



Accelerating Pharmaceutical R&D: The Role of Generative Artificial Intelligence in Modern Drug Discovery

Abhijat Mishra¹, Saurabh Sarkar², Radha Raman Chandan³, Shashi Bhushan⁴ and Basu Dev Shivahare^{1,*}

¹School of Computer Science and Engineering, Galgotias University, Greater Noida 203201, India

²Chicory AI, Greater Seattle, United States

³Department of Computer Science, School of management & sciences, Varanasi 221011, India

⁴Department of Computing, Universiti Teknologi PETRONAS, Seri Iskandar, Malaysia

Abstract

Exorbitant expenses, lengthy development periods, and a high incidence of drug candidate attrition plague the conventional pharmaceutical R&D pipeline—a problem sometimes referred to as “Eroom’s Law.” By radically reorganizing the discovery process, generative artificial intelligence (AI), which has emerged as a transformational force, promises to buck this tendency. Through data synthesis on key performance metrics, this review offers a thorough analysis of the effects of AI-enhanced methodologies. We explore how a new set of tools is changing the paradigm from experimental screening to in silico design. These tools include graph neural networks (GNNs)—a class of neural architectures that operate directly on graph-structured data by recursively aggregating information from neighbouring nodes—for molecular modelling. Additionally, large language

models (LLMs)—Transformer-based neural networks trained on massive text corpora that learn contextual representations of biological sequences and scientific literature—are revolutionizing target identification. According to our analysis, integrating AI results in previously unheard-of benefits: clinical success rates for AI-discovered candidates are expected to rise from a baseline of 7.9% to as high as 90%, costs are predicted to be cut by 45–80%, and early-stage discovery timelines are compressed by up to 62.5% (e.g., reducing target-to-lead time from 24 to 9 months). These improvements stem from a sharp rise in molecular-level prediction accuracy. We come to the conclusion that generative AI is a crucial tool for accelerating the development of new treatments, allowing for a quicker, more economical, and more successful strategy that will characterize the next phase of pharmaceutical innovation.

Keywords: generative AI, drug discovery, target identification, machine learning, AlphaFold,



Submitted: 06 April 2025

Accepted: 03 September 2025

Published: 27 September 2025

Vol. 1, No. 2, 2025.

10.62762/BISH.2025.789201

*Corresponding author:

✉ Basu Dev Shivahare

basuiimt@gmail.com

Citation

Mishra, A., Sarkar, S., Chandan, R. R., Bhushan, S., & Shivahare, B. D. (2025). Accelerating Pharmaceutical R&D: The Role of Generative Artificial Intelligence in Modern Drug Discovery. *Biomedical Informatics and Smart Healthcare*, 1(2), 67–78.



© 2025 by the Authors. Published by Institute of Central Computation and Knowledge. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>).

pharmaceutical research, computational chemistry.

1 Introduction

Ironically, one of the most important and ineffective activities in contemporary research is the creation of new medications. With startling failure rates, the conventional drug research and development pipeline is a protracted, difficult, and expensive procedure. It can take 10 to 15 years from the first spark of a biological theory to the final approval of a marketable medicine, and the expenses of each licensed therapy sometimes surpass \$2.5 billion [1, 17]. Known as “Eroom’s Law” (Moore’s Law spelled backward), this phenomenon, in which the cost of creating a new medication almost doubles every nine years, highlights a decades-long pattern of falling R&D productivity despite enormous scientific and technical advancements. The main causes of this inefficiency are the high rates of drug candidate attrition; more than 90% of medications that start clinical trials end up not receiving regulatory clearance, frequently as a result of unexpected toxicity or ineffectiveness. The preclinical phases—target identification, hit discovery, and lead optimization—which have traditionally depended on a mix of serendipity, brute-force HTS, and laborious, repeated chemical synthesis, are where the bottlenecks are most noticeable. In addition to being lengthy and expensive, this paradigm drastically restricts the amount of chemical space that can be investigated, leaving large areas of potentially therapeutic compounds unexplored. With an estimated 10^{60} molecules, the chemical universe is so vast that it is impossible to conduct a thorough experimental investigation. Despite their strength, HTS campaigns are only able to survey a small portion of this space (usually 10^6 to 10^7 compounds), and the molecules they find frequently exhibit subpar drug-like qualities that need time-consuming and expensive improvement. This conventional, empirical method’s inherent inefficiency and expense may be explained by the fact that it is essentially a filtering process, a hunt for a needle in an impossible haystack.

The scientific community has resorted more and more to computational techniques in response to this crisis in R&D output. Foundational methods like molecular docking and QSAR modeling were introduced during the CADD era, offering the first indications of a more rational, data-driven strategy. However, the accuracy of their underlying physical models and processing capacity frequently placed limitations on these early approaches. The introduction of AI,

and more especially the development of generative models and deep learning, marked the beginning of the real paradigm change. Generative AI differs from discriminative algorithms, which only categorize or forecast characteristics based on available data. Rather, these models may produce completely new, feasible, and optimal outputs by learning the fundamental patterns and principles of complicated data distributions, such as the grammar of molecular structures or the language of protein sequences. Generative models can create new molecules with unique properties, predict the three-dimensional structure of proteins from their amino acid sequence with remarkable accuracy, and find promising therapeutic targets from intricate biological networks by using advanced algorithms to learn from massive datasets of biological and chemical data. As seen in Figure 1, this signifies a shift from a screening paradigm to a design paradigm.

A thorough examination of the revolutionary effects of generative AI in the field of drug development is provided in this review. To measure the advancements in the three crucial R&D axes of timelines, costs, and success rates, we combine data from influential scholarly works, technical reports, and comparative industry analyses. Our results demonstrate a significant speedup in the computationally-driven early stages of discovery. According to some reports, AI integration results in a significant increase in clinical success rates from a historical baseline of 7.9% to as high as 90% for candidates found by AI. It also reduces the target-to-lead timeframe by 62.5% (from 24 months to 9 months), as well as the associated early-stage costs by 45% [1]. By combining and contextualizing these metrics, this work offers a comprehensive picture of how particular AI technologies—such as LLMs like BioGPT [3], GNNs for molecular modeling, reinforcement learning for property optimization [4], and AlphaFold’s ground-breaking protein-folding capabilities [5]—are working together to transform medicine and provide a potent remedy for the problems posed by Eroom’s Law.

2 Related Work

Although the use of AI in medicine has a long history, a recent explosion of innovation has been sparked by the combination of large datasets (genomic, proteomic, and chemical), innovative deep learning architectures, and scalable computer technology. QSAR was the focus of early machine learning applications in drug discovery, which created regression models




Figure 1. Note that “Figure” is abbreviated. There is a period after the figure number, followed by two spaces. It is good practice to explain the significance of the figure in the caption.

to forecast a compound's activity based on its physicochemical descriptors. Although helpful, these models' prediction capacity and generalizability to new chemical scaffolds were frequently constrained. A new class of tools was brought about by the deep learning revolution. Originally created for image recognition, CNNs have been modified to examine medical images and 3D molecular structures. Recurrent neural networks (RNNs) and their contemporary offspring, Transformers—which were first created for natural language processing, as pioneered by models like BERT for pre-training deep bidirectional transformers [14]—were repurposed to treat proteins as biological language and molecules as sentences (using representations like SMILES) in order to generate new sequences and structures, as demonstrated in transformer-based models for antiviral drug design [9].

Generative models dominate the state-of-the-art at the moment. Since they may learn a condensed “latent space” of chemical properties from which new molecules can be sampled, VAEs and GANs have been extensively investigated for de novo molecular

design [6]. By exploring this latent space, VAEs, for example, are excellent at producing a smooth, continuous representation that enables the incremental tuning of molecular attributes. Although GANs may produce very realistic and unique molecules through their adversarial training process, they may also be vulnerable to training instability [7]. Diffusion models have demonstrated remarkable potential in producing high-quality 3D protein structures and molecular geometries in more recent times [8]. These algorithms generate complicated, high-dimensional data by methodically introducing noise to the input and then training a neural network to reverse the process. This approach has been shown to be highly reliable and effective. Since molecules are naturally represented as graphs of atoms (nodes) and bonds (edges), GNNs have become the industry standard for working with molecular data. This makes it possible to extract features and predict molecular characteristics and interactions more easily and effectively.

DeepMind's AlphaFold, a system that use a deep learning network to predict a protein's 3D structure from its 1D amino acid sequence with accuracy

comparable to experimental approaches, may have been the most significant innovation [5]. Its success comes from its ability to solve a 50-year-old biological grand challenge by using an attention-based network to reason about the spatial relationships between amino acid pairs and interpreting co-evolutionary data within multiple sequence alignments. For many proteins whose structures were previously unknown, this has opened the door to structure-based medication creation. This feature is further expanded in the most recent version, AlphaFold 3, which offers an even more potent tool for target validation and drug discovery by modeling intricate interactions between proteins, nucleic acids, and ligands [5]. In order to evaluate their synergistic effect on the key performance metrics that characterize success in the pharmaceutical sector, our study expands upon the fundamental work in these separate but related fields.

3 Methodology

This review synthesizes and analyzes the effects of generative AI on drug development using a methodical, multifaceted approach. The strategy consists of two steps: first, a thorough literature analysis was carried out to compile information from technical reports, industry white papers, and peer-reviewed publications. Second, in order to organize the collected data and make insightful inferences on the performance variations between conventional and AI-enhanced processes, a framework for comparative analysis was developed.

3.1 Literature Review and Data Aggregation

A methodical, multifaceted approach is used in this study to summarize and examine how generative AI affects drug development. There are two parts to the strategy: first, a thorough literature review was done to compile information from technical reports, industry white papers, and peer-reviewed articles. To organize the collected data and make insightful inferences on the performance variations between conventional and AI-enhanced processes, a comparative analysis framework was designed.

3.2 Comparative Framework and Metrics

In this analysis, we identified the key stages of the drug discovery process, alongside the metrics used to evaluate these stages, ensuring a systematic approach for comparing the impact of artificial intelligence (AI) on the process. The framework focuses on critical aspects such as timelines, costs, and success rates,

which are essential for assessing the effectiveness of AI in drug development.

The timeline of the drug discovery process is measured in months or years and is typically divided into distinct phases. The first phase, Target Identification and Validation, focuses on identifying a biological molecule, such as a protein, that plays a causal role in a disease. This is followed by Hit Discovery, where large libraries of chemical compounds are screened to find those that interact with the identified target. Once a hit is found, the Hit-to-Lead phase begins, which involves optimizing the chemical properties of the hits to generate compounds with higher potency and drug-like characteristics. In the subsequent Lead Optimization phase, further refinement is carried out to enhance efficacy, safety, and pharmacokinetic properties. Finally, the drug proceeds through the Preclinical and Clinical Phases, where it undergoes testing in animals and humans to assess safety, efficacy, and overall potential for therapeutic use.

The cost metric is measured either in absolute monetary terms, such as USD, or as a percentage of savings. AI's most direct influence on cost occurs during the early stages of drug development, especially in the preclinical phase. Through AI-driven optimizations, such as reducing the number of compounds that need to be synthesized and eliminating the need for costly high-throughput screening (HTS) campaigns, significant cost savings can be achieved. Therefore, our analysis focuses on these early-stage cost reductions. Additionally, when calculating the cost of failures, it is important to consider both immediate cost savings and the decrease in capitalized costs that would result from late-stage drug development failures, which are often costly.

Success and accuracy are multi-dimensional metrics that are crucial in evaluating the performance of AI in the drug discovery process. One of the key metrics is the Clinical Success Rate, which measures the likelihood that a drug, after entering clinical trials, will ultimately receive regulatory approval. Alongside this, Predictive Accuracy is assessed using standard machine learning metrics. For classification tasks, such as predicting toxicity, the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) is commonly used. For regression tasks, such as predicting binding affinity, Root Mean Square Error (RMSE) is employed. Additionally, the Global Distance Test (GDT_TS) is the gold standard used to

evaluate the accuracy of protein structure predictions, an essential part of drug development. Another important metric is the Hit Rate, which reflects the percentage of compounds tested in the screening phase that exhibit the desired level of activity, providing an indication of the efficiency and success of the screening process.

3.3 AI Technologies Analyzed: A Deeper Dive

The "AI-Enhanced" workflow is not a singular, unified process but rather a collection of specialized technologies, each playing a critical role in advancing drug discovery. These technologies leverage various AI models to optimize different stages of the drug development process, from molecular design to protein structure prediction, and network-based target discovery.

A variety of generative models are explored for molecular design in this review. Variational autoencoders (VAEs) are particularly useful for optimization; however, they have limitations. Although VAEs can generate molecular structures, they may occasionally produce invalid structures because they learn a continuous latent representation of molecules, which is not always ideal for generating chemically valid compounds. Generative adversarial networks (GANs), though more challenging to train, are highly effective at generating lifelike molecules. This is achieved through a two-player game between a discriminator, which tries to distinguish real from generated molecules, and a generator, which creates the molecules. More controlled and target-aware generation is made possible by transformer-based models, which utilize the attention mechanism to capture long-range dependencies in sequences—such as those found in molecular structures. Finally, diffusion models have shown remarkable success in generating intricate 3D molecular conformations. These models work by gradually introducing noise to data and learning to reverse the process, ultimately resulting in highly refined molecular structures.

A significant area of focus in AI-driven drug discovery is protein structure prediction, with AlphaFold standing out as a key player. This deep neural network approach predicts the 3D coordinates of a protein's atoms with unparalleled accuracy. AlphaFold combines multiple sequence alignment, co-evolutionary analysis, and a unique attention-based architecture to achieve its remarkable performance. This breakthrough makes structure-based drug creation a universal strategy, extending its applicability

from just a few thousand proteins with known structures to millions of proteins. This development has profound implications for drug discovery, as it allows researchers to target proteins that were previously beyond reach.

Network-based target discovery involves modeling complex biological systems as networks, such as networks of protein-protein interactions, using graph neural networks (GNNs). These networks learn intricate system-level patterns through a message-passing mechanism, where nodes (proteins) aggregate information from their neighbors in an iterative process. This ability enables GNNs to identify crucial nodes within these networks that could serve as promising therapeutic targets. The insights derived from GNNs are invaluable for discovering novel drug targets and improving therapeutic strategies.

In addition to generative models, predictive models for ADMET (Absorption, Distribution, Metabolism, Excretion, and Toxicity) characteristics play a pivotal role in the AI-driven workflow. These models are typically ensemble methods, such as Random Forest or XGBoost, which combine the predictions of multiple models to enhance robustness and accuracy. Trained on large datasets of experimental ADMET data, these models are capable of predicting crucial features, including solubility, permeability, metabolic stability, and various forms of toxicity. By predicting these parameters *in silico*, these models enable the early deselection of compounds that are likely to fail in later stages, significantly accelerating the drug development process and reducing costs.

3.4 Mathematical Framework of Core AI Technologies

In order to offer more in-depth technical understanding, we briefly describe the mathematical underpinnings of the main generative models that are being explored.

- **Variational Autoencoders (VAEs):** VAEs are trained by maximizing the Evidence Lower Bound (ELBO), which balances two objectives: reconstructing the input data and regularizing the latent space. The loss function L_{VAE} for a data point x is:

$$\begin{aligned} \text{Variational Autoencoder (VAE)} \\ = E_q(z|x) [\log p(x|z)] \\ \quad \log p(x|z)] \\ - D_{KL}(q(z|x) \parallel p(z)) \end{aligned} \quad (1)$$

Here, the first term is the reconstruction loss, encouraging the decoder $p(x|z)$ to accurately rebuild the input x from the latent variable z . The second term is the Kullback-Leibler (KL) divergence, which regularizes the encoder's output distribution $q(z|x)$ to be close to a prior distribution $p(z)$ (typically a standard normal distribution), ensuring a smooth and useful latent space for generation.

- **Generative Adversarial Networks (GANs):** A Generator (G) and a Discriminator (D), two rival networks trained in a minimax game, make up a GAN. The discriminator attempts to separate real data (x) from synthetic data, while the generator generates synthetic data ($G(z)$) from random noise (z). $V(D, G)$, their objective function, is:

$$\begin{aligned} &< \min_G < \max_D \\ V(D, G) = &E_x / \text{simp_data}[\log D(x)] \\ &+ \\ &E_z / \sin p_{-z} \sqrt{\log(1 - D(G(z)))} \end{aligned} \quad (2)$$

The discriminator D works to maximize this function (correctly labeling real and fake), while the generator G works to minimize it (fooling the discriminator), driving the system towards generating highly realistic data.

- **Graph Neural Networks (GNNs):** GNNs aggregate information from their neighbors iteratively to learn representations of nodes in a graph. The following is an expression for a general update rule for the feature vector h_v of a node v at layer $k + 1$:

$$\begin{aligned} h_v^{(k+1)} = &\text{UPDATE}^{(k)} \left(h_v^{(k)}, \right. \\ &\left. \text{AGGREGATE}^{(k)} \left(\{h_u^{(k)} : u \in \mathcal{N}(v)\} \right) \right) \end{aligned} \quad (3)$$

where AGGREGATE (such as sum or mean) and UPDATE (such as a neural network layer) are learnable functions and $\mathcal{N}(v)$ is the collection of node v 's neighbors. GNNs can capture intricate topological information that is essential for simulating molecular interactions thanks to this message-passing mechanism.

4 Comparative Analysis: Performance and Impact

Significant, measurable gains are obtained in all examined domains when generative AI is incorporated

into the drug discovery process. The findings, which are compiled from a large number of studies and reports, show that AI-enhanced techniques are clearly and frequently dramatically superior, especially in the critical, resource-intensive early phases of discovery.

4.1 Timeline Acceleration Across the Discovery Pipeline

AI has the potential to dramatically compress research and development (R&D) timelines, especially in the early stages of the drug discovery pipeline. By replacing labor-intensive, sequential, and often manual experimental procedures with rapid, parallel, and automated in silico computations, AI significantly accelerates key processes, enhancing both efficiency and cost-effectiveness.

One of the most notable areas where AI has made an impact is in Target Identification, which has traditionally been a lengthy and resource-intensive phase. Historically, the process of identifying and validating a new target could take two to three years of exhaustive biological research. However, AI has drastically shortened this timeline, achieving a reduction of up to 81.5%. Large language models (LLMs) [10] such as Med-PaLM 2 and BioGPT are now capable of quickly searching through and synthesizing vast amounts of biological literature to suggest new target-disease connections. This capability accelerates the discovery process by orders of magnitude. Additionally, graph neural networks (GNNs), which analyze protein-protein interaction networks, can identify crucial nodes in disease pathways within just a few days. As a result, the typical target identification phase has been reduced to just three to six months, a remarkable improvement over traditional timelines.

The Hit Discovery and Lead Generation phase has also seen significant time savings. In the traditional workflow, high-throughput screening (HTS) of millions of compounds followed by a lengthy hit-to-lead optimization process typically spans around 24 months. AI has transformed this process into a more efficient and rational design approach. Rather than conducting extensive screening, generative models are employed to design a small, targeted library of a few hundred novel molecules, each optimized specifically for the identified target. Combined with AI-powered virtual screening and ADMET prediction tools, the target-to-lead timeline has been reduced by 62.5%, now taking approximately 9 months. This AI-driven process not only accelerates

discovery but also ensures a more focused and precise selection of lead candidates.

The acceleration of these early-stage processes has had a substantial impact on the overall drug development timeline. With candidates moving through the pipeline faster, more promising drug candidates can enter clinical trials sooner. However, the clinical phases of development, constrained by regulatory restrictions and the biological complexities of human trials, have not seen as significant reductions in time. Nevertheless, the front-end acceleration is crucial for identifying better candidates, enabling the clinical phases to begin with higher-quality compounds, thus improving the overall chances of success.




Figure 2. Comparison of traditional and AI-enhanced drug discovery metrics across time, cost, and success rate.

As shown in Figure 2, the comparison between traditional and AI-enhanced drug discovery metrics highlights the substantial improvements in time, cost, and success rate that AI has brought to the early stages of drug development. While clinical trials remain relatively unaffected by these technological advances, the time and cost savings in earlier phases can dramatically improve the efficiency of the entire drug discovery process.

4.2 Economic Impact and Cost Reduction

AI’s time-saving capabilities have a direct and profound impact on cost reductions in pharmaceutical research and development. The maxim "time is money" holds particular significance in this field, as the costs associated with maintaining large research teams and laboratory facilities are substantial. AI’s ability to accelerate various stages of drug discovery helps mitigate these costs, especially during the early phases of the process.

During the early stages of drug development, AI systems contribute to cost savings of 45-80%. One of the key factors in these savings is the reduction in the number of compounds that require physical synthesis and testing. AI-driven optimization allows researchers to focus on a smaller, more targeted set of compounds, which significantly reduces the resources allocated to large-scale high-throughput screening (HTS) efforts. Furthermore, AI enhances the accuracy of ADMET prediction, enabling the prediction of absorption, distribution, metabolism, excretion, and toxicity properties in silico. This not only eliminates the need for extensive animal testing but also reduces the costs associated with preclinical studies, which are traditionally both costly and ethically challenging. Some evaluations suggest that computational discovery phases can achieve even greater savings, ranging from 60 to 80 percent. These reductions in resource usage and time investment make a significant impact on the overall cost structure of drug development, particularly in the early stages.

As shown in Figure 3, AI-enhanced drug discovery methods lead to considerable improvements in both time and cost metrics when compared to traditional methods. The figure illustrates how AI’s impact accelerates the drug discovery timeline while simultaneously lowering associated costs, making the process more efficient and cost-effective.

Beyond the early phases, overall development costs

Table 1. Comparative Analysis of Key Drug Discovery Metrics: Traditional vs. AI-Enhanced Methods.

Metric	Traditional Method	AI-Enhanced Method	Improvement Factor	Metric
Timeline				
Target-to-Lead Time (months)	24	9	2.7x Faster	Target-to-Lead Time (months)
Cost				
Early Stage Cost Savings (%)	~0%	45-80%	Significant Savings	Early Stage Cost Savings (%)
Success/ Accuracy				
Clinical Success Rate (%)	7.90%	up to 90%	>10x Increase	
Protein Structure Accuracy (GDT_TS)	30-50%	>90%	2-3x Increase	
Target ID Accuracy (%)	35%	90%	>2.5x Increase	
Virtual Screening Hit Rate (%)	1.50%	up to 70%	>45x Increase	




Figure 3. Comparative Performance: AI-Enhanced vs Traditional Drug Discovery Methods.

are also significantly affected by AI. While clinical trials still account for the majority of the budget, the efficiency gains made during the earlier stages have a substantial impact on later stages. Specifically, AI reduces the “cost of failure”—the money spent on candidates that ultimately fail during clinical trials. By improving the quality of candidates that progress through trials, AI ensures that fewer resources are wasted on compounds that do not meet the required standards. According to estimates, this contributes to a reduction in the average capitalized cost per approved drug, decreasing from approximately \$2.5 billion to \$1.75 billion, resulting in savings of around \$750 million per successful drug.

The comparative analysis of key metrics between traditional and AI-enhanced drug discovery methods is further summarized in Table 1. The table highlights the significant improvements in metrics such as target-to-lead time, early-stage cost savings, clinical success rate, protein structure accuracy, and hit rates in virtual screening. For example, the target-to-lead time is reduced by over 2.7 times, and early-stage cost savings can range from 45% to 80%. Additionally, AI-driven methods lead to a greater than 10x increase in clinical success rates and a dramatic improvement in the accuracy of protein structure predictions and virtual screening hit rates.

4.3 Enhancements in Predictive Accuracy and Success Rates

Artificial intelligence (AI) is not only accelerating drug discovery but also improving the quality and likelihood of success of medication candidates [12]. While speed and cost reduction are crucial, AI’s most significant contributions may lie in enhancing the accuracy of predictions at each stage of the molecular

design process, ultimately improving the chances of successful drug development.

One area where AI has made considerable strides is protein structure prediction. AI models like AlphaFold have revolutionized this aspect by achieving greater than 90% accuracy as measured by the GDT_TS score. In comparison, traditional methods such as homology modeling typically reach only 30-50% accuracy. The high-resolution 3D maps produced by AlphaFold are invaluable for rational drug design, enabling the creation of compounds with vastly improved binding affinities and specificities. A notable example of AlphaFold’s impact was in the prediction of the SARS-CoV-2 main protease (Mpro) structure, which was released on March 5, 2020. Within just three weeks, structure-enabled virtual screening campaigns (e.g., Diamond Light Source XChem fragment screen) identified potent non-covalent inhibitors that progressed to animal testing in less than 90 days, further exemplified by its utility in broader viral research [13]. This highlights AI’s ability to rapidly accelerate the identification of effective drug candidates, especially in response to emerging infectious diseases.

AI also brings substantial improvements to target identification and validation. By analyzing sequence and network data, AI models are able to accurately differentiate between druggable and non-druggable targets, achieving accuracy rates of over 90%. In contrast, traditional, target-centric techniques have shown significantly lower success rates, highlighting the efficiency of AI in identifying viable targets earlier in the process. This ability to prioritize the most promising targets is essential for streamlining drug discovery and increasing the likelihood of success.

In the realm of ADMET (Absorption, Distribution, Metabolism, Excretion, and Toxicity) prediction, ensemble machine learning models are now routinely achieving AUC-ROC scores ranging from 0.85 to 0.95 across various toxicity and pharmacokinetic tasks. This level of predictive accuracy allows researchers to filter out problematic compounds well before they undergo costly and time-consuming experimental testing. As a result, AI enables a more reliable and efficient drug discovery pipeline by preemptively identifying candidates with unfavorable characteristics.

As shown in Figure 4, AI-driven improvements in predictive accuracy and efficiency contribute to a significant reduction in timelines across drug discovery




Figure 4. Timeline Reduction by AI Across Drug Discovery Phases.

phases. These improvements not only optimize the development process but also enhance the quality of candidates that progress to clinical testing.

The most important metric reflecting these advancements is the clinical success rate, which has shown a notable increase due to AI's ability to improve the predictive power of drug candidates. Historically, the success rate of a new drug entering Phase I trials has been dismal, with only 7.9% of drugs achieving success. However, AI-driven methods have led to a prospective study of 158 AI-designed compounds, which revealed an overall success rate of 25% (95% confidence interval of 18-32%), representing a 3.2-fold improvement over the historical average [2]. This success rate is achieved by ensuring that only the most promising, well-optimized compounds, with a high predicted safety and efficacy profile, progress to human testing. While this figure is based on a relatively small number of AI-designed drugs that have advanced through trials, it represents a revolutionary leap in drug development. Even with a more cautious estimate, suggesting a 3-6x improvement (reaching a 24-48% success rate), AI's impact on the clinical success rate could dramatically alter the economic model of the pharmaceutical industry, making the drug discovery process more efficient and cost-effective.

5 Discussion

Although the quantitative findings provide a convincing picture of change, a more thorough examination of their consequences is necessary. Generative AI has a profound influence on the

scientific method and financial model of drug discovery, going beyond simple enhancements. The results are interpreted, the clinical and scientific implications are discussed, and the major obstacles and constraints that need to be overcome in order to fully utilize this technology's potential are examined in this part.

5.1 Interpretation of Findings: A Paradigm Shift in R&D

The information makes it abundantly evident that AI has a significant influence on the discovery process. AI's capacity to change drug development from an experimental screening method to an in silico design process is directly responsible for the significant time and cost savings for target selection and lead optimization. Starting with millions of chemicals and sifting them down is the standard subtractive strategy. The artificial intelligence paradigm is additive, creating a limited number of perfect chemicals from the ground up. This is a profound change in perspective. There is a causal relationship between the >45x rise in virtual screening hit rates and the >10x increase in clinical success rates; these are not independent variables. The final attrition rate is significantly lower when the initial projections and design are improved. This solves the fundamental issue with Eroom's Law, which is that the likelihood of each experiment succeeding has been decreasing rather than that experiment costs are increasing. By significantly raising the likelihood of success for every molecule that is created and tested, AI bucks this tendency.

5.2 Clinical and Scientific Implications

The advancements brought about by AI in drug discovery have far-reaching clinical and scientific implications, revolutionizing not only the development of therapies but also the strategies employed to target and treat diseases. These changes are paving the way for more effective treatments, personalized care, and faster responses to emerging health crises.

One significant advancement is the ability to tackle "undruggable" targets, which have long been a major challenge in drug discovery. Many proteins involved in disease processes have been deemed "undruggable" because they lack the distinct binding pockets necessary for conventional small-molecule drugs to interact effectively. However, with the structural insights provided by AlphaFold, combined with generative AI models, new opportunities are emerging. These AI tools can design novel compounds

or even entirely new therapeutic modalities, such as protein degraders, which can target and modify these challenging proteins. This innovation is creating entirely new treatment possibilities for diseases like cancer and neurological illnesses, offering hope for conditions that were once considered difficult or impossible to treat, through perspectives on multi-target drug discovery and design tailored for such complex diseases [15].

Additionally, AI is driving the future of personalized and precision medicine. Through the integration of patient-specific genetic or proteomic data, AI models can help develop treatments that are tailored to specific individuals or patient groups, leveraging enabling technologies that facilitate such customized therapies [11]. This ability to create customized therapies represents a significant step forward in the pursuit of precision medicine. By incorporating detailed biological data into the drug design process, AI helps ensure that treatments are not only more effective but also safer, reducing the trial-and-error approach that has traditionally characterized drug development.

Moreover, AI has proven invaluable in the context of pandemic preparedness and rapid response. The COVID-19 pandemic highlighted the urgent need for fast-track research and treatment development. AI platforms significantly reduce the time required to identify disease targets and design effective inhibitors for novel pathogens. The speed with which AI can assist in target identification and drug design makes it an essential tool for accelerating the response to future pandemics, ensuring that the pharmaceutical industry is better equipped to handle global health crises swiftly and efficiently.

5.3 Challenges and Limitations of AI in Drug Discovery

Despite the immense potential of generative AI in drug discovery, its widespread adoption is accompanied by a number of challenges and limitations that must be addressed to fully harness its capabilities. These hurdles range from data-related issues to difficulties in model interpretability and generalizability, all of which can hinder the effectiveness and reliability of AI models in the complex field of drug development.

One of the most significant challenges is data quality and scarcity. The performance of AI models is directly linked to the quality of the data on which they are trained. While large public databases exist, such as ChEMBL, these datasets often suffer from issues of

noise, diversity, and incompleteness, as evidenced by efforts toward direct deposition of bioassay data to enhance its utility [16]. For many novel targets or rare diseases, there is a lack of high-quality experimental data, which can limit the prediction power of AI models and lead to inaccurate or biased conclusions. For example, the ChEMBL-32 database contains experimental data for only about 2.2 million distinct compounds, a tiny fraction of the estimated 10^{60} drug-like chemical space. This gap significantly limits the ability of AI models to generalize to new, unexplored chemical scaffolds and compounds.

Another considerable limitation is the "black box" problem inherent in many deep learning models. These models, which drive much of the current AI revolution in drug discovery, often operate in a manner that is not transparent. As a result, it can be extremely difficult to understand how they make decisions or why they reach specific conclusions. This lack of interpretability presents a significant barrier to clinical uptake and regulatory approval, as both physicians and regulators need to not only know what a model predicts but also why it makes those predictions, particularly concerning the safety and efficacy of substances. Consequently, there is ongoing research into developing Explainable AI (XAI) techniques to improve transparency and trust in AI-driven decisions.

Generalizability is another challenge that arises in AI applications for drug discovery. A model trained on data from one biological or chemical space may not perform as well when applied to new targets or different types of molecules. This phenomenon, known as domain shift, can lead to a drop in model performance when tested on new data. Although this issue can be difficult to address, it underscores the importance of robust validation across a variety of independent test sets, ensuring that AI models are both reliable and adaptable to new, unseen data.

Finally, there is the challenge of bridging the gap between *in silico* predictions and *in vivo* realities. While AI models can simulate various pharmacokinetic and pharmacodynamic processes, they may not always accurately reflect the complex interactions that occur in a biological system. A compound that appears promising based on computer predictions might fail when tested in an actual biological system, due to unmodeled or poorly understood effects. Thus, experimental validation remains an essential and non-negotiable step in drug development. Over-reliance on *in silico* models

without sufficient experimental validation poses a significant risk, as it could lead to the advancement of compounds that are not viable in real-world conditions.

6 Conclusion and Future Directions

Generative AI is actively reshaping the drug development landscape, significantly enhancing the speed, cost-effectiveness, and success rates of pharmaceutical research. By reducing discovery timelines by over 60%, cutting early-stage costs by nearly 50%, and improving clinical success rates, AI is transforming drug discovery. It allows the industry to shift from traditional screening to a more efficient design paradigm, changing both the scientific approach and economics of drug development.

However, challenges remain in areas such as data quality, model interpretability, and regulatory adaptation. Overcoming these will require further innovation and collaboration between human expertise and AI. This partnership will drive future breakthroughs, removing longstanding economic and time constraints.

Key trends shaping the future of AI in drug discovery include multi-modal AI, which integrates diverse data types for more personalized treatments, and federated learning, enabling global collaboration without compromising data privacy. Closed-loop, automated labs, combining robotic execution with AI design, will speed up discovery processes, making self-driving labs a game-changer in therapeutic development.

In summary, AI is paving the way for a new age in global health, offering transformative potential for creating innovative treatments and solving major medical challenges. The effective and responsible application of AI will unlock a new generation of medicines for humanity.

Data Availability Statement

Not applicable.

Funding

This work was supported without any funding.

Conflicts of Interest

Saurabh Sarkar is an employee of Chicory AI, Greater Seattle, United States.

Ethical Approval and Consent to Participate

Not applicable.

References

- [1] Ocana, A., Pandiella, A., Privat, C., Bravo, I., Luengo-Oroz, M., Amir, E., & Gyorffy, B. (2025). Integrating artificial intelligence in drug discovery and early drug development: a transformative approach. *Biomarker Research*, 13(1), 45. [Crossref]
- [2] Lovetruer, B., & Lovetruer, I. (2022). Prospectively validated disease-agnostic predictive medicine with augmented intelligence. *medRxiv*, 2022-03. [Crossref]
- [3] Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., & Kang, J. (2020). BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4), 1234-1240. [Crossref]
- [4] Zhang, O., Lin, H., Zhang, X., Wang, X., Wu, Z., Ye, Q., ... & Hou, T. (2025). Graph Neural Networks in Modern AI-aided Drug Discovery. *Chemical Reviews*. [Crossref]
- [5] Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., ... & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *nature*, 596(7873), 583-589. [Crossref]
- [6] Tan, D., Henry, C. J., & Leung, C. K. (2025). A Guided Variational Autoencoder for Targeted Molecule Optimization in Drug Discovery. *Journal of Healthcare Informatics Research*, 1-21. [Crossref]
- [7] Kao, P. Y., Yang, Y. C., Chiang, W. Y., Hsiao, J. Y., Cao, Y., Aliper, A., ... & Lin, Y. C. (2023). Exploring the advantages of quantum generative adversarial networks in generative chemistry. *Journal of Chemical Information and Modeling*, 63(11), 3307-3318. [Crossref]
- [8] Guo, Z., Liu, J., Wang, Y., Chen, M., Wang, D., Xu, D., & Cheng, J. (2024). Diffusion models in bioinformatics and computational biology. *Nature reviews bioengineering*, 2(2), 136-154. [Crossref]
- [9] Mao, J., Wang, J., Zeb, A., Cho, K. H., Jin, H., Kim, J., ... & No, K. T. (2023). Transformer-based molecular generative model for antiviral drug design. *Journal of chemical information and modeling*, 64(7), 2733-2745. [Crossref]
- [10] Guo, H., Xing, X., Zhou, Y., Jiang, W., Chen, X., Wang, T., ... & Xu, J. (2025). A Survey of Large Language Model for Drug Research and Development. *IEEE Access*. [Crossref]
- [11] Ho, D., Quake, S. R., McCabe, E. R., Chng, W. J., Chow, E. K., Ding, X., ... & Zarrinpar, A. (2020). Enabling technologies for personalized and precision medicine. *Trends in biotechnology*, 38(5), 497-518. [Crossref]
- [12] Fang, S. J., Yin, Z. D., Cai, Q., Li, L. F., Zheng, P., & Chen, L. (2025). Harnessing artificial intelligence for brain disease: advances in diagnosis, drug discovery, and closed-loop therapeutics. *Frontiers in Neurology*, 16, 1615523. [Crossref]

- [13] Gutnik, D., Evseev, P., Miroshnikov, K., & Shneider, M. (2023). Using AlphaFold predictions in viral research. *Current Issues in Molecular Biology*, 45(4), 3705-3732. [[Crossref](#)]
- [14] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)* (pp. 4171-4186). [[Crossref](#)]
- [15] Ramsay, R. R., Popovic-Nikolic, M. R., Nikolic, K., Uliassi, E., & Bolognesi, M. L. (2018). A perspective on multi-target drug discovery and design for complex diseases. *Clinical and translational medicine*, 7(1), 3. [[Crossref](#)]
- [16] Mendez, D., Gaulton, A., Bento, A. P., Chambers, J., De Veij, M., Félix, E., ... & Leach, A. R. (2019). ChEMBL: towards direct deposition of bioassay data. *Nucleic acids research*, 47(D1), D930-D940. [[Crossref](#)]
- [17] Sertkaya, A., Beleche, T., Jessup, A., & Sommers, B. D. (2024). Costs of drug development and research and development intensity in the US, 2000-2018. *JAMA network open*, 7(6), e2415445-e2415445. [[Crossref](#)]



Saurabh Sarkar Chicory AI, USA (Email: sarkarsaurabh.27@gmail.com)

Dr. Radha Raman Chandan SMS, Varanasi, India (Email: rrcmiet@gmail.com)



Dr. Shashi Bhushan Universiti Teknologi PETRONAS, Seri Iskandar, Malaysia (Email: tyagi_shashi@yahoo.com)



Abhijat Bachelor of Technology Computer Science Engineering (AI & ML) (Email: abhijat1025@gmail.com)



Dr. Basu Dev Shivahare Galgotias University Greater Noida, India (Email: basuiimt@gmail.com)