PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mb9-xvn0-tenm
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-0282
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Yale University

General Comment

See attached file(s)

Attachments

NSF RFI AI priorities SUBMIT 052925

Docket ID No. NSF-2025-OGC-0001
Submitted by Yale University
May 29, 2025

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.

The United States' global competitiveness depends on our ability to effectively create, harness and utilize AI technologies broadly to drive discovery, innovation, and decisions. We suggest the following strategic approaches to ensure US R&D leadership in AI, with a focus on federal support for areas not likely to be addressed by industry.

Computing Infrastructure

The national AI strategy should prioritize multi-tiered, distributed computing infrastructure and standards for AI research and development. The scale of computing and infrastructure resources needed to store, train, and analyze data with large language models (LLMs) and other AI methods is well beyond the reach of many institutions and organizations. Scaled investments in AI-focused technical computing and data platforms are needed to ensure broad access to the capabilities driving economic innovation and discovery. While the High-Performance Computing (HPC) and mass storage resources at national laboratories such as Los Alamos and Sandia provide potential pathways to collaboration for businesses and universities, the US needs more resources of this type. An expanded network of national HPC laboratories with clear research collaboration pathways would accelerate US advances in artificial intelligence by leveraging the big ideas coming from small and medium sized institutions. Advances in telecommunications infrastructures will also be required to keep pace with increased cloud traffic.

Data Infrastructure

Quality data is the cornerstone of effective AI training and validation, enabling more accurate, reliable, and innovative AI models. Expanded access to reliable federal data sets from agencies like NASA, the CDC, the US Census, US Department of Commerce, the Library of Congress and others will accelerate the pace of innovation and scientific discovery by providing a solid foundation of American-generated data sources for experimentation and development. To advance this aim, we recommend a commitment to comprehensive federal data collection and expanding access to machine-readable federal datasets. Doing so will provision high-quality, trustworthy data and remove existing barriers to progress.

We also recommend that the Federal Government affirm the Library Copyright Alliance's recommendations to institutionalize fair use for AI and curb restrictive licensing terms. By treating the ingestion of content necessary to power AI as fair use under 17 U.S.C. § 107, we safeguard the foundational principle of fair use that drives innovation and research forward. Ensuring that licensing agreements do not undermine fair use rights is critical for fostering a robust and dynamic research environment. Implementing these recommendations will catalyze research innovation and build a strong and sustainable research infrastructure essential for future progress.

Open-Source Foundation Models

The Federal government should prioritize the development of open-source foundation models led by US academic institutions, to ensure national competitiveness and feedback loops with workforce development. These models should be trained on diverse, high-quality datasets and made openly available with full transparency on architecture, training procedures, and intended use. Emphasizing domains underrepresented in commercial development, such as basic scientific research, education, and public interest technologies where societal impact is high but market incentives are insufficient, is crucial. Dedicated funding, scalable computing infrastructure, and centralized governance mechanisms are necessary to ensure responsible development, reproducibility, and alignment with national values.

Technical Standards

Establishing America as the global leader in AI technical standards is a critical priority. By supporting NIST's plan for global engagement and promoting industry adoption of international AI standards, we position the United States as the leader in setting benchmarks for AI technology. Technical standards will drive technological progress, ensure interoperability across systems, and create a cohesive framework for managing risk. Moreover, these standards will enhance collaboration networks and create pathways for ongoing cross-sector education, training, and information dissemination, fostering a globally integrated AI ecosystem.

Workforce Development

The national AI strategy should prioritize infrastructure support for the AI R&D community alongside expanding an AI-savvy technical workforce. As AI hardware and software evolve rapidly, both utilizing these tools effectively and maintaining the underlying infrastructure require specialized expertise that is currently in short supply across academia, government, and industry. Universities can address this workforce gap through expanded federal funding for pilot certificate and degree programs in AI, including options designed for adult learners transitioning careers. Beyond technical training, we need sustainable investments in agile manufacturing facilities that can adapt as chip and GPU technologies advance, moving beyond traditional factory models toward more nimble production capabilities. Expanding this critical workforce is foundational to the stability and security of AI-enabled systems. To support transformation-friendly design methodologies, curriculum updates should also extend to MBA, architecture, and engineering programs, fostering the agile skillsets necessary for rapid technological adaptation.

Mitigation of Environmental Impact

The national AI strategy should prioritize R&D to reduce AI's environmental footprint while meeting growing computational demands. Research priorities include energy-efficient hardware design, optimized algorithms, sustainable facility infrastructure, renewable power systems, and innovative cooling technologies. Academia-industry partnerships would accelerate breakthroughs by combining theoretical research with practical implementation.

Advancing AI capabilities requires not only greener hardware but substantially more power generation overall. New power sources and grid infrastructure should integrate renewable energy from the outset, creating a unique opportunity to build next-generation energy systems optimized for AI workloads rather than retrofitting existing infrastructure. Federal investment in green AI research positions the United States as a global leader in sustainable computing technologies, creating exportable innovations and economic opportunities. Establishing mandatory energy efficiency standards for government AI systems would create market incentives driving private sector adoption across the entire AI ecosystem.

Incremental Funding for AI Research and Education

To enhance AI research and education, it will be necessary to allocate new funding to federal grant-making agencies. This funding should bolster research, infrastructure development, and workforce training on AI. Investments in these areas are essential to nurturing the next generation of AI leaders and innovators. Additionally, federal funds should be earmarked for resources and infrastructure that encourage the public sharing of data and code resulting from federally funded research activities. This openness will catalyze further innovation, ensuring that the benefits of AI development are widely disseminated and accessible.

Healthcare Applications

To lead in the application of AI in healthcare, the United States must invest in the development of a workforce with both clinical and methodological training. Integrating AI into health professions education is essential, facilitating the incorporation of AI and data science into core health curricula. Dedicated federal funding streams should support dual domain training pathways, including MD/PhD, PhD, MD/MS, and MS tracks in AI, biostatistics, and biomedical informatics. This approach will address an existing gap and cultivate clinician-investigators trained to develop and critically evaluate AI systems.

Federal R&D efforts should place emphasis on supporting AI applications in high-impact domains where commercial investment is limited but national needs are urgent. In healthcare, this includes areas such as rural health, public health surveillance, scientific discovery, and rare diseases. These deep domains require sustained investments, domain-specific expertise, and close collaboration with practitioners and communities. The Federal government should create dedicated funding mechanisms and public-private-academic partnerships to foster innovation in these areas, ensuring AI addresses societal challenges and improves quality of life across diverse populations.

Infrastructure is crucial for the development, validation, and equitable deployment of AI tools across diverse healthcare settings. Federal investment should establish regional and national compute hubs accessible to academic and non-academic institutions. Furthermore, it is important to support federated research networks by incentivizing the development of secure, interoperable data ecosystems that facilitate model training and evaluation across institutions without necessitating data centralization. Enhancing health system digital infrastructure by supporting healthcare delivery organizations in upgrading their data infrastructure also plays an essential role in AI deployment.

AI integration into healthcare must be accompanied by mechanisms that ensure its safety, equity, and continued effectiveness over time. Federal research funding should prioritize the development of tools and frameworks for continuous monitoring of AI systems post-deployment, including performance drift and failure detection. Additionally, supporting fairness auditing and transparency tools is vital for routinely assessing differential performance across populations; this enables the detection and mitigation of bias in deployed systems. Promoting regulatory science for clinical AI through collaboration across federal agencies ensures accountability, safety, and transparency in AI-driven care delivery.

Human-Computer Interaction

As AI systems become more autonomous and embedded in critical decision-making, it is essential to reimagine human-computer interaction (HCI) principles and paradigms. The Federal AI R&D strategy should promote research that redefines the roles and boundaries between humans and AI agents, supporting the development of collaborative, controllable, and accountable AI systems. Future research should encompass broader considerations such as transparency, trust, adaptability, human oversight, and safety. Interdisciplinary research integrating cognitive science, human factors, systems design, ethics, and cybersecurity is vital to ensuring AI agents augment human intelligence while remaining aligned with user intent and societal norms.

Basic Sciences

To advance applications of artificial intelligence in the basic sciences, we recommend prioritizing several key areas that bolster both the depth and breadth of AI capabilities. Enhancing support for computational graduate programs will equip the next generation of scientists with robust computational skills, fostering innovation and interdisciplinary research. Additionally, developing AI-powered systems capable of analyzing high-throughput multimodal data will transform the way complex datasets are interpreted, driving progress in areas such as environmental science, neuroscience, and biochemistry. It is equally important to establish programs that increase AI literacy in basic statistics and machine learning, achieved through practical AI training initiatives developed in collaboration with industry and academia. These strategic investments will not only enhance the proficiency of researchers in utilizing AI across the basic sciences but will also ensure a workforce adept at navigating the challenges and opportunities presented by cutting-edge AI technologies.

Applications in Particle Physics and Astrophysics

Advancing research in AI and Machine Learning (ML) in particle physics and astrophysics has the potential to transform our understanding of the universe. Prioritizing these areas could lead to pivotal scientific breakthroughs and technological innovations. Integration of Quantum Machine Learning techniques in high-energy physics experiments holds immense potential for revolutionizing data analysis and pattern recognition. AI and ML play a crucial role in optimizing the processing and interpretation of vast amounts of data generated in these experiments, paving the way for discoveries beyond the current Standard Model of particle physics.

AI and ML are increasingly vital in neutrino and particle physics experiments, aiding in data analysis, event reconstruction, and sophisticated signal/background discrimination techniques in detectors. These technologies enhance image reconstruction, object tracking, and noise reduction capabilities. Leveraging AI and ML tools offers deeper understandings of neutrino properties and particle interactions, essential for uncovering dark matter, antimatter, and other fundamental aspects of the universe.

AI and ML techniques are instrumental in analyzing high-energy particle collisions to study the quark-gluon plasma (QGP) and the conditions of the early universe. Developing and applying advanced AI tools in these domains could provide profound insights into the fundamental forces of nature and the evolution of the universe. Consequently, this provides a deeper understanding of its origins and structure.

In detecting and analyzing astronomical phenomena like supernovae and other high-energy galactic events, AI and ML are pivotal. These technologies enable efficient processing of large volumes of image data from wide-field sky surveys and facilitate analysis of type IA supernovae as standard candles. Accurate measurements of these events are crucial for probing the history of cosmic expansion and understanding dark energy, offering deeper insights into the universe's evolution and fate. Investing in these research areas will drive significant scientific innovation and provide humanity with unprecedented knowledge of the cosmos.

Trust in AI Systems

Benchmarking: Current benchmarks fail to capture real-world complexity, long-term reasoning, and multi-agent interaction. New evaluation methods are needed to assess AI model behavior in dynamic, adversarial, and high-stakes environments. Evaluation is one of the hardest problems in modern AI and significant advances are needed to better evaluate models.

Certified AI: Currently, the technology industry is focused on deploying AI techniques for a wide range of applications. Even though AI algorithms can make non-trivial errors, the overall benefit provided by AI algorithms has led to their deployment with the understanding that human users will take these errors into account and make the appropriate accommodations–in essence such solutions operate under human supervision. However, such an approach is fraught with peril in the context of unsupervised AI algorithms.

A critical area that industry is unlikely to address on their own is the development of AI systems with certifiable guarantees—models whose behavior can be formally verified to meet safety, robustness, and performance criteria. As AI is increasingly deployed in high-stakes domains such as healthcare, infrastructure, and national security, the lack of rigorous guarantees poses risks that private-sector incentives may not fully internalize. The mathematical aspects of machine learning and AI play a foundational role in this effort, enabling the development of frameworks for verifiable learning and understanding worst-case behavior. A shorter-term possibility is investments in certified techniques that can detect when an AI system is producing trustworthy output. Progress in this area requires sustained investment in long-term, high-risk research that bridges ML theory, formal methods, and systems design.

Integrating Data-Driven and Physical Modeling: Development of robust machine learning strategies that fully integrate data-driven learning and physically based models for a variety of classification/ regression tasks--providing advantages of implicit interpretability while addressing difficult computational steps with large amounts of data (as available).

Reasoning about AI Systems: To make AI useful in a reliable sense, one needs to understand what it does and/or how it works. Present AI models are black boxes that train on enormous data sets with huge numbers of free parameters and “learn” something which is unknown to the trainer of the model (and the users as well). While this may be useful for practical applications where some degree of error or "hallucination" is tolerable, it is not really the way forward to build a robust type of AI. Ideally, one would like an AI that is able to "answer" questions about its reasoning and where it may have gone wrong (which is a tall order requiring true intelligence). A less ambitious aim is to make the AI models that are trained more interpretable: namely, the trainer can understand what the AI has "learned" and thus be able to see what it will do well and what it will do poorly. This is an area industry is very unlikely to fund as it is hard, a long-shot, and does not generate an immediate return.

Resource-Efficient AI

Neuroscience-inspired AI: Modern machine learning techniques can trace their origins back to early neuroscientific discoveries such as the threshold neuron model of McCulloch and Pitts from 1943. This area of research has delivered countless dividends to the private sector in the long term, even though this is not widely accepted in the short term. This is exactly the type of investment that requires sustained levels of government support. Significant investments are needed to further understand the computational principles that underpin brain function that can be translated into a combination of software and hardware for efficient AI. Better understanding of neuroscientific principles could lead to intelligent machines that exhibit some of the hallmarks of biological systems: robustness under uncertainty, efficiency, and continual learning to name a few.

Novel algorithms beyond scaling: Blind scaling of AI models is something that will be explored in industry. Academia should explore alternative architectures, modularity, symbolic integration, and continual learning to drive the next wave of breakthroughs.

AI in the Public-Interest: Market incentives overlook low-resource use-cases for AI and public-sector needs. Investment in inclusive infrastructure will preserve cultural heritage and expand access to AI technologies.

Other Topics

Multi-Agent Learning: In recent years, Deep Learning and the broader AI technology have delivered significant advances in learning challenges across various data modalities including speech, images, and language. On the algorithmic front, much of that progress has been fueled by the success of gradient descent-based optimization methods in computing good solutions to non-convex optimization problems. Going forward, many outstanding challenges in Deep Learning lie at its interface with Game Theory. These challenges include playing complex board games at super-human level, robustifying Deep Neural Net-based classifiers against adversarial attacks, training deep generative models, and multi-robot interactions.

More generally, as AI technology is getting broadly deployed, it is becoming increasingly important to train AI agents that can reason about their interaction with other AI agents and/or humans. On the multi-agent learning front, however, the Deep Learning paradigm has been less successful. Here, the role of single-objective optimization is played by equilibrium computation, but Gradient-descent based methods struggle to converge, let alone to good solutions. Even worse, due to non-convexities, standard game-theoretic equilibrium concepts either fail to exist and/or are intractable. Key research challenges include How does one train DNN- based agents to be strategic? And what is the goal of their training?

Large Language Models (LLMs) and Natural Language Processing: Large Language Models (LLMs) and Natural Language Processing (NLP) are among the most critical areas within AI. Work is needed in the core science and technologies surrounding language models. This includes: 1) studying mechanisms of generalization in LLMs to understand to what extent they memorize knowledge versus their abilities to reason about and generalize to unseen scenarios. 2) Developing improved algorithms for LLM training and inference to make them more efficient and reliable. 3) Developing open source and transparent LLM technologies. 4) Efficient extension of LLM capabilities to novel scenarios and domains, including applications in the scientific domain to improve scientific workflows.

Computer systems: Progress in computer systems is one of the key foundations of modern AI. However, astronomical growth in AI model complexity has resulted in unsustainable compute costs for AI. This leads to two natural questions: (i) how do we co-design new hardware and software systems for more energy-efficient and sustainable AI? and (ii) are there better ways to build Intelligent systems by drawing deeper inspiration from nature?

Social Robotics: Social and embodied Artificial Intelligence (AI) is a crucial foundational topic in AI. One of the main goals of this research is to advance robot autonomy in order to enable new robotics applications that help people in positive ways.

Intelligent Autonomy: The next frontier of AI systems will traverse the physical world to support humanity in solving real-world problems in the wild. The biggest challenge facing the development of such AI in the physical world is the need for internet-scale data to effectively train the AI systems. Today, engineers are required to collect, curate, and siphon large amounts of data to produce what we currently see in AI–an unsustainable effort. A key area of research should explore mathematical principles and algorithms for deploying reliable intelligent robotic and AI systems in the wild. This would involve algorithmic reliability and formal guarantees of physically embodied AI systems through the development of optimal control and learning theory. These algorithms would enable AI systems to explore, become curious and play through interacting with their environment to collect data to learn dexterous and agile skills.