

PUBLIC SUBMISSION

Received: May 29, 2025 Tracking No. nb9-lkj4-j531 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-0200
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Technuf, LLC

General Comment

On behalf of Technuf LLC, a Maryland based company providing leading-edge and proven technologies and an SBA certified Small Disadvantaged Business (SDB) holding STARS III contract vehicle, engaged in artificial intelligence (AI) research, cybersecurity, and federal innovation support, we thank the Office of Science and Technology Policy (OSTP) and the National Coordination Office (NCO) for inviting public comment on the 2023 National Artificial Intelligence Research and Development Strategic Plan. We support the plan's overall direction and welcome the opportunity to enhance its impact by offering targeted feedback and recommendations grounded in both our applied experience and current academic and policy research. The attached file has the detailed comments.

Attachments

Technuf_NSF-2025-OGC-0001



Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan

National Science Foundation

Comments to RFI#: NSF-2025-OGC-0001

Response Deadline: May 29, 2025 - 11:59 PM EST

Submitted to:

Contact Person:

Faisal D'Souza

ostp-ai-rd-sp-rfi@nitrd.gov

Submitted by:

Company Name: Technuf LLC

Name: Faisal Quader, PhD

Website: <https://www.technuf.com>

GSA Contracts:

GSA MAS – GS-35F-017DA

STARS III – 47QTCB22D0337



This document is approved for public dissemination. This 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.

Table of Contents

Response To Request for Information on the Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan	1
1 Elevate Deep Learning Transparency as a Strategic Priority	1
2 Align Federated Learning with Edge AI to Enhance Privacy	1
3 Standardize Synthetic Data to Support Digital Twin Reliability	1
4 Address Responsible AI in Education as a National Priority	2
5 Institutionalize AI Ethics Review Boards	2
6 Institutionalize “AI for Social Good” Tracks in AI Conferences and Professional Societies	2
7 Expand Benchmarks to Include Dataset Bias, Sustainability, and Trust	3
8 Bridge Research and Deployment via Small Business Innovation	3
9 Update Open Knowledge Network Strategy with Recent Progress	4
10 Public-Private Partnerships (Strategy 8)	4
Concluding Remarks	4

RESPONSE TO REQUEST FOR INFORMATION ON THE DEVELOPMENT OF A 2025 NATIONAL ARTIFICIAL INTELLIGENCE (AI) RESEARCH AND DEVELOPMENT (R&D) STRATEGIC PLAN

On behalf of Technuf LLC, a Maryland based company providing leading-edge and proven technologies and an SBA certified Small Disadvantaged Business (SDB) holding STARS III contract vehicle, engaged in artificial intelligence (AI) research, cybersecurity, and federal innovation support, we thank the Office of Science and Technology Policy (OSTP) and the National Coordination Office (NCO) for inviting public comment on the 2023 National Artificial Intelligence Research and Development Strategic Plan. We support the plan's overall direction and welcome the opportunity to enhance its impact by offering targeted feedback and recommendations grounded in both our applied experience and current academic and policy research.

1 Elevate Deep Learning Transparency as a Strategic Priority

We recommend the addition of a sub-strategy within Strategy 1 that directly targets transparency, interpretability, and theoretical understanding of deep learning systems. Deep learning is a dominant paradigm across academia and industry, yet recent trends signal a concern over its opacity and lack of theoretical clarity. This black-box nature of AI models requires extensive federal research. Furthermore, its omnipresence at the present time cannot be denied because of its large integration in the private sector and effectiveness.

We recommend OSTP to support research that explains how deep models construct representations, identify uncertainty, and fail under domain shifts. Such work should include auditing generative models, improving explainability tools, and documenting generalization limits. A dedicated strategy on deep learning transparency would reflect its ubiquity and practical significance across scientific and commercial domains.

2 Align Federated Learning with Edge AI to Enhance Privacy

While federated learning is well-positioned in the 2023 plan, we recommend explicitly integrating edge computing and edge devices R&D into the same strategic track. Federated learning alone does not address the constraints of on-device learning, bandwidth, or intermittent connectivity, all of which are central to real-time, privacy-first applications in healthcare, defense, and smart infrastructure.

We suggest that the updated plan includes a joint roadmap for federated and edge AI, with federal support for energy-efficient edge inference, secure on-device learning, and hybrid deployment models. Further edge computing has modernized federated learning. Together, these approaches can fulfill the broader goals of privacy, latency reduction, and data sovereignty.

3 Standardize Synthetic Data to Support Digital Twin Reliability

We support the plan's inclusion of digital twin systems under Strategy 1 and recommend placing stronger emphasis on the role of "synthetic data" in ensuring their success. Many digital twin deployments, ranging from aerospace simulations to urban planning, rely on high-quality synthetic data to replicate real-world dynamics. However, the lack of standardized methods to generate, validate, and benchmark synthetic data undermines confidence in their outputs.

We recommend:

- Establishing federal standards for synthetic data quality and representativeness,

- Funding research into synthetic data bias and domain fidelity, and
- Creating shared repositories that include provenance metadata and domain alignment metrics.

Digital twin success depends on trustworthy simulations, which in turn depend on realistic, reproducible synthetic data inputs. Hence, due to the nature of its critical real-world implications and usefulness, further research and funding efforts would go a long way in emerging technologies.

4 Address Responsible AI in Education as a National Priority

The plan currently omits one of the most urgent and visible societal domains affected by AI: education. We propose the addition of a subsection focused on Ethical and Responsible AI in Education. The rise of generative AI tools like ChatGPT has introduced new ethical tensions between students and instructors. One example from Northeastern University involved a student complaint after discovering a professor used ChatGPT to generate course materials, despite institutional guidelines prohibiting students from doing the same (as reported in Newsweek and Forbes, among many outlets in 2025). This underscores a growing inconsistency in how AI is used and monitored within educational settings.

Recent academic literature [1][2][3] confirms these tensions, highlighting institutional blind spots, unclear policies, and student skepticism. We recommend that the federal government provide support for:

- Development of fair-use frameworks for AI in teaching and assessment,
- Research on the social and psychological effects of AI in academic settings, and
- Collaboration with small businesses that practice transparent, and audit-friendly AI.

These actions will help create effective academic environments and support ethical integration of AI into classrooms.

5 Institutionalize AI Ethics Review Boards

To align with the plan's call for responsible R&D, we strongly support the creation of domain-specific Ethics and Scientific Integrity Review Boards for federally funded AI projects. Much like Institutional Review Boards (IRBs) in biomedical research, these bodies would review project proposals for bias, downstream societal impact, and misuse risks.

We recommend piloting such boards within NSF, NIH, and DHS-funded AI research and requiring funded researchers to submit structured social impact statements and mitigation plans. Review templates and audit protocols should be shared across agencies and made adaptable for academic labs, nonprofits, and small businesses.

6 Institutionalize “AI for Social Good” Tracks in AI Conferences and Professional Societies

While top-tier AI conferences (e.g., International Conference on Computer Vision (ICCV), and Conference on Computer Vision and Pattern Recognition (CVPR)) now include “AI for Social Good” or “AI for Societal Advancement” tracks, many others do not. We recommend mandating such tracks in all national and international AI conferences hosted in the U.S. or supported by federal agencies. These tracks serve as critical platforms for advancing research that addresses sustainability, accessibility, ethics, and community-led AI design, domains often underrepresented in mainstream technical programs.

Given that major industry stakeholders (e.g., NVIDIA, Meta, Google), alongside academics, routinely participate in these conferences to showcase state-of-the-art research, requiring dedicated tracks on societal impact would meaningfully influence community norms, elevate ethical considerations, and reinforce alignment with national priorities.

To institutionalize this practice, we also recommend that federal guidance formally engage leading professional societies, such as Institute of Electrical and Electronics Engineers (IEEE), Association for Computing Machinery (ACM), and Association for the Advancement of Artificial Intelligence (AAAI), to embed “AI for Social Good” and “AI for Societal Advancement” as required components of their sponsored conferences. Doing so will not only encourage consistent inclusion of socially impactful work but also signal a cultural shift in how AI research is evaluated and disseminated across both public and private sectors.

7 Expand Benchmarks to Include Dataset Bias, Sustainability, and Trust

The plan correctly identifies the importance of evaluation frameworks, but current benchmarks focus too narrowly on accuracy. We recommend expanding federal support for benchmarks that capture bias and human trust dimensions. Recent work has shown that ImageNet and other canonical datasets contain social and semantic biases that influence model behavior which used to be a future-work in the 2023 plan, but now we have results [4]. These limitations persist even in state-of-the-art models and must be addressed through bias-aware evaluation standards. Additionally, benchmarking trust, especially in decision-support applications, should include user-centered evaluations such as perceived objectivity, interpretability, and reliability.

8 Bridge Research and Deployment via Small Business Innovation

The gap between academic AI research and real-world deployment remains a persistent problem. Universities often produce cutting-edge models but lack the infrastructure or incentives to turn those ideas into usable tools. Small businesses are ideally positioned to fill this gap, but they require targeted support, particularly during the Technology Readiness Level (TRL) 4–7 range when prototypes mature into deployable systems.

We recommend that the federal government:

- Creates streamlined transition programs from academic AI research into SBIR more assertively,
- Promotes co-development and licensing opportunities through public innovation portals, and
- Highlights successful transitions in reports and outreach, and lower administrative barriers to funding for small and minority-owned AI firms.

This would turn publicly funded research into public benefit, while stimulating local innovation ecosystems. This is evidenced by the research and innovation at Technuf LLC’s Innovation Center of Excellence (ICE) where we have successfully productized academic research, that are traditionally housed only in the theoretical circles (e.g., CVPR, ICCV, ACM etc.), to solutions for everyday Americans. This includes face recognition-based student management systems including crucial school bus applications powered by AI, deepfake detection through generative AI, automated ticketing system via voice synthesis, only to name a few.

9 Update Open Knowledge Network Strategy with Recent Progress

The plan's reference to Open Knowledge Network (OKN) development should be updated to reflect recent NSF initiatives. Since 2022, the Open Knowledge Network Roadmap [5] has outlined major goals, and two major follow-ups, namely, the Proto-OKN investment [6] and the “Regional Hubs Prototype Programme” [7], have launched to prototype data interoperability and semantic infrastructure. These developments demonstrate federal progress and reveal ongoing gaps in coordination, metadata governance, and interagency application.

We recommend that the plan updates:

- Acknowledges these programs and frame the OKN as a backbone for explainable AI,
- Clarifies how federal R&D programs will contribute to and benefit from open, interoperable data ecosystems, and
- Regularly assesses and communicates the implementation status and lessons learned from the OKN rollout.

10 Public-Private Partnerships (Strategy 8)

Federal agencies should have increased flexibility to bring in Industry partners to showcase their AI capabilities, with specific focus on how they will accelerate solutioning to known challenges. To enable this, industry partners should be required to understand existing IT and Data infrastructure along with an understanding of Agency specific problems, as these are in the public domain, e.g., Annual Financial Report. Additionally, increased collaboration among federal agencies should be required to tackle common challenges such as payment integrity. Advancing these efforts requires more than the CFO, CIO Councils; it requires strategic and tactical team formation with measurable goals.

Concluding Remarks

In summary, the 2023 National Artificial Intelligence Research and Development Strategic Plan lays an essential foundation for guiding the future of U.S. leadership in AI. We appreciate the opportunity to contribute to its refinement and expansion. Our recommendations reflect both on-the-ground experience in responsible AI deployment and emerging insights from current research and federal activity. By strengthening commitments to transparency, ethical governance, educational accountability, equitable workforce development, and research-to-deployment translation, particularly through small business support, the U.S. can secure its position as the global leader in trustworthy and socially beneficial AI. We look forward to continued collaboration with OSTP, NITRD, NSF and partner agencies to ensure the Strategic Plan evolves to meet the urgent needs and opportunities of the next five years.