PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mb9-t7sf-a1ye
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-0251
Comment on FR Doc # 2025-07332

Submitter Information

Organization: University of Texas at San Antonio

General Comment

Please see attached RFI for the Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan. Thank you.

Attachments

RFI_AI

University of Texas at San Antonio
Response to the
Request for Information on the
2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan

May 23, 2025

On behalf of UTSA, we welcome the opportunity to respond to the National Science Foundation's (NSF) and Office of Science and Technology Policy's (OSTP) request for information (RFI) regarding the "2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan".

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. The content of the document can be reused by the government in developing the AI R&D strategic plan and associated documents without attribution.

1 Introduction

The University of Texas at San Antonio (UTSA) has made a strategic commitment to advancing artificial intelligence (AI) through the establishment of the MATRIX AI Consortium for Human Well-Being, launched in July 2020. This initiative represents a bold move into transdisciplinary AI research, bringing together 87 research scientists with multifaceted expertise across academia and industry.

The MATRIX Consortium serves as a model for successful cross-sector partnerships, fostering collaborations between academia, industry, and government. Its primary aim is to harness AI in ways that promote human well-being, addressing societal challenges through deep technological solutions.

By leveraging the expertise across various domains—engineering, health sciences, social sciences, and computing—UTSA's initiative demonstrates how a multi-disciplinary approach to AI can lead to innovative breakthroughs and scalable real-world applications.

UTSA is also launching the College of Artificial Intelligence, Cyber, and Computing, set to open in Fall 2025. This new academic unit places UTSA at the cutting edge of AI education and research, unifying key disciplines—computer science, computer engineering, statistics, and information systems—into a cohesive, interdisciplinary curriculum. By serving as a nexus for these fields, the college is poised to accelerate innovation across industries, driving the development of smarter, more secure, and efficient systems. This integrated approach will prepare students and researchers to tackle complex, real-world challenges where AI, data, and security intersect.

2 Opportunities for R&D Investment

We anticipate that this input will inform the development of the 2025 National Artificial Intelligence R&D Strategic Plan to reflect the administration's commitment to foster an environment where AI drives economic growth, technological advancement, human flourishing and national security.

2.1 Overcoming the Energy Barrier of AI

Neuromorphic Computing. Training state-of-the-art AI models requires significant energy and compute resources. This level of energy consumption is increasingly unsustainable as the scale of foundational models continues to grow. The challenge becomes even more acute when these models are deployed on resource-constrained platforms such as edge devices, where compute, memory, and power budgets are severely limited. To address this challenge, the nation must invest in the design and development of new AI hardware systems that can have a transformative impact on energy resources. These technologies include neuromorphic computing, near-memory computing, high-density 3D architectures, and emerging CMOS technologies. Examples at UTSA include neuromorphic computing hardware accelerators operating under ultra low power, infrastructure to emulate AI models on neuromorphic platforms such as the neuromorphic commons for at-scale computing, and research on the use of emerging technologies such as memristors, memcapacitors as a platform to develop low energy and computationally efficient neuromorphic devices.

Energy Security. Energy security is critical for national security especially in support of the expansion of AI where energy supply chain security and efficiency must increase dramatically. Cybersecurity Manufacturing Innovation Institute (CyManII) at the University of Texas at San Antonio (UTSA) is developing cyber secure defensible architectures, cyber-physical passports, and other CyManII innovations to be deployed directly into devices such as sensors, actuators, and controllers, and to measure the impact of cybersecurity using formal methods. CyManII will enable cyber secure data management that will lead to a) energy efficiency scaling that directly translates into increased productivity with a fixed, or lower, power budget; and b) a 50% reduction in the cost of domestic manufacturing. CyManII will also build data libraries for industry sectors to enhance energy efficiency, and also focus on developing and deploying digital twins and AI methods to model and optimize these energy efficiencies across automated processes, production lines, and entire supply chain networks. In aggregate, these CyManII innovations will support the creation of a verifiable supply chain ecosystem and support the expansion of the use of AI in national security.

2.2 Advances in AI Systems that are Dynamic, Adaptive, and Learn Continually

One of the central challenges in current AI research is enabling systems to assimilate new information while preserving previously learned knowledge, a capability known as continual or never-ending learning. While this ability is intrinsic to humans and other mammals, replicating it in artificial systems remains a significant hurdle.

Cross-species studies of cortical brain connectivity highlight several key mechanisms critical to robust lifelong learning: local and global plasticity (including neurogenesis, neuromodulation, and homeostasis), hierarchical cell-assembly architectures that encode individualized perception and memory, and large-scale population dynamics that give rise to multi-level abstractions across diverse temporal and contextual scales. These neurobiological principles suggest hybrid representational approaches that can enhance self-regulation, adaptability, and robustness in AI networks.

Recent research indicates that neuro-inspired AI systems can harness their own accumulated metadata to assess the relevance and utility of prior knowledge in novel domains. This capability represents a departure from traditional teaming and static model architectures, opening pathways toward more autonomous, context-aware systems, and positioning a competitive advantage for the nation.

At UTSA, several ongoing initiatives exemplify this interdisciplinary approach:

- Neurogenesis-inspired AI models that simulate the integration of new neural elements to support adaptive learning.
- Sleep-inspired algorithms that emulate memory consolidation processes for continual learning.
- Neocortex-based anomaly detection algorithms tailored for complex, dynamic environments.
- Temporal learning architectures that capture time-sensitive patterns in data streams.
- Insect-brain-inspired meta-learning frameworks that enhance problem-solving through lightweight yet effective learning strategies.

These projects bring together AI engineers, neuroscientists, psychologists, and mathematicians to explore the frontiers of cognition and computation. The collaborative efforts are not only pushing the boundaries of individual disciplines but are also forging new interdisciplinary domains with the potential to transform AI systems into more resilient, generalizable, and efficient learners.

2.3 Resilient and Trustworthy AI Systems

To ensure broad societal benefit and sustainable deployment, the 2025 National AI R&D Strategic Plan must prioritize trustworthy AI. Future systems must be interpretable, aligned with human values, robust to uncertainty and imperfect data, and secure against adversarial threats. We recommend federal investment in:

- Research and infrastructure for measurable fairness, explainability, and trustworthiness, including benchmarks, audit tools, and standards;
- Techniques to test, interpret, and ensure the safety and reliability of AI models in high-stakes domains such as healthcare, defense, and infrastructure;
- AI systems that remain dependable when trained on limited, imbalanced, incomplete, corrupted, or continually varying data and tasks;
- Methods for reliable inference in high-dimensional settings under constraints on amounts of data, compute, memory, or latency;
- Privacy-preserving, secure architectures enabling cross-institutional learning without sharing raw data (essential in healthcare and defense, among other sectors);
- AI methods resilient to adversarial attacks, such as data poisoning and model inversion;
- AI systems aligned with human intent and values and existing decision-support tools.

These challenges are often unmet by the private sector, which relies on massive compute and curated data. Such models perform poorly in resource-limited, data-scarce, or adversarial environments typical of public deployments. At UTSA, we develop computationally efficient AI that remains reliable under imperfect data. Our focus is on interpretable, generalizable systems that do not depend on large-scale infrastructures.

2.4 AI for Human Well-Being

Beyond the ability of AI to automate processes, augmenting human capabilities would open the door for AI to impact human well-being more broadly. However, augmentation demands new skillsets: AI systems today have a limited ability to perform complex cognitive tasks like designing, planning and adaptation to unknown events. They also lack the agility, dexterity and regenerative capacity intrinsic to the human body, and rarely adapt rapidly to subtle social and environmental cues. Achieving these key characteristics for AI-driven well-being can improve the quality of human life at all stages and levels (individual, family, and societal). Example applications include real-time decisions on the oxygen infusion rate for a newborn in the ICU, augmented breathing, or designing a smart-health home whose floors notify family members of an elderly relative's well-being. Integrating diverse data-streams from passive and active sensors, real-time social and biomedical feedback, and biologically-inspired algorithms will position AI systems to augment human capabilities.

Secure Implantable Biomedical Devices for Health Care. Trustworthy AI is essential for implantable biomedical devices—such as neurostimulators, cardiac pacemakers, and insulin pumps—that interact directly with critical physiological systems and increasingly rely on autonomous decision-making. Without strong safeguards, these devices are vulnerable to cyber threats and algorithmic failures that could compromise patient safety, starting from device manufacturing to its operational phase.

Secure AI for Time-Critical Emergencies. Trustworthy AI is equally essential for securing AI for point-of-care medical devices, sensors at home and in the field, and resilient emergency medical systems across U.S. cities. UTSA together with UT Health Sciences and UT Tyler has been leading a Texas-wide initiative to revolutionize AI for emergency care. The iRemedyACT project includes the development of AI real-time assistants that can guide clinical decision-making in rural, underserved areas and/or in military deployment. These AI decision-support systems are poised to be combined with real-time passive sensors of both individuals' health and of local environments (e.g., traffic, weather, resources). Securing this ecosystem of AI devices - and ensuring its robustness - is of utmost importance to optimal medical care and national security.

National Center for Secure and Trustworthy Biomedical Devices. To address these urgent needs for secure AI for Biomedicine, we propose the development of a National Center for Secure and Trustworthy Biomedical Devices, focused on advancing AI reliability, cybersecurity, and regulatory science to ensure secure manufacturing, and safe, ethical, and effective integration of intelligent medical implants, on-site sensors and AI assistants into clinical care.

A foundational model for this initiative is the CyManII institute at UTSA, which leads national efforts in securing U.S. manufacturing through applied research, workforce training, and public-private collaboration. Furthermore, CyManII's cybersecurity technologies become instrumental for biomedical devices. For example, the cyber-physical passport (CPP) framework generates a digital "passport" for every manufactured part and verifies its integrity and security, leveraging formal verification and trustworthy AI. Furthermore, CyManII's Compositional Attack and Defense Annex (CADA) offers the ability to conduct comprehensive safety and security analyses without the need for physical assessments that could potentially risk damaging devices. CADA's models capture the system's architecture and behavior accurately supported by generative AI models, allowing for the simulation of various operational scenarios and the identification of potential vulnerabilities or safety risks.

Resilience and Repair - Neuromorphic Systems and the Human Nervous System Both the human nervous system and neuromorphic systems exhibit forms of resilience, yet their capacities for repair differ significantly. The human nervous system relies on plasticity to support learning and response to external stresses. Recovery after a brain injury can be remarkable, particularly during early development, and in brain regions like the cerebellum, where neurogenesis occurs throughout life. However, long-range axons represent a key point of vulnerability. Damage from stroke or spinal cord injury often leads to permanent deficits, because of the inability to repair or regrow axons. Neuromorphic AI systems, while not capable of biological regeneration, can be designed with redundancy, fault tolerance, and adaptive reconfiguration. We propose interdisciplinary working groups of AI researchers, neuroscientists, and clinicians to identify shared vulnerabilities and develop strategies to improve resilience and recovery in both biological and artificial systems.

2.5 Strategic Partnerships

Transdisciplinary partnerships that advance foundational research. AI research teams often operate within disciplinary silos, resulting in a lack of convergent methodologies that hinder progress toward foundational breakthroughs. This lack of convergence has created persistent knowledge gaps—particularly in areas such as generalizable intelligence, context-aware systems, and human-centered AI. To promote research advancements that are transformative, sustainable, transdisciplinary collaborations and coordinated engagement with a broad community of stakeholders is critical. To drive transformative and sustainable research advancements, it is essential to cultivate transdisciplinary collaborations that bridge disciplines such as computer science, neuroscience, behavioral science, engineering, and ethics. These collaborations must be complemented by coordinated engagement with a broad spectrum of stakeholders, including academic researchers, industry practitioners, policymakers, and the public. The proposed National Center for Secure and Trustworthy Implantable Devices would cover this broad spectrum of stakeholders. Such collaborative ecosystems are critical to shaping the future of AI in ways that are innovative and robust. At UTSA, the MATRIX AI Consortium exemplifies this approach. It serves as a dynamic model for transdisciplinary research, where 87 experts from multiple domains work together to address complex challenges and pioneer novel AI paradigms.

Workforce training through public-private partnerships. The development of a system to generate a ready-to-hire workforce in a field that is changing rapidly is a fundamental challenge for the US to maintain AI dominance. As such, it is necessary to develop workforce training across all the AI research enterprise. We propose that this can be achieved by developing widely adopted examples, platforms, and protocols in public-private partnerships at all levels of education. An example of this approach is taking place at UTSA through the CyManII. Here, projects and platforms are used by industry and researchers to develop and to train the workforce. By having clear shareable resources that can be adapted to teaching, training, and research will allow for the development of a ready-to-hire workforce. Further, programs across the new UTSA College of AI, Cyber, and Computing offer Senior Design, a year-long capstone course in which our computing majors engage directly with industry *clients* to gather requirements and develop real solutions. The UTSA School of Data Science offers a summer version of this initiative, engaging teams of students across disciplines to solve real-world data problems for local non-profits.