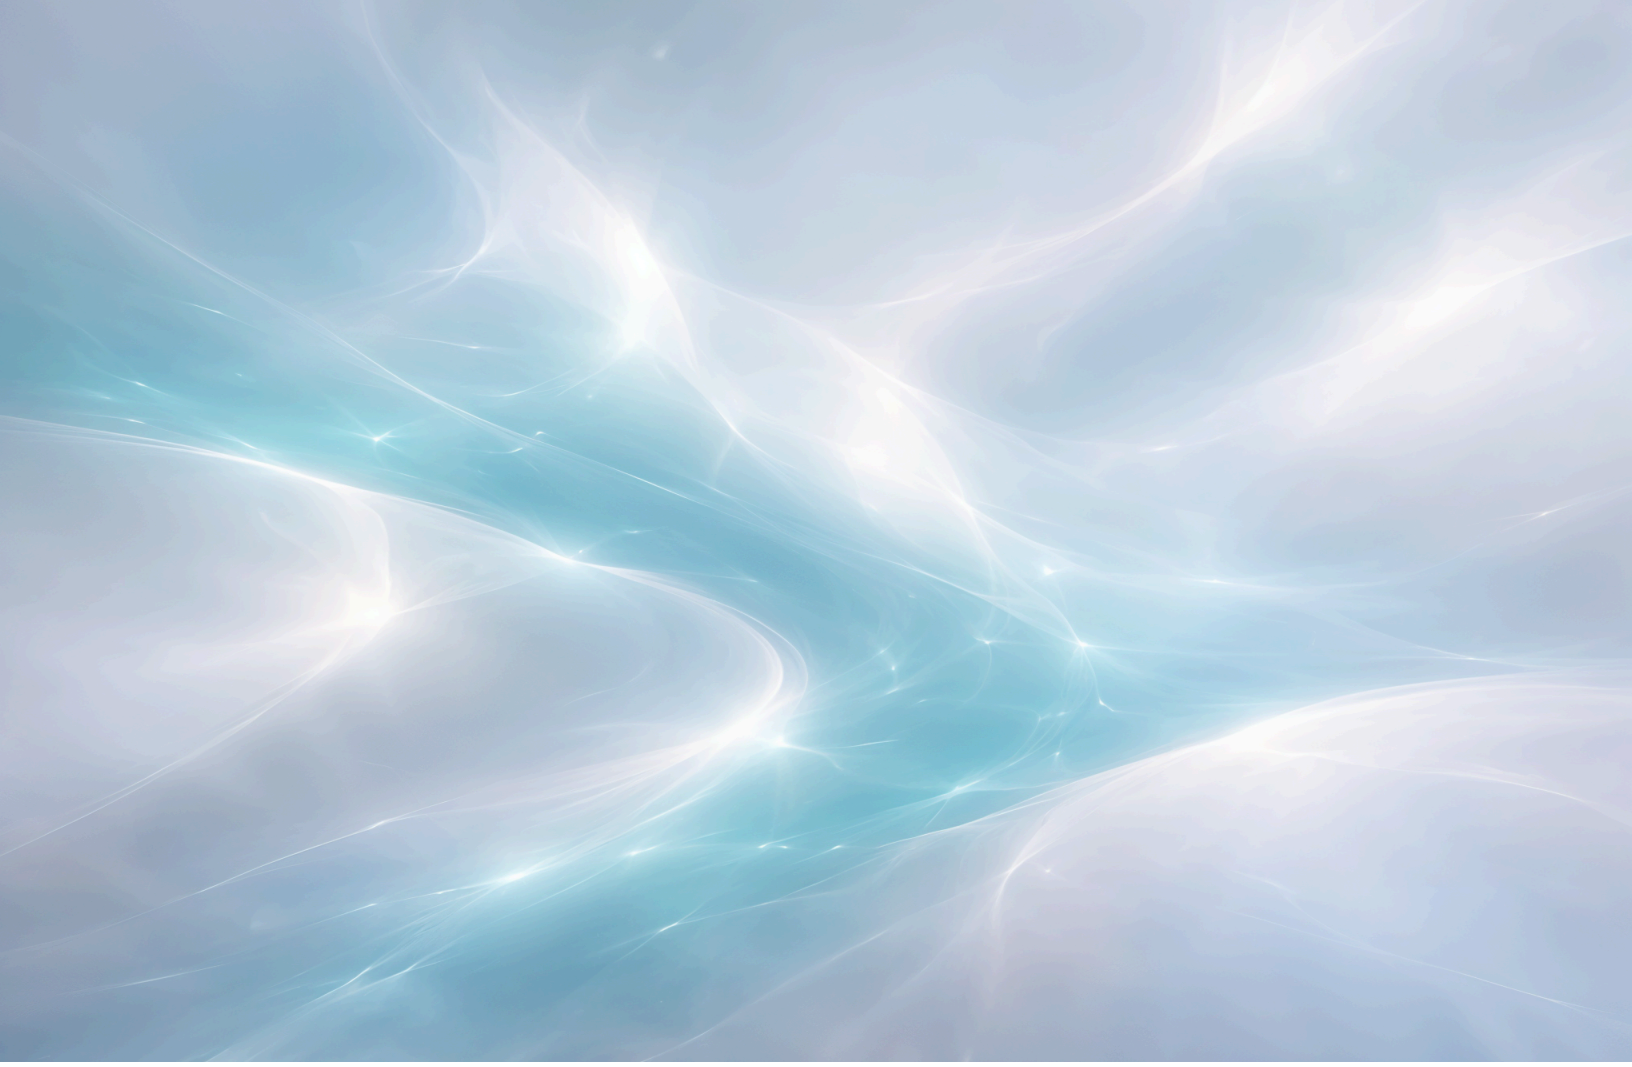

Reigniting the Discovery Engine for Tomorrow's Cures

Why AI Matters for the Life Sciences

April 2026

OpenAI

Introduction

The life sciences have done more than any other branch of applied science to improve human welfare, if you measure by lives saved, suffering prevented, and years gained.

In 1900, infectious disease caused roughly one-third of all deaths in the United States. Pneumonia, tuberculosis, and diarrheal disease were the leading killers. A minor cut could turn into fatal sepsis. In war, infected wounds killed more soldiers than actual combat. Medicine had no reliable cure for bacterial infection.

The century that followed changed the human condition. Antibiotics, beginning with the mass deployment of penicillin in the 1940s, have saved hundreds of millions of lives. The golden age of antibiotic discovery, from the 1940s through the 1960s, produced nearly two-thirds of the drug classes still in use today. Vaccines against just 14 pathogens have saved another 154 million lives over the past 50 years, roughly six every minute. Smallpox, which killed 300 million people in the 20th century alone, was eradicated. Global life expectancy rose, on average, by more than 30 years.

Yet the engine of biomedical discovery is slowing. Drug development costs roughly double every nine years. Clinical trials now take an average of nearly 7.5 years to complete. And the remaining burden of our most lethal conditions is still immense: cancer, heart disease, stroke, diabetes, and other noncommunicable diseases account for 74 percent of global deaths. Faster progress would save and prolong countless lives.

AI is beginning to jumpstart that progress, with approximately 200,000 global weekly users – from the United States and Great Britain, to India and South Korea – already relying on ChatGPT for advanced life sciences work that helps researchers integrate knowledge and translate across the disciplinary gaps that act as barriers to discovery.

These users include a seemingly terminal patient who used ChatGPT to outpace his rare cancer while processing massive amounts of personal data to identify experimental treatments; rare-disease families that have used ChatGPT to go beyond their diagnostic odyssey and surface possible treatments, communities, and experts; and AI-native startups that have built lean platforms for novel drug candidates.

Powerful New Tools Lead to New Discoveries

Over many centuries, medicine has cleared a series of hurdles to better understand the human body. Each threshold expanded what doctors could understand, as well as the treatments and cures at their disposal. Vesalius, in 1543, made the body readable at the level of organs. Cell theory, developed in the mid 1800s, established that disease is ultimately a cellular dysfunction, and must be treated at the cellular level. Each advance opened possibilities that the earlier view could not imagine: first in surgery, then pathology, germ theory, and oncology.

Genes as Code: The Genomic and Computational Revolution

We are crossing a new threshold, which allows us to see more deeply and act more effectively. Cells are governed by code, and biology has entered its genomic and computational era. The Human Genome Project, completed in 2003 at a cost of roughly $3 billion, produced the first essentially complete sequenced human genome. Since then, sequencing costs have fallen by seven orders of magnitude, outpacing Moore's Law, to a few hundred dollars per genome. Biology has become an information science. DNA is a four-letter digital code that can now be read at industrial scale for trivial marginal cost.

The change is larger than cheaper measurement. For the first time, the body is readable, writable, and computable within the same scientific frame. Vesalius could see organs, but he could not redesign them. Virchow could describe cells, but he could not reprogram them. With CRISPR and its successors, genetic code can now be edited with precision. With AI, the likely consequences of those edits can increasingly be predicted. The body has become, in an operational sense, programmable.

The Productivity Crisis: Why Breakthroughs Have Slowed

Eroom's Law: The Inertia of Greater Knowledge

Despite this revolution in capability, biomedical R&D has become steadily less productive. Scannell and colleagues coined "Eroom's Law," the reverse of Moore's Law on the rate of progress in transistors, to describe this discouraging pattern: the number of new drugs approved per billion dollars of inflation-adjusted R&D has fallen by roughly 80-fold since 1950. Development cost per approved drug has risen from about $5 million in the 1960s to about $2.6 billion today. That decline has persisted despite an 800-fold improvement in combinatorial chemistry, a billion-fold improvement in DNA sequencing, and a tripling of health R&D's share of total spending. Technical power has expanded, but therapeutic output has not kept pace, to the detriment of patients in need.

New drug development costs more time as well as more money. Median clinical development for innovative drugs takes 8.3 years, and the full path from research to approval spans 12 to 15 years. Clinical trials have lengthened rather than shortened: 89.8 months on average in 2014 to 2018, up from 83.1 months in 2008 to 2013. Only 7.9 percent of drugs that enter Phase I trials ultimately reach approval. By contrast, the FDA's own review time has fallen from 26.6 months to 9.9 months since 1992. The bottleneck lies in the clinical pipeline, where time, cost, and attrition compound.

The Burden of Knowledge and the Great Endarkenment

The slowdown in life-science progress is not confined to pharmaceuticals. Bloom et al showed that across industries, research effort keeps rising while per-researcher productivity falls sharply. U.S. research productivity has declined 41-fold since the 1930s. The reason is structural. As knowledge accumulates, each generation of scientists must spend more time training, only to specialize more narrowly. The polymath who once moved across multiple domains – a Pasteur, a Darwin, a Mendel – has largely disappeared from modern science.

In the 19th century, Pasteur moved from crystallography to fermentation to infectious disease to vaccination while inventing the germ theory of disease. Darwin combined geology, comparative anatomy, and biogeography into a single explanatory framework that we call natural selection. Their breakthroughs depended on integration across fields. Today, what Darwin knew as natural history has split into evolutionary biology, population genetics, genomics, phylogenetics, ecology, paleobiology,

developmental biology, and a long list of further specialties. Biomedicine alone now encompasses more than 20 distinct fields, each with its own journals, professional networks, and mutually incomprehensible vocabularies.

Philosopher Elijah Millgram has given this condition a name, The Great Endarkenment. He argues that hyperspecialization not only raises coordination costs, but also creates mutual unintelligibility. Specialists adopt vocabularies, methods, and standards of evidence so particular to their own domains that they struggle to evaluate reasoning in closely neighboring fields. Historically, the Enlightenment assumed that more knowledge would produce better collective judgment. Millgram's counterpoint is that beyond a certain threshold, the collective grows darker because no one can see enough of the whole. In the life sciences, where decisive problems span scales from gene to cell to tissue to organism, and where progress depends on many disciplines moving together, that condition is especially severe. The problems to be solved do not recognize disciplinary boundaries.

ChatGPT and Life Sciences

Approximately 200,000 worldwide users each week rely on ChatGPT for graduate-level or professional life sciences work. Most of that activity is concentrated in literature review: about 75 percent of identified life sciences users use ChatGPT to find, synthesize, and compare scientific literature. Roughly 25 percent use it to support analysis. All shares below are approximate.

Top Countries by Share of Life Sciences Usage	Top U.S. States by Share of Life Sciences Usage	Top Countries, as a percentage of Weekly Average Users (WAU)	Top U.S. States, as a percentage of Weekly Average Users (WAU)
1. United States	1. California	1. Singapore	1. Vermont
2. India	2. New York	2. South Korea	2. Maine
3. Great Britain	3. Texas	3. United States	3. Massachusetts
4. South Korea	4. Massachusets	4. Great Britain	4. New Hampshire
5. Germany	5. Illinois	5. Canada	5. California

AI at the Frontier of the Life Sciences

This is what makes the life sciences such an important domain to which we can apply AI. They combine the largest measurable effect on human welfare, the weakest long-run productivity trend, and, because of the genomic revolution, an unusually deep well of digital biological data. What follows is a survey of places where AI is already producing tangible results.

Protein Structure and Molecular Prediction

In 2020, DeepMind's AlphaFold 2 solved the 50-year grand challenge of protein structure prediction and then released predictions for more than 200 million structures, a task that would have required hundreds of millions of researcher-years if done experimentally. AlphaFold has been cited more than 20,000 times, and the work behind it was recognized with the 2024 Nobel Prize in Chemistry. In one large-scale 2025 impact study, researchers building on AlphaFold 2 submitted experimental protein structures at rates roughly 45 to 49 percent higher than comparable peers, with evidence of greater novelty, and papers linked to AlphaFold 2 were about twice as likely to be cited in clinical articles.

AlphaFold 3, released in 2024, extended these capabilities to the structure and interactions of all major biological molecules, including proteins, DNA, RNA, and drug-like small molecules. It outperformed the best physics-based tools for predicting drug-target binding. Isomorphic Labs, a DeepMind spinoff, is now applying these models directly to rational drug design.

Genome Editing: From CRISPR to Programmable Biology

The gene-editing method CRISPR-Cas9, demonstrated by Doudna and Charpentier in 2012, made the genome writable by repurposing a bacterial immune system to cut DNA at specified sites. The first CRISPR therapy, targeting sickle cell disease, won FDA approval in late 2023. Clinical trials are now underway for cancers, lupus, HIV, and inherited blindness. Yet CRISPR-Cas9 is only the first stone in the foundation of a broader toolkit.

Base editing, developed by David Liu in 2016, changes a single DNA letter without cutting the double strand, an important advance because roughly 30,000 known single-nucleotide mutations are associated with human disease. Prime editing, introduced in 2019, rewrites longer stretches of sequence using an RNA template and supports all base conversions as well as small insertions and deletions. Both approaches have entered human clinical trials. Epigenetic editing, which uses

deactivated Cas9 fused to regulatory proteins, can raise or lower gene expression without changing the DNA sequence itself. In animal models, a single administration targeting the PCSK9 gene reduced LDL cholesterol by roughly 70 percent.

In early 2025, a multi-institutional team developed the first personalized CRISPR therapy for an infant in just six months, a proof of concept for on-demand gene medicine. In a 2025 Nature paper, large language models trained on CRISPR diversity designed entirely novel gene editors that outperformed natural ones. AI is moving from helping scientists use gene-editing tools to helping design the tools themselves.

Self-Driving Labs and the Automation of Discovery

A new class of self-driving laboratories now combines AI with robotic automation to run the full cycle of scientific experimentation, from hypothesis generation and experimental design to execution, analysis, and iteration, continuously and at machine speed. These systems compress the design-make-test-learn loop from months to days.

Insilico Medicine has reported using AI to identify a target for idiopathic pulmonary fibrosis and nominate a preclinical candidate in 18 months, later advancing an AI-discovered and AI-designed program into Phase I in under 30 months. Exscientia produced the first AI-designed molecule to enter human clinical trials in under 12 months. More broadly, companies such as LabGenius and Recursion are building increasingly automated drug-discovery workflows that combine machine learning, experimental iteration, and large-scale biological data.

Clinical Trials: Attacking the Real Bottleneck

Clinical trials are the longest, costliest, and most failure-prone part of the bottleneck. Because most candidates fail here, even modest gains in speed or selection can have large downstream effects on overall productivity. AI is beginning to intervene directly by improving trial design, identifying eligible patients, stratifying populations, monitoring safety signals in real time, and predicting which trial arms are likely to fail early. One analysis estimated that these tools could cut clinical-phase timelines by more than 20 percent.

Regulators are adapting to this shift. In December 2025, the FDA qualified its first AI drug development tool, AIM-NASH, for liver disease assessment in clinical trials. In May 2025, the agency began deploying AI in its own review processes and issued risk-based guidance for sponsors using AI in drug development. As of late 2023, AI-native biotechs and their pharma partners had brought 75 molecules into the clinic, with 67 still in ongoing trials; by mid-2025, no AI-discovered or AI-designed drug had yet progressed through phase 3 trials.

AI as Translator: Reversing the Great Endarkenment

AI's deepest effect on the life sciences may be structural rather than tool-specific. Hyperspecialization created silos because no individual could absorb enough adjacent knowledge to see the connections among them. AI systems can be trained across the literature, data, and methods of multiple disciplines at once. In practice, they function as a shared cognitive substrate across fields.

Before AlphaFold, a drug discovery team that needed a protein structure might wait months for a specialist pipeline. AlphaFold reduced that wait to minutes for any researcher with an internet connection. Structural biology did not disappear; the bottleneck it created for neighboring fields did. The same pattern is appearing elsewhere. A geneticist can predict mutation effects without becoming a biophysicist. A drug designer can screen billions of candidates without owning a chemistry lab. Each case removes a coordination barrier. Each is a small reversal of Millgram's Great Endarkenment.

The 19th-century polymath achieved integration by carrying several disciplines in one mind. The 20th-century research team achieved it through collaboration, though often at high cost in time, coordination and money. The emerging model is the AI-augmented team, in which AI performs some of the cross-domain translation that the polymath once handled internally. This does not restore the conditions of the 19th century – the knowledge base is far too large for that – but it does accelerate teams and their research.

Use case profiles

The use cases below show how AI is already helping life-sciences work clear hurdles that would normally slow it down – from patient records research literature and genomic data, to wet-lab automation, model reasoning and narrow specialties inside small teams.

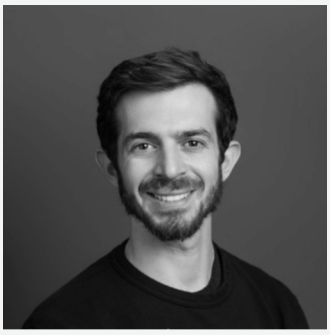

Sid Sijbrandij and Valius: Patient-driven medicine

When Sid Sijbrandij's rare cancer, an osteosarcoma, came back and standard options ran out, he went into founder mode. His team pulled together imaging, monthly blood tests, immune monitoring, RNA sequencing, single-cell data, and specialist review so that each new signal could be read against the rest. ChatGPT helped the team, led by Jacob Stern, handle much of the connective work: writing glue code, doing deep research across a wide range of disparate topics, and turning fresh results into questions for the next expert conversation. AI enables the team to move fast enough to keep ahead of a lethal disease. To support other patients in their race against rare cancers, Jacob and Sid joined up with Edward Larkin to found Valius Sciences. Most cancer diagnostics are still built for broad categories of patients, not for individuals with rare cancers. A patient often tests only a short list of genes, and receives a label based on where the cancer started, such as lung, colon, or brain. The result is usually a static report: a snapshot of their cancer at one moment in time. But the clues that matter most may lie elsewhere, e.g. in gene expression, which shows which genes the tumor is actually using, or even in cell-surface targets, the proteins on the outside of cancer cells that newer drugs are designed to recognize. Valius is built to perform that deeper read for patients, and translate what that means

for care. In its first eight months, the company has analyzed 15 client tumors and found actionable treatment leads in every case, meaning real therapies doctors could plausibly pursue. Four clients were already on single-patient INDs, the FDA pathway for using an experimental drug in one person. In one glioblastoma case, an aggressive brain cancer, standard DNA testing found no mutation that clearly pointed to a treatment. But Valius's broader analysis instead found high DLL3 expression, meaning the tumor was making a large amount of a cell-surface protein that some newer drugs can target. That put a real option on the table for the patient: using a drug off-label, outside the specific cancer it was originally approved for, or seeking FDA permission to use it for that one patient. Instead of handing patients one more report, the company is trying to turn molecular data into a dynamic guide of what they can do next, aided by ChatGPT.

Sequencing.com: Conversations with genomes

Sequencing.com makes genome sequencing valuable even before a family manages to find the right specialist. One family that relied on Sequencing's product, Kate and her daughter Ellen, illustrate how this works: Ellen was born very small, with microcephaly, low muscle tone, and later developmental delays. Even after a NICU stay, her family still had no clear explanation. Tests for a possible infection during pregnancy were negative, and the specialists they saw mostly told them to wait and see. Kate heard about direct-to-consumer genome sequencing in a Facebook group and paid herself to get it done with Sequencing.com. When the results came back, Sequencing.com surfaced a CTNNB1 finding linked to a rare neurodevelopmental disorder. The platform's AI tool, ChatGPT, helped turn that result into something Kate could understand well enough to act on. While the finding was upsetting, it gave Kate a direction: she sought out other families, pushed for a formal genetics evaluation, and clinical exome sequencing later confirmed Ellen's condition. Sequencing.com and ChatGPT helped move the family out of diagnostic limbo and into the right medical conversation.

Dr. Oral Alpan:
Repurposing FDA-approved drugs

Immunologist Dr. Oral Alpan is using GPT-5 Pro to find new uses for FDA-approved drugs in diseases that still lack solid treatments. The work began with a patient under his care for severe eczema who also had Food Protein-Induced Enterocolitis Syndrome, or FPIES, a rare allergy that causes delayed, severe inflammation in the gut. For two decades, even trace amounts of wheat brought vomiting, cramping, or diarrhea. Then the patient started dupilumab for eczema, accidentally ate wheat while traveling, and had no reaction. Back in clinic, Dr. Alpan supervised an oral food challenge approaching 50 grams of wheat protein. Again, no reaction. When insurance later cut off access to dupilumab, the symptoms returned; when the drug restarted, they cleared. Amerimmune then identified seven more patients with favorable food-allergy responses while taking dupilumab for approved uses. Before those findings had entered the literature, Alpan's colleague Dr. Derya Unutmaz gave GPT-5 Pro a de-identified version of the case. The model ranked dupilumab first and identified patient profiles that might face side effects. Dr. Alpan is now extending that approach through IRB-approved in vitro blood-assay studies at his combined clinic and research lab, testing how existing drugs affect different immune phenotypes and feeding each round of de-identified results back into ChatGPT to refine the hypothesis, narrow in on the right patient subgroup, and guide the next experiment. He hopes the same cycle can surface overlooked therapies in other conditions, including subtypes of POTS.

Anton Maximov:
AI as a daily research collaborator

Anton Maximov's lab at Scripps Research epitomizes how AI can help a scientist integrate many methods and disciplines at once. His neuroscience work spans

genetics, sequencing, biochemistry, optical imaging, 3D electron microscopy, electrophysiology, and behavior. In a 2025 Science paper, specialized AI image analysis made a large 3D-EM reconstruction tractable that would have taken years by hand. ChatGPT plays a different role in that environment, but an important one. Maximov uses it as an editor, coding aide, literature synthesis tool, and sounding board. In Maximov's lag, AI helps researchers move between neighboring literatures fast enough to keep the whole research question in view, even as they conduct experiments the old-fashioned way.

Nasha Fitter and Citizen Health: Rare-disease families as community intelligence

Nasha Fitter and her daughter Amara experienced something rare-disease families know well: diagnosis is the end of one uncertainty, and the start of many others, including the search for treatment. After Amara was diagnosed with FOXG1 syndrome, a rare and severe neurodevelopmental disorder, Nasha helped build the research infrastructure the field lacked. Through the startup she co-founded, Citizen Health, she turned that same instinct into a platform where rare-disease families can bring medical records together, extract clinically important details, capture symptoms and visits in real time, and ask questions that draw on their own data, broader medical knowledge, and insights from similar patients. At home, that helped Nasha reason through Amara's gastrointestinal pain by linking low muscle tone to foods that were hard to digest. In research, Citizen's medical-record-derived historic symptom data will be used as a comparator for a FOXG1 trial, with FDA alignment on study design. With ChatGPT, Citizen is connecting better care and evidence generation to support new treatments for the families who need them.

Junevity: A biotech team moving fast toward the clinic

Junevity was built to run with a team much smaller than a traditional biotech. Co-founded by Janine Sengstack, Rob Cahill and John Hoekman, the company has roughly fifteen people and core consultants, most with little overlap in expertise. Its lead program, JUN_01, is an siRNA therapy for type 2 diabetes, a gene-silencing treatment the company plans to take into clinical trials as early as Q4 2026. Cahill says building around AI from day one let Junevity move toward the clinic at a fraction of the usual cost. Inside the company, AI is less a discovery engine than a working language: it helps people in clinical, toxicology, regulatory, and target-discovery roles stay in sync without routing every question through another specialist.

Ginkgo Bioworks and Retro Biosciences: Model-to-lab biological design

Ginkgo Bioworks connected GPT-5 to a cloud lab for cell-free protein synthesis, a way to make proteins outside living cells. The model proposed reaction conditions, the lab ran them, the results came back, and the next round was planned from those measurements. Over six rounds, the collaboration tested more than 36,000 conditions and reduced protein-production cost by 40 percent. Separately, Retro Biosciences used a similar loop for cellular reprogramming. OpenAI generated candidate variants of OSKM factors, the four proteins commonly used to reset cell identity, and Retro tested them experimentally. In both collaborations, the work moved through the same sequence: model design, physical experiment, measured result, revised design. Scientists in the lab used AI to shorten the distance between an idea and the next experiment.

Policy

With AI, the life sciences are moving back toward a fusion of disciplines that the burden of knowledge has driven apart. Genomics turned the genetic code into a digital stream of information. CRISPR and its successors, including base editing, prime editing, and epigenetic editing, made that code writable and programmable with growing precision. AI systems trained on massive biological datasets began to interpret the code's complexity at a scale no specialist could match. Robotic automation is closing the loop between computation and wet-lab validation. Together, these advances create a compounding cycle: reading genomes, computing predictions, writing changes, and validating them in automated labs.

Policy should treat AI in the life sciences as a new scientific instrument, widening access to the inputs that make that cycle work: data, compute, infrastructure, and talent.

Open and securely connect life sciences data

Policymakers should make high-value scientific and medical datasets more usable for AI-enabled discovery by facilitating machine-readable access, reducing information-blocking, and extending portability beyond traditional electronic health records to laboratory results, pharmacy data, wearable and home-monitoring data, and care-coordination information. The goal should be a system in which patients, researchers, and authorized institutions can assemble fragmented information into usable context for diagnosis, trial matching, and discovery. That openness should be paired with strong privacy protections: de-identification for sensitive data and strong limits on the use of sensitive health data for model training.

Treat frontier AI access as a shared national research resource

To help the scientific community advance AI-enabled life sciences research, it is important to give universities, national laboratories, and nonprofit research institutions sustained access to advanced AI systems for experimentation and validation. Public-private research partnerships should complement that access, including new collaboration vehicles that pair AI labs, academic researchers, and mission-driven scientific institutions on ambitious problems in biology, drug discovery, and adjacent fields.

Build the physical stack

AI can accelerate target discovery, molecular design, and trial optimization, but those gains will stall if the United States lacks the physical capacity to test, validate, and scale what the models produce. Policy should treat energy, compute, chips, lab automation, preclinical testing capacity, biomanufacturing, and modern clinical-trial infrastructure as a connected innovation stack. That means investing in AI infrastructure hubs, streamlining permitting for high-capacity compute, and expanding

the experimental and trial capacity needed to move promising experiments from simulation into real-world validation.

AI-skilling the next generation of scientists and life-sciences workers

The life sciences will increasingly reward people who can work across biology, code, statistics, automation, and evidence generation. Policy should support upskilling at all levels of education, expand certification programs, apprenticeships, and community-based AI training so the benefits of AI-enabled science are shared broadly. The strategic aim should be to broaden the pool of people who can use AI tools to accelerate discovery.