December 5, 2023

OMB RFI: 88 FR 75625
Organization: The Center for Security and Emerging Technology (CSET)
Respondent type: Organization>Academic institution / Think tank
Primary POC: Tessa Baker, Director of Communications (tgb5@georgetown.edu)

The Center for Security and Emerging Technology (CSET) at Georgetown University offers the following comments in response to OMB's Request for Comments on the *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum.*

A policy research organization within Georgetown University, CSET provides decision-makers with evidence-based analysis on the security implications of emerging technologies, focusing on artificial intelligence, advanced computing, cybersecurity, and biotechnology. We appreciate the opportunity to offer these comments. This response was compiled by the Primary POC, but represents contributions from Mia Hoffmann, Mina Narayanan, Dahlia Peterson, Emelia Probasco, Jack Corrigan, William Hannas, Huey-Meei Chang, Dewey Murdick, Margarita Konaev, and others.

We address OMB's eight questions below. In addition to OMB's efforts to enable AI adoption within the USG, we think it's important that the USG devise a resourced plan to actively monitor the adoption of these technologies by other governments (friend and foe) to ensure we can learn from others as well as ensuring that our technological approaches (and security practices) are calibrated to the potential offensive use of these technologies by other state and non-state actors.

Question 1: Chief AI Officer Roles and Responsibilities

As described in the draft memo, the roles, responsibilities, seniority, position, and reporting structures for CAIO's appear sufficiently flexible. However, the anticipated breadth of responsibilities is vast. Depending on the size of the agency, a single individual or small team may struggle to engage effectively on all of the anticipated activities, which include coordinating agency use of AI (and reporting), communicating the value of AI to each agency's mission, promoting innovation, and managing risks.

Additionally, the absence of clear budget or approval authority may hamstring CAIO's from ensuring compliance with the EO, preventing risky use cases, or monitoring adoption across agencies. We recommend:

- Requiring CAIOs at Chief Financial Officer (CFO) Act agencies, at minimum, not be dual-hatted with other significant roles like CTO or CIO given the breadth of their responsibilities.
- Include AI use cases and relevant categorical information in future updates to Technology Business Management (TBM) Taxonomy or other IT taxonomies and include clear instructions on how agencies should report on investments in AI for mission delivery in future revisions of A-11 Guidance (Section 55) and/or A-130 Guidance, to enable better tracking of investments in AI technologies.
- Requiring CAIOs be involved in agency Capital Planning and Investment Control (CPIC) procedures and milestone reviews in order to enable visibility into agency projects of sufficient scale that may already leverage AI or might benefit from AI advancements. In addition:
    - Strongly consider granting CAIOs approval authority for new CPIC investments in AI-related projects to ensure that they are involved in all aspects of project development and implementation.
    - Consider integrating agency AI use case inventories into existing dashboards like ITdashboard.gov so that both use cases and funding/investment levels can be interrogated by the public.
- Consider revising agency guidance in A-123 or comparable instructions to ensure CAIO's risk review responsibilities re: AI are codified and incorporated into Enterprise Risk Management Procedures to strengthen the CAIO's position and ensure appropriate internal agency governance.

Question 2: Coordination Mechanisms for AI Governance Bodies

In light of the desired expediency in implementing EO 14410, convening new governance bodies rather than repurposing or leveraging existing governing bodies seems likely to increase the administrative burden and delay meaningful results.

We recommend:

- To the extent that agencies have already established effective Data Governance Bodies, as contemplated in M-19-23, leveraging these or similar CPIC or IT-related governance bodies rather than creating standalone AI Working Groups would be a good starting point. Since AI is an enabling tool that requires data, compute, and people to be effective, piggybacking on existing governance infrastructure would enable CAIOs to more easily engage with relevant stakeholders, including Chief Data Officers, Chief Human Capital Officers, Chief Information Officers, and Chief Technology Officers.
- Similarly, we'd also recommend that existing infrastructure, like the CIO Council, serve as a launching point for any contemplated interagency coordination of AI activities, rather than creating standalone coordination and clearinghouse functions on day one in addition to the AI Development Council (EO14410, Section 10.1.a), AI Technology Task Force (EO14410, Section 10.2.b), and the White House AI Policy Council (EO14410, Section 12).

Question 3: OMB and Responsible AI Innovation

AI offers tremendous opportunity, but it may not be the solution for every mission or every government agency, and it requires people to be able to discern use cases and mitigate implementation hurdles. In order to best enable responsible AI innovation, OMB can do two things now:

- Prioritize Workforce Development and Recruitment/Retention: Section 10.2 of EO 14410 demonstrates the administration's commitment to talent as a critical enabler of the potential of AI for government. Hiring, (re)training existing employees, providing relevant career tracks for existing and new technical talent, and identifying gaps in capability are essential for the long-term, sustainable implementation (and maintenance) of AI systems in agencies. Without the talent to execute this memorandum or the EO, federal efforts to regulate and deploy AI systems may falter.
    - As noted in Section 4.b.iv of the draft memorandum, both technical and non-technical workers are essential to AI implementation. CSET research demonstrates the vital importance of not just the data scientists, engineers, or IT

workers who can develop and fine-tune AI, but also those, including contracting officers and product managers that enable AI adoption.

- Focus Agencies on Requirements, Not Hype: OMB can ensure agencies appropriately adopt AI for relevant mission areas by encouraging agencies to first identify capability needs, including legacy IT retirement/modernization requirements. By focusing first on requirements, we are hopeful that agencies can steward the public's trust and resources to the areas where AI might be most helpful, while also protecting the ability of agencies to continue executing their missions.

Question 4: Leveraging Generative AI (safely) to Accomplish Agency Missions

Given the varied missions of USG agencies, we cannot prescribe specific use cases in this response. However, we would like to emphasize the importance of clear requirements, establishing anticipated concepts of operations, and ensuring issues and harms are reported appropriately and promptly to support the responsible use of these technologies. We believe agencies should have clearly defined use cases for incorporating AI into their missions or business operations, be it generative AI or topic modeling **before** adopting these tools. If agencies determine that their goals are best served by taking advantage of AI, they should be ready to clarify the goal of the AI in context of use, how the goal aligns with agency objectives, how the AI achieves the goal, and who – if anyone – is tasked with enabling the AI to achieve the goal. NIST AI RMF Profiles may be a helpful starting point.

Question 5: Additional Use Cases for Consideration

The AI applications that are deemed safety-impacting and rights-impacting in Section 5 appear to constitute a fairly comprehensive list. We recommend:

- Add to Safety Impacting list (Memo Section 5.b.i): AI used to control/meaningfully impact the function of critical digital infrastructure (including emergency broadcasting systems and critical information delivery platforms);
- Move from Safety (Section 5.b.i) to Rights-Impacting (Section 5.b.ii): The functioning of ….”integrity of elections and voting infrastructure” because the draft memo specifies additional impact and rights assessments for rights-impacting technologies, including user testing and public comment (Section 5.c.iv).
- Add additional use cases to Rights-Impacting technologies, including:
    - ii.B (Law Enforcement or Surveillance-related): Identifying or predicting location of crime through techniques like gunshot detection algorithms;

- ii.C (Deciding immigration, asylum, or detention status): Verification of travel documents and migration analytics, including the forecasting of migration patterns and/or border crossings.
- ii.E (Education): Assessing student performance and learning outcomes.
- ii.G (Employment): Real-time task allocation/scheduling.

Question 6: Minimum Practices for Safety and Rights Impacting AI

The proposed practices and requirements are an excellent starting point, we recommend the following additions and modifications:

- Impact and Risk Assessments (iv.A.2) should (specifically) include evaluations of forces outside of an AI system, such as its suppliers or downstream applications creating undesirable impacts or negative system effects, and should include evaluating the cybersecurity of the AI system.
- Data Assessment (iv.A.3): Minimum documentation requirements should also include:
  - a description of the data collection process, and steps taken to prepare or pre-process the data;
  - Any assumptions made about the information the data represents; and,
    - For example, if healthcare expenses are used as a proxy for healthcare needs, this reflects an underlying assumption that all people seek out medical care at the same rates.
  - A description of relevant data gaps and how they are addressed.
- Monitoring (iv.D): Agencies (or ideally, OMB) should set up a mechanism to identify, record and report AI incidents. Incidents can be realizations or near-realizations of previously identified or unforeseen risks and can emerge as a result of AI system behavior or failure of human-AI-teams, among other reasons. Incidents should be investigated to determine their cause and the resulting findings used to improve risk mitigation strategies across the agency going forward. While the authorization and responsibility for this decision may lie with the CAIO, a clear, mandatory reporting and accountability structure should facilitate the identification and reporting of those conditions by staff directly interacting with the AI systems.
- Mitigation Approaches (iv.E): Before deploying an AI, agencies should define circumstances and conditions that would trigger immediate review and/or suspension of the AI system and prepare a response plan, especially for AI systems that are tightly integrated with other entities. Among the considerations should be the criteria under which any fail-safe or override modes are triggered, as well as the criteria under which

the system should be recalled or shut down. Agencies should also document the circumstances in which changes to an AI system or its operating environment necessitate reassessment or re-testing.

In addition, we'd recommend providing clear policy guidance to all agencies regarding the proper supervision and control of AI deployed at agencies, including minimum training on standards for employees involved in adjudicating AI recommendations. In particular, for systems that are safety- or rights-impacting it is prudent to assess both AI and human capabilities and limitations (together) for a given AI use case and implement comprehensive risk mitigation measures rather than evaluating the AI system in isolation. It is also worth considering standards for enabling timely intervention mechanisms and override authority in cases where AI failure is causing unacceptable harm. DoD policies on autonomous weapons systems (DODD 3000.09) provide a helpful model.

Question 7: Materials and Resources To Enable Agency Implementation in Contracts

We suspect that agencies will need to reallocate funding and/or deprioritize existing efforts to meet the aggressive timelines contemplated in this memorandum and the attendant EO. Guidance on flexibilities in existing services contracts, to enable agencies to re-task existing contractors would seem helpful. Additionally, templated language regarding minimum AI system compliance standards (as documented in 5.iv of the memo) for vendor contracts and GSA-established contracting vehicles (like IDIQs or BPAs) that support CAIO and other task execution would be useful.

Question 8: Key Data Elements for Annual Agency AI Use Case Inventories

We believe agency use cases should document:

- **Intended Goal:** the goal of the AI in context of use and how it would improve the status quo, how the goal aligns with agency objectives, how the AI achieves the goal, and who (if anyone) is tasked with enabling the AI to achieve the goal.
- **Anticipated Concept of Operations (CONOPs):** More detailed documentation can stem from agency goals, such as clarifying whether the AI system will interact with people in pursuit of agency objectives or unpacking how the AI achieves its goals by documenting its tasks, methods, and the data it relies on.
- **Cost and Resourcing Information:** As stated elsewhere, including budget and resourcing information in public disclosures would permit the public (and researchers

such as ourselves) to evaluate the relative efficacy of AI use cases to improve the efficiency and performance of government services.

    - **Staff Profiles and Contractor Counts:** Descriptions of the skills and qualifications of people charged with implementing AI systems would be helpful for transparency and to inform AI workforce development.

- **Risk Rating:** Use cases should also include a description of risks that AI systems pose and a grading of these risks so it is easy to determine which AI use cases in an agency are higher risk than others – the existing IT Portfolio Management Risk categories may be a useful starting point for this categorization.
- **AI Incidents:** As stated above, effective monitoring will require capturing any instances of AI harms or near-misses to ensure the USG and public have a comprehensive view into the kinds of issues emerging from the wholesale adoption of AI-technologies in government.

From an AI assessment perspective, all of this information will make it easier to evaluate these systems and make sure they are working as intended.

Additional Assistance

Additional CSET research on workforce, AI assessment, and cybersecurity, may also be helpful as agencies implement the EO and associated guidance. CSET is happy to brief on our prior work and engage, as helpful, with the USG as they undertake this important work. Please reach out to the primary POC or cset@georgetown.edu for further assistance.