

OpenAI releases GPT-Rosalind for life sciences research through trusted access

Source copy kept by wheresthe.ai for the story “OpenAI releases GPT-Rosalind for life sciences research through trusted access” (<https://wheresthe.ai/d/openai-gpt-roosalind-life-sciences-2026-04/>).

Source: <https://openai.com/index/introducing-gpt-roosalind/>

Text saved 18 Oct 2026 (the page text through a reader, because the site refused a direct fetch).

SHA-256 of the saved text: 30e60c2ddb12e413228d3cad814df82d72a22764a9b365de9f83228ce1d0c488

****Update on September 11, 2026**** **__GPT-Rosalind__**, our specialized model for life sciences research, is coming out of research preview and is now available globally to eligible organizations through our trusted-access program. All eligible organizations will continue to get access to the latest Rosalind models as they’re released. Published pricing will take effect on **__October 5, 2026__**. Review **__GPT-Rosalind pricing__** (opens in a new window) **__**.

Today, we’re introducing GPT-Rosalind, our frontier reasoning model built to support research across biology, drug discovery, and translational medicine. The life sciences model series is optimized for scientific workflows, combining improved tool use with deeper understanding across chemistry, protein engineering, and genomics.

On average, it takes roughly 10 to 15 years to go from target discovery to regulatory approval for a new drug in the United States. Gains made at the earliest stages of discovery compound downstream in better target selection, stronger biological hypotheses and higher-quality experiments. Progress in the life sciences is constrained not only by the difficulty of the underlying science, but by the complexity of the research workflows themselves. Scientists must work across large volumes of literature, specialized databases, experimental data, and evolving hypotheses in order to generate and evaluate new ideas. These workflows are often time-intensive, fragmented, and difficult to scale.

We believe advanced AI systems can help researchers move through these workflows faster—not just by making existing work more efficient, but by helping scientists explore more possibilities, surface connections that might otherwise be missed, and arrive at better hypotheses sooner. By supporting evidence synthesis, hypothesis generation, experimental planning, and other multi-step research tasks, this model is designed to help researchers accelerate the early stages of discovery. Over time, these systems could help life sciences organizations discover breakthroughs that wouldn’t otherwise be possible, with a much higher rate of success.

GPT-Rosalind is available in ChatGPT, Codex, and the API for qualified customers through our trusted access program. We’re also introducing a freely accessible Life Sciences research plugin for Codex, helping scientists connect models to over 50 scientific tools and data sources. We are working with customers like Amgen, Moderna, the Allen Institute, Thermo Fisher Scientific, and others to apply GPT-Rosalind across workflows that accelerate research and discovery.

The model is named after Rosalind Franklin, whose rigorous research helped reveal the structure of DNA and laid foundations for modern molecular biology.

From raw data to grounded discovery decisions, see how our purpose-built model accelerates research workflows.

Built for scientific workflows

The GPT-Rosalind life sciences model series is built for modern scientific work across published evidence, data, tools, and experiments. In our evaluations, it delivers the best performance on tasks that require reasoning over molecules, proteins, genes, pathways, and disease-relevant biology, and it is more effective at using scientific tools and databases in multi-step workflows such as literature review, sequence-to-function interpretation, experimental planning, and data analysis.

This is the first release in our GPT-Rosalind life sciences model series, and we will continue to expand the frontiers of the model’s biochemical reasoning capabilities across long-horizon, tool-heavy scientific workflows. OpenAI’s compute infrastructure gives us the ability to continue training, evaluating, and improving increasingly capable domain models against real scientific tasks—helping these systems become more useful as the

workflows themselves become more complex.

From evidence-based discovery insights to high-impact experiments, see how our suite of solutions translate into measurable improvements in your research workflows.

Customers and ecosystem

We are working with leading pharmaceutical, biotechnology, and research customers, as well as life sciences technology organizations, to apply GPT-Rosalind across workflows that drive discovery.

> “The life sciences field demands precision at every step. The questions are highly complex, the data are highly unique, and the stakes are incredibly high. Our unique collaboration with OpenAI enables us to apply their most advanced capabilities and tools in new and innovative ways with the potential to accelerate how we deliver medicines to patients.”

—Sean Bruich, Senior Vice President of Artificial Intelligence and Data, Amgen

Performance and evaluation

We evaluated GPT-Rosalind across a range of capabilities fundamental to scientific discovery and industry research. These evaluations measure core reasoning across scientific subdomains, including chemical reaction mechanisms; protein structure, mutation effects, and interactions; and phylogenetic interpretation of DNA sequences. They also assess whether models can support real research workflows by interpreting experimental outputs, identifying expert-relevant patterns, and synthesizing external information to design follow-up experiments. Finally, they test whether models can select and use the right computational tools, databases, and domain-specific capabilities to augment their reasoning. Taken together, these evaluations show progress across the end-to-end process of scientific research and suggest a stronger ability to help researchers work through challenging discovery tasks.

Industry evaluations

We evaluated GPT-Rosalind on a series of public benchmarks. On BixBench, a benchmark designed around real-world bioinformatics and data analysis, GPT-Rosalind achieved leading performance among models with published scores.

On LABBench2, a benchmark measuring performance on a range of research tasks such as literature retrieval, database access, sequence manipulation and protocol design, GPT-Rosalind outperforms GPT-5.4 on 6 out of 11 tasks. The most notable improvement comes from CloningQA, which requires end-to-end design of DNA and enzyme reagents for molecular cloning protocols.

We also partnered with Dyno Therapeutics, a company pioneering AI-designed gene therapies, to evaluate the model on an RNA sequence-to-function prediction and generation task using unpublished, uncontaminated sequences. Performance was compared against 57 historical scores from human experts in the AI-bio field. When evaluated directly in the Codex app, best-of-ten model submissions ranked above the 95th percentile of human experts on the prediction task and around the 84th percentile of human experts on the sequence generation task.

These evaluations provide a meaningful signal of performance on the kinds of workflows scientists rely on every day to generate evidence, analyze complex data, and move toward defensible biological conclusions.

Connecting to the tools scientists use

Scientists can use our new Life Sciences research plugin(opens in a new window) for Codex, available today in GitHub. This package includes a broad set of modular skills for most common research workflows, designed to help users work across human genetics, functional genomics, protein structure, biochemistry, clinical evidence, and public study discovery.

These skills act as an orchestration layer that helps scientists work through broad, ambiguous, and multi-step questions more effectively. They provide access to more than 50 public multi-omics databases, literature sources, and biology tools, and offer a flexible starting point for common repeatable workflows such as protein structure lookup, sequence search, literature review, and public dataset discovery.

Eligible Enterprise users can leverage this plugin in research workflows with GPT-Rosalind for deeper biological reasoning, while all users can use the plugin package with our mainline models.

Trusted access

We want to make these capabilities available to the scientists and research organizations best positioned to advance human health, while maintaining strong safeguards against biological misuse. The Life Sciences model is launching through a trusted-access deployment structure for qualified Enterprise customers in the U.S. to start, with controls around eligibility, access management, and organizational governance. At the same time, we are making a set of connectors and the Life Sciences Research Plugin available more broadly, so researchers can use our mainline models more effectively for life sciences research tasks.

The Life Sciences model was developed with heightened enterprise-grade security controls and strengthened access management, enabling professional scientific use in governed research environments. We evaluate access based on three core principles: beneficial use, strong governance and safety oversight, and controlled access with enterprise-grade security. In practice, this means participating organizations must be conducting legitimate scientific research with clear public benefit; maintain appropriate governance, compliance, and misuse-prevention controls; and restrict access to approved users within secure, well-managed environments.

Getting started

Organizations can request access([opens in a new window](#)) through our qualification and safety review process.

The Life Sciences model is built to help scientific organizations do higher-quality work, faster, in environments that require both technical capability and operational control. Our dedicated Life Sciences team—as well as advisory partners including McKinsey & Company, Boston Consulting Group (BCG), and Bain & Company—help organizations identify high-impact use cases, integrate the model into enterprise environments, and drive measurable outcomes. If you'd like to explore ways OpenAI Life Sciences can support your work, you can contact our Life Sciences team.

What's next

This is the first release in our Life Sciences model series, and we view it as the beginning of a long-term commitment to building AI that can accelerate scientific discovery in areas that matter deeply to society, from human health to broader biological research. We will continue improving the model's biological reasoning, expanding support for tool-heavy and long-horizon research workflows, and working closely with leading scientific institutions to evaluate real-world impact. That includes ongoing partnerships with national laboratories such as Los Alamos National Laboratory, where we are exploring AI-guided protein and catalyst design, including the ability of AI systems to modify biological structures while preserving or improving key functional properties.

Over time, we expect these systems to become increasingly capable partners in discovery—helping scientists move faster from question to evidence, from evidence to insight, and from insight to new treatments for patients.