

GUEST ESSAY

We Didn't Build the Atomic Bomb This Way

July 3, 2026

**By Dan Rockmore**

Dr. Rockmore is a professor of mathematics and computer science at Dartmouth College.

See more of our coverage in your search results.

[Add The New York Times on Google](#) ↗

Artificial intelligence, funded overwhelmingly by private capital, has careened forward despite immense concerns about the effects it will have on labor, education, science, defense and civic life. A.I. companies have outpaced public oversight and, at times, successfully lobbied against it. The central achievements of the industry, the proprietary “frontier models” developed by companies like Anthropic, OpenAI and Google, are guarded intellectual properties even as they are incorporated into schools, offices, hospitals, courts, commerce and our everyday devices. The public did not ask for these A.I. tools and now can hardly opt out of them.

There is no one-size-fits-all solution to addressing the impacts that may be coming. But the scale of the concerns requires ambitious responses that serve the public. The United States can start by building a national A.I. laboratory.

Think about the atomic age. The early work in the United States was carried out across a classified but federally accountable archipelago with national laboratories in four sites: Berkeley, Oak Ridge, Hanford and, most famously, Los Alamos. They were expensive and sometimes secretive, but connected to universities as well as the rest of our nation’s research apparatus.

From its inception, the nuclear enterprise was regulated, imperfectly but deliberately, by the federal government. The Manhattan Project that successfully created the atomic bomb gave birth to the Atomic Energy Commission, which watched over weapons development and seeded a nascent nuclear energy industry. The Defense Department, the Department of Energy and intelligence agencies now keep an eye on the weapons, while the Nuclear Regulatory Commission oversees the civilian nuclear sector. Research in applied nuclear physics continues apace both in the open and in secret.

Imagine if, instead of a national effort to make an atomic bomb, a small number of private companies had gotten there first and began selling powerful nuclear technologies back to the nation, with the government trying to catch up afterward. Would we have accepted that arrangement? Would we have said that because innovation was fast, regulation should wait? Would we have trusted a few executives and technologists to decide what the public needed to know?

Sign up for the Opinion Today newsletter Get expert analysis of the news and a guide to the big ideas shaping the world every weekday morning. [Get it sent to your inbox.](#)

The analogy is imperfect — A.I. is not a bomb and we're not fighting a world war. But like all technologies, it can be exploited for dark purposes. Yet the United States apparently has no public institution capable of building, testing and understanding frontier A.I. at the same scale as the companies that now define it.

A federal A.I. lab could change this by ensuring that the American public has at least one institution that can build, see and test the most advanced A.I. systems in the public interest. Such a lab should be built by, and provide controlled access for, university researchers, and it would be a training ground for our next generation of A.I. architects. It could publish research openly where possible, while classifying work only where necessary. It could study risks without needing to protect a new product launch. It could build models for science, education, medicine and national security — not as corporate afterthoughts, but as public missions that aren't subverted by investor concerns.

The government has made modest moves in this direction, but nothing that seems to come close to the scale of a national A.I. lab. The most ambitious effort is the White House's "Genesis Mission." This initiative tasks the Department of Energy and its national laboratories to use federal data sets to build an A.I. "platform" that would encompass a suite of A.I.-related software and hardware to accelerate scientific research. This is a public-private collaboration with participation by some of the leading A.I. firms, but the advanced models the companies bring are likely to be still proprietary, and the terms of access are at best currently vague. To date, engagements with these proprietary models in other national lab programs appear to have largely been at the level of prompting them, rather than getting a look under the hood of the technology itself. The real learning can take place only with the access that comes from building a frontier model from soup to nuts.

The country can't afford to drag its feet here. A.I. is advancing and embedding itself faster than perhaps any technology before it. The more firmly private systems become entrenched, protected and indispensable, the harder it becomes for any public institution to catch up.

Would this be too expensive? Not really — or at least, not relatively. Sandia National Laboratories, the home of nuclear stockpile testing and assurance, lists total fiscal-year 2025 funding of about \$5.8 billion. National military spending has approached \$1 trillion a year. Estimates for standing up a lab capable of building frontier A.I. models vary widely, running from \$6 billion to \$20 billion.

The talent exists, too. Even when the private sector is humming, not all computer scientists opt for industry rewards. Research freedom and a love of educating and mentoring a next generation of scientists drive many of us to look to academia. Recent hiring slowdowns across the tech sector mean there are more top-flight recent graduates eager for ambitious work, making a noncorporate path more attractive. A robust national A.I. lab would offer another path in public service for great computer scientists and mathematicians, the sort of path that gifted lawyers, economists and scientists have long found in working for the greater good.

There are also geopolitical pressures. “A.I. sovereignty” initiatives — which incorporate national priorities and values in the development, deployment and governance of these technologies — are underway in Britain, the United Arab Emirates and China, the last of which has set explicit goals for global A.I. leadership by 2030. Without a U.S. national lab, we rely on private companies for these judgments and then hope that public oversight can be added later. That order is backward, and the hope naïve. A nation and its people cannot govern a technology they do not deeply understand, and they cannot deeply understand a technology from the outside.

We are letting a handful of executives develop one of the most consequential technologies in human history, sell it into public life and define its future for us. The public absorbs the risks while surrendering not only oversight but understanding and participation. We did not run the atomic age that way. We shouldn’t run the A.I. age that way either.

Dan Rockmore directs Dartmouth College’s Neukom Institute for Computational Science and the Dartmouth Kalaniyot program. He is an external faculty member at the Santa Fe Institute.

The Times is committed to publishing a diversity of letters to the editor. We’d like to hear what you think about this or any of our articles. Here are some tips. And here’s our email: letters@nytimes.com.

Follow the New York Times Opinion section on Facebook, Instagram, TikTok, Bluesky, WhatsApp and Threads.