

Anthropic launches Claude Science, a research workbench for scientists

Source copy kept by wheresthe.ai for the story "Anthropic launches Claude Science, a research workbench for scientists" (<https://wheresthe.ai/d/anthropic-claude-science-workbench-2026-06/>).

Source: <https://www.anthropic.com/news/claude-science-ai-workbench>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: 9105da8a79085c606c0d1bb8d97a22389e45a09527cfd6c9fb9daca3a6eb34ce

Announcements Beneficial Deployments Claude Science, an AI workbench for scientists, is now available Jun 30, 2026 Get started with Claude Science

AI has the potential to dramatically accelerate the pace of scientific discovery and the development of healthcare interventions. Since launching our efforts in the life sciences last fall, we've worked to improve our model capabilities, make connections to the scientific ecosystem via MCPs and skills, and launch partnerships in an effort to realize this potential. Today, we're introducing our most significant expansion of these efforts: Claude Science, an AI workbench for scientists. Claude Science is an app that integrates the tools and packages that researchers most commonly use, produces auditable artifacts, and provides flexible access to computing resources. Introducing Claude Science Scientific research is often tedious. Researchers must work across dozens of databases, each with their own schema, contend with file formats that require bespoke data pipelines and viewers, and transition between a roster of tools: PubMed, Jupyter, R, a cluster terminal, and more. Claude Science brings these fragmented tools into a single research environment where scientists can conduct all stages of their work. It helps you analyze literature and execute multi-step research, produces detailed artifacts, and lets you iteratively refine figures and manuscripts until they're ready for publication. Every output carries an auditable history of how it was made, so you can validate and reproduce the results. Like a Jupyter Notebook, you can access Claude Science wherever you already work—locally on macOS or Linux, or on a remote machine over SSH or with an HPC login node. Users interact with a generalist coordinating agent with access to over 60 curated skills and connectors pre-configured for genomics, single-cell, proteomics, structural biology, cheminformatics, and more. These agents can spin up others and engage with specialist agents created by users. And a reviewer agent checks citations and calculations, flagging and correcting errors. We are releasing Claude Science today in beta for Claude Pro, Max, Team, and Enterprise users, and will continue to refine the platform as we collect feedback from users. How it works Claude Science displays proteins, structures, and molecules natively, with every result reproducible and traced to its code. Rich scientific artifacts, fully reproducible. Scientific research is inherently visual, so Claude Science generates figures and manuscripts alongside the code that created them. It natively renders rich scientific artifacts, including 3D protein structures, genome browser tracks, chemical structures, and more. You can chat with the agent about any detail, annotating figures and manuscripts in-line so the agent knows what to address to make them publication-ready. When it generates a figure, Claude Science includes the exact code and environment that produced it, a plain-language description of how it was created, and the full message history. This allows you to understand the inputs, making the work easier to validate and reproduce even months later. You can ask Claude Science to make edits to figures in plain language—removing gridlines, for example, or changing an axis to log scale—and the agent will edit its own code. Claude Science builds environments and manages compute on your laptop, your cluster, or GPUs on demand. Manages your compute and scales on demand. Large analyses—folding a protein, for example, or running a genomics pipeline over a massive dataset—often require researchers to shift their focus to setting up a computing job, waiting while it's sent to a cluster, checking whether it succeeded or failed, and pulling the results back. Claude Science handles this process for you. It drafts a plan, asks before reaching new resources, and lets you review or revoke any decision before writing and submitting the job to the computing resources your lab already uses (your own HPC cluster over SSH, or your Modal account for compute on demand), scaling the analysis from a single GPU to hundreds as needed. Because its agents work inside a running session that holds context in memory, even massive datasets only need to be loaded once. It runs on your lab's own infrastructure—your laptop, Linux box, or HPC login node—so large or sensitive datasets never have to leave the systems they're already on, and only the context needed for each step of the analysis is sent to Claude. As the pipeline runs, a reviewer agent inspects the outputs, flagging incorrect citations, untraceable numbers, and figures that don't match their underlying code, and self-correcting as it goes. You can fork the session at any point to compare two approaches without losing the original thread. Claude Science is pre-configured for genomics, single-cell, proteomics, and cheminformatics,

backed by more than 60 scientific databases. Domain-ready on day one . Scientific knowledge is scattered across hundreds of specialized sources. In biology, for example, relevant data might sit across resources such as UniProt, PDB, Ensembl, Reactome, ClinVar, ChEMBL, GEO—each with its own schema and query language—as well as in journals and preprint servers, and domain-specific open models. When you ask Claude Science a question in plain language, specialist agents query and synthesize across all of these sources so you don't have to navigate them individually. Claude Science uses the skills in NVIDIA's BioNeMo Agent Toolkit to connect natively to the life sciences models and libraries in BioNeMo , including Evo 2, Boltz-2, and OpenFold3. Scientists already have models, datasets, and pipelines they trust. Claude Science can connect to these as well, saving any pipeline as a reusable skill or accessing your lab's preferred tool using a connector, with future sessions inheriting them automatically. This customizability allows you to access Claude, your proprietary data, and the validated tools you already rely on in one conversation. Claude Science benefits from our partners' specialized expertise and platforms, while more scientists reach their tools through Claude. What scientists are doing with Claude Science Over the past few months, researchers have worked with Claude Science in beta for tasks like single-cell RNA sequencing analysis, CRISPR screen design, protein structure prediction, cheminformatics, and more. Manifold Bio designs tissue-targeting medicines—which home to a specific organ or cell type, so the drug acts where it's needed and spares the rest of the body—and tests how millions of candidate binders corresponding to hundreds of targets distribute through a living body at once. Manifold used Claude Science to nominate the targets for its latest experiments. For each tissue and target, Claude Science assessed surface expression, trafficking, and safety, ranking candidates against the criteria Manifold has learned from its own internal proprietary data. What set Claude Science apart from a general coding assistant, Manifold said, was that it could do this end-to-end, gathering the right data and applying the right judgment with the context of past programs built in. Jérôme Lecoq, a neuroscientist at the Allen Institute, used Claude Science to build a multi-agent “computational review template” comprising about 20 custom skills geared towards writing long-form reviews. The sub-agents read through thousands of papers, pulling the central claim and the key quantitative finding, and storing them in an evidence state database. Then the pipeline constructs a narrative arc, writing the review section by section and delegating each to its own specialized sub-agent. Within each section, dedicated agents generate quantitative cross-study figures directly from the evidence database. A key component of the workflow, enabled by Claude Science, is the use of actor-critic pairs: one agent creates content while a separate reviewer agent evaluates it for accuracy and citation fidelity. Before Claude Science, it could take Lecoq's team as many as two years to write such a review. He now has about 10 reviews, many more than 100 pages, with citations that were checked over by reviewer agents. The team is now working with domain experts to further refine the AI-based critic agents. And Stephen Francis, an associate professor and epidemiologist at the UCSF Brain Tumor Center, has used Claude Science to support studies on the molecular epidemiology of glioma, a type of primary tumor that begins in the glial cells of the brain. His lab investigates the genetic basis for how thousands of small-effect germline variants combine to shape individual susceptibility. Although this work predated Claude Science, Francis said the app has dramatically accelerated the analysis, enabling comprehensive germline workups across multiple approaches in roughly one-tenth the time it previously took. His group independently validated Claude Science's results, confirming that it can produce both rapid and robust analyses. Getting started with Claude Science The Claude Science app is available in beta on macOS and Linux for Pro, Max, Team, and Enterprise plans. We're sharing it early so scientists can start to use it on real problems and tell us how to refine it. Team and Enterprise users will need their admin to enable Claude Science. We now have a Team plan offering discounted seats for active scientific labs at academic institutions and nonprofit research organizations; learn more here . We'll also be supporting up to 50 Claude Science AI for Science projects, providing up to \$30,000 in credits. Modal will also be providing up to \$2,000 in compute for select projects. We are looking for projects that span domains and explore the boundaries of science, with an early focus on biology and biomedical research. Applications are open through July 15, 2026, with award notifications sent out by July 31. Projects will run from September 1 to December 1, 2026— apply here . To stay up-to-date on product announcements, provide feedback, and learn from others in the Claude Science community, join the AI for Science Discourse community . Get started with Claude Science at claude.com/science .

Related content

An update on independent evaluators To evaluate the safety of our systems, Anthropic is giving employee-like access to a team from METR, a nonprofit that evaluates frontier AI models, and we're in early discussions with several other organizations, including Hugging Face. Read more Preparing for the 2026 US midterms Earlier this year, we shared an update on our election safeguards. Here, we outline how we've continued to strengthen these safeguards ahead of this year's US midterms. Read more Anthropic opens Singapore office and signs MOU with

the Government of Singapore Anthropic has opened an office in Singapore and signed an MOU with the country's government to advance safe, responsible AI adoption. [Read more](#)