

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.



SOFTWARE & INFORMATION INDUSTRY ASSOCIATION

COMMENTS

ON THE

2016 NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN

OCTOBER 26, 2018

The Software & Information Industry Association (SIIA) is the principal trade association for the software and digital content industries worldwide. The association provides global services in government relations, business development, corporate education, and intellectual property protection to its members, the leading companies that are setting the pace for the digital age.

SIIA appreciates the opportunity to comment on whether the [2016 National Artificial Intelligence Research and Development Strategic Plan](#) should be revised and, if so, the ways in which it may be improved. The [request for information](#) notes that comments could include suggestions as to the addition, removal, or modification of strategic aims.

SIIA agrees with the goal of creating, maintaining and updating a strategic plan to coordinate Federal government AI funding programs to achieve a urgent and significant national priorities. We further agree that Federal AI funding should supplement private sector research and application development to fulfill needs that the private cannot or should not satisfy. As the 2016 strategic plan says on p. 15:

“The Federal government is the primary source of funding for long-term, high-risk research initiatives, as well as near-term developmental work to achieve department- or agency-specific requirements or to address important societal issues that private industry does not pursue. The Federal government should therefore emphasize AI investments in areas of strong societal importance that are not aimed at consumer markets—areas such as AI for public health, urban systems and smart communities, social welfare, criminal justice, environmental sustainability, and national security, as well as long-term research that accelerates the production of AI knowledge and technologies.”

SIIA also agrees with the additional comment on p. 15 that the strategic plan should not mandate specific agency research plans.

“The AI R&D Strategic Plan, however, does not define specific research agendas for individual Federal agencies. Instead, it sets objectives for the Executive Branch, within which agencies may pursue priorities consistent with their missions, capabilities, authorities, and budgets, so that the overall research portfolio is consistent with the AI R&D Strategic Plan.”

Moreover, SIIA strongly agrees with the optimistic vision of AI that underlies the national strategy. AI is indeed poised to increase economic prosperity, to improve educational opportunity, to upgrade the quality of life for ordinary citizens and consumers and to enhance national and homeland security. Federal

AI research should supplement and promote these beneficial uses of AI.

Other countries have developed their own national strategies for the promotion and development of AI for the public good and the U.S. urgently needs to focus its own governmental research programs in the areas of greatest national need and greatest potential for public good. We encourage the administration to continue the largely successful strategic plan.

SIIA has two recommendations for updating and improving the strategic plan.

First, the strategic plan should include guidance to Federal agencies that their research programs on AI should explicitly call for an assessment of potential impacts on jobs.

Second, the plan should provide guidance that each Federal research program contain its own domain-specific design for ethical assessment as an intrinsic part of the research program itself.

In addition, the Administration should embrace ancillary objectives that have the effect of bolstering the national plan. The Administration should favor full and timely government funding and public-private partnerships for urgent AI research priorities. In particular, the Administration should commit itself to full funding of research on quantum computing.

SIIA discusses these recommendations in more detail below.

An Updated Strategic Plan Should Call for Assessment of AI on Jobs

The 2016 strategic plan explicitly says (on pp. 15-16) that it will not consider the potential impact of AI

“The AI R&D Strategic Plan also does not set policy on the research or use of AI technologies nor does it explore the broader concerns about the potential influence of AI on jobs and the economy. While these topics are critically important to the Nation, they are discussed in the Council of Economic Advisors report entitled “Is This Time Different? The Opportunities and Challenges of Artificial Intelligence.” The AI R&D Strategic Plan focuses on the R&D investments needed to help define and advance policies that ensure the responsible, safe, and beneficial use of AI.”

The only reason proffered for this omission is that the topic was covered in a separate report from the Council of Economic Advisors. But this CEA report is now over two years old and much research has been done since then on the important question of how AI will affect the number and the design of jobs. There is a serious on-going debate over whether job displacement associated with the introduction of AI into the workplace will be counterbalanced by an increase in employment opportunities elsewhere in the economy and if so how quickly.

Moreover, a government study on the issue does not focus the attention of Federal research entities on this crucial national issue. A crucial element in developing a coherent national AI strategy is to ensure that the impact on the economy and jobs is part of the research agenda of all Federal AI research programs.

There is a tension in the 2016 strategic plan between the explicit refusal early in the report to provide guidance to government agencies focusing AI research on job implications and an indication later in the

report (p. 19) that the plan will discuss how to mitigate potential disruption and risk:

“AI technologies also present risks, such as the potential disruption of the labor market as humans are augmented or replaced by automated systems, and uncertainties about the safety and reliability of AI systems. Subsequent sections of this AI R&D Strategic Plan discuss high-priority, strategic areas of AI R&D investments that will support this vision, while mitigating potential disruption and risk.”

It is impossible to have an AI strategic plan that promises to “mitigate potential disruption and risk: while ignoring the risk of job displacement.

We strongly recommend that the strategic plan be revised to resolve this tension in favor of assessing potential disruption and risk of job displacement. SIIA recommends that the national plan should include guidance to Federal agencies that their research programs on AI should explicitly call for an assessment on potential impacts on jobs.

An Updated Strategic Program Should Favor Domain-Specific AI Ethical Assessment

The 2016 strategic plan includes a strategic goal (p. 3) of focusing on the ethical implications of AI:

“Strategy 3: Understand and address the ethical, legal, and societal implications of AI. We expect AI technologies to behave according to the formal and informal norms to which we hold our fellow humans. Research is needed to understand the ethical, legal, and social implications of AI, and to develop methods for designing AI systems that align with ethical, legal, and societal goals.”

This is an important part of the national strategy and should be maintained in any revision. There are many broad ethical principles to guide the development and use of AI. SIIA has its own set of ethical principles for AI and advanced data analytics ([link](#)) and just recently ACM (Association for Computing Machinery) updated their code of conduct for computer science professionals ([link](#)) for the first time in 20 years to take into account the new challenges of AI.

But the consensus of AI practitioners and experts working in the area is that abstract principles can go only so far. The key ethical decisions are heavily context-dependent and specific to the domain of application. This has been well-understood and acknowledged in the policy sphere, where universal AI policy principles are less attractive than adapting domain-specific policy approaches to the new challenges of AI. For instance, the 2016 report on the state of AI research and policy by the [AI Study Group](#), a panel of industry and academic experts, endorsed the “sector-specific” approach to AI policy currently in place in the U.S. and warned that:

“attempts to regulate “AI” in general would be misguided, since there is no clear definition of AI (it isn’t any one thing), and the risks and considerations are very different in different domains.

In the same way, research into ethic and AI should be domain specific, and not just a separate AI research area.

Unfortunately, much of the discussion in the 2016 strategic plan seemed to suggest an abstract approach. It describes program of investigation into how to make AI systems as such conform to ethical principles in the abstract, that is, how to incorporated ethics in general into AI in general. In doing so, it focuses on

several high-level approaches that might develop a universal ethical framework that could be incorporated into each AI application.

It is clear however that the different domains of AI application such as autonomous weapons, autonomous cars, and AI-driven diagnosis and treatment raise different issues policy and ethical issues that need separate and specific assessment.

It is true that AI raises some unique issues such as the difficulty of providing explanations of machine learning predictions that people can understand and use. But this difficulty cannot be approached solely by developing tools and techniques to make machine learning algorithms in general more understandable in general. What type and scope of explanation is needed differs depending on whether the applications are focused on marketing, granting credit, diagnosing diseases, teaching students, or assigning liability for a wrongful combat death.

SIIA recommends that the strategic plan be revised to provide guidance to Federal research agencies that their AI research programs should contain their own element of ethical assessment as an intrinsic part of the programs themselves. In this way, ethical issues can be designed into systems at the beginning, as part of building the system itself, rather than as part of an external evaluation after the fact.

For instance, a research program that calls for AI-systems that can aid areas such as education, security and law enforcement, social welfare and criminal justice should provide that grantees in the area explicitly conduct a disparate impact analysis to determine whether the applications or the research itself would have a disproportionate adverse impact on legally protected classes, and if so whether there are alternative applications that could achieve the same results with less discriminatory impact. Each of these subject matter areas are the sites of significant ethical controversies in which the fairness of algorithms used to achieve valuable public purposes have been called into question.

SIIA Urges the Administration to Support Full Funding for Quantum Computing.

Quantum computing is a major new development in the computer science field. It is widely expected to be operational within 5 to 8 years and will be a disruptive force in every domain from cryptography to AI. The national AI research strategy should instruct agencies to prioritize projects that foster the development of this breakthrough technology.

Congress is supportive of additional research funds. The House recently passed [H.R. 6227, the National Quantum Initiative Act](#) that outlines a 10-year program to set goals, fund research, promote coordination between different agencies, and create partnerships between industry and academia. The program would allocate \$400 million to the National Institute of Standards and Technology, \$625 million to the Department of Energy, and \$250 million to the National Science Foundation to train scientists, expand research, and establish up to 10 quantum research and education centers.

This initiative could not be more important. China, Europe, and Russia are investing heavily in quantum research and technology development. The United States must take concrete steps to achieve and maintain global leadership in this area. The application of quantum computing to AI could enable a breakthrough to achieve genuine autonomous agents that can augment human capabilities in important areas of national defense, manufacturing, education and medicine.

The [Center for a New American Security](#) recently emphasized the strategic importance of funding research

in quantum computing. It examined the national security issues raised by the rise of China to prominence in quantum computing and recommended that the United States “must ensure that basic and applied research and development in quantum science and technology receive adequate, sustained funding, while seeking to attract and retain top talent.”