IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

AI in Preclinical Development: Emerging Biotech Guide 2026

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-06-22 · ai in preclinical development · ai in drug discovery · machine learning preclinical · biotech ai strategy · generative ai biotech · preclinical ai platforms · fda ai guidance · organ-on-chip · drug discovery ai · life sciences ai

Original article: https://intuitionlabs.ai/articles/ai-in-preclinical-development-biotech



In this article

1. Executive Summary
2. Introduction and Background
3. Predictive Toxicology, Pharmacology, and ADMET Modeling
4. Digital Twins, Organ-on-Chip Systems, and Animal-Testing Alternatives
5. Preclinical Data Infrastructure and Generative AI Research Assistants
6. Data Analysis and Evidence
7. Case Studies and Real-World Examples
8. Insilico Medicine: Rentosertib and an Early Phase 2a Efficacy Signal
9. Ainnocence and the Broader Chinese and Vertically Integrated Platform Landscape
10. Bayer and Thoughtworks: PRINCE as an Internal Preclinical Knowledge Engine
11. Eli Lilly and NVIDIA: Industrial-Scale Compute for Preclinical Modeling
12. Recursion Pharmaceuticals and Exscientia: Consolidation
13. Emulate, Inc.: Organ-Chip Technology and the Animal-Testing Roadmap
14. Implications and Future Directions
15. Frequently Asked Questions (FAQs)
16. Conclusion

Executive Summary

Artificial intelligence (AI) has moved from a peripheral research tool to a core operating layer in preclinical drug development, the stage between target discovery and first-in-human dosing where safety, efficacy, and pharmacokinetics are established in laboratory and animal (or human-relevant) models. As of July 2026, the global AI in drug discovery market was valued at **$2.3 billion in 2025** and is projected to grow to **$2.9 billion in 2026** and **$13.8 billion by 2033**, a **24.8% compound annual growth rate (CAGR)**, with North America holding a **52.8% revenue share** (www.grandviewresearch.com). A 2023 modeling exercise by the **Boston Consulting Group (BCG)** and the **Wellcome Trust** found that applying AI to early drug research and development could yield time and cost savings of **25 to 50%** through the preclinical stage (wellcome.org), while the **McKinsey Global Institute** separately estimated that generative AI could unlock **$60 billion to $110 billion** in annual economic value across the pharmaceutical and medical-product industries (www.mckinsey.com).

Regulators have moved in parallel. The U.S. **Food and Drug Administration (FDA)** issued its first draft guidance on AI in drug and biological product development in **January 2025**, informed by its review of over **500 AI-component submissions between 2016 and 2023** (www.fda.gov), and followed in **January 2026** with "Guiding Principles of Good AI Practice in Drug Development" (www.fda.gov). On **April 10, 2025**, the FDA announced a roadmap to phase out animal testing requirements for monoclonal antibodies and other drugs, explicitly promoting **AI-based computational models of toxicity** alongside organoid and organ-on-a-chip systems as substitutes (www.fda.gov).

Real-world evidence of impact is accumulating but remains mixed. **Insilico Medicine** advanced its AI-discovered and AI-designed antifibrotic candidate, now named **rentosertib**, from target discovery through IND-enabling preclinical work to a Phase I trial in under **30 months**, versus an industry average closer to six years, at a preclinical budget of roughly **$2.6 million** (insilico.com). In June 2025, the drug's Phase IIa results, published in *Nature Medicine*, showed a mean forced vital capacity (FVC) improvement of **+98.4 mL** at the highest dose versus a decline in the placebo arm (www.nature.com), an encouraging preliminary signal from a small, 12-week phase 2a study whose primary endpoint was safety. A first systematic academic analysis found AI-discovered molecules clear **Phase I at an 80-90% success rate**, well above historic industry norms, though Phase II success (roughly **40%**) tracks closer to historical averages (pubmed.ncbi.nlm.nih.gov).

Inside large pharmaceutical organizations, AI in preclinical development increasingly means data infrastructure and generative assistants layered atop decades of study archives, exemplified by **Bayer's PRINCE platform**, built with Thoughtworks, which integrates more than **18,000 in-house safety and toxicology study reports** through a multi-agent large language model (LLM) architecture (pmc.ncbi.nlm.nih.gov). Compute investment is scaling in parallel: **Eli Lilly** and **NVIDIA** announced in October 2025 the construction of the pharmaceutical industry's most powerful supercomputer, powered by more than **1,000 NVIDIA DGX B300 GPUs** (www.prnewswire.com). This report examines how AI is reshaping predictive toxicology, organ-on-chip and digital-twin modeling, high-throughput screening, and preclinical data management; quantifies the market, regulatory, and performance evidence; profiles five real-world deployments; and outlines the implementation considerations biotech and pharmaceutical organizations should weigh as they build AI-enabled preclinical capability.

Introduction and Background

Preclinical development is the phase of the drug development pipeline in which a candidate compound, having emerged from target identification and lead optimization, is subjected to systematic evaluation of safety, pharmacokinetics (how a drug moves through the body), pharmacodynamics (how it acts on the body), and toxicology before it can be administered to humans in a Phase I clinical trial. Historically, this stage has combined in vitro (cell-based) assays, in vivo (animal) studies, and computational modeling, and it has functioned as one of the industry's most persistent bottlenecks. Traditional drug development timelines run **10 to 15 years** from initial discovery to regulatory approval, roughly **90% of candidates fail somewhere in clinical trials**, and total research and development (R&D) investment per successful drug can exceed **$2.5 billion**. A large share of that cost and time is consumed before a molecule ever reaches a human volunteer.

Artificial intelligence, and specifically machine learning (ML), a set of techniques that trains algorithms to improve at a task based on data, has been used in pockets of preclinical research for over a decade, but 2023 through 2026 marked an inflection point. The convergence of three developments explains why: first, generative AI models capable of proposing novel molecular structures rather than merely screening known ones matured rapidly following the 2024 Nobel Prize in Chemistry awarded to the developers of AlphaFold (www.prnewswire.com); second, regulators began actively codifying how AI-derived evidence could be used in submissions rather than treating it as an unregulated black box; and third, mounting political and scientific pressure to reduce reliance on animal testing created a policy tailwind for AI-based and human-relevant alternatives.

The FDA's **Center for Drug Evaluation and Research (CDER)** has documented this shift directly, noting a "significant increase in the number of drug application submissions using AI components over the past few years," spanning the **nonclinical, clinical, postmarketing, and manufacturing** phases of the drug product life cycle (www.fda.gov). CDER's January 2025 draft guidance, "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products," was built on feedback from over **800 public comments** and the agency's own experience reviewing more than **500 AI-containing submissions** from 2016 to 2023 (www.fda.gov).

Framed against these numbers, "AI in preclinical development" is best understood as sitting at the intersection of three overlapping but distinct fields that this report treats as related but not interchangeable. "AI in drug discovery" is the broader umbrella covering target identification, hit generation, and lead optimization; "machine learning in preclinical drug development" refers specifically to the predictive and generative models applied once a candidate exists and must be evaluated for safety and pharmacokinetics; and "AI enablement" or "biotech AI strategy" refers to the organizational, data-governance, and infrastructure decisions, covered at length in this report's Implications section, that determine whether either of the first two categories can actually be deployed inside a regulated R&D organization. This report surveys how AI is applied across preclinical toxicology, safety pharmacology, organ-on-chip and digital-twin modeling, high-throughput screening, and data infrastructure; grounds the discussion in the clearest available market and performance data; profiles named deployments at Insilico Medicine, Bayer, Eli Lilly, Recursion Pharmaceuticals, and Emulate; and closes with implementation guidance for biotech and pharma teams building this capability, addressing the search intents behind "AI in drug discovery," "machine learning in preclinical drug development," "biotech AI strategy," "generative AI for biotech," "AI enablement life sciences," and "preclinical AI platforms."

Predictive Toxicology, Pharmacology, and ADMET Modeling

The most mature application of AI in preclinical development is predictive toxicology: using machine learning models trained on chemical structures and historical assay results to forecast a compound's toxicity, off-target effects, and **ADMET properties** (absorption, distribution, metabolism, excretion, and toxicity) before it is ever synthesized or dosed in an animal. By flagging unsafe or ineffective candidates early, these models reduce the number of iterative animal studies a program requires.

The technical sophistication of these tools has advanced quickly. **ADMETPred**, a platform described in a March 2026 paper in *Science China Life Sciences*, was trained on a rigorously curated dataset of **120,616 compounds** and combines **189 separate models** using LightGBM, XGBoost, Random Forest, and graph attention network architectures to predict **27 distinct pharmacokinetic, metabolism, and toxicity endpoints** (link.springer.com). Separately, a July 2026 paper in *Scientific Reports* describing the **3Br-MGD** framework applies a three-branch deep encoder with meta-learning for few-shot toxicity prediction, addressing one of the field's core constraints: many toxicity endpoints have too few labeled examples for conventional deep learning (www.nature.com). Another March 2026 paper describing **META-Tox** reported that combining large-language-model-derived semantic representations with explicit topological (structural) constraints amplified fine-tuning

performance gains nearly threefold (a **5.4% gain versus 1.9%** for standard fine-tuning alone), reaching a new benchmark **area-under-the-curve (AUC) of 0.772** on an independent external dataset for in vivo toxicity prediction (doi.org).

Beyond generalized ADMET modeling, AI is being applied to specific, historically difficult toxicity endpoints. **Drug-induced liver injury (DILI)** remains, in the words of researchers publishing in *Frontiers in Toxicology* in March 2026, "a leading cause of clinical trial attrition and post-marketing drug withdrawals," in part because conventional preclinical models capture limited interindividual variability (www.frontiersin.org). That team's "clinical trial-in-a-dish" platform combines a human-serum-derived spheroid system with a proprietary AI-driven algorithm that integrates severity and incidence metrics to classify DILI risk, distinguishing dose-dependent (intrinsic) from idiosyncratic toxicity patterns that binary preclinical assays routinely miss.

At the platform level, AI-driven screening compresses the "design-make-test-analyze" (DMTA) cycle that dominates lead optimization. Vendors including **Insilico Medicine's Pharma.AI** and **Atomwise's AtomNet** use deep learning to predict biological activity and prioritize which candidate analogues are worth synthesizing, reducing the number of compounds that must be physically tested to reach a viable lead (insilico.com). Iktos markets a similar generative chemistry platform, **Makya**, paired with its **Spaya** retrosynthesis engine and in-house robotics, claiming the combined DMTA loop can shorten the discovery phase to under two years for some programs (iktos.ai). Terray Therapeutics reports that its EMMI platform has yielded unique chemistry starting points for over **90% of difficult targets screened** and advanced **eight AI-driven programs** against previously undruggable targets in the past four years (www.terraytx.ai). These are vendor-reported figures rather than independently audited benchmarks, and readers should treat them as directional claims pending third-party replication, a caveat that applies broadly across this fast-moving vendor landscape.

The broader vendor and platform landscape for AI-enabled preclinical screening has continued to diversify. **Merck KGaA, Darmstadt, Germany** (operating as MilliporeSigma in the U.S. and Canada) launched its **AIDDISON** drug discovery software in December 2023, the first software-as-a-service platform to integrate virtual molecule design with real-world manufacturability, combining machine learning, generative AI, and computer-aided drug design with the **Synthia** retrosynthesis application programming interface (API) (www.grandviewresearch.com). More recently, in November 2025, **Ginkgo Datapoints** launched the **Virtual Cell Pharmacology Initiative (VCPI)**, an open-source platform for standardized virtual cell modeling in AI drug discovery, designed to test more than **100,000 compounds** while generating over **12 billion data points** using its V-Ref293 engineered cell line and DRUG-seq RNA profiling technology (www.grandviewresearch.com). In December 2025, **Quotient Sciences** and **Intrepid Labs** announced a multi-year strategic partnership giving Quotient access to Intrepid's ANDROMEDA machine-learning model for optimizing drug product formulations, with the explicit aim of "reducing drug substance requirements in early development, shorten[ing] development cycles, increas[ing] speed to clinic" (www.grandviewresearch.com). These deals collectively illustrate that "biotech AI strategy" in the preclinical space is no longer solely about proprietary discovery engines built by a handful of AI-native biotechs; it increasingly includes narrower, task-specific machine learning models embedded into contract development and manufacturing workflows, licensed or partnered into existing R&D organizations rather than built in-house.

Two further companies illustrate the range of generative-chemistry business models now competing for pharma and biotech partnerships. **Variational AI** markets its **Enki** platform as a way to bypass conventional hit identification and hit-to-lead stages entirely, generating "optimized, synthesis-ready, lead-like compounds tailored to your target product profile" directly from a target product profile in weeks (variational.ai). **Genesis Molecular AI** takes an "operating system" framing with its **GEMS** platform, built on its **Pearl** foundation model, which the company says lets chemists "generate candidates, predict their properties, interrogate the predictions, and decide what to synthesize," with agents orchestrating routine steps so chemists focus on design decisions (www.genesis.ml). As with the screening-platform vendors discussed above, these are self-reported capabilities rather than independently benchmarked performance, and the appropriate posture for a biotech evaluating any of them is the same one recommended throughout this report: request independently reproducible evidence, not marketing collateral, before committing a discovery program to a single proprietary platform.

Digital Twins, Organ-on-Chip Systems, and Animal-Testing Alternatives

A parallel and increasingly regulator-endorsed application combines AI with **New Approach Methodologies (NAMs)**, non-animal models such as microphysiological systems (organ-on-chip platforms) and computational "digital twin" simulations of human biology. On **April 10, 2025**, the FDA announced what it called "a groundbreaking step to advance public health by replacing animal testing in the development of monoclonal antibody therapies and other drugs with more effective, human-relevant methods" (www.fda.gov). The agency's roadmap encourages developers to use **AI-based computational models of toxicity** alongside **cell-line and organoid toxicity testing** as substitutes for some animal studies, with implementation beginning immediately for Investigational New Drug (IND) applications where NAMs data is now explicitly encouraged (www.fda.gov).

Then-FDA Commissioner **Martin A. Makary** framed the shift in terms of both efficiency and ethics: "By leveraging AI-based computational modeling, human organ model-based lab testing, and real-world human data, we can get safer treatments to patients faster and more reliably, while also reducing R&D costs and drug prices. It is a win-win for public health and ethics" (www.fda.gov). The roadmap specifically envisions AI software models simulating how a monoclonal antibody distributes through the human body to predict side effects from its molecular composition, alongside organ-on-a-chip systems mimicking liver, heart, and immune-organ responses to reveal toxic effects that "could easily go undetected in animals" (www.fda.gov).

Organ-chip vendors have moved quickly to align with this regulatory opening. **Emulate, Inc.**, a leading provider of organ-chip systems, has publicly welcomed the FDA's animal-testing roadmap. Emulate's Chief Scientific Officer, Dr. Lorna Ewart, noted: "Traditional animal and reductionist models are limited because they simply do not capture the complexity of the human body. With Organ-Chips, we can explore biological mechanisms of health and disease in a setting that more closely mirrors human physiology" (www.biospace.com). In June 2025, Emulate also launched the **AVA Emulation System**, a self-contained instrument that cultures, incubates, and images up to 96 independent organ-chips simultaneously, aimed at increasing the throughput of human-relevant preclinical testing (www.ddw-online.com).

Target-discovery deals are also increasingly structured around AI platforms that generate preclinical, patient-derived evidence before a candidate ever enters the clinic. In December 2025, **Novartis** agreed to pay UK-based **Relation Therapeutics** up to **$1.7 billion** in preclinical, development, regulatory, and commercial sales milestones (plus a **$55 million** upfront cash, equity, and R&D package) for access to Relation's "Lab-in-the-Loop" platform, which the company describes as integrating "state-of-the-art AI with patient-derived multi-omic data and proprietary experimental systems to uncover causal genes and refine target hypotheses" for atopic (allergic) skin diseases (www.fiercebiotech.com). Under the deal, Relation will use patient tissue to build "functional cell atlases" of disease "to help maximize the chances of identifying successful targets before any candidate actually enters the clinic" (www.fiercebiotech.com), illustrating how AI-driven target discovery and preclinical de-risking are increasingly bundled into a single upstream platform rather than treated as sequential, separately contracted stages.

It bears noting that the shift toward NAMs and AI-based toxicology is a stated regulatory direction rather than a completed transition. Validation, standardization, and acceptance of these methods across all therapeutic modalities remains a multi-year process coordinated through the **Interagency Coordinating Committee on the Validation of Alternative Methods (ICCVAM)**, working with the National Institutes of Health, the National Toxicology Program, and the Department of Veterans Affairs (www.fda.gov). Sponsors should expect a phased rollout, beginning with a pilot program for select monoclonal antibody developers, rather than an immediate wholesale replacement of animal studies.

Preclinical Data Infrastructure and Generative AI Research Assistants

A less visible but arguably higher-volume application of AI in preclinical development is internal data infrastructure: the use of large language models (LLMs) to make decades of unstructured toxicology and safety study reports searchable, interpretable, and actionable. **Bayer AG**, in collaboration with the consultancy **Thoughtworks**, has been developing the **Preclinical Information Center (PRINCE)** since 2020, described in an August 2025 *Frontiers in Artificial Intelligence* paper as "a cloud-hosted data integration platform" that "integrates decades of structured and unstructured safety study reports, leveraging a multi-agent architecture based on Large Language Models (LLMs)" (pmc.ncbi.nlm.nih.gov). PRINCE integrates a repository of over **18,000 in-house studies**, spanning toxicology, safety pharmacology, and drug metabolism and pharmacokinetics endpoints (pmc.ncbi.nlm.nih.gov).

The platform's evolution illustrates a broader pattern of staged AI adoption inside regulated R&D organizations: a keyword-search data platform (available since a 2021 minimal viable product) evolved into a chatbot in **March 2024**, and then into a full multi-agent system in **November 2024** capable of drafting regulatory documents and designing new studies (pmc.ncbi.nlm.nih.gov). Bayer's toxicology safety lead, quoted in Thoughtworks' case study, described the chatbot's value in blunt operational terms: "The Chatbot helps to unravel the wealth of our internal knowledge, which cannot be easily retrieved by conventional query strategies using established key words and search fields... A safety-search engine is at the gates" (www.thoughtworks.com). A Bayer researcher separately reported using the tool to "find a study on DNA strand breaks with higher exposure in female rats than in male rats, which greatly aided in planning my current

genotoxicity study" (www.thoughtworks.com). Technically, PRINCE's chatbot pairs a text-to-SQL pipeline (translating natural-language questions into structured database queries against Amazon Athena, using **Claude 3.5 Sonnet**) with a retrieval-augmented generation (RAG) pipeline for unstructured study reports, embedded via OpenAI's text-embedding-3-large model into Amazon OpenSearch (pmc.ncbi.nlm.nih.gov).

This pattern, a bespoke, governed LLM layer built on top of existing regulated data systems rather than a wholesale platform replacement, is a template with implications well beyond Bayer. Roche has developed an analogous internal hub, Safety Data Integration (SDI), for the same purpose (pmc.ncbi.nlm.nih.gov), suggesting that large pharmaceutical R&D organizations are converging on similar architectures: preserve the underlying regulated system of record, and add a governed conversational and agentic layer with source-level traceability.

Data Analysis and Evidence

Quantifying the impact of AI on preclinical development requires triangulating vendor claims, independent academic analysis, and market-research estimates, each of which carries different biases. *Table 1* below summarizes the principal quantitative estimates gathered for this report, organized by originating source and category.

Metric	Value	Source	As of
Global AI in drug discovery market size	$2.3B (2025) growing to $2.9B (2026) and $13.8B by 2033, 24.8% CAGR	Grand View Research (www.grandviewresearch.com)	2025-2033 forecast
North America market share	52.8% of global revenue (2025)	Grand View Research (www.grandviewresearch.com)	2025
Time/cost savings from AI in early R&D	25-50% through the preclinical stage	Boston Consulting Group / Wellcome Trust (wellcome.org)	June 2023
Annual economic value of generative AI in pharma/medtech	$60B-$110B per year	McKinsey Global Institute (www.mckinsey.com)	January 2024
Phase I success rate, AI-discovered molecules	80-90% (versus ~52% historical industry average, per secondary sources)	Peer-reviewed clinical-pipeline analysis, PubMed (pubmed.ncbi.nlm.nih.gov)	2024
Phase II success rate, AI-discovered molecules	~40% (comparable to historical industry averages, small sample)	Same as above (pubmed.ncbi.nlm.nih.gov)	2024
Insilico rentosertib preclinical-to-Phase-I timeline	Under 30 months (versus ~6 years average); ~18-month discovery-to-candidate budget of ~$2.6M	Insilico Medicine (insilico.com)	2022 announcement
Rentosertib Phase IIa lung-function effect	+98.4 mL mean FVC change (60mg QD) vs. placebo decline	*Nature Medicine* (www.nature.com)	June 2025
Bayer PRINCE preclinical study repository	18,000+ in-house studies integrated	Frontiers in AI / PMC (pmc.ncbi.nlm.nih.gov)	August 2025

Metric	Value	Source	As of
Eli Lilly-NVIDIA supercomputer scale	1,000+ NVIDIA DGX B300 GPUs, world's first DGX SuperPOD of its kind	Eli Lilly / PR Newswire (www.prnewswire.com)	October 2025

The most important interpretive point in *Table 1* is the divergence between the BCG/Wellcome and McKinsey estimates, which are directional and modeled rather than measured, and the PubMed clinical-pipeline analysis, which tracks actual trial outcomes. The Wellcome-commissioned report is explicit about this limitation, noting that "while there are examples of early clinical successes, a more thorough analysis in the coming years will be needed to truly evaluate the impact of AI on drug discovery" (wellcome.org). Market-size figures for AI in drug discovery vary by research firm, so this report treats Grand View Research's figures as the reference case and flags that market-sizing in this nascent category is still an estimate, not a settled fact.

On regulatory throughput, the FDA's own account shows the trajectory clearly: over **500 AI-component submissions** were reviewed by CDER between 2016 and 2023 (www.fda.gov). Corporate deal activity in 2025 alone included Novartis's up to **$1.7 billion** partnership with Relation Therapeutics, Eli Lilly's **TuneLab** platform sharing models trained on over **$1 billion** of proprietary Lilly data with partner biotechs, and NVIDIA's Gefion supercomputer collaboration with Novo Nordisk, according to Grand View Research's tracked developments log (www.grandviewresearch.com). Recursion Pharmaceuticals, following its late-2024 merger with Exscientia, reported having received approximately **$450 million in upfront and realized milestone payments** from partners, with the potential for more than **$20 billion** in additional milestones, and disclosed access to over **60 petabytes** of proprietary biological, chemical, and patient-centric data (ir.recursion.com).

The preclinical toxicology data underlying these platforms also carries measurable technical performance benchmarks worth reporting alongside the market and deal figures above. ADMETPred's underlying training set spans **120,616 compounds** scored across **27 endpoints** by **189 constituent models** (link.springer.com), while the META-Tox framework's reported **AUC of 0.772** on an independent external toxicity dataset, and its threefold amplification of standard LLM fine-tuning gains (**5.4%** versus **1.9%**), represent some of the only third-party-publishable, peer-reviewed accuracy figures available for generative-AI-augmented toxicity prediction as of mid-2026, in contrast to the vendor-reported efficiency claims that dominate the commercial platform landscape (doi.org). Readers evaluating any AI toxicology tool for internal adoption should weight peer-reviewed, externally validated benchmarks such as these considerably more heavily than unaudited vendor marketing figures, even when the latter are more widely circulated in trade press.

Case Studies and Real-World Examples

Insilico Medicine: Rentosertib and an Early Phase 2a Efficacy Signal

Insilico Medicine's antifibrotic candidate, now designated **rentosertib** (formerly ISM001-055), is the most cited real-world example of AI compressing preclinical timelines. The company's own account describes bringing the drug "from start to Phase 1 in 30 months," with the target and molecule both discovered using its **Pharma.AI** platform (insilico.com). The full preclinical program, from novel target discovery to preclinical candidate

nomination, took "just under 18 months to complete at a budget of around $2.6 million," followed by an additional 12 months to progress through Phase 0 and into Phase I (insilico.com). Rentosertib is a small-molecule inhibitor of **TNIK** (Traf2- and Nck-interacting kinase), a target identified de novo using generative AI tools for **idiopathic pulmonary fibrosis (IPF)**, a progressive lung disease with a median survival of two to four years after diagnosis and no therapies that reverse its course (www.nature.com).

In its Phase 2a trial, reported in *Nature Medicine* in June 2025, patients receiving the highest dose of rentosertib (60 mg once daily) showed a mean forced vital capacity (FVC) improvement of **+98.4 mL** (95% confidence interval 10.9 to 185.9), compared with a change of -20.3 mL (95% CI -116.1 to 75.6) in the placebo group, in a 71-patient, 12-week study whose primary endpoint was safety (FVC was a secondary endpoint) (www.nature.com). Treatment-emergent adverse events were reported in comparable proportions across dose arms and placebo (72.2% to 83.3% versus 70.6% for placebo). This case matters for the preclinical AI conversation specifically because it is one of the only publicly reported instances where an AI-identified target and an AI-designed molecule have together produced a preliminary human efficacy signal in a small, early-phase trial published in a peer-reviewed journal, providing an evidentiary anchor for otherwise largely preclinical or vendor-reported claims. It should also be read alongside the caveat below: not every headline AI drug-discovery program has cleared this bar. Rentosertib has since moved into a Phase III trial (NCT07687459, estimated enrollment 320, actual start date 9 September 2026), for which no efficacy results are yet available.

Rentosertib's affected patient population underscores why the preclinical timeline compression matters beyond a single company's business case. IPF affects between **10 and 60 per 100,000 individuals** in the United States, with an approximately tenfold greater incidence in people over 65, occurring at a rate similar to stomach, brain, and testicular cancers, and median survival after diagnosis of only **two to four years** despite existing antifibrotic therapies (www.nature.com). A preclinical development program compressed from an industry-typical multi-year timeline to under two years, as Insilico reports for the target-to-candidate phase of this program, therefore represents a meaningfully faster path to testing new mechanisms in a disease area where existing drugs only slow progression rather than reverse it.

Ainnocence and the Broader Chinese and Vertically Integrated Platform Landscape

Beyond the US- and Europe-centered examples above, AI-native platform companies have also emerged with large-scale, vertically integrated computational claims. **Ainnocence**, describing itself as an "AI-driven pharmaceutical and material IP-generating platform company," reports running **10 billion virtual screens per day**, having generated **over 100 drug design projects**, and having filed more than **20 Patent Cooperation Treaty (PCT) patents**, built on a **1-billion-parameter foundation model** trained on 16 years of proprietary computational and wet-lab data (ainnocence.com). As with the other platform vendors profiled in this report, these figures are self-reported rather than independently benchmarked, and should be read as an indicator of the scale at which some AI-native drug discovery companies now claim to operate rather than as validated performance data.

Bayer and Thoughtworks: PRINCE as an Internal Preclinical Knowledge Engine

Bayer's PRINCE platform illustrates a different, less headline-grabbing but arguably more representative application: using generative AI not to design new molecules, but to make an existing multi-decade archive of preclinical safety data usable. Developed with Thoughtworks starting in 2020, PRINCE evolved through three stages, described by its developers as "Search," "Ask," and "Do," culminating in a multi-agent system able to draft regulatory documents such as IND submissions (pmc.ncbi.nlm.nih.gov). The platform has been applied to real Bayer research workflows, including cross-study vehicle toxicity analysis that previously required significant manual effort from researchers (pmc.ncbi.nlm.nih.gov). Because the system retains a human-in-the-loop review step and cites its original source documents for every answer, it addresses one of the most persistent objections regulated organizations raise about generative AI: unverifiable, non-traceable outputs.

Eli Lilly and NVIDIA: Industrial-Scale Compute for Preclinical Modeling

In October 2025, Eli Lilly announced it was building, in collaboration with NVIDIA, "the most powerful supercomputer owned and operated by a pharmaceutical company," an "AI factory" managing the full AI lifecycle from data ingestion through fine-tuning and high-volume inference (www.prnewswire.com). The system, the world's first NVIDIA DGX SuperPOD built with DGX B300 systems, is powered by more than **1,000 B300 GPUs** on a unified networking fabric (www.prnewswire.com). Lilly's chief information and digital officer, Diogo Rau, framed the investment around proprietary data rather than raw compute alone: "As a 150-year-old medicine company, one of our most powerful assets is decades of data. With purpose-built AI models and AI, we can set a new scientific standard that accelerates innovation to deliver medicines to more patients, faster" (www.prnewswire.com). Select models trained on this infrastructure are being made available to partner biotechs through Lilly's **TuneLab** federated-learning platform, allowing external companies to fine-tune Lilly-trained models on their own data without sharing raw datasets (www.prnewswire.com).

Recursion Pharmaceuticals and Exscientia: Consolidation

The late-2024 merger of Recursion Pharmaceuticals and Exscientia, described in the combined company's own announcement as bringing together "two leaders in the AI drug discovery space," created an entity with more than **10 clinical and preclinical programs**, roughly **10 partnered programs**, and access to over **60 petabytes** of proprietary data (ir.recursion.com). Recursion's CEO, Chris Gibson, stated the combined companies had received "over $450M in upfront and realized milestone payments received from partners to date out of more than $20B possible" (ir.recursion.com).

Emulate, Inc.: Organ-Chip Technology and the Animal-Testing Roadmap

Emulate's experience illustrates how a non-AI-native but AI-adjacent preclinical technology, organ-on-a-chip systems, is part of the broader shift toward non-animal methods that also includes computational toxicology. Following the FDA's April 2025 animal-testing roadmap, Emulate's CEO, Jim Corbett, stated: "By embracing Organ-Chips and other innovative, human-based methods, the FDA is simultaneously promoting patient safety, accelerating the delivery of new therapies, and reducing animal use. We believe this is the start of a new paradigm" (www.biospace.com).

Implications and Future Directions

The near-term trajectory for AI in preclinical development runs along three converging tracks: regulatory formalization, computational scale, and disciplined skepticism about clinical translation. On the regulatory front, the FDA's evolution from a discussion paper (May 2023) to a draft guidance with a formal seven-step credibility assessment framework (January 2025) to a further "Guiding Principles" document (January 2026) signals that AI use in nonclinical and clinical development is receiving more structured regulatory attention, although the January 2025 guidance remains a draft (www.fda.gov). Biotech and pharma teams building AI-enabled preclinical workflows should expect this credibility framework, not raw model performance metrics alone, to become the operative currency in regulatory interactions.

On computational scale, the direction of travel is toward proprietary, vertically integrated infrastructure: Eli Lilly's supercomputer, NVIDIA's parallel collaboration with Novo Nordisk on the Gefion system, and Recursion's in-house compute and wet-lab scale all point toward a bifurcation between a handful of well-capitalized players building foundation-scale infrastructure and a longer tail of biotechs accessing similar capability through federated or API-based partnerships, such as Lilly's TuneLab. This bifurcation has direct strategic implications: smaller biotechs will increasingly need to decide between building bespoke discovery capability, partnering with a platform vendor, or accessing shared infrastructure, a decision that should be driven by therapeutic focus and data assets rather than a generic desire to "have an AI strategy."

Table 2 below summarizes the four broad archetypes of AI-enabled preclinical capability referenced throughout this report, to help organizations orient a build-partner-buy decision against their own data assets and regulatory posture.

Archetype	Representative Examples	Primary Value Proposition	Typical Adopter
Generative small-molecule design platforms	Insilico Medicine Pharma.AI, Iktos Makya, Variational AI Enki, Genesis Molecular AI GEMS	De novo molecule generation against a target product profile, compressing hit-to-lead cycles	AI-native biotechs; large pharma via partnership or license
Predictive ADMET and toxicology models	ADMETPred, ADMET-AI, admetSAR3.0, META-Tox, 3Br-MGD	In silico safety and pharmacokinetic screening before synthesis or animal dosing	Discovery and preclinical teams across biotech and pharma

Archetype	Representative Examples	Primary Value Proposition	Typical Adopter
Organ-chip and NAMs hardware	Emulate Liver-Chip and AVA system, CN Bio microphysiological systems	Human-relevant toxicity data as an animal-testing alternative	Preclinical safety and toxicology functions, especially for DILI-prone modalities
Internal data infrastructure and generative research assistants	Bayer PRINCE, Roche SDI	Making decades of proprietary structured and unstructured safety data searchable, explainable, and reusable	Large pharma R&D organizations with substantial legacy study archives

Each archetype in *Table 2* carries a different build-partner-buy calculus. Generative design platforms and predictive ADMET models are increasingly available as licensed software or partnership arrangements, making them accessible to biotechs without in-house AI teams, whereas organ-chip hardware requires physical lab infrastructure and sustained regulatory engagement to translate into submission-ready evidence. Internal data infrastructure, by contrast, is inherently bespoke: PRINCE could not simply be purchased off the shelf because its core asset, Bayer's own multi-decade study archive, is not transferable, meaning any organization pursuing this fourth archetype should expect a multi-year internal build regardless of how mature the underlying LLM technology becomes.

On clinical translation, the honest picture as of mid-2026 is one of real but uneven progress, a theme this report has returned to across the Insilico (www.nature.com), Recursion-Exscientia (ir.recursion.com), and Bayer PRINCE (www.thoughtworks.com) case studies. The peer-reviewed clinical-pipeline analysis showing an **80-90% Phase I success rate** for AI-discovered molecules is a genuinely strong signal about AI's capacity to design or identify molecules with drug-like properties (pubmed.ncbi.nlm.nih.gov), and it is corroborated by rentosertib's own progression from IND-enabling studies through a favorable Phase 0 microdose result to encouraging preliminary Phase IIa results. But Phase II success rates for the same cohort of companies sit closer to historical industry norms. The reasonable inference is that AI's advantage in preclinical development is currently strongest at generating and de-risking candidates with favorable early safety and pharmacokinetic profiles, a genuinely valuable capability, rather than at solving the harder, longer-standing problem of predicting human efficacy, which remains gated by biology the field does not yet fully understand regardless of computational method.

For biotech and pharmaceutical organizations evaluating how to build AI capability into preclinical operations, several considerations recur across the case studies in this report:

- **Data governance precedes model sophistication.** Bayer's PRINCE platform has been under development since 2020, with a minimal viable product in 2021, a chatbot in March 2024 and a multi-agent system in November 2024, progressing from a searchable data repository to a multi-agent research assistant (pmc.ncbi.nlm.nih.gov); the foundational data aggregation and governance work, not the LLM layer itself, consumed the bulk of that time (pmc.ncbi.nlm.nih.gov).
- **Source-level traceability is non-negotiable in regulated environments.** Every credible internal AI deployment profiled in this report, from PRINCE's citation of original study paragraphs (www.thoughtworks.com) to the FDA's credibility framework itself (www.fda.gov), treats explainability and provenance as a prerequisite rather than an optional feature.

- **Vendor-reported performance claims require independent scrutiny.** Figures such as Terray's "90% of difficult targets" (www.terraytx.ai) or Ainnocence's "80% cost reduction" (ainnocence.com) are useful directional indicators but are not third-party validated benchmarks, and should be treated with the same skepticism applied to any single-source marketing claim, unlike the peer-reviewed META-Tox and PubMed clinical-pipeline figures cited above (doi.org) (pubmed.ncbi.nlm.nih.gov).
- **Regulatory engagement should begin early and be proportionate to model risk.** The FDA's own guidance explicitly favors early sponsor engagement and risk-tiered validation over a one-size-fits-all documentation burden (www.fda.gov).
- **Organizational readiness, not just tooling, determines outcomes.** The FDA's 2025 NAMs roadmap creates regulatory permission to reduce animal testing (www.fda.gov), but translating that roadmap into practice requires internal validation capability, cross-functional buy-in.

For consultancies and technology advisors operating adjacent to this space, the durable opportunity is less about building a competing discovery platform and more about the unglamorous but essential work of data engineering, regulatory-aligned governance, and integration between generative AI tools and the validated systems of record that pharmaceutical R&D already depends on.

Frequently Asked Questions (FAQs)

What is AI actually used for in preclinical drug development? AI is primarily used for predictive toxicology and ADMET modeling (forecasting a compound's safety and pharmacokinetic properties before synthesis) (www.fda.gov), high-throughput virtual screening and lead optimization (www.terraytx.ai), digital-twin and organ-on-chip data analysis (www.fda.gov), and increasingly for internal knowledge management, using large language models to search and synthesize decades of safety study reports, as demonstrated by Bayer's PRINCE platform (pmc.ncbi.nlm.nih.gov).

Is AI in drug discovery the same as AI in preclinical development? Not exactly. "AI in drug discovery" typically spans the earlier, upstream stages of target identification and hit-to-lead generation, and Grand View Research reports that the drug optimization and repurposing application segment "led the market with the largest revenue share of 52.4% in 2025" (www.grandviewresearch.com), while "AI in preclinical development" refers more specifically to the safety, toxicology, and pharmacokinetic evaluation stage that follows candidate selection and precedes first-in-human trials (www.fda.gov). In practice, the same generative and predictive AI platforms, such as Insilico's Pharma.AI, are often used across both stages (www.nature.com).

Has the FDA approved any AI-discovered drugs? As of mid-2026, no AI-discovered drug had completed the full FDA approval process, though Insilico Medicine's rentosertib had reported an encouraging preliminary signal from a small, 12-week phase 2a study published in *Nature Medicine*, whose primary endpoint was safety (www.nature.com) (pubmed.ncbi.nlm.nih.gov).

Is AI reducing the need for animal testing in preclinical studies? The FDA's April 2025 roadmap explicitly promotes AI-based computational models alongside organoid and organ-on-chip systems as substitutes for some animal testing, beginning with monoclonal antibodies (www.fda.gov), but implementation is proceeding

through pilot programs and phased validation rather than an immediate wholesale replacement, coordinated through ICCVAM together with the National Institutes of Health and the National Toxicology Program (www.fda.gov).

How should a biotech company approach an AI enablement strategy for preclinical R&D? Available evidence suggests prioritizing data governance and infrastructure before model sophistication, following the multi-year path Bayer took with PRINCE (pmc.ncbi.nlm.nih.gov), insisting on source-level traceability for any generative AI outputs used in regulatory-facing work (www.thoughtworks.com), engaging regulators early with a risk-tiered validation plan aligned to the FDA's credibility framework (www.fda.gov), and treating vendor performance claims, including market-size and success-rate figures cited throughout this report, with appropriate scrutiny given the wide variance across research firms and vendors (www.grandviewresearch.com).

What does "generative AI for biotech" mean in a preclinical context, as distinct from predictive AI? Generative AI creates new candidate molecules, protein structures, or synthetic study data rather than merely classifying or predicting outcomes for existing compounds, as illustrated by platforms like Iktos's Makya and Genesis's GEMS that propose novel, synthesizable structures rather than only scoring existing ones (iktos.ai) (www.genesis.ml).

Conclusion

AI in preclinical development has progressed from an experimental add-on to a documented, regulator-engaged, and commercially consequential layer of pharmaceutical R&D. The evidence assembled in this report points to genuine acceleration in specific, well-bounded tasks: predictive toxicology and ADMET modeling, which the FDA's April 2025 roadmap places among the approaches intended to reduce reliance on some animal testing (www.fda.gov), high-throughput screening platforms compressing design-make-test cycles (iktos.ai), and generative AI research assistants that make decades of internal safety data usable for the first time (pmc.ncbi.nlm.nih.gov). Insilico Medicine's rentosertib program stands as a prominent published example of these capabilities compounding into an encouraging preliminary signal from a small, 12-week phase 2a study whose primary endpoint was safety (www.nature.com), while the more measured Phase II performance of AI-discovered molecules broadly (pubmed.ncbi.nlm.nih.gov) is an equally important data point against uncritical enthusiasm. The FDA's parallel construction of a formal AI credibility framework (www.fda.gov) and its April 2025 roadmap toward reduced animal testing (www.fda.gov) indicate that regulatory infrastructure is being built in step with, rather than lagging far behind, the technology itself. For biotech and pharmaceutical organizations, and for the consultancies and advisors who support them (www.grandviewresearch.com), the practical task ahead is less about chasing the most impressive-sounding model and more about the disciplined work of data governance, source traceability, and proportionate validation that turns AI's demonstrated preclinical potential into evidence regulators and clinicians can trust.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-drug-discovery-market

2. https://wellcome.org/insights/reports/unlocking-potential-ai-drug-discovery
3. https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality
4. https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development
5. https://www.fda.gov/news-events/press-announcements/fda-announces-plan-phase-out-animal-testing-requirement-monoclonal-antibodies-and-other-drugs
6. https://insilico.com/phase1
7. https://www.nature.com/articles/s41591-025-03743-2
8. https://pubmed.ncbi.nlm.nih.gov/38692505
9. https://pmc.ncbi.nlm.nih.gov/articles/PMC12401995
10. https://www.prnewswire.com/news-releases/lilly-partners-with-nvidia-to-build-the-industrys-most-powerful-ai-supercomputer-supercharging-medicine-discovery-and-delivery-for-patients-302597285.html
11. https://www.prnewswire.com/il/news-releases/converge-bio-raises-25m-to-bring-generative-ai-drug-discovery-to-every-biotech-and-pharmaceutical-company-302659690.html
12. https://link.springer.com/article/10.1007/s11427-025-3166-8
13. https://www.nature.com/articles/s41598-026-60814-1
14. https://doi.org/10.1021/acs.jpclett.6c00635
15. https://www.frontiersin.org/journals/toxicology/articles/10.3389/ftox.2026.1740791/full
16. https://iktos.ai
17. https://www.terraytx.ai/emmi-platform
18. https://variational.ai
19. https://www.genesis.ml/ai-platform
20. https://www.biospace.com/press-releases/emulate-applauds-fdas-roadmap-to-reduce-animal-testing-and-embrace-organ-chip-technologies
21. https://www.ddw-online.com/emulate-launches-unique-high-throughput-organ-chip-platform-35326-202506
22. https://www.fiercebiotech.com/biotech/novartis-pens-17b-dermatology-pact-ai-enabled-british-biotech
23. https://www.thoughtworks.com/en-us/clients/bayer-ag
24. https://ir.recursion.com/news-releases/news-release-details/recursion-and-exscientia-two-leaders-ai-drug-discovery-space
25. https://ainnocence.com

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.