International Journal of Pharmaceutical Sciences and Drug Research

[ISSN: 0975-248X; CODEN (USA): IJPSPP]

journal home page : http://ijpsdronline.com/index.php/journal

Review Article

Decade of AI in Drug Discovery: Hype or Reality?

Leela Laroia*

Singapore American School, 40 Woodlands Street 41, Singapore 738547

ARTICLE INFO

Article history:

Received: 29 March, 2026

Revised: 13 April, 2026

Accepted: 28 April, 2026

Published: 30 May, 2026

Keywords:

Artificial intelligence, Drug discovery, AlphaFold, Clinical trials, Phase II failure, Target validation, Mendelian randomization, ADMET prediction

DOI:

10.25004/IJPSDR.2026.180303

ABSTRACT

AI promised to reinvent drug discovery: compress timelines, reduce attrition, and unlock previously undruggable drug targets. A decade of substantial capital investment has produced more than 173 clinical programs and landmark structural biology breakthroughs, including the 2024 Nobel Prize–winning AlphaFold platform. Yet as of early 2026, no AI-discovered drug has received regulatory approval. This review analyses the clinical evidence systematically, examining Phase I and Phase II trial outcome data across AI-native biopharma pipelines. The data reveal a bifurcated picture: AI demonstrably outperforms historical industry norms in Phase I (80–90% *vs.* ~52% success), reflecting genuine superiority in molecular safety engineering and ADMET prediction, while Phase II success (~40%) remains statistically indistinguishable from conventional drug development (~37%). The persistent Phase II failure rate maps onto a single root cause: target hypothesis validity. AI has not yet improved the probability of selecting the correct biological premise. This paper identifies where AI adds value (protein structure prediction, lead optimization, timeline compression), where it does not (target validation, patient stratification, animal model-to-human translation), and charts the emerging frontier of human-genetics- multi-omics integration and biomedical foundation models as the next phase of AI-enabled progress.

Introduction

On 9 October 2024, the Royal Swedish Academy of Sciences announced that the Nobel Prize in Chemistry would be awarded to David Baker, Demis Hassabis, and John Jumper, recognizing, in the Committee's own framing, "the computational design of proteins" and "protein structure prediction." The protein folding problem, predicting the three-dimensional conformation of a protein from its amino acid sequence alone, had resisted solution for more than 50 years. The challenge was formidable: a 300-residue protein has an astronomical number of possible conformations, and the energy landscape governing folding is extraordinarily complex. Conventional experimental methods, such as X-ray crystallography, cryo-electron microscopy, and nuclear magnetic resonance spectroscopy, could resolve structures with atomic precision but at enormous cost in time and material. The structural data gap was a direct bottleneck in rational drug design: one cannot design a molecule to fit a binding pocket one cannot see.

When DeepMind's AlphaFold 2 system dominated the CASP14 protein structure prediction competition in late 2020, achieving a level of accuracy previously associated only with experimental methods, it was immediately apparent that something fundamental had changed. The landmark Nature paper by Jumper et al. in July 2021 formalized what the competition had demonstrated: a deep learning architecture combining multiple sequence alignments and geometric reasoning could predict protein structures at atomic resolution, at proteome scale, and at negligible computational cost per structure once trained. [1]

The practical implications for drug discovery were immediate. The bottleneck in target-based drug discovery

***Corresponding Author:** Ms. Leela Laroia

Address: Singapore American School, 40 Woodlands Street 41, Singapore 738547

Email ✉**:** laroia771876@sas.edu.sg

Tel.: +65 9170 4242

Relevant conflicts of interest/financial disclosures: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

has always had two distinct phases: first, identifying which proteins are implicated in disease; second, understanding their structural biology well enough to design molecules that modulate them selectively. AlphaFold addressed the second phase with unprecedented comprehensiveness. The release of AlphaFold DB, hosting predicted structures for over 200 million proteins across virtually all cataloged organisms, covering 98.5% of the human proteome at high confidence [2], transformed structure availability from a limiting factor into a near-universal resource.

It was this transformation that drew my attention to the intersection of AI and drug discovery as an area warranting systematic evaluation. Rational chemistry could be applied where previously only phenotypic screening had been possible. Compounds could be designed to exploit allosteric pockets visible only in particular conformational states, which AlphaFold's conformational sampling was beginning to make accessible.

AlphaFold 3, released in 2024, extended the paradigm further still, enabling co-prediction of protein–DNA, protein–RNA, and protein–ligand interactions with accuracy sufficient to guide real experimental campaigns. [3] The Nobel Committee's recognition was, in this context, a statement not merely about scientific achievement but about therapeutic consequence: structural biology, and with it rational drug design, had been democratized.

This democratization raised an obvious and important empirical question. If AI can now predict where a molecule should bind, can it predict whether doing so will cure a disease? The answer, as the clinical record now documents, is considerably more nuanced and the gap between structural prediction and therapeutic success reveals precisely where the next decade's scientific investment must be directed.

Was the Hype Rational?

The excitement was not irrational. Drug discovery is, in large part, a pattern-recognition problem and pattern-recognition is precisely where modern machine learning excels. The chemical space of potential drug-like molecules is estimated at 10^{60}, a number so large it makes brute-force exploration meaningless. A foundational 2019 review in Nature Reviews Drug Discovery mapped how machine learning could systematically navigate that space, learning structure–activity relationships from decades of assay data to propose molecules more likely to bind a target, avoid metabolic liabilities, and survive the immune system. [4]

The early proof points were real (Table 1). Companies like Exscientia demonstrated that AI could compress the hit-to-lead optimization phase from years to months. Insilico Medicine designed a fibrosis drug candidate in under 18 months at a reported cost below $2.6 million. [5] Then, in July 2021, AlphaFold's landmark paper in Nature solved protein structure prediction at a scale that had eluded structural biology for five decades, opening new avenues for rational drug design against previously undruggable targets.[1] AlphaFold 3, released in 2024, extended this to protein–DNA, protein–RNA, and protein–ligand co-prediction.[3] These were genuine scientific achievements indeed.

Investors saw the compression of timelines, the reduction of chemistry costs, and the possibility of applying AI systematically across entire therapeutic areas. They priced in a future where AI-native companies could run many parallel programs at a fraction of traditional pharmaceutical costs, with higher hit rates producing a fundamentally superior pipeline. It was a coherent thesis. It was also, as it turns out, a thesis about the wrong bottleneck.

What the Data Actually Shows

By early 2026, more than 173 AI-discovered drug programs have entered some stage of clinical development.[8,9] As of that date, none have received FDA approval. The pattern of attrition maps cleanly onto known chokepoints. A 2024 analysis by the Boston Consulting Group (Tables 2 and 3), the first systematic examination

Table 1: Key milestones in AI drug discovery, 2012–2026

Year	*Event*	*Detail*
2012	AlexNet/Deep Learning Revolution	ImageNet victory catalyzes life sciences interest in deep learning for quantitative structure-activity relationship (QSAR) and target identification.
2014–19	Formation of AI-native drug companies	Exscientia, BenevolentAI, Recursion, and *In-silico* Medicine raise funding
2020	DSP-1181 enters first human trials	Exscientia's AI-designed OCD drug dosed in humans. Phase I was completed with favorable safety; development was later discontinued.
2021	AlphaFold 2 published in Nature	Jumper et al. achieve atomic-level protein structure prediction at the proteome scale. Awarded the 2024 Nobel Prize in Chemistry
2022–23	First wave of Phase II failures	BEN-2293 fails vs. placebo; REC-994 discontinued; EXS-21546 wound down
2025	First Phase IIa proof-of-concept	Insilico Medicine's rentosertib (ISM001-055) published in Nature Medicine: dose-dependent Forced Value Capacity (FVC) improvement in Idiopathic Pulmonary Fibrosis (IPF) patients

Source: Jumper *et al*. 2021 [1]; Jayatunga *et al*. 2024 [6]; Drug Target Review 2026 [7]

of AI-native biopharmaceutical clinical pipelines, found that AI-discovered molecules achieve 80–90% success in Phase I, substantially above historic industry averages of roughly 52%. Phase II success, however, was tracked at approximately 40%, comparable to conventional drug industry norms.[6] A complementary analysis in Clinical Pharmacology & Therapeutics confirmed that no AI-discovered drugs had attained clinical approval despite several companies emerging between 2014 and 2019.[10]

What Worked and What Did Not

The distinction between success and failure points towards the quality of the underlying target hypothesis and the approach to patient selection used to test it.[11,12] A 2025 landscape review (Tables 4 and 5) that compared five leading AI platform types found that physics-augmented machine learning and conformational dynamics modeling required well-validated targets to be successful.[13]

Where AI has demonstrably worked is against well-understood targets. Relay Therapeutics' RLY-2608, a PI3Kα mutant-selective inhibitor for breast cancer, entered Phase III in 2025. Its trajectory is instructive: the target (a specific PIK3CA activating mutation) is causally established in human cancer genetics; patients are biomarker-selected; and AI was applied to the genuinely hard structural problem of conformational selectivity. [14] Schrödinger's zasocitinib (TAK-279), a TYK2 inhibitor leveraging physics-based free energy calculations, similarly advanced to Phase III, for a well-validated immunology target.

The 2020 landmark paper in Cell by Stokes *et al*., identifying halicin as a novel antibiotic via graph neural networks[14], further demonstrates that the technology is genuinely powerful when the assay is direct and the biological premise well-grounded. The problem is not the chemistry. The problem is the hypothesis it rests on.

The Next Phase of Improvement

If Phase II failure is fundamentally a target biology problem, the next frontier for AI is not better generative chemistry; it is target validation. This is harder, more uncertain, and less amenable to the labeled datasets that made molecular property prediction tractable. But it is where the leverage lies (Table 5).

Three overlapping developments make cautious optimism reasonable. First, large-scale human genetics databases, UK Biobank, FinnGen, NIH's All of Us initiative, are generating phenotype–genotype data at unprecedented scale. These datasets enable Mendelian randomization (MR) analyses that proxy the lifelong consequences of drug target perturbation.[15-17] When a genetic variant that naturally reduces expression of a target also reduces disease incidence in tens of thousands of individuals, that constitutes causal evidence of categorically different quality than any cell line or mouse model. Second, the integration of multi-modal biological data, transcriptomics, proteomics, metabolomics, and imaging, is creating richer

Table 2: Phase I and Phase II clinical trial success rates: AI-discovered vs. traditional drug candidates

Trial phase	*AI-Discovered Candidates (Jayatunga et al., 2024 [6])*	*Traditional Drug Candidates (Wong et al., 2019 [11])*
PHASE I *Safety & tolerability*	**80 - 90%** Significantly above industry average *ADMET prediction, PK/PD optimisation, off-target liability avoidance*	**~52%** Historical industry baseline *Failures: safety signals, tolerability, PK liabilities*
PHASE II *Efficacy proof-of-concept*	**~40%** Comparable to industry average *Failures: target hypothesis errors, patient heterogeneity, translation gap*	**~37%** Historical industry baseline *Same target biology and translation challenges*

Key finding: AI delivers a meaningful Phase I advantage (80-90% vs. ~52%), reflecting genuine superiority in molecular safety engineering and ADMET prediction. This advantage disappears at Phase II (~40% vs. ~37%), where efficacy depends on target hypothesis validity, a domain AI has not yet improved.

Source: Jayatunga *et al*. [6]; Drug Target Review [7]; Wong *et al*. [11]

Table 3: AI drug pipeline status by phase (Early 2026) vs. historical industry baseline

Clinical Phase	*AI Programs*	*Traditional Success Rate*	*AI Success Rate*
Phase I	~94 programs	~52%	80–90% ↑
Phase II	~56 programs	~37%	~40% →
Phase III	~15 programs	~59%	Data pending
FDA Approved	0	—	0

↑ = significantly above industry average; → = comparable to industry average.
Source: Jayatunga *et al.* [6], Drug Discov Today [7], Pharmaceutics [8], Wong *et al.* 2019 [10]

Table 4: Notable AI drug failures and their mechanistic interpretation

Drug	*Company*	*Indication*	*Stage*	*Failure Mode*
DSP-1181	Exscientia / Sumitomo	OCD	Phase I	Favourable safety; discontinued 2022. Efficacy and business criteria unmet
EXS-21546	Exscientia	Oncology	Phase I/II	Poor therapeutic index in humans
BEN-2293	BenevolentAI	Atopic dermatitis	Phase IIa	No efficacy vs. placebo. Target hypothesis incorrect
REC-994	Recursion	Cerebral cavernous malformation	Phase II	No patient benefit despite early MRI signals. The preclinical signal did not translate
ABBV-CLS-7262	Calico / AbbVie	ALS	Phase II/III	Missed primary and secondary endpoints. Disease biology exceeded preclinical model capacity.
BEN-8744 (PDE10A)	BenevolentAI	Ulcerative colitis	Phase II	No efficacy signal on the primary endpoint

Source: Drug Target Review 2026 [7]; LoonBio [11]; Guerrero-González, *et. al.*[12]

Table 5: Where AI adds value vs. where it does not

AI demonstrably works	*AI has not improved this*
Molecular design vs. validated targets: ADMET prediction, lead optimization, QSAR, generative chemistry against structurally characterized proteins with known causal disease links.	Target hypothesis validation: Determining whether a biological pathway causally drives disease in the patient population being treated. Phase II failure rates unchanged from industry norms.
Protein structure prediction: AlphaFold covers 98.5% of the human proteome at high confidence, enabling rational design against previously undruggable targets.	Patient heterogeneity: Real patient populations confound preclinical signals. Biomarker selection for stratification remains largely intuition-driven, blunting trial power.
Timeline compression: Preclinical to IND in 13–18 months vs. 3–4 years conventionally. Cost reduction of early-stage chemistry by 30–40%.	Animal model-to-human translation: The fundamental gap between rodent models and human disease biology is not addressed by better molecule design.

Synthesis from: Jayatunga *et al.* [6]; Drug Target Review [7]; Wilczok *et al.* [9]; Stokes *et al.* [14]

representations of disease state. A 2025 review described how AI-driven multi-omics integration is beginning to enable genotype–environment–phenotype modeling at a resolution that may identify responder subpopulations before trial enrolment.[18] Third, foundation models trained on the biomedical literature are developing the capacity to synthesize contradictory evidence, identifying where a target hypothesis rests on shaky mechanistic ground, or where patient heterogeneity is likely to swamp a trial signal.

None of this will arrive cleanly or quickly. Biology remains the most complex system AI has attempted to model, and human disease is biology at its most heterogeneous. The $60 billion hypothesis was not wrong about AI's potential; it was premature about the timeline and perhaps imprecise about the bottleneck. The industry learned, at considerable cost, that you cannot automate your way around a bad target. The next generation of AI drug discovery will be defined by whether the technology can help researchers choose better ones.

Conclusion

A decade of AI in drug discovery yields a verdict that is neither the triumph of early investors nor the dismissal of late skeptics: it is a story of genuine, bounded progress. AI has fundamentally changed what is possible in the early stages of drug design. Protein structure prediction has been democratized through AlphaFold and its successors, ADMET-optimized molecules consistently reach Phase I with superior safety profiles, and preclinical timelines have compressed by years. These are durable contributions to medicine.

What AI has not changed is Phase II attrition. The ~40% Phase II success rate for AI-designed molecules is statistically indistinguishable from the historical industry baseline of ~37%, and the failure modes are identical: incorrect target hypotheses, patient population heterogeneity, and the irreducible gap between animal models and human disease. The central lesson of the 2022–2025 failure wave is that computational power cannot substitute for biological insight. A perfectly optimized molecule against the wrong target will fail, and it will fail at exactly the same rate as a conventionally designed one. The next chapter will be won or lost on target validation. The convergence of large-scale human genetics databases, Mendelian randomization methodology, multi-omics integration, and biomedical foundation models is creating, for the first time, a computational pathway to causal inference in disease biology. This is harder science than generative chemistry, operates on noisier

data, and will require different validation frameworks than those currently applied to molecular design AI. But it is also where the remaining leverage lies. If AI can shift the probability that a Phase II trial is testing the right hypothesis in the right patient population, the compounding effect on pipeline economics will be transformative in ways that better molecule factories, alone, never could be.

The $60 billion invested in AI drug discovery over the past decade bought the industry an important and expensive lesson: the bottleneck was not chemistry alone[11]. It involved biology. The technology that spent a decade learning to build better molecules must now learn to ask better questions pertaining to biology.

References

1. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596:583–589. Available from: https://doi.org/10.1038/s41586-021-03819-2
2. Tunyasuvunakool K, Adler J, Wu Z, Green T, Zielinski M, Žídek A, et al. Highly accurate protein structure prediction for the human proteome. Nature. 2021;596:590–596. Available from: https://doi.org/10.1038/s41586-021-03828-1
3. Abramson J, Adler J, Dunger J, Evans R, Green T, Pritzel A, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature. 2024;630:493–500. Available from: https://doi.org/10.1038/s41586-024-07487-w
4. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nature Reviews Drug Discovery. 2019;18:463–477. Available from: https://doi.org/10.1038/s41573-019-0024-5
5. Xu Z, Ren F, Wang P, Cao J, Tan C, Ma D, et al. A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomized phase 2a trial. Nature Medicine. 2025;31:2602–2610. Available from: https://doi.org/10.1038/s41591-025-03743-2
6. Jayatunga MKP, Ayers M, Bruens L, Jayanth D, Meier C. How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons. Drug Discovery Today. 2024;29:1–6. Available from: https://doi.org/10.1016/j.drudis.2024.104009
7. 2026: The year AI stops being optional in drug discovery. Drug Target Review [Internet]. 2026. Available from: https://www.drugtargetreview.com/2026-the-year-ai-stops-being-optional-in-drug-discovery/682243.article
8. Dermawan D, Alotaiq N. From lab to clinic: How artificial intelligence (AI) is reshaping drug discovery timelines. Pharmaceuticals. 2025;18:981. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12298131/
9. Wilczok D., Zhavoronkov A. Progress, pitfalls, and impact of AI-driven clinical trials. Clinical Pharmacology & Therapeutics. 2025;117:887–890. Available from: https://doi.org/10.1002/cpt.3542
10. Wong CH, Siah KW, Lo AW. Estimation of clinical trial success rates and related parameters. Biostatistics. 2019;20:273–286. Available from: https://doi.org/10.1093/biostatistics/kxx069
11. AI drug discovery's $60 billion reality check: Hype, failures, and the market access blindspot. LoonBio Reflections [Internet]. 2025. Available from: https://loonbio.com/reflections/ai-drug-discoverys-60-billion-reality-check-hype-failures-and-the-market-access-blindspot
12. Niazzi SK. Artificial Intelligence in Small-Molecule Drug Discovery: A Critical Review of Methods, Applications, and Real-World Outcomes. Pharmaceuticals. 2025;18:1271. Available from: https://doi.org/10.3390/ph18091271
13. Dharmasivam M, Kaya B, Akinware A, Azad MG, Richardson DR. Leading artificial intelligence-driven drug discovery platforms: 2025 landscape and global outlook. Pharmacological Reviews. 2026;78:1. Available from: https://doi.org/10.1016/j.pharmr.2025.100102
14. Stokes JM, Yang K, Swanson K, Jin W, Cubillos-Ruis A, Donghia NM, et al. A deep learning approach to antibiotic discovery. Cell. 2020;180:4.688–702. Available from: https://doi.org/10.1016/j.cell.2020.01.021
15. Daghlas I, Gill D. Mendelian randomization as a tool to inform drug development using human genetics. Cambridge Prisms Precision Medicine. 2023;1:1–12. Available from: https://doi.org/10.1017/pcm.2023.5
16. Nelson MR, Tipney H, Painter JL, Shen J, Nicoletti P, Shen Y, et al. The support of human genetic evidence for approved drug indications. Nature Genetics. 2015;47(8):856–860. Available from: https://doi.org/10.1038/ng.3314
17. Minikel EV, Painter J, Dong CC, Nelson MR. Refining the impact of genetic evidence on clinical success. Nature. 2024;629:624–629. Available from: https://doi.org/10.1038/s41586-024-07316-0
18. Wu Y, Xie L. AI-driven multi-omics integration for multi-scale predictive modeling. Computational and Structural Biotechnology Journal. 2025;27:265–277. Available from: https://doi.org/10.1016/j.csbj.2024.12.030

HOW TO CITE THIS ARTICLE: Laroia L. Decade of AI in Drug Discovery: Hype or Reality?. Int. J. Pharm. Sci. Drug Res. 2026;18(3):144-148. DOI: 10.25004/IJPSDR.2026.180303