PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mb9-jc36-e8k5
Comments Due: May 28, 2025 **Submission Type:** API

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-0191
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Hugging Face

General Comment

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

Attachments

2025_HF_AI_Research_and_Development_Strategic_Plan_RFI

HUGGING FACE

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

Submitted to: Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO), National Science Foundation
May 2025

About Hugging Face

Hugging Face is a community-driven U.S. company that democratizes responsible artificial intelligence (AI) and machine learning (ML). Our platform is the most widely used for sharing and collaborating on ML systems, fostering open-source and open-science initiatives. We host machine learning models and datasets within an infrastructure that enables efficient data processing, analysis, and research. Additionally, we provide educational resources, courses, and tooling to lower the barrier to AI participation for individuals from all backgrounds.

Executive Summary

The 2025 National AI R&D Strategic Plan represents a defining moment for American leadership in artificial intelligence. This plan must prioritize federal investments that advance **public goods** and **scientific discovery**, address **structural market blind spots**, democratize access to **tools, infrastructure, and opportunity**, ensure **trustworthy and resilient AI systems**, and anticipate and manage AI's **societal and labor impacts**.

Open, accessible, and efficient AI ecosystems already demonstrate how inclusive innovation can match and often exceed closed systems. What once required models with over 100 billion parameters two years ago can now be accomplished with 2 billion parameter models. Research shows that open technical systems act as economic force multipliers, with an estimated 2000x multiplier effect, meaning strategic federal investment in open systems could generate trillions in value across the American economy by distributing the benefits of AI development and the control of advanced systems from being concentrated among a handful of big private companies to all Americans.

The federal government must lead by investing in AI that is **open and auditable**, **efficient and accessible**, **scientifically empowering**, **governable and safe**, and **responsive to labor and**

societal needs. By grounding the 2025 National AI R&D Strategic Plan in public mission, democratic access, and real-world resilience, the U.S. can ensure that the most powerful technology of our time is a force for broader accessibility, competition, innovation, and shared prosperity.

We urge the federal government to deepen its commitment to open science, public infrastructure, and socially aligned AI through four strategic priorities that complement rather than compete with private investment: (1) Fundamental AI research for efficiency, transparency, and universal access; (2) AI for science, health, climate, and national resilience; (3) Trustworthy and secure AI for public institutions; and (4) Strengthening AI infrastructure.

Key Federal AI R&D Investment Priorities

1. Fundamental AI Research for Efficiency, Transparency, and Universal Access

Current research and investment trends prioritize larger models and proprietary architectures. While these can be credited for notable performance increases on specific tasks over the last couple of years, they only address part of the demand for better AI across the country. For AI to serve rural clinics, local governments, small businesses, or public universities, among others, it must be significantly more **efficient, transparent, and accessible**.

The open ecosystem naturally drives efficiency innovations that well-resourced commercial organizations have little incentive to pursue. Resource-constrained teams consistently demonstrate that necessity drives innovation, achieving remarkable performance with dramatically fewer resources. The open AI research community has produced systems like AI2's OLMO2, which matches the performance of leading proprietary models while being fully transparent, and OlympicCoder, which exceeds Claude 3.5's performance on complex coding problems despite being much smaller.

The implications of investing in open research for the American public are far-reaching. For instance, rural hospitals can improve patient outcomes using existing hardware like phones and laptops, small biotechs can discover drugs independently, and community colleges can offer industry-standard AI training to under-resourced students via free courses.

Recommended Research Directions:

- **Doing More with Cheaper Models:** Continue to develop compression techniques achieving massive reduction ratios while preserving model performance, enabling advanced AI to run on mobile phones and edge computing devices that small businesses and rural communities can afford. Smaller models also allow finetuning and modification for specific use cases on local hardware.
- **New Technical Paradigms Beyond Transformers:** Invest in research into alternative AI architectures, such as selective state space models inspired by biology (Mamba) or

diffusion language models that can run extremely fast, such as Google's Gemini Diffusion Model, as well as sparse and conditional computation approaches like Mixture of Expert models or systems that offload tasks to tools. Transformers, the current dominant generative AI architecture, are extremely resource-intensive and show diminishing returns, while these new paradigms could potentially reduce computational costs by orders of magnitude.

- **Mathematical Foundations for Predicting Capabilities:** Invest in research into scaling laws, aiming to form rigorous mathematical models that estimate the resource needs, model architecture, and training data size required to create models with the required capabilities.
- **Robotics Research for National Applications:** Support AI-powered open robotics systems for emergency response, public infrastructure maintenance, and scientific automation. Long-horizon federal investment will boost breakthrough capabilities and lower hardware costs.
- **Task-Specific Model Selection:** Most users do not require an expensive general-purpose model for their specific tasks. Research into choosing the most efficient model for context-specific tasks can be extremely cost-effective in the long run.

2. AI for Science, Health, Climate, and National Resilience

Critical domains like public health, energy systems, and environmental monitoring face a persistent market failure: high public benefit but uncertain commercial returns make them unattractive for private sector AI investment. This leads to general-purpose models geared towards consumer usage that may not be well-suited for specific public benefit tasks. This creates strategic vulnerabilities where America's most pressing challenges lack adequate AI research support.

The government has a unique role in supporting AI applications that serve essential national functions but may not generate immediate profits. Historical precedent shows that federal investment in foundational technologies, such as GPS and the internet, creates capabilities that eventually benefit the general economy. Today's scientific challenges require similar long-horizon capital and a mission-driven research focus.

Federal investment in open scientific AI could accelerate medical research to cure complex diseases, discover materials critical to energy independence, and provide early warning systems for environmental threats that affect American infrastructure. These capabilities represent essential infrastructure for American resilience and competitiveness in the 21st century.

Recommended Research Directions:

- **AI-Accelerated Materials Discovery:** Create specialized AI systems for discovering new materials critical to batteries, semiconductors, and strategic technologies, which are capabilities essential for reducing dependence on foreign supply chains.

- **Domain-Specific Foundation Models:** Develop and maintain large-scale AI models specifically trained on scientific literature, experimental data, and simulations for biology, medicine, chemistry, physics, and other basic sciences openly available to American researchers.
- **Climate & Environmental AI Tools:** Develop AI systems for processing satellite imagery, sensor networks, and climate simulations to track environmental changes, predict extreme weather events, and optimize resource usage.
- **Pandemic and Health Preparedness:** Build AI systems for early disease detection, drug discovery, protein folding research, epidemiological modeling, and privacy-preserving surveillance systems that can be rapidly deployed during health emergencies.

3. Trustworthy and Secure AI for Public Institutions

As AI systems become integrated into critical infrastructure, like transportation networks, healthcare systems, and emergency response, they must meet the same rigorous standards of **security, reliability, and auditability** that we demand from existing public systems. Unlike consumer applications that can fail gracefully, critical infrastructure requires AI that works reliably under attack and in unexpected conditions.

Current AI development often prioritizes capability over trustworthiness, creating systems that are powerful but opaque. This approach has tensions with public sector requirements where decisions must be explainable, systems must be auditable, and failures can have life-or-death consequences. Building on decades of cybersecurity experience while addressing AI's unique challenges requires coordinated federal action that no single company can provide.

The implications extend beyond technical requirements to transparent governance itself. When AI systems make decisions such as benefit allocation, algorithmic hiring or other critical decisions that can impact citizens' lives, the public has a right to understand and verify those decisions. Trustworthy AI for public institutions is therefore both a technical and civic imperative.

Recommended Research Directions:

- **Verifiable and Auditable Reasoning:** Develop AI that can explain its reasoning in ways humans can verify and trust, including neurosymbolic AI that combines neural networks with formal reasoning and chain-of-thought mechanisms that can be audited.
- **Resilience to Adversarial Threats:** Research defenses against sophisticated attacks including data poisoning, model extraction, and adversarial examples designed by adversaries including nation-state actors.
- **Documentation, Provenance, and Transparency:** Develop comprehensive requirements tracking all components of AI systems including training data sources, model architectures, and computational requirements, enabling security auditing and ensuring systems can be verified and reproduced.

- **Privacy-Preserving Methods:** Develop techniques enabling AI analysis of sensitive public data without exposing individual information, including homomorphic encryption and secure multi-party computation for collaborative learning.

4. Strengthening AI Research Infrastructure and Understanding Societal Outcomes

America's AI leadership depends not just on breakthrough frontier AI systems, but on ensuring broad access to the tools, data, and computational resources needed for innovation. Currently, if only elite institutions and for-profit entities have access to cutting-edge compute and expertise, AI will concentrate advantages and miss opportunities for broader innovation. This concentration of resources threatens both American competitiveness and democratic values.

Simultaneously, we face a critical knowledge gap about AI's societal impacts. As AI reshapes work, education, and civic life, we lack systematic understanding of these changes and their implications for policy. Without rigorous research on AI's effects on employment, regional development, and economic outcomes, we cannot make informed decisions about this technology's role in American society. Federal investment must therefore serve a dual purpose: democratizing access to AI capabilities across American institutions while generating the knowledge needed to govern AI's societal impacts wisely.

Recommended Research Directions:

- **Fund National Public AI Research Infrastructure:** Expand NAIRR beyond the pilot phase and expand public research infrastructure such as the Argonne National Laboratory with sustained funding for compute, storage, and networking accessible to researchers nationwide, including dedicated allocations for high-risk research and mechanisms for sharing successful models across institutions.
- **Public Knowledge Digitization and Access:** Launch large-scale initiatives to digitize and make accessible government documents, scientific literature, and cultural heritage materials currently locked in analog formats, creating training resources that serve the public interest.
- **Accessible AI Interfaces:** Create natural language and visual programming environments that make advanced AI capabilities available to researchers in social sciences, humanities, policy research, and journalism who lack programming expertise.
- **Longitudinal Societal Impact Research:** Conduct rigorous studies on how AI deployment affects employment, skill requirements, job quality, and economic distribution across different industries, demographics, and regions, with particular attention to early indicators of technological displacement.
- **AI Workforce Transition Analysis:** Develop methodologies for measuring AI's impact on work satisfaction, income inequality, and regional economic development to inform evidence-based policy responses and workforce development programs.

HUGGING FACE

Open Research Partnerships for National Impact

The public must benefit from the research it funds. Federal AI R&D requires new collaboration models that combine government domain expertise with private sector capabilities while ensuring public benefit:

- **Public-Private Evaluation Partnerships:** Establish frameworks combining government use cases with industry technical expertise, with explicit mandates to prioritize open evaluation data and tooling that boost the open research ecosystem.
- **Distributed Research Networks:** Enable secure collaboration across institutions using federated learning and privacy-preserving techniques that protect sensitive data while advancing collective knowledge.
- **Open Development Mandates:** Require federally-funded research to release not just papers but usable artifacts, such as code, models, datasets, and reproducibility materials, with sustained support for maintenance and public access.

Conclusion

The 2025 National AI R&D Strategic Plan should prioritize areas where federal investment uniquely advances American interests and leadership—fundamental research, public goods, and long-term challenges that market forces cannot adequately address. Open approaches to AI development consistently demonstrate their ability to match or exceed proprietary systems while enabling broader participation, better security through transparency, and rapid innovation.

American leadership in AI requires more than innovation, it also demands accessibility, security, and self-reliance. When researchers, industries, and institutions across the country can build and deploy advanced AI openly, safely, and in service of the American public, we strengthen both our economic foundation and national resilience. By investing in AI that is transparent, efficient, and aligned with real-world needs, the federal government can ensure this technology drives opportunity, protects our sovereignty, and secures America's position as the undisputed global leader in the next era of innovation.

Submitted by:
Avijit Ghosh, Applied Policy Researcher, Hugging Face
Yacine Jernite, ML & Society Lead, Hugging Face
Irene Solaiman, Head of Global Policy, Hugging Face