PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mba-3c1a-p52j
Comments Due: May 28, 2025 **Submission Type:** API

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-0311
Comment on FR Doc # 2025-07332

Submitter Information

Name: John Thickstun

General Comment

See attached file(s)

Attachments

Priorities

AI-Priorities

Cornell Bowers CIS
College of Computing and Information Science

Priorities for federal funding in AI

Contributing authors: John Thickstun, Bharath Hariharan, Nate Foster, Thorsten Joachims, Kilian Weinberger

Introduction

Artificial intelligence (AI) is reshaping economic systems, geopolitics, and society—and its transformative influence is set to deepen in the years ahead. The United States' leadership in AI follows a similar blueprint to previous technological revolutions—such as semiconductors and the Internet—where federal investments played a catalytic role. Namely, the U.S. federal government complements the private sector by making strategic investments, in partnership with universities, to support research with high potential but no immediate profit incentive. The federal government can maintain U.S. dominance in AI by strengthening its investments in fundamental AI research at universities that foster AI engineering and research talent. We articulate several areas of research below that we believe are particularly high-leverage areas for supporting U.S. innovation, talent, and leadership.

Fundamental advances in AI algorithms and architectures: Algorithmic and architectural innovation are essential to the development of next generation AI systems. Current AI systems are expensive to train and run, difficult to control, and have limited multi-modal capabilities. Innovations in AI architecture design enable the development of models that scale more efficiently with data and computation, paving the way for more powerful models. Algorithmic innovations will be needed to make AI more controllable, facilitating its deployment across a range of high-stakes applications where its behavior must conform to precise specifications. Future AI systems must also natively accommodate and fuse diverse input types—language, video, audio, time series, etc.—rather than relying on brittle post-hoc solutions that attempt to align these disparate domains. Advances in these fundamental technologies will support the development of stronger, more reliable models with more contextual awareness, enabling the applications of AI for scientific discovery, education, agentic and embodied reasoning, and interactive behaviors that we elaborate on below.

Interdisciplinary research for AI-accelerated fundamental scientific discovery: AI has the potential to dramatically accelerate discovery in science, mathematics, and engineering. However, this potential cannot be realized simply by applying existing AI techniques. Instead, we need to develop new AI techniques tailored for the problem of scientific discovery. In particular, the data hungry nature of modern AI systems is a challenge in scientific domains where annotated data can be limited. Beyond learning from data, scientific discovery also requires AI systems to synthesize existing scientific knowledge, autonomously perform experiments and produce formal mathematical reasoning. Crucially, these challenges must be solved in the context of the scientific domain in question. Thus, interdisciplinary teams of AI researchers and domain scientists are crucial. Such interdisciplinary teams are outside the scope of industrial operations. They must therefore be a focus of federally funded research at universities which host multiple scientific disciplines in close proximity.

Advances in agentic and physically embodied AI: An open challenge for current AI systems is the physical world, especially in cluttered, dynamic, “open world” and potentially adversarial environments. For example, a robot in the kitchen might need to reach into cluttered cabinets with unknown organization while navigating around hot stoves and running kids. Such challenges are also pervasive in medical settings, rescue operations and even on the factory floor. These open research challenges preclude commercialization. They also require fundamental advances that are beyond the near-term focus of industrial research, requiring new multi-sensory perception systems, new algorithms for fine-grained control in robots and novel AI architectures that integrate perception with planning and reasoning.

High-risk, high-reward AI research relevant for domains critical to future U.S. competitiveness, including human-AI interaction: Many transformative applications of AI require systems that can communicate and adapt fluidly in partnership with humans. As AI becomes embedded in decision-making processes across industries, from healthcare to defense to education, interaction with these systems is no longer optional. We need AI that can communicate its goals, adapt to user preferences, respond to uncertainty, and engage in collaborative activities with human partners. Effective interaction underwrites many of the systems discussed elsewhere in this document: from human-robot interaction that ensures safe and effective behaviors in shared physical spaces, to AI education systems that effectively meet the needs of their students.

Educating an AI Research Workforce: For the AI research that will be done in industry, the nation needs a well-educated workforce that is trained to do cutting edge research. US universities are uniquely positioned to supply the talent that is educated in scientific methodology and AI, especially at the PhD level. Any AI research that is done at universities will not only produce new knowledge, but also train the PhD-level workforce that has driven technological leadership of the US tech industry in the past, and that will be required for AI leadership in the future.

Conclusion

Above, we have laid out high-priority areas where federal funding of academic research is vital for the U.S. to maintain leadership in AI. These include directions (such as physically embodied AI as well as fundamental advances in AI algorithms) where technology is not currently viable for commercialization. They also include directions (such as AI for scientific discovery and education) that U.S. universities are particularly well-poised to explore because of their role as multidisciplinary educational institutions.

However, crucially, this list above is not meant to be exhaustive. Many of today's AI breakthroughs are rooted in decades of federal funded research. To sustain this history of innovation and leadership, the U.S. must continue to invest robustly and broadly in AI research. After all, who would have predicted in the 1990s that the niche research area of language modeling would be the breakthrough technology of the 2020s. Ongoing federal support is essential to maintain the nation's global leadership in AI.