

News and Perspectives

From Robots to Organizations: The Role of the AI Scientist

Cliff Dominy, JMIR correspondent

Abstract

AI is increasingly used in science—and increasingly *doing* science. In this *News and Perspectives* article, JMIR Correspondent Cliff Dominy reports on the rise of “AI scientists” and the role human scientists will need to continue to play.

Key Takeaways:

- Sophisticated AI scientists can accelerate laboratory research by integrating experimental design, control, and execution with wet lab robotics.
- Current research is focused on improving “closed-loop” automation and advanced multiagent complexity to expand the functionality of the field.
- AI scientist systems, under human governance, will shift the discovery bottleneck from the lab bench to data validation by experts.

In 2023, Gabe Gomes, PhD, and colleagues at Carnegie Mellon University used an AI system named Coscientist to recreate a Nobel Prize–winning chemical reaction. The mechanism—a Suzuki-Miyaura coupling—the second-most used reaction in medicinal chemistry, took its discoverer Akira Suzuki 11 years to develop. Coscientist, without detailed instruction, compiled Suzuki’s protocol in under 4 minutes before executing the experiment using a cloud laboratory service at a remote location.

Unlike the humans they emulate, AI scientist technologies never tire, can work around the clock, and excel at repetitive tasks. Today, with the merger of modern wet lab robotics and AI automation, thousands of experiments can be run in parallel and automatically recorded without human intervention.

How can this technology assist scientists, and what is their role in the process?

The First Robot Scientists: Closing the Loop

The purpose of laboratory automation is to improve data throughput, reproducibility, and safety and to free highly skilled researchers from repetitive tasks. Developments toward autonomous robots began in 2009, when Ross King, PhD, at Aberystwyth University in Wales, developed the first robot scientist, [Adam](#). Adam could generate hypotheses about gene functions in yeast, test them in a laboratory, and report the outcomes. It was the first machine to discover new scientific knowledge—albeit within the bounds of the resources it was supplied with. Adam was much more than an automated instrument but not quite a fully autonomous, self-correcting scientist. Alex Tobias, PhD, a Caltech-trained engineer with automation experience at DuPont Corporation, notes that before Adam, “the robots were flashy and

attractive, but they couldn’t *do* science.” Adam was becoming autonomous—but could Adam think?

From Automation to Autonomy

Technology began focusing on integrating automated robotics with the reasoning capabilities of machine learning algorithms. The new platforms excelled at hands-free closed-loop automation of very defined tasks. Enhanced speed was followed by scale—evolving toward smaller hardware and greater impact—with King’s latest AI scientist, [Genesis](#), incorporating advances in modern [microfluidic engineering](#). Microfluidics downsizes laboratory experiments onto silicon chips, enabling thousands of chemical reactions to be run simultaneously.



The fanciest hardware is not going to advance science [much further] ... it’s really the software that has the opportunity to achieve the kinds of results that we’re hoping for.

Alex Tobias, PhD

Hardware aside, Tobias anticipates even greater gains from AI software, saying “The fanciest hardware is not going to advance science [much further]... it’s really the software that has the opportunity to achieve the kinds of results that we’re hoping for.”

Selected AI scientists

Platform	Year	Field	Contribution
Adam	2009	Functional genomics	First closed-loop robot scientist to generate and test hypotheses
Eve	2015	Drug discovery	Early machine learning identified drug candidates against malaria
Mobile robotic chemist	2020	Materials chemistry	Autonomous robot identified photocatalysts 6-fold more active than starting formulations
BacterAI	2023	Microbiology	Identified the growth requirements of uncharacterized bacteria
A-Lab	2023	Inorganic chemistry	Iterative cycling with active machine learning
Coscientist	2023	Synthetic chemistry	Large language model coupling for semiautonomous research, design, and optimization
Genesis	2024	Systems biology	Microfluidic scale automated with advanced analytics
ChemCrow	2024	Molecular design	Expert chemistry tool integration with GPT-4
AI Scientist	2024	Machine learning research	Automated idea generation
Virtual Lab	2025	Protein engineering	Automated nanobody design for improved binding efficacy to new viruses
Co-Scientist	2025	Biomedical	Hypothesis generation and ranking across multiple agents
Virtual Biotech	2026	Biomedical	Virtual biotechnology company estimating probability of clinical trial success
Anthropic	2026	Protein chemistry	Claude Mythos Preview and Opus 4.8 models

From Isolated Skills to Whole Organizations

Automated AI scientists are expanding beyond specific tasks to manage sophisticated biomedical programs. Drug development is a complex field combining molecular biology and medicinal chemistry with clinical studies in humans. AI agents excel at completing individual tasks within clinical trials, but in order to work together, this large biomedical orchestra will need a conductor.

James Zou, PhD, and colleagues at Stanford University set out to address this by developing Virtual Biotech, a large language model (LLM)-guided manager and analyst of clinical trials. Zou explains, “Most of the existing [AI models] are about having a single or small number of agents to automate a particular task... [With Virtual Biotech,] it’s really about... tens of thousands of AI agents... trying to solve a much more complex problem.”

Like typical biotechnology companies, Virtual Biotech has a Chief Scientific Officer (CSO)—in this case, an AI agent responsible for coordinating and managing the operations of specialized subagents. The CSO can break a question down into problems, dispatch the relevant specialist to retrieve the required data, and integrate its answers into the results. Zou describes the underlying Anthropic LLM as the “brains of the operation,” and the specialist AI agents his team is developing as “the arms and legs” of Virtual Biotech, which access the relevant datasets.

Virtual Biotech [deployed](#) 37,075 AI agents to analyze outcome data from 55,984 clinical trials and integrate the findings with datasets from the biomedical record. When it comes to clinical trials, Zou notes, “Often we have more failures than successes. And it would be very useful if [Virtual Biotech] can teach us some lessons from those failures that are actionable and can help us to improve things in the future.” Within hours, it delivered 3 key insights:

- Target prioritization: Therapies targeting specific cancer cell types were 48% more likely to receive Food

- and Drug Administration approval, with 32% fewer adverse events.
- Failure analysis: Virtual Biotech learned from clinical trial failures by identifying potential drug pathway redundancies; it proposed testable hypotheses and clinical markers to improve future success.
- Therapeutic strategy: The system proposed a new therapeutic intervention for lung cancer by connecting multidomain biological evidence from the scientific record to related clinical trial data.

Virtual Biotech’s proposals still require human scrutiny and wet lab validation, but importantly, it produced testable hypotheses—potentially saving time and money before resources are committed to a clinical study. The concept is not intended to be fully autonomous, with Zou stressing that experts must always provide “guardrails and feedback, define research priorities, and review the agents’ work.”

Human at the Helm, Not on the Oars

What is the exact role of the scientist in this new era? In a [recent review](#), Tobias raised an important concern: the training of junior scientists and protection of their roles, recommending that “publicly funded scientific bodies and universities... spend their grant dollars on people.” He states that, “in [his] view, training the next generation of scientists is [still] an academic’s primary job.”

AI will redefine the traditional division of labor in science. King and others have [proposed](#) that AI scientists should be given responsibility for well-defined, data-intensive, low-risk repetitive tasks while humans, in turn, retain control of problem formulation, societal relevance, ethics, and the interpretation and translation of findings.

Another reason for keeping a professor in the loop: [epistemic collapse](#), the point at which, in operating beyond the capabilities of human understanding, computers may become recursive, relying on their own outputs and contributing to “science itself becoming self-referential and narrow.” Keeping our scientific knowledge base tethered

to our humanity—with its associated priorities, ethics, and values—is key to maintaining scientific integrity and trust in science itself. Another powerful incentive for humans in science: money. Patent law in many countries only allows for the ideas of “[natural persons](#),” not those of AI systems, to be patentable.

Keeping Science Honest

The unprecedented speed and scope at which machine learning operates will only accelerate in the future, hastening new discoveries and opportunities for science. But human governance of the process will be critical, as the potential for misadventure will advance with it. For example, Anthropic

recently released an unreviewed article where the Claude Mythos LLM “discovered” three unknown reverse transcriptase gene families in jumbo phages—viruses that infect bacteria. The findings were immediately challenged by Mario Rodriguez Mestre, a doctoral student in Denmark who had used Claude as part of his [unpublished research](#) that reported the new genes. Did Claude discover new knowledge or did it simply learn from information provided by a user?

AI scientists promise to reduce the material and time costs of conducting science. In this new era, old principles still apply: with expert human oversight, data validation, and external review, AI scientists may retain our trust while helping to expand our knowledge base.

Keywords: artificial intelligence; automation; robotics; scientific research; laboratory research; data validation; research personnel; AI scientists; agentic AI; AI

Please cite as:

Dominy C

From Robots to Organizations: The Role of the AI Scientist

J Med Internet Res 2026;28:e114408

URL: <https://www.jmir.org/2026/1/e114408>

doi: [10.2196/114408](https://doi.org/10.2196/114408)

© JMIR Publications. Originally published in the Journal of Medical Internet Research (<https://www.jmir.org>), 23.Oct.2026