PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mba-5w4w-818f **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-0321
Comment on FR Doc # 2025-07332

Submitter Information

Organization: MIT Computer Science and Artificial Intelligence Laboratory

General Comment

See attached file(s)

Attachments

2025-05-28-CSAIL-response-final

MIT CSAIL response to OSTP Request for Information on the Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan

Executive Summary Artificial Intelligence (AI) is rapidly reshaping global technological leadership, economic competitiveness, and national security. U.S. leadership in AI—like earlier innovations such as the internet, cloud computing, and foundational science—has been built on sustained federal investment in university research and an open innovation ecosystem. To extend that leadership, the United States must ensure it has the best and most advanced AI. The MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) offers seven key recommendations to guide the 2025 National AI R&D Roadmap and Strategic Plan:

1. **Invest in Basic Research Toward Artificial Superintelligence (ASI):** To lead in the next era of AI, the U.S. must go beyond current deep learning and transformer architectures and invest in fundamentally new methods and models capable of providing assurance and combine perception, reasoning, common-sense understanding, and physical interaction. This includes agentic AI, physical AI, and hybrid neuro-symbolic approaches that enable adaptive, explainable, energy-efficient systems. These models will form the backbone of future capabilities in national security, healthcare, infrastructure, science, and more.
2. **Advance AI for Scientific Discovery:** AI is poised to become an active partner in science—not just analyzing data, but generating hypotheses, designing experiments, and uncovering new knowledge. The U.S. should invest in "AI Co-Scientists" and synthetic science infrastructure to accelerate breakthroughs in energy, medicine, and beyond. Supporting AI in healthcare, personalized medicine, and diagnostics will also be critical to future health resilience and innovation.
3. **Ensure AI Ready Data**: AI software is similar to other software with one big exception: *data* continuously writes AI software. AI data readiness is often 80% of the effort in an AI undertaking. In the USA, vast collections of digital logs are produced, but very little of this is AI ready data. Pivoting American Science to producing AI ready data by default will accelerate AI innovation.
4. **Understand the theoretical foundations of intelligence:** We suggest that significant progress can be made in a theoretical understanding of intelligence and that this progress will lead to new architectures for ASI. As AI systems begin to rival human capabilities in language and reasoning, we face an urgent need for a theory that explains how minds learn and generalize. Foundational principles of intelligence that may underlie both biological and artificial intelligences may emerge from mathematics, computer science and neuroscience. The interplay between experiments and theory in the development of nuclear physics offers a useful analogy for this pursuit.
5. **Invest in AI for Strengthening the Nation:** AI should be developed to directly address the operational and strategic needs of the U.S. government and public. This includes AI for national security and defense, cybersecurity, public health, logistics, and infrastructure. Targeted

investments are needed to support secure-by-design systems, real-time decision-making, and cross-agency coordination. AI can also improve the efficiency and responsiveness of government operations through predictive analytics and intelligent automation.

6. **Invest in AI Education, Training, Skilling, and Reskilling:** The AI-driven future will demand a rethinking of education and workforce development. The U.S. must build strong pipelines from K–12 through higher education and develop scalable reskilling programs for workers across sectors. Investments in foundational math, AI literacy, adaptive learning platforms, and role-specific training will ensure a broad, inclusive, and agile AI-capable workforce. Education must also address the ethical, societal, and governance dimensions of AI to build public trust and informed leadership.
7. **Establish a National Frontier Lab (NFL):** To catalyze AI progress across mission-critical areas, the U.S. should launch the **National Frontier Lab**—a new model for public-private-academic partnerships. The lab would focus on a specific domain (e.g., security, healthcare, scientific discovery), combining university research expertise with industry resources, compute infrastructure, and real-world users. The Frontier Lab would support next-generation AI, address assurance and alignment, and provide government agencies with direct pathways for research translation. This initiative would ensure that U.S. AI leadership remains mission-driven, grounded in scientific excellence, and responsive to national priorities.

Overall Goal: Accelerating USA Leadership What if cloud computing, search engines, the World Wide Web, or the Internet had been invented in China? How about the transistor, the Black Scholes derivatives pricing model, the medical MRI or thousands of other innovations we depend on in our daily lives and for global leadership. All of these innovations, and so many more, are the product, either directly or indirectly, of the academic science and engineering ecosystem based at US universities and funded by Federal government research dollars since the US victory in World War Two. They are the material basis for American prosperity and global leadership. Now imagine what could happen the next generation of LLM-like tools are invented and commercialized in China instead of the United States? Would US tech companies like Microsoft have to negotiate licenses with Chinese firms under the thumb of the Chinese government? How about the thousands of US firms and tens of millions of US students, teachers and small businesses? The question of US global technical leadership goes far beyond just this generation of AI.

American innovation has been world-leading not only because of our scientific and commercial prowess, but also because of the open society and rule of law that is at the heart of the United States. These innovations have benefited from contributions from scientists who came to the United States from all around the world. Great research ideas led to robust products and services thanks to the open marketplace guaranteed by the US legal system, protecting property, assuring fair treatment of all

citizens, and protecting each individual's right to free speech under the Constitution. These elements are vital to our nation's preeminent leadership in innovation as they are the basis of both science and commerce. So, if we want the next generation of new technology to develop in a way that supports human flourishing, it is essential that we strengthen our university-to-marketplace innovation ecosystem in the free, rule-of-law based society of the US and our democratic allies. China has shown it can bring strong technical skills to the table, but the US can retail vital global leadership by developing new AI technologies that work best in our open, rule-of-law driven marketplace.

Roadmap for Continued AI Leadership The United States currently leads the world in AI development and applications, but other countries such as China are vying for leadership. The US will need to invest in foundational research to progress and maintain its leadership. We provided a response to the OSTP Request for Information on the Development of an Artificial Intelligence (AI) Action Plan (FR Doc. 2025-02305), AI action plan recommendations from the MIT Computer Science and Artificial Intelligence Lab (April 10, 2025)). We build on this response for our recommendations for the USA AI Roadmap here.

Recommendation 1: Invest in Basic Research Toward Artificial Super Intelligence

A core long-term aim of US AI research should be to lead in artificial superintelligence (ASI). Our current model architectures are capable of doing pattern recognition and language processing well, but they lack general understanding, common-sense reasoning, and a clear view of comprehensive world models. So, a key research challenge before us is to integrate real-world awareness, common-sense reasoning, agentic computation, and physical intelligence. American AI leadership should champion the development of AI frameworks that combine the pattern-recognition power of neural networks with formal reasoning and symbolic representation. These architectures are key to enabling adaptation, abstraction, causal inference, compositional reasoning, and more transparent decision-making. Moving beyond correlations to true semantic understanding will be essential for AI systems to operate safely and effectively in high-stakes domains like medicine, national security, science, and law. This requires new research investments in:

- **Assurance:** Research is needed to develop technical tools, including safety metrics and explainability techniques, that can enable the use of AI while adhering to legal requirements and supporting public confidence in the reliability of AI systems.
- **Resilience:** A key research priority is building technical tools and resilience metrics to ensure the consistent performance of models across diverse use cases and in challenging conditions.
- **Generalization**: AI systems need to be able to generalize effectively from limited training data and mimic human capabilities where we can learn from fewer examples and apply our knowledge in novel situations.

- **Energy optimization**: Research should continue to address how we can make our AI systems more cost effective from an energy perspective. This involves understanding the power supply and consumption needs to accommodate AI growth, developing less-energy-intensive processes, and building more efficient AI models.

To maintain global leadership in artificial intelligence, the United States must not only advance current models, but also invest in the bold, exploratory directions that will define the next era of AI. Today's deep learning systems, while powerful, are limited in generalization, reasoning, efficiency, and adaptability. To push past these boundaries, we need targeted, sustained investment in high-risk, high-reward AI research that explores fundamentally new architectures, capabilities, and learning paradigms. The opportunities are in:

- **Physical AI**: The U.S. should also be making bold investments in new AI paradigms that move well beyond today's dominant approaches. This includes significantly advancing agentic AI systems—models that not only generate responses but can perceive, decide, act, and adapt autonomously over time. These systems must learn from direct interaction with the physical world, developing what we call Physical AI: intelligence that is the result of integration of digital data and reasoning with physical embodiment. Physical AI enables autonomy and automation in dynamic, unstructured environments—reasoning *about* the world, *within* it. These capabilities are foundational for the future of human-machine collaboration and will be essential in sectors defined by physical complexity and real-world constraints, including national security, manufacturing, infrastructure, logistics, energy, and telecommunications. Investing in this new class of AI will unlock new applications and will secure American leadership in the technologies that support economic competitiveness and national resilience.
- **Neuro-Symbolic AI**: Today's deep learning systems are powerful pattern recognizers but struggle with abstraction, reasoning, and compositionality. To overcome these limitations, the future of AI will involve neuro-symbolic architectures, which are hybrid systems that combine the statistical strengths of neural networks with the structured, logic-based reasoning of symbolic AI. These models can be designed to explicitly represent causal relationships, rules, and concepts, enabling more interpretable and trustworthy decision-making. This line of work is critical for domains where explanations, accountability, and guarantees are essential—such as legal reasoning, financial auditing, and safe autonomous systems. Modular reasoning systems will also enable plug-and-play components for vision, language, planning, and control, supporting more flexible and scalable AI ecosystems. By bridging learning and logic, neuro-symbolic AI offers a path to intelligent systems that can see, predict, and understand.
- **Adaptive AI**: A next frontier in AI lies in systems that are adaptive capable of learning after training based on the inputs they see, modifying their own learning processes, architectures, or

objectives over time. Unlike today's models, which are trained once and frozen for deployment, self-evolving AI systems can grow in capability, adjust to new environments, and respond to long-term shifts in context. These systems will develop meta-learning abilities that enable them to autonomously optimize their own performance across tasks and domains. Inspired by biological evolution and developmental learning, this paradigm opens the door to lifelong AI, where agents evolve new capabilities through experience rather than retraining. In dynamic sectors like defense, disaster response, and autonomous exploration, self-evolving systems could dramatically extend operational lifespans and reduce the need for constant human reprogramming—delivering agility and resilience far beyond current approaches.

- **AI-Quantum hybrids**: AI and quantum computing are two of the most transformative technologies of our time—and their convergence offers profound new possibilities. Quantum-AI hybrids could tackle computational problems that are currently intractable for even the largest classical models, such as simulating molecular dynamics, solving combinatorial optimization at scale, or enabling new forms of probabilistic inference. Quantum-enhanced machine learning could reduce training times, compress model representations, and solve high-dimensional problems with radically greater efficiency. On the flip side, AI can accelerate the design and calibration of quantum algorithms and error correction protocols. Together, these two domains create a mutually amplifying frontier, with the potential to reinvent industries ranging from pharmaceuticals and materials to cryptography and logistics. Strategic investment now could place the U.S. at the forefront of this emerging space—shaping the future of AI hardware and unlocking computational paradigms beyond Moore's Law.

These forward-looking AI solutions are critical for technical progress and for strategic leadership. The U.S. must invest now to shape the trajectory of AI toward systems that are general, adaptive, trustworthy, and deeply integrated with the physical and social world. That means funding ambitious research, supporting open testbeds, and building new coalitions across academia, industry, and government that are willing to challenge today's assumptions and invent what comes next.

Recommendation 2: Invest in AI for Scientific Discovery

Scientific discovery has always proceeded at the pace at which humans can generate hypotheses, design experiments, and interpret complex, high-dimensional data. As challenges in fields like climate, energy, and medicine grow more urgent—and datasets grow exponentially—we need new tools to expand the boundaries of human insight. AI now stands at the threshold of transforming how science is done, not by automating routine tasks, but by becoming an active partner in the scientific process. To unlock this potential, the U.S. must invest in the infrastructure, talent, and foundational research needed to develop AI systems that do more than analyze: they discover. These systems will be

essential for accelerating progress and for opening up entire new domains of inquiry that are currently out of reach.

- **AI Co-Scientists**: The rise of foundation models has paved the way for a new class of AI systems to enable scientific discovery: AI agents capable of forming hypotheses, designing experiments, and autonomously extracting insights from complex data. This goes far beyond data analytics—these systems will simulate, reason, and generate new knowledge. Early success stories like AlphaFold have already demonstrated the potential of AI to leapfrog traditional scientific discovery timelines. In the future, AI models could autonomously explore new materials, discover fundamental laws, or propose novel therapies. These systems would operate not only as tools for scientists, but as collaborative engines of discovery, augmenting human creativity with machine-scale exploration. Investing in synthetic science infrastructure will accelerate breakthroughs in energy, medicine, climate modeling, and beyond, changing both how we do science.
- **AI for healthcare**: research should support AI applications for personalized medicine, diagnostics, drug discovery, and optimizing patient care. This would allow the US to move beyond generalized approaches to disease treatment to a much more personalized approach centered on an individual.

Recommendation 3: Ensure AI Ready Data

To systematically transition scientific data collection toward AI readiness, agencies should establish standardized data pipelines that incorporate preprocessing, normalization, and metadata annotation at the point of collection rather than as post-processing steps. This requires implementing automated data validation schemas that ensure consistent formatting, completeness checks, and semantic labeling using controlled vocabularies and ontologies relevant to each scientific domain. Technical infrastructure should include ETL (Extract, Transform, Load) processes that automatically generate machine-readable datasets with standardized APIs, proper versioning, and documented lineage tracking. Additionally, agencies must adopt common data models and interchange formats (such as JSON-LD, Parquet, or domain-specific standards like NetCDF for climate data) while embedding quality metrics and uncertainty quantification directly into datasets. This approach necessitates training scientific personnel in data engineering principles, establishing cross-agency data governance committees to maintain consistency standards, and creating federated data catalogs with AI/ML-specific metadata tags that enable automated discovery and integration across government scientific repositories. This technical approach addresses the core challenge of making the "80% effort" of data preparation a built-in feature rather than a downstream bottleneck in AI development projects.

Recommendation 4: Advance the Theoretical Foundations of Intelligence

To ensure long-term U.S. leadership in artificial intelligence—and to responsibly guide the development of Artificial Superintelligence (ASI)—we must invest in developing a rigorous theoretical framework for intelligence. As AI systems begin to rival human performance in language, reasoning, perception, and planning, the absence of a formal theory leaves us without a principled understanding of how these systems learn, generalize, adapt, or fail. We propose a research initiative that treats intelligence as a fundamental scientific question—one that cuts across computer science, mathematics, neuroscience, and cognitive science—and seeks to identify the core computational and informational principles that underlie both natural and artificial minds, and further extending into tackling the intimate yet mysterious connection between consciousness and cognition and then cognition and computation.

This effort must move beyond empirical benchmarking and architecture scaling. While data-driven and heuristic models have driven recent breakthroughs, they lack interpretability, reliability, and formal guarantees. A theoretical framework could unify currently fragmented insights—such as representation learning, causal inference, abstraction, compositionality, and goal-directed behavior—into a cohesive understanding of intelligent computation. Potential directions include mathematical formalizations of learning, dynamical systems and state-space theories of cognition, theoretical neuroscience and brain-inspired computation. This effort could take inspiration from the development of nuclear physics—a field where tight interplay between theory and experiment yielded deep understanding and transformative technologies. Similarly, progress in AI will increasingly depend on theory-informed design of architectures, training regimes, and objective functions. The goal to build systems we understand, whose capabilities, limitations, and alignment with human values can be precisely reasoned about. A national initiative focused on the science of intelligence would position the United States to shape the conceptual foundations of ASI. It will create the intellectual infrastructure needed for safety, alignment, and innovation—ensuring that the trajectory of AI development is guided by insight, not just scale.

Recommendation 5: Invest in AI for Strengthening the Nation

AI research should also address targeted application areas such as national security and healthcare that enhance decision-making ability and improve operational efficiency. In national security, research can focus on AI applications for advanced threat protection, cyber defense, secure communications, logistical planning, and improving the robustness and resilience in complex environments.

In an era defined by geopolitical uncertainty, accelerating technological change, and increasingly complex threats, AI is emerging as a cornerstone of national security and national strength. From cyber defense to strategic forecasting and supply chain resilience, the nation that leads in trustworthy, mission-ready AI will define the terms of global leadership. AI has the potential to strengthen the entire national fabric—from protecting critical infrastructure and enabling faster disaster response, to

improving public health, modernizing transportation, and supporting energy transition. To realize this vision, we must invest in AI systems that are capable of functioning in adversarial conditions, coordinating across agencies, and supporting complex decision-making in real time.

This effort demands a national commitment to developing next-generation AI technologies that are built not only to compete, but to protect, empower, and unify. That means designing AI to operate safely in high-stakes domains, integrating seamlessly with human teams, and grounding decisions in transparency and explainability. It also means investing in secure AI infrastructure, advancing AI assurance frameworks, and building the workforce needed to develop, deploy, and govern these technologies. This way, we can ensure that AI serves as a force for national resilience, economic vitality, and democratic strength**.**

- **AI for Security and Cybersecurity**: AI can enhance national defense and resilience by detecting cyber threats in real time, anticipating attack vectors, and autonomously responding to system breaches. Next-generation models can support secure-by-design architectures, monitor critical infrastructure for anomalies, and help protect both civilian and military networks from evolving adversarial threats. Investing in AI for cybersecurity is essential for maintaining trust in digital systems, protecting sensitive data, and defending against both state and non-state actors in an increasingly contested information space.
- **AI for Organizational Efficiency**: In an era with a specific focus on government efficiency, new AI applications can support improved government operations. AI research could support predictive analytics for public services, automate administrative tasks, and generally enhance the effectiveness and responsiveness of government functions.

Recommendation 6: Invest in AI Education, Training, Skilling, and Reskilling

Clearly the advancement of AI models will shift the nature of work. We will need deeper research to understand and prepare for this transformation. For example, research is needed to map the coming growth in AI capabilities to predictions about the skills and tasks that will be automated. At the same time as some skills are devalued by AI, others will become more valuable because they are complementary to new AI capabilities. As new technologies hit the labor market, “creative destruction” occurs; new jobs are created and old ones are made obsolete. The economy is used to such churn, but it is not costless. Workers must retrain, and social institutions must exist to give them the time and security to do so.

The widespread integration of AI into the economy and society will require new strategic research on workforce development including comprehensive training, retraining, and re-skilling initiatives. New educational programs spanning primary, secondary, and tertiary levels must move beyond traditional

computer science instruction to teach AI literacy and effective AI use. Research can focus on curriculum design that equips individuals with the skills that they need to interact with, interpret, and leverage AI tools in both professional and social settings. We need to understand and identify the most promising pedagogical approaches, and AI itself can be used to tailor learning for different age groups and learning styles. Additionally, a significant portion of the workforce will require reskilling to use AI technologies effectively, similarly to the digital literacy reskilling introduced by early Internet adoption. Reskilling will be especially critical as AI systems transform industries such as manufacturing, healthcare, logistics, finance, and public services. We need research to understand the most effective methods for this large-scale reskilling. Research can include a focus on adaptive learning platforms, modular training programs and credential systems that can deliver and evaluate AI skills at scale.

To ensure long-term national leadership in AI, the United States must make strategic, sustained investments in AI education and workforce development at all levels. This includes expanding AI literacy in K–12 education, integrating AI fundamentals into undergraduate and graduate curricula across disciplines, and creating robust programs and pathways for reskilling and upskilling workers across sectors of the economy. The future AI workforce is not limited to PhDs in computer science. We need diverse talent pipelines that prepare engineers, policymakers, health professionals, educators, and skilled tradespeople to develop, use, deploy, and govern AI technologies responsibly. This means building educational programs tailored to different roles—whether you're designing models, using them in applied settings, or making strategic decisions about their impact.

A strong foundation in mathematics—particularly in areas such as linear algebra, probability, and statistics—is essential to understanding and innovating in AI. Investing in math education early and broadly will expand access to AI careers and ensure that more students are equipped to engage meaningfully with AI technologies in every field. Finally, education must also focus on AI ethics, safety, and societal impact, so that the next generation of AI leaders understands not only what AI can do, but what it should do. Building an AI-literate society—from classrooms to boardrooms—is essential to sustaining public trust, promoting innovation, and preserving democratic values in the age of intelligent systems.

Recommendation 7: Invest in a National Frontier Lab

To secure long-term leadership in artificial intelligence and ensure its transformative power serves the public good, the United States should establish a National Frontier Lab program—a dedicated, mission-driven research initiative that sits at the intersection of AI and the nation’s greatest challenges. This lab would bring together the strengths of academia, government, and industry to tackle problems no single sector can solve alone—from energy efficiency and healthcare to national security,

infrastructure, and scientific discovery. Modeled on successful public-private research ecosystems, the Frontier Lab would enable deep collaboration across disciplines, provide shared compute and data resources, and support high-risk, high-reward research with long-term impact. CSAIL's own experience with industry-funded research partnerships has demonstrated the value of private-sector perspective in guiding and participating in research at the transition from basic science to commercialization. Critically, it would serve as a national platform for developing trustworthy, robust, and application-ready AI systems that reflect American values, advance strategic priorities, and accelerate innovation across society.

The National Frontier Lab will be a new research funding model that can bring leading academic research expertise together with industry partners and user organizations in order to design and train a new generation of foundation models. Industry has a well-established path for training language models, but reaching full Artificial Super Intelligence will require new research to train models with better real-world awareness, genuine reasoning capabilities, as well as more comprehensive techniques for assurance and alignment. These research endeavors require well-organized access to compute, data and industrial and user partnerships. The National Frontier Lab (NFL) we propose is a new structure for leveraging government research funds, combined with private sector financial support and in-kind contributions of expertise and data. The Lab would be targeted at making progress on specific, mission-critical application areas such as national security (both cybersecurity and cyber physical systems), healthcare and government operations efficiency. The NFL would bring together a unique combination of academic research expertise, industrial marketplace experience and user organizations in order to frame real-world challenges for researchers to solve.

The National Frontier Lab can be created by a call for one or more existing science funding agencies. Each Lab would be focused on a defined research challenge, and constituted by a competition in which academic labs, industries and AI-user organizations would bid for government seed funds to be matched with industry resources, compute infrastructure, identified data repositories and other partner organizations. The government funding agency would award seed research funds to support university student and faculty researchers to work in collaboration with the rest of the NFL team. A key role for government leadership is to assure that research is focused on the highest priority challenges necessary to advance US global AI leadership. The NFL could also play a role in the commercialization of the most effective ideas.

NOTEs
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.