Artificial Intelligence in Pharmaceutical R&D: Impact and Use Across Development Stages

By: Jennifer Kao, Melissa Ouellet, Ariel Stern

Extended Abstract

Introduction

In January 2026, the U.S. Food and Drug Administration and the European Medicines Agency jointly issued *Guiding Principles of Good AI Practice in Drug Development,* a set of principles for the use of artificial intelligence (AI) across the drug development life cycle (*1*). The principles reflect a transatlantic recognition by regulators: the question is no longer whether AI will play a role in drug development, but how its use should be overseen. Industry appears to have reached the same expectation and is investing accordingly. In a recent survey, 95% of healthcare executives anticipate that generative AI will have a transformative effect (*2*). Investment in AI in the pharmaceutical industry is projected to grow from more than $4 billion in 2025 to $25.7 billion by 2030 (*3*).

At the same time, there has been growing concern that AI's impact on drug development may be more limited than expected. For example, a key reason that drugs fail in development is not because the molecule is poorly engineered but because the mechanism it targets is wrong (*4*). However, much of AI-related efforts have focused on the first problem rather than the second. As Daphne Koller, founder and chief executive of insitro, a machine learning-driven drug discovery company describes, the industry is "doing a pretty good job at manufacturing keys, but they are often for the wrong locks" (*5*). She notes that identifying the mechanism that drives a disease is a key bottleneck: until it is addressed, AI may simply produce the same failures, only faster.

Despite this attention, documentation and analyses of the nature and frequency of AI use in drug development remain limited. Beyond the stakeholders directly leading these projects and individual anecdotes, relatively little is known about how AI is being used in practice. Moreover, the phrase "AI in drug development" is used expansively and encompasses a wide range of activities, from target discovery to clinical trial site optimization. Some of these algorithm-driven activities–in particular simpler forms of predictive modeling–rely on methods that machine learning experts might not characterize as AI (*6*). But implications for industry leaders, patients, physicians, and regulators depend on the types of applications for which AI is used.

A Framework for AI in Drug R&D

Against this backdrop, we develop a framework to distinguish categories of AI use based on *where* in the development process AI operates and *what* it is intended to accomplish: research applications vs. development applications. Research applications (e.g., drug molecule design or drug target discovery) aim to lower the cost of research at the earliest stages and widen the pool of candidate therapies. If successful, they can result in treatments based on novel mechanisms, targets, and molecules that conventional methods might overlook. Development applications (e.g., clinical trial patient stratification or clinical outcome prediction) aim to move existing candidates through the development pipeline and toward clinical translation faster and/or at lower cost. The stakes are considerable: estimates suggest that bringing a new drug to market can cost more than $2 billion and take over a decade (*7*). Given the scale of these costs, even modest efficiency improvements could result in substantial savings. In one recent survey, industry participants reported that applying AI reduced development times by an average of 18% (*8*).

Documenting the Use of AI

To measure how AI is used in practice, we examined the *Citeline Pharmaprojects Drugs Database*, a comprehensive drug development database. First, we identified candidate drugs by searching the database's full-text fields for AI-related terms already established in the literature (including artificial intelligence, deep learning, generative adversarial network, machine learning, neural network, and random forest, along with "AI"-prefixed terms). Next, we manually inspected and confirmed each candidate to keep only drugs with genuine documented AI use. Following prior work (*9*), we assigned each AI application to one or more of 11 non-mutually exclusive purposes (from drug molecule discovery to clinical outcomes prediction) and classified each as either a research or development application.

Table S1 documents that between January 1, 2015 and January 2, 2025, 363 drug candidates had documented use of AI in R&D activities ("AI-drugs"). Relative to the broader pipeline of contemporaneous drug candidates, AI-drugs were more likely to be small molecules than biologics: 70% of AI-drugs were small molecules, compared with 40% of drugs without documented AI use. They also concentrated in a few therapeutic areas: oncology (33%) and neurological conditions (24%), with a smaller share used to treat immunological disorders (14%) and metabolic conditions (12%). Relative to the broader pipeline, AI is disproportionately in neurological (24% vs. 16% for non-AI-drugs) and immunological conditions (14% vs. 6%).

Table 1 shows that across the two broad categories of AI use, research applications were the most common: 313 (86%) of the drugs had at least one research application, while 50 (14%) drugs had only development applications. Among research applications, two approaches were most frequently used: drug molecule discovery (appearing in 61% of AI-drugs) and molecular design (52%). AI, for instance, was used to design novel candidate molecules for Parkinson's disease and for oncology (see examples in the rightmost column of Table 1). In contrast, target

discovery, an application that is closely linked to whether a drug ultimately proves to be effective in patients, appears in only 19% of AI-drugs. Development applications were substantially less commonly described. Many applications, including patient stratification, clinical outcomes analysis, and outcomes prediction, each appeared in fewer than 2% of AI-drugs. One example is the use of AI in the development of elraglusib to identify trial subgroups more likely to respond to treatment.

Across development phases, drugs with AI applications clustered in the early stages. Most had not yet entered a Phase 1 trial and only 7 drugs had reached Phase 3. The contrast between application types is notable (though sample sizes remain small): just 3 of the 313 drugs (1%) with research applications had reached Phase 3, compared with 4 of the 50 drugs (8%) with development applications. This difference is consistent with the chronology of R&D: development applications are used by candidates already advancing through clinical testing, while research applications largely target research activities upstream of entry into clinical trials.

These patterns also reflect the relatively recent adoption of AI in drug R&D. For example, many of the generative models used for drug molecule discovery only became practically useful in recent years. The candidates associated with these approaches have therefore had only limited time to advance through a process that often takes more than a decade. As such, what is true of the AI-drug R&D pipeline in this sample need not be true a decade from now–that is, after the current cohort of early-stage drug candidates has had time to progress and newer AI applications have entered the pipeline. Therefore, the concentration of AI-drugs in early development should neither be interpreted as evidence that these applications are ineffective, nor as an indication that these are the most meaningful applications. Rather, the available track record reflects the fact that the use of AI in drug R&D has not yet had enough time to mature enough to understand AI's effects on drug development or patient outcomes beyond the early years.

Implications for Manufacturers, Physicians, Patients, and Regulators

Even at this early stage, emergent patterns have distinct implications for manufacturers, patients, and regulators. With so few drug candidates having advanced through the pipeline, manufacturers cannot yet rely on a meaningful body of evidence to justify adoption. As a result, firms must be disciplined about clarifying what specific problem AI is being used to solve and why AI is a differentiated tool for solving it. Our framework makes this distinction concrete. A firm should first determine whether it intends to use AI to broaden its pool of candidates and separately whether and how to use AI to move existing candidates through clinical development. Strategically, manufacturers should clarify how and for which tasks AI is most likely to outperform traditional methods and define how its performance will be assessed. The relevant outcomes and endpoints for such an assessment differ across applications: for research applications such as molecule discovery and design, an important outcome is likely to be the rate

of clinical success. For development applications such as trial site optimization, a key outcome would be trial recruitment durations or higher completion rates. In each setting, methods for evaluating the impact of AI tools will need to be use-case specific and fit-for-purpose.

For patients and physicians, these findings suggest that the benefits of AI in drug R&D remain prospective. Most drugs with documented AI applications remain in preclinical or early clinical development, and their effects on patient outcomes at scale have yet to be observed. Since AI activity concentrates in a few therapeutic areas (most notably, oncology and neurological conditions, as described above) these are where physicians should expect to encounter AI-therapies soonest. Notably, the involvement of AI should not itself be interpreted as evidence of greater clinical value. That value must ultimately be established through evidence on time-to-access, alongside safety and effectiveness.

Finally, for regulators, a key priority should be on reducing uncertainty about how the use of AI tools and their outputs will be evaluated (*10*). Such uncertainty increases case-by-case work for regulators themselves and raises the cost and risk of AI drug development for manufacturers. Indeed, such uncertainty may explain some of the relatively low number of candidates that have advanced and/or some hesitation to adopt such tools in the first place (obviously unobservable in our data).

The joint good-AI-practice principles published earlier this year nevertheless represent a positive signal that regulators recognize these challenges and are working toward a more predictable framework. Regulators can build on this foundation by developing the technical expertise to assess such applications, clarifying what evidence and documentation they will expect from manufacturers, and incorporating feedback from external stakeholders (*11*). This will become increasingly important as adoption of AI continues and early-stage drug candidates, such as those identified here, begin to approach regulatory review.

Looking Ahead

More broadly, this analysis provides a starting point for understanding the growing use of AI in pharmaceutical R&D. Since most of the activity to-date remains early stage, an important next step for researchers will be to follow these drugs as they enter clinical testing, advance through the pipeline, or are discontinued. This will facilitate a clearer understanding of whether and under what conditions AI candidates are more likely to enter the clinic and, conditional on entering, whether they are more likely to succeed. Ongoing surveillance will also allow researchers to assess whether these outcomes differ across types of AI applications and whether AI affects the time and cost required to move candidates through development. Following these products closely through R&D will help move the discussion beyond expectations about AI's

potential and toward evidence about the applications translate into meaningful scientific, clinical, and economic value.

Supplementary Materials: Table S1

Table 1. AI-drugs (n=363) categorized by their research (n=313) & development purposes (*12*) - categories are not mutually exclusive.

AI drug purpose	N (%)	Definition	Examples
Drug molecule discovery	223 (61.4%)	This refers to using Artificial Intelligence (AI) / Machine Learning (ML) methods to identify or generate novel chemical or biological entities (small molecules, antibodies, oligonucleotides, etc.) that show potential as therapeutic candidates. This often involves screening vast virtual libraries, predicting molecule properties, or using generative models to create new structures.	triple-negative breast cancer therapy, Jiangsu Chia Tai Fenghai *(anticancer)*; Parkinson's disease therapy, Celeris Therapeutics *(neurological)*
Drug target discovery	70 (19.3%)	This involves using AI/ML to analyze biological data (genomic, proteomic, transcriptomic, literature, clinical data) to identify and/or validate specific biological targets (e.g., proteins, genes, pathways) implicated in a disease, which can then be pursued for drug development.	T-cell therapy, Bluebird Bio *(anticancer)*; neurodegeneration therapy, C4X Discovery *(neurological)*
Drug molecule design	190 (52.3%)	Using AI/ML, often generative approaches or structure-based methods informed by AI, to specifically design or optimize the structure of drug molecules de novo or modify existing ones to achieve desired properties (e.g., binding affinity, selectivity, pharmacokinetic profile).	cancer therapeutics, AstraZeneca *(anticancer)*; COVID-19 vaccine, HaloVax *(anti-infective)*
Protein folding	0 (0.0%)	Using AI/ML algorithms (like AlphaFold or similar methods) to predict the three-dimensional structure of proteins based on their amino acid sequence. Understanding protein structure is critical for rational drug design and target identification.	PD-L1/TGF-beta bi-specific antibody, Gigagen *(anticancer)*; AAV-gene therapy, Novartis *(sensory)*
Personalized treatment	10 (2.8%)	Using AI/ML to guide the selection or dosing of therapies for individual patients based on their unique biological or clinical profiles. This can include identifying patient-specific neoantigens for cancer vaccines or predicting individual responses to different drugs.	TG-4050 *(anticancer)*; Lewy body dementia therapy, Path Biotech *(neurological)*
Biomarker development	14 (3.9%)	Applying AI/ML to analyze diverse datasets (e.g., clinical, genomic, proteomic, imaging) to identify and validate biomarkers. These biomarkers can help diagnose diseases, predict patient prognosis, monitor disease progression, or predict response to therapy.	T-cell therapy, Bluebird Bio *(anticancer)*; blarcamesine hydrochloride *(alimentary/metabolic; cardiovascular; neurological)*
Drug repurposing	38 (10.5%)	Utilizing AI/ML techniques to systematically scan existing drugs (approved or investigational) or known compounds to identify potential new therapeutic uses for different diseases based on predicted interactions, pathway analysis, literature mining, or analysis of real-world data.	mebendazole, BullFrog AI *(anticancer)*; bavisant *(neurological)*
Clinical trial patient stratification	4 (1.1%)	Employing AI/ML models to analyze baseline patient data (e.g., demographics, clinical history, genetics, biomarkers) to identify subgroups within a trial population who are more likely to respond positively or negatively to the investigational therapy, potentially leading to more efficient trials or personalized medicine approaches.	elraglusib *(anticancer; blood and clotting)*; neurodegeneration therapy, C4X Discovery *(neurological)*
Clinical outcomes analysis	3 (0.8%)	Applying AI/ML methods (potentially more advanced than standard statistics) to analyze data generated during or after a clinical trial to better understand treatment effects, identify complex response patterns, analyze imaging data, or interpret real-world evidence related to the drug's performance.	inaleucel *(anticancer)*; blarcamesine hydrochloride *(alimentary/metabolic; cardiovascular; neurological)*
Clinical outcomes prediction	1 (0.3%)	Developing AI/ML models based on historical or simulated data to forecast the likelihood of success, duration, cost, or specific patient outcomes for a planned or ongoing clinical trial.	IGV-001 *(anticancer)*; CIL-001 *(alimentary/metabolic; neurological; sensory)*
Clinical development challenges	3 (0.8%)	Applying AI/ML to specifically address obstacles in the clinical development pipeline, such as optimizing trial design (e.g., selecting endpoints, patient populations), improving patient recruitment, predicting adverse events, or analyzing factors contributing to trial failures.	triple-negative breast cancer therapy, Jiangsu Chia Tai Fenghai *(anticancer)*; cardiovascular disease therapies, Bayer *(cardiovascular)*

References and Notes

1. U.S. Food and Drug Administration, European Medicines Agency, “Guiding principles of good AI practice in drug development” (FDA and EMA, 2026); https://www.fda.gov/media/189581/download.
2. E. Berger, C. Dowling, S. Guerra, S. Kraus, “The healthcare AI adoption index” (Bain & Company, 2025); https://www.bain.com/insights/the-healthcare-ai-adoption-index/.
3. A. Devereson, N. Nagra, “How pharma is rewriting the AI playbook: Perspectives from industry leaders” (McKinsey & Company, 2025); https://www.mckinsey.com/industries/life-sciences/our-insights/the-synthesis/how-pharma-is-rewriting-the-ai-playbook-perspectives-from-industry-leaders.
4. R. K. Harrison, Phase II and phase III failures: 2013–2015. *Nat. Rev. Drug Discov.* **15**, 817–818 (2016).
5. D. Koller, “Drug discovery has no magic wands” (Deep Phenotype, 2026); https://deepphenotype.substack.com/p/drug-discovery-has-no-magic-wands.
6. European Commission, Directorate-General for Health and Food Safety, “Artificial intelligence in healthcare” (European Commission, accessed 2 November 2026); https://health.ec.europa.eu/ehealth-digital-health-and-care/artificial-intelligence-healthcare_en.
7. Congressional Budget Office, “Research and development in the pharmaceutical industry” (CBO, 2021); https://www.cbo.gov/publication/57126.
8. M. J. Lamberti, M. I. Florez, H. Do, S. Rosner, T. Menard, C. Nielson, A. Donovan, J. Ye, S. Kaveripakam, B. Schoeberl, A. R. Hunt, H. Yeardley, The adoption and use of artificial intelligence and machine learning in clinical development. *Ther. Innov. Regul. Sci.* **59**, 1074–1086 (2025).
9. L. C. Druedahl, W. N. Price II, T. Minssen, A. Sarpatwari, Use of artificial intelligence in drug development. *JAMA Netw. Open* **7**, e2414139 (2024).
10. D. Wilczok, A. Zhavoronkov, Progress, pitfalls, and impact of AI-driven clinical trials. *Clin. Pharmacol. Ther.* **117**, 887–890 (2025).
11. H. J. Warraich, T. Tazbaz, R. M. Califf, FDA perspective on the regulation of artificial intelligence in health care and biomedicine. *JAMA* **333**, 241–247 (2025).
12. Z. Chen, X. Liu, W. Hogan, E. Shenkman, J. Bian, Applications of artificial intelligence in drug development using real-world data. *Drug Discov. Today* **26**, 1256–1264 (2021).

Supplementary Materials for Artificial Intelligence in Pharmaceutical R&D: Impact and Use Across Development Stages

By: Jennifer Kao, Melissa Ouellet, Ariel Stern

Table S1. Descriptives of our drug sample, non-AI (n=58,923) and AI- (n=363) drugs.

	Non-AI (%)	AI (%)
Drug Type		
Biologic	46.7	25.9
Small molecule	40.4	70.2
Reformulation	8.7	2.8
Natural product	2.0	0.6
Diagnostic	1.1	0.0
Not applicable/Unidentified	1.0	0.6
Therapeutic Class		
Anticancer	36.7	32.5
Neurological	15.6	23.7
Anti-infective	13.5	6.9
Alimentary/Metabolic	11.8	11.6
Musculoskeletal	6.6	5.0
Immunological	6.1	13.5
Cardiovascular	5.7	5.0
Respiratory	4.7	3.9
Sensory	4.6	3.3
Dermatological	4.5	6.9
NA/Unspecified	3.5	6.3
Blood and Clotting	3.2	1.7
Genitourinary (including sex hormones)	3.1	2.2
Miscellaneous	1.4	0.0
Antiparasitic	0.7	0.0
Hormonal (excluding sex hormones)	0.6	0.0
AI purpose		
Drug molecule discovery		61.4
Drug molecule design		52.3
Drug target discovery		19.3
Drug repurposing		10.5
Biomarker development		3.9
Personalized treatment		2.8
Clinical trial patient stratification		1.1
Clinical development challenges		0.8
Clinical outcomes analysis		0.8
Clinical outcome prediction		0.3
Protein Folding		0.0
AI keyword		
Artificial intelligence		38.3
AI-based		20.4

Machine learning	15.4
AI-driven	11.0
AI-powered	4.4
AI-enabled	3.0
Algorithm	2.5
AI-discovered	2.2
Generative AI	2.2
AI-guided	1.7
AI-designed	1.7
AI-derived	1.7
AI-integrated	1.7
Deep learning	0.8
AI-generated	0.8
Generative Adversarial Network	0.6
AI-augmented	0.6
AI-developed	0.3
Decision tree	0.3
AI-predicted	0.3
Deep clustering algorithm	0.0
Deep generative model	0.0
Random forest	0.0
AI-assisted	0.0
Image recognition	0.0
Neural network	0.0
AI-focused	0.0
AI-empowered	0.0
AI-native	0.0
AI-optimized	0.0

Notes: The ‘Drug Type’ panel presents mutually exclusive categories, while the other panels feature categories that overlap.