PUBLIC SUBMISSION

Received: May 29, 2025
Tracking No. mba-6q0s-dy6h
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-0329
Comment on FR Doc # 2025-07332

Submitter Information

Government Agency Type: State
Government Agency: University of Michigan

General Comment

Thank you for the opportunity to provide feedback on the development of a 2025 National Artificial Intelligence Research and Development Strategic Plan. The attached document includes our response, making the case for "Foundational AI as a Catalyst for Real-World Impact". This response is submitted by Rada Mihalcea, Jenna Wiens, Michael Wellman, H.V. Jagadish, and Jing Liu, on behalf of the Michigan AI Lab, MIDAS, AI & DHI and e-HAIL at the University of Michigan

Attachments

UM.RFI.AI.Response.2025

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

Submitted by: Rada Mihalcea, Jenna Wiens, Michael Wellman, H.V. Jagadish, and Jing Liu, on behalf of the Michigan AI Lab, MIDAS, AI & DHI and e-HAIL at the University of Michigan

Foundational AI as a Catalyst for Real-World Impact

AI remains an evolving scientific discipline, closer to a "rough draft" than a finished technology. Large models today often default to one-size-fits-all behaviors, struggling to recognize the subtle context shifts that matter in real-world use. In a hospital, for example, a radiology assistant must integrate multimodal data (e.g., images, patient history, clinician notes), while gracefully adapting to a resident's learning needs one moment, and a senior specialist's diagnostic the next. Advancing from such brittle uniformity to truly context-aware systems that have real-world impact demands fundamental research in (1) alignment (fine-grained adaptation to diverse user goals and expertise), (2) grounding (multimodal, dynamic representations of the world), and (3) instructability (long-term, co-evolving human-AI collaboration), all alongside with interdisciplinary collaborations spanning computer science, cognitive science, domain sciences, and the social sciences, across global research communities. Pursuing these questions is essential to scientific progress as it is to industrial competitiveness, and a focus on the foundations of AI has the potential to lead to the next wave of transformative applications.

Importantly, we need to recognize that AI does not advance in a vacuum. Progress comes with parallel innovations in areas such as statistics, information theory, cognitive science, and domain-specific expertise covering a broad space of disciplines such as healthcare, education, transportation, and more. Investing in close collaborations with these "supporting sciences" will expand the search space of algorithmic ideas and ground AI in empirical reality. When materials scientists discover new battery chemistries, they will enable smarter energy-constrained edge models; when ethicists and HCI scholars will uncover user patterns, they will inform better training protocols. A holistic strategy that nurtures these upstream or peer disciplines, alongside core AI research, will build a robust pipeline of collaborative knowledge and will ensure far-reaching real-world impacts of AI.

Specific Directions in Foundational AI Research

Recent years have witnessed a shift from "scientific discovery" as the primary objective of AI development to an "industry competition" environment dominated by leaderboards and benchmark-driven model accuracy. While this shift has driven impressive gains in the performance of large-scale models, it has also reinforced a narrow development paradigm centered around a small set of high-performing models that largely adopt a one-size-fits-all

approach. These models often lack the flexibility to accommodate the diversity of real-world contexts, user needs, and dynamic environments. We recommend renewed and sustained investments in AI research that refocus on scientific discovery and foundational understanding. This includes prioritizing models that are not only powerful but also (1) adaptable to diverse users and evolving conditions, (2) context-aware in their interpretation and response to situated inputs, and (3) instructable in ways that enable users to effectively guide and customize AI behavior for real-world problem solving and decision-making. While industry will undoubtedly play a role in addressing these challenges, academia is uniquely positioned to lead long-term, high-risk research; to explore ideas unconstrained by immediate commercial viability; and to integrate interdisciplinary perspectives that deepen our understanding of AI's societal implications and potential.

1. We need AI systems that can adapt with minimal supervision. We need more research into adaptive AI systems that can align with a wide range of user needs and preferences. Current techniques, such as reinforcement learning from human feedback (RLHF), often rely on static, single reward models and require extensive user-labeled data, making them difficult to scale or personalize for new or infrequent users. To address these limitations, future research should explore new approaches that enable rapid personalization to new users, and broader inclusivity by capturing a wider range of interaction styles, goals, and constraints. Further, AI systems should be designed to operate efficiently with limited user supervision. We need research into interactive protocols (e.g., multi-choice preference elicitation) that allows users to easily convey their goals and preferences, thus enabling AI systems that are robust and aligned (or alignable) with the full spectrum of human needs.

2. We recommend increased support for research into context-aware AI through multimodal and causally-grounded adaptation. We must pursue foundational research on AI systems capable of interpreting and adapting to dynamic, real-world environments through multimodal and temporally-aware mechanisms. Most current AI systems are trained on static datasets and fail to reflect the evolving nature of user contexts, environmental cues, and interactions over time. As AI becomes increasingly embedded in daily life, its ability to recognize and respond to context changes (whether in physical surroundings, user needs, or multi-user scenarios) will be essential for safe and effective deployment. Future research should explore the development of generative and predictive models that integrate information across modalities, such as visual inputs, spoken language, and written text, to detect and respond to temporal changes in both the environment and user behavior. This includes identifying early indicators of evolving needs, modeling shared context in multi-user interactions, and simulating future states of interaction to proactively adapt responses.

NSF should also support work that incorporates causal reasoning into training objectives, for example, encoding physiological or behavioral mechanisms that inform how user needs may evolve. Causally motivated approaches could improve the efficiency of the AI data use and support generalization to underrepresented or novel scenarios. By investing in these directions, NSF can help catalyze the development of robust, context-aware AI systems that maintain relevance and utility in changing environments, while respecting the complexity of real-world human experiences.

3. We recommend targeted investments to advance the robustness of instruction-tuned AI systems. Current instruction-tuned models have been shown to be fragile, often producing

inconsistent outputs in response to semantically equivalent instructions that differ in wording. This lack of robustness limits the accessibility and usability of these systems, especially for users with different life experiences, communication modes, or interaction needs.

We recommend more research on developing models that are resilient to variation in how instructions are phrased and that can adapt to individual users' backgrounds. This includes compositional modeling of user characteristics, which can respect the privacy of the users, while supporting scalable personalization and enhancing the ability of AI systems to serve a broader and more diverse user base.

Investments in Frameworks and Resources

Advancing foundational AI and translating it into real-world impact requires strategic investments in key infrastructures. These include support for use-inspired research, resilient and trustworthy information ecosystems, pipelines for responsible AI adoption, scalable AI education and workforce development, and frameworks that support interdisciplinary and global collaboration.

1. Prioritize the development of a robust national pipeline for use-inspired research. To ensure that advances in AI research translate into tangible societal benefits, NSF should prioritize the development of a robust national pipeline for use-inspired research. This pipeline should include mechanisms that actively engage domain experts, community stakeholders, and end users throughout the research lifecycle, from problem formulation and data collection all the way to system design, deployment, and evaluation.

One core component of this pipeline is the creation of sustained partnerships with practitioner communities and organizations that serve populations with a wide range of lived experiences and needs. These partnerships should not only inform research priorities but also provide access to real-world use cases, facilitate iterative feedback, and help avoid solutions that are technologically impressive but practically irrelevant. By grounding AI research in the lived experiences of communities, researchers can identify problems that matter, develop systems that are contextually appropriate, and ensure technologies are designed with, rather than for, the people they aim to serve. We recommend investment in formal structures that support such collaboration, such as funded partner programs or shared governance models, that support the collaboration of AI researchers and the stakeholders of the AI solutions being developed.

2. Build and maintain reliable and trustworthy information ecosystems. AI is already transforming how we access and engage with the accumulated knowledge of our society: its scientific discoveries, archives of observation and experience, artistic creations, cultural wisdom, and traditions of thought and teaching. Leveraging this access promises enormous benefits in education, invention, and economic productivity. With this promise, however, comes significant risk that we lose control of the integrity of our information ecosystem. As the most critical functions of a complex society depend on accurate information--the financial system, public health, and civic decision-making, to name a few--understanding the effects of AI on the information environment is fundamental to ensuring that AI operates to the benefit of humanity.

Research on technology's effect on information and misinformation is sensitive, and controversial. There is legitimate concern that the science could be distorted by judgments

reflecting the personal biases of researchers working in this area. We must find ways to address these concerns, as the underlying problem is too urgent to ignore. The development of objective standards for research in information integrity is itself a ripe subject for investigation. NSF can lead the way in this effort, through the establishment of careful guidelines that ensure judgments of misinformation are not tainted by the subjective beliefs of individual scientists.

3. Provide institutional-level support to facilitate the integration and sustained adoption of AI in research, ensuring our nation's scientific competitiveness. We recommend that the national AI strategy address not only the development of new models and computational infrastructure but also the adoption and integration of AI across the research enterprise because such adoption and integration is critical for our nation's competitiveness in science. The adoption of AI is not just about adding new tools; it requires the rethinking of research workflows and how research teams operate, and establishing mechanisms to ensure the trustworthiness and interpretability of AI-driven results. Without such institutional capacity, even the most powerful AI tools will have limited scientific impact.

An example of such institutional capacity building is the work carried out by the MIDAS Institute at Michigan, which focuses on building research capacity and enabling research breakthroughs across research fields through the adoption of AI. Such efforts have allowed for the development of AI-embedded research workflows that help scientists make effective use of AI without needing to become AI experts. Currently, funding agencies have prioritized the development of AI technologies and the creation of training programs, but have largely overlooked the institutional scaffolding that allows AI to take root and flourish in the day-to-day practice of science. We recommend investments into novel institutional support models, such as centralized staff teams, cross-disciplinary AI connectors, and frameworks for responsible use, that can help the institutions move from isolated demonstrations to sustained AI adoption.

4. Incorporate AI education at all educational levels and across disciplines to equip future professionals in diverse fields with AI understanding and skills. We need to prioritize scalable AI education that reaches learners at all stages and across all disciplines. AI education should begin early in students' academic journeys, integrated into K–12 curricula in age-appropriate ways, and continue through higher education, where it should be infused into all fields beyond computer science, including health, law, agriculture, social sciences, the arts, and more. This integration of AI education across disciplines will ensure that future professionals across different sectors will be equipped to understand, evaluate, and responsibly use AI in their domains. Additionally, given the rapidly evolving nature of AI technologies, we recommend the support of programs that develop life-long AI learning opportunities, including continuing education programs, certifications, and accessible resources for mid-career professionals, policymakers, and workers in transitioning industries.

5. Support sustained interdisciplinary and international collaboration in alignment with the national AI strategy to maintain and enhance U.S. leadership in AI. While the RFI appropriately emphasizes the importance of a national AI strategy to secure U.S. leadership, it is equally important to recognize that the challenges and opportunities posed by AI are global in nature. As inhabitants of a shared planet, "we are in it together," facing common concerns that demand collaborative, cross-border solutions. Historically, many of the most transformative scientific advances have emerged from interdisciplinary teams that brought together a wide range of perspectives, often across national and institutional boundaries. Supporting such collaboration

will not only amplify the reach and relevance of U.S.-led research but also helps build trust and interoperability that benefit everyone. We strongly recommend that the national AI strategy explicitly support sustained interdisciplinary and international collaboration, ensuring that U.S. innovation contributes meaningfully to shared worldwide progress, while also benefiting from the collective intelligence and creativity of the broader scientific community.

About AI at Michigan

The University of Michigan is at the forefront of advancing artificial intelligence, from its core enabling technologies to impactful, interdisciplinary use-inspired research that aim to address today's complex challenges. Our pioneering research builds on longstanding leadership that goes back to the dawn of the information age. Some of the early work around the time of the "Dartmouth Summer Research Project on Artificial Intelligence" workshop where the concept of AI was introduced included work on "learning machines", followed soon after by work on computer vision, robotics, and cognitive architecture, and later the establishment of the Michigan AI Lab in 1988. Since then, our faculty have continued to advance the models that underpin AI technology, including generative AI, machine learning and deep-learning models that uncover patterns in complex data, as well as advanced computer-vision and natural language models.

Today, AI research at Michigan spans a broad and growing ecosystem. In addition to the Michigan AI Lab, we are home to MIDAS (Michigan Institute for Data and AI in Society), e-HAIL (e-Health and Artificial Intelligence initiative), AI & DHI (AI & Digital Health Innovation initiative), MICDE (Michigan Institute for Computational Discovery and Engineering), a new Department of Robotics, among other campus-wide initiatives. These units serve as hubs for cross-disciplinary collaboration, bringing together computer scientists, engineers, clinicians, ethicists, social scientists, and artists, among others, to push the boundaries of what AI can do and for whom.

Our faculty collaborate extensively across the University and beyond, leveraging U-M's breadth of expertise to explore novel applications in healthcare through partnerships with Michigan Medicine; in sustainability and energy; in transportation and mobility; in semiconductors and materials; and in education, policy, and the humanities. We also maintain strong partnerships with industry to ensure our research translates into real-world impact, and we advise policymakers on how to build trustworthy and safe AI systems. Meanwhile, our students graduate to become leaders across academia, industry, and the government.

The AI research at Michigan has been consistently ranked among the top in the nation, with strong placements in the U.S. News & World Report and CSRankings, among others.