

PUBLIC SUBMISSION

Received: May 29, 2025 Tracking No. mba-74qk-g6hb Comments Due: May 28, 2025 Submission Type: Web

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0334
Comment on FR Doc # 2025-07332

Submitter Information

Name: Mihir Kshirsagar

General Comment

See attached file(s)

Attachments

Princeton Researchers Comment May 29 2025

May 29, 2025

Docket ID No. NSF-2025-OGC-0001

Networking and Information Technology Research and Development (NITRD)

National Coordination Office (NCO), National Science Foundation

via regulations.gov

**Comments of Researchers at Princeton University on the 2025 National AI R&D
Strategic Plan, NSF-2025-OGC-0001**

Thank you for the opportunity to comment on the 2025 National AI R&D Strategic Plan. We are academic researchers associated with the Princeton Laboratory for Artificial Intelligence (“AI Lab”) and the Center for Information Technology Policy (“CITP”) at Princeton University and write to share our thoughts on the research needs and development challenges in AI that the Federal government should prioritize over the next 3 to 5 years.¹

Like the Federal government, our AI Lab is dedicated to research in areas that industry is unlikely to address. We focus on advancing fundamental research and incubating ambitious AI-related research initiatives across the natural sciences, engineering, social sciences, and humanities to better serve the public interest. Similarly, at CITP we work to understand and improve the relationship between digital technologies and society, and have submitted comments in a number of regulatory proceedings to assist the government’s consideration of complex policy choices.

¹ In keeping with Princeton’s tradition of service, the CITP Clinic provides nonpartisan research, analysis, and commentary to policy makers, industry participants, journalists, and the public. This response is a product of the CITP Clinic and reflects the independent views of the undersigned scholars.

We use “AI” to refer to a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions, influencing real or virtual environments. We use the term “[foundation models](#)” to describe powerful models that can be fine-tuned and used for multiple purposes. Recent advances in AI, in particular with foundation models, are poised to have a transformative effect on our society. The Federal government can play a critical role in harnessing the benefits of AI-related technologies and mitigating their risks by making substantial investments in research and development.

Our comment makes three core points. *First*, the government should prioritize promoting the diffusion of AI-related technologies. Specifically, this requires prioritizing research to aid public sector and government applications of AI. In addition, research funding should support the infrastructure, capabilities, and institutional adaptations that enable productive AI adoption. *Second*, the government should invest in supporting the development of open models to democratize access to technology. This includes research infrastructure for the AI communities at universities. *Third*, the government should prioritize research on the impact of AI on the workforce, anticipate potential disruptions, and develop strategies to address them.

More fundamentally, universities stand ready to serve as essential partners with the Federal government in realizing the societal benefits of AI. The NSF has long recognized the value of universities conducting research that may be less likely to happen in commercial settings. It should continue to support basic, foundational research that has a long term horizon, which helps us understand how complex technical systems work, and improve on them. Universities have also had considerable success translating such research into commercial products that benefit society. In particular, Princeton’s AI-related research is helping make advances in other fields, including healthcare, engineering, and science. We have found that this progress requires interdisciplinary work that encompasses expertise across multiple fields – a role that universities, with their commitment to public service, are uniquely positioned to serve. Finally, not only do universities

provide education and training for workforce development, but they also help create experts in these fields. NSF funding for students through programs like the Graduate Research Fellowship Program and grants for early career researchers is therefore essential to maintaining U.S. leadership in AI development.

A. Promote diffusion, not just innovation

The Federal AI R&D strategy should recognize a fundamental distinction that shapes how technologies transform economies: the difference between invention and diffusion. The invention of new technologies is the specific point at which a new type of technology is created, such as the creation of the steam engine, dynamo, or the new types of generative AI models. Diffusion, by contrast, is how these innovations spread through the economy and are adopted at scale. It encompasses the messy, gradual process by which organizations learn to use new tools, restructure their operations, and develop the complementary skills needed to realize productivity gains.

This distinction matters because we risk focusing too heavily on winning the innovation race while neglecting the diffusion. As Jeffrey Ding [argues](#) in his [analysis](#) of technological revolutions, countries that excel at adopting and adapting technologies often surpass those that pioneered them. For general-purpose technologies, being first to invent matters less than being best at implementation.

AI [follows](#) this pattern. Despite rapid advances in AI methods, diffusion proceeds glacially. We have also seen this with past general-purpose technologies. For example, factories saw no productivity gains for 40 years after the invention of the dynamo because realizing benefits of electrification required redesigning the entire factory.

To avoid this outcome, the Federal government should prioritize investments in what we call the “complements of automation”—the infrastructure, capabilities, and institutional adaptations that enable productive AI adoption. This includes:

AI literacy and workforce training: Promote comprehensive AI literacy programs and workforce training initiatives across both public and private sectors. This means not just training AI specialists, but ensuring workers in all occupations—from healthcare to manufacturing—can effectively work alongside AI systems.

Digitization and open government data: Accelerate the digitization of government services and expand access to high-quality, open datasets. Previously inaccessible government data can enable AI applications that serve public needs in diverse areas such as urban planning, disaster preparedness, healthcare analytics, or policy development. This is unaddressed by the private sector. Federal investment in data infrastructure, standardization, and secure sharing frameworks would be helpful.

Energy infrastructure improvements: Invest in grid reliability and energy infrastructure to support both AI innovation and diffusion. As AI adoption spreads across the economy, reliable power becomes increasingly critical.

Public sector adoption: Support careful, evidence-based adoption of AI in government services. This means avoiding both extremes: neither rushing to deploy [inadequately tested systems](#) nor moving so slowly that citizens turn to private sector alternatives for basic services. Fund pilot programs, rigorous testing protocols, and best practice sharing across agencies.

Without deliberate attention to diffusion, AI's benefits could concentrate among large technology companies and well-resourced organizations while bypassing small businesses, rural communities, and public institutions. By investing in diffusion infrastructure, the Federal government can ensure that AI serves as a tool for broadly shared benefits.

B. Support open models

Open models can provide significant benefits to society, and the Federal government should focus on research that sustains and expands these benefits. In a parallel context, open source software has proven to be invaluable in designing secure systems and promoting innovation. According to recent estimates, open source software is worth more than \$8 trillion in value and is a part of 96% of commercial software. Openness in AI can provide similar benefits and ensure the competitiveness of U.S. companies today.

One of the biggest concerns with the current trajectory of AI development is the concentration of power in the hands of a few tech companies. Open models can be an antidote to this threat. They allow for smaller companies and individual researchers to develop innovative models for novel circumstances. Indeed, they have already enabled a vast amount of research on AI that could not be done without being able to download and examine the model's internals. Open models also benefit research that uses AI to study other scientific questions, say in chemistry, mathematics, material science, or social science. In particular, open models help accelerate scientific research because they can be less expensive, easier to fine-tune, and more secure. For example, one research group at Princeton has made extensive use of open models to explore the limits of prediction. Open models also have the added advantage of supporting reproducible research. This is in contrast to closed model developers who often deprecate or remove access to their older models, which leads to research based on these models being impossible to reproduce.

Open models can also increase transparency, education, testing, and trust around the use of AI, enabling researchers to audit the system. And, they also lower the barrier for stakeholders outside of large tech companies to shape the future of AI, enabling more AI services to be built by and for diverse communities with different needs that larger companies may not always address.

Open models advance safety and security by accelerating our understanding of AI capabilities, risks, and harms through independent research, collaboration, and knowledge sharing. In turn, this supports regulators and researchers who need the latest methods, tools, and understanding to effectively monitor and test large scale AI systems. Our [research](#) highlights the importance of evaluating the risks of open models not in a vacuum, but in comparison to the risks and benefits from closed models and pre-existing technologies like the internet.

To support research and development of open models, the Federal government must make a substantial and sustained effort to invest in computing infrastructure that supports universities’ research using such models.

C. Research workforce transformation

AI will not simply eliminate or create jobs—it will [change the nature of work](#) across occupations. While popular discourse on the economic impacts of AI focuses on fears of mass unemployment—reminiscent of [past concerns](#) about technologies like the copy machine, radio, and internet—recent research shows more gradual effects. The Federal AI R&D strategy can play a critical role in supporting research that understands how AI transforms the workforce. There are at least four areas for further systematic research:

- *How does AI impact productivity across different occupations?*

Early evidence suggests that AI’s impact on workers depends on characteristics of the occupation. [Felten et al. \(2023\)](#) suggest that highly-educated, highly-paid, white-collar occupations—like legal services—face significant exposure to AI. However, exposure does not mean replacement if workers use AI to complement their tasks. The potential for positive transformation becomes clearer when examining within-job task transformations. In one randomized trial, an AI assistant increased customer service productivity by 14%—mainly by helping support agents handle routine chats faster and in parallel ([Brynjolfsson et al.](#),

[2023](#)). Gains were largest for newer, lower-skilled workers (34%), while experienced agents saw little change. In another randomized control trial, programmers' use of an AI coding assistant increased task completions by 26% ([Cui et al. 2025](#)). However, other studies show minimal productivity gains. Another study in Denmark found chatbots reduced work time by just 3%, with no wage impact ([Humlum & Vestergaard, 2025](#)). More broadly, it appears AI adoption has not significantly boosted overall productivity, in part, due to limited diffusion across industries ([Brynjolfsson & Mitchell, 2025](#)). With mixed evidence from isolated examples, we need extensive, longitudinal research examining how AI impacts vary across different occupations in diverse contexts and when it can lead to positive transformation.

- *How can AI be designed to promote augmentation over automation?*

There is some evidence that AI helps workers augment their capacity to take on more expert tasks. For example, studies of business consultants found that below-average performers improved much more than above-average performers when using AI ([Dell'Acqua et al., 2023](#)). [David Autor \(2024\)](#) argues that AI's near-term potential likely lies in extending the relevance, reach and value of human expertise. These middle-skill roles—such as paralegals and nurse practitioners—depend heavily on procedural knowledge. With AI, these workers could take on more advanced responsibilities, approaching tasks traditionally performed by elite experts such as lawyers or doctors. AI can enable this by combining procedural knowledge with workers' expert judgment to tackle high-stakes tasks. Current usage patterns support this: AI use is highest in mid-to-high-wage roles and mainly augments expertise across specific tasks ([Anthropic, 2025](#); [LinkedIn, 2025](#); [Cheng et al., 2025](#)). However, there is also evidence that the opposite effect—deskilling—can occur ([Macnamara et al., 2025](#)). When workers become overly reliant on AI, they may lose opportunities to develop or maintain their own expertise, potentially leading to a degradation of skills over time. We need systematic research to better understand how to design AI systems

that effectively support this elevation of foundational expertise across occupations and minimize deskilling.

- *What are effective workforce training programs?*

We see a significant skills gap threatening AI adoption and slowing diffusion. By some estimates, AI will be a driver of changes to core job skills in the next 5 years ([World Economic Forum, 2025](#); [LinkedIn, 2025](#)). Yet, fewer than 1% of workers have advanced AI skills, and 60% of companies lack basic AI literacy ([LinkedIn, 2025](#)). This creates dual challenges: employers struggle to fill AI-related roles, while new college graduates face unusually high unemployment. Preliminary research shows that job sector-specific training outperforms general retraining programs, and employer-linked upskilling proves highly effective ([Bürgisser, 2023](#)). But we need more extensive research on developing scalable infrastructure that partners government, industry, and academia for AI literacy and reskilling.

- *How can we develop programs that allow workers to retain agency?*

AI tools deployed without worker input often worsen conditions. For example, gig workers face wage discrimination ([Dubal, 2023](#)) and opaque pay structures, undermining fairness and agency ([Nagaraj Rao et al. 2025](#)). However, efforts to co-design with workers ([Calacci et al. 2025](#); [Nagaraj Rao et al. 2025](#)) show potential. The Writers Guild of America and Screen Actors Guild [secured](#) contract protections ensuring AI augments rather than replaces their work. Similarly, research has found that engaging UNITE HERE hospitality workers in AI innovation addresses adoption challenges ([Spektor et al., 2023](#)). We need research on systematic approaches to collaboration with worker organizations in AI development to ensure AI use reflects worker voices.

D. Other topics for research

In addition to the three core points outlined above, we recommend Federal government prioritize the following research needs:

- *AI for low-resource languages*

AI systems are predominantly trained on high-resource languages, leaving out over 5,000 low-resource languages ([Microsoft, 2025](#)). When AI systems process queries in low-resource languages, outputs tend to be of lower quality ([Asai et al. 2024](#)), more expensive ([Ahia et al. 2023](#)), and less culturally relevant ([Bhutani et al. 2024](#)). Voice capabilities are limited due to lack of local context data and mono-lingual training ([Babu et al. 2022](#)). While community-driven initiatives for workers across [Africa](#), [Southeast Asia](#), [Wales](#), and [India](#) are developing localized datasets and models ([Nekoto et al. 2020](#)), the US—with arguably the best expertise and talent in AI—should lead research on developing AI systems tailored to diverse languages and local contexts. Indeed, at Princeton we [recently released](#) African-centric Lugh-Llama open models, which achieve high performance on challenging benchmarks for African languages. There is a lot more fruitful research to do in this domain.

- *Novel research and evaluation mechanisms for foundation models*

Commercial model developers largely conduct their own red-teaming and testing of their systems. While the AI Safety Institute has an important role to play in developing evaluation standards, independent academic researchers can help expose flaws and suggest improvements that might be ignored by the commercial imperatives of the model developers. One potential solution is to fund the development of a central resource that supports testbeds and research access to these systems. This would help remove barriers to access resulting from uneven distribution of financial and computing resources across research institutions. It

would also help quell concerns about potential legal challenges to researchers who seek to independently evaluate proprietary systems.

Respectfully submitted,

Nathaniel Daw,

*Huo Professor of Theoretical and Computational Neuroscience,
Princeton Neuroscience Institute and Department of Psychology*

Christiane Fellbaum,

*Lecturer with the Rank of Professor, Program in Linguistics and
Department of Computer Science*

Thomas L. Griffiths,

Professor of Psychology and Computer Science, AI Lab Director

Andrés Monroy-Hernández,

Associate Professor of Computer Science

Sayash Kapoor,*

Graduate Student, Department of Computer Science

Steven Kelts,

Lecturer in Public Policy, Integrated Ethics in CS Lead

Mihir Kshirsagar,

Technology Policy Clinic Lead

Sophie Luskin,

Emerging Scholar, CITP

Arvind Narayanan,
Professor of Computer Science, CITP Director

Varun Nagaraj Rao,*
Graduate Student, Department of Computer Science

Matthew Salganik,
Professor of Sociology

Madelyne Xiao,
Graduate Student, Department of Computer Science

* Provided substantial drafting assistance.

Contact:

Websites: <https://citp.princeton.edu>; <https://ai.princeton.edu/ai-lab>