PUBLIC SUBMISSION

Received: May 27, 2025
Tracking No. mb6-tnat-mwih
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-0119
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Cognitive Science Society

General Comment

See attached file

Attachments

CognitiveScienceSociety_ AIActionPlan

Executive summary

Modern AI has its intellectual roots in the vibrant, multidisciplinary research space of *cognitive science*, where work from multiple traditions in the mind and brain sciences is integrated to advance the science of intelligence. The Cognitive Science Society, the world's leading professional organization for this community, has recognized many of these contributions through its flagship Rumelhart Prize, honoring a pioneer in the development of neural network architectures, David Rumelhart. Winners of this prize have developed fundamental algorithms and architectures underlying modern deep learning systems (Jeffrey Elman; Geoffrey Hinton, also winner of the 2024 Nobel Prize in Physics as well as the 2018 Turing Award; Michael Jordan; James McClelland; Paul Smolensky). Critically, these contributions reflect the deeply multidisciplinary orientation of cognitive scientists. For example, convolutional neural networks (a fundamental step in the development of deep learning systems) by Yann LeCun (2018 Turing Award) and Neocognitron by Kunihiko Fukushima resulted from integrating insights from the neuroscience of mammalian vision with computer science.

Over the next 3-5 years, taking AI research to the next level will require an **evidence-based approach to AI**. Cognitive scientists are uniquely positioned to inform such efforts at all stages of AI system development, testing, and deployment. **This requires policies that foster the development of multidisciplinary research projects** so that we can effectively capitalize on the insights of different research traditions in the study of human intelligence.

Cognitive Science's key role in the development of AI solutions

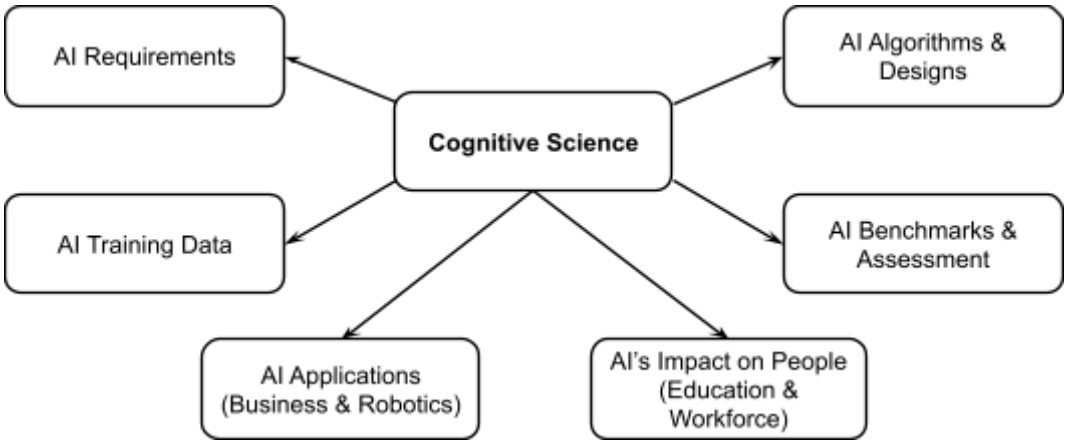

Cognitive Science provides the cross-disciplinary foundation for theoretical and empirical approaches to understanding natural intelligence and is therefore ideally situated to offer evidence-based insights into future requirements, architectures,

designs, training data and benchmarks for the development and assessment of artificial intelligence applications. Below, we outline how a cognitive science perspective on policy can advance US research into AI at this critical moment.

Responses

Fundamental advances in AI algorithms, architectures, mathematical foundations and computing paradigms; research on architectures beyond deep learning

Cognitive science has been a key player in the development of AI algorithms and architectures beyond deep learning, pioneering the development of symbolic models (including Rumelhart Prize winners John Anderson, Aravind Joshi), Bayesian approaches (Michael Jordan, Judea Pearl), and computational neuroscience (Peter Dayan, Shimon Ullman). This diversity of approaches reflects a fundamentally cognitive science orientation: a focus on understanding the nature of human intelligence rather than utilizing a particular AI architecture. Such an orientation is essential for pushing AI development to the next level – to imagine AI beyond the limits of current deep learning systems.

Current training for AI developers does not typically reflect this orientation. Developers are trained in the sciences of software engineering, artificial intelligence and mathematical techniques but lack the understanding of intelligent systems which cognitive science describes. These include the sciences of adaptive behaviour, problem solving, decision making, learning, action control, language and reasoning.

Policy recommendations:

- Rather than focusing strictly on specific computational methods, ensure calls for research proposals target modeling of capabilities underlying intelligence and intelligent behavior.
- Provide support for team science approaches that explicitly incorporate research insights from multiple disciplines.
- Support development of cross-disciplinary training opportunities for AI developers to broaden their understanding of core issues in human intelligence.

Research into AI standards and reliability; advances in AI systems capable of reasoning, adaptability, and robustness in dynamic environments

Cognitive scientists currently play a crucial role in evaluation of AI systems. For example, cognitive scientists in academia and industry have played a central role in the development of AI benchmarks. Researchers with a background in the cognitive science of language have developed key benchmarks for large language models;

researchers with a background in vision science have played a similar role in developing benchmarks for object recognition systems. Furthermore, while AI leaderboards typically focus on static assessments of model performance – how well does a model perform after receiving massive amounts of training? – cognitive scientists have extensively studied how humans learn, reason, and adapt in dynamic environments. These insights will prove invaluable in developing the next generation of AI benchmarks; current systems fall far short of human capabilities in these contexts.

Policy recommendations:

- Require AI research proposals to incorporate research-informed benchmarks grounded in the science of intelligence.
- Ensure that the next generation of AI systems face the challenge of matching human performance in dynamic environments.

Research on AI systems and education supporting American workers; high-risk, high-reward AI research relevant for domains critical to future U.S. competitiveness, including human-AI interaction

What are the implications of AI for human society? It is essential that we understand how workers *actually use* AI, rather than how designers intend for it to be used; that we understand whether AI tools *actually improve* learning and education, rather than our hopes for how children might benefit from technology. Cognitive scientists are perfectly positioned to examine these issues. Using a rich set of techniques, honed by extensive studies of human behavior, cognitive science is the ideal source for the evidence base policy makers need to accurately assess the appropriateness and effectiveness of AI interventions in human organizations.

Policy recommendations:

- Couple investments in AI system development for the workforce with rigorous, evidence-based assessment of their impact on human cognition.
- Require AI interventions in the educational system to be subject to rigorous, evidence-based assessment.