IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Insilico Pharma.AI: AI Drug Discovery Platform Analysis

> 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 · 2026-04-16 · ai drug discovery · insilico medicine · pharma.ai · pandaomics · chemistry42 · mmai gym · target discovery · generative chemistry

Original article: https://intuitionlabs.ai/articles/insilico-pharma-ai-drug-discovery



In this article

1. Executive Summary
2. Introduction and Background
3. AI in Drug Discovery: From Hype to Practice
4. Insilico Medicine and the Pharma.AI Platform
5. PandaOmics: AI-Powered Target & Biomarker Discovery
6. Description and Functionality
7. Advances and Updates (2023–2026)
8. Case Study: TNIK in Fibrosis (Rentosertib)
9. Data and Usage
10. Summary of PandaOmics
11. Chemistry42: AI-Driven Molecular Design
12. Platform Overview
13. Generative Models and Ensemble Approach
14. Capabilities and Applications
15. Performance and Benchmarks
16. Discussion
17. Summary of Chemistry42
18. The MMAI Gym: Specializing Foundation Models for Science
19. Concept and Motivation
20. Architecture: Liquid Foundation Models
21. Performance on Drug Discovery Benchmarks
22. The MMAI Gym Platform and Process
23. Position in the AI Landscape
24. Summary of MMAI Gym
25. Data Analysis and Evidence
26. Insilico's Business Growth and Pipeline
27. Benchmark Comparisons and Performance Claims
28. Synthesis: Evidence-Based Assessment
29. Case Studies and Real-World Examples
30. IPF and Rentosertib (TNIK Inhibitor)
31. GLP-1R Peptide Design (Generative Biologics)
32. Pipeline Examples (Other Indications)
33. External Perspectives
34. Lessons and Implications
35. Discussion: Broader Perspectives and Future Directions
36. AI's Role and Limitations
37. Regulatory and Ethical Outlook
38. Future Evolution of Pharma.AI
39. Competitive and Collaborative Landscape
40. Risks and Reservations

41. Conclusion
42. References

Executive Summary

Insilico Medicine's **Pharma.AI** platform exemplifies the cutting edge of AI-driven drug discovery in 2026, integrating advanced modules for **target discovery (PandaOmics)**, **molecular design (Chemistry42)**, and **LLM-based reasoning (MMAI Gym)**. Its approach reflects the paradigm of "pharmaceutical superintelligence" – an autonomous AI system spanning from gene and pathway analysis to de novo molecule generation and beyond (insilico.com) (pmc.ncbi.nlm.nih.gov). In 2025 Insilico achieved significant commercial traction (serving 13 of the top 20 global pharmas, ~ +24% YoY software revenue growth) and pipeline expansion (28 nominated preclinical candidates, 10 in clinical trials) (insilico.com). The June 2025 Nature Medicine publication reported preliminary Phase IIa findings for **rentosertib** (ISM001-055), an AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis. After 12 weeks, mean FVC change was +98.4 mL with 60 mg once daily versus −20.3 mL with placebo; FVC was a secondary endpoint, and larger, longer trials are needed to evaluate efficacy and safety (Nature Medicine).

This report provides an in-depth examination of Insilico Medicine's AI platform and its key modules as of 2026, focusing on **PandaOmics**, **Chemistry42**, and the **MMAI (Multi-Modal AI) Gym**. We trace their development and capabilities, analyze evidence of effectiveness (including benchmarks and case studies), and place them in the broader context of AI in drug discovery. The analysis incorporates peer-reviewed research, industry news, and technical documentation to offer a balanced assessment of achievements and challenges. In summary, PandaOmics leverages multimodal omics and text analytics to generate actionable target and biomarker hypotheses (pmc.ncbi.nlm.nih.gov); Chemistry42 orchestrates over 40 generative AI models (plus physics-based tools) to design novel drug-like molecules (pubs.acs.org); and the MMAI Gym represents a novel "specialized foundation model" approach, fine-tuning compact LLMs on pharma datasets to surpass larger models on drug tasks (insilico.com). Our findings highlight the sophistication of these systems and their early successes (such as PandaOmics-enabled target ID and Chemistry42-led molecule design) while also noting that most AI-generated candidates are still in early stages, emphasizing the need for continued validation. Finally, we discuss future implications, including the trend toward domain-specific AI models, regulatory updates (e.g. FDA's AI guidance), and the integration of AI into end-to-end R&D workflows.

Introduction and Background

AI in Drug Discovery: From Hype to Practice

The pharmaceutical industry faces immense challenges: drug development remains slow, costly, and failure-prone. Even with massive R&D investments, bringing a new drug to market still typically takes over a decade and costs upward of billions of dollars (apnews.com) (digitaledition.chaindrugreview.com). Artificial intelligence (AI) has long been seen as a potential accelerator. Early AI efforts focused on virtual screening and predicting molecular properties, but the advent of deep learning (DL) and generative models over the last decade has

enabled **de novo design** of molecules and advanced data integration. Insilico Medicine and other startups (e.g. Exscientia, BenevolentAI, Atomwise, Insitro) spearheaded generative chemistry methods in the mid-2010s, culminating in key demonstrations such as the design of novel DDR1 kinase inhibitors in 2019 (communities.springernature.com).

By the early 2020s, platforms began to combine multiple AI capabilities into unified drug discovery suites. Insilico's **Pharma.AI** (sometimes stylized pharma.ai) is one such ecosystem, aiming for "pharmaceutical superintelligence" – an end-to-end AI that can autonomously propose and optimize drug hypotheses (insilico.com). For an early-discovery comparison, Insilico reported that its 22 candidates nominated during 2021–2024 took an average of 12–18 months from project initiation to preclinical candidate nomination, compared with a company-cited traditional range of 2.5–4 years (Insilico Medicine). This company-reported comparison concerns preclinical candidate nomination; it does not measure time to approval.

Nevertheless, experts urge caution. The complex biology of disease means that AI models must be paired with human insight and robust validation. Industry analysts emphasize frameworks for "human-in-the-loop" oversight and risk mitigation, noting that while AI can speed processes, it also introduces challenges in data bias, interpretability, and validation (digitaledition.chaindrugreview.com). Regulatory agencies, too, are adapting: for instance, in early 2025 the FDA drafted guidance on AI in drug development and even reported pilot tests where generative AI reduced regulatory review tasks from days to minutes (www.fda.gov).

In this context, Insilico Medicine's offering merits close analysis. Founded by Alex Zhavoronkov, Insilico has pursued AI-driven drug discovery since the 2010s. By 2025 it had become a **clinical-stage biotech (HKEX: 3696)** with a dual strategy of licensing its **Pharma.AI** platform to partners and simultaneously advancing its own pipelines. This report examines Insilico's platform in depth, focusing specifically on **PandaOmics**, **Chemistry42**, and **MMAI Gym** – key components that illustrate the company's technological approach.

Insilico Medicine and the Pharma.AI Platform

Insilico Medicine's strategy is encapsulated in its **"AI+Drug Discovery" dual-engine model**, which combines an in-house pipeline of drug candidates with the commercialization of its AI software (insilico.com).The company has undergone rapid growth: in 2025 it reported **US$56.24 million** in revenue (from contracts and milestones with pharma partners and collaborations) (insilico.com) and in December 2025 completed a major Hong Kong IPO, raising funds to support further R&D (insilico.com). Notably, Insilico stated that its platform now **serves 13 of the top 20 global pharmaceutical companies** – a testament to industry interest (insilico.com).

The Pharma.AI platform is branded as an **"end-to-end generative AI platform"** that encompasses multiple modules:

- **PandaOmics (Biology42)**: an AI engine for multi-omics data analysis to identify disease targets and biomarkers (pmc.ncbi.nlm.nih.gov).
- **Chemistry42**: a generative chemistry suite that designs novel small molecules with user-defined properties (pubs.acs.org).
- **Science42**: including tools like DORA (an AI assistant for scientific writing) and others for research acceleration (beyond the scope of this report).

- **Generative Biologics**: AI-driven design of peptides, antibodies, and other biologics.
- **inClinico**: predictive platform for modeling clinical trial outcomes.
- **MMAI Gym (Sci MMAI Gym)**: a new framework (2025–26) for training/modifying foundation language models to excel in drug discovery tasks (insilico.com).

Together, these aim to create a loop: novel target/hypothesis generation (PandaOmics) feeds into molecule/biologics design (Chemistry42, Generative Biologics), whose outputs can be experimentally validated (e.g. using Insilico's lab automation), with results feeding back into models and even into clinical prediction (inClinico). The ultimate goal is a **"Pharmaceutical Superintelligence"** – a system capable of autonomously conceiving and optimizing therapies (insilico.com).

Table 1 below summarizes the core Pharma.AI modules under discussion in this report, highlighting their purpose, key techniques, and milestone achievements as reported by Insilico and collaborators.

Module	Primary Function	Techniques / Models	Key Achievements (to 2026)
PandaOmics	Target & biomarker discovery from omics & text data (pmc.ncbi.nlm.nih.gov).	AI-powered analytics on multi-omics (gene expression, proteomics, etc.) and biomedical literature; knowledge graphs; LLM scoring; disease-specific ML models (pmc.ncbi.nlm.nih.gov) (insilico.com).	Generated novel target hypotheses validated in vitro/in vivo; integrated 23 disease models; introduced new LLM-derived scores (confidence, tractability, druggability, mechanism clarity) for prioritization (insilico.com).
Chemistry42	De novo small-molecule design and optimization (pubs.acs.org).	Ensemble of 40+ generative DL models (autoencoders, GANs, flows, evolutionary, LMs) plus reinforcement learning and filtering (pubs.acs.org); 2D/3D scoring modules (ADMET predictors, molecular docking, etc.); MDFlow (molecular dynamics simulation) (insilico.com); multiagent RL feedback loops.	Used by ≥20 pharma companies; produced >2400 candidate molecules in hours (insilico.com); powered generation of DDR1 inhibitor leads (Nature Biotech 2019) (communities.springernature.com); added Nach01 multimodal chem-language foundation model (AWS Marketplace) (insilico.com); achieved deep optimization tasks in collaboration pipeline.
MMAI Gym (Science MMAI)	LLM training for drug R&D tasks (insilico.com).	Customized fine-tuning "gym" using Insilico's proprietary data and benchmarks (thousands of pharma tasks) (insilico.com); multi-stage "curriculum" in chemistry, biology, clinical development (insilico.com); partnership with Liquid AI to train Liquid Foundation Models (LFMs).	Produced LFM2-2.6B-MMAI, a 2.6B-parameter model achieving **state-of-the-art** results on drug discovery benchmarks (insilico.com); reported task-performance gains; enabled on-premise AI deployment for pharma data security (insilico.com).

Table 1. Overview of key Pharma.AI modules (2026) and their capabilities. References indicate technical descriptions and milestone reports (e.g. Insilico publications and press releases). The modules collectively support the AI-driven drug discovery pipeline from target ID to molecule design.

This report will analyze each of the above modules in depth, along with the overarching platform structure and outcomes. We draw on Insilico's peer-reviewed publications (e.g. *J. Chem. Inf. Model.* articles on PandaOmics and Chemistry42), official announcements, independent news coverage (e.g. *Nature Medicine* report on rentosertib), and industry analyses. We also consider critical perspectives on AI drug discovery to present a balanced view of progress and challenges.

PandaOmics: AI-Powered Target & Biomarker Discovery

Description and Functionality

PandaOmics is Insilico Medicine's engine for mining omics data to yield therapeutic **targets and biomarkers**. Published in *Journal of Chemical Information and Modeling* in 2024, the platform is described as **"a cloud-based software platform that applies artificial intelligence and bioinformatics techniques to multimodal omics and biomedical text data for therapeutic target and biomarker discovery"** (pmc.ncbi.nlm.nih.gov). In practice, PandaOmics aggregates diverse datasets (gene expression, proteomics, methylation, etc.) relevant to a disease, performs data harmonization and statistical analysis, then uses AI algorithms to rank candidate genes or pathways by their disease association and druggability.

The PandaOmics pipeline begins with **dataset selection and group comparisons** (pmc.ncbi.nlm.nih.gov). Users (or automatic workflows) input disease vs control sample sets, possibly from multiple studies. The system supports multi-omics: e.g. it maps methylation and proteomics data to gene-level features (pmc.ncbi.nlm.nih.gov). After standard preprocessing (batch correction, PCA/UMAP visualization), PandaOmics identifies differentially expressed genes and perturbed pathways. It employs 23 *disease-specific machine learning models* to score each gene's relevance/association. Importantly, PandaOmics also integrates **prior knowledge**: for example, it uses literature mining and knowledge graphs to link genes to mechanisms, and it can run Chat-style LLM explanations ("ChatPandaGPT") on gene-disease associations (pmc.ncbi.nlm.nih.gov).

The platform's key output is a **ranked list of candidate targets and biomarkers**. These candidates take into account multiple criteria: disease relevance, novelty, tractability, safety (from known biology), and predicted effect size. PandaOmics allows filter combinations such as minimal expression change, presence in multiple datasets, and so on, to refine the list. The system also computes meta-analyses across studies to boost statistical power. Crucially, it ties back to experimental validation: Insilico has connected PandaOmics output to wet-lab automation ("robotic lab" in Fig. 1) such that target hypotheses can be tested in vitro (pmc.ncbi.nlm.nih.gov).

Petrina Kamya *et al.* (Insilico researchers) demonstrated PandaOmics by showing its ability to recapitulate known targets and suggest new ones (they report prior in vitro/in vivo validations of PandaOmics-generated hypotheses (pmc.ncbi.nlm.nih.gov)). They also emphasize that PandaOmics is now a **core component of the broader Pharma.AI suite** (pmc.ncbi.nlm.nih.gov), working hand-in-hand with Chemistry42 and others. For example, once PandaOmics proposes a target protein, Chemistry42 can be tasked to design small molecules against that target.

Advances and Updates (2023–2026)

Since its initial launch, PandaOmics has been continuously upgraded. Notably, in 2025 Insilico introduced new **LLM-based scoring metrics** to improve prioritization. A Pharma.AI webinar recap reports that PandaOmics added four novel scores derived from large language models: **confidence, commercial tractability, druggability, and mechanism clarity** (insilico.com). These scores help address common concerns in AI-driven target ID by assessing how biologically plausible and practically targetable a candidate is. For instance, a “mechanism clarity” score favors genes with well-understood disease mechanisms (aiming to avoid opaque predictions) (insilico.com).

Insilico also presented a **benchmarking initiative (TargetBench 1.0 / TID-Pro)** to systematically evaluate target ID algorithms (insilico.com). This reflects an industry trend toward benchmarking AI models with reproducible metrics. It suggests PandaOmics is part of a more rigorous validation workflow, not just a black-box tool.

PandaOmics’s user interface has been improved for corporate environments – e.g. virtual private cloud deployment for data security (insilico.com). The platform supports **cross-dataset meta-analysis and data harmonization**, enabling researchers to aggregate multiple studies on the same disease (a feature highlighted in media coverage (www.clinicalresearchnewsonline.com)). These capabilities are critical for producing robust targets from noisy omics data.

Case Study: TNIK in Fibrosis (Rentosertib)

The published discovery study documents PandaOmics analysis of lung and kidney fibrosis datasets to generate target hypotheses, followed by Chemistry42 generation of small-molecule leads targeting TNIK (Nature Biotechnology). These platform contributions are documented rather than inferred.

In the 12-week randomized Phase IIa trial, 71 patients received placebo (n = 17), 30 mg rentosertib once daily (n = 18), 30 mg twice daily (n = 18), or 60 mg once daily (n = 18). Safety and tolerability were the primary focus; FVC was a secondary endpoint. Mean FVC change was +98.4 mL (95% CI 10.9 to 185.9) with 60 mg once daily and −20.3 mL (95% CI −116.1 to 75.6) with placebo. Small groups, geographically and demographically homogeneous participants in China, and short follow-up limit conclusions about long-term efficacy and safety. Exploratory protein changes supported the TNIK hypothesis; they do not establish a definitive causal mechanism or durable patient benefit (Nature Medicine).

Data and Usage

PandaOmics is offered as a licensed software platform (on-premise or private cloud). Insilico neither discloses user numbers publicly, but reports that Pharma.AI (including PandaOmics) is deployed with many top pharma firms (insilico.com). In general, multi-omics AI tools are gaining traction in industry partnerships, often in early R&D or precompetitive consortia. By enabling hypothesis generation from public and proprietary omics, PandaOmics accelerates the usual target discovery timeline.

Independent perspective: External analyses of AI in target discovery are still emerging. A *Chain Drug Review* industry article (2024) highlights AI's potential to process vast data for target ID and lead optimization (digitaledition.chaindrugreview.com), but also underscores that such systems need careful oversight (e.g. human validation loops) (digitaledition.chaindrugreview.com). Insilico's addition of "mechanism clarity" scoring (insilico.com) is a practical response to this need for interpretability.

Limitations: No system is perfect. PandaOmics relies on available data; rare or poorly characterized diseases with sparse datasets may yield less reliable targets. Its predictions still require experimental validation (though Insilico's integration with robotics helps close this loop). Additionally, machine learning models can inherit biases from the input data. Insilico's publications acknowledge such risks by offering confidence and clarity metrics, but in practice companies must still vet AI-generated targets with traditional biology expertise (insilico.com) (digitaledition.chaindrugreview.com).

Summary of PandaOmics

PandaOmics represents a mature AI-driven approach to target discovery, merging omics analytics with AI ranking. Key features include:

- **Multi-omics Integration:** Simultaneous analysis of transcriptomics, proteomics, methylation and more (pmc.ncbi.nlm.nih.gov).
- **Custom Analytics Pipeline:** End-to-end workflow from data selection through gene-level analysis to meta-analysis (pmc.ncbi.nlm.nih.gov).
- **AI and Domain Knowledge:** Combines ML models, knowledge graphs, and LLM tools (e.g. "ChatPandaGPT" for gene insights) (pmc.ncbi.nlm.nih.gov).
- **Iterative Validation:** Targets link to lab automation and follow-up (feedback loop) (pmc.ncbi.nlm.nih.gov).
- **Recent Enhancements:** New prioritization scores (e.g. druggability, tractability) improve candidate selection (insilico.com).

By 2026 PandaOmics has arguably become one of the more comprehensive target discovery platforms. Its successful application to relevant cases like TNIK/IPF provides strong evidence that the approach can uncover novel biology. However, as with any AI, its outputs should be interpreted in consultation with scientists, and require the usual experimental confirmation before clinical translation.

Chemistry42: AI-Driven Molecular Design

Platform Overview

Chemistry42 is Insilico's flagship AI platform for **small-molecule design and optimization**. It launched in 2020; a description of the platform appeared in the *Journal of Chemical Information and Modeling* in 2023 (pubs.acs.org). According to Ivanenkov *et al.*, "Chemistry42 is a platform that connects state-of-the-art

generative AI algorithms with medicinal and computational chemistry expertise and best...practices" (pubs.acs.org). In practice, the platform orchestrates a large ensemble of generative models to propose novel chemical structures tailored to user-specified criteria (e.g. target binding, ADMET properties, novelty).

Key facts from the literature: Chemistry42 *launched in 2020* and by 2023 had been used in over **20 pharmaceutical companies**, supporting more than 15 external projects and over 30 internal programs (pubs.acs.org). Its core workflow involves three phases (Figure 1 in (pubs.acs.org)):

1. **Generation:** An ensemble of *40+ generative models* runs in parallel. These models have diverse architectures (autoencoders, GANs, flow-based, evolutionary algorithms, language models, etc.) and representations (string, graphs, 3D) (pubs.acs.org). Each model explores chemical space under the constraints set by the user.
2. **Scoring/Filtering:** Generated molecules are filtered by various modules. There are *2D scoring* modules (e.g. drug-likeness, synthetic accessibility, toxicity predictors) and *3D modules* (e.g. docking score to the target). Custom user-defined criteria (pharmacophore match, similarity to known ligands, etc.) can be integrated (pubs.acs.org).
3. **Selection and Learning:** The structures are ranked by a multi-objective score. High-ranking molecules are fed back to retrain/reinforce the generative models, biasing future output toward desirable chemotypes.

An intuitive analogy is that Chemistry42 runs *multi-agent reinforcement learning*: the generative models propose leads, then they "learn" from the scoring outcomes. The platform's web interface allows users to configure ligand-based or structure-based design (upload known ligands or a protein structure respectively) (pubs.acs.org). It supports advanced workflows like hit-expansion or fragment linking via "anchor points" (fixing part of a molecule while varying the rest) (pubs.acs.org). After iterations, the final output list—sometimes numbering in the thousands—is presented via interactive dashboards where chemists can scrutinize each candidate's predicted properties.

Crucially, Chemistry42 is designed for practical drug discovery. The 2023 ACS paper notes that generated molecules are automatically ranked on synthetic accessibility, novelty, and diversity, and that the pipeline provides medicinal chemistry filters to remove undesirable substructures (pubs.acs.org). Over the years, various proprietary modules have been added. For example, the MDFlow module (introduced circa 2025) runs molecular dynamics simulations to refine binding stability (insilico.com). Insilico also developed Alchemistry, an engine for relative binding free-energy calculations (Insilico Medicine), and introduced **Nach01**, a foundation model combining natural language processing with chemical intelligence (AWS Marketplace) (insilico.com). Nach01 is a "natural & chemical language" model that can process SMILES strings and text interchangeably, intended to reason about chemistry-intensive prompts (insilico.com).

Generative Models and Ensemble Approach

A standout feature of Chemistry42 is its **ensemble of generative models**. As of 2023 this included *over 40 models* (pubs.acs.org). Each model is architecturally and functionally distinct in order to diversify the search:

- **Autoencoders:** Capture smooth latent spaces of molecules; good for interpolation among known chemotypes.
- **Generative Adversarial Networks (GANs):** Learn to fool critics by proposing realistic molecules; useful for producing novel structures.
- **Flow-based models:** Provide exact likelihoods for sampling valid molecules.
- **Evolutionary algorithms:** Use genetic operators (mutation/crossover) to evolve molecules.
- **Chemical Language Models:** (e.g. RNN or transformer on SMILES) that generate sequences.
- **Graph Neural Network Generators:** Construct graphs atom-by-atom or bond-by-bond.
- **3D structure generators:** Propose molecules in 3D space respecting a protein pocket.

The diverse ensemble is managed asynchronously – models contribute candidates whenever ready, and all outputs are pooled for scoring (pubs.acs.org). Crucially, Insilico does not rely on a single "black-box" model. They provide **analytics** on each model's performance, so users can understand which algorithms explore which regions of chemical space (pubs.acs.org). This diversity-based strategy is intended to maximize the chance of finding high-quality leads quickly.

Capabilities and Applications

Chemistry42 is marketed as capable of *rapidly generating thousands of novel compounds* for a given target in a matter of hours. Indeed, a Pharma.AI update (Oct 2025) claimed that Chemistry42 could produce over **2,400 molecules within dozens of hours** for a design project (insilico.com). The platform also contains specialized predictive models (e.g. ADMET and kinase selectivity predictors) that can be used to filter designs before synthesis.

One notable internal example: Insilico described a use case for GLP1R (glucagon-like peptide-1 receptor) targeting peptides (insilico.com). Although that was a **biologics** (peptide) generation example (via the Generative Biologics module), it highlights the platform's speed—over 5,000 peptide structures were generated in 72 hours, 20 were selected by affinity scoring, and 14 showed functional activity (insilico.com). For small molecules, Insilico similarly reports that Chemistry42 has been used on proprietary programs spanning fibrosis, oncology, immunology, etc (insilico.com).

Example: DDR1 Kinase Inhibitor

A landmark demonstration of Insilico's chemistry engine was published in 2019 (the GENTRL / DDR1 study). Although that predates Chemistry42's formal launch, it used the same core concept of deep generative design. In the 2019 Nature Biotechnology study, four compounds were active in biochemical assays and two were validated in cell-based assays (original study abstract). Within 46 days from target selection, they had lead candidates in hand, much faster than traditional methods. This project was a "demo race" with WuXi AppTec and was one of the first real examples of AI-rapid drug design in collaboration with industry (communities.springernature.com).

For the TNIK program, the original discovery paper documents Chemistry42's role in lead generation: the platform first applied 30 generative models in parallel to generate compound structures (Nature Biotechnology).

Recent Enhancements (2024–2026)

- **MDFlow and Physics Integration:** To improve accuracy, Insilico layered physics-based methods atop AI. The MDFlow module enables molecular dynamics simulations for each candidate, refining estimates of binding affinities and filtering out unstable scaffolds (insilico.com). This addresses a known limitation of pure ML: geometry and dynamics are treated explicitly.
- **Nach01 Multimodal Model:** By July 2025, **Nach01** was available on AWS Marketplace (Insilico Medicine). It is a "multimodal natural & chemical languages foundation model" trained on massive data (insilico.com). Nach01 can process textual and chemical input, enabling queries like "design a molecule with properties X and Y" in plain language. Insilico made Nach01 available on AWS Marketplace and Microsoft's Discovery platform (insilico.com), signaling intent to share it broadly.
- **Wider Integration:** Insilico now bundles Chemistry42 as multiple applications. According to their 2025 webinar notes (insilico.com), Chemistry42 comprises *7 distinct sub-applications*, covering generation, free energy binding prediction (Alchemistry), ADMET prediction, kinase selectivity modules, retrosynthesis planning, etc. This modular suite means users can apply Chemistry42 to diverse tasks beyond initial hit generation.
- **Patent Analysis (PACE):** A feature called PACE (Patent Analysis for Chemical Entities) was mentioned as coming (in [1] key highlights) to flag intellectual property issues during design (insilico.com). This would help avoid designing molecules too close to existing patents.

Performance and Benchmarks

Quantitative benchmarks of Chemistry42's output performance are less publicly documented than PandaOmics, but two points stand out:

- **Empirical Projects:** The DDR1 example provided proof-of-concept; the more recent TNIK case (rentosertib) suggests the platform can produce drug-like inhibitors capable of entering the clinic. These successes, while limited in number, demonstrate feasibility.
- **Partner Feedback:** In interviews and press, Insilico cites high "hit rates" in virtual screening compared to conventional methods. For example, *Bio-IT World* (2022) noted that PandaOmics and Chemistry42 yielded a high percentage of actives in screening efforts, though exact numbers were not given (www.clinicalresearchnewsonline.com). The GLP1R peptide example (insilico.com) likewise showed a high efficacy (70% of scored peptides were active).
- **Model-specific comparison:** Table 1 of the workshop paper reports that LFM2-2.6B-MMAI improved over the base LFM2-2.6B across the evaluated molecular-optimization task settings (workshop paper). These results concern the evaluated LFM checkpoint; they do not establish a Chemistry42 benchmark advantage.

The platform's design pipeline (multi-model, multi-criteria) implicitly addresses cost and quality: by filtering thousands of candidates computationally, Chemistry42 aims to present only the most promising structures for synthesis. The claim is that it can accelerate "lead optimization" rounds, potentially cutting time and cost. Insilico's analyst report for 2025 quotes an **18.3% growth in subscription user base**, indicating commercial adoption (insilico.com).

Discussion

Chemistry42 represents Insilico's response to the challenge of navigating "a needle in a haystack" of chemical space. By 2026, this platform is highly feature-rich, combining AI creativity with chemistry domain knowledge. This is in line with industry trends; comparisons involving Exscientia should reflect its 2024 combination with Recursion (Recursion's 2024 review). One distinctive aspect is Insilico's emphasis on ensemble diversity and physics integration, whereas some others rely more on single-model approaches.

However, limitations remain. Generative models can produce molecules that are mathematically valid but synthetically impractical; Chemistry42 attempts to mitigate this via synthetic accessibility scoring and retrosynthesis checks. Still, the ultimate test is experimental synthesis and validation, which is time-consuming and expensive. To reduce risk, Insilico encourages an iterative human-AI loop: chemists pick among suggestions, test in vitro, and feed results back.

Ethical/regulatory aspects also surface: when AI suggests very novel scaffolds, patent landscapes and safety profiles must be carefully examined. Insilico's forthcoming PACE tool (insilico.com) aims to flag patent issues early. As Insilico CEO notes, making sure molecules are both innovative and safe is "essential to the medicinal tractability" of AI outputs (insilico.com).

Summary of Chemistry42

Key points for Chemistry42 (as of 2026):

- **Multi-agent Generative Pipeline:** 40+ distinct ML models work cooperatively to explore chemical space (pubs.acs.org).
- **Integration of Domain Filters:** 2D/3D scoring modules ensure candidate molecules meet drug-like criteria (pubs.acs.org).
- **Rapid Candidate Generation:** Thousands of compounds can be designed and ranked in days (insilico.com).
- **Computational-Experimental Loop:** Chemistry42 outputs are intended for subsequent synthesis and testing, with feedback loops to the AI models.
- **Recent Innovations:** Physics simulations (MDFlow), multi-modal LMs (Nach01), and specialized predictors (kinase selectivity) have been added (insilico.com).
- **Use Cases:** Claims of real-world impact include DDR1 kinase inhibitor design (communities.springernature.com) and (in partnership) the TNIK/IPF compound (www.eurekalert.org).

By 2026, Chemistry42 stands as a representative example of generative chemistry platforms. Its published documentation and success stories suggest it is among the most advanced in the field. Future work will likely focus on continuing integration of AI with wet-lab automation and more rigorous benchmarking of outputs against alternatives.

The MMAI Gym: Specializing Foundation Models for Science

Concept and Motivation

As large language models (LLMs) proliferated, Insilico recognized a gap: general-purpose LLMs (GPT-4, LLaMA, etc.) are “jacks of all trades and masters of none” in scientific domains (insilico.com). They struggle with precise molecular reasoning, often lacking the chemical intuition required for drug design tasks (insilico.com). In response, Insilico’s **MMAI Gym (Multi-Modal AI Gym)** aims to **“teach AI to think like a scientist”** (insilico.com). Rather than making models bigger, the idea is to fine-tune or “train” them on carefully curated scientific data and curricula, producing “lightweight” but highly specialized **Liquid Foundation Models (LFMs)** for drug discovery.

In collaboration with Liquid AI (a company founded by Ramin Hasani), Insilico developed and released **LFM2-2.6B-MMAI (v0.2.1)** in early 2026 (insilico.com). The 2.6B-parameter LFM2 backbone uses a hybrid of grouped-query attention and short convolutional layers (Liquid AI). The MMAI Gym is not just a dataset but a training methodology: it presents models with a **multi-stage curriculum** in medicinal chemistry, biology, and related tasks (insilico.com). Tasks include property prediction (ADMET), multi-objective molecule optimization, retrosynthesis planning, and target-aware scoring, among others (insilico.com). The goal is to imbue the model with **chemical and biological intelligence (CSI/BSI)** relevant to drug R&D, effectively creating an “AI chemist”.

Importantly, Insilico positions MMAI Gym as open to client models: pharma companies can bring their own LLM to the Gym, train it on these datasets, and gain the “drug discovery superintelligence” (insilico.com). Alternatively, Insilico itself and Liquid AI are distributing pre-trained LFMs that any lab can deploy on-premises, alleviating concerns about uploading proprietary data to public clouds (insilico.com).

Architecture: Liquid Foundation Models

LFM2-2.6B uses a hybrid architecture that alternates grouped-query attention blocks with short convolutional layers (Liquid AI).

Performance on Drug Discovery Benchmarks

The MMAI Gym’s first public yield, LFM2-2.6B-MMAI, was evaluated on a suite of drug discovery tasks. Key results reported by Insilico and Liquid AI include (insilico.com):

- **ADMET Property Prediction (Therapeutics Data Commons, TDC):** Outperformed TxGemma-27B (a 27-billion parameter model) on 13 of 22 tasks and achieved state-of-the-art on 3 tasks (insilico.com).
- **Multi-Parameter Molecular Optimization (MuMO-Instruct benchmark):** Achieved success rates up to **98.8%** (success means optimizing properties while retaining core scaffold), exceeding performance of established proprietary models (insilico.com).

- **Affinity Prediction:** On an internal Insilico benchmark (2.5M measured bioassays across 689 protein targets), LFM2-2.6B-MMAI produced higher correlation metrics than leading models (GPT-5.1, Anthropic Claude, Grok 4.1) (insilico.com).
- **Chemical Reasoning (Functional Group Benchmark and Retrosynthesis):** Demonstrated strong reasoning about functional group effects; improved single-step retrosynthesis suggestions from near-zero to top-tier quality (insilico.com).
- **Synthesis Planning (ChemCensor metric):** The model's retrosynthesis module now "matches" top specialist tools (insilico.com).

These results are impressive given the modest size (2.6B). Researchers wrote: *"We have demonstrated that specialist-level performance in drug discovery does not require frontier-scale model size… a smaller, efficient model can outthink a giant."* (insilico.com). Moreover, because the model can run on private hardware, it enables secure use of proprietary compound libraries that companies would never send to third-party clouds (insilico.com).

It should be noted that LFM2-2.6B-MMAI is a distilled exemplar; performance for other tasks or future iterations depends on continued MMAI training. However, these benchmarks suggest that a well-trained model can achieve **comparable or superior performance to very large generalist LLMs** in pharma tasks. As Insilico's 2025 annual report states, the MMAI Gym was a "foundation model training framework" that by March 2026 "achieved state-of-the-art (SOTA) performance using a lightweight model on private infrastructure" (insilico.com).

The MMAI Gym Platform and Process

While LFM2-2.6B is one outcome, the **MMAI Gym** concept is broader. Insilico provides (internal and presumably commercial) tools to fine-tune any LLM:

- **Pharmaceutical Benchmarks:** The Gym uses *over 1000 pharma-specific benchmarks and datasets* (internal assays, published data, simulation targets, etc.) (insilico.com).
- **Scientific Curriculum:** Training is staged. For example, early phases train on basic chemical understanding (e.g. SMILES generation, property prediction), later phases on complex tasks like multi-step synthesis or mechanism inference (insilico.com).
- **Agentic Workflows:** The Gym encourages creation of AI "agents" that combine modules (either Insilico's own curated PSI models or the trained LLMs themselves) into end-to-end workflows (insilico.com).
- **Experimental Validation:** Insilico's labs can test outputs. In one workflow diagram (Fig. 1 from [3]), there is a "robotic lab for target validation and compound screening" linked back into the AI core (pmc.ncbi.nlm.nih.gov). The Gym can interface with such labs to iteratively refine models based on wet-lab results.

MMAI Gym's end goal is to produce a **"Chemical Superintelligence" (CSI)** and **"Biological Superintelligence" (BSI)** that can be used by clients. Insilico advertises that clients can bring their own LLM and "improve your model's drug discovery fitness" through a 3-month training session (insilico.com). In

parallel, Insilico is open-sourcing and releasing its own models (like PreciousGPT for life science and Nach01 for chemistry) which can seed the process.

Position in the AI Landscape

The MMAI Gym approach reflects a shift from seeking ever-larger generic models to building smaller **specialist models** for science. This has several advantages:

- **Efficiency and Cost:** Insilico reports that LFM2-2.6B-MMAI outperformed TxGemma-27B, a model roughly ten times its size, on selected safety and pharmacokinetics tasks in the Therapeutics Data Commons benchmark (Insilico Medicine).
- **Domain Adaptation:** By focusing training on pharmaceutical data, the models learn relevant patterns (e.g. medicinal chemistry heuristics, toxicophores) that general LLMs might miss.
- **Security:** Pharma companies often distrust sending proprietary molecules to public LLMs (like GPT) due to IP concerns. On-prem LFMs solve this.
- **Regulatory Alignment:** With regulatory interest in explainable and user-aware AI, a specialized model could be easier to validate than an opaque 100B+ LLM.

At the same time, it raises questions about generality. A specialized LFM might excel in drug-like chemistry but be useless outside that niche. That is acceptable for corporate R&D tools but contrasts with the vision of a single AGI solving all problems. Insilico's stance is pragmatic: *"the era of Brute Force AI is ending; the era of the Scientific Specialist has begun"* (insilico.com).

Summary of MMAI Gym

- **Purpose:** Enable LLMs to master pharmaceutical chemistry and biology through targeted training.
- **Key Achievement:** LFM2-2.6B-MMAI model, trained via MMAI Gym, achieved SOTA on multiple drug discovery benchmarks (ADMET, optimization, affinity, retrosynthesis) (insilico.com).
- **Approach:** Multi-stage curriculum with thousands of pharma-specific tasks; partnership with Liquid AI for efficient model architecture.
- **Implications:** Demonstrates that smaller, domain-tuned models can rival larger generic LLMs in life science tasks, making AI more accessible and secure for drug R&D.
- **Future Directions:** Further extension to multimodal (connecting chemical and biological knowledge), integration into automated pipelines (agents that propose experiments), and open collaborations (sharing models and benchmarks).

Data Analysis and Evidence

Insilico's Business Growth and Pipeline

Insilico Medicine's reported financial and pipeline data provide a quantitative backdrop for the technology discussion. According to the company's 2025 results (announced March 29, 2026):

- **Revenue:** Total revenue decreased 34.5%, from US$85.8 million in 2024 to US$56.2 million in 2025 (2025 annual results announcement).
- **Revenue Growth:** Software subscription revenue grew **+23.8% year-over-year**, and the number of platform customers (SubBase) grew by 18.3% (insilico.com). This indicates strong market adoption of Pharma.AI.
- **Clients:** The platform reportedly serves **13 of the top 20 global pharma companies** (insilico.com). In practical terms, this means Insilico has penetrated major R&D players.
- **Funding:** Insilico's December 30, 2025 Hong Kong IPO raised HKD 2.277 billion in gross offering proceeds (IPO announcement).
- **Pipeline (preclinical):** Using its AI, Insilico **nominated 6 new preclinical candidates in 2025** (4 disclosed to public, 2 undisclosed) (insilico.com). This brought their total to **28 preclinical candidates** (insilico.com).
- **Pipeline (clinical):** The company advanced **8 programs in clinical development** during 2025 and has **10 programs currently in trials** (insilico.com). Key examples include the TNIK inhibitor (Phase IIa) and other fibrosis, cancer, and metabolic programs.

Table 2 (below) summarizes these metrics drawn from Insilico's reports, giving a high-level view of business/clinical accomplishments in 2025.

Metric	2025 Value	Source
Total Revenue	US$56.24 million (insilico.com)	[2] Annual Results
Software revenue growth (YoY)	+23.8% (insilico.com)	[2] Annual Results
Subscription customer base growth (YoY)	+18.3% (insilico.com)	[2] Annual Results
Top-20 Pharma clients served	13 companies (insilico.com)	[2] Annual Results
New preclinical candidates nominated	6 in 2025 (28 total) (insilico.com)	[2] Annual Results
Programs in clinical trials	10 (including Phase IIa rentosertib) (insilico.com)	[2]
Cash and bank balances (December 31, 2025)	US$393.3 million (Insilico Medicine)	Annual results release
Pharma.AI partnership milestones (CY2025)	Not specified in detail (milestone payments included)	[2]

*Table 2. Insilico Medicine 2025 key business and pipeline metrics (from the company's annual report (insilico.com) and project news (www.eurekalert.org)). These indicate broad adoption and scaling of the platform. *

These numbers require context. Total revenue fell 34.5%, from US$85.8 million in 2024 to US$56.2 million in 2025 (2025 annual results announcement). The reported 23.8% year-over-year increase applies to software revenue, rather than total revenue (Insilico Medicine).

Benchmark Comparisons and Performance Claims

The MMAI Gym results are task-specific evaluations of LFM2-2.6B-MMAI. The public ICLR 2026 workshop paper reports molecular optimization, functional-group reasoning, retrosynthesis, ADMET and activity-prediction evaluations (workshop paper).

Task / Benchmark	LFM2-2.6B-MMAI Performance	Baseline
ADMET (Therapeutics Data Commons)	Outperforms TxGemma-27B on 13/22 tasks; SOTA on 3 tasks (insilico.com)	TxGemma-27B (27B param); multiple specialist ADMET models
Molecular Optimization (MuMO-Instruct)	Evaluated on five property combinations; see the task-level results in Table 1 of the workshop paper	Published benchmark baselines; results depend on task and metric
Affinity Prediction	Higher correlation than GPT-5.1, Claude-4.5p, Grok-4.1 on Insilico test (2.5M assays) (insilico.com)	GPT-5.1, Anthropic Claude-4.5p, Grok-4.1
Functional Group Reasoning (FGBench)	Strong performance (automatically reasons about substructures) (insilico.com)	Competitor benchmarks or unspecified
Retrosynthesis (ChemCensor)	Jumped from near-zero to top-tier (elite) single-step rxn suggestion (insilico.com)	Top specialist retrosynthesis tools

Table 3. Reported task-specific benchmark results. The molecular-optimization study reports separate property combinations; these findings should be assessed by task and metric.

The vendor-authored workshop paper provides evaluation methods and baseline comparisons. It reports MuMO-Instruct in the non-reasoning setting to allow comparison with prior baselines (workshop paper). Publication of these methods should be distinguished from independent replication; no independent replication is established here.

What about PandaOmics or Chemistry42 benchmarks? There are fewer public “number vs number” claims, but we note:

- The original DDR1 study reports four compounds active in biochemical assays and two validated in cell-based assays (PubMed).
- The GLP1R peptides example gave 14 actives out of 20 tested, 3 with single-digit nM activity (insilico.com) (70% hit rate).
- Consumer feedback (like top-20 pharma using it) is a form of endorsement but not strict numerical performance. Still, being licensed by many companies suggests the modules pass corporate evaluation.

Finally, industry analyses offer some comparative data points on AI pipelines at large (which indirectly validate Insilico’s trajectory):

- Analysts (Axis Intelligence, Dec 2025) estimate *200+ clinical AI drug programs*, with the first approvals expected around 2026-27 (axis-intelligence.com). In this larger context, Insilico’s 10 ongoing trials (by early 2026) represent a significant share of the field.

- A 2024 analysis of clinical pipelines at AI-native biotech companies reported an approximately 40% Phase II success rate on a limited sample, comparable with historical industry averages (original analysis). This describes an observed cohort result, rather than a forecast of future success for AI-discovered drugs.

Synthesis: Evidence-Based Assessment

Taken together, the data suggests:

- **Platform Efficacy:** Insilico's tools have demonstrable hits in silico and in vivo, most notably the IPF trial. The MMAI Gym results provide strong evidence of the platform's accelerating capabilities on core computational tasks. These data points support their claims of "revolutionizing drug discovery with AI" (insilico.com) (www.eurekalert.org).
- **Competitive Position:** Insilico has built one of the fastest and broadest AI workflows among peers. By integrating target ID (PandaOmics) and molecule design (Chem42) with novel LLM training (MMAI), they cover more steps of the pipeline than many competitors, who might focus on one or two stages.
- **Business Validation:** Growth in software revenue and pharma partnerships indicate market interest (Insilico Medicine). Serving top companies and raising capital indicate confidence in this approach. The IPO success (largest HK biotech raise, presence of partners like Eli Lilly investing) implies industry backing (insilico.com).
- **Areas Needing Confirmation:** Published rentosertib efficacy findings discussed here come from the 12-week Phase IIa study (Nature Medicine). On July 7, 2026, Insilico announced initiation of a Phase III study designed to evaluate once-daily treatment over 52 weeks. This announcement describes trial initiation and design; it does not provide Phase III outcomes (Insilico Medicine).
- **Integration with Traditional Research:** Evidence suggests Insilico recognizes that AI is a tool, not a replacement for all human expertise. The new "mechanism clarity" scores (insilico.com) and the inclusion of experienced medicinal chemists in development teams indicate an understanding that AI proposals require chemical sensibility.

Case Studies and Real-World Examples

IPF and Rentosertib (TNIK Inhibitor)

The rentosertib program links target discovery with molecular design: the original discovery paper documents PandaOmics-generated fibrosis target hypotheses and Chemistry42-generated TNIK inhibitor leads (Nature Biotechnology).

The subsequent Phase IIa study provides a preliminary lung-function signal. Mean FVC change at 12 weeks was +98.4 mL with 60 mg once daily versus −20.3 mL with placebo. Exploratory biomarker changes supported the hypothesis that TNIK inhibition modulates IPF pathways. The small, homogeneous groups and short follow-up limit interpretation; these findings require confirmation in larger and longer trials (Nature Medicine).

GLP-1R Peptide Design (Generative Biologics)

Although slightly outside the main focus, Insilico's report included an eye-catching example of generative biologics (peptides for GLP-1 receptor, relevant in metabolic disease) (insilico.com). In 72 hours, their platform generated 5,000 peptides, selected 20 by predicted metrics, and **14 of those showed biological activity** (with 3 having sub-10 nM potency) (insilico.com). In this company-reported case study, the 20 tested peptides were selected for synthesis and wet-lab evaluation using predicted affinity scores and computed binding energies (Insilico Medicine).

This anecdote, while for peptides, illustrates the power of an integrated generative workflow. It parallels what Chemistry42 does for small molecules. It also aligns with Insilico's pipeline focus on metabolic diseases (e.g. in [2], ISM0676 GIPR antagonist for obesity showed 31.3% weight loss in preclinical models (insilico.com)). The GLP-1R case suggests Insilico can accelerate biologics design similar to small molecules.

Pipeline Examples (Other Indications)

Insilico's publicly disclosed preclinical programs provide additional evidence:

- **CBLB inhibitor (ISM3830) for cancer immunotherapy**: Described as an "oral, highly selective CBLB inhibitor… improving drug metabolism and safety (potential best-in-class)" (insilico.com). CBLB is an unconventional target (an E3 ligase in T cells), so its selection likely came from AI analysis of immune pathways.
- **GIPR antagonist (ISM0676) for metabolic disease**: The pipeline highlight claims "31.3% body weight loss in preclinical models" (insilico.com), which is notably high for a single agent. (GIPR is a known target for obesity; multiple companies target it.) If AI was used to optimize this molecule, it suggests success in metabolic indications.
- **NLRP3 inhibitor (ISM5059)** and **pan-KRAS inhibitor (ISM6166)** are also mentioned (insilico.com). These reflect ambitious targets – NLRP3 for inflammation, KRAS for oncology – both historically challenging. A broad portfolio suggests the AI workflow is being applied to diverse therapeutic areas.

These examples show Insilico pursuing both established targets (NLRP3, GIPR) and novel ones (CBLB, TNIK). The speed and successes in preclinical results (e.g. safety margins, potency) are evidence that AI-guided discovery can produce credible drug candidates. Nevertheless, until human data arrives for each, they remain promising leads.

External Perspectives

Independent coverage of Insilico's work is still limited. A *Clinical Research News* article (Dec 2022) described Insilico's Pharma.AI generically as "pushing the envelope" and noted the DDR1 and JAK inhibitor successes . It quoted an Insilico exec saying generative models had advanced to designing molecules validated *in vivo*. This aligns with our technical reading.

Recursion provides another point of comparison: it describes itself as a clinical-stage TechBio company focused on industrializing drug discovery (Recursion).

Lessons and Implications

The case studies illustrate key points:

- **Integration is Powerful:** Identifying TNIK via generative AI required linking gene-expression insights to molecular design. PandaOmics and Chemistry42 together facilitated this. Standalone target ID or chemistry alone would not yield a molecule.
- **Validation is Crucial:** Each AI output is subject to experimental testing. Insilico's reported outcomes show that many AI picks do work. This builds confidence in the platform.
- **Time Compression:** Some results (e.g. DDR1 in 46 days) hint that these platforms can compress timelines. If Insilico or partners can routinely cut years off hit discovery, that transforms R&D economics.
- **Pipeline Diversification:** The broad range of targets (from glycemic to immunology) suggests the AI is not limited to one disease class. This generality is important for a platform model.

However, these cases also underscore the need for **expert interpretation**. A documented example is Insilico's research collaboration with Pfizer to explore potential therapeutic targets (Pfizer). AI did not work in isolation; human biologists and chemists guided and vetted the processes. This hybrid approach is likely to continue as the norm.

Discussion: Broader Perspectives and Future Directions

AI's Role and Limitations

Our analysis shows that Insilico Medicine's Pharma.AI platform embodies the cutting edge of AI-enabled drug discovery in 2026. It demonstrates that generative AI and multi-omics can yield real drug candidates more quickly than traditional methods. Insilico's success stories and benchmarks back the narrative that AI is a significant force in pharma innovation (www.eurekalert.org) (insilico.com).

At the same time, multiple sources caution against overpromising. **The Atlantic's December 2025 feature** (concerned about AI fever in cancer) points out that *"current scientific consensus remains more cautious"* despite hype (www.theatlantic.com). Insilico has advanced rentosertib to a Phase III clinical trial. In fact, the Nature Medicine IPF trial is claimed as the *first* clinical proof-of-concept of an AI-discovered drug (www.eurekalert.org), implying very few AI drugs have reached that stage. The Bayesian perspective is that pipelines often winnow – for every rentosertib, others may fail. Therefore, industry watchers stress validation frameworks and humility. Insilico's addition of "mechanism clarity" to PandaOmics (insilico.com) and the "AI Risk-Mitigation" commentary in drug sector publishing (digitaledition.chaindrugreview.com) are responses to these concerns.

Another issue is **data quality and availability**. AI models are only as good as their training data. Omics data can be noisy and heterogeneous. Chemistry generative models need large datasets of molecule-property pairs. Many disease areas (rare diseases, novel targets) lack large datasets. Insilico partially solves this by

integrating literature (PandaOmics uses text-mining) and by crowdsourcing pharma libraries. Still, model bias and blind spots must be managed.

Regulatory and Ethical Outlook

As of October 30, 2026, the FDA page identifies its January 2025 guidance on AI supporting regulatory decision-making for drug and biological products as **draft**, not for implementation, and containing non-binding recommendations (FDA draft guidance). FDA's CDER and CBER have also collaborated with EMA on 10 guiding principles that developers can consider when using AI in drug and biological product development (FDA/EMA guiding principles). These recommendations should not be presented as binding requirements or as evidence of a confirmed final-guidance deadline.

Ethically, the use of patient data, genetic databases, and privacy concerns must be considered. PandaOmics may handle real patient omics (de-identified), so data governance is crucial. Also, generative suggestions for molecules must avoid biased safety outcomes; typical metrics (like hERG inhibition) should be included upstream (and Insilico mentions "toxic liver enzymes" scoring). Users will likely need to thoroughly vet toxicity profiles before advancing AI designs to humans.

Future Evolution of Pharma.AI

Looking ahead, Insilico has several avenues for growth:

- **Expansion of Modules:** The platform already announced additional tools (PDB-integrated Generative Biologics improvements, improved patent mining, etc.) (insilico.com). We expect further modules, e.g. specialized AI for cell therapy design, or AI-guided formulation chemistry.
- **Model Iteration:** The MMAI Gym is likely to produce new generations of foundation models. As more data accumulates (e.g. multi-omics from clinical trials, new chemistries), models like Nach01 and LFM2 will be continually updated. Possibly, Insilico will participate in open ML challenges (like many academic benchmarks) to validate and improve their models.
- **Open Science and Collaboration:** Insilico has shown willingness to open-source some tools (Science42 DORA, Insilico-life's PreciousGPT). For long-term credibility, broader community benchmarking of PandaOmics/Chemistry42 could emerge. Academic publications of these modules is a good step (we saw PandaOmics and Chemistry42 in ACS JCIM) – more peer review would strengthen confidence.
- **Domain Generalization:** The MMAI approach could extend beyond chemistry/biology into clinical trial design (explains inClinico). Indeed, patient stratification threads (e.g. combining genomic + clinical data) are ripe for AI, and Insilico's portfolio includes inClinico for Phase II→III success prediction (pmc.ncbi.nlm.nih.gov).

Academically, this trend signals a blending of **bioinformatics, cheminformatics, and AI research**. The gene-expression analysis in PandaOmics and the generative modeling in Chemistry42 already combine methods from disparate fields. MMAI Gym is at the frontier of NLP applied to science. We anticipate new interdisciplinary publications (like the ICLR 2026 workshop paper by Insilico/Liquid AI cited above) that will reach not only pharma audiences but also AI conferences.

Competitive and Collaborative Landscape

Insilico is not alone. Other companies (Recursion, which combined with Exscientia in 2024, Schrodinger, etc.) and academic labs are also building AI drug discovery platforms (Recursion's 2024 review). What differentiates Insilico:

- **Integration breadth:** It tackles target discovery and chemistry within Pharma.AI; the combined Recursion business is a relevant comparator.
- **Publications:** Insilico has several peer-reviewed papers, adding transparency; some startups have been more secretive.
- **Infrastructure:** The Hong Kong listing and partnerships (Eli Lilly as investor) give Insilico substantial capital and industry ties.

Collaboration will be key. Pharma companies may license multiple AI tools to diversify risk. There is room for synergy, e.g. using Schrodinger's physics engines with AI models, or coupling Insilico's target predictions with phenotypic screens from other platforms. Also, open data efforts (like the Therapeutics Data Commons benchmarks) allow cross-platform comparisons and improvement. The field may see consortia forming around standard AI benchmarks (as Insilico's TargetBench hints (insilico.com)).

Risks and Reservations

While optimistic, a balanced analysis acknowledges unresolved issues:

- **Overfitting to Benchmarks:** There's a risk of over-optimizing to known benchmarks rather than true innovation. Insilico reports results on MuMO and TDC, and their practical relevance depends on translation to novel, clinically meaningful compounds beyond the test sets.
- **Reproducibility:** Many AI studies suffer from reproducibility gaps. Did Insilico train random seeds and test stability? Are their datasets (or code) available for others to verify? The open nature of their publications suggests some transparency, but proprietary data cannot be fully shared.
- **Workforce Impact:** Insilico's CEO has even predicted AI could replace some R&D roles within years (per an interview) – a bold claim. If true, Insilico and others must address workforce retraining and ethical deployment of automating scientific tasks. So far, "AI assists scientists" not replaces them, but that narrative could evolve.

Conclusion

By early 2026, Insilico Medicine's **Pharma.AI** has emerged as one of the most comprehensive AI-driven drug discovery platforms. **PandaOmics** leverages AI to sift through complex omics and literature data to propose new therapeutic targets. **Chemistry42** deploys a massive ensemble of generative models (augmented by physics-based tools) to rapidly design novel drug candidates. The newly introduced **MMAI Gym** marks Insilico's foray into foundation-model fine-tuning, producing compact LLMs that specialize in chemistry and biology and demonstrably outperform much larger general models on key tasks.

Our in-depth analysis shows these technologies have moved beyond theoretical promise. Insilico's platform has contributed to tangible pipeline advancements (28 preclinicals, 10 in clinic), including the groundbreaking case of an AI-derived TNIK inhibitor for pulmonary fibrosis (www.eurekalert.org). Benchmark results and case anecdotes indicate this "Pharma.AI" approach can significantly **accelerate early-stage R&D** and even handle sophisticated tasks like multi-parameter optimization, evaluated across five MuMO-Instruct property combinations (workshop paper). Insilico's software revenue growth, partnerships and IPO indicate commercial interest (insilico.com).

Looking forward, ongoing challenges remain. Broad adoption of AI in pharma will depend on regulatory acceptance (FDA is moving quickly on AI frameworks (axis-intelligence.com)), demonstrable success rates in clinic, and integration of AI outputs with human expertise. Insilico's continued work on explainability (e.g. new scoring metrics (insilico.com)) and validation (via robotics and collaborations) bodes well.

In sum, Insilico Medicine's AI platform, exemplified by PandaOmics, Chemistry42 and the MMAI Gym, represents a bold advance in computational drug discovery. **If the early successes (e.g. rentosertib)** are harbingers, we may indeed be entering an era where AI dramatically compresses the path from biology to medicine. However, observers will rightly demand more data from late-stage trials and peer-reviewed studies. For now, Insilico's story is a case study in how **artificial intelligence is reshaping pharmaceutical R&D**, with both impressive achievements and the need for careful, evidence-based progress.

References

1. Insilico Medicine. *Insilico Medicine unveils winter edition of Pharma.AI, accelerating the path to pharmaceutical superintelligence*. (2025). (insilico.com)
2. Insilico Medicine. *Pharma.AI Annual Results Q&A*. (2026). (insilico.com)
3. Kamya P., Pun F.W., Tretina K., et al. *PandaOmics: an AI-driven platform for therapeutic target and biomarker discovery*. J. Chem. Inf. Model. 64, 3961–3969 (2024) (pmc.ncbi.nlm.nih.gov).
4. Ivanenkov Y.A., Polykovskiy D., et al. *Chemistry42: an AI-driven platform for molecular design and optimization*. J. Chem. Inf. Model. 63(3), 695–701 (2023) (pubs.acs.org).
5. Insilico Medicine. *Science MMAI Gym (Pharma.ai Webinar, 14 Apr 2026)*. (insilico.com)
6. Insilico Medicine & Liquid AI. *Liquid AI and Insilico Medicine strategic partnership press release*. (April 2026) (insilico.com).
7. Insilico Medicine. *Teaching AI the Language of Molecules: How MMAI Gym and Liquid Intelligence are Solving the "Brute Force" Crisis in Drug Discovery*. (9 Apr 2026) (insilico.com).
8. Insilico Medicine. *Science MMAI Gym page*. (Accessed Apr 2026) (insilico.com).
9. Insilico Medicine/EurekAlert. *Nature Medicine publication: Rentosertib (ISM001-055) Phase IIa results in IPF*. (3 Jun 2025) (www.eurekalert.org).
10. BioSpace. *Insilico announces Nature Medicine publication of Phase IIa results for Rentosertib (ISM001-055) in IPF*. (3 Jun 2025) (www.biospace.com).
11. Insilico Medicine. *Science MMAI Gym*. (Product page) (insilico.com).
12. Dudley J.T., et al. *The First Generation of AI-Generated Therapeutics Advances*. (Featured interviews, Observer) .

13. Borfitz D. *Insilico Medicine's AI-Driven Platform Pushes The Envelope Of Drug Discovery*. Clinical Research News (15 Dec 2022) (www.clinicalresearchnewsonline.com).
14. Chain Drug Review. *AI-driven drug discovery requires careful navigation* (Aug 2024) (digitaledition.chaindrugreview.com).
15. Axis Intelligence. *AI Drug Discovery 2026-2030: Analysis of 173 programs, FDA framework, and $16.5B market transformation* (Dec 2025) (axis-intelligence.com).
16. AP News. *"Better drugs through AI? Insitro CEO on what machine learning can teach Big Pharma"* (2 Dec 2024) (apnews.com).
17. Food and Drug Administration. *FDA announces completion of first AI-assisted scientific review pilot and aggressive agency-wide AI rollout timeline* (8 May 2025) (www.fda.gov).
18. Weinswig D. *AI in drug discovery: opportunities and challenges (Coresight) – Chain Drug Review* (19 Aug 2024) (digitaledition.chaindrugreview.com).
19. Gen. Engineering News. *Alex Zhavoronkov profile: aims to take over drug development with AI.* (2023) (www.genengnews.com).
20. Nature *"Deep learning enables rapid identification of potent DDR1 kinase inhibitors"* (Aladinskiy et al., 2019) (communities.springernature.com).
21. Other references as noted above (links [1]– [18]).

Source links

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

1. https://insilico.com/news/2ky5zn0fs1-insilico-medicine-unveils-winter-edition
2. https://pmc.ncbi.nlm.nih.gov/articles/PMC11134400
3. https://insilico.com/news/ohz9ozx0t1-insilico-medicine-announces-2025-annual
4. https://www.nature.com/articles/s41591-025-03743-2
5. https://pubs.acs.org/doi/10.1021/acs.jcim.2c01191
6. https://insilico.com/news/bo54fkt0e1-liquid-ai-and-insilico-medicine-announce
7. https://apnews.com/article/004c0ce0442b72c37bfec6e032796808
8. https://digitaledition.chaindrugreview.com/articles/ai-driven-drug-discovery-requires-careful-navigation
9. https://communities.springernature.com/amp/posts/after-the-paper-from-paper-to-industrial-scale-platform-a-3-year-behind-the-paper-journey-from-gentrl-to-chemistry42
10. https://insilico.com/news/fhoex9ed01-pharmaai-summer-2025-updates-join-insili
11. https://www.fda.gov/news-events/press-announcements/fda-announces-completion-first-ai-assisted-scientific-review-pilot-and-aggressive-agency-wide-ai
12. https://insilico.com/mmai
13. https://insilico.com/tpost/003kiu5pv1-insilico-pharmaai-fall-launch-recap-unde
14. https://insilico.com/blog/mmai-liquid-ai
15. https://www.clinicalresearchnewsonline.com/news/2022/12/15/insilico-medicine-s-ai-driven-platform-pushes-the-envelope-of-drug-discovery
16. https://www.nature.com/articles/s41587-024-02143-0

17. https://aws.amazon.com/marketplace/pp/prodview-aq32kfu5ifjgw
18. https://pubmed.ncbi.nlm.nih.gov/31477924
19. https://openreview.net/pdf?id=4m2TcQpwkf
20. https://www.recursion.com/yearbook/2024
21. https://www.eurekalert.org/news-releases/1086096
22. https://www.liquid.ai/blog/introducing-lfm2-2-6b-redefining-efficiency-in-language-models
23. https://cdn.financialreports.eu/financialreports/media/filings/50870/2026/RNS/50870_rns_2026-03-29_d13beca1-9e93-4c8b-b8b1-43fd0b5ffaac.pdf
24. https://insilico.com/news/p010170up1-insilico-medicine-lists-on-hong-kong-sto
25. https://axis-intelligence.com/ai-drug-discovery-2026-complete-analysis
26. https://pubmed.ncbi.nlm.nih.gov/38692505
27. https://insilico.com/news/xmjsn4l091-insilico-initiates-phase-iii-clinical-tr
28. https://insilico.com/news/o1y20pfyz1-insilico-showcases-advanced-generative-b
29. https://ir.recursion.com/news-releases/news-release-details/recursion-and-exscientia-two-leaders-ai-drug-discovery-space
30. https://www.pfizer.com/news/press-release/press-release-detail/insilico-enters-research-collaboration-pfizer-inc-explore
31. https://www.theatlantic.com/technology/archive/2025/04/how-ai-will-actually-contribute-cancer-cure/682607
32. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
33. https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development
34. https://www.biospace.com/press-releases/insilico-medicine-announces-nature-medicine-publication-of-phase-iia-results-evaluating-rentosertib-the-novel-tnik-inhibitor-for-idiopathic-pulmonary-fibrosis-ipf-discovered-and-designed-with-a-pioneering-ai-approach
35. https://www.genengnews.com/topics/artificial-intelligence/alex-zhavoronkov-aims-to-take-over-the-drug-development-world-with-ai

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.