

## AI RFI Responses, October 26, 2018

---

### Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan RFI Responses

**DISCLAIMER:** The [RFI public responses](#) received and posted do not represent the views and/or opinions of the U.S. Government, National Science and Technology Council (NSTC) Select Committee on Artificial Intelligence (AI), NSTC Subcommittee on Machine Learning and AI, NSTC Subcommittee on Networking and Information Technology Research and Development (NITRD), NITRD National Coordination Office, and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality or content of all external links included in this document.

**Response to Request for Information (RFI) on:  
“AI Research and Development Strategic Plan”**  
*Indiana University*

Overall, the Strategic Plan presents a compelling vision for advancing AI and identifies a set of areas which remain crucial areas for progress, each highly worthy of support and in their totality providing good coverage. Particularly commendable is the inclusion of perspectives from many angles including technical, social, ethical, workforce, etc. The strategies in the report remain timely. The overall message that U.S. leadership depends on focusing “its investments on high-priority fundamental and long-term AI research” is even clearer today than when the report was first released.

**General suggestions**

We consider it likely that achieving desired performance in rich task areas will require the development of complicated integrated systems drawing on multiple AI paradigms. Integrated AI systems involve their own challenges; significant advances will require focused effort through initiatives aimed at integrated systems. One step towards addressing this could be increased focus on end-to-end problems and systems.

AI architectures are only mentioned under architectures for ethical AI (p. 27). The study of additional architectures, such as cognitive architectures or architectures for meta-reasoning (an area we consider worthy of additional emphasis), are important and could be mentioned.

Explainability is mentioned primarily in a single paragraph on p. 28. As explanation is now recognized as an even more important area, especially in applications where humans and computers interact or collaborate, explanation could receive more emphasis.

**Specific suggestions**

P. 7, transportation: This includes the use of AI in assistive roles but should be updated to include fully autonomous vehicles.

P 10, medicine: AI can also contribute to capture and mining of electronic health records and proactive health applications to encourage/support healthy behaviors.

P 20, fostering research on human-like AI: This section mentions important aspects such as explanation, but exactly what is meant by “human-like” in the first two sentences could be clarified. AI systems may work very differently than humans but still have the need to interact and communicate with them. Also, it could be clarified whether “human” is meant to imply human-like capabilities, or human-like reasoning processes, or both. We believe both of these are important and worthy of support.

P 22, categories of functional role divisions between humans and AI systems: The three categories describe AI supporting the human or replacing the human. An important future fourth category would include intermediate level tasks, as well as interaction modes closer to true peer interactions, for example in human-computer collaborative teams, or interactions with fluid shifting of roles based on current needs and capabilities. We see considerable space to fill between AI systems addressing small and more menial tasks and fully autonomous AI systems, and we believe that many successful applications will fall between those endpoints.

P 30, strategy 5: We agree that large, high-quality datasets are essential to progress in modern AI, and support efforts to make sure datasets are widely available to entire research communities instead of held and controlled by companies and other individuals. In addition, large scale computation has also become essential. Deep learning research, for example, requires specialized high-end Graphics Processing Units (GPUs). Companies and premiere universities that can invest in massive GPU compute clusters have a significant advantage over small companies and academic research groups. While cloud services can provide these resources on-demand, the cost of these services can be prohibitive. We believe there is a need to study and invest in shared computational resources for AI research at the national level.

P 35-36, strategy 7: The study of the AI R&D workforce is clearly vitally important. However, making AI work in real applications will need an additional workforce component that integrates research and technical communities, e.g. smart manufacturing needs people who know manufacturing and also know AI. Understanding this need is especially important because it is less recognized. In general, we believe that meeting workforce needs will require expanding the pool of expertise in AI and STEM fields as widely as possible, by better including people of diverse, educational backgrounds, technical training, genders, socio-economic groups, nationalities, etc.