



Using X-Labs to Unleash AI-Driven Scientific Breakthroughs

How to adapt our science funding mechanisms to the unique infrastructure needs of large-scale AI project | **Caleb Watney**

Using X-Labs to Unleash AI-Driven Scientific Breakthroughs

How to adapt our science funding mechanisms to the unique infrastructure needs of large-scale AI projects | Caleb Watney

This essay is part of [The Launch Sequence](#), a collection of concrete, ambitious ideas to accelerate AI for science and security.

Summary

Publicly funded science has long been an engine of prosperity and national power. AI for science now promises to supercharge this engine. Foundation models have already delivered major breakthroughs in protein folding prediction, materials discovery, and genomic variant analysis. As this series explores, similar advances could transform everything from pathogen detection to software security and drug design.

Many of the most important AI for science efforts could slip through the cracks of federal funding — too multidisciplinary, infrastructure- and engineering-heavy, or speculative for traditional grants. To operationalize this new scientific engine, the US needs to update its scientific funding infrastructure to support institution-scale efforts.

We propose funding 25 “X-Labs”: independent research organizations funded through a new four-part X Series award framework.

- **X01 (Excellence):** Breakthrough basic science institutions.
- **X02 (Execution):** Focused nonprofits building critical tooling with startup-like agility.
- **X03 (Experimentation):** Portfolio-based regranteeing organizations.

- **X04 (Exploration):** Planning grants to test a proof of concept.

Each X-Lab would receive between \$10 to \$50 million per year for seven years, supporting team-based, exploratory, and infrastructure-heavy work, with a hard cap on the percentage of renewals to maintain a dynamic portfolio.

The program could launch immediately via Other Transaction Authority (OTA), allowing science funding agencies to establish X-Labs without new legislation. Each participating agency (likely NSF, NIH, and DOE) would retain control over its own awards while coordinating within a unified X-Labs framework. Congressional appropriations could further expand the program's scale, and philanthropic matching funds could amplify its impact through public-private partnerships.

Motivation

Scientific innovation is the primary driver of economic growth, accounting for about [50%](#) of annual GDP growth, [25%](#) of post-World War II productivity growth, and the marginal \$1 of public R&D returns approximately [\\$2-5 in GDP](#). Advanced AI has the potential to supercharge this growth engine. For instance:

- [AlphaFold 2](#) now predicts protein structures in minutes rather than the months or years required experimentally.
- [GNoME](#) predicted 2.2 million potential crystal structures, including about 380,000 that appear thermodynamically stable, in just 17 days.
- [Evo 2](#) analyzed over 9 trillion base pairs and can identify pathogenic BRCA1 mutations with over 90% accuracy and perform other large-scale genomic analyses.

As this series highlights, similar breakthroughs, from self-driving labs to metagenomic biosurveillance to secure chip verification, could transform health, energy, and national security. But these projects require a fundamentally different model of scientific work.

AI for science efforts are engineering-heavy, infrastructure-intensive, and multidisciplinary. They often involve full-time teams of ML researchers, biologists, roboticists, and software engineers, organized around shared tools, open-ended

goals, and high-variance bets. These teams often require the flexibility to pivot as experiments unfold, as the most promising research direction is frequently unknown at the outset.

To replicate these breakthroughs at scale, we must address a structural mismatch in our research funding ecosystem. The majority of federal science funding flows through mechanisms designed for individual scientific investigators pursuing discrete, short-term projects — a model fundamentally ill-suited for transformative AI initiatives.

The NIH R01 grant, the workhorse of biomedical research funding, exemplifies this mismatch. R01s typically fund individual principal investigators to work on a specific, pre-specified project for 3–5 years with budgets of around \$600,000 in total over that period. This structure creates multiple barriers to AI breakthroughs:

- **Resource constraints:** Training foundational AI models requires computational infrastructure and datasets costing far more than entire R01 budgets. AlphaFold's development required years of sustained investment that no traditional project-based grant could support.
- **Expertise fragmentation:** AI for science requires teams that span machine learning, domain science, software engineering, and data infrastructure — expertise that rarely exists within single university labs funded through project-based grants.
- **Structural limitations:** R01s weren't intended to fund the shared computational infrastructure, curated datasets, and specialized platforms that AI initiatives require, leaving researchers to cobble together inadequate resources through overhead or philanthropic support.

Traditional labs also rely heavily on graduate students and postdocs — a model that prioritizes training but undermines continuity and specialization. In contrast, X-Labs would provide stable, full-time roles for staff scientists, engineers, and technicians, enabling institutional memory and technical depth more typical of industrial R&D.

Even when researchers manage to force AI for science work into existing grant formats, the process is [slow](#) and often [ineffective](#). Reviewers tend to favor low variance proposals that closely resemble past work. And our best investigators,

the ones who should be maximizing their time in the lab, spend nearly [half of their time](#) on grant writing.

It's telling that many of the most important AI for science breakthroughs have come not through traditional R01-style grants, but through newer institutional models — often backed by philanthropy or bespoke public-private partnerships. Evo 2 was developed at the philanthropically funded Arc Institute. AlphaFold came out of DeepMind. GNoME was built at Google. These projects succeeded not because they conformed to a traditional grantmaking structure, but because they had the luxury to bypass it, combining deep infrastructure, interdisciplinary teams, and stable, long-term support. Even when more traditional funding mechanisms are used, they are often cobbled together, as with the [Brain Initiative Cell Census Network](#), which required coordination across multiple NIH centers using a patchwork of U19s, U01s, and R01s.

Other programs haven't been able to fill the gap. DARPA and ARPA-H still rely on project-based grants and do not offer the long-term, institution-scale support needed for exploratory, infrastructure-heavy work. National labs are powerful but bureaucratically encumbered and limited in topical scope. NIH center grants and NSF research hubs require specifying projects upfront, and, in practice, often end up funding piecemeal collections of individual PIs' projects, rather than focused, integrated teams. Even when university research succeeds, its downstream impact is often bottlenecked by tech transfer offices. These offices, designed for risk-averse IP management, frequently impose burdensome licensing fees and equity terms that deter external partners. X-Labs could bypass these constraints, enabling public-interest science to be shared and scaled more freely.

In the mid-20th century, industrial labs like Bell Labs and Xerox PARC served as hubs of long-term, exploratory research. Today, with most industrial R&D focused on short-term product development, the public sector must take the lead. Many AI for science opportunities create public goods or operate years ahead of commercial readiness. Even well-funded startups struggle to finance foundational research or shared infrastructure. And while commercial AI labs like Google and OpenAI have partnered with scientists on specific challenges like protein folding, their core incentives ultimately revolve around product development and shareholder value, not long-term public benefit or foundational science. If the US wants to lead in this next era of discovery, we need publicly funded institutions

built for AI-native science from the ground up. Without them, we risk missing world-changing breakthroughs and ceding leadership to nations with more adaptive funding systems.

Solution

Launching the X-Labs Initiative

To close the growing mismatch between legacy funding mechanisms and the needs of AI-native science, the federal government should fund 25 X-Labs — independent research institutions awarded competitive block grants through a new “X-Series” program. Each X-Lab would receive an **X01**, **X02**, or **X03** award, funded at \$10–50 million per year over a seven-year cycle. No more than 70% of labs would renew into a second term, ensuring continuous dynamism, institutional experimentation, and room for new entrants.

These X-Labs would fill a longstanding structural blind spot in the US research ecosystem: work that is infrastructure-intensive, team-based, exploratory, or oriented around critical bottlenecks — and therefore poorly served by the individual PI model. Rather than replace existing mechanisms like the NIH R01 or NSF core programs, X-Labs would complement them by enabling types of research that traditional university grants weren’t built to support.

To support future entrants, the program would also offer **X04 (Exploration)** grants: \$1–3 million over two years to help new teams refine their vision, build partnerships, and complete early proof-of-concept work before applying for full institutional funding.

The four X-Series award mechanisms

Each X-Series award corresponds to a distinct gap in the research ecosystem:

- **X01 (Excellence)** funds minimally constrained basic science institutions — the closest federal analogue to models like [Janelia](#), the [Arc Institute](#), or the [Allen Institute](#). These labs are structured to support long-term discovery, deep collaboration, and institutional memory, giving world-class teams the freedom

and flexibility to pursue open-ended research unconstrained by bureaucracy. The core bet behind X01s is on people, not projects — the goal is to assemble the best team in the world to pursue open-ended scientific inquiry with minimal bureaucratic constraint.

- **X02 (Execution)** targets clearly defined bottlenecks: missing datasets, infrastructure, or platform technologies. These awards fund time-limited, high-impact teams that operate like [nonprofit startups](#) or Focused Research Organizations, executing against well-scoped scientific challenges that are critical but commercially underincentivized. Think of these as scalable, public-interest analogues to focused AI labs or instrumentation teams. The fundamental selection principle is the challenge: funding a talented group with a nimble organizational structure to execute against a clearly defined bottleneck in the scientific ecosystem.
- **X03 (Experimentation)** supports [portfolio-based](#) regranteeing and incubation organizations — meta-level institutions that fund early-stage ideas before they become consensus picks. Inspired by models like [Convergent Research](#), [Speculative Technologies](#), and [Science Angels](#), these labs operate outside traditional grantmaking channels and experiment with new science-funding methodologies. Some awards would explicitly support metascience to evaluate and improve selection practices. The core idea is to empower scientific scouts — individuals or teams with the judgment, network, and conviction to spot promising talent or research directions early.
- **X04 (Exploration)** provides seed funding of between \$1M to \$3M over a few years to support the formation and planning of new scientific institutions, enabling teams to refine their vision, build key partnerships, and develop initial proof-of-concept work before applying for full X01, X02, or X03 funding.

Historical analogues to X-Labs

Organization	X-Labs analogue	Years of operation	Primary funding source	Area of science	Key features	Results
Bell Labs	X01	1925 – 1996	Industry (AT&T)	Physics, computing, materials, telecom	Interdisciplinary tool-rich lab with massive impact	Transistor, laser, cellular networks, UNIX operating system
MRC Laboratory of Molecular Biology	X01	1962 – present	Government (UK)	Molecular biology, biochemistry	Long-term, staff-driven institute with Nobel-heavy output	DNA structure discovery, monoclonal antibodies, x-ray crystallography
Janelia Research Campus	X01	2006 – present	Philanthropy (HHMI)	Neuroscience, imaging, bioengineering	Flat, engineer-heavy lab focused on tool-building and discovery.	Light-sheet microscopy, GCaMP neural imaging tools, brain mapping advances
Arc Institute	X01	2021 – present	Philanthropy (multiple donors)	Genomics, cell biology, systems biology	Research institute prioritizing top talent and technical infrastructure	Evo 2 genomic AI model, disease mutation prediction tools
Allen Institute	X01	2003 – present	Philanthropy (Paul Allen)	Neuroscience, cell biology, immunology	Family of research institutes, mission-driven, tool and dataset-generating	Allen Brain Atlas, open neuroscience resources
Broad Institute	X01	2004 – present	Philanthropy + Government (NIH)	Genomics, precision medicine, computational biology	Large-scale genomics hub with strong data/platform focus	CRISPR tools, Cancer Dependency Map, human genome resources
Los Alamos (Manhattan Project era)	X02	1942 – 1946	Government (US)	Nuclear physics, applied engineering	Wartime team-science model focused on urgent engineering challenge	Nuclear technology, advanced computing, radiation detection systems
Wyss Institute	X02	2009 – present	Philanthropy (Hansjörg Wyss)	Biological engineering, diagnostics	Harvard-based bioengineering and translation center	Organs-on-chips, SHERLOCK diagnostics, biologically inspired engineering
DeepMind (early years)	X02	2010 – 2018	Venture Capital, then Industry (Google)	Artificial intelligence, ML, neuroscience	Mission-driven AI lab with major early breakthroughs	Deep reinforcement learning, AlphaGo, AI for scientific discovery

Using X-Labs to Unleash AI-Driven Scientific Breakthroughs 9 | Watney

E11 Bio	X02	2022 – present	Philanthropy (Convergent Research)	Neurotechnology, brain–computer interfaces	FRO building open neurotech platforms	Developing: accessible neural interface technologies
Cultivarium	X02	2022 – present	Philanthropy (Convergent Research)	Synthetic biology, experimental platforms	FRO tackling microbial tooling in synthetic biology	Developing: platform for previously uncultivable microbes
Convergent Research	X03	2020 – present	Philanthropy (multiple donors)	Biotech, tools, infrastructure	Nonprofit incubator for FROs solving tooling bottlenecks.	Launched 8+ FROs including E11 Bio, Cultivarium, [C]Worthy
Speculative Technologies	X03	2023 – present	Philanthropy (multiple donors)	Materials science, manufacturing, electronics	High-risk R&D programs for 'big-if-true' science	Developing: programs on molecular additive manufacturing and nanomodular electronics
Experiment Foundation	X03	2022 – present	Philanthropy (multiple donors)	Regranting, talent scouting, experiments	Scout-driven regranting and incubation platform	Science Angels program, 100+ grants to early-career researchers
DARPA (initially ARPA)	X03	1958 – present	Government (US DoD)	Defense and advanced technology R&D	Coordinated R&D programs run by empowered program officers	ARPANET, GPS, stealth technology, mRNA vaccine research

Implementation

X-Labs would be implemented via Other Transactions Authority (OTA), a flexible funding tool that allows agencies to bypass many traditional constraints on grants and contracts. Both NIH and NSF already have OTA authority through their Director's Offices and the TIP Directorate. But effective use of OTA requires specialized expertise. Agencies should bring in experienced talent — e.g., from DARPA and NASA — to structure agreements correctly and ensure the mechanism is used to its full potential. X-Labs should be deliberately structured for agility, with lean governance, minimal reporting overhead, and mission-driven cultures that empower small, high-trust teams. Like startups, their advantage is focus: aligning technical capability with a tightly scoped scientific challenge.

Awards should be open to applicants outside of academia, explicitly encouraging independent research organizations, public-private partnerships, and novel institutional models. This shift would reduce overreliance on university structures and expand the pool of high-functioning scientific institutions in the US.

Selection should prioritize scientific vision, execution capability, and institutional leadership — not just past funding success. With fewer, more targeted awards, agencies can prioritize quality over volume, recruiting a small bench of truly world-class reviewers with deep domain expertise to make qualitative, taste-based evaluations about novelty, feasibility, and potential impact. This is particularly important for **X03s**, where success may be better evaluated at the portfolio level, rather than by the individual success of each grant, mirroring how early-stage investors measure returns across a set of bets. The goal is to create a new tier of scientific selectors — the “angels” of science — capable of spotting and backing ideas years ahead of the mainstream.

If the model works, it should grow. The initial cohort of X-Labs would represent just ~1% of combined NSF, NIH, and DOE science budgets. But if these institutions demonstrate transformative impact, Congress could expand appropriations to scale the program 5-10x over time. The goal is not to constrain X-Labs, but to pilot a new model for scientific institutions that earns the right to grow.

Further resources

- Ben Reinhardt, "[Fund Organizations, Not Projects](#)," Institute for Progress, 2022.
- Adam Marblestone et al., "[Unblock Research Bottlenecks with Non-Profit Start-Ups](#)," Nature, 2022.
- Michael Nielsen and Kanjun Qiu, "[A Vision of Metascience](#)," Scienceplusplus, 2022.
- Ben Reinhardt, "[Unbundling the University](#)," Speculative Technologies, 2025.