Stanford University
Human-Centered
Artificial Intelligence

Issue Brief
HAI Policy & Society
December 2024

Expanding Academia's Role in Public Sector AI

Kevin Klyman, Aaron Bao, Caroline Meinhardt, Daniel Zhang, Elena Cryst, Russell Wald, Fei-Fei Li

Introduction

ACADEMIC RESEARCHERS INVENTED NEARLY ALL OF THE CORE TECHNOLOGIES UNDERPINNING AI. From the 1963 founding of Stanford's AI Lab, the home of many foundational AI breakthroughs, to the creation of the groundbreaking image classification architecture AlexNet at the University of Toronto in 2012, the academy has been at the heart of the field from the very beginning. For AI to be developed responsibly and in the public interest, academia must continue to play a central role.

In the last decade, however, the field has been increasingly dominated by the private sector. Building and deploying AI systems has become hugely resource intensive, often requiring billions of dollars in investment, custom supercomputing clusters, and enormous datasets containing much of the available data on the internet. This shift has created a significant power imbalance, where academic talent and government support flows to private companies that now produce the vast majority of the world's most powerful AI systems.

As AI systems have grown more capable, the costs of developing foundation models has become a substantial barrier to entry. In 2017, Google spent less than $1,000 to build its first transformer-based AI model—by 2023, it cost Google almost $200 million worth of computational resources to develop its state-of-the-art model Gemini Ultra. In a world where the cost of building a single foundation model is equivalent to the annual operating budget of an entire university, few academics can meaningfully participate in the development of state-of-the-art AI models.

Key Takeaways

Academia is falling behind industry in frontier AI research. Today, no university in the world can build a frontier AI system on par with industry.

Industry is dominating AI development due to its massive datasets, unprecedented computational power, and top-tier talent. Companies have over 1,000x more compute than universities, and they produce AI models that are 50x larger.

Governments should continue investing in public sector AI. Academia must be at the forefront of training the next generation of innovators and advancing cutting-edge scientific research in the public interest.

This disparity undermines not only the future of academic research but also the potential for a public sector AI ecosystem that serves the public interest. Unlike industry, academic research is driven not by profit but by the pursuit of scientific knowledge. Time and time again, pathbreaking AI innovation has come from the curiosity-driven research of academics who have the freedom to pursue ideas that are not immediately commercializable. Academia must play a leading role in developing frontier AI to ensure that we can understand and safely deploy the technology.

In 2017, Google spent less than $1,000 to build its first transformer-based AI model—by 2023, it cost Google almost $200 million to develop its state-of-the-art model Gemini Ultra.

The Industry-Academia Divide in AI

Until 2014, academia produced the largest number of notable machine learning models each year. Since then, academia has continued to advance, while industry has raced ahead. In 2023, industry produced 51 notable machine learning models, far outstripping academia's 15; that same year, there were just two notable machine learning models produced by governments and none from nonprofits or government collaborations with academia.

The gap between academia and industry can be measured across three dimensions: funding, compute, and talent. In each of the three, the resource divide between researchers in academia and industry continues to grow.

The funding disparity between academia and industry is long-standing, but it has expanded significantly in the last decade. Since 2013, aggregate private investment in American AI companies has exceeded $300 billion, whereas many universities have seen a decline in funding in real terms as government support has not recovered following the 2008 global financial crisis. A decade after the Great Recession, U.S. states spent on average 16 percent less on higher education, and many states made further reductions as part of their budget cuts associated with the COVID-19 pandemic. Academic institutions do not have the resources to keep pace with Big Tech or many startups, and the trend lines are heading in a troubling direction.

Computational resources are an essential building block of AI. Building large AI models requires advanced chips capable of carrying out trillions of operations in parallel. Without chips and the energy to run them, it is impossible to build AI systems like ChatGPT. In recent months, top American universities have announced sizable purchases of the graphics processing units (GPUs) needed to build foundation models. Consider access to Nvidia H100s, the current best-in-class GPU. In 2023, Harvard announced its purchase of 384 H100s, while this year UT Austin announced it would purchase 600 and Princeton 300.

These universities are now among the world's top academic institutions in terms of computational power.

Nevertheless, these purchases pale in comparison to those of major AI companies. The same week that Princeton announced it would purchase 300 H100s, Meta announced it would buy 350,000. Microsoft plans to have 1.8 million H100s by the end of this year, and the startup xAI is using a supercomputing cluster of 100,000 H100s to train its latest Grok model. Among academics with access to compute, it is typical for them to have access to between 1 and 8 GPUs, whereas industry researchers may have access to thousands. This gap in raw computational power has led to a situation where many foundation models from industry are more than 50 times larger than those from academics.

The same week that Princeton announced it would purchase 300 H100s, Meta announced it would buy 350,000. Microsoft plans to have 1.8 million H100s by the end of this year.

Academia must also contend with issues related to attracting and retaining talent. In 2011, PhD graduates in AI were equally likely to pursue careers in academia or industry. By 2020, the balance had shifted decisively in favor of industry, with nearly 70 percent of new AI PhDs pursuing careers in the private sector. The compute divide paired with funding disparities has contributed to an increase in top talent moving from academia to industry. Access to computational power determines the scale and complexity of experiments that researchers can conduct, leading top researchers to gravitate to organizations that can offer them sufficient compute to carry out cutting-edge research. Without access to adequate computational resources, academics cannot complete even basic research projects that involve building small foundation models.

In addition to extra compute, industry positions pay substantially more. Whereas the average computer science professor in the United States earns $113,000 per year, research scientists in AI at companies such as Meta earn $335,000 in total compensation—nearly triple their academic counterparts. The earnings differential has led many faculty members to leave the academy for the private sector. Nationwide, the net flow of talent from academia to industry has doubled in recent years. The brain drain from academia to industry is noticeable across the country, with computer science departments losing many of their brightest researchers and professors to industry.

Why This Matters: Less Innovation and a Weaker Talent Pipeline for the Future

High barriers to entry due to the cost of compute also restrict independent, public-interest AI research. Historically, academic institutions have been crucial sources of fundamental scientific breakthroughs

that prioritize public benefit over commercial gain. The internet, GPS, and the core technologies underlying modern AI all emerged from academic and government research labs focused on long-term impact rather than short-term returns for shareholders. The current resource disparity makes cutting-edge AI research difficult to conduct in academic settings and inaccessible to research groups without substantial outside funding.

Today's AI industry is highly concentrated, with a handful of powerful companies controlling the most powerful models, the largest data centers, and top talent. As in other industries, reduced competition due to market concentration can have negative effects, such as higher prices for consumers, supply chain inefficiencies, and reduced innovation. In the case of AI, powerful companies have attracted a disproportionate amount of funding, compute, and talent, preventing many organizations in academia, government, and civil society from making a contribution to the development and deployment of leading-edge AI systems.

This type of concentration of power among private sector AI developers has profound implications for the field and for academia in particular. When even top U.S. research universities like Stanford and Princeton can access only a fraction of one percent of the computing resources available to major companies, they cannot offer students hands-on experience in building frontier AI systems. As a result, students are heavily incentivized to take their talents to industry, where they can master cutting-edge techniques for building AI systems that shape society. If the next generation can only learn how powerful AI systems actually work by working for industry, academic training grounds for future scientists will be diminished.

Today's AI industry is highly concentrated, with a handful of powerful companies controlling the most powerful models, the largest data centers, and top talent.

Today, smaller and less well-resourced academic institutions often cannot meaningfully participate in the development of cutting-edge AI systems, leaving immense value on the table that could be captured by including thousands of smaller groups. There is potential for transformative work through this model: Academic research can take place over the long term as it is not tied to rapid product development cycles—this methodical, long horizon research often results in more innovation than rapid A/B tests conducted by companies rushing to market.

Without adequate resources for academic AI research, discussions of AI policy have also become skewed. Big Tech companies are among the most active and well-funded lobbyists in Washington, reducing the amount of airtime for public interest research and perspectives. With adequate resources, academia could act as a more neutral arbiter in policy discussions of the costs and benefits of AI systems—without many of the conflicts of interest that plague industry advocates.

A Stronger Role for Academia Can Transform Public Sector AI

We need a new generation of institutions to create a public option for safe and trustworthy AI.

Universities and policymakers should adopt a three-pronged approach to ensure a stronger role for academia in AI and a whole-of-government approach to public sector AI.

1. Scaling public investment in academic AI research

In 2021, the global AI industry invested over $340 billion in AI development, while U.S. government agencies allocated just $1.5 billion to academic AI research—a ratio of more than 225 to 1. This massive funding gap between public and private investment is part of why universities do not have a seat at the table when it comes to frontier AI.

Recent policy proposals suggest growing recognition of this challenge. The National Security Commission on AI called for $32 billion in annual federal spending on non-defense AI innovation, which has been taken up by the Bipartisan Senate AI Working Group. Still, even this would amount to less than 10 percent of industry's current investment. More ambitious funding models are needed and could draw inspiration from successful examples like Switzerland's National Supercomputing Centre, which has created a national research infrastructure with over 10,000 H100s available to academics. The U.S. State Department's Critical Language Program offers another example of federal investment in talent development to advance national security that has equipped thousands of students with essential skills for public service. A similar model should be pursued to develop AI talent with an eye toward public service and national security as the future of AI is central to the competition for global economic leadership.

Effective public investment needs to go beyond funding. Government initiatives should focus on creating shared infrastructure and resources that maximize impact. The National AI Research Resource (NAIRR) pilot is an excellent start, but it must be adequately resourced in order to meet the moment. Bills like the CREATE AI Act would represent transformational, once-in-a-generation investment in the future of public sector AI, helping academia scale the impact of public interest research. Funding for universities to acquire training data, build interdisciplinary partnerships, and develop talent is also a critical piece of a winning public sector AI strategy.

2. Embracing "team science"

The complexity of modern AI development demands a shift in how academic research is organized. Solving the toughest scientific questions of the 21st century will require working across disciplines to make use of massive computational and data resources. The traditional model of individual principal investigators leading small teams of graduate students cannot

match the pace and scale of industry research. Even generous grants cannot keep up with multibillion-dollar investments by venture capital into AI startups. And academic administrators, unlike company executives, do not see their primary role as helping to advance deployment of AI systems, even if doing so would strengthen their institutions. For academia to participate in the development of frontier AI, universities need to reinvent how they carry out research to be less bureaucratically burdensome and more hospitable to cutting-edge AI research.

Universities should consider adopting a "team science" approach that brings together interdisciplinary expertise in order to advance the state of the science for AI. Academia can learn from success stories in the private sector, such as Google DeepMind. DeepMind has brought together top experts in machine learning, chemistry, ethics, data science, and other fields to produce some of the most impactful AI systems of the 21st century. DeepMind's AlphaFold model, which earned its developers a Nobel Prize for its ability to predict protein structures, helped solve one of the most challenging open problems in science in a matter of weeks.

Physics has also made remarkable advances through team science. From CERN to Fermilab, some of the most important discoveries about the nature of our universe have come from major investments in research infrastructure and broad academic collaborations. NASA's Frontier Development Lab offers another promising model. In collaboration with the private sector and government partners, NASA offers support for interdisciplinary academic research at the intersection of AI and space. This had led to new discoveries across astrophysics, planetary science, and climate adaptation.

Implementing team science requires universities to rethink traditional academic structures. This includes developing new models for credit attribution, tenure evaluation, and resource allocation that motivate and reward collaborative research. Universities must also create physical infrastructure that enables large-scale collaboration, following examples like Stanford's SLAC National Accelerator Laboratory, which has successfully integrated AI research across multiple scientific domains to make fundamental breakthroughs in theoretical physics.

3. A new generation of government-backed academic institutions

While public funding and organizational reform are essential, universities must also develop new models for collaboration with government on public sector AI. Strategic partnerships with government can provide academia with access to crucial resources while maintaining research independence and focus on the public interest.

Government-backed research organizations have made fundamental scientific advances in the past, from RAND and MITRE to Lawrence Livermore and Los Alamos national laboratories. We need a new generation of institutions to create a public option for safe and trustworthy AI. These partnerships must go beyond traditional research sponsorship and facilitate access to industry compute resources and datasets while preserving academic freedom and intellectual property rights. With government support, new academic research institutions could receive computing resources at discounted rates, though this will not make up for a lack of in-house compute. With adequate government support for a new generation of institutions, the academy can make a difference in shaping the future of frontier AI.

Conclusion

The industry-academia divide in AI is widening dramatically. Expanding academia's role in public sector AI through additional funding and institutional reform would be a transformational shift in the future of public interest technology and innovation as a whole. Investing in public AI infrastructure will empower academia and the public sector to drive innovation that prioritizes societal benefit, ensuring that the next generation of AI technologies aligns with the public interest and sets the United States on a path to maintaining its competitiveness in AI while strengthening its leadership in science and technology for decades to come.

Stanford University's Institute for Human-Centered Artificial Intelligence (HAI) applies rigorous analysis and research to pressing policy questions on artificial intelligence. A pillar of HAI is to inform policymakers, industry leaders, and civil society by disseminating scholarship to a wide audience. HAI is a nonpartisan research institute, representing a range of voices. The views expressed in this policy brief reflect the views of the authors. For further information, please contact **HAI-Policy@stanford.edu**.

Kevin Klyman is an MA candidate in international policy at Stanford University and a graduate fellow at the Stanford Center for Research on Foundation Models (CRFM). His work on this brief was completed prior to becoming a government employee.

Aaron Bao is a research assistant at the Stanford Institute for Human-Centered Artificial Intelligence (HAI).

Caroline Meinhardt is the policy research manager at Stanford HAI.

Daniel Zhang is the senior manager for policy initiatives at Stanford HAI.

Elena Cryst is the director of policy and society at Stanford HAI.

Russell Wald is the executive director of Stanford HAI.

Fei-Fei Li is the Sequoia Professor of Computer Science at Stanford University and the Denning Co-Director of Stanford HAI.

Stanford University
Human-Centered
Artificial Intelligence

Stanford HAI: 353 Jane Stanford Way, Stanford CA 94305-5008

T 650.725.4537 **F** 650.123.4567 **E** HAI-Policy@stanford.edu **hai.stanford.edu**