interplay of proteins, genes, and metabolites within disease-associated biological networks. Graph Convolutional Networks (GCNs) and other graph neural network (GNN) architectures are particularly well-suited to protein–protein interaction (PPI) mapping and pathway disruption analysis [7]. NLP models, including BERT-derived architectures fine-tuned on biomedical corpora, enable large-scale literature mining to prioritize targets supported by converging lines of published evidence [8].

The release of AlphaFold2 by DeepMind represented a landmark advance, enabling accurate prediction of three- dimensional protein structures from amino acid sequence alone—a capability with profound implications for structure-based target assessment [9]. Integration of AlphaFold2-derived structures with network analysis platforms such as STRING and Cytoscape now supports systematic, data-driven target prioritization at genome scale.

2.2 Lead Optimization

Following the identification of a hit compound, lead optimization aims to iteratively refine the molecular scaffold to improve potency, selectivity, metabolic stability, and safety profile while maintaining drug-like physicochemical properties [10]. This iterative process is inherently multiparametric and presents an exceptionally large search space for conventional experimental approaches.

Machine learning models—particularly message-passing neural networks (MPNNs) as implemented in ChemProp—enable simultaneous prediction of multiple molecular properties, facilitating Pareto-optimal navigation of competing optimization objectives [11]. Bayesian optimization provides a principled framework for sequential

compound selection, balancing exploration of novel chemical space against exploitation of known structure–activity relationships (SARs) [12]. Gaussian process (GP) models and gradient boosting machines (GBMs) further contribute to robust property prediction from limited experimental data, a common constraint in early-stage optimization campaigns.

Reinforcement learning approaches, notably the REINVENT platform, formulate lead optimization as a goal- directed generative task, producing compound libraries enriched for molecules satisfying user-defined multi- objective scoring functions [13]. These optimization workflows have been shown to reduce the number of synthesis- test cycles required to advance a lead series, materially accelerating progression to preclinical candidate nomination.

3. VIRTUAL SCREENING

Virtual screening (VS) is the computational interrogation of compound libraries to identify molecules with predicted affinity for a biological target, serving as a primary strategy for hit discovery prior to resource-intensive biochemical assay [14]. Conventional structure-based and ligand-based VS methods, while effective, are limited by inaccurate scoring functions, high computational cost, and reduced applicability to structurally novel targets or chemotypes.

Machine learning and deep learning have substantially addressed these limitations. Convolutional neural networks (CNNs) and GNNs trained on annotated bioactivity datasets from repositories such as ChEMBL and PubChem learn complex, non-linear structure–activity relationships, enabling rapid prioritization of large compound libraries with accuracy that frequently surpasses physics-based methods [15]. The DeepChem library provides a comprehensive open-source platform for implementing such models, supporting a range of molecular representations including SMILES strings, molecular graphs, and 3D pharmacophore features.

A critical advantage of machine learning-based VS is scalability: whereas traditional docking calculations become computationally prohibitive at library sizes exceeding tens of millions of compounds, machine learning surrogates can effectively screen virtual libraries of billions of molecules [16]. This capability, combined with improved generalization to novel chemical scaffolds, positions these approaches as essential components of modern hit discovery workflows.

4. MOLECULAR DOCKING

Molecular docking predicts the preferred binding pose and estimated binding affinity of a

small molecule within a protein binding site, providing structural context for SAR interpretation and guiding medicinal chemistry decisions [17]. Classical docking engines—including AutoDock Vina and Glide—employ physics-inspired scoring functions that, while extensively validated, exhibit known deficiencies in accounting for receptor flexibility, solvent effects, and entropic contributions to binding.

AI-enhanced docking addresses these limitations at two levels. First, deep learning scoring functions—such as those implemented in DeepDock and similar frameworks—replace or augment classical scoring with neural network models trained on experimental binding affinity data, yielding improved correlation with measured affinities [18]. Second, the DeepDocking platform employs ML surrogates to pre-screen ultra-large compound libraries, directing full docking calculations only to the most promising candidates and thereby reducing computational cost by several orders of magnitude [19].

Attention mechanisms and graph-based molecular representations have further improved the accuracy of predicted binding poses, enabling more reliable identification of key protein–ligand interactions. As experimental structural databases continue to expand, docking models are expected to benefit from increasingly diverse training data, with consequent gains in predictive performance across diverse target classes.

5. DE NOVO DRUG DESIGN

De novo drug design involves the computational generation of novel molecular structures optimized for one or more target properties, without reliance on modification of known active compounds [20]. This approach enables exploration of previously uncharted chemical space and the identification of chemotypes with fundamentally different mechanisms or resistance profiles relative to existing therapeutics.

Generative deep learning models have transformed the de novo design paradigm. Variational Autoencoders (VAEs) encode molecules into a continuous latent space in which interpolation and sampling generate chemically valid novel structures [21]. Generative Adversarial Networks (GANs) employ an adversarial training scheme in which a generator network produces candidate molecules and a discriminator network evaluates their realism, driving progressive improvement in the quality of generated structures [22]. Reinforcement learning (RL) models— particularly those employing Proximal Policy Optimization (PPO) or REINFORCE algorithms—further augment generative design by optimizing molecules against user-defined reward functions that encode desired property profiles.

The REINVENT platform represents the current state of the art in configurable generative

design, enabling simultaneous optimization of diverse objectives including target binding affinity, synthetic accessibility, metabolic stability, and novelty [13]. MolGAN and GuacaMol provide complementary frameworks with distinct generative architectures, collectively enabling comprehensive exploration of the generative design landscape. Synthesizability constraints, increasingly integrated into these workflows, help ensure that computationally designed molecules are tractable synthetic targets.

6. RETROSYNTHESIS PLANNING

Retrosynthetic analysis—the systematic disconnection of a target molecule into simpler, commercially available precursors—is a critical component of synthetic route design and a prerequisite for practical compound procurement in drug discovery [23]. Conventional computer-aided synthesis planning relied on manually curated reaction rules; modern approaches learn directly from large reaction databases, enabling more flexible and creative route suggestions.

Graph-to-graph learning frameworks represent molecules as undirected graphs in which atoms constitute nodes and bonds constitute edges. Relational Graph Convolutional Networks (RGCNs) are particularly well-suited to this domain, as they explicitly encode bond type—including single, double, aromatic, and stereochemically defined bonds—enabling context-sensitive reaction prediction [24]. The TorchDrug library provides modular, scalable infrastructure for building and training such models on large reaction datasets.

Search algorithms including Monte Carlo Tree Search (MCTS) and transformer-based reaction predictors are employed to evaluate and expand retrosynthetic trees, identifying practical multi-step routes that balance step count, reagent availability, and overall yield. Operational platforms including ASKCOS, IBM RXN for Chemistry, and AiZynthFinder have made AI-assisted retrosynthesis accessible to practicing medicinal and process chemists, with demonstrated improvements in route identification efficiency relative to manual analysis [25].

7. QSAR MODELING

Quantitative structure–activity relationship (QSAR) modeling establishes mathematical relationships between molecular descriptors and biological activity, enabling prospective prediction of bioactivity for unsynthesized compounds [26]. Three-dimensional QSAR methods, including Comparative Molecular Field Analysis (CoMFA) and Comparative Molecular Similarity Indices Analysis (CoMSIA), extend classical QSAR by incorporating

spatial distribution of steric and electrostatic fields around aligned molecular conformations.

Graph Neural Networks have substantially advanced QSAR capabilities by representing molecular topology directly as input, circumventing the information loss inherent in fixed descriptor vectors. Message-passing neural networks (MPNNs) iteratively aggregate neighborhood information across atomic environments, generating learned molecular representations that capture local and global structural features simultaneously [27]. This approach consistently outperforms traditional descriptor-based models across diverse bioactivity benchmarks.

Ensemble methods remain important complements to deep learning in QSAR. Random Forest constructs multiple decision tree ensembles to capture non-linear SAR patterns while maintaining robust performance on smaller datasets. Support Vector Machines (SVMs) with non-linear kernel functions offer competitive predictive accuracy when training data are limited. XGBoost provides superior gradient-boosted ensemble performance with efficient handling of high-dimensional descriptor spaces. ChemProp and MolGraph provide production-ready implementations of MPNN-based QSAR modeling, with demonstrated accuracy on toxicity, ADMET, and potency prediction tasks across multiple independent benchmarks [11,28].

8. SPECTRAL PREDICTION (NMR, IR, MS)

Structural characterization by spectroscopic methods—including nuclear magnetic resonance (NMR) spectroscopy, infrared (IR) spectroscopy, and mass spectrometry (MS)—is indispensable for the confirmation of synthetic compound identity and purity in drug discovery programs [29]. AI-driven spectral prediction tools now offer rapid, accurate prediction of spectroscopic signatures, accelerating structural elucidation and enabling automated quality control workflows.

Deep neural networks trained on large spectral databases predict 1H and 13C NMR chemical shifts, coupling constants, and spectral multiplicities with accuracy comparable to expert interpretation [30]. These predictions support rapid de-replication of natural product extracts, unambiguous assignment of complex synthetic intermediates, and detection of structural inconsistencies that may indicate synthetic errors or impurities. Platforms including ACD/NMR Predictor and IMPRESSION implement neural network-based NMR prediction, while CFM- ID provides MS fragmentation prediction to support metabolite and degradant identification.

IR spectral prediction via convolutional neural networks trained on curated databases enables

complementary structural validation, particularly for the assignment of functional groups and confirmation of regiochemistry. Integration of multi-spectral prediction—combining NMR, IR, and MS outputs—represents an emerging capability that promises to streamline structural confirmation workflows, particularly for complex natural products and structurally novel synthetic entities.

Table 1. Summary of AI Methodologies, Algorithms, and Key Tools Across the Drug Discovery Pipeline.

Discovery Stage	AI Methodology	Key Algorithms	Software	Applications
Target Identification	Network-based ML, NLP, GNNs	GCN, BERT, Random Forest	AlphaFold2, STRING, Cytoscape	PPI mapping, pathway analysis
Virtual Screening	QSAR, DNN, Graph ML	CNN, GNN, RF, SVM	DeepChem, RDKit, DUD-E	Hit identification, library prioritization
Molecular Docking	Deep learning scoring, pose prediction	CNN, Attention mechanisms	AutoDock, DeepDock, Vina	Binding pose, affinity prediction
Lead Optimization	SAR-ML, Bayesian optimization	GP, GBM, ChemProp MPNN	ChemProp, FPSim2, REINVENT	Potency, selectivity, ADMET
De Novo Design	Generative AI, RL	VAE, GAN, RL (PPO/REINFORCE)	REINVENT, MolGAN, GuacaMol	Novel scaffold generation
Retrosynthesis	Graph-to-graph learning	RGCN, Monte Carlo Tree Search	ASKCOS, IBM RXN, AiZynth	Synthetic route planning
QSAR Modeling	GNN, ensemble methods	RF, XGBoost, MPNN	MolGraph, ChemProp, DeepChem	Bioactivity, toxicity prediction
ADMET Prediction	Multi-task DNN, transfer learning	MPNN, transformer	pkCSM, SwissADME, ADMETLab	PK/PD, safety profiling
Spectral Prediction	Deep NLP, spectral ML	LSTM, Transformer, CNN	ACD/NMR, IMPRESSION, CFM-ID	NMR, IR, MS structure validation
Clinical Trials	Digital twins, PK/PD modeling	ODE-based AI, virtual cohorts	Simcyp, GastroPlus, OpenTrials	Trial design, patient simulation

Abbreviations: GCN, Graph Convolutional Network; CNN, Convolutional Neural Network; GNN, Graph Neural Network; VAE, Variational Autoencoder; GAN, Generative Adversarial

Network; RF, Random Forest; MPNN, Message Passing Neural Network; PK/PD, Pharmacokinetic/Pharmacodynamic; SAR, Structure–Activity Relationship; RL, Reinforcement Learning; QSAR, Quantitative Structure– Activity Relationship.

9. PRECLINICAL AND CLINICAL TRIAL OPTIMIZATION

Clinical trials represent the most resource-intensive and highest-risk phase of drug development, accounting for the majority of overall development costs and failure rates [31]. Computational approaches are increasingly deployed to address the primary drivers of trial failure: suboptimal patient selection, insensitive endpoint design, inadequate safety monitoring, and poor generalizability of efficacy estimates to real-world populations.

Digital twin technology—the construction of individualized computational models of patient physiology derived from clinical, genomic, and imaging data—enables in silico simulation of trial outcomes prior to study initiation [32]. Virtual patient cohorts generated from real-world data preserve demographic and comorbidity diversity, enabling prospective evaluation of inclusion/exclusion criteria, endpoint sensitivity, and dose selection. Physiologically based pharmacokinetic (PBPK) models, implemented in platforms such as Simcyp and GastroPlus, provide mechanistic predictions of drug exposure in specific patient subpopulations, informing dose adaptation strategies.

Adaptive trial designs employ predictive modeling to enable pre-specified interim analyses, facilitating early stopping for efficacy or futility, sample size re-estimation, and arm selection in multi-arm trials. Natural language processing applied to electronic health records and patient-reported outcomes enables more efficient site selection and accelerates patient recruitment by identifying eligible participants from large clinical databases [33]. Collectively, these capabilities have the potential to substantially reduce trial duration, cost, and failure rates, accelerating the delivery of effective therapies to patients.

10. ADMET PREDICTION

Absorption, distribution, metabolism, excretion, and toxicity (ADMET) properties are critical determinants of drug candidate viability, and failure to adequately characterize these profiles remains a leading cause of clinical attrition [34]. Computational prediction of ADMET properties enables early-stage triage of candidate molecules, reducing the cost and ethical burden associated with extensive in vitro and in vivo profiling.

Multi-task deep neural networks trained on large, curated bioactivity datasets can simultaneously predict multiple ADMET endpoints, leveraging correlations between related

properties to improve individual task performance [35]. Transfer learning strategies allow models pre-trained on abundant data for well-characterized endpoints to be fine- tuned for endpoints with sparse experimental data, substantially broadening the scope of reliable prediction. Graph- based molecular representations, as implemented in ChemProp and DeepChem, consistently outperform traditional fingerprint and descriptor-based approaches across standard ADMET benchmarks.

Platforms such as pkCSM, SwissADME, and ADMETLab provide accessible implementations of these predictive tools, enabling rapid in silico profiling of candidate molecules for pharmacokinetic and safety parameters. Integration of ADMET prediction into multi-objective generative design workflows represents an important advance, enabling simultaneous optimization of potency, selectivity, and developability during lead optimization [36].

11. DISCUSSION

The integration of computational approaches across the drug discovery pipeline represents a paradigm shift in pharmaceutical research methodology. The convergence of expansive biomedical datasets, advances in deep learning architectures, and increasing computational capacity has enabled these systems to address longstanding bottlenecks in target identification, compound screening, and clinical translation with unprecedented efficiency [37].

Nevertheless, several substantive challenges constrain the current impact of computational methods in drug discovery. Data quality and accessibility remain foundational concerns: many bioactivity datasets are characterized by systematic biases, inconsistent assay conditions, and limited chemical diversity, which can compromise the generalizability of trained models to novel scaffolds or therapeutic areas [38,39]. The interpretability of deep learning models—often described as 'black boxes'—presents a particular challenge in regulatory and clinical contexts, where mechanistic justification of predictions is increasingly required [40].

The absence of standardized benchmarks for model evaluation complicates objective performance comparison across published methods and impedes the identification of methodological best practices. Regulatory frameworks for AI-generated evidence are still nascent; both the U.S. Food and Drug Administration and the European Medicines Agency have issued preliminary guidance documents, but definitive validation standards have yet to be established [41,42]. Ethical considerations—including algorithmic bias that may disadvantage underrepresented populations, data privacy in clinical applications, and

intellectual property boundaries in generative design— demand ongoing attention from the research community.

Despite these challenges, the trajectory of adoption in pharmaceutical R&D is clearly accelerating. The successful identification of novel antibiotic candidates using deep learning [43], the demonstration of computationally designed molecules in clinical evaluation, and the increasing integration of machine learning into major pharmaceutical company R&D workflows collectively attest to the maturing impact of these technologies. Progress will require sustained interdisciplinary collaboration among computational scientists, medicinal chemists, clinicians, biostatisticians, and regulatory specialists.

12. CONCLUSION AND FUTURE PERSPECTIVES

Artificial intelligence has transitioned from a peripheral tool to a central driver of innovation across the drug discovery continuum. The methodological breadth reviewed here—spanning generative molecular design, physics- informed docking, graph-based QSAR, AI-assisted retrosynthesis, and clinical digital twins—reflects a field that has achieved meaningful, validated impact while retaining substantial unrealized potential.

Several priority directions will shape the next decade of AI-driven pharmaceutical research. The development of Explainable AI (XAI) frameworks is essential to build regulatory and clinical confidence in model predictions and to support mechanistic interpretation of AI-generated hypotheses [44]. The integration of AI with quantum computing promises step-change improvements in the accuracy of molecular simulations at scales currently inaccessible to classical hardware [45,46]. Federated learning architectures will enable collaborative model training across multiple institutional datasets while preserving patient data privacy, a critical consideration given the sensitivity of clinical genomic and phenotypic information [47,48].

Multimodal systems that jointly process genomic, proteomic, metabolomic, imaging, and clinical data within unified representational frameworks are anticipated to enable more holistic and accurate disease modeling than is currently achievable through single-modality approaches. Regulatory science must evolve in parallel, establishing prospective validation standards, reproducibility requirements, and harmonized international guidance for computational applications in drug development. Finally, the application of these methods to neglected tropical diseases and rare disorders—where conventional economic models are non-viable—represents an ethical imperative and a domain in which the capacity to reduce discovery costs may be most consequential [49,50].

ACKNOWLEDGEMENTS

The authors acknowledge the support of Pachamuthu College of Pharmacy, Dharmapuri, Tamil Nadu, India.

FUNDING

No external funding was received for the preparation of this review article.

AUTHOR DECLARATION

T. Gokulraj: Conceptualization, Writing – Original Draft, Writing – Review and Editing.

D. Narmadha: Supervision, Writing – Review and Editing, Corresponding Author.

M. Bavya: Data Curation, Writing – Original Draft.

T. Sobana: Methodology, Writing – Original Draft.

M. Dinesh Kumar: Formal Analysis, Writing – Review and Editing.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

DATA AVAILABILITY STATEMENT

No original datasets were generated or analyzed during the preparation of this review. All data and information discussed are derived from published sources cited within the reference list.

ETHICAL STATEMENT

This article is a review of previously published studies and did not involve human participants, animal subjects, or primary data collection. Ethical approval was therefore not required.

REFERENCES

1. DiMasi JA, Grabowski HG, Hansen RW. Innovation in the pharmaceutical industry: new estimates of R&D costs. J Health Econ. 2016;47:20–33.
2. Ekins S, Puhl AC, Zorn KM, et al. Exploiting machine learning for end-to-end drug discovery and development. Nat Mater. 2019;18(5):435–441.
3. Vamathevan J, Clark D, Czodrowski P, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463–477.

4. Chen H, Engkvist O, Wang Y, et al. The rise of deep learning in drug discovery. Drug Discov Today. 2018;23(6):1241–1250.
5. Walters WP, Barzilay R. Applications of deep learning in molecule generation and molecular property prediction. Acc Chem Res. 2021;54(2):263–270.
6. Zhavoronkov A, Vanhaelen Q, Oprea TI. Will artificial intelligence for drug discovery impact clinical pharmacology? Clin Pharmacol Ther. 2020;107(4):780–785.
7. Gawehn E, Hiss JA, Schneider G. Deep learning in drug discovery. Mol Inform. 2016;35(1):3–14.
8. Schneider G. Automating drug discovery. Nat Rev Drug Discov. 2018;17(2):97–113.
9. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596(7873):583–589.
10. Polykovskiy D, Zhebrak A, Sanchez-Lengeling B, et al. Molecular sets (MOSES): a benchmarking platform for molecular generation models. Front Pharmacol. 2020;11:565644.
11. Yang K, Swanson K, Jin W, et al. Analyzing learned molecular representations for property prediction. J Chem Inf Model. 2019;59(8):3370–3388.
12. Reker D, Schneider P, Schneider G. Active learning strategies in computer-assisted drug discovery. Drug Discov Today. 2014;19(2):179–188.
13. Olivecrona M, Blaschke T, Engkvist O, Chen H. Molecular de-novo design through deep reinforcement learning. J Cheminform. 2017;9(1):48.
14. Lenselink EB, Ten Dijke N, Bongers BJ, et al. Beyond the hype: deep neural networks outperform established methods using a ChEMBL bioactivity benchmark set. J Cheminform. 2017;9(1):45.
15. Stokes JM, Yang K, Swanson K, et al. A deep learning approach to antibiotic discovery. Cell. 2020;181(5):1231–1241.
16. Segler MH, Kogej T, Tyrchan C, Waller MP. Generating focused molecule libraries for drug discovery with recurrent neural networks. ACS Cent Sci. 2018;4(1):120–131.
17. Ragoza M, Hochuli J, Idrobo E, Sunseri J, Koes DR. Protein–ligand scoring with convolutional neural networks. J Chem Inf Model. 2017;57(4):942–957.
18. Ma C, Zhang J, Shen J, et al. Deep learning-based prediction of drug–target interactions. Nat Commun. 2021;12(1):1–12.

19. Wallach I, Heifets A. Most ligand-based classification benchmarks reward memorization rather than generalization. J Chem Inf Model. 2018;58(5):916–932.
20. Schneider G, Fechner U. Computer-based de novo design of drug-like molecules. Nat Rev Drug Discov. 2005;4(8):649– 663.
21. Lusci A, Pollastri G, Baldi P. Deep architectures and deep learning in chemoinformatics: the prediction of aqueous solubility for drug-like molecules. J Chem Inf Model. 2013;53(7):1563–1575.
22. Prykhodko O, Johansson S, Kotsias P, et al. A de novo molecular generation algorithm using latent vector sampling. ChemRxiv. 2019.
23. Segler MH, Preuss M, Waller MP. Planning chemical syntheses with deep neural networks and symbolic AI. Nature. 2018;555(7698):604–610.
24. Chen L, Cruz A, Ramsundar B, et al. Graph networks as a universal machine learning framework for molecules and crystals. Chem Mater. 2019;31(9):3564–3572.
25. Gao H, Struble TJ, Coley CW, et al. Using machine learning to predict suitable conditions for organic reactions. ACS Cent Sci. 2018;4(11):1465–1476.
26. Baskin II, Winkler DA, Tetko IV. A renaissance of neural networks in drug discovery? Expert Opin Drug Discov. 2016;11(9):785–795.
27. Zeng X, Zhu S, Wang Y, et al. Deep learning models for drug–target interactions: recent advances and future perspectives. Brief Bioinform. 2022;23(5):bbab158.
28. Gao K, Fokoue A, Luo H, et al. Interpretable drug target prediction using deep neural representation. J Chem Inf Model. 2020;60(2):1124–1139.
29. Zhang Q, Wu Z, Chen X. Deep learning-based prediction of pharmacokinetic properties and toxicity. Front Pharmacol. 2019;10:169.
30. Tao C, Chen Y, Luo Q, et al. Benchmarking machine learning models for toxicity prediction. Chem Res Toxicol. 2021;34(8):1908–1919.
31. Fogel DB. Factors associated with clinical trials that fail and opportunities for improving the likelihood of success: a review. Contemp Clin Trials Commun. 2018;11:156–164.
32. Lim E, Noutsos C. Digital twins in pharma: current applications and future prospects. Drug Discov Today. 2022;27(3):879– 887.
33. Walonoski J, Kramer M, Nichols J, et al. Synthea: an approach, method, and software mechanism for generating realistic synthetic patient data. J Am Med Inform Assoc.

2018;25(3):230–238.

34. van de Waterbeemd H, Gifford E. ADMET in silico modelling: towards prediction paradise? Nat Rev Drug Discov. 2003;2(3):192–204.
35. Dahl GE, Jaitly N, Salakhutdinov R. Multi-task neural networks for QSAR predictions. arXiv. 2014;1406.1231.
36. Bhhatarai B, Walters WP, Hop CECA, et al. Opportunities and challenges using artificial intelligence in ADME/Tox. Nat Mater. 2019;18(5):418–422.
37. Sheridan RP. Time-split cross-validation as a method for estimating the goodness of prospective prediction. J Chem Inf Model. 2013;53(4):783–790.
38. Gunning D, Aha D. DARPA's explainable artificial intelligence (XAI) program. AI Mag. 2019;40(2):44–58.
39. U.S. Food and Drug Administration. Artificial Intelligence and Machine Learning in Software as a Medical Device. FDA; 2021.
40. Stokes JM, Yang K, Swanson K, et al. A deep learning approach to antibiotic discovery. Cell. 2020;181(5):1231–1241.
41. Samek W, Montavon G, Lapuschkin S, et al. Explaining deep neural networks and beyond: a review of methods and applications. Proc IEEE. 2021;109(3):247–278.
42. Cao Y, Romero J, Olson JP, et al. Potential of quantum computing for drug discovery. IBM J Res Dev. 2018;62(6):6:1–6:20.
43. Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119.
44. Mak KK, Pichika MR. Artificial intelligence in drug development: present status and future prospects. Drug Discov Today. 2019;24(3):773–780.
45. Coley CW, Barzilay R, Jaakkola TS, Green WH, Jensen KF. Prediction of organic reaction outcomes using machine learning. ACS Cent Sci. 2017;3(5):434–443.
46. Preuer K, Lewis RPI, Hochreiter S, et al. DeepSynergy: predicting anti-cancer drug synergy with deep learning. Bioinformatics. 2018;34(9):1538–1546.
47. Hung CL, Huang YF, Lin WC. A novel approach to predict drug-drug interactions based on chemical interactions and network topology. Comput Biol Med. 2019;107:10–18.
48. Atz K, Grisoni F, Schneider G. Prospective de novo drug design with deep interactome learning. Nat Commun. 2021;12(1):3543.

49. Popova M, Isayev O, Tropsha A. Deep reinforcement learning for de novo drug design. Sci Adv. 2018;4(7): eaap7885.

50. Aliper A, Plis S, Artemov A, et al. Deep learning applications for predicting pharmacological properties of drugs and drug repurposing using transcriptomic data. Mol Pharm. 2016;13(7):2524–2530.