CCC

Computing Community Consortium
Catalyst

# The Computing Community Consortium's Response to the Office of Management and Budget's Request for Comments on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum

**December 5, 2023**

*Written by: David Danks (University of California, San Diego), Haley Griffin (Computing Community Consortium), David Jensen (University of Massachusetts Amherst), Chandra Krintz (University of California Santa Barbara), Daniel Lopresti (Lehigh University), Rajmohan Rajaraman (Northeastern University), Matthew Turk (Toyota Technological Institute at Chicago), and Holly Yanco (University of Massachusetts Lowell)*

**This response is from Computing Research Association (CRA)'s Computing Community Consortium (CCC). CRA is an association of nearly 250 North American computing research organizations, both academic and industrial, and partners from six professional computing societies. The mission of the CCC, a subcommittee of CRA, is to enable the pursuit of innovative, high-impact computing research that aligns with pressing national and global challenges.**

**Please note any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the authors' affiliations.**

This response pertains to questions 1-6, and 8 (each numbered and copied below in italics) from the Request for Comments:

*1. The composition of Federal agencies varies significantly in ways that will shape the way they approach governance. An overarching Federal policy must account for differences in an agency's size, organization, budget, mission, organic AI talent, and more. Are the roles, responsibilities, seniority, position, and reporting structures outlined for Chief AI Officers sufficiently flexible and achievable for the breadth of covered agencies?*

We strongly endorse the notion that the Chief AI Officers should be at the senior level in order to maximize their impact. They will be responsible for helping to shepherd our government through one of the most significant technological transitions in our nation's history. Our impression is that it will be a dual-hatted position, so we also emphasize that they should be provided with all necessary resources, including sufficient staffing and a highly qualified deputy with expertise in a wide range of emerging AI topics.

*2. What types of coordination mechanisms, either in the public or private sector, would be particularly effective for agencies to model in their establishment of an AI Governance Body? What are the benefits or drawbacks to having agencies establishing a new body to perform AI governance versus updating the scope of an existing group (for example, agency bodies focused on privacy, IT, or data)?*

A new body to perform AI governance is the right approach because there is danger in following the status quo given the speed and scale of AI developments today. It is essential to look for new ways of thinking, and this can be actualized through AI Governance Bodies that consist of individuals who are studying pressing AI issues both broadly and deeply. We also propose that the AI Governance Bodies at each agency have several members who are Special Government Employees (SGEs), bringing insights from outside of each agency while still understanding key aspects of each agency's mission. In particular, we encourage the Bodies to include SGEs such as researchers participating in NSF AI Institutes, professional organizations with long-standing expertise in AI including the Computing Research Association (CRA), Association for the Advancement of Artificial Intelligence (AAAI), Institute of Electrical and Electronics Engineers (IEEE), Association for Computing Machinery (ACM), etc., and members of potentially impacted communities. Given the breadth of AI's impact, interdisciplinary contributions from a wide range of fields will be critical in order to provide appropriate perspectives about research and impact on communities. AI Governance Bodies should not only include experts in computing, but also those with expertise in other disciplinary approaches, such as health and human services, transportation, and the social sciences.

Additionally, we believe that there needs to be a transparent mechanism for interactions across AI Governance Bodies of different agencies. While AI Councils are mentioned as a potential vehicle for coordination, the draft policy lacks a description of how the collaboration between these bodies will take effect. We believe such collaborations will be vital for reducing duplication of effort and, ultimately, for success of this effort, but these benefits will not be realized without careful design.

*3. How can OMB best advance responsible AI innovation?*

It is essential for responsible AI innovation advancement to…

- Bring in non-agency and non-vendor expertise, perhaps through outside individuals on advisory boards for agencies, both technical experts and average citizens. The work of such boards should be observable and transparent where possible (so that the public can give feedback as needed).
- Prioritize that the outcomes from using AI are trustworthy, fair, and reliable. While using good, clean data is important in developing AI systems, the results and impacts on individuals and communities need to be top of mind especially since even good data can be misused.
- Implement monitoring mechanisms that recognize that even if a large number of individuals have used an AI system over a period of time and have not yet reported any issues, that does not mean that the AI system is safe or reliable. Rare but impactful edge cases are common across all fields of computing.
- Reevaluate an AI/ML system each time a software update is implemented, or a dataset is changed. We recognize that there are difficult questions about exactly when a software or dataset update is significant enough to require re-evaluation. We do not propose a hard-and-fast rule, but we contend that any such criterion should err on the side of too much re-evaluation, rather than too little.
- Encourage independent evaluation and testing of all AI/ML systems so as not to become dependent on vendors who may employ differing criteria.
- Actively plan and prepare for unforeseen deleterious impacts because the behavior of many AI systems remain unpredictably complex.
- Release metrics and reasons for determinations that an AI is neither “safety impacting” nor “rights impacting.” The public should have access to these measures given the significant incentives for companies to claim their systems are neither safety- nor rights-impacting in order to minimize scrutiny.
- Use a more fluid characterization of an AI system, as the impact of such technology can change rapidly and needs to be reconsidered often.
- Encourage and support public scrutiny of the AI use case inventories that Agencies are releasing. There is always a possibility they are choosing to display the work they know will please the public instead of a random/average sample.
- Require that AI use cases be published with a sufficient level of detail and, where appropriate, representative data, so that outside independent experts can evaluate potential risks and provide constructive feedback for ongoing improvement of AI deployments.

*4. With adequate safeguards in place, how should agencies take advantage of generative AI to improve agency missions or business operations?*

At this early stage in the development of generative AI, we urge a cautious, carefully considered approach to employing generative AI in agency missions and business operations. We applaud current research on generative AI, and we are optimistic about the tremendous possibilities that it offers, but caution should be taken. Systems and processes built around generative AI should be prepared to deal with errors and unanticipated outputs. The private sector is going through processes of trial and error in the generative AI space, and will continue to do so for the immediate future; there is wisdom in letting them work through the growing pains of this new technology before taking any large scale measures as a government. Additionally, at this time business operations can oftentimes be more improved by non-generative AI solutions; agencies should not presume that generative AI systems are the only, or even the best, way to advance their missions and operations.

*5. Are there use cases for presumed safety-impacting and rights-impacting AI (Section 5 (b)) that should be included, removed, or revised? If so, why?*

The following challenges should be considered when categorizing use cases:

- *The challenge of AI software components* — One fundamental challenge is that a vast number of computer systems could be deemed to be "AI" because one (perhaps largely incidental) component of that system uses some form of AI technology. For example, consider
  - a video-conferencing system that uses a computer vision model to identify the background of an image (to add blur to that background).
  - an email system that uses a simple Bayesian classifier for spam detection.
  - a word processor that uses a simple language model for autocomplete.
- *The challenge of update frequency* — Software components based on machine learning have the potential to dramatically increase the frequency, scope, and potential impact of updates to software that is routinely used by agencies.

*6. Do the minimum practices identified for safety-impacting and rights-impacting AI set an appropriate baseline that is applicable across all agencies and all such uses of AI? How can the minimum practices be improved, recognizing that agencies will need to apply context-specific risk mitigations in addition to what is listed?*

There should be an open mechanism for affected groups/users of an AI system to have a redress path that leads back to the appropriate AI experts in the agency. When a decision is being made by a human, there is a human to complain to, and a process for the complaint to be taken up the chain of command as needed. It is crucial that especially for rights or safety-impacting AI, that there is a clear process for perceived harms to be documented and addressed.

On a similar note, if possible there should be a notice to an individual if they were negatively impacted by an AI system. If an error was made and is being fixed, they should be updated on this change. This is similar to what happens today when there is a data breach and private information is stolen from a website.

Additionally, in regards to data usage in these AI systems, legacy data should automatically be considered suspect until proven otherwise since so much data from the past contains inherent biases, as has been demonstrated through numerous published studies. All data used to develop AI needs to newly pass the criteria outlined in the minimum practices, just like the AI systems themselves will have to do.

*8. What kind of information should be made public about agencies' use of AI in their annual use case inventory?*

Given the huge impact AI can have on our society and competing interests, many of which are powerful and well-funded, special efforts will be needed to make sure the public receives a true and accurate picture of what AI can and cannot do, and how it is being applied in government. So use cases should include an interpretation, explanations, and examples that the average citizen can follow, understand, and question when necessary. There should also be a process to make sure that the case the agency is providing is either random or average so they can not choose to just display their best examples.