

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.



October 26, 2018

Ms. Kamie Roberts
Director, National Coordination Office
Networking and Information Technology Research and Development Program
4121 Wilson Boulevard
Arlington, VA 22230

Comments submitted at: [Regulations.gov](https://www.regulations.gov)

Re: Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan

Ms. Roberts:

AMIA appreciates the opportunity to comment on Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan (AI R&D Plan).

AMIA is the professional home for more than 5,500 informatics professionals, representing frontline clinicians, researchers, public health experts, and educators who bring meaning to data, manage information, and generate new knowledge across the research and healthcare enterprise. As the voice of the nation's biomedical and health informatics professionals, AMIA provides a forum for members to evaluate health informatics interventions, innovations, and public policy across care settings and patient populations. AMIA also supports a vibrant and growing data science community within our membership who have expertise in artificial intelligence and machine learning methodologies and tools as applied to biomedicine and care delivery.

We commend the administration for designating Artificial Intelligence, Quantum Information Sciences, and Strategic Computing, as Administration R&D priorities in its FY20 Budget¹ and for chartering a Select Committee on Artificial Intelligence under the National Science and Technology Council.² We see these as necessary and important steps towards ensuring that the federal government takes a coordinated approach to AI R&D.

As the Select Committee considers its charge to coordinate AI activities across the Executive Branch, AMIA recommends it takes the following actions related to the AI R&D Plan:

- Proceed with developing an implementation framework based on the existing strategic areas and associated aims, supplementing the existing strategy;
- Emphasize AI R&D investments in areas of strong societal importance that are not aimed at consumer markets, especially relating to human-AI collaboration and the ethical, legal, and societal implications of AI with additional emphasis on malicious AI;

¹ Office of Management and Budget. FY 2020 Administration Research and Development Budget Priorities. July 31, 2018. Available at: <https://www.whitehouse.gov/wp-content/uploads/2018/07/M-18-22.pdf>

² Charter available at: <http://bit.ly/2z68RI6>

October 26, 2018

- Direct cabinet-level Departments produce annual reports articulating how their AI R&D efforts contribute to the federal AI R&D strategy; and
- Study the national landscape for creating and sustaining a healthy AI R&D workforce with an emphasis on both general-purpose AI R&D experts as well as domain-specific AI R&D experts.

Below, we offer details and rationale for these high-level recommendations. We appreciate NITRD's work in this important area, and we are eager to work with the Select Committee to bring the expertise of health informatics professionals to this national priority. Thank you for considering our comments. Should you have questions about these comments or require additional information, please contact Jeffery Smith, Vice President of Public Policy at jsmith@amia.org or (301) 657-1291. We look forward to continued partnership and dialogue.

Sincerely,

Douglas B. Fridsma, MD, PhD, FACP, FACMI
President and CEO
AMIA

(Enclosed: Detailed comments related to Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan)

October 26, 2018

After reviewing the 2016 AI R&D Plan,³ **AMIA strongly supports the existing strategic areas and associated aims. Therefore, we recommend that any updates to the 2016 strategy be informative to the existing R&D strategic areas, rather than a wholesale rework.** We note the AI R&D Plan rightly articulates a need to “emphasize AI investments in areas of strong societal importance that are not aimed at consumer markets,” and Strategy 1 correctly identifies the imperative to make long-term investments in AI research with federal funds. AMIA shares this view of the federal government’s role in AI R&D and we encourage the Select Committee to ensure that AI R&D funding prioritizes public good over private sector profit – which should be a secondary (rather than primary) byproduct of federal investments. Specifically, we encourage Federal investment where the market is less likely to focus: Strategy 2, Human-AI Collaboration, and Strategy 3, Ethical, Legal, and Societal Implications of AI.

Additionally, **AMIA recommends the Select Committee deliver on the 2016 AI R&D Plan’s two primary recommendations to (1) develop an AI R&D implementation framework and (2) study what is needed to sustain a robust AI R&D workforce.** The 2016 Plan articulated a vision that must be implemented, and the Select Committee should have a framework to understand how federal investments in AI R&D have aligned with the 2016 Plan since its release, as well as how such investments should be modified moving forward.

Focus AI R&D Strategy on Strategies 2 and 3

In medicine, we tend to frame AI as “augmented intelligence,” given that there is surely no better example of a scientific discipline so enmeshed with and influenced by the human condition. Given this view, the art and science of medicine will surely be impacted greatly by AI. Questions regarding how clinicians interact with AI or how AI will influence clinical decision-making represent daunting challenges for which federal R&D funding should be leveraged. We strongly support Strategy 2, “Develop Effective Methods for Human-AI Collaboration,” and we recommend the Select Committee ensure this area of the strategy is prioritized.

AI also presents certain societal challenges, as noted in the recommendations established by the House Subcommittee on Information Technology, Committee on Oversight and Government Reform in its September 2018 report, “Rise of the Machines.”⁴ In particular, we attach significance to the observations on bias and on possible malicious use. When AI is constructed by abstraction from ordinary human decisions, there is a danger that biases present in those decisions may survive the translation process and be built into the AI itself; credit decisions in banking may be such an example. And while the danger that AI may be used for nefarious purposes is no different from any other technology, the possibility that a malicious goal may be built into an AI program, leading it to make subtle undermining decisions (e.g. in the field of security) is of a different nature and should

³ 2016 National Artificial Intelligence Research and Development Strategic Plan. October 2016. Available at: https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf

⁴ Subcommittee on Information Technology Committee on Oversight and Government Reform. U.S. House of Representatives. Rise of the Machines Artificial Intelligence and its Growing Impact on U.S. Policy. September 2018. <https://oversight.house.gov/wp-content/uploads/2018/09/AI-White-Paper-.pdf>

October 26, 2018

remain a source of concern. Methods and tools for building AI free of biases and ethical AI that can explain and justify its decisions should be promoted.

These two broad categories of inquiry – human-AI collaboration and ethical, legal, and societal implications of AI – are fundamental to the operationalization of AI and they represent areas where private sector investment will surely be lacking.

Develop an AI R&D Implementation Framework

As an initial step to facilitate this implementation framework, AMIA recommends that cabinet-level Departments produce annual reports articulating how their AI R&D efforts contribute to the Strategic Plan. These reports should articulate not only where investments are being made, but where gaps in research remain from the respective Departments' point of view. Department-level reports will provide the Select Committee with a sense of progress to-date and illuminate areas for needed guidance so key agencies can align and implement the strategy articulated in the AI R&D Plan.

For instance, the National Institutes of Health (NIH) recently finalized a Data Science Strategic Plan,⁵ wherein it articulated a set of goals and implementation tactics to modernize the NIH-funded biomedical data-resource ecosystem through “storing data efficiently and securely; making data usable to as many people as possible (including researchers, institutions, and the public); developing a research workforce poised to capitalize on advances in data science and information technology; and setting policies for productive, efficient, secure, and ethical data use.”⁶ Specifically, Goal 3 of the NIH Plan is “Support the Development and Dissemination of Advanced Data Management, Analytics, and Visualization Tools,” which includes AI, and Goal 4 is to “Enhance Workforce Development for Biomedical Data Science.”⁷ We are confident that other Executive Branch Departments have similar strategic plans in place or in development. For the AI R&D Plan to be successful, the Select Committee must understand such efforts and encourage alignment.

Further, we highlight a December 2017 report commissioned by HHS asking an independent group of elite scientists, known as JASON, to answer the question: How might AI shape the future of public health, community health, and health care delivery?⁸ This report identified six core challenges related to this question and offered several recommendations to address these challenges. It is unclear what HHS has done in pursuit of these recommendations, but we could imagine a future JASON report investigating domain-specific AI questions – at the direction of the Select Committee’s strategic plan – with more direction and potentially more impact.

⁵ NIH Data Science Strategic Plan

⁶ Ibid. Pg. 5

⁷ Ibid. Pg. 16 and 20, respectively

⁸ Artificial Intelligence in Health and Healthcare. December 2017. JASON. JSR-17-Task-002. Available at: https://www.healthit.gov/sites/default/files/jsr-17-task-002_aiforhealthandhealthcare12122017.pdf

October 26, 2018

Develop a Robust AI R&D Workforce

AMIA believes that when applied to the broad domains of health and healthcare, AI should: facilitate discovery and translation of research findings; deliver insights to improve patient outcomes; manage and prevent disease; reduce clinician and researcher burdens; and increase value (lower costs) associated with the research enterprise and healthcare delivery system. However, these goals will not be realized without an educated and trained workforce.

The 2016 AI R&D Plan recommended that a study review the national landscape for creating and sustaining a healthy AI R&D workforce. AMIA recommends the Select Committee initiate this study, and we recommend that both domain-independent and domain-specific workforce training be examined.

While we acknowledge that the AI R&D Plan is sector-agnostic, we would encourage the Select Committee to develop a workforce strategy that promotes general-purpose AI R&D experts as well as AI R&D experts with more specialized knowledge in a specific area of application. Domain-independent AI experts will be necessary to advance R&D, but as is the case with a strategy that goes unimplemented, research that goes unapplied is useless. Specific domains of application will likely have drastically different challenges associated with the complexity of its data and the complexity of the scientific inquiry. As is the case with Biomedical Data Science, we believe the best AI R&D work will come from those scientists who are knowledgeable in specific domains and are able to make appropriate decisions about how to apply their craft.⁹

⁹ Altman R. What is Biomedical Data Science and Do We Need an Annual Review of It? Annual Review of Biomedical Data Science. Vol. 1:i-iii. July 2018. <https://doi.org/10.1146/annurev-bd-01-041718-100001>