Received: May 28, 2025
Tracking No. mb8-6ykc-n9qr
Comments Due: May 28, 2025 **Submission Type:** Web

PUBLIC SUBMISSION

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-0140
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Brown University

General Comment

Please find attached Brown University's comment.

Attachments

Brown's Response to RFI _ Development of 2025 National AI Research and Development Strategic Plan - Google Docs

Response to the Request for Information on the Development of a 2025 National AI Research and Development Strategic Plan

Submitted by: Greg Hirth, Vice President for Research, Brown University

---

Thank you for the opportunity to provide feedback on the development of a National AI Research and Development Strategic Plan. Given the exciting progress in Artificial Intelligence and the opportunities for profound impacts on our nation, a revised strategic plan is timely and important.

AI is a topic that has arisen from fundamental, curiosity-driven academic research. Many of the recent industry breakthroughs that have brought AI into the mainstream have been driven by a world class workforce of researchers and engineers trained in federally-supported academic labs and universities. The continued health of the field and its future impact require that these investments remain a national priority.

We have identified the following areas we believe should be highlighted in the National AI R&D Strategic Plan. For simplicity, we align these areas with the categories used in the previous strategic plan, but we believe new strategy labels would be appropriate.

Make Long-Term Investments in Fundamental and Responsible AI Research

Science of Deep Learning: Deep learning (DL) underlies the most powerful AI technology that the world has ever seen, and yet we have very little understanding of *why* it works. We have a reproducible recipe, but no generalizable theories or principles. We need investment in both theoretical and empirical work that seeks to illuminate the basic mathematical, scientific, and cognitive principles that underlie AI's success, so that we can build systems more reliably, efficiently, and deliberately. As is consistently the case, this type of foundational, conceptual work is best performed in the open scientific literature and is therefore likely to be too high-risk and long-term for industry to carry.

Holistic Evaluation of AI: LLMs are currently compared using a variety of standardized and easy-to-run benchmarks. Scores on these benchmarks do not necessarily equate with downstream usefulness, but basic research on how to better compare models is a long-term and open-ended endeavor and thus not prioritized by industry (who have a vested interest in only using evaluations that place their models on top). We need fundamental work on evaluating and comparing models, including when to assign significant labels such as "intelligence" to systems. To fully assess the applicability of models, evaluation methodologies should include elements from research communities such as human-computer interaction (HCI), cognitive science and psychology, education and economics, as such a holistic evaluation is likely to result in better measures of how people use and respond to AI systems and how they are effectively integrated into real-world workflows.

Open Source AI tools: The current state-of-the-art AI systems require large amounts of data and expensive specialized equipment to produce, which makes it impossible for the majority of experts to contribute to its design and development. As a result, we are leaving the majority of America's potential for innovation and creativity untapped. Open source tools will accelerate the country's ability to address major open challenges such as reducing the time and cost of developing models, and making them more reliable and secure so they can be applied to high-stakes problems such as national defense. Industry has little incentive to invest in open source infrastructure, so federal funding is the best avenue to achieve these goals.

Streamlined funding process: To maximize the effectiveness of research dollars invested in AI, refresh the Federal Grant funding model. We estimate that the average faculty member spends 10 hours per week on the task of proposal writing. Similarly, each proposal is massively time-intensive to the government from the reviewing perspective. Meanwhile, many more scientific proposals provide innovative contributions than our current federal funding mechanisms can support. There is thus a tremendous opportunity to increase efficiency in federal grantmaking. One option would be to, in addition to ranking grant proposals based on their quality, assess whether each proposal is above or below a threshold for innovation. Then, the awards could be distributed at random or perhaps round-robin to proposals that exceed the threshold. The result would be a more even distribution of federal research funding across institutions and states and less effort in creating a complete ranking for all submitted proposals. This approach, sometimes called a "partial lottery", has a number of benefits and has been enacted with some success. It is worth experimenting with this model at a larger scale.

Develop Effective Methods for Human-AI Collaboration

Instructable interfaces: The academic research community is well positioned to be involved with designing, testing, and developing prototypes of AI-enabled interfaces that allow end users to tailor the behavior of applications using instructions, demonstrations, and/or incentives, as needed. Providing such a facility empowers members of society to make use of computing to accomplish tasks that are most personally valuable to them. To maximize the economic benefits of such systems, the research community should develop standards that allow similar interfaces to be applied across a wide variety of systems. It is unlikely the industry alone will engage in broad interoperability efforts because it is not in their direct benefit to do so. But pre-competitive research to identify the most useful interface approaches will benefit both individuals and companies by avoiding the need for people to learn a new interface for each system they interact with.

Explainable AI: There is no obvious incentive for industry to increase the transparency of their systems, especially since one suspects that the commercial LLMs are in fact hybrid systems with many rule-based responses. On the other hand, explainability remains a key impedance to the real-world applicability of AI. There seems to be a growing effort in the open source community to make explainable AI more usable and to some degree this relates to the prior topic. For example, many researchers use Shapley values as an effort towards explainability, but these are neither immediately intuitive nor completely reliable. Further R&D is needed in this area.

Ensure the Safety and Security of AI Systems

Secure and people-centered AI systems: Secure, efficient, and privacy-preserving software infrastructure for AI workloads can be created by placing people---developers, verification engineers, end users, and policymakers----at the center of system design. AI infrastructure should empower developers and engineers to extract significantly more useful results from the same hardware in environmentally sustainable ways. It should also protect individuals' data and institutional assets by embedding security, vulnerability detection, and regulatory compliance into the design and lifecycle of AI systems. Pre-competitive research is needed to establish standardized practices and toolchains that help identify and repair software vulnerabilities, detect breaches in multi-tenant AI environments, and prevent costly data leaks. Such research will ensure that AI systems can be trusted by the public and responsibly adopted at scale across sectors.

Develop Shared Public Datasets and Environments for AI Training and Testing

Public infrastructure and democratization for AI: To continue American dominance in AI, we must continue to grow local expertise. Current AI systems are expensive, difficult to access, and for the most part behind corporate walls. The government must build on the success of the National AI Research Resource (itself a product of the National AI Initiative Act of 2020) and invest in public infrastructure for building, experimenting with, and innovating new AI applications.

AI testing at universities: To advance the implementation of AI beyond the laboratory and niche sectors of society, it is important that AI-based technologies be reliable such that the response of these technologies to a given set of input parameters be well-understood and intended. To facilitate the development of reliable AI, universities could partner with efforts by national laboratories such as NIST's Generative AI evaluation program and others, such as the Coalition for Health AI (CHAI), to develop complementary strategies for 3rd party testing of AI-based technologies. For example, research universities are well-positioned to address fundamental research questions pertaining to the ground-truth in model testing such as *how can laboratory-based models be developed to better capture the dynamic nature of human interaction with AI tools?* Another example is with red-teaming efforts to test AI tools for their limitations. *How can red-teaming exercises be developed for meaningful-duration experiments that better simulate the routine use by the target demographic for which the tools are developed?* These questions suggest another unique role for universities as AI testing centers. *How can universities develop such Centers to perform at-scale, realistic AI testing that involves all stakeholders including the technology developers and representatives from the end-user communities?*

Coordinated data collection for AI for science and engineering: Scientific experiments are often prohibitively expensive either because of their size and material costs or because of the need for many experiments to search for the correct discovery space. Nonetheless, most current experimental work is done by many independent experimentalists in a duplicative or random fashion guided by taste. More efficient would be using active learning like frameworks---en masse---to guide which experiments would yield the most information. Achieving this goal would require opening up and sharing data on vast scales in large repositories and the use of AI models to highlight which portions of discovery space have been less traversed or may yield new insights. While researchers still need to retain independence, the point is that much money is wasted on unintegrated efforts guided by preference rather than likelihood of discovery. Example: It is well-known that chemists have not been able to access large parts of chemical space, instead running experiments in the same organic space for quite some time. What if AI could help direct where in organic or other spaces different chemistries were most likely to be found?

Governance and Cyberinfrastructure for Data and AI: The development of AI models depends on access to good quality data. At a foundational level, data governance frameworks need to be in place to ensure secure sharing, access, and use of data for AI research and education. Robust cyberinfrastructure, including high-performance computing environments, secure enclaves, and open-source software, also needs to be available to support development of standardized AI models and applications.

> Example: Challenges and considerations for health AI: [Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril]() (National Academy of Medicine, 2022). This topic would naturally dovetail with Congressional efforts to create AI for science infrastructure (https://www.heinrich.senate.gov/asap).

Measure and Evaluate AI Systems through Standards and Benchmarks

Cooperative Group Trials: The government should create the ecosystem needed to conduct robust prospective trials of AI interventions. This would be in the same mold as the cancer cooperative groups that drove much of the innovation in cancer treatment from the 1970s-1990s. Now, industry dominates the space, but that wouldn't have happened without the best practices and rigors of the cooperative groups, which continue to have a role especially in working with already approved therapeutics. Most AI currently in development for healthcare applications is being deployed in clinical settings without adequate evaluation, and this risks repeating mistakes of the past and harming patients. Rigorous prospective randomized controlled trials (RCTs) of AI interventions would greatly increase their acceptance and prove their worth, but are expensive and complex to carry out. Relatedly, cooperative groups have advanced trial methodology from the simple (such as block stratification) to the more complex (new statistical techniques for non-proportional hazards, for example). The European Commission is supporting something like this: https://digital-strategy.ec.europa.eu/en/policies/european-ai-research .

Better Understand the National AI R&D Workforce Needs

AI Education: With the wealth of big data and advancements in technology has come the rapid growth of AI across disciplines. There is a need to build the AI workforce in academia, industry, and government through transdisciplinary training and education programs. These programs should balance theoretical concepts with practical experience, fostering a generation of researchers equipped to not only understand the foundations of AI but also to apply these principles to real-world challenges. Curricula should be tailored to the needs of different roles (developer, implementer, user, etc.) and the specific competencies they require.

Example: Artificial Intelligence for Health Professions Educators