International Journal of Biological and Pharmaceutical Sciences
Vol 1 No I, April, 2026, Page 1-25
An International Open Access, Peer Reviewed, Referred Journal

AI and ML in drug discovery: a comprehensive review related to virtual screening, molecular modelling & accelerating clinical trials.

Bhaskar Kumar Gupta; Ashish Kumar;

Principal, School of Pharmacy and Research, People's University, Bhopal, Madhya Pradesh,Assistant Professor, School of Pharmacy and Research, People's University, Bhopal, Madhya Pradesh

Corresponding Email: bhaskar.peoples1003@gmail.com,

Article History:
Received: 15/04/2026
Accepted: 18/04/2026
Online First: 30/04/2026

Keywords:
Computer-aided drug design, Artificial intelligence, Machine learning, Deep learning, Drug discovery and development, Virtual Screening, Molecular Modelling

This is an open access article under the CC-BY-NC-ND license

Abstract:

This chapter examines how Artificial Intelligence (AI) and Machine Learning (ML) have been transforming traditional drug discovery between 2019 and 2025, addressing the high costs, long timelines, and low success rates associated with this process. Highlighting recent advancements in AI/ML across the entire drug discovery pipeline, from target identification to clinical development. It explores various AI techniques, including deep learning, graph neural networks, and transformers, and their applications in key areas such as Target identification, Lead discovery, Hit optimization, and preclinical safety assessment. A comparative analysis of the advantages, limitations, and practical challenges of different AI approaches. It emphasizes crucial factors for successful implementation, including data quality, model validation, and ethical considerations. In this chapter, we study how we synthesize current applications, identify ongoing gaps (especially in data accessibility, interpretability, and clinical translation), and propose future directions. The ultimate goal is to unlock AI's full potential to create safer, more effective, and more accessible medicines by emphasizing transparent methodologies, robust validation, and ethical frameworks for responsible AI integration into pharmaceutical research and development. Drug-target identification: AI can pinpoint potential drug targets more effectively by analysing complex datasets. Through virtual screening, we studied molecular properties and compound analysis, Drug development and quality assurance, and Drug toxicity assessment. The integration of AI and ML offers a promising strategy to overcome the complexities of the pharmaceutical industry, accelerating the entire process from research to clinical trials and ultimately bringing safer, more effective medicines to patients faster and at a lower cost. While facing challenges such as data quality, interpretability of models, and regulatory hurdles, the collaborative efforts and continuous advancements in AI technology promise to unlock its full potential in revolutionizing pharmaceutical research and development. CADD has significantly impacted this area of research. Further, the combination of CADD with Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) technologies to handle enormous amounts of biological data has reduced the time and cost associated with the drug development process. This review will discuss how CADD, AI, ML, and DL approaches help identify drug candidates and various other steps of the drug discovery process. It will also provide a detailed overview of the different in silico tools used and how these approaches interact.

Introduction

The conventional drug development process remains time-consuming, expensive, and fraught with complications, frequently requiring more than a decade and costing upwards of $2 billion to bring a single therapeutic agent to market [1-3]. This substantial investment of

resources is largely attributable to the sequential nature of drug development phases, which demand extensive resources and rigorous validation at every stage—from target identification and hit discovery through lead optimization, preclinical testing, and protracted clinical trials. Critically, the process suffers from alarmingly low success rates, with approximately only 10% of drug candidates that enter clinical trials ultimately receiving regulatory approval. This high attrition rate is frequently exacerbated by safety concerns and lack of efficacy observed during late-stage development [2,4]. Furthermore, high-throughput screening (HTS), a popular approach for identifying initial hit compounds, yields only a 2.5% hit rate on average, further extending timelines, inflating costs, and squandering valuable research resources [5]. These persistent challenges underscore the urgent need for more efficient and effective approaches to drug discovery and development.

The pharmaceutical industry has witnessed a paradigm shift with the emergence and maturation of computational approaches. Computer-Aided Drug Design (CADD) emerged as a promising solution, leveraging computational power to model molecular interactions, predict binding affinities, and screen vast chemical libraries efficiently [6,7]. However, despite its significant contributions, traditional CADD approaches still face limitations in handling the complexity and scale of biological data. The advent of Artificial Intelligence (AI) and Machine Learning (ML) has presented a transformative opportunity to address these constraints and unlock new possibilities in drug development [8,9]. By implementing AI/ML methodologies, the pharmaceutical industry can not only address existing limitations but also create novel opportunities through sophisticated model implementations based on AI-driven parameters [10,11].

1.1 The Drug Discovery Pipeline: Challenges and Opportunities

Drug research and discovery represents an inherently complex, multi-dimensional process comprising several critical stages. The journey begins with target discovery and validation—the identification and confirmation of biological targets implicated in disease pathophysiology. This initial phase requires comprehensive proteomics, genomics, and bioinformatics studies to identify cellular and genetic targets amenable to therapeutic intervention. Following target validation, the process progresses to compound screening and hit discovery, where the first molecule demonstrating activity against the designated target is identified. This can be accomplished through screening chemical libraries or isolating natural compounds from diverse sources including fungi, plants, and bacteria [12].

The subsequent stage involves identifying the lead chemical—the compound with the greatest potential to be developed into a therapeutic agent. Lead optimization then commences, involving systematic chemical modification of the selected lead to enhance its specificity, efficacy, and drug-like properties, ideally achieving therapeutic effects at lower dosages. This iterative cycle integrates structure-activity relationships and cellular assays to progressively enhance the functional characteristics of newly synthesized therapeutic candidates [13]. Following successful optimization, in vivo investigations involving pharmacokinetic and toxicity evaluations are conducted using animal models to assess the compound's behaviour in living systems.

After rigorous preclinical testing, the drug candidate advances to clinical trials—the most expensive and time-consuming phase of drug development. Clinical trials are essential for determining whether a drug candidate can provide desired medical benefits while maintaining patient safety. This multi-phase process involves escalating numbers of human participants and typically spans several years [14,15]. Pharmaceutical enterprises therefore continuously seek strategies to reduce costs and accelerate project timelines, making AI-driven approaches increasingly attractive.

1.2 The Emergence of AI in Pharmaceutical Sciences

The ability of machines to mimic human cognitive processes in learning and problem-solving—known as Artificial Intelligence—has found increasingly diverse applications across pharmaceutical sciences. Technology-based AI systems, utilizing an array of sophisticated tools and neural networks, can effectively simulate and augment human intelligence in ways previously unimaginable [16]. To save time, reduce costs, and increase profitability, AI-based technologies are being increasingly deployed across various phases of the drug discovery process. These applications span diverse tasks including real-time cell sorting, cell classification, quantum mechanics-based compound attribute calculation, computational organic synthesis, de novo compound creation, and numerous others [17,18]. The integration of AI into drug discovery represents a fundamental departure from traditional approaches. Rather than merely accelerating existing processes, AI enables entirely new methodologies—from predicting protein structures with near-experimental accuracy to generating novel molecular structures optimized for specific biological targets. As highlighted in recent literature, "AI can identify hit and lead compounds, enabling faster drug-target validation and optimization of drug structure design" [19,20]. This transformative potential has attracted substantial investment from both established pharmaceutical companies and emerging biotechnology firms, driving rapid innovation in the field.

1.3 Scope and Objectives of This Review

This comprehensive review examines the transformative impact of AI and ML on drug discovery, with particular focus on the period from 2019 to 2025. It synthesizes current applications, identifies ongoing gaps and challenges, and proposes future directions for the field. The review places special emphasis on three key areas: virtual screening, molecular modelling, and clinical trial acceleration. By providing a detailed overview of different in silico tools and approaches, this review aims to serve as a valuable resource for medicinal chemists, computational scientists, and pharmaceutical researchers seeking to leverage AI technologies in their drug discovery projects. Furthermore, it explores how CADD, AI, ML, and Deep Learning (DL) approaches interact and complement each other in identifying drug candidates and advancing various steps of the drug discovery process.

2. COMPUTER-AIDED DRUG DESIGN: FOUNDATIONS AND EVOLUTION

Computer-Aided Drug Design (CADD) has fundamentally reshaped the landscape of pharmaceutical research over the past three decades. By leveraging computational power to model, predict, and optimize molecular interactions, CADD has substantially reduced the time and cost associated with traditional drug discovery approaches [21]. The field has evolved from simple molecular modeling tools to sophisticated platforms integrating quantum mechanics, molecular dynamics, and, most recently, artificial intelligence. Understanding the foundations of CADD is essential for appreciating the transformative impact of AI/ML integration [22].

2.1 Structure-Based Drug Design (SBDD)

Structure-Based Drug Design (SBDD) represents a cornerstone of modern computational drug discovery. This approach leverages three-dimensional structural information of biological targets—typically proteins—to guide the design and optimization of therapeutic compounds. The fundamental premise of SBDD is that knowledge of a target's atomic-level structure enables rational design of molecules that bind specifically and effectively to that target [23].

2.1.1 Target Preparation

The SBDD workflow begins with obtaining a reliable three-dimensional structure of the target protein. While experimental methods including X-ray crystallography, Nuclear Magnetic Resonance (NMR) spectroscopy, and cryo-electron microscopy (cryo-EM)

provide gold-standard structures, they are not always available for all targets of interest. Computational approaches have therefore been developed to generate protein structures when experimental data is unavailable [24,25].

Homology modeling, also known as comparative modeling, remains the most widely used computational method for protein structure prediction. This approach constructs a three-dimensional model of a target protein based on its sequence similarity to one or more proteins with known experimental structures (templates). The accuracy of homology models correlates strongly with sequence similarity—models based on templates with >50% sequence identity typically achieve near-experimental accuracy [26]. The success of tools like SWISS-MODEL and MODELLER has democratized access to protein structure prediction.

Fold recognition or threading methods represent a more advanced approach applicable when no suitable template exists for straightforward homology modeling. These methods detect compatible folds by evaluating how well a target sequence fits into known structural folds, even when sequence similarity is minimal. Threading algorithms evaluate the compatibility of amino acid sequences with different structural environments, identifying potential folds that may not be detectable through sequence alignment alone [27,28].

Ab initio or de novo modeling represents the most challenging approach, predicting protein structure from sequence alone without relying on known templates. These methods simulate the physical forces governing protein folding to predict tertiary structure. While historically limited to small proteins, recent advances in deep learning—particularly DeepMind's AlphaFold and RoseTTAFold—have revolutionized ab initio prediction, achieving near-experimental accuracy for a vast range of proteins [29,30].

2.1.2 Binding Site Identification and Characterization

Once a target structure is available, the next critical step involves identifying and characterizing the active binding site—the region where potential drug molecules will interact with the target. Binding sites typically possess distinctive characteristics including surface clefts, hydrophobic patches, and specific hydrogen-bonding patterns. Computational tools for binding site prediction include CASTp, Fpocket, and SiteMap, which analyze protein geometry, electrostatic potential, and evolutionary conservation to identify potential binding pockets [31,32].

2.1.3 Molecular Docking

Molecular docking constitutes the central technology of SBDD, predicting the preferred orientation, conformation, and binding affinity of a small molecule (ligand) when bound to a target protein. Docking algorithms must efficiently sample possible ligand poses and orientations within the binding site while accurately scoring the predicted binding affinity [33]

Classification Methods for Protein-Ligand Docking**

The accurate representation of molecular flexibility remains one of the greatest challenges in molecular docking. Several approaches have been developed to address this complexity at varying computational costs.

Protein Flexibility: Traditional docking approaches treated proteins as rigid bodies, assuming that ligand binding does not significantly alter protein conformation. However, this assumption often fails in practice, as induced fit effects can dramatically influence binding. Modern approaches address protein flexibility through several strategies.

Soft docking represents a computationally efficient compromise, employing softened potential energy functions that tolerate small degrees of steric overlap between protein and ligand atoms. This approach allows minor side-chain movements without explicit conformational sampling [34].

Side-chain flexibility methods explicitly model the conformational freedom of amino acid side chains within the binding site. Rotamer libraries—collections of energetically favourable side-chain conformations—enable systematic or stochastic sampling of side-chain positions during docking [35].

Molecular relaxation techniques, such as energy minimization or short molecular dynamics simulations, refine docked complexes by allowing both protein and ligand to relax into energetically favorable conformations. While more computationally demanding, these approaches better capture induced fit effects [36].

Protein ensemble docking represents a rigorous approach to protein flexibility, docking ligands against multiple protein conformations derived from experimental structures or molecular dynamics simulations. This ensemble approach can identify cryptic binding sites and capture conformational selection mechanisms [37,38].

Ligand Sampling: Efficient sampling of ligand conformational space is essential for accurate docking. Systematic search methods explore ligand degrees of freedom exhaustively through incremental construction, torsional scanning, or library-based approaches. While systematic search ensures comprehensive sampling, computational costs grow exponentially with ligand flexibility, limiting applicability to relatively rigid molecules [39].

Stochastic algorithms, including genetic algorithms, Monte Carlo methods, and particle swarm optimization, provide more efficient sampling of high-dimensional conformational spaces. These methods employ probabilistic search strategies to identify low-energy conformations without exhaustive enumeration. AutoDock's genetic algorithm and Glide's Monte Carlo searches exemplify this approach [40,41].

Scoring Functions: Scoring functions predict the binding affinity of docked poses, enabling ranking of potential ligands and identification of promising candidates. Three main classes of scoring functions are widely employed.

Force field (FF) scoring functions calculate binding affinity as the sum of molecular mechanics energies—van der Waals interactions, electrostatic energies, and internal ligand strain. These physically rigorous functions, implemented in AutoDock and DOCK, provide detailed energetic descriptions but are computationally intensive [42,43].

Empirical scoring functions approximate binding free energy as a weighted sum of individual energy terms including hydrogen bonds, hydrophobic contacts, and rotatable bond penalties. These functions, exemplified by GlideScore and ChemScore, are parameterized using experimental binding affinity data for training sets of protein-ligand complexes [44,45].

Knowledge-based scoring functions derive statistical potentials from databases of known protein-ligand complexes, assuming that observed interatomic distances reflect energetically favourable interactions. These functions, including DrugScore and PMF, capture implicit contributions from solvation and entropy without explicit calculation [46,47].

2.2 Virtual Screening (VS)

Virtual screening has emerged as a powerful computational technique for identifying promising drug candidates from large compound libraries, dramatically reducing the number of compounds requiring experimental testing. Virtual screening approaches fall into two main categories, each with distinct advantages and applications [48,49].

Structure-Based Virtual Screening (SBVS): SBVS employs molecular docking to computationally screen large compound libraries against a target protein structure. Each compound is docked into the binding site, scored for predicted binding affinity, and ranked to identify top candidates for experimental validation. SBVS can screen millions of compounds in silico, representing a massive efficiency gain over experimental high-throughput screening [50].

Ligand-Based Virtual Screening (LBVS): When protein structural information is unavailable, LBVS provides an alternative approach based on known active ligands. LBVS identifies compounds similar to known actives using molecular fingerprints, pharmacophore models, or shape-based comparisons. Methods including similarity searching, quantitative structure-activity relationship (QSAR) models, and machine learning classifiers have been successfully applied in LBVS campaigns [51,52].

2.3 Molecular Dynamics (MD) Simulation

Molecular dynamics simulation represents the most rigorous computational approach for studying biomolecular systems, providing atomic-level insights into protein dynamics, ligand binding, and conformational changes over time. MD simulations numerically solve Newton's equations of motion for all atoms in a system, generating trajectories that reveal the dynamical behaviour of proteins and their complexes with ligands [53,54].

Several force fields—parameter sets describing the potential energy of atoms as a function of their positions—have been developed for biomolecular simulation. AMBER Assisted Model Building with Energy Refinement) and CHARMM (Chemistry at Harvard Macromolecular Mechanics) represent two of the most widely used and validated force fields for protein and nucleic acid simulation. OPLS (Optimized Potentials for Liquid Simulations) and its all-atom variant OPLS-AA provide excellent descriptions of organic molecules and have been extensively parameterized for drug-like compounds. GROMOS (Groningen Molecular Simulation) offers a united-atom representation that balances accuracy with computational efficiency. Coarse-grained force fields reduce computational costs by grouping atoms into interacting beads, enabling simulations of larger systems over longer timescales, though with reduced atomic detail [55-58].

2.4 Ligand-Based Drug Design (LBDD)

When structural information about the target protein is unavailable, Ligand-Based Drug Design (LBDD) provides powerful alternatives for guiding compound optimization. LBDD approaches leverage information from known ligands to build predictive models and guide the design of new compounds with improved properties [59].

Quantitative Structure-Activity Relationship (QSAR): QSAR modelling represents a cornerstone of LBDD, establishing mathematical relationships between molecular structure and biological activity. The QSAR workflow involves several essential steps: obtaining a congeneric set of ligands tested in similar biological assays; identifying molecular descriptors that capture relevant physicochemical properties; calculating descriptor values for training set compounds; developing statistical models correlating descriptors with biological activity; and validating model performance using external test sets [60,61].

Pharmacophore Modeling: Pharmacophore models represent the three-dimensional arrangement of chemical features essential for biological activity. The pharmacophore modelling workflow involves selecting a training set of ligands, including both active and inactive compounds; generating low-energy conformations likely to include bioactive conformations; superimposing possible combinations of molecular conformations; abstracting an abstract representation of overlaid functional groups; and validating the resulting pharmacophore hypothesis [62,63].

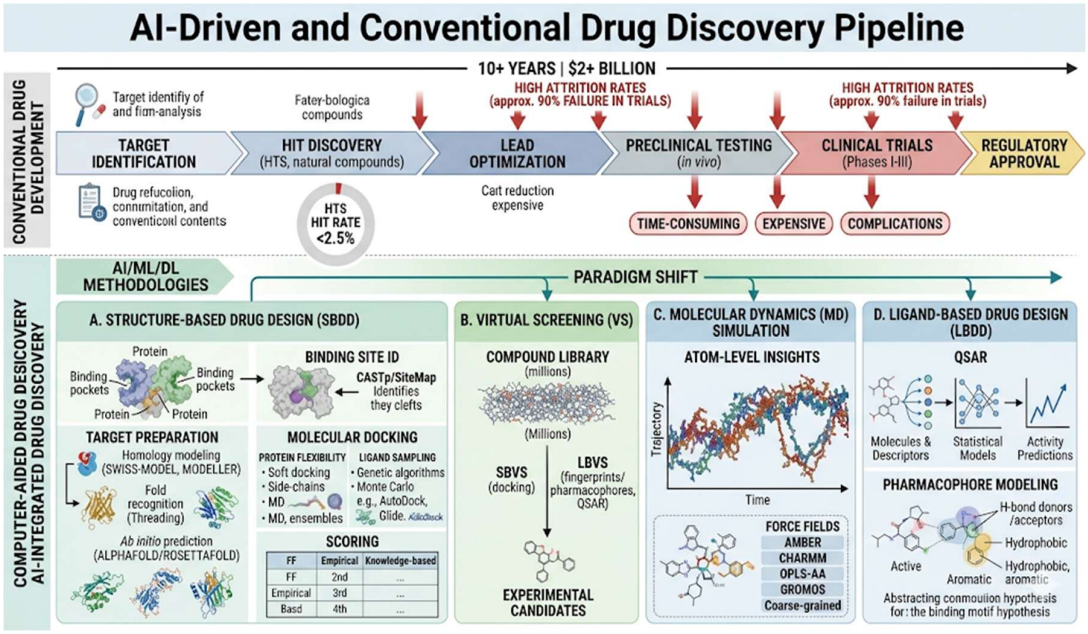

Figure 1 Conventional Drug Development: A sequential, resource-intensive pipeline spanning target identification to regulatory approval that takes over 10 years, costs upwards of $2 billion, and suffers from a high 90% failure rate in clinical trials.Computer-Aided & AI-Integrated Paradigms: A transformative approach that shifts traditional methods toward sophisticated computational tracks divided into Structure-Based Drug Design (SBDD), Virtual Screening (VS), Molecular Dynamics (MD) Simulations, and Ligand-Based Drug Design (LBDD). Advanced In Silico Techniques: Methodologies leveraging tools like AlphaFold for target preparation, molecular docking/scoring for binding site interactions, trajectory simulations across diverse force fields (like AMBER and CHARMM), and QSAR/pharmacophore modeling to systematically predict and optimize experimental candidates.

3. ARTIFICIAL INTELLIGENCE IN DRUG DISCOVERY

3.1 Defining Artificial Intelligence in Pharmaceutical Contexts

Artificial Intelligence encompasses a broad set of technologies enabling machines to perform tasks that typically require human intelligence. In pharmaceutical contexts, AI systems are being deployed across the entire drug discovery and development pipeline, from target identification through clinical development. As defined in recent literature, "Artificial intelligence is a scientific discipline dedicated to intelligent machine learning, primarily encompassing advanced computer programs that replicate human cognitive functions" [16,64].

The integration of AI across pharmaceutical sectors has created new opportunities for enhancing medication discovery, patient care, and overall healthcare efficiency. Key applications span drug discovery and development (expediting novel pharmaceutical identification by scrutinizing extensive databases), drug interaction analysis (improving prediction precision), personalized medicine (enabling customized therapeutic strategies based on genetic composition), drug safety monitoring (enhancing pharmacovigilance), clinical trial optimization (transforming trial design and analysis), and pharmacy automation (improving operational efficiency) [16,65].

3.2 Key AI Techniques in Drug Discovery

Artificial Neural Networks (ANNs): Artificial Neural Networks are computational models inspired by the neuronal network architecture of the human brain. ANNs, in their most basic configuration, comprise fully connected or feed-forward networks including three layers: input layer, hidden layer(s), and output layer. Each layer contains computational units called neurons that function as non-linear modifications of incoming data [66,67].

ANNs employ two primary learning paradigms. Unsupervised learning provides input data exhibiting known patterns utilized for organizational purposes, with the Self-Organizing Map or Kohonen network being a prominent algorithm. This approach is highly effective for identifying relationships within complex datasets. Supervised learning employs corresponding inputs and outputs to learn relationships between these variables, proving useful for modelling cause-and-effect connections. The backpropagation learning algorithm is the most utilized supervised learning method, demonstrating exceptional performance for prediction and classification tasks [16,68].

Applications of ANNs in pharmaceutical contexts are extensive, including data analysis, pharmaceutical quality control modelling, molecular modeling, QSAR investigations, formulation optimisation, and biopharmaceutical studies, including pharmacokinetic and pharmacodynamic modelling [69,70].

Support Vector Machines (SVMs): Support Vector Machines represent a class of supervised learning algorithms that find optimal hyperplanes separating different classes of data. SVMs excel at handling high-dimensional data and have found extensive applications in drug discovery including compound classification, toxicity prediction, and QSAR modelling. The ability of SVMs to handle non-linear relationships through kernel functions makes them particularly valuable for complex biological datasets [71,72].

Deep Learning: Deep learning extends traditional neural networks by incorporating multiple hidden layers, enabling hierarchical feature extraction from raw data. Deep learning architectures including Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Transformers have achieved remarkable success across diverse drug discovery applications. CNNs excel at extracting spatial patterns from molecular graphs and protein structures, RNNs process sequential data including SMILES strings and protein sequences, while Transformers capture long-range dependencies and have revolutionized protein language modelling [73-75].

Graph Neural Networks (GNNs): GNNs have become central to cheminformatics because molecules inherently form graph-structured data where atoms represent nodes and chemical bonds represent edges. GNN architectures including Graph Convolutional Networks (GCNs), Graph Attention Networks (GATs), and Message Passing Neural Networks (MPNNs) learn molecular representations by iteratively aggregating information from neighbouring atoms. These approaches have demonstrated state-of-the-art performance in property prediction, toxicity assessment, and drug-target interaction prediction [76,77].

Recent advances have produced sophisticated frameworks such as **GS-DTI** (Graph Structure-based Drug-Target Interaction prediction), which integrates molecular graph transformers, protein language models, and protein tertiary structure to achieve robust and interpretable DTI predictions. GS-DTI extracts drug features from SMILES-derived molecular graphs using knowledge-guided pre-trained transformers while protein features derive from both sequence and predicted 3D structure. Importantly, GS-DTI achieved more than 10% improvement in Matthew's Correlation Coefficient (MCC) over previous methods in drug-target pair cold start tests, demonstrating superior generalization to unseen compounds and targets [78].

Generative Models: Generative models enable exploration of the vast chemical space—estimated to encompass 10^{30} to 10^{60} potential small molecules. Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and diffusion models learn the underlying distribution of known molecular structures and generate novel compounds with desired properties. The GENTRL model demonstrated this capability by designing potent DDR1 inhibitors in radically reduced timeframes [79,80]. Similarly, reinforcement learning approaches refine compounds by rewarding novelty and drug-likeness while optimizing for specific biological activities [81].

3.3 The Black Box Problem and Explainable AI

Despite the remarkable predictive performance of deep learning models, their inherent opacity poses significant challenges for pharmaceutical applications. The "black-box" nature of these models limits interpretability and acceptance within the pharmaceutical research community [82]. As noted in recent literature, "although these AI models yield highly accurate results, the basis for their reasoning is obscured by the highly complex mathematical processes that underpin these models" [83].

Explainable Artificial Intelligence (XAI) has emerged as a crucial solution for enhancing transparency, trust, and reliability by clarifying the decision-making mechanisms that underpin AI predictions. XAI techniques bridge the gap between AI model outcomes and the underlying reasoning behind those outcomes, establishing foundations for trusting models that assist in drug design pipelines [84,85].

Two widely accepted explainability methods are SHAP (Shapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations). SHAP values, grounded in cooperative game theory, quantify the contribution of each molecular feature to model predictions. LIME approximates complex model behaviour locally with interpretable surrogate models. These techniques enable researchers to identify which molecular features or descriptors contribute most significantly to predictions, estimate marginal feature contributions to outputs, and highlight specific substructures associated with predicted outcomes [86,87].

XAI applications in drug discovery encompass molecular modelling, therapeutic target identification, ADME prediction, clinical trial design, personalized medicine, and molecular property prediction. By offering human-interpretable explanations for model predictions—such as identifying substructures responsible for poor absorption or toxicity—XAI enables medicinal chemists to rationally prioritize or modify molecular scaffolds, improve candidate selection, and enhance lead optimization [83,88].

Table 1: Comparison of Major AI Techniques in Drug Discovery

Technique	Core Principle	Key Applications	Advantages	Limitations
Artificial Neural Networks (ANNs)	Multi-layer perceptrons with backpropagation	QSAR, property prediction, formulation optimization	Handles non-linear relationships; integrates diverse data	Requires large training sets; potential overfitting
Deep Neural Networks	Hierarchical feature extraction through multiple layers	Molecular property prediction; image analysis; sequence analysis	Automatic feature learning; high accuracy	Computationally intensive; black-box nature
Graph Neural Networks (GNNs)	Message passing on molecular graphs	Property prediction; toxicity assessment; DTI prediction	Naturally represents molecular structure; captures topological features	Graph construction sensitivity; scaling challenges

Transformers	Self-attention mechanisms for sequence processing	Protein language modelling; SMILES generation; DTI prediction	Captures long-range dependencies; parallel processing	Large parameter counts; extensive training data
Generative Models (VAEs, GANs)	Learning data distribution for novel generation	De novo drug design; scaffold hopping; lead optimization	Explores vast chemical space; generates novel structures	Mode collapse (GANs); validity control challenges
Reinforcement Learning	Reward-based policy optimization	Molecular optimization; reaction prediction	Goal-directed generation; multi-property optimization	Reward design complexity; sample efficiency
XAI (SHAP, LIME)	Feature attribution and local approximation	Model interpretation; feature importance; toxicity explanation	Enhances transparency; builds trust; regulatory support	Computational overhead; approximation accuracy

4. MACHINE LEARNING IN DRUG DISCOVERY

4.1 Fundamentals of Machine Learning

Machine learning (ML) enables computers to learn predictive patterns from data without being explicitly programmed for specific tasks. The development of novel algorithms and models with the ability to comprehend massive volumes of data constitutes the essence of machine learning. While not all AI approaches are ML techniques, machine learning is defined as "an AI technique used to design and train software algorithms to learn from and act on data" [16,89].

By utilizing diverse algorithms, machine learning can predict the physical, biological, and chemical characteristics of new molecules based on patterns learned from existing data. ML approaches are broadly categorized into supervised, unsupervised, semi-supervised, and reinforcement learning paradigms, each suited to different types of problems and data availability [90,91].

4.2 Supervised Machine Learning

Supervised learning aims to train algorithms to generate predictions by identifying patterns and testing hypotheses using labelled training data. Supervised learning algorithms learn the relationship between input features (molecular descriptors, fingerprints, or graphs) and output labels (activity, toxicity, or property values) from examples where both are known [92].

Classification algorithms categorize data into discrete classes based on training dataset patterns. Common applications in bioinformatics include identifying genomic regions coding for genes and classifying compounds as active or inactive against a target. Classification algorithms achieve high accuracy when trained on sufficiently large and representative datasets [93].

Regression algorithms predict continuous values such as binding affinities, IC50 values, or solubility measurements. Recent applications have focused on predicting new targets or structures, including locations of protein-protein interactions. Studies on regression methods have shown encouraging results with accuracy exceeding 80% in proteomics structure identification [94].

The major advantage of supervised learning lies in using optimal decision boundaries derived from labelled training data. However, supervised approaches face several limitations including complexity and time consumption when classifying large datasets, risk of overtraining decision boundaries due to insufficient or unrepresentative examples, and challenges in data preparation and pre-processing [95].

4.3 Unsupervised Machine Learning

Unsupervised ML does not rely on predetermined labels or phenotypes to interpret or learn abstract representations of provided data. Instead, these methods extract useful biological information by clustering data points into patterns based on inherent similarities. Two popular clustering techniques are hierarchical clustering and k-means clustering [96].

Clustering algorithms partition unlabeled datasets into groups based on shared characteristics. When working with massive datasets, k-means clustering is frequently employed to group small molecule profiles into clusters according to similarity degrees. This approach has proven valuable for identifying chemical series and scaffold hopping opportunities [97].

Association rule learning discovers relationships between variables in large datasets, identifying patterns such as frequent substructures associated with particular activities or toxicity profiles. Association methods have been applied to identify combinations of molecular features predictive of desired properties [98].

A major advantage of unsupervised learning techniques over supervised algorithms is their reduced complexity, as they do not require labelled training data and are useful for sorting raw data and understanding learning models in real-time. Additionally, unlabeled data is substantially easier to obtain automatically compared to labelled data requiring human annotation. However, unsupervised techniques also face significant disadvantages including imprecise data sorting due to lack of labelled data, leading to less accurate and potentially unpredictable results [99].

4.4 Semi-Supervised and Reinforcement Learning

Semi-supervised learning bridges supervised and unsupervised approaches by leveraging small amounts of labelled data together with larger quantities of unlabelled data. This paradigm is particularly valuable in drug discovery where experimental labelling (e.g., activity determination) is expensive and time-consuming, while unlabelled compound data is abundant. Semi-supervised methods have been successfully applied to virtual screening and toxicity prediction [100,101].

Reinforcement learning (RL) trains agents to make sequential decisions by rewarding desirable actions and penalizing undesirable ones. In drug discovery, RL has been applied to de novo molecular design, where agents generate molecules atom-by-atom or bond-by-bond, receiving rewards for satisfying desired property profiles. Positive reinforcement learning rewards the agent for generating molecules with favourable properties (high predicted activity, drug-likeness), while negative reinforcement learning penalizes undesirable characteristics (toxicity, poor synthetic accessibility) [102,103]. RL-guided optimization has demonstrated significant success in improving potency while maintaining drug-like properties [81].

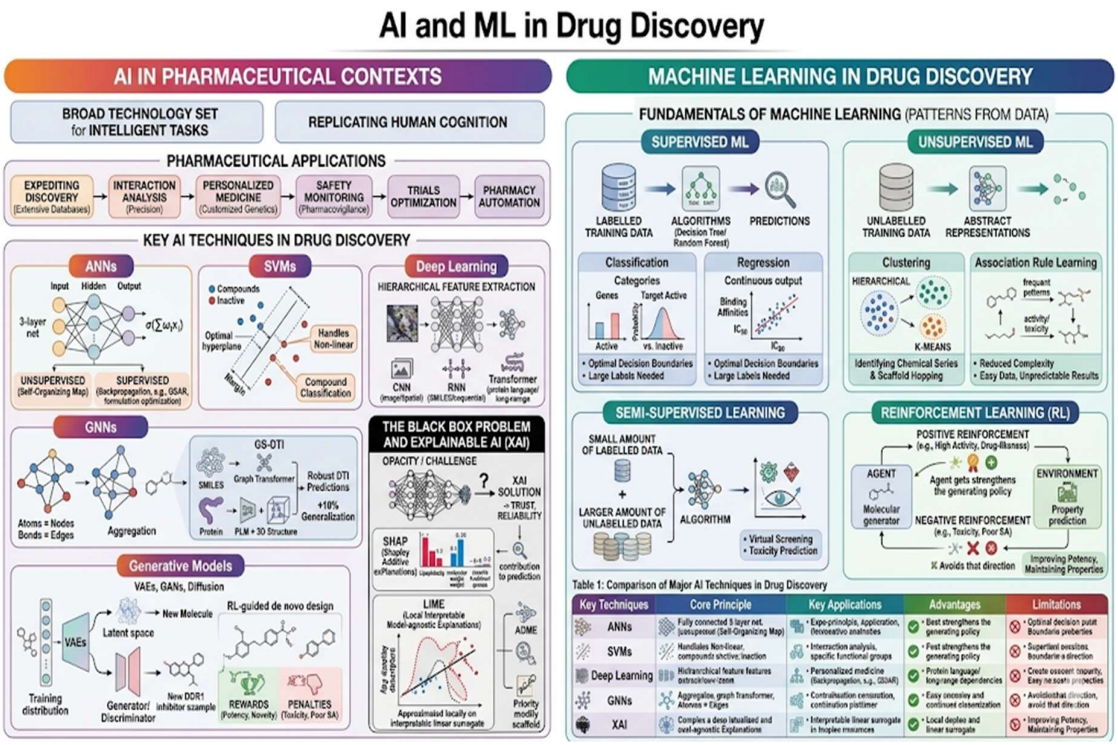

Figure 2 - AI in Pharmaceutical Contexts: This section illustrates how intelligent systems replicate human cognition to expedite tasks like interaction analysis, personalized medicine, and safety monitoring, utilizing core architectures like ANNs, SVMs, Deep Learning (CNNs/RNNs/Transformers), GNNs, and Generative Models. The Black Box Problem & XAI: It highlights the opacity challenge of highly complex deep learning models and presents Explainable AI (XAI) solutions—specifically SHAP and LIME—which provide human-interpretable feature attributions to safely predict and modify ADME or toxicity properties. Machine Learning Paradigms: The right half details the fundamentals of processing data through Supervised ML (Classification/Regression), Unsupervised ML (Clustering/Association Rules), Semi-Supervised ML, and Reinforcement Learning, which optimizes molecular generation using goal-directed rewards and penalties.

5. APPLICATIONS ACROSS THE DRUG DISCOVERY PIPELINE

5.1 Target Identification and Validation

Target identification—the process of identifying biological molecules (typically proteins) whose modulation is expected to produce therapeutic benefit—represents the critical first step in drug discovery. AI and ML have dramatically enhanced target identification capabilities by integrating and analyzing multi-omics data including genomics, transcriptomics, proteomics, and metabolomics [104].

Machine learning models analyze complex biological networks to identify disease-relevant targets based on their position and connectivity within pathways. Ensemble learning and network-based deep learning models help prioritize targets with high druggability potential—the likelihood that a target can be modulated by a small molecule drug. Bayesian classifiers and random forest algorithms support phenotypic correlation-based target selection by linking genetic variations or expression patterns to disease phenotypes [105,106].

Deep learning approaches have been particularly successful at predicting target-disease associations by integrating heterogeneous data sources. Graph neural networks model biological systems as graphs where nodes represent genes, proteins, or metabolites and edges represent interactions or relationships. These models can predict previously unknown

disease associations for potential targets and identify targets likely to respond to therapeutic intervention [107,108].

5.2 Virtual Screening and Hit Discovery

Virtual screening has been dramatically enhanced by machine learning approaches, which often outperform classical docking methods in both accuracy and speed. ML-based scoring functions, such as RF-Score and DeepDock, learn to predict binding affinity from protein-ligand complex structures, capturing interaction patterns that fixed scoring functions may miss [109,110].

Deep learning has transformed structure-based virtual screening through convolutional neural networks that extract spatial protein-ligand interaction patterns from three-dimensional grids or distance maps. Graph neural networks operating on molecular graphs have demonstrated state-of-the-art performance in predicting binding affinities and identifying promising hit compounds. Importantly, AI-enabled virtual screening can evaluate billions of compounds in silico, exploring chemical space at scales impossible with experimental approaches [111,112].

Large language models (LLMs) and foundation models pre-trained on massive chemical datasets have recently emerged as powerful tools for virtual screening. Models including ChemBERTa, Galactica, and MolFormer learn rich molecular representations that transfer effectively to diverse downstream tasks including property prediction and virtual screening [113,114].

5.3 Lead Optimization and Property Prediction

Lead optimization—the iterative process of modifying hit compounds to improve potency, selectivity, and drug-like properties—has been substantially accelerated by AI/ML approaches. QSAR models and deep neural networks predict ADMET (Absorption, Distribution, Metabolism, Excretion, and Toxicity) properties with increasing accuracy, enabling early identification and elimination of compounds with poor pharmacokinetic profiles [115,116].

Deep learning models specifically developed for toxicity prediction, including DeepTox and ADMET-AI, have demonstrated impressive performance in identifying compounds likely to cause adverse effects. These models learn from large databases of compounds with known toxicities, identifying structural alerts and property combinations associated with safety risks. By enabling early toxicity prediction, AI helps eliminate problematic compounds before costly preclinical and clinical studies [117,118].

Multi-property optimization—simultaneously optimizing multiple often-competing properties—represents a particular strength of AI approaches. Generative models guided by multi-objective reward functions can design molecules balancing potency, selectivity, solubility, permeability, metabolic stability, and safety. Reinforcement learning frameworks have successfully optimized lead compounds by iteratively proposing modifications and receiving feedback on predicted properties [119,120].

5.4 Protein Structure Prediction

The breakthrough achievement of AlphaFold—DeepMind's AI system for protein structure prediction—has revolutionized structural biology and drug discovery. AlphaFold achieved near-experimental accuracy in predicting three-dimensional protein structures from amino acid sequences alone, solving a 50-year grand challenge in biology. The subsequent AlphaFold2 and AlphaFold3 extended capabilities to predict protein-ligand complexes, protein-protein interactions, and the effects of mutations on structure [121,122].

The impact of accurate protein structure prediction on drug discovery cannot be overstated. For the first time, researchers can obtain high-quality structural models for thousands of previously uncharacterized proteins, enabling structure-based drug design for targets lacking experimental structures. AlphaFold predictions have been successfully applied to identify

cryptic binding sites, predict ligand binding poses, and guide virtual screening campaigns [123,124].

Complementary AI approaches predict protein-ligand docking, binding pocket identification, and the effects of mutations on protein stability and function. These tools, integrated into platforms like Flare™ and Schrödinger, enable computational exploration of protein-ligand interactions at unprecedented scale and accuracy [125].

5.5 Drug Repurposing

Drug repurposing—identifying new therapeutic uses for existing approved drugs—offers significant advantages including reduced development timelines, lower costs, and established safety profiles. AI models have proven exceptionally effective at identifying repurposing opportunities by integrating diverse data sources [126].

Large language models including BioBERT and SciBERT extract drug-disease relationships from the biomedical literature, identifying previously unrecognized connections. Network-based approaches model drug-target-disease relationships as graphs, propagating information from known associations to predict novel indications. Deep learning models predict drug-disease associations by integrating chemical structures, genomic data, and clinical information [127,128].

The COVID-19 pandemic provided a striking demonstration of AI-enabled drug repurposing. BenevolentAI's knowledge graph identified baricitinib, a rheumatoid arthritis drug, as a potential COVID-19 treatment. This prediction was rapidly validated, and baricitinib subsequently received emergency use authorization. Similar AI-driven repurposing efforts identified other candidates including remdesivir, demonstrating the power of computational approaches to rapidly respond to emerging health threats [129,130].

5.6 De Novo Drug Design

De novo drug design—generating entirely new chemical entities with desired properties from scratch—represents the frontier of AI-driven drug discovery. Generative models learn the underlying distribution of drug-like chemical space and sample novel molecules optimized for specified property profiles [131].

The GENTRL model demonstrated remarkable capabilities by designing potent inhibitors of DDR1, a kinase target implicated in fibrosis. From target selection to experimental validation, the entire process was completed in just 21 days—a fraction of the time required by traditional approaches. Six compounds synthesized based on GENTRL designs showed nanomolar potency in biochemical assays, with one demonstrating favourable pharmacokinetic properties in mice [79,132].

Generative chemistry platforms, such as those integrated into Cresset's Flare™ V11, combine AI generation with physics-based evaluation. Researchers can re-train generative AI models on specific biological targets, generating focused libraries of molecules likely to bind productively. These approaches dramatically accelerate hit generation, exploring chemical regions that medicinal chemists might not intuitively consider [125].

6. ACCELERATING CLINICAL TRIALS WITH AI

6.1 Patient Recruitment and Trial Design

Clinical trials represent the most expensive and time-consuming phase of drug development, typically accounting for 60-80% of total development costs. AI and ML are transforming clinical trial operations across multiple dimensions, offering substantial opportunities for acceleration and efficiency improvement [133,134].

Patient recruitment—identifying and enrolling eligible participants—remains a major bottleneck in clinical trials, with approximately 80% of trials failing to meet enrolment timelines. AI addresses this challenge by automatically screening electronic health records (EHRs) to identify patients matching trial eligibility criteria. Natural language processing

extracts relevant clinical information from unstructured physician notes, while ML models predict patient likelihood of meeting inclusion/exclusion criteria [135].

Site selection—choosing the most effective locations for trial conduct—benefits from AI analysis of historical trial performance, patient population characteristics, and investigator expertise. ML models predict site-specific recruitment rates and data quality, enabling optimized trial footprint design that reduces costs and accelerates completion [136].

Adaptive trial design, enabled by AI, allows pre-specified modifications to ongoing trials based on accumulating data. ML models continuously monitor trial outcomes, enabling early termination of ineffective arms, sample size re-estimation, or enrichment of responding patient populations. These dynamic approaches substantially reduce the number of patients exposed to ineffective treatments while accelerating identification of effective therapies [137,138].

6.2 Patient Stratification and Enrichment

Patient heterogeneity—the substantial variability in how different individuals respond to the same treatment—contributes significantly to clinical trial failures. AI enables sophisticated patient stratification approaches that identify subgroups most likely to benefit from experimental therapies [139].

Response prediction models integrate baseline patient characteristics including genomics, proteomics, imaging, and clinical data to predict individual treatment response. Deep learning models, particularly those combining multiple data modalities, have demonstrated impressive accuracy in predicting which patients will respond to specific therapies. This capability enables enrichment strategies that selectively enrol predicted responders, reducing required sample sizes and increasing statistical power to detect treatment effects [140,141].

Biomarker discovery—identifying measurable indicators of treatment response or disease progression—has been dramatically accelerated by AI analysis of high-dimensional molecular data. ML models identify transcriptomic, proteomic, or metabolomic signatures predictive of clinical outcomes, enabling development of companion diagnostics that guide treatment selection in clinical practice [142].

6.3 Real-World Data and External Control Arms

Real-world data (RWD)—health information collected outside of traditional clinical trials, including EHRs, claims databases, and patient registries—offers substantial opportunities to augment and accelerate clinical development. AI enables effective utilization of RWD by addressing challenges of data heterogeneity, missingness, and bias [143,144].

External control arms use RWD from patients receiving standard of care to supplement or replace concurrent control groups. By leveraging AI to match trial patients to similar real-world patients, external control arms reduce the number of patients required to receive placebo or standard therapy, accelerating enrolment and reducing costs. Regulatory agencies have increasingly accepted external control arms in registration trials, particularly for rare diseases or oncology settings [145].

Historical trial emulation using AI enables "virtual twins" of trial participants, predicting outcomes for patients not receiving treatment based on similar patients in historical trials. These approaches enable more efficient trial designs and provide context for interpreting open-label or single-arm study results [146].

6.4 Clinical Outcome Prediction

Predicting which patients will experience clinical events—disease progression, adverse events, or treatment response—enables more efficient trial monitoring and early identification of safety signals. AI models trained on clinical trial data can identify patients at elevated risk of protocol violations, dropout, or adverse events, enabling proactive intervention [147].

Adverse event prediction models integrate patient characteristics, laboratory measurements, and trial data to predict individual patient risk of specific adverse events. Deep learning approaches, particularly those incorporating longitudinal patient data, have demonstrated superior performance compared to traditional statistical methods. These models enable risk-based monitoring strategies that focus resources on high-risk patients [148,149].

Digital biomarkers derived from wearable devices and sensors provide continuous, objective measurements of patient status throughout clinical trials. AI analysis of these high-dimensional time-series data can detect subtle changes in patient condition earlier than traditional assessments, enabling more sensitive detection of treatment effects and earlier identification of safety signals [150].

6.5 Regulatory Perspectives and the FDA Roadmap

Regulatory acceptance is essential for AI-driven clinical trial innovations to impact drug development. The US Food and Drug Administration (FDA) has demonstrated increasing openness to AI-enabled approaches, releasing guidance documents and establishing specialized review groups [151,152].

In April 2025, the FDA published a roadmap for leveraging New Approach Methodologies (NAMs), including in silico AI approaches, to reduce and replace animal testing in preclinical safety studies. This roadmap acknowledges that "animal models on drug safety and efficacy not only present ethical concerns but are also inadequately representative of human response," noting that more than 90% of animal-tested drugs do not reach the market. The FDA intends to create comprehensive, open-access databases of animal and human toxicity data for AI model training and validation [153].

The UK's OpenBind consortium similarly aims to curate the world's largest open dataset of drug-protein interactions to support AI training and validation, strengthening the country's position in AI-led drug discovery. These initiatives reflect growing recognition that robust, diverse, and accessible datasets are essential for advancing AI applications in drug development [153,154].

7. CHALLENGES AND LIMITATIONS

7.1 Data Quality, Availability, and Standardization

Despite significant advances, data-related challenges remain the most substantial barrier to AI/ML adoption in drug discovery. The performance of AI models is fundamentally limited by the quality, quantity, and diversity of training data [155,156].

Data scarcity affects many therapeutic areas and target classes. For rare diseases, limited biological samples and patient populations restrict available training data. Similarly, for emerging targets or novel chemical matter, insufficient historical data may exist to train robust predictive models. Data augmentation, transfer learning, and generative approaches can partially address scarcity but cannot fully compensate for limited information [157].

Data quality presents pervasive challenges. Experimental data contains measurement noise, systematic biases, and inter-laboratory variability. Assay conditions, endpoint measurements, and data reporting standards vary substantially across datasets, complicating integration. Machine learning models trained on low-quality data learn spurious correlations, limiting generalization to new experimental contexts [158,159].

Data accessibility remains limited despite growing open science initiatives. Pharmaceutical companies appropriately protect proprietary data as competitive assets, but this data remains unavailable for model training. While public databases including ChEMBL, PubChem, and BindingDB provide valuable resources, these represent only a fraction of total pharmaceutical data. Federated learning approaches—training models across distributed data sources without centralizing sensitive data—offer promising solutions but remain technically challenging to implement [160,161].

Data bias poses particular risks for AI applications intended to guide clinical development. Training data may overrepresent certain chemical scaffolds, target classes, or patient populations, leading models to perform poorly on underrepresented examples. Models trained on biased data risk perpetuating or amplifying existing disparities when deployed in new contexts. Careful attention to dataset composition and bias mitigation strategies is essential [162].

7.2 Model Interpretability and Trust

The "black box" nature of deep learning models continues to limit adoption in pharmaceutical contexts where understanding mechanistic basis is essential. As noted in recent literature, "clinical decisions must be founded on well-established principles. Although the conclusion is accurate, flawed reasoning is unacceptable, especially in safety-critical applications such as healthcare" [83,84].

Interpretability challenges vary across model architectures. While simple models including logistic regression and decision trees produce inherently interpretable outputs, their predictive performance typically lags behind more complex approaches. Deep neural networks, which achieve state-of-the-art performance across most tasks, remain fundamentally challenging to interpret due to their distributed, non-linear representations [163].

XAI methods provide partial solutions but face their own limitations. Feature attribution methods including SHAP and LIME identify which input features most influenced predictions, but do not explain why those features are important or whether models have learned valid chemical principles. Attention visualization in transformers highlights attended regions but does not guarantee that attention corresponds to mechanistic importance. Confidence calibration—ensuring that model uncertainty estimates reflect true prediction accuracy—remains an active research area [164,165].

Trust building requires demonstration that AI models generalize beyond training distributions and maintain performance in real-world deployment settings. Retrospective validation on held-out test data provides necessary but insufficient evidence of reliability. Prospective validation—testing model predictions with new experiments—provides stronger evidence but requires substantial investment. Establishing standardized benchmarking protocols and continuous monitoring systems will be essential for building stakeholder confidence [166].

7.3 Generalization and Domain Shift

Machine learning models assume that training and deployment data are drawn from similar distributions—an assumption frequently violated in drug discovery applications. Domain shift—differences between training data and the data on which models are deployed—remains a fundamental challenge [167,168].

Temporal shift occurs when models trained on historical data are applied to novel chemical matter or biological targets. Medicinal chemistry evolves over time, with new scaffolds, synthetic methodologies, and property profiles emerging. Models trained on compounds from one era may not generalize to compounds from subsequent eras [169].

Chemical domain shift occurs when models trained on one chemical series or target class are applied to structurally distinct compounds or different targets. Deep learning models can fail catastrophically when encountering molecular features absent from training data, even when these features are common in deployment contexts [170].

Biological domain shift occurs when models trained on one assay format, cell line, or species are applied in different biological contexts. Assay interference, batch effects, and context-dependent biology all contribute to poor generalization across experimental conditions. Validation in relevant biological systems is essential before deploying models for critical decisions [171].

Recent advances including GS-DTI have demonstrated improved generalization through careful architecture design. By leveraging pre-trained models and contrastive learning, GS-DTI achieved more than 10% improvement in MCC on drug-target pair cold start tests, showing promising generalization to unseen compounds and targets. These approaches—combining large-scale pre-training with domain-specific fine-tuning—represent a promising direction for improving generalization [78,172].

7.4 Computational Costs and Infrastructure Requirements

State-of-the-art AI models require substantial computational resources, creating barriers for academic laboratories and smaller biotechnology companies. Training large language models or deep graph neural networks typically requires multiple high-end GPUs and days to weeks of computation [173].

Foundation models including protein language models (e.g., ESM-2 with 15 billion parameters) and molecular representation models (e.g., ChemBERTa-2) demand substantial computational investment for pre-training. However, once trained, these models can be fine-tuned for specific tasks with modest additional computation—a pattern enabling broader access to state-of-the-art capabilities [174].

Infrastructure requirements extend beyond raw computing power. Managing large-scale training data, tracking model versions and hyperparameters, reproducing results across computing environments, and deploying models for production use all require substantial engineering investment. Cloud computing platforms reduce hardware barriers but introduce their own management complexities and ongoing costs [175].

Green AI considerations—reducing the environmental impact of large model training—have gained attention. Model efficiency, including pruning, quantization, and knowledge distillation approaches, can substantially reduce computational requirements while preserving predictive performance [176].

7.5 Regulatory and Ethical Considerations

The integration of AI into drug discovery raises important regulatory and ethical questions that remain incompletely addressed. Regulatory frameworks developed for traditional drug discovery may not adequately capture the unique considerations of AI-driven approaches [177].

Regulatory guidance for AI/ML in drug development remains limited, though rapidly evolving. The FDA has released discussion papers and draft guidance on topics including the use of AI for drug development, good machine learning practices, and the total product lifecycle approach for AI-enabled devices. However, substantial gaps remain, particularly regarding model validation requirements, change management for continuously learning models, and evidence standards for AI-driven predictions [153,178].

Intellectual property questions surrounding AI-generated inventions remain unresolved across major jurisdictions. Patent offices have grappled with whether AI systems can be named as inventors, who owns inventions generated by AI, and whether training data attribution creates freedom-to-operate risks. Clear legal frameworks are needed to incentivize investment while appropriately rewarding AI contributions [179].

Ethical considerations extend beyond legal compliance to questions of appropriate AI use. Bias and fairness concerns are particularly acute for AI applications affecting patient care or trial participation. Transparency about AI limitations, model uncertainty, and potential failure modes is essential for responsible deployment. Accountability frameworks specifying who is responsible when AI-guided decisions produce adverse outcomes remain underdeveloped [180,181].

8. FUTURE PERSPECTIVES

8.1 Integration of AI with Physics-Based Methods

The future of computational drug discovery lies not in choosing between AI and physics-based methods, but in their intelligent integration. Each approach offers complementary strengths: physics-based methods provide thermodynamic rigor and generalization, while AI methods offer speed and pattern recognition from large datasets [182].

AI-accelerated molecular dynamics combines deep learning with MD simulation to achieve longer timescales and greater sampling efficiency. Neural network potentials trained on quantum mechanical data enable accurate simulations at classical force field speeds. Similarly, AI prediction of slow collective variables enables enhanced sampling methods to explore conformational landscapes orders of magnitude faster than traditional approaches [183,184].

Free energy perturbation (FEP) enhanced by AI reduces the substantial computational cost of rigorous binding affinity calculations. Machine learning models predict perturbation outcomes, guide selection of intermediate states, and accelerate convergence. These hybrid approaches maintain the thermodynamic rigor of FEP while reducing time-to-results from weeks to days [185].

8.2 Automated Laboratories and Closed-Loop Discovery

The integration of AI with laboratory automation promises to create closed-loop discovery systems where computational predictions are rapidly tested experimentally, with results feeding back to improve subsequent predictions. This virtuous cycle dramatically accelerates the design-make-test-analyse (DMTA) cycle [186,187].

Robotic synthesis and testing platforms can execute experimental workflows with minimal human intervention. AI proposes compounds to synthesize, robotic systems execute chemical synthesis or biological assays, and results automatically update predictive models. These systems operate continuously, dramatically reducing the time between design and experimental feedback [188].

Self-driving laboratories represent the ultimate expression of closed-loop discovery, integrating AI planning with automated experimentation across multiple modalities. These systems formulate hypotheses, design experiments, execute them, interpret results, and plan subsequent experiments without human intervention. While still emerging, self-driving labs have demonstrated impressive results in catalyst discovery, materials science, and increasingly drug discovery [189].

8.3 Quantum Computing Applications

Quantum computing offers theoretical advantages for simulating molecular systems—the central challenge of computational drug discovery. Quantum computers naturally represent quantum mechanical systems, potentially enabling exact calculations of molecular properties, reaction mechanisms, and binding interactions that challenge classical approaches [190].

Near-term quantum applications include quantum machine learning models that may offer advantages for certain molecular property prediction tasks. Variational quantum eigensolvers and quantum approximate optimization algorithms have been applied to small molecular systems as proofs of concept. However, current quantum computers remain limited by qubit count, coherence times, and error rates [191].

Long-term potential includes exact quantum chemistry calculations for large molecular systems, simulation of enzymatic reaction mechanisms, and quantum-enhanced sampling of conformational space. Realizing this potential will require fault-tolerant quantum computers with millions of qubits, likely decades away [192].

8.4 Large Language Models and Foundation Models

Large language models (LLMs) pre-trained on vast chemical and biological corpora are transforming molecular representation learning. Models such as ChemBERTa, MolFormer, and Galactica learn rich representations that transfer effectively to diverse downstream tasks, achieving state-of-the-art performance with modest fine-tuning [113,114].

Multimodal foundation models that integrate chemical structures, protein sequences, omics data, and text offer particular promise. These models can answer natural language questions about molecular properties, generate novel molecules meeting specified criteria, and propose experimental plans. As models scale and training data expands, emergent capabilities continue to surprise the field [193].

Smaller, specialized models may ultimately prove more valuable than generic foundation models for many drug discovery tasks. Knowledge distillation, model pruning, and efficient architecture design enable powerful models that can be trained and deployed on commodity hardware, democratizing access to state-of-the-art AI [176].

8.5 Collaborative and Open Science Initiatives

Realizing the full potential of AI in drug discovery requires collaboration across sectors and open sharing of data, models, and best practices. No single organization possesses all necessary capabilities, data, and expertise [194,195].

Public-private partnerships bring together pharmaceutical companies, academic laboratories, technology providers, and non-profit organizations. Initiatives such as the Accelerating Therapeutics for Opportunities in Medicine (ATOM) consortium and the Machine Learning for Pharmaceutical Discovery and Synthesis (MLPDS) consortium have demonstrated the value of precompetitive collaboration for developing and validating AI methods [196].

Open datasets are essential for training robust, generalizable models. The FDA's planned open-access toxicity database, the UK's OpenBind drug-protein interaction dataset, and initiatives like Therapeutics Data Commons provide critical infrastructure. Continued investment in high-quality, well-curated, accessible datasets will accelerate progress across the field [153,197].

Benchmarking and reproducibility efforts establish objective standards for evaluating AI methods, enabling fair comparison and identifying promising approaches. Platforms like MoleculeNet, Therapeutics Data Commons, and TDC provide standardized datasets, evaluation protocols, and leaderboards. However, benchmarking studies show that modest methodological variations often produce larger performance differences than model architecture choices—highlighting the importance of rigorous, reproducible research practices [198].

9. CONCLUSION

The integration of Artificial Intelligence and Machine Learning into drug discovery represents a paradigm shift with profound implications for pharmaceutical research and development. The remarkable progress witnessed between 2019 and 2025 has transformed computational approaches from supporting tools to central drivers of discovery. AI/ML methods now contribute across the entire drug development continuum—from initial target identification through clinical trial optimization—demonstrating tangible impacts on timelines, costs, and success rates.

Deep learning, graph neural networks, and generative models have achieved particular success in molecular property prediction, virtual screening, and de novo drug design. The breakthrough of AlphaFold for protein structure prediction has democratized access to structural information, enabling structure-based design for thousands of previously intractable targets. Graph-based approaches including GS-DTI have demonstrated superior performance in drug-target interaction prediction, with improved generalization to novel

compounds and targets. Generative models are beginning to realize the promise of designing entirely new chemical entities optimized for complex property profiles.

However, substantial challenges remain to be addressed. Data quality, availability, and standardization continue to limit model performance and generalization. The black-box nature of deep learning models impedes interpretation and trust, though emerging explainable AI methods offer partial solutions. Domain shift—the failure of models trained on historical data to generalize to new chemical and biological contexts—remains a fundamental challenge requiring continued methodological innovation. Regulatory frameworks for AI-driven drug development remain incomplete, creating uncertainty for implementation.

The future of AI in drug discovery lies in intelligent integration: combining AI with physics-based methods, integrating computational predictions with automated experimentation, and fostering collaboration across sectors through open science initiatives. As foundation models scale, self-driving laboratories mature, and quantum computing emerges, the capabilities of AI-driven drug discovery will continue to expand. Realizing the full potential requires not only technical advances but also continued attention to ethical considerations, including bias, transparency, and accountability.

Table 2: Summary of AI Applications Across Drug Discovery Stages

Stage	AI Applications	Key Techniques	Impact
Target Identification	Target-disease association; Druggability prediction	GNNs; Ensemble learning; Network analysis	Reduced false targets; Novel target discovery
Hit Discovery	Virtual screening; De novo generation	Deep learning; GANs; Transformers	1000x faster screening; Novel scaffolds
Lead Optimization	Multi-property optimization; ADMET prediction	RL; GNNs; QSAR	Reduced attrition; Improved drug-likeness
Preclinical Safety	Toxicity prediction; Organ modelling	Deep learning; In silico models	Reduced animal testing; Earlier safety signals
Clinical Trials	Patient recruitment; Adaptive design; Response prediction	NLP; ML classifiers; RWD integration	Faster enrolment; Smaller sample sizes
Drug Repurposing	Indication prediction; Literature mining	LLMs; Network analysis; Knowledge graphs	Reduced timelines; Established safety

The ultimate measure of success for AI in drug discovery will be the number of safe, effective medicines reaching patients and the speed and cost with which they are developed. Early evidence suggests that AI can deliver on this promise—compounds discovered and optimized through AI have entered clinical trials, and AI-enabled efficiencies are being realized across the pharmaceutical industry. By emphasizing transparent methodologies, robust validation, and ethical frameworks for responsible AI integration, the research community can unlock AI's full potential to create safer, more effective, and more accessible medicines. The transformation of pharmaceutical research and development through AI has begun; the coming decade will reveal its full scope and impact.

Acknowledgments:

NA

Authors' contributions:

Sole Authership

Data availability statement:

NA

Disclosure statement:

NA

Ethical approval:

NA

References:

[1] J. A. DiMasi, H. G. Grabowski, and R. W. Hansen, "Innovation in the pharmaceutical industry: New estimates of R&D costs," Journal of Health Economics, vol. 47, pp. 20-33, 2016.

[2] D. Sun, W. Gao, H. Hu, and S. Zhou, "Why 90% of clinical drug development fails and how to improve it?" Acta Pharmaceutica Sinica B, vol. 12, no. 7, pp. 3049-3062, 2022.

[3] O. Wouters, M. McKee, and J. Luyten, "Estimated Research and Development Investment Needed to Bring a New Medicine to Market, 2009-2018," JAMA, vol. 323, no. 9, pp. 844-853, 2020.

[4] D. C. Swinney and J. Anthony, "How were new medicines discovered?" Nature Reviews Drug Discovery, vol. 10, no. 7, pp. 507-519, 2011.

[5] J. P. Hughes, S. Rees, S. B. Kalindjian, and K. L. Philpott, "Principles of early drug discovery," British Journal of Pharmacology, vol. 162, no. 6, pp. 1239-1249, 2011.

[6] S. J. Y. Macalino, V. Gosu, S. Hong, and S. Choi, "Role of computer-aided drug design in modern drug discovery," Archives of Pharmacal Research, vol. 38, no. 9, pp. 1686-1701, 2015.

[7] G. Sliwoski, S. Kothiwale, J. Meiler, and E. W. Lowe, "Computational methods in drug discovery," Pharmacological Reviews, vol. 66, no. 1, pp. 334-395, 2014.

[8] A. Lavecchia, "Machine-learning approaches in drug discovery: methods and applications," Drug Discovery Today, vol. 20, no. 3, pp. 318-331, 2015.

[9] H. Chen, O. Engkvist, Y. Wang, M. Olivecrona, and T. Blaschke, "The rise of deep learning in drug discovery," Drug Discovery Today, vol. 23, no. 6, pp. 1241-1250, 2018.

[10] K. K. Mak and M. R. Pichika, "Artificial intelligence in drug development: present status and future prospects," Drug Discovery Today, vol. 24, no. 3, pp. 773-780, 2019.

[11] D. Paul, G. Sanap, S. Shenoy, D. Kalyane, K. Kalia, and R. K. Tekade, "Artificial intelligence in drug discovery and development," Drug Discovery Today, vol. 26, no. 1, pp. 80-93, 2021.

[12] L. G. Klingmann, "The development of new drugs," Journal of Clinical Pharmacology, vol. 42, no. 9, pp. 948-951, 2002.

[13] D. E. Clark, "What has computer-aided drug design taught us about drug design?" Expert Opinion on Drug Discovery, vol. 4, no. 4, pp. 345-349, 2009.

[14] I. Kola and J. Landis, "Can the pharmaceutical industry reduce attrition rates?" Nature Reviews Drug Discovery, vol. 3, no. 8, pp. 711-715, 2004.

[15] M. J. Waring, J. Arrowsmith, A. R. Leach, et al., "An analysis of the attrition of drug candidates from four major pharmaceutical companies," Nature Reviews Drug Discovery, vol. 14, no. 7, pp. 475-486, 2015.

[16] "Machine learning & deep learning tools in pharmaceutical sciences: A comprehensive review," Intelligent Pharmacy, vol. 3, no. 3, pp. 167-180, 2025.

[17] R. Gupta, D. Srivastava, M. Sahu, S. Tiwari, R. K. Ambasta, and P. Kumar, "Artificial intelligence to deep learning: machine intelligence approach for drug discovery," Molecular Diversity, vol. 25, no. 3, pp. 1315-1360, 2021.

[18] X. Yang, Y. Wang, R. Byrne, G. Schneider, and S. Yang, "Concepts of artificial intelligence for computer-assisted drug discovery," Chemical Reviews, vol. 119, no. 20, pp. 10520-10594, 2019.

[19] S. Kumar, K. Bala, and M. S. Yadav, "Artificial Intelligence in Drug Discovery: Integrating Algorithms, Data, and Innovation for Accelerated Therapeutics," International Journal of Pharmaceutical Sciences, vol. 3, no. 11, pp. 3369-3373, 2025.

[20] A. Lavecchia and C. Di Giovanni, "Virtual screening strategies in drug discovery: a critical review," Current Medicinal Chemistry, vol. 20, no. 23, pp. 2839-2860, 2013.

[21] T. I. Oprea, "Current trends in lead discovery: Are we looking for the appropriate properties?" Molecular Diversity, vol. 6, no. 1, pp. 1-10, 2003.

[22] B. K. Shoichet, "Virtual screening of chemical libraries," Nature, vol. 432, no. 7019, pp. 862-865, 2004.

[23] I. D. Kuntz, "Structure-based strategies for drug design and discovery," Science, vol. 257, no. 5073, pp. 1078-1082, 1992.

[24] A. M. Davis, S. J. Teague, and G. J. Kleywegt, "Application and limitations of X-ray crystallographic data in structure-based ligand and drug design," Angewandte Chemie International Edition, vol. 42, no. 24, pp. 2710-2736, 2003.

[25] A. C. Anderson, "The process of structure-based drug design," Chemistry & Biology, vol. 10, no. 9, pp. 787-797, 2003.

[26] N. Eswar, B. John, N. Mirkovic, et al., "Tools for comparative protein structure modeling and analysis," Nucleic Acids Research, vol. 31, no. 13, pp. 3375-3380, 2003.

[27] D. T. Jones, "Protein structure prediction in the post-genomic era," Current Opinion in Structural Biology, vol. 10, no. 3, pp. 371-379, 2000.

[28] J. Moult, "A decade of CASP: progress, bottlenecks and prognosis in protein structure prediction," Current Opinion in Structural Biology, vol. 15, no. 3, pp. 285-289, 2005.

[29] J. Jumper, R. Evans, A. Pritzel, et al., "Highly accurate protein structure prediction with AlphaFold," Nature, vol. 596, no. 7873, pp. 583-589, 2021.

[30] M. Baek, F. DiMaio, I. Anishchenko, et al., "Accurate prediction of protein structures and interactions using a three-track neural network," Science, vol. 373, no. 6557, pp. 871-876, 2021.

[31] J. Liang, H. Edelsbrunner, and C. Woodward, "Anatomy of protein pockets and cavities: measurement of binding site geometry and implications for ligand design," Protein Science, vol. 7, no. 9, pp. 1884-1897, 1998.

[32] P. Schmidtke, C. Souaille, F. Estienne, N. Baurin, and R. T. Kroemer, "Large-scale comparison of four binding site detection algorithms," Journal of Chemical Information and Modeling, vol. 50, no. 12, pp. 2191-2200, 2010.

[33] G. M. Morris and M. Lim-Wilby, "Molecular docking," Methods in Molecular Biology, vol. 443, pp. 365-382, 2008.

[34] M. L. Lamb, "Software review: Soft docking and scoring," Drug Discovery Today, vol. 4, no. 3, pp. 147-148, 1999.

[35] S. B. Nabuurs, M. Wagener, and J. de Vlieg, "A flexible approach to induced fit docking," Journal of Medicinal Chemistry, vol. 50, no. 26, pp. 6507-6518, 2007.

[36] P. H. Colizzi, R. V. Guido, and A. D. Andricopulo, "Advances in molecular modeling and docking as a tool for modern drug discovery," Current Medicinal Chemistry, vol. 17, no. 4, pp. 327-346, 2010.