1350 Eye Street NW, Suite 350, Washington, DC 20005 | 202-768-8950 | fpf.org

April 29, 2024

Via Electronic Submission
U.S. Office of Management and Budget
Attn: Samantha Hubner
New Executive Office Building
725 17th Street NW
Washington, DC 20503
(202) 395-3039

Re: Request for Information: Responsible Procurement of Artificial Intelligence in Government

Dear Ms. Hubner,

The Future of Privacy Forum (FPF)[1] welcomes the opportunity to submit comments in response to the Office of Management and Budget's (OMB) Request for Information regarding responsible procurement of artificial intelligence (AI) in government, particularly regarding the intersection of the procurement of AI tools and systems[2] with other risks posed by the development and use of AI tools and other emerging technologies.[3]

As one of the largest purchasers of AI tools and systems, the U.S. government has the power to set procurement policies with respect to AI that can become potent drivers for privacy, transparency, and equitable outcomes.[4] Federal agencies are also obligated to ensure their use of AI tools and systems is lawful, ethical, and privacy-protective. By establishing clear guidelines around the procurement of AI tools and systems, OMB has a material opportunity to address and mitigate potential risks to personal data that some AI tools and systems may pose, either through design or use. Successful guidance will acknowledge the dynamic, evolving nature of AI technology and take into account that such guidance may have a standard-setting impact on the development of AI systems more generally beyond the intended scope of the public sector.

---

[1] FPF is a global non-profit organization dedicated to advancing privacy leadership, scholarship, and principled data practices in support of emerging technologies. The opinions expressed herein do not necessarily reflect the views of FPF's supporters or the Advisory Board.
[2] We consider AI tools to be any AI-driven tool that can either form a system in its own right or be integrated into another system, for example some agencies may consider procuring an AI tool to add to an existing system of their own, or alternatively procure an entire AI-driven system.
[3] In response to Questions 6 and 10 in the "Request for Information: Responsible Procurement of Artificial Intelligence in Government," 89 Fed. Reg. 22196 (Mar. 29, 2024), https://www.federalregister.gov/documents/2024/03/29/2024-06547/request-for-information-responsible-procurement-of-artificial-intelligence-in-government
[4] *Recommendations: AI's Procurement Challenge*, The National Artificial Intelligence Advisory Committee, (Oct. 2023), https://ai.gov/wp-content/uploads/2023/11/Recommendations_AIs-Procurement-Challenge.pdf

Further, given the current boom in AI, it is important to consider both generative and non-generative (or so-called traditional AI) when analyzing any potential risks of procured AI systems. Generative AI has the ability to generate content, and has gained popularity in the past two years in the consumer space not least because of leaps in development and scale. However non-generative AI, including automated decision making technologies (ADMT), robotics and pattern recognition,[5] has a longer history and was already widely used across numerous data-driven use cases before the generative AI boom of 2023, and indeed may already be used by some government agencies.[6] Therefore any analysis of AI systems under procurement consideration should be made in the context of not just what type of AI system is being procured, but what data that system is using, how the underlying model is trained and what it is designed to do. Some federal agencies may be developers or deployers of AI tools or systems,[7] including but not limited to those procuring AI from third parties under existing contracts. It will be important to define both the scope of AI that is to be captured under the OMB's future procurement rules as well as defining developer and deployer obligations with care and nuance given that strict binary definitions could have unintended disruptive effects on agencies' ability to operate using existing solutions.

Given these considerations, FPF recommends:

1. OMB should ensure that contractual responsibilities and requirements for transparency, testing, evaluation, and impact assessments in procured AI systems are based on clear definitions and roles, taking into account the risk profile of the AI system;

2. OMB should ensure that agencies procure AI systems or services that meet the existing data protection standards that apply to federal agencies when they handle personal data; and

3. OMB should ensure that agencies procure AI systems or services that support, rather than undermine equitable outcomes, by requiring agencies to analyze the particular risks these systems may pose to people, especially marginalized individuals and communities.

1. OMB should ensure that contractual responsibilities and requirements for transparency, testing, evaluation, and impact assessments in procured AI systems are based on clear definitions and roles, taking into account the risk profile of the AI system.

[5] Facial Recognition is a prominent example of pattern recognition and use case of AI.
[6] AI as a technology has primarily been developed by the private sector rather than the government so there may be many instances of AI in use in government that was developed by a third party.
[7] Depending on how the government chooses to define the common terms of AI developer or deployer, some U.S. government agencies may find themselves captured by the scope of those definitions.

Transparency, testing, evaluation, and impact assessments are useful tools to hold AI developers and deployers accountable and help prevent harm in AI systems. Procurement contracts for AI should include clear and appropriate requirements for developers and deployers of AI systems while ensuring that those requirements are proportionate to the level of risk identified.

Prescribed requirements should take into account the covered entity's relationship with the AI tool or system when determining its obligations. Such an approach would recognize that developers and deployers will likely have different responsibilities and capabilities regarding operation and oversight, and that vendors and agencies may not always occupy the same roles. By calibrating obligations to an entity's status as a developer or deployer - and taking into account the specific use case or deployment environment - obligations can be tailored to the ways each entity is best situated to evaluate and address issues including non-discrimination, responsible AI governance, transparency, data security, and privacy.[8]

It is worth noting that not all vendors will be the true developers of the underlying AI model behind their AI system offerings: many contracting vendors add business value to AI models developed by third parties to create proprietary AI-driven systems. It is therefore crucial to ensure that both procurement contracts and Service Level Agreements (SLAs) for procured AI tools or systems outline clear responsibilities between the contracting parties and do not rely solely on "developer" and "deployer" terminology, but should explicitly define roles and responsibilities of named parties, including on a contract by contract and case by case basis. This is especially important where an agency foresees a high risk deployment environment or context.

Clarity in these roles and respective responsibilities will be of great importance for AI system procurement. In the case of the government, where much of the data being handled may be sensitive in nature, a developer may not have foreseen their technology being deployed in a situation that requires a higher degree of care than normal, such as processing highly sensitive data inputs that may have implications for privacy or discrimination. In such a scenario, a developer may not be positioned appropriately to mitigate contextual risks associated with that environment, compared to the deployer of the AI system. Similarly, a deployer may lack insight into an AI tool's known efficacy limits or understand how it was developed and trained, and will rely on the developer being as transparent as is reasonable about the capability of the system. Understanding the system's capabilities will be essential to the deployer's ability to identify and therefore mitigate against risks from the system's operation.

To address the information asymmetry between developers and deployers, OMB may want to consider what information should be shared between these entities in order to facilitate compliance.

---

[8] There may be cases where developers and deployers of an AI system are the same entity. In that case, the entity would need to be assessed through both lenses.

For example, *FPF's Best Practices for AI Systems Used in Employment Decisions* (chart below) focuses on the use case of AI and employment decisions, a relatively high risk use case that can have meaningful consequences for people's lives. We identified a list of disclosures that developers should make to deployers in that specific use case, and a different list of disclosures that deployers should make to individuals.[9]

Developer disclosures to deployers	**Deployer disclosures to individuals**
• the intended purposes of the AI tool • purposes for which the AI tool is not intended • known efficacy limits of the AI tool • how the AI tool was trained • whether the AI tool was assessed for potential discriminatory bias • whether the AI tool uses information from deployers or Individuals to further train or otherwise improve the tool • how the AI tool is intended to be deployed • uses of the AI tool that are not intended • what choices the AI tool provides to deployers regarding anti-discrimination, governance, transparency to Individuals, privacy, security, and human oversight • what choices the AI tool provides to deployers to communicate to Individuals about how they implement the tool, and how the tool fits into the deployer's overall decision-making processes regarding Consequential Impacts	• the fact that Individuals are interacting with an AI tool • the intended use of the AI tool (e.g., to evaluate job candidates, make compensation decisions, or consider employees for promotion) • how the AI tool was trained • how an AI tool may have a consequential Impact and how the tool fits into the deployer's overall decision-making processes • the extent to which Individuals' Personal Data is shared with third parties or used to train or improve the AI tool • what alternative options are available to all Individuals, and how Individuals with disabilities may seek accommodations

When it comes to testing, evaluating and assessing the ultimate impact of AI systems, it will be essential for the procuring agency to ensure they have explicitly apportioned respective responsibilities between the vendor (as developer or deployer or both) and themselves (as a deployer). The procuring agency will need to play some role in testing, evaluating and assessing the impact of the procured AI tool or system as it will need to be responsible for assessing the suitability of the procured tool or system for the particular environment or context insofar as is reasonable. This includes the conducting of impact assessments.

[9] *Best Practices for AI and Workplace Assessment Technologies*, Future of Privacy Forum (Sept. 2023), https://fpf.org/wp-content/uploads/2024/02/FPF-Best-Practices-for-AI-and-WA-Tech-FINAL-with-date.pdf

Vendors can further assist agencies to understand the AI tools or systems they are procuring to help establish their suitability for the use case the agency has in mind. They can also establish a formal feedback loop, such as is already well established by software as a service (SaaS) and Cloud vendors for the use of non-AI related tools and systems where the procuring party can flag bugs, as well as elevate other issues regarding the performance, accuracy and efficiency of the procured tool or system. In addition, in SaaS and cloud procurement scenarios, procuring clients typically already engage in a formal feedback loop as to whether the tool or system is fit for purpose or operating as expected. This is especially the case with bespoke client solutions, and we can assume that in many instances agencies will demand bespoke rather than off-the-shelf solutions, tailored to specific use cases. Again, the relative risk of the use case will need to be taken into account when deciding on requirements and apportioning responsibilities.

2. **OMB should ensure that agencies procure AI systems or services that meet the existing data protection standards that apply to federal agencies when they handle personal data.**

OMB should ensure that procured AI systems comply with existing federal regulations and guidance such as the Privacy Act of 1974, which establishes a code of fair information practices that governs how federal agencies collect and use records about individuals.[10] The Privacy Act and related guidance govern how agencies handle personal data, provide rights to individuals, and ensure that each agency "maintain[s] in its records only such information about an individual as is relevant and necessary to accomplish a purpose of the agency required to be accomplished by statute or by executive order of the President."[11] The Act proceeds from a premise of strong data minimization requirements,[12] prohibits unauthorized disclosures,[13] mandates appropriate administrative and technical privacy safeguards,[14] guarantees individuals the rights to access and correct information about them,[15] and requires that all records used to make a determination about an individual be accurate, relevant, timely, and complete.[16]

The Privacy Act typically applies when agencies process personal data, and some AI systems and tools procured by agencies will implicate this sort of personal information in either their development or deployment. Personal information might be used to train generative AI systems, generative AI tools might be used to create documents that include personal information, and non-generative AI technologies might be used to make determinations about individuals' qualification for government programs; each of these uses would implicate Privacy Act requirements. Given the technical complexities of AI, there can be uncertainty regarding how

[10] 5 U.S.C. § 552(a).
[11] 5 U.S.C. § 552a(e)(1).
[12] 5 U.S.C. § 552a(e)(1).
[13] 5 U.S.C. § 552a(b).
[14] 5 U.S.C. § 552a(e)(10).
[15] 5 U.S.C. § 552a(d)(1)-(2).
[16] 5 U.S.C. § 552a(e).

best to apply Privacy Act guarantees in the context of AI. For example, it is not obvious how an agency could effectively delete an individual's inaccurate personal information from a training data set that was used in the past to train an AI tool that remains operational. At the same time, the application of Privacy Act requirements to some AI use cases is clear: if an AI tool contributes to agency determinations about an individual, agencies must ensure that the data is accurate, relevant, timely, and complete. FPF therefore recommends that the procurement of third party AI systems or services by agencies should incorporate a "fitness test" that ensures procured systems and services meet relevant Privacy Act and related standards.

3. **OMB should ensure that agencies procure AI systems or services that support, rather than undermine equitable outcomes, by requiring agencies to analyze the particular risks these systems may pose to people, especially marginalized individuals and communities.**

The use of personal data in any AI system can create or exacerbate privacy risks for individuals, communities, and beyond. It is worth noting that regulators abroad have been focussed on automated decision-making tools (ADMTs) for some time, including the regulation of automated decision making under the EU's GDPR.[17]

In 2017, FPF documented a wide range of potential harms related to the risks of using personal data in automated decision-making.[18] ADMTs analyze personal data to make decisions about individuals by employing AI or other analysis techniques. These harms fall into four broad categories: loss of opportunity, economic loss, social detriment, and loss of liberty.

As OMB continues to formulate new guidance for AI and modify existing rules to take account of AI, we recommend adding a specific requirement for agencies to ensure their analysis includes the identification and analysis of risks to individuals and communities including those constituting protected classes under U.S. law. In particular, marginalized communities have been documented as being disproportionately impacted by inaccurate or biased information in commercial and government data sets,[19] which can lead to greater bias and discrimination.[20] These risks can be heightened when personal data is used as training or other inputs to systems that make

---

[17] Sebastião Barros Vale and Gabriela Zanfir-Fortuna, *Automated Decision-Making Under the GDPR: Practical Cases from Courts and Data Protection Authorities*, Future of Privacy Forum (May 2022), https://fpf.org/wp-content/uploads/2022/05/FPF-ADM-Report-R2-singles.pdf

[18] *Unfairness by Algorithm: Distilling the Harms of Automated Decision-Making*, Future of Privacy Forum (Dec. 11, 2017), https://fpf.org/blog/unfairness-by-algorithm-distilling-the-harms-of-automated-decision-making/

[19] Will Knight, *Inside a Misfiring Government Data Machine*, Wired (Mar. 26, 2043), https://www.wired.com/story/algorithmic-bias-government/

[20] Levi Kaplan, Alan Mislove, and Piotr Sapiezynski, *Measuring Biases in a Data Broker's Coverage*, FTC PrivacyCon 2022 (July 2017), https://www.ftc.gov/system/files/ftc_gov/pdf/PrivacyCon-2022-Kaplan-Mislove-Sapiezynski-Measuring-Biases-in-a-Data-Brokers-Coverage.pdf

consequential decisions about individuals. Not only are the risks more significant, but when they are realized, the harms related to the use of these systems are often particularly acute for individuals or communities. Such consequential decisions can include decisions affecting housing, employment, education, credit, and even the administration of justice. Federal law recognizes the unique sensitivity of consequential decisions,[21] and leading proposed AI legislation in the states also acknowledges the unique risks arising from automated decisions regarding these matters.[22]

Importantly, the possibility exists for personal data to be used at the development stage, deployment stage, or by any user of the AI system; therefore the cooperation of the developer to detect potential bias in the model may be recommended and required but will not necessarily be sufficient to mitigate this risk in end use. In summary, developers and deployers as well as end users of the model will likely all have roles to play in identifying and mitigating AI system-associated risk given their unique perspectives.

The Future of Privacy Forum appreciates this opportunity to comment on these issues and OMB's efforts to mitigate privacy risks, particularly those exacerbated by AI and other technologies. We welcome any further opportunity to provide resources or information to assist in this vital effort. If you have any questions regarding these comments and recommendations, please contact Anne J. Flanagan at ai@fpf.org (cc: info@fpf.org).

Sincerely,

Anne J. Flanagan, Vice President for Artificial Intelligence
Daniel Berrick, Policy Counsel
Amber Ezzell, Policy Counsel
Beth Do, Christopher Wolf Diversity Law Fellow

---

[21] E.g. Fair Credit Reporting Act (FCRA), 15 U.S.C. § 1681 et seq.; Fair Housing Act, 42 U.S.C. § 3601 et seq.; Equal Credit Opportunity Act (ECOA), 15 U.S.C. § 1691 et seq.; Americans with Disabilities Act (ADA), 42 U.S.C. § 12101.

[22] Cal. A.B. 2930 (2024) (regulating automated decision tools, formerly known as A.B. 331), https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202320240AB2930; Conn. S.B. 2 (2024) (regulating private-sector developers and deployers of high-risk AI systems), https://www.cga.ct.gov/asp/cgabillstatus/cgabillstatus.asp?selBillType=Bill&bill_num=SB00002&which_year=2024; Colo. S.B. 205 (2024) (modeled after CT S.B. 2), https://leg.colorado.gov/bills/sb24-205; Ill. H.B. 3773 (2024) (regulating the use of predictive data analytics for employment and credit decisions), https://ilga.gov/legislation/BillStatus.asp?GA=103&SessionID=112&DocTypeID=HB&DocNum=3773