



December 5, 2023

To: Director Shalanda D. Young
The Office of Management and Budget
725 17th St. NW
Washington, DC 20503

Submitted via regulations.gov

Re: Request for Comments: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum | 88 FR 75625

Table of Contents

Introduction	2
RFC Question 1: Are the roles, responsibilities, seniority, position, and reporting structures outlined for Chief AI Officers sufficiently flexible and achievable for the breadth of covered agencies?	3
The Chief AI Officer should have responsibility for, and issue guidance for, agencies' grantmaking for AI uses.	3
RFC Question 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?	4
Periodically review the presumed use cases; provide clearer guidance for agencies to use the list.	4
Clarify the list of presumed rights-impacting use cases to include further law enforcement-related activities in immigration; uses in public housing; and uses targeting labor and collective bargaining.	6
RFC Question 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?	7
Issue public guidance on analyzing AI risks to identify, evaluate and mitigate safety- and rights- impacting use cases, subject to stakeholder input.	7
Require AI Impact Assessments to document agency justifications for their AI uses relative to non-AI options.	9

Provide further guidance to agencies on steps to assess and mitigate bias.	9
Strengthen and concretize the waiver process to clarify when waivers may be granted and increase transparency.	12
RFC Question 7: What types of materials or resources would be most valuable to help agencies, as appropriate, incorporate the requirements and recommendations of this memorandum into relevant contracts?	13
Guide agencies towards AI-specific procurement policies that provide for pre-award vendor evaluation and post-award governance.	13
RFC Question 8: What kind of information should be made public about agencies' use of AI in their annual use case inventory?	18
Clarify guidance to ensure agencies' compliance with disclosure obligations.	18
Ensure the inventories' utility by establishing submission standards, quality checks, and a centralized inventory.	20
Establish AI inventories as the central place to report whether a use-case is deemed safety- or rights- impacting; whether a waiver was secured; and to publish any AI Impact Assessment	22
Prepare agencies for public requests for information	24
Conclusion	24

Introduction

The Center for Democracy & Technology (CDT) respectfully submits these comments in response to the Office of Management and Budget (OMB) Request for Comments regarding the draft memorandum for agency use of Artificial Intelligence (hereinafter “draft Memo”). CDT is a nonprofit 501(c)(3) organization that works to advance civil rights and civil liberties in the digital age. Among our priorities, CDT advocates for the responsible and equitable use of technology by government agencies, and promotes the adoption of robust, technically-informed solutions for the effective regulation and governance of AI systems.¹

We commend OMB for its leadership in setting AI accountability policy. The draft Memo shows a continued commitment to advancing equity, civil rights, and innovation, and recognizes the importance of providing implementing guidance to ensure that commitment is carried out effectively and consistently across the federal government. CDT welcomes OMB’s incorporation of key elements from the Administration’s Blueprint for an AI Bill of Rights in the draft Memo’s required minimum practices, such as requiring AI impact assessments, human consideration

¹ Center for Democracy & Technology, *Comments to Office of Science and Technology Policy on National Priorities for Artificial Intelligence* (Jul. 7, 2023), <https://cdt.org/wp-content/uploads/2023/07/CDT-comment-to-OSTP-on-national-AI-strategy23.pdf>.

and accountability, and notice and plain language documentation, as well as heightened requirements for uses impacting people’s rights such as disparate impact assessments, opt-outs, and consultation with affected groups. These requirements are essential to effectuating the Administration’s guiding principles for AI as set forth in the Executive Order on Safe, Secure & Trustworthy Development and Use of Artificial Intelligence.

Despite the strength of the Memo’s guidance, there are several places where it should be bolstered to increase transparency to help agencies and the public better understand federal uses of AI systems, and to improve consistency and ease agencies’ burden in interpreting the Memo. In addition, OMB must act quickly to issue further guidance on the key work streams needed to achieve the Memo’s goals, including more detailed guidance on AI use case inventories, disparate impact testing, and the waiver process, as well as necessary updates to federal procurement processes and guidance.

RFC Question 1: *Are the roles, responsibilities, seniority, position, and reporting structures outlined for Chief AI Officers sufficiently flexible and achievable for the breadth of covered agencies?*

Recommendation #1

The Chief AI Officer should have responsibility for, and issue guidance for, agencies’ grantmaking for AI uses.

The draft Memo does not extend to grants and federal financial assistance, and in so doing fails to provide guidance on a central way in which federal agencies cause the development and use of AI systems. Already, federal agency grantmaking has directly led to AI misuses: for example, when public housing authorities used funding from the Department of Housing & Urban Development to purchase surveillance cameras not only to monitor crime—as HUD intended—but also to monitor and discipline residents.²

OMB should clarify that the responsibilities of the CAIO includes managing risks and providing resources for responsible grantmaking with respect to AI, and provide agencies with guidance to mitigate these risks. For example, OMB could:

- clarify that Chief AI Officers’ responsibilities include managing risks from the use “*and federal funding of*” (draft Memorandum sections (3)(b)(ii)(J); (3)(b)(ii)(L));

² Dan Bateyko, *Taken for Granted: Where’s the Oversight of AI and Federal Funding?*, Center for Democracy & Technology (Aug. 7, 2023), <https://cdt.org/insights/taken-for-granted-wheres-the-oversight-of-ai-and-federal-funding/>.

- clarify that the Chief AI Officer’s responsibility to “maintain awareness of agency AI activities” includes AI-related grantmaking activities (section (3)(b)(ii)(B));
- clarify that agencies’ compliance plans under the draft Memorandum should include grantmaking programs (section(3)(b)(ii)(C));
- provide, or direct a relevant office (such as 18F or USDS) to provide, guidance on supporting responsible AI development and use for agency grantmakers – for example by directing agencies to work with legal counsel to review grants that implicate AI before the award of funds using the substantive practices set forth in the draft Memorandum, and convening experts to grant review boards.³

Translating OMB’s guidance on AI Accountability policy to grantmaking is in keeping with the Executive Order on Further Advancing Racial Equity,⁴ which aims for agencies to have comprehensive equity strategy in their programs and policies, and OMB M-21-27, which urges the use of high-quality evidence in government decision-making.⁵ Such guidance would reduce information silos and support grant reviewers that welcome timely consideration of grant applications by other agency staff.⁶

RFC Question 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?*

Recommendation #1

Periodically review the presumed use cases; provide clearer guidance for agencies to use the list.

The current list of presumed use cases is effective, and should not be altered by removing any items (we recommend several additions and clarifications below). The list reflects the Administration’s extensive work and public consultation in developing the Blueprint for an AI Bill

³ *Id.*

⁴ The White House, *Executive Order on Further Advancing Racial Equity and Support for Underserved Communities Through The Federal Government* (2023), <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/02/16/executive-order-on-further-advancing-racial-equity-and-support-for-underserved-communities-through-the-federal-government/>.

⁵ Office of Management and Budget, M-21-27, EVIDENCE-BASED POLICYMAKING: LEARNING AGENDAS AND ANNUAL EVALUATION PLANS, (2021), <https://www.whitehouse.gov/wp-content/uploads/2021/06/M-21-27.pdf>.

⁶ For more on agency grantmaking reforms, see, e.g. Eloise Pasachoff, *Federal Grant Rules and Realities in the Intergovernmental Administrative State: Compliance, Performance, and Politics*, YALE J. REGUL. (2020), <https://openyls.law.yale.edu/handle/20.500.13051/8307>.

of Rights, and is backed by a strong evidence base of harms documented in news coverage, policy reports, court cases, and academic literature.⁷

However, OMB should take further steps to ensure that agencies are consistently and appropriately interpreting the Memo. Specifically, OMB should:

- **Revisit and update the list of presumed use cases regularly**, with an opportunity for relevant experts and the public to recommend changes and provide comments (see further response on the process for soliciting public input in Q6).
- **Underscore that the list of presumed safety- and rights- impacting uses is not exhaustive** in any guidance, training, or other interactions, , and that agencies must conduct their own assessments to ensure that uses are appropriately designated.
- **Include specific examples of the presumed use cases** to help agencies understand them and recognize uses raising similar concerns (this could leverage the list of examples documented in the Administration’s Blueprint for an AI Bill of Rights).⁸
- **Provide examples of how data that implicates protected characteristics – like disability or race – can be processed in ways that create rights-impacting concerns**, even when that data is not ultimately used in a decision-making context. Such examples would clarify to agencies how AI systems that seem to function in a facially neutral way may in fact harm particular marginalized communities, and would affirm that agencies are obligated to scrutinize such AI uses for unanticipated harms.⁹
- **Require that agencies identify in their AI inventories whether or not a use case has been deemed safety- or rights-impacting in its AI inventory, with a justification, by making this a mandatory field in the AI inventory.** This requirement would build off existing reporting by agencies, which already must report on outcomes of stakeholder

⁷ See Data & Society’s regulatory submission in response to this RFC for further detail. *See also, e.g.,* Elizabeth Laird and Maddy Dwyer, Center for Democracy & Technology, *Off Task: EdTech Threats to Student Privacy and Equity in the Age of AI* (2023),

<https://cdt.org/insights/report-off-task-edtech-threats-to-student-privacy-and-equity-in-the-age-of-ai/>; Michael Yang, Center for Democracy & Technology, *The Promise and Peril of Data & Technology Use By Public Housing Agencies* (2023),

<https://cdt.org/wp-content/uploads/2023/03/2023-03-17-Civic-Tech-Technology-in-Public-Housing-Agencies-final.pdf>; Lydia X.Z. Brown et al., Center for Democracy & Technology, *Challenging the Use of Algorithm-Driven Decision-Making in Benefits Determinations Affecting People with Disabilities* (2020),

<https://cdt.org/insights/report-challenging-the-use-of-algorithm-driven-decision-making-in-benefits-determinations-affecting-people-with-disabilities/>; Hannah Quay-de la Vallee, Center for Democracy & Technology, *Public Agencies’ Use of Biometrics to Prevent Fraud and Abuse: Risks and Alternatives* (2022), <https://cdt.org/insights/public-agencies-use-of-biometrics-to-prevent-fraud-and-abuse-risks-and-alternatives/>.

⁸ Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights: Examples of Automated Systems* (2022), <https://www.whitehouse.gov/ostp/ai-bill-of-rights/examples-of-automated-systems/>.

⁹ For example, the guidance lists "detecting or measuring emotions, thought, or deception in humans." An AI system might perform such functions to help detect and prevent suspicious behavior based on speech, eye contact, bodily movements, and similar data points. Because of how the system might be designed and how its outputs might be used, it could adversely affect people with disabilities, Black and brown people, and transgender people.

engagements that identify AI applications that are within the agencies’ regulatory authorities.¹⁰ Agencies’ AI inventories should cite the regulatory authorities that apply to each use case, as this will provide further insights into whether the use case meets the guidance’s definitions of rights- or safety-impacting AI.

Recommendation #2

Clarify the list of presumed rights-impacting use cases to include further law enforcement-related activities in immigration; uses in public housing; and uses targeting labor and collective bargaining.

OMB should clarify the draft Memo’s application to certain AI uses in the immigration system which raise similar concerns to the use-cases listed in paragraphs 5(b)(ii)(B) and (C) but are not squarely covered by either paragraph. Immigration and Customs Enforcement (ICE) engages in law enforcement activities to investigate not just criminal, but civil and administrative law violations.¹¹ These activities include the use of AI for surveillance through license plate recognition, face recognition scans, DNA testing, and other functionalities.¹² Such uses are rights-impacting and far-reaching: one report finds ICE has used facial recognition technology on driver’s license photographs belonging to 1 in 3 adults in the U.S.¹³ The draft Memo should be amended to include the proposals in Just Futures Law’s comment for Section 5(b)(ii) to expressly include processes in the immigration system, biometric collection, and the use of credit bureaus.¹⁴

The draft Memo should also expand to cover uses in public housing, such as the surveillance of tenants. Earlier this year, investigative reporting revealed that public housing authorities used “Emergency Safety and Security Grants” administered by the Department of Housing and Urban Development to purchase surveillance cameras—many equipped with face surveillance technologies. Those cameras were then used by housing authorities to surveil tenants, detect violations of lease terms (including those unrelated to tenant safety), and evict tenants.¹⁵ As HUD has since advised, such uses were not eligible for the emergency grants.¹⁶

¹⁰ Office of Management and Budget, M-21-06, *Guidance for Regulation of Artificial Intelligence Applications* (2020), <https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-06.pdf>.

¹¹ Nina Wang et al., *American Dragnet: Data-Driven Deportation in the 21st Century* (2022), <https://www.americandragnet.org/>.

¹² *Id.*

¹³ *Id.*

¹⁴ See Just Futures Law’s regulatory submission in response to this RFC for further detail.

¹⁵ Douglas MacMillan, *Eyes on the Poor: Cameras, Facial Recognition Watch Over Public Housing*, Washington Post (May 16, 2023), <https://www.washingtonpost.com/business/2023/05/16/surveillance-cameras-public-housing/>.

¹⁶ Dept. of Housing and Urban Development, Notice PIH 2023-10, *Emergency Safety and Security Grants Annual Funding Notification and Application Process* (2023), <https://www.hud.gov/sites/dfiles/PIH/documents/2023PIH10.pdf>.

OMB should also amend the list of rights-impacting use cases to include “collective bargaining, workplace organizing, union membership, or concerted activity” in part G to further ensure the protection of workers from the use of automated worker surveillance systems that interfere with the right to organize for better and safer working conditions.¹⁷

RFC Question 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?*

Recommendation #1

Issue public guidance on analyzing AI risks to identify, evaluate and mitigate safety- and rights- impacting use cases, subject to stakeholder input.

The minimum practices require agencies to identify and analyze safety- and rights-impacting use cases. This welcome policy will face a persistent challenge: people fail to recognize high-risk AI uses. Even experts struggle. In one study, graduate-level computer scientists asked to pinpoint risks in an AI system made poor judgment calls over who was at risk, what the harms were, and what caused the harm.¹⁸

Agencies need to rely on more than their best guesses when identifying AI risks. Compounding the problem, agency officials have been given mixed signals through varying government definitions of what counts as “AI.” Without good, consistent guidance, agency officials will struggle to classify AI,¹⁹ as an Inspector General’s report on NASA’s AI management strategy illustrates:

“Personnel we interviewed stated they reported AI based on their own individual understanding of what the term means rather than a formal definition provided by the

¹⁷ See Center for American Progress’s regulatory submission in response to this RFC for further detail.

¹⁸ Emily Black et al., *Toward Operationalizing Pipeline-Aware ML Fairness: A Research Agenda for Developing Practical Guidelines and Tools*, in EQUITY AND ACCESS IN ALGORITHMS, MECHANISMS, AND OPTIMIZATION 1 (2023), <https://dl.acm.org/doi/10.1145/3617694.3623259>.

¹⁹ Chloe Autio et al., *A Snapshot of Artificial Intelligence Procurement Challenges* (2023), <https://files.thegovlab.org/a-snapshot-of-ai-procurement-challenges-june2023.pdf> (finding, e.g. “Civilian agencies like the Department of Transportation (DOT) describe a more conservative approach to AI procurement, given the agency’s mandate to protect public safety.”)

Agency. As a result, NASA does not have a singular designation or classification mechanism to accurately classify and track AI or to identify AI expenditures within the Agency's financial system, making it difficult for the Agency to meet federal requirements to monitor its use of AI."²⁰ (emphasis added)

OMB should issue further guidance, provide support and establish oversight mechanisms to help agencies accurately categorize their uses. Any such guidance should be publicly available, and agency AI system classification accuracy should be regularly reviewed to evaluate the guidance's effectiveness. Additionally, OMB should implement any of the following to further support agency staff:

- **Stand up a cross-agency working group.** A working group, supported by OMB, can discuss implementation across agencies. Cross-agency information sharing will be vital, especially with statistical, scientific, and cybersecurity agencies with experience in managing datasets, AI tools, and their attending cybersecurity risks. Any products from these working groups could be made accessible to all agencies.
- **Designate a point-of-contact at OMB in charge of compliance and support for agencies in determining AI use cases.**
- **Consult with NIST on a standard for Algorithmic Impact Assessments (AIA).** There is no formal standard for an AIA. NIST could help develop and define this standard based on the level of risk presented, similar to how NIST established federal standards for information processing to determine whether risks have low, moderate, or high security impacts.²¹ Agencies should be encouraged to add sector-specific additions to the AIAs. Such a standard, too, must be open to stakeholder feedback.
- **Improve guidance through public consultations and workshops inviting academics, civil society, and other impacted members of the public.** When engaging with those affected by an AI use case, OMB must meet them where they are by providing multiple ways to provide input.
- **OMB should provide clearer guidance to help agencies interpret the list of presumed safety- and rights-impacting use cases, and require agencies to disclose in the AI inventory whether an AI use case has been designated safety- or rights-impacting, with an explanation,** to help OMB and the public ensure consistency across agency interpretations.

²⁰ Office of Inspector General, National Aeronautics and Space Administration, *NASA'S MANAGEMENT OF ITS ARTIFICIAL INTELLIGENCE CAPABILITIES* (2023), <https://oig.nasa.gov/docs/IG-23-012.pdf>.

²¹ See, e.g. National Institute of Standards and Technology, *FIPS PUB 199, Standards for Security Categorization of Federal Information and Information Systems* (2004), <https://doi.org/10.6028/NIST.FIPS.199>; National Institute of Standards and Technology, *NIST SP 800-53, Security and Privacy Controls for Information Systems and Organizations* (2020), <https://doi.org/10.6028/NIST.SP.800-53r5>.

Recommendation #2

Require AI Impact Assessments to document agency justifications for their AI uses *relative to non-AI options*.

The draft Memo requires agencies to document the expected benefit of an AI system, supported by specific metrics or qualitative analysis. This appropriate requirement pushes agencies to provide evidence that the AI system is a good fit to accomplish agency goals and related tasks. **To build on this culture of evidence,²² OMB should require agencies to document the expected benefit *relative to other, non-AI options*.**

Documenting the benefit of AI relative to non-AI options has several upsides. First, agencies that document alternative options in an AI Impact Assessment will be prompted to consider efficacy risks particular to AI.²³ Second, agencies would provide a useful baseline to compare the AI's benefits and drawbacks for any affected groups providing feedback on whether to remove the AI from use when called for consultation under Section 5(v)(B). Lastly, instead of simply determining whether an AI system achieves its goal, having this additional point of comparison would allow agencies and people to determine if the goal could be better accomplished without AI.

This recommendation may require only a simple change to AI or Algorithmic Impact Assessments: for example, the Algorithmic Impact Assessment questionnaire hosted at CIO.gov asks respondents what other approaches to solving the problem were considered.²⁴ This and any similar sections should be amended to ask agencies to provide details about any non-AI approaches that were considered.

Recommendation #3

Provide further guidance to agencies on steps to assess and mitigate bias.

As part of the additional minimum practices required for rights-impacting AI, agencies must take steps to ensure that the AI will advance equity, dignity, and fairness, including proactively identifying and removing factors contributing to algorithmic discrimination; assessing and

²² M-21-27, *supra* note 5.

²³ For risks particular to AI, see, e.g. Inioluwa Deborah Raji et al., *The Fallacy of AI Functionality*, in *2022 ACM Conference on Fairness, Accountability, and Transparency* 959 (2022), <https://dl.acm.org/doi/10.1145/3531146.3533158>.

²⁴ Federal Chief Information Officers Council, *Algorithmic Impact Assessment*, <https://www.cio.gov/aia-eia-is/#/> (last visited Nov 29, 2023).

mitigating disparate impacts; and using representative data. For many agencies, these steps will involve creating wholly new processes requiring additional support from OMB.

For example, **OMB should issue guidance to agencies to address the lack of adequate demographic data used in identifying and removing algorithmic discrimination.** The Racial Equity EO 14091 directed agencies to create Equity Action Plans and for OMB to consider how the President’s Budget can support these plans. One review of those plans found that 21 out of 25 agencies saw their lack of demographic data as a barrier to implementation.²⁵ While navigating this challenge is not unfamiliar in regulated and sensitive contexts such as employment, credit, and healthcare,²⁶ agencies have incentives not to collect demographic data or correct bias. In a Government Accountability Office report, one testimonial reflects how employees “prefer to remain unaware [of bias] because they consider this the safest way of proceeding” given risks of litigation and other anti-discrimination liabilities.²⁷ OMB guidance may help to structure and place appropriate guardrails around the gathering of demographic data to alleviate these concerns. To this end, a policy report from the Stanford Institute for Human-Centered AI offers OMB-relevant proposals, including:

- **Interpreting the Paperwork Reduction Act to speed up the gathering of demographic data** on surveys,²⁸
- **Interpreting The Privacy Act to exempt bias assessments**, which fall under the exempt “statistical research” category,²⁹ and
- **Promulgating guidance for agencies around how to place technical and institutional protections around demographic data** to ensure it is only used for bias measurement and lawful bias mitigation efforts, and to build agency personnel’s confidence in using the data safely.³⁰

OMB should consider **providing advice on when agencies or their vendors should consider the use of imputation techniques such as BISG (Bayesian Improved Surname Geocoding)**³¹ or

²⁵ Arushi Gupta et al., *The Privacy-Bias Tradeoff: Data Minimization and Racial Disparity Assessments in U.S. Government*, in 2023 ACM Conference on Fairness, Accountability, and Transparency 492 (2023), <https://dl.acm.org/doi/pdf/10.1145/3593013.3594015>.

²⁶ Miranda Bogen et al., *Awareness in Practice: Tensions in Access to Sensitive Attribute Data for Antidiscrimination*, in 2020 Conference on Fairness, Accountability, and Transparency (2020), <https://arxiv.org/abs/1912.06171v1>.

²⁷ U.S. Government Accountability Office, GAO-21-519SP, *Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities* (2021), <https://www.gao.gov/products/gao-21-519sp>.

²⁸ GUPTA ET AL., *supra* note 25.

²⁹ *Id.*

³⁰ *Id.*

³¹ See, e.g. Marc N. Elliott et al., *A New Method for Estimating Race/Ethnicity and Associated Disparities Where Administrative Records Lack Self-Reported Race/Ethnicity*, 43 HEALTH SERV. RES. 1722 (2008), <https://pubmed.ncbi.nlm.nih.gov/18479410/>; Kosuke Imai et al., *Addressing Census Data Problems in Race Imputation via Fully Bayesian Improved Surname Geocoding and Name Supplements*, 8 SCI. ADV. eadc9824 (2022), <https://www.science.org/doi/10.1126/sciadv.adc9824>.

proxies such as zip code to conduct bias analyses, taking into account the benefits and limitations of such approaches.³² For some protected characteristics (such as race and ethnicity), such techniques may be readily available due to established use by industry and regulators for nondiscrimination purposes, while for others (such as sexual orientation or religion) there may be compelling reasons to avoid such an approach.³³

Additionally, OMB should consider adding force to the requirement that **agencies test AI for significant disparities by requiring agencies (or their vendors) to search for a “less discriminatory algorithm” (LDA).**³⁴ Current computer science findings suggest that an “equally accurate, but less discriminatory alternative algorithm almost always exists.”³⁵ These alternative algorithms are not apparent without proactive testing, but such testing generally imposes little cost, as the steps taken to compare model candidates before selecting one for use are already standard in the development of AI systems.³⁶ Agencies looking for mitigations to disparate impact might start by performing such a search or requiring an AI vendor show sufficient proof that such a search has been conducted.

In no case should a lack of demographic data be considered a free pass for agencies to avoid confronting potential discriminatory treatment or impacts of a system. **Agencies should be obligated to make best efforts to understand whether such impacts are likely**, even if doing so requires manual review of input data for obvious proxies or for qualitative analysis to identify potentially detrimental outcomes likely to fall most heavily on marginalized communities. In clarifying these minimum practices further, OMB will be following its concurrent obligation under the Racial Equity EO 14091 to embed equity into government-wide processes.³⁷

Recommendation #4

³² Under some conditions, methods like Bayesian Improved Surname Geocoding (BISG) may over- or under-estimate demographic disparities, but they are still useful in identifying directional disparities. See, e.g. Aaron Rieke et al, *Imperfect Inferences: A Practical Assessment*, ACM Conference on Fairness, Accountability, and Transparency (2022), <https://dl.acm.org/doi/10.1145/3531146.3533140>. However, improper use of imputation approaches can risk misrepresentation and reinforcement of social categories. McKane Andrus et al, *Fairer Algorithmic Decision Making & Its Consequences Interrogating the Risks and Benefits of Demographic Data Collection, Use, and Non-Use*, Partnership on AI (2022), <https://partnershiponai.org/paper/fairer-algorithmic-decision-making-and-its-consequences>.

³³ Fairer Algorithmic Decision Making, *supra* note 32.

³⁴ Emily Black et al., *Less Discriminatory Algorithms*, Washington University in St. Louis Legal Studies (2023), <https://papers.ssrn.com/abstract=4590481> (last visited Nov 22, 2023).

³⁵ *Id.*

³⁶ *Id.*

³⁷ *Executive Order on Further Advancing Racial Equity and Support for Underserved Communities Through the Federal Government*, *supra* note 4.

Strengthen and concretize the waiver process to clarify when waivers may be granted and increase transparency.

Under the draft Memo, an agency may request a waiver to the minimum practices through its Chief AI Officer if fulfilling the minimum practices “would increase risks to safety or rights overall or would create an unacceptable impediment to critical agency operations.” Without further checks, such a waiver incentivizes agencies to skip the procedural requirements of the draft Memo. OMB can look to existing laws and practices for examples of how to improve the waiver process.

OMB can look to principles from administrative law for standards to review waivers. In reviewing agency justifications for skipping rulemaking procedures due to “good cause” or emergencies, courts have rejected certain arguments that would incentivize end-runs around procedural requirements. Those arguments include bald assertions of an emergency, the need to meet an agency’s own efficiency goals, and the agency’s aim to comply with a statutory deadline.³⁸ An agency-induced delay to meet its mission should not be cause for a waiver. OMB can also look to the Federal Communications Commission (FCC)’s forbearance rules. Under these rules, the FCC reviews telecom carriers’ petitions asking the FCC to waive enforcement of its regulations, and will grant those petitions only if all requirements are met, including meeting the three statutory requirements.³⁹ The forbearance rules enumerate procedural requirements⁴⁰ relevant to waivers submitted to OMB.

OMB should further clarify its process for granting waivers and provide transparency into the process. The review process should include a **strong presumption against issuing waivers**. Further, OMB should provide **strong and clear standards for bulk waivers** contemplated at Footnote 25 of the draft memo, such as granting them only in extraordinary circumstances.

A related exclusion that requires further guardrails is the one allowing AI research and development to skip the minimum practices. Agencies that conduct **research and development**

³⁸ Jared P Cole, *The Good Cause Exception to Notice and Comment Rulemaking: Judicial Review of Agency Action*, CONGR. RES. SERV. (2016), <https://crsreports.congress.gov/product/pdf/R/R44356>.

³⁹ 47 U.S. Code § 160; Pub. L. 104–104, title IV, § 401, Feb. 8, 1996, 110 Stat. 128.

⁴⁰ Federal Communications Commission, Forbearance Procedural Rules, <https://www.fcc.gov/general/forbearance-procedural-rules>. Those three statutory requirements are “(1) enforcement of such regulation or provision is not necessary to ensure that the charges, practices, classifications, or regulations by, for, or in connection with that telecommunications carrier or telecommunications service are just and reasonable and are not unjustly or unreasonably discriminatory; (2) enforcement of such regulation or provision is not necessary for the protection of consumers; and (3) forbearance from applying such provision or regulation is consistent with the public interest.” 47 U.S. Code § 160; Pub. L. 104–104, title IV, § 401, Feb. 8, 1996, 110 Stat. 128.

that involves human subjects should not skip the minimum practice, given that such research may affect the welfare of human subjects.⁴¹

RFC Question 7: *What types of materials or resources would be most valuable to help agencies, as appropriate, incorporate the requirements and recommendations of this memorandum into relevant contracts?*

Recommendation #1

Guide agencies towards AI-specific procurement policies that provide for pre-award vendor evaluation and post-award governance.

Modern federal procurement policy is inadequate to ensure responsible AI practices. Improving contracting will necessarily involve improving acquisition planning.

The comments below begin with an illustrative table of how AI-specific risks will require the procurement process to change. We next present some high-level considerations for pre-award vendor evaluation and post-award governance. These **acquisition planning considerations can be implemented by formal changes to the Federal Acquisition Regulation (FAR), OMB central guidance, or at the discretion of individual agencies.**

The following chart from CDT’s forthcoming report, *How to Not Buy Magic: Updating Procurement Policies to Support the Federal Government’s Responsible Use of AI*, is illustrative of how procurement processes will need to change to manage AI-specific risks, as characterized by NIST.⁴²

How NIST characterizes the AI-specific risk	How the risk might affect the procurement process
<i>The data used for building an AI system may not be a true or appropriate representation of the context or intended use of the AI system, and the ground truth may either not exist or not be available. Additionally, harmful bias and other data quality issues can affect AI system trustworthiness, which could lead to negative</i>	A hallmark principle of public procurement is that up-front competition will lead to the best vendor selection, but adequate testing and monitoring of a solution in production will need to be built into procurement processes because full evaluation is not possible <i>ex ante</i> .

⁴¹ For how research may impact welfare, see, e.g. Office for Human Research Protections, *Dept. of Health and Human Services, The Belmont Report* (1979), <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/index.html>.

⁴² HOW TO NOT BUY MAGIC: UPDATING PROCUREMENT POLICIES TO SUPPORT THE FEDERAL GOVERNMENT’S RESPONSIBLE USE OF AI, (Center for Democracy & Technology, Forthcoming).

How NIST characterizes the AI-specific risk	How the risk might affect the procurement process
<i>impacts.</i>	
<i>Intentional or unintentional changes during training may fundamentally alter AI system performance.</i>	Procurement rules require formal processes (such as change orders or modifications) to address changes after “acceptance” of a service or supply. These processes should be triggered when intentional changes are made, and when unintentional changes are detected or suspected due to system performance.
<i>Datasets used to train AI systems may become detached from their original and intended context or may become stale or outdated relative to deployment context.</i>	Procurements are often structured toward the creation of long-term contracts (typically, 5 years). Retraining AI will need to be built into procurements.
<i>AI systems may require more frequent maintenance and triggers for conducting corrective maintenance due to data, model, or concept drift.</i>	Without effective management practices, there is a significant likelihood that “Operations and Maintenance” budgets will compete with “Development, Modernization, and Enhancement” budgets or incentivize agencies to underinvest in post-award monitoring and improvement.
<i>Increased opacity and concerns about reproducibility.</i>	A lack of transparency can create distrust and prevent accountability between government and industry and harm the public’s trust in government. Appropriate transparency and explainability requirements will need to be built into contracts with internal technical expertise to analyze disclosures.
<i>Difficulty in performing regular AI-based software testing, or determining what to test, since AI systems are not subject to the same controls as traditional code development.</i>	Structuring procurements with appropriate evaluation criteria is already difficult. Developing and tailoring appropriate evaluation criteria for AI will require additional technical capacity within the government.
<i>Inability to predict or detect the side effects of AI-based systems beyond statistical measures.</i>	Without mature post-award governance models, it may not be clear who has the legal and financial burden to deal with unintended consequences.

i. Pre-award vendor evaluation

Agencies rely on vendors to evaluate performance, but vendors are exploiting confusion over what AI can accomplish.⁴³ The market for AI tools, particularly in the generative AI space, is filled with hyperbole.⁴⁴ In light of the unique risks from AI-based systems and the challenging market, OMB guidance should help prepare agencies for acquisition planning.

1. Advise agencies to prepare a template RFI or questionnaire in advance for market research.

OMB should lead in publishing a template questionnaire which would lay out requirements for responsible AI that agencies could then adapt according to their assessment of the needs they want an AI tool to meet. Such a template would serve two purposes: a template will help government agencies evaluate the market and signal to industry the government's goals for responsible AI. For example, the questionnaire could require vendors to explain:

- Their procedures for assessing and documenting their products' accuracy rates and disparate outcomes for marginalized groups
- Their governance over how their systems are used, and the systems' reliability for intended and likely actual uses
- The sourcing of training data and actual inputs, and representation of different demographics in this data
- Data minimization and other data protection measures applied to their systems, and how vendors provide for agencies to have and retain control over the data
- Procedures for monitoring and improving system performance, including through stakeholder engagement, as well as processing of complaints and appeals.⁴⁵

Including such requirements in the template questionnaire would give agencies insight into vendors' practices that would streamline agencies' implementation of minimum practices when proceeding with a chosen system. Agencies would use responses to the questionnaire to refine their request for proposals (RFPs) and contracts. **OMB should lead not only by publishing this**

⁴³ *Id.*

⁴⁴ Sharon Goldman, *Princeton University's 'AI Snake Oil' Authors Say Generative AI Hype Has 'Spiraled Out of Control,'* VENTUREBEAT (Aug. 23, 2023), <https://venturebeat.com/ai/princeton-university-ai-snake-oil-professors-say-generative-ai-hype-has-spiraled-out-of-control/>.

⁴⁵ For further detail about these and other requirements agencies should consider, see Hannah Quay-de la Vallee and Natasha Duarte, Center for Democracy & Technology, *Algorithmic Systems in Education: Incorporating Equity and Fairness When Using Student Data* 27 (2019), <https://cdt.org/wp-content/uploads/2019/08/2019-08-08-Digital-Decision-making-Brief-FINAL.pdf>.

template, but also by educating agencies in how to use and evaluate responses to the questionnaire for AI procurement.

2. Empower agencies to use their authority to test systems before purchase.

Unlike in traditional procurement, vendor evaluation of AI relies more heavily on the vendor's assertions that a system works as described. These assertions are difficult to check without thorough vetting of the system in the context in which it is intended to be used. For that reason, agencies should consider "pre-evaluating" systems. Agencies already have the authority to conduct technical evaluations before making an award. **OMB should encourage agencies to use this authority to incorporate more rigorous evaluation methods such as "show, not tell" demonstrations.**⁴⁶

3. Obligate Chief AI Officers to ensure all contracts for AI systems receive legal review.

Contracting offices have triggers and thresholds at which a contract must receive legal review. **A safety- or rights-impacting use of AI should be considered another trigger for a contract to be subject to legal sign-off.** Agency attorneys should ensure that acquired AI is consistent with civil rights laws, and EO 14091's requirement for agencies to deliver equitable outcomes.⁴⁷ Due to the harms arising from the use of personally identifiable information (PII) in AI, agency attorneys should also review how AI contracts implicate PII or otherwise pose privacy concerns.

4. Ensure procurement offices have adequate staffing and training.

In our forthcoming report in which we interview experts working with the federal government, multiple interviewees noted that there are not enough individuals within government with the time or expertise to oversee AI vendors.⁴⁸ **The government must grow its capacity to oversee AI vendors in order to innovate on AI uses.**

As the government grows to accommodate these **new skills, training modules, ideally incorporated into existing procurement curriculum, will help.** For example, the Digital IT Acquisition Profession (DITAP) program is one training model that might bring in new training

⁴⁶ See, e.g. Dept. of Homeland Security Procurement Innovation Lab, *PIL Boot Camp Innovation Technique 2 - Product or Technical Demonstration*, https://www.fai.gov/sites/default/files/periodic_table/PIL_Boot_Camp_workbook-Innovation_Technique_2.pdf. As we write in a forthcoming report, "if a vendor is unable to meet the government's requirements, the government would be better off eliminating them *pre-award*. To accomplish this requires the government to conduct pre-award evaluations and technical demonstrations." HOW TO NOT BUY MAGIC: UPDATING PROCUREMENT POLICIES TO SUPPORT THE FEDERAL GOVERNMENT'S RESPONSIBLE USE OF AI, *supra* note 42.

⁴⁷ *Executive Order on Further Advancing Racial Equity and Support for Underserved Communities Through the Federal Government*, *supra* note 4.

⁴⁸ HOW TO NOT BUY MAGIC: UPDATING PROCUREMENT POLICIES TO SUPPORT THE FEDERAL GOVERNMENT'S RESPONSIBLE USE OF AI, *supra* note 42.

specific to AI procurement. In this effort, the United States Digital Service and the Office of Federal Procurement Policy could work with DITAP providers to incorporate AI accountability policy into the curriculum, ensuring a broader reach of training.⁴⁹

ii. Post-award governance

Post-award governance is critical to agency innovation in the use of AI. The realities of managing AI systems—realities such as training AI systems and assigning legal and financial responsibilities to any unintended consequences—must be handled with post-award governance practices. Such practices are often either unfamiliar to agencies or deprioritized.⁵⁰ Yet successful adoption of AI requires foresight and planning throughout the agency’s use case. The responsibilities and costs for managing AI systems should be developed early in procurement. Agencies should also establish *redress processes*, to assign responsibility for fixing issues that occur in deployment. IEEE’s forthcoming AI procurement standards for AI and automated systems may also be instructive in this stage.⁵¹

1. Advocate for the addition of specific reporting requirements that correspond to the AI RMF and AI Bill of Rights in the FITARA scorecard.

Agencies have improved their IT adoption and inventorying under the Federal Information Technology Acquisition Reform Act (FITARA) in large part due to the success of the FITARA scorecard, which grades agencies performance on implementation.⁵² A **similar scorecard for responsible AI practices would likely similarly improve AI management**, because, as detailed in a forthcoming report, “every [expert] interviewee expressed doubt that agencies are meaningfully planning for post-award monitoring of AI.”⁵³

2. Develop and publish an “oversight guide” for reviewing agency acquisition activities.

OMB can **develop an “oversight guide” like that developed by GAO⁵⁴ for use by agency officials, congressional representatives, inspectors general, and civil society.** Such a guide would help these government actors and stakeholders monitor the responsible use of AI by offering example questions to ask and corrective actions to take.

⁴⁹ *Id.*

⁵⁰ *Id.*

⁵¹ IEEE Standards Association, *Standard for the Procurement of Artificial Intelligence and Automated Decision Systems*, <https://standards.ieee.org/ieee/3119/10729/> (last visited Nov 12, 2023).

⁵² U.S. Government Accountability Office, GAO-22-105659, *Information Technology: Biannual Scorecards Have Evolved and Served as Effective Oversight Tools* (2022), <https://www.gao.gov/products/gao-22-105659>.

⁵³ HOW TO NOT BUY MAGIC: UPDATING PROCUREMENT POLICIES TO SUPPORT THE FEDERAL GOVERNMENT’S RESPONSIBLE USE OF AI, *supra* note 42.

⁵⁴ U.S. Government Accountability Office, *supra* note 27.

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

Recommendation #1

Clarify guidance to ensure agencies' compliance with disclosure obligations.

The AI Executive Order and draft Memo emphasize transparency as a key element of agencies' responsible AI governance. However, the Memo must be strengthened to achieve effective transparency, and to ensure that agencies interpret the guidance consistently in a manner that makes their disclosures useful to government actors and the public.

Currently, AI inventories are falling short of expectations, with agencies missing the mark on compliance, quality, completeness, and accessibility. Many agencies failed to publish inventories despite a previous Executive Order mandate: a review by Stanford Law scholars of 220 agencies found that 3 out of every 4 agencies potentially required to publish an AI use case inventories did not.⁵⁵ Moreover, agencies that did publish inventories often left them incomplete: a review by CDT in summer 2023⁵⁶ showed that agencies' entries ranged considerably from one-line descriptions to full paragraphs, with case summaries often using vague and inaccessible language that fails to convey the purpose of a tool, or whom the tool may affect. Entries typically lack dates, making it difficult to know when updates happened. Delays in reporting abound.⁵⁷ Agencies varied widely in where they host the inventory online and in which format, despite attempts by The Chief Information Officers Council (CIO Council) to standardize this process.⁵⁸

OMB must take warning from this flawed start and issue clear, specific directives to make the AI inventories into the effective transparency tool the Executive Order and draft Memo envision.

⁵⁵ Christie Lawrence et al., *The Bureaucratic Challenge to AI Governance: An Empirical Assessment of Implementation at U.S. Federal Agencies*, in Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society 606 (2023), <https://dl.acm.org/doi/pdf/10.1145/3600211.3604701>. Among 49 agencies with publicly known AI uses—including AI uses core to the agency's mission—nearly 50% did not publish an AI inventory. Notably, agencies with missions that significantly bear on equity concerns, like the Department of Housing and Urban Development, Small Business Administration, and U.S. Postal Service, have failed to publish known AI uses.

⁵⁶ Bowman Cooper, *Like Looking for a Needle in an AI-Stack*, Center for Democracy & Technology (Jul. 21, 2023), <https://cdt.org/insights/like-looking-for-a-needle-in-an-ai-stack/>.

⁵⁷ Lawrence et al., *supra* note 55 (see, Appendix D.)

⁵⁸ Federal Chief Information Officers Council, *Guidance for Creating Agency Inventories of AI Use Cases Per EO 13980* (2023), <https://www.cio.gov/assets/resources/2023-Guidance-for-AI-Use-Case-Inventories.pdf>.

Specifically OMB should:

1. Clarify agencies’ obligation to disclose AI uses and give specific guidance on permissible exceptions.

OMB must emphasize agencies’ obligation to disclose AI uses and give clearer guidance on permissible exceptions. Currently, the draft Memo states that “Agencies must only publicly report use cases to *the extent practicable and consistent with applicable law and governmentwide guidance*, including those concerning the protection of privacy and of sensitive law enforcement, national security, and other protected information.”⁵⁹ (emphasis added)

This language, drawn from the 2023 Advancing American AI Act,⁶⁰ provides insufficient guidance for agencies to consistently interpret and effectuate the transparency goals of the Memo and AI Executive Order. OMB must guide agencies on how to interpret the language of the statute, by:

- **Clarifying in what circumstances reporting could be deemed “not practicable”** and ensure this language is narrowly construed.
- Giving **further explanation on how guidance regarding “sensitive law enforcement, national security, and other protected information”** should be interpreted. Such guidance is essential to avoid this exception being misused to avoid reporting, when the draft Memo expressly recognizes that many law enforcement uses are presumptively safety- and/or rights-impacting and warrant heightened scrutiny.
- **Clarifying the reporting obligations of sub-agencies and departments.**
- **Reinforcing to agencies that these inventories support stated objectives of the President** to support agency transparency, accountability, coordination, and innovation – and are not mere compliance exercises.
- **Reviewing, alongside other relevant parties and on a regular basis, agencies’ determinations that systems may be exempt from public disclosure** and challenge overbroad interpretations where they are identified.

2. Develop an annex for sensitive and classified use cases.

The executive office should never be left in the dark about classified AI uses by federal agencies. The draft Memo currently exempts intelligence agencies from AI inventory requirements, and (as noted above) permits agencies to withhold reporting when consistent with law and guidance concerning sensitive law enforcement, national security and other protected

⁵⁹ Office of Management and Budget, Exec. Office of the President, *Proposed Memorandum for the Heads of Executive Departments and Agencies, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence* (2023), <https://ai.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-Public-Comment.pdf>.

⁶⁰ Pub. L. No. 117-263, div. G, title LXXII, subtitle B, §§ 7224(a), 7224(d)(1)(B), and 7225 [Sec. 7225(a)(2)].

information. To ensure agency accountability and meet the information-sharing purposes of the AI Inventory while preserving secrecy obligations, **OMB should work with agencies to create a classified annex to the Inventory.** Such an annex should be made available to relevant oversight bodies like the House and Senate Select Committees on Intelligence, GAO, and the National Security Council. The annex at minimum should include the same kind of information the Department of Defense provides to OMB for its AI Uses not subject to inventory: namely, a summary of the AI Use, along with relevant metrics as laid out in Section 3(a)(v) of the draft Memo. Independent overseers, like the Privacy and Civil Liberties Oversight Board, should have a role in reviewing these disclosures.

3. Lead by example.

OMB and other offices within the Executive Office of the President should lead by example by publishing their own AI use cases in the consolidated inventory. The White House can set a precedent in transparent and responsible AI by proactively publishing details of their AI applications.

Recommendation #2

Ensure the inventories' utility by establishing submission standards, quality checks, and a centralized inventory.

1. Establish basic requirements for AI uses case summaries

A reader should understand the *what*, *why*, *when*, and *how* of each AI use case from its summary. Summaries in use case inventories currently miss that mark: some summaries tell us *what* the agency's goal is⁶¹ but not *how* AI is involved and vice versa.⁶²

To meet this criteria, **an AI use case summary should include the following:**

- What inputs the AI system uses
- What the AI system predicts or decides
- What is the agency's stated goal in using an AI system to fulfill the task
- How the agency will measure the success of the AI system

⁶¹ For example, listings for the National Institutes of Health provide descriptions that simply state "Protection of Windows and Mac endpoints from Cyberthreats" or "This project will help to reduce the energy usage for producing chilled water to cool the NIH campus." Available at: Department of Health and Human Services: Artificial Intelligence Use Cases Inventory (last updated Aug. 2023), <https://www.hhs.gov/about/agencies/asa/ocio/ai/use-cases/index.html>.

⁶² For example, a listing in the Department of Agriculture's inventory states: "AI-type statistical techniques are used to model predictive relationships between variables. We routinely use modeling approaches such as random forest, artificial neural networks, k-nearest neighbor clustering, and support vector machines, for statistical prediction." Available at: Dept. of Agriculture, *Inventory of USDA Artificial Intelligence Use Cases* (last updated May 2023), https://www.usda.gov/data/AI_Inventory.

- When the agency uses AI in a certain process
- How long the agency plans to use the AI system
- What departments or agency actors will use the AI system⁶³

OMB can include such requirements in its revised memo, as well as change fields in the current inventories like “*where did/does the training data originate*” and “*What specific AI techniques were used*” from “optional” to “required”.⁶⁴

Providing guidance in the revised memo to this effect will better serve the Memo’s transparency goals, promote consistency, and make OMB’s role in administering the policy easier because agencies will adopt more uniform approaches from the start.

2. Establish quality checks on agencies’ Inventory submissions

OMB must ensure agency submissions are complete and easily understood by the broader public. To maximize efficiency for OMB and agencies alike, this should be achieved through clear guidance and parameters at the data entry phase, as well as auditing. OMB should draw on guidance from Inspectors General for managing spending records,⁶⁵ including:

- **Provide example language and structured questions**
- **Ensure that entries (and modifications) are dated**
- **Consider minimum length requirements** for fields that warrant longer explanations, such as the AI Use Case Summary (clarifying that a 1-2 sentence submission is insufficient)
- **Instruct agencies to avoid jargon** and abbreviations⁶⁶
- **Develop automated checks and controls** in the submission system, for example via required fields.

3. Require the maintenance of a consolidated inventory and long-term recordkeeping

The website *ai.gov* currently hosts an inventory consolidating agencies’ disclosures, but the inventory does not appear to be mandated by the draft Memo or other documents (CIO Council

⁶³ For risks stemming from cross-departmental technical tool use, see, e.g. Mary Pat Dwyer and Rachel Levinson-Waldman, Brennan Center for Justice, *A Realignment for Homeland Security Investigations* (2023), <https://www.brennancenter.org/our-work/research-reports/realignment-homeland-security-investigations>.

⁶⁴ *Guidance for Creating Agency Inventories of AI Use Cases Per EO 13980*, supra note 58.

⁶⁵ U. S. Government Accountability Office, *DATA Act: OIGs Reported That Quality of Agency-Submitted Data Varied, and Most Recommended Improvements* | U.S. GAO, (2020), <https://www.gao.gov/products/gao-20-540>.

⁶⁶ See, e.g. Office of Inspector General, Dept. of Homeland Security, *DHS Has Made Progress in Meeting DATA Act Requirements, But Challenges Remain* (2020), <https://www.oig.dhs.gov/reports/2020/dhs-has-made-progress-meeting-data-act-requirements-challenges-remain/oig-20-62-aug20> (“DHS did not use plain English language when reporting award descriptions as required. Officials used shorthand descriptions, abbreviations, or terminology that could only be understood by officials at the Department or component that made the award.”)

instructions state only that use cases “will be linked at ai.gov.”).⁶⁷ OMB should support efficient, centralized information-finding for the public and government actors alike, consistent with OMB’s long standing open data policies.⁶⁸ This should include:

- OMB should **require the creation and maintenance of a central, consolidated inventory, with requirements to ensure consistency of data, ease of access, and the capacity to grow** to accommodate more frequent reporting.
- OMB should **publish a list of agencies that have yet to submit their required inventories** on AI.gov or another centrally findable location.
- **Create a mechanism for users to receive alerts** about updates and modifications to the consolidated inventory.
- **Make use cases accessible through programmatic means, like an API**⁶⁹
- **Post the consolidated inventory data at data.gov**
- **Maintain an archive or other permanent historical record** of the inventory for users to see changes to the inventory over time
- **Support agencies that wish to voluntarily disclose or update** their uses on a regular basis.

Recommendation #3

Establish AI inventories as the central place to report whether a use-case is deemed safety- or rights- impacting; whether a waiver was secured; and to publish any AI Impact Assessment

1. **Require agencies to report in the AI inventory whether or not an AI use case was deemed safety- or rights-impacting, with a justification**

As noted in our response to Question 5, the draft Memo or implementing guidance should require agencies to **report in their AI inventories whether or not a use case has been deemed safety- or rights-impacting in its AI inventory, with a justification**. Such a step will help OMB, other agencies, and the public understand how agencies are interpreting their obligations under the Memo and consolidate this information in a centralized location for ease of use. Agencies’ AI inventories should also cite the regulatory authorities that apply to each use case, as this will provide further insights into whether the use case meets the guidance’s definitions of rights- or safety-impacting AI.

⁶⁷ *Guidance for Creating Agency Inventories of AI Use Cases Per EO 13980*, supra note 58.

⁶⁸ See, e.g. Office of Management and Budget, Exec. Office of the President, M-13-13, *Open Data Policy—Managing Information as an Asset* (2013), https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/memoranda/2013/m-13-13.pdf.

⁶⁹ See, e.g. api.data.gov, <https://api.data.gov/> (last visited Nov 29, 2023).

2. Require agencies to report in the AI inventory waivers from minimum practices, and OMB to conduct annual look-back of waiver use

Under the draft memo, agency Chief AI Officers may grant waivers to the minimum practices, and must report to OMB detailing the waivers' scope, justifications, and supporting evidence. To ensure waivers are being appropriately and consistently evaluated across agencies, OMB should:

- Require agencies to report in the AI inventory entry for a use case whether a waiver was sought from minimum practices. Reports for rights- or safety-impacting systems that are granted waivers should be accompanied by an explanation of why such a waiver was granted. By publishing their written determinations, agencies will be acting with a commitment to transparency and accountability to the public, and help other federal agencies see what practices meet the standard set by OMB for waiver.
- OMB should evaluate the reports it receives from each agency to confirm whether reported uses are rights- or safety-impacting, and **publish an annual report analyzing the extent to which waivers are being granted appropriately for such uses.**

3. Require agencies to publish AI Impact Assessments in a manner that is findable via their AI inventory submission

For safety- and rights-impacting uses, OMB should require agencies to **publish their AI Impact Assessment (AIA) in a manner that is findable via their inventory submission** (for example, with a permanent link to an AIA for each use case included in their inventory submission). Public AIAs should, among other things, report any raw figures related to the agency's independent evaluation, like the number of items in each training dataset used by the system being evaluated, a breakdown of the demographic groups for whom the agency measured the system's impacts, and aggregated statistics about the system's outcomes specific to each demographic group. An abbreviated version may also suffice, should it contain answers to questions of high public interest like any prohibited uses of the AI, who has access and ownership, the number of users, any known causes of failure, and, if applicable, the vendor's terms of service.⁷⁰

Recommendation #4

Prepare agencies for public requests for information

Lastly, agencies should be prepared for public follow up for further information. While OMB already instructs agencies to include an email address for public input in their inventories, it is

⁷⁰ See, e.g. a proposal for Ecosystem Graphs, which includes detailed fields for transparency in foundation models, Rishi Bommarani et al., *Ecosystem Graphs: The Social Footprint of Foundation Models*, Stanford Institute for Human-Centered Artificial Intelligence (2023), <https://arxiv.org/pdf/2303.15772v1.pdf>.

important that one or more individuals within each agency is formally designated to monitor that inbox and respond to inquiries and feedback in a timely manner. Agencies need to be responsive to public feedback on agency use cases.

Conclusion

We applaud the Office of Management and Budget (OMB)'s thoughtful work in developing the draft Memo, and its openness to receive public feedback on the proposed memorandum.

In revising the draft Memo, OMB should recognize the value and significance of this guidance, which, with effective execution, will play a significant role in effectuating the Administration's commitment to advancing safe, secure and trustworthy AI. The modifications suggested in this document are each designed to more effectively achieve the Memo and Executive Order's stated goals, by providing clarity and actionable guidance for agency decisionmakers, supporting consistency across agency interpretations, and easing administrative burden for OMB by prioritizing upfront guidance rather than post-hoc compliance efforts once agencies' work has already begun.

OMB's Memorandum and the resulting agency practices have the opportunity to establish the Federal government as a leader in responsible AI development, funding, and use – creating a model that will be looked to by state, localities, and other governments around the world. CDT thanks OMB for focusing on the important issue of advancing robust AI governance and strengthening U.S. leadership on responsible, rights-respecting innovation.