PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mba-1e80-ao8l
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-0299
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Oregon State University

General Comment

See attached file(s)

Attachments

Response_to_NSF-2025-OGC-0001

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.

Response to Docket: NSF-2025-OGC-0001

Recommendations on Priorities to inform the Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan

Oregon State University (OSU) is Oregon's land grant university and the state's largest research university. Our research spans areas that include agricultural sciences, nuclear energy, artificial intelligence, bioengineering, and more. Artificial intelligence is a foundation of OSU's strategic plan, and the university has a long history of national leadership in AI research, development, and instruction. The university's research excellence is matched by a commitment to engagement with industry, through industry sponsored research, technology transfer, spinouts, and other means. Based on this experience, members of the OSU research community offer the following recommendations for the development and success of the 2025 National AI R&D Plan. We welcome opportunities for further discussion of these ideas.

Recommendations:

#1. Address Urgent National Priorities through Embodied AI

The next 3–5 years represent a pivotal opportunity to shape the trajectory of embodied AI, particularly for robotics and autonomy. Many of the United States' most urgent national priorities—from modernizing public infrastructure to increasing housing supply, securing supply chains, and scaling healthcare and elder care—are fundamentally constrained by a shortage of physical labor. The scale and complexity of these tasks vastly exceed the capacity of the current workforce, and demographic trends suggest that the gap will only widen in the coming decades.

The United States needs a physical labor multiplier—a way to dramatically expand the available labor force without proportionally increasing its size. General-purpose physical intelligence, embodied in mobile, semi-autonomous robots and autonomy enabled systems that can adapt across environments and workflows, is the most viable long-term path to achieving that goal. These AI enabled systems can enable one worker to direct many semi-autonomous entities and accelerate progress on nationally strategic efforts such as:

- Rebuilding transportation and energy infrastructure
- Constructing affordable housing
- Strengthening domestic logistics, manufacturing, and warehousing
- Supporting care work in hospitals and homes
- Responding to disasters or national security events.

Achieving this kind of capability requires AI enabled systems, particularly embodied systems, that are flexible, instructible, general, able to work in teams, and able to operate safely and usefully in environments that have not been retrofitted specifically for automation. Along with research in the areas listed in recommendation #1 above, the 2025 AI R&D Plan should include support for adaptable general-purpose robots.

This is an area where universities can complement industry developments. Industry is currently advancing embodied AI through available narrow, near-term solutions that do not prioritize flexibility, generality, or long-term autonomy. The reason is simple: ***today's constrained, high-cost AI learning pipelines are still vastly better than the alternative of full infrastructure re-engineering.*** Even if it takes 30 days of data collection and teleoperation to train a robot to perform a specific workflow, that investment is often far cheaper and faster than building a custom robotic assembly line. These factors lead to industry solutions that are:

- Heavily scripted or prompt-tuned, specialized to tightly controlled workflows.
- Trained via teleoperated, labeled demonstrations, requiring significant expert effort per task.
- Marginally generalizable, with limited applicability beyond the deployment environment.

In addition to economic factors, with current solutions available to industry, **liability and safety concerns drive industry away from the most socially critical applications**— particularly those involving human-rich, unstructured environments such as homes, hospitals, construction sites, and disaster zones. The current safety toolkits for embodied AI are inadequate to guarantee that a robot won't, for example, misclassify a person as a movable object or damage fragile infrastructure through a planning error. Universities have the ability to conduct research on general purpose AI that can function in these environments and be adopted by industry.

#2. Leverage University Supercomputers for Complex AI Simulations

The 2023 Update included the priority: *Developing Shared Large-Scale and Specialized Advanced Computing and Hardware Resources*. We recommend the new plan add the priority of support for bringing communities and companies together around sharing highimpact university-based AI supercomputing infrastructures. These resources and those groups can aid research in:

- massively large robot swarms
- large, heterogenous multiagent systems for firefighting or other disaster response (e.g. evacuation post-tsunami on coastlines)
- AI systems that operate over extended periods of time (power grid management or home care for older adults) where the goals might change, and the systems will need to balance multiple stakeholders with potentially conflicting objectives.

Research in these critical AI directions requires simulations that are difficult, if not impossible, within the operational or financial constraints associated with the computing resources of a national supercomputing facility or commercial cloud resources. Because of their expense, or competition for their capacity, those resources do not aid the development of simulations that persist, e.g., digital twins, need to be run repeatedly, or offer the flexibility of permitting modifications to a simulation. University supercomputing facilities can offer greater flexibility and are integrated with regional communities and companies.

#3. Priority Areas for Research

Once the initial wave of adoption of narrowly scoped AI applications saturates, **market differentiation will shift toward flexibility and speed of deployment.** Companies that can rapidly adapt their systems to new environments, tasks, and users with minimal additional training will have a decisive competitive edge. Without foundational research into this adaptability, U.S. companies risk becoming locked into brittle, one-off solutions, restricted to a narrow application area, while global competitors build on more adaptable, general-purpose platforms.

Much of today's commercial AI ecosystem—from deep learning to reinforcement learning to foundation models—originated through federally supported academic research. The same approach is needed to support research on adaptable embodied and non-embodied AI. The techniques required to enable adaptability—robust generalization, efficient transfer, modular system design, learning from multiple objectives, and real-time learning from sparse datasets—are exactly those that demand long-horizon, federally supported research.

Universities are uniquely positioned to lead solving these long-horizon, high-risk challenges by:

- pursuing interdisciplinary research across AI, robotics, language, control, and safety.

- building open-source tools, datasets, and benchmarks that accelerate industry adoption.
- training the next generation of embodied AI researchers and practitioners.
- being accountable to public needs, not just short-term market demand.

To enable the AI capabilities the USA needs to remain an international leader, we recommend the 2025 National Artificial Intelligence Research and Development Strategic Plan (2025 AI R&D Plan) focus federal investment focus on basic research in the following core capabilities that are underdeveloped in industry due to their complexity, long time horizons, or lack of near-term commercial payoffs:

<u>Predictive Modeling and Physical Reasoning</u>:

Developing AI capabilities can anticipate and evaluate outcomes before acting—using learned models of physics, causality, and task dynamics to enable more efficient training, safer behaviors, and human-in-the-loop plan refinement without extensive trial-and-error.

<u>Robust Safety and Skill Verification</u>:

Build systems that understand their own limits, detect likely failure modes, and verify safety constraints during planning and execution. This includes minimizing unintended disturbances and ensuring actions align with high-level goals and physical norms.

<u>Long-Term Memory and Adaptation</u>:

Support AI frameworks that can retain and apply knowledge across tasks, environments, and users. Key needs include persistent memory architectures, incremental skill refinement, and the ability to update workflows through natural instruction over time.

<u>Low-Supervision Learning from Language and Demonstration</u>:

Enable learning new tasks from a few demonstrations, simple language commands, and sparse non-IID datasets. This includes grounding abstract concepts (e.g., “fragile,” “left side of the truck”) in perception and action, and generalizing across similar but unseen tasks.

<u>Modular and Extensible Architectures</u>:

Design AI systems that evolve over time, adding new skills, concepts, and behaviors without retraining from scratch. Priorities include composability, stable integration of new components, and scalable representations that support continual improvement.

<u>Multi-objective AI systems</u>:

Design AI systems that learn from multiple rewards (or objectives) across different time scales. Priorities include determining which objectives are relevant at what time, dynamically shifting performance as mission needs change, and balancing short- and long-term goals.

Human-AI Interaction:

Enabling humans in differing interaction roles to work with AI-enabled systems as in situ partners, which requires developing new hands-free, but not necessarily natural language-based, interaction modalities that result in bi-directional communication, understanding, decision support, and task achievement.

Levels of Intelligence:

Different AI systems will need different levels of intelligence depending on their morphologies, applications, environments, and associated risks. For example, delivery drones are highly autonomous and handle almost all faults and issues themselves with no human intervention, including detecting and avoiding other aircraft. Priorities include understanding the variation in forms of intelligence appropriate to particular domain relevant contexts.

#4. Build a Competitive AI Talent Pipeline

To sustain US competitiveness in AI, particularly against strategic rivals like China, ensuring a robust and continuous supply of highly skilled AI researchers entering the workforce is a paramount concern. The 2025 AI R&D Plan should prioritize developing a pipeline of talent for AI research. There should be support for attracting capable students to AI careers and creating environments where they can excel, learn using state-of-the-art resources, and gain recognition. While the US cannot match China's sheer population size (China has roughly 15 million 18-year-olds annually compared to the combined total of the US, South America, and Europe), it can strive to maximize talent yield from its diverse population through outreach and support programs.

Key initiatives to maximize that talent yield should include:

1. Support for regular workshops, conferences, and competitive events (e.g., AI challenges or "Grand Challenges") to stimulate idea exchange and motivate students to push the boundaries of the field.
2. Federally supported AI-targeted fellowships and scholarships for both grads (M.S. and Ph.D.) and undergrads.
3. Opportunities for recognition: prestigious credentials or honors (inspired by systems like the Chess Grandmaster title, perhaps a "National AI Scholars" program) to recognize top talent, enhance their visibility, and facilitate career progression into academia and industry.
4. Access to Advanced Tools: Fewer students may pursue AI research if they perceive a lack of access to top-tier educational tools and resources, shrinking the potential talent pool needed for long-term competitiveness. Reliance on potentially lagging

open-source models or expensive proprietary systems can create a training gap: the next generation of AI researchers may not gain hands-on experience with the most advanced algorithms unless they rely on commercial platforms, potentially limiting their foundational understanding and innovation capacity and deterring students.

Compiled by David Barber, Director, Strategy for Research Computing, Office of Research Advancement, Division of Research and Innovation, Oregon State University.