PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mb9-q3eh-ouwi
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-0224
Comment on FR Doc # 2025-07332

Submitter Information

Government Agency Type: State
Government Agency: Virginia Polytechnic Institute and State University

General Comment

See attached file(s)

Attachments

AI RD Strategic Plan_VT Response

RESEARCH AND INNOVATION
VIRGINIA TECH

Burruss Hall, Suite 340
800 Drillfield Drive
Blacksburg, Virginia 24061
p: (540) 231-6077 f: (540) 231-4384
research.vt.edu

Virginia Tech Response to Request for Information on the Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan

Docket ID No. NSF-2025-OGC-0001
May 29, 2025

The rapid diffusion of AI across every sector of society demands a national research agenda that couples technical excellence with human-centered values, resilience, and security. Drawing on expertise that spans multiple Virginia Tech colleges and departments, we identify research domains in which sustained, mission-driven federal investment is both essential and uniquely catalytic.

The themes and recommendations below were sourced from across our university with a view to identifying (i) high-impact opportunities unlikely to be funded by short-term commercial incentives yet promise transformative societal impact, and (ii) areas that require advances in foundational research, multi-disciplinary collaboration, and/or long evaluation horizons—conditions under which academia excels and where sustained, mission-driven federal funding has the greatest capacity to drive innovation and direct progress towards national interests.

As a land-grant institution dedicated to service to society, Virginia Tech has a long tradition of use-inspired research and a deep appreciation for the harmony between translational innovation and fundamental discovery. Anchored in that dual mission, researchers across all nine of the university's colleges are actively engaged in AI research, building on expertise in artificial intelligence, data science, systems engineering, neuroscience, human factors, robotics, and immersive visualization to accelerate human-technology partnerships toward seamless, effective, ethical augmentation. In particular, the university's Sanghani Center for Artificial Intelligence and Data Analytics is a hub for collaborative frontier research by nationally celebrated faculty members addressing critical challenges in network analysis, explainable AI, adversarial AI, natural language processing, and other domains that underpin the transformative capacity of AI and are foundational to both economic competitiveness and national security.

That research expertise and interdisciplinary fluency informed the selection of the research themes outlined below.

Theme 1 – Physically-Grounded, Resource-Efficient AI for Cyber-Physical & Edge Systems
Goal: Create safe, explainable AI that operates under real-world data and compute constraints, and is grounded in the environmental contexts in which they operate.

- **Physically-grounded algorithms**: The recent history of AI research has focused primarily on using large and larger datasets to improve performance. In the real world, however, data is scarce and continuous reliance on massive volumes of data is unsustainable. Moreover, data-intensive frameworks like large language models are often not grounded in realistic operational

environments. We need a new generation of AI algorithms and frameworks that can operate with minimal training and that can understand the intuitive physics underlying the environment in which they operate. These capabilities are prerequisites for these systems to handle unforeseen cases and unfamiliar domains. Without such physical grounding, we cannot ensure the safety and security of AI when deployed in the real world.

- **Real-world AI systems under resource constraints***:* Many engineering systems, like wireless communication systems, the Internet of Things, and the power grid, require AI frameworks that can operate under stringent compute and communication resources; devices in such systems cannot run large-scale models like today's large language models. Global leadership in AI development demands systems that can be deployed over compute-constrained devices and achieve high accuracy with minimal compute, energy, and memory resources.
- **AI-powered digital twins***:* Digital twins are digital representations of real-world autonomous systems like vehicles, factories, or robots that can facilitate control and optimization and therefore offer the potential to enhance productivity across multiple commercial and industrial domains. Off-the-shelf AI frameworks, which often rely on blackbox tools like neural networks without explainability of performance guarantees, are not suitable for constructing effective digital twins. There is a need for twin-centric AI frameworks that have inherent explainability and are amenable to mathematical analysis and performance guarantees.
- **AI-Hardware Co-Design for Edge Intelligence:** Many high-value applications for AI, including robotics, wearable health systems, and environmental monitoring, require specific and unique hardware specifications, and there is a critical need for scalable, explainable AI models that are jointly optimized with their deployment platforms. New research must bridge AI algorithm design and hardware implementation, especially for edge devices where compute, power, and communication bandwidth are highly constrained.
- **National Edge-AI Testbeds:** National testbeds coupling embedded sensors, wireless networks, and low-power accelerators with open-source model stacks — similar to Virginia Tech's O-RAN infrastructure — would be invaluable resources to enable and coordinate research on all the topics described above.

Theme 2 – Trustworthy & Secure AI Ecosystems
Goal: Protect critical infrastructure and AI systems against evolving cyber threats.

- **AI for cybersecurity**: To safeguard national security, there is a need for AI solutions to improve the security of systems, including intrusion detection and fraud detection through advanced behavioral analytics, AI for zero-trust systems, generative biometric models for authentication systems, and large language models for malware detection and prevention of social engineering attacks.
- **Cybersecurity for AI**: We are now also seeing cyberattacks specifically targeting AI systems, such as data poisoning attacks, backdoor attacks, adversarial attacks on cyber-physical systems, and jailbreak attacks on large language models. Some of these attacks present substantial risks, including the release of highly sensitive data and, in the case of cyber-physical systems, risks to human life. It is crucial that we develop strategies to prevent and mitigate against such attacks.
- **Distributed Sensing for Critical Response:** AI+Cybersecurity research should be contextualized in distributed intelligent sensing networks supporting emergency response and national infrastructure. These systems should operate under tight latency, energy, and reliability constraints, requiring new approaches for co-designing AI with embedded hardware and physical sensors. Successful state-level initiatives — for example, Virginia's Commonwealth Cyber Initiative — can be scaled nationally to leverage existing assets for rapid innovation.

Theme 3 – AI for Health, Biology & Epidemiology
Goal: Democratize data access and causal reasoning to advance precision health and epidemiological modeling.

- **Amalgamating high-quality, de-identified diverse health datasets for AI/ML Approaches**: No single health care organization (HCO) captures the totality of natural history and human disease experience necessary to build generalizable models. If AI is going to truly impact human health, sharable datasets of sufficient size and quality are needed across the age and disease spectra. Computational approaches for automating such data across a range of HCOs is needed to scale generalizable AI models for a range of health-focused approaches.
- **Scaling causal reasoning for healthcare approaches:** Predictive models have found utility in the healthcare space for making disease diagnosis, ascertaining variables of interest for classification tasks and for patient phenotyping. The impact of real-time causal frameworks in healthcare could drive improvements in patient safety, biomarker discovery and precision care. However, despite the benefits towards personalized and precision health, causal reasoning is largely not a primary focus for healthcare AI. Research investment is needed to overcome barriers to leveraging causal analysis in healthcare, including integrating complex models into health information environments, compute resources needed, data sparsity and noise inherent to electronic health record data and insufficient EHR functionality.
- **Rare Diseases & Public Health AI:** Federated learning and transfer-learning frameworks could leverage sparse clinical cohorts to accelerate rare-disease immunology research, real-time epidemiological forecasting, and multi-omic biomarker discovery—coupling AI models with wet-lab validation pipelines and global surveillance networks.

Theme 4 – AI-Enabled Built Environment & Construction Automation
Goal: Leverage generative AI and robotics to build adaptive, resource-efficient infrastructure.

- **Construction Automation & Generative Design:** Construction automation research that integrates robotics, generative AI, and sustainable materials will accelerate the development of adaptive, accessible, and resource-efficient building systems, especially in public infrastructure.
- **Shared Fabrication Testbeds:** Mid-scale fabrication labs and pilot sites would allow researchers, contractors, and regulators to prototype, stress-test, and certify AI-guided construction workflows under realistic safety and code-compliance constraints.
- **AI-Driven Circularity & Disassembly:** AI-assisted design-for-disassembly and circular material flows can enhance resilience and reduce dependence on offshore or otherwise fragile supply chains.
- **Enhancing Built-Environment AI**: Dynamic and responsive building elements powered by AI can facilitate spatial awareness and indoor navigation for individuals across the spectrum of physical capability. Ensure that focus groups, user testing, and advisory panels involved in building design reflect the real user base, including individuals with disabilities.

Theme 5 – Human-Centered & Inclusive AI for Creative, Accessible Societies
Goal: Integrate the arts, accessibility, and lived experience into AI research and education.

- **Interdisciplinary AI-Design Testbeds:** Interdisciplinary AI + design testbeds can enable architects, educators, computer scientists, and artists to co-develop ethical applications with real-world constraints where AI augments human expression rather than replacing professional or student agency.

- **Spatial-Computing Tools for Accessibility:** Multimodal AI-powered spatial-computing tools can improve indoor navigation, hazard recognition, and design literacy for visually -impaired people, merging computer vision, voice interfaces, and the built environment.
- **Arts-Integrated AI Education & Workforce:** Arts-integrated AI curricula, research centers, and workforce programs will cultivate AI-literate artists and embed creative perspectives within technical R&D — reinforcing the principle that human use creates the context for the meaning, purpose, and impact of AI.
- **Creative Talent Pipelines:** Incentivize K-16 pipelines, internships, and residency programs positioning artists and designers as full partners in AI innovation.

Theme 6 – National Workforce & R&D Ecosystem: Standards, Testbeds, and Partnerships
Goal: Deploy federal levers that amplify cross-sector collaboration and talent pipelines.

- **Catalytic Pre-Competitive Funding:** Provide catalytic R&D funding for high-risk, pre-competitive AI science that industry alone will not undertake.
- **Public-private Consortia:** Convene public-private consortia to publish agile safety benchmarks, red-team protocols, accessibility guidelines, and arts-integrative best practices.
- **National AI Clouds & Data Trusts:** Operate secure National AI Research Clouds and domain-specific data trusts (health, CPS telemetry, built-environment scans).
- **Grand-Challenge Hubs & Innovation Networks:** Launch grand-challenge programs and regional AI innovation hubs anchored by land-grant universities to unite academia, industry, government, and community stakeholders.