Guidance on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum: Explanation and Response to Public Comments

March 28, 2024

On November 1, 2023, the Office of Management and Budget (OMB) published a draft memorandum titled *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence* and solicited public comments through written submissions on Regulations.gov. Alongside the draft guidance, OMB published a Request for Comment that listed eight questions suggesting potential topics for feedback.[1] OMB received 208[2] written submissions during the public comment period, which lasted from November 2, 2023 to December 5, 2023.[3] OMB has now issued a memorandum that incorporates many of the perspectives and suggestions offered during the comment period. This document summarizes the comments received and OMB's responses.

I. Scope and Applicability

A. Applicability to National Security Systems, the Department of Defense, and the Intelligence Community

The draft guidance explained that OMB's memorandum does not cover AI when it is being used as a component of a National Security System (NSS)[4]. A number of commenters characterized the definition of NSS as being too broad and expressed concerns that an agency could exempt systems that serve a de minimis national security purpose from implementing the guidance. For example, commenters offered that a potential misuse of the rule would shield AI deployed for rights-impacting or safety-impacting use cases, such as those in law enforcement and immigration.

Further, some commenters asserted a need to clarify how "dual-use systems," or agency systems that (1) perform both national security and non-national security functions; and (2) are used by both civilian agencies and the Intelligence Community, fit into the scope of the guidance. Some commenters urged OMB to adjust the guidance to explicitly cover such systems, subject to the limited and existing carveouts for when the AI is used by an element of the Intelligence Community. Other commenters called for the removal of any carveouts for the Department of Defense or the Intelligence Community.

[1] The Request for Comment can be viewed at https://www.regulations.gov/document/OMB-2023-0020-0001.
[2] A small number of submitted comments were not publicly posted due to the presence of certain information not suitable for public posting (*e.g.*, social security numbers, addresses, passwords) or were withdrawn by the submitter.
[3] Although a few commenters requested an extension of this comment period, OMB determined that the period was sufficient, as reflected by the numerous, varied, and detailed comments, including from entities that requested an extension.
[4] National Security Systems has the meaning established in 44 U.S.C. § 3552(b)(6).

OMB takes seriously these comments and acknowledges the public's concern on this topic. OMB strongly believes that issues of AI governance, innovation, and risk for NSS must be managed appropriately. But given the legal authorities under which this guidance is being issued, it would be inappropriate to include NSS within its scope. In particular, Section 10.1(i) of Executive Order 14110 directs that, "[t]he initial means, instructions, and guidance issued pursuant to subsections 10.1(a)-(h) of this section"—which include this OMB memorandum—"shall not apply to AI when it is used as a component of a national security system, which shall be addressed by the proposed National Security Memorandum described in subsection 4.8 of this order." The Executive Order further directs that the National Security Memorandum shall require "specific AI assurance and risk-management practices for national security uses of AI." Additionally, the designation of a system as an NSS is not a determination for an agency to make lightly. Federal law defines and limits the grounds upon which a system may be classified as an NSS. Statutes and governmentwide policies, such as National Security Directive 42 and National Security Memorandum 8, establish governance requirements and stringent cybersecurity controls for any system so classified.

In addition, none of the requirements imposed or the authorities conferred by the Advancing American AI Act reach the Intelligence Community. The Act's inventory requirement for AI use cases also does not apply to the Department of Defense.[5] Given these express limitations and the direction provided in Executive Order 14110, OMB has not implemented commenters' suggestions to alter these exemptions.

B. Scope of OMB Authority and its Oversight Role

One commenter expressed doubt about OMB's statutory authority to require agencies to stop the use of AI that is non-compliant with the memorandum's risk management practices. As the memorandum notes, its directions to agencies, including requirements to stop non-compliant use, are supported by the AI in Government Act of 2020, which in relevant part directs OMB to issue a memorandum to the head of each agency regarding various aspects of their use of AI,[6] and the Advancing American AI Act, which directs OMB to "develop an initial means by which to … address any other issue or concern determined to be relevant by the Director to ensure appropriate use and protection of privacy and Government data and other information."[7] Those specific and more recently enacted statutory authorities complement OMB's longstanding authority to establish binding government-wide policies for the management of information resources, including information technology.[8]

Other commenters recommended that OMB impose additional consequences for agencies' failure to comply with the memorandum's requirements. Relatedly, commenters suggested that OMB

[5] Pub. L. No. 117-263, div. G, title LXXII, subtitle B, §§ 7225(d), 7228 (codified at 40 U.S.C. 11301 note), https://www.congress.gov/117/plaws/publ263/PLAW-117publ263.pdf.
[6] Pub. L. No. 116-260, div. U, title 1, § 104(a) (codified at 40 U.S.C. § 11301 note).
[7] Pub. L. No. 117-263, div. G, title LXXII, subtitle B, § 7224(a), (d)(1)(B) (codified at 40 U.S.C. 11301 note).
[8] *See, e.g.*, 44 U.S.C. §§ 3501–3521.

should condition agency budget approvals and funding recommendations on (1) compliance with the memorandum's substantive requirements; and (2) the agencies' submission of documentation that they are using dedicated funding directly and solely toward AI risk management.

OMB has multiple methods for conducting oversight of agency activity, particularly when it comes to tracking implementation of OMB memoranda. However, OMB does not view it as necessary or appropriate to precommit to the imposition of these particular consequences, particularly when they are not specified by the relevant statutes or executive orders governing agency use of AI.

C. Overall Scope of the Memorandum

A few commenters focused on the definition of AI, as it determines the scope of the memorandum. One commenter characterized the memorandum's definition of AI as too narrow and inconsistent with the definition in Executive Order 14110. Relatedly, other commenters recommended that the definition of AI be clarified or adjusted to expressly include items such as automated decision-making systems (ADS).

Regarding the memorandum's definition of AI, OMB has chosen to use the definition from the AI in Government Act of 2020, which in turn cross-references section 238(g) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019 (10 U.S.C. § 2358). OMB has done so to ensure alignment with the definition contained in the AI in Government Act of 2020, *i.e.*, the statute directing OMB to issue this guidance. OMB also did not incorporate the recommendation to include ADS in the definition of AI, as ADS does not always rely on AI to perform a certain task. When agencies use AI-enabled ADS, such systems would fall within scope of the memorandum.

Additionally, one commenter stated that the memorandum's scope should be expanded to expressly protect all persons in the United States regardless of citizenship or immigration status. This memorandum, however, regulates agencies' own use of AI, rather than members of the public, and thus the memorandum deliberately does not define (let alone limit) its scope by reference to members of the public who may be affected by agencies' use of AI.

D. Definitions and Presumed Rights-Impacting and Safety-Impacting Use Cases

Some commenters requested that OMB revisit the definitions for rights-impacting AI and safety-impacting AI, characterizing some of the provisions as too broad and capable of being interpreted in a way that sweeps in innocuous use cases. Several commenters also suggested additions to the lists of presumed rights-impacting or safety-impacting uses in the draft guidance. These recommended additions included specific applications of AI, such as its use in a drug safety mechanism or to predict the location of a crime, as well as whole categories of uses, such as any use of AI in healthcare. Based on these recommendations, OMB has included some additional items in its lists of uses presumed to impact rights or safety, removed other uses, and adjusted

certain language in the definitions and appendix to clarify the intended scope of the rights- and safety-impacting categories. When assessing the potential inclusion of additional uses, OMB took into account factors such as (1) whether the action is likely to be undertaken by a Federal agency; (2) whether the use is sufficiently discrete and concrete, so as to be susceptible to risk-mitigation measures, as opposed to a general reference to a broad category of uses; and (3) whether the use is already adequately captured in the draft lists. And although OMB declined to include several of the proposed uses in the lists, the memorandum makes clear that those lists are not designed to catalog the entire range of uses that may be rights- or safety-impacting.

Additionally, commenters requested that OMB establish a process to review, adjust, and remove presumed uses to account for the evolving nature of risks and AI's future technological advancements. OMB has existing processes in place to update its memoranda as a general matter, and the statutes directing OMB's issuance of this particular memorandum require periodic updates.[9] There is accordingly no need for the memorandum itself to provide for updates. However, OMB has moved the lists of presumed rights-impacting AI and safety-impacting AI uses from the memorandum's main text to an appendix, where they can be more easily updated if revisions are warranted.

Other commenters recommended OMB take a different approach entirely to the identification of rights-impacting AI and safety-impacting AI. Rather than have a baseline definition for rights-impacting AI and safety-impacting AI, with a further set of presumptions for particular uses, commenters suggested an exhaustive list of rights-impacting or safety-impacting AI that would focus on consequential decisions. Alternatively, other commenters recommended an adjustment of the memorandum's structure to instead focus on a tiered risk structure (e.g., high risk, limited risk, low risk), similar to that found in the European Union's Artificial Intelligence Act. OMB did not adopt the suggestion to maintain an exhaustive list of rights-impacting or safety-impacting AI. AI technology and use cases continue to evolve rapidly. Because the Federal Government is responsible for a vast array of missions, it would be impossible to have a fully comprehensive list of all use cases that could impact rights or safety. Risk from the use of AI is heavily context dependent, making the categorical designation of a particular use case as safety-impacting or rights-impacting across all agencies an inferior alternative. Similarly, OMB did not implement the recommendation to shift to a tiered risk structure. The methodology adopted in the memorandum is a more streamlined approach, reflecting OMB's initial analysis to determine particular uses that meet a presumptive "high-risk" threshold, as they significantly impact rights and safety.

E. Intersection of OMB's Guidance and Existing Law

Some commenters expressed concerns about the relationship between OMB's risk management practices—particularly the actions that focused on notice, explanation, and remedy— and existing constitutional and statutory provisions, such as the Due Process and Equal Protection Clauses of the U.S. Constitution and the Administrative Procedure Act. For example, some commenters

[9] Pub. L. No. 116-260, div. U, title 1, § 104(d); Pub. L. No. 117-263, div. G, title LXXII, subtitle B, § 7224(d)(3).

expressed concern that a waiver of the memorandum's requirement to "notify negatively impacted individuals" could enable the agency to waive notice requirements mandated by other sources of law. Other commenters proposed stricter or more detailed notice requirements, while still others expressed concern that the draft memorandum's notice requirements would be overly burdensome to agencies.

As the memorandum explains, the required risk management practices supplement—rather than supersede or modify—existing legal requirements, including, for example, those regarding notice to affected individuals that might be required by due process. To help clarify, the final memorandum includes language that encourages agencies to make use of any existing processes related to these requirements and offers recommendations on how the guidance's AI-specific reporting may interact with these processes.

To mitigate cited concerns, some commenters recommended bolstering notice requirements for individuals prior to agency use of rights-impacting AI or safety-impacting AI, to include expanding on the criteria for what agencies are expected to provide in such notices. However, given the varied nature of government decisions that rely on AI, a one-size-fits-all notice requirement is impractical. Moreover, as noted above, agencies still remain subject to other applicable laws governing notice, and so agencies may be subject to more stringent notice requirements depending on the agency action at issue or the degree to which AI played a role in the action.

F. Applicability to Federal Agency Inspector Generals

One commenter recommended that the guidance be more explicit in its application to and effect on each agency's office of the Inspector General (OIG) and OIG's authorities under the Inspector General Act of 1978, 5 U.S.C. §§ 401-424, as amended. In particular, the commenter emphasized the need for OIGs to retain their independence and the ability to administer the use of AI independent of an agency's Chief Artificial Intelligence Officer's oversight and control. While the memorandum does not explicitly address the role of OIGs within Federal agencies, it notes that "[a]ll agency responsible officials retain their existing authorities and responsibilities established in other laws and policies."

II. Agency AI Governance and the Role of the Chief AI Officer

G. Chief Artificial Intelligence Officer Qualifications

A few commenters recommended that OMB provide additional standards or requirements for the qualifications of an agency Chief Artificial Intelligence Officer. Some asserted that agencies with a significant law enforcement mission should require their Chief Artificial Intelligence Offers to have expertise in civil rights and civil liberties, digital safety, and technology-related ethics. Other commenters requested a more specific list of qualifying educational experience and technical

expertise. Lastly, commenters suggested that while agencies can initially "dual hat" existing officials in the CAIO role, that should be temporary until agencies can create AI offices led by CAIOs whose primary roles are leading that function.

OMB recognizes the importance of having qualified experts serve in an agency's senior leadership role and the value of interdisciplinary expertise when managing a technology such as AI, which has the potential for large societal impact. OMB also understands that the memorandum applies to a diverse set of agencies that vary in size, mission, budget, expertise, and many additional dimensions. OMB's goal is to provide guidance that is applicable across the entire Federal Government and empower agencies to designate Chief Artificial Intelligence Offers that can perform the responsibilities outlined in the memorandum. Further, OMB's memorandum recommends, and in some cases directs, engagement and consultation with relevant experts both inside and outside the agency to bring in a range of diverse perspectives.

H. Chief Artificial Intelligence Officers' Waiver Authority

Some commenters expressed concern that a Chief Artificial Intelligence Officer would hold too much discretion to waive the applicability of risk management requirements to particular AI uses cases. Commenters offered a range of recommendations to address this, including the public release of agency waiver requests, periodic reviews of waivers, time limits for waivers, or the elimination of waivers entirely. To address these concerns, OMB has updated the memorandum's language to improve the transparency of agency waivers and introduce additional considerations for oversight. For example, OMB is requiring the agency's Chief Artificial Intelligence Officer to annually certify the ongoing validity of any granted waivers and publicly release a summary detailing each individual waiver and its justification. OMB will be providing further instructions regarding these summaries.

I. Agency Compliance Plans

Commenters offered a range of opinions on the content and release of agency compliance plans. Some commenters recommended that the memorandum require agencies' compliance plans to identify the officer directly responsible for carrying out each element of the plan, and as the plan is implemented, disclose the officer directly responsible for each significant decision, action, or omission that occurs in implementing the plan. Other commenters expressed the view that the memorandum should not direct public disclosure of agencies' compliance plans, inventories of AI use cases, or agency strategies, as such disclosures could enable threat actors to conduct malicious activities.

OMB did not incorporate the recommendation to require identification of directly accountable individuals. The memorandum already assigns agency Chief Artificial Intelligence Officers responsibility for developing the compliance plan. Further, a static list of individuals in a public document could cause confusion and would not meaningfully increase accountability as the

workforce fluctuates or responsibilities change. Agencies are encouraged to capture such information internally, as is consistent with the recommendations in the National Institute of Standards and Technology's AI Risk Management Framework.

Regarding commenters' proposed restrictions on sharing agency compliance plans, the AI in Government Act of 2020 requires their publication, so OMB's memorandum must do the same.[10] Regarding the sharing of AI use cases, agency strategies, or other materials, the memorandum contains caveats reflecting that agencies have existing processes governing the public release and withholding of agency documents, which may include assessments for potential adversarial use.

III. Innovation

J. Strengthening Innovation Through Transparency and Government Reuse

Some commenters noted that open approaches to innovation are critical for responsible, accountable, and transparent use of AI for the Federal Government. These commenters requested that agencies share AI code, models, and data for public benefit, in some instances suggesting that the memorandum explicitly rely on OMB Memorandum M-16-21[11]. OMB appreciates commenters' focus on open source. The final memorandum includes a new subsection stating that, subject to specified exceptions, agencies must share custom-developed code for AI (including models and model weights) to promote collaboration Government-wide and with the public. This additional subsection aligns with requirements in M-16-21 and the Open, Public, Electronic and Necessary (OPEN) Government Data Act.[12]

K. Harmonizing Compliance Requirements and Integrating AI into Federal Security Authorization Processes

Many commenters focused on the benefits of harmonizing and standardizing agency implementation of the draft memorandum's risk management practices. Particularly, commenters shared that through standardization, both vendors and agencies (1) would have greater clarity on compliance with the risk management practices; and (2) could focus resources on rights-impacting and safety-impacting use cases. In an effort to avoid duplication of efforts, commenters encouraged a "do once, use many" approach, such as through the use of common forms where a vendor can complete a single form and it would be accepted by any agency to demonstrate compliance with Federal AI policy. In response, the final memorandum contains new provisions that promote such efficiencies as well as opportunities for sharing resources and best practices among agencies. This includes an action for OMB to coordinate, through an interagency council of Chief Artificial

[10] Pub. L. No. 116-260, div. U, title 1, § 104(c).
[11] *See* OMB Memorandum M-16-21, *Federal Source Code Policy: Achieving Efficiency, Transparency, and Innovation through Reusable and Open Source Software* (Aug. 8, 2016), https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/memoranda/2016/m_16_21.pdf.
[12] Pub. L. No. 115-435, 132 Stat. 5529 (2019), https://www.congress.gov/115/bills/hr1770/BILLS-115hr1770ih.pdf.

Intelligence Officers, knowledge and sharing of resources such as templates, examples of key documents and user guides, lessons learned, and technical information on testing and evaluation.

Additionally, commenters requested that the memorandum's AI risk management requirements be integrated into Federal cybersecurity authorization processes, such as FedRAMP, to work through a centralized approving authority. FedRAMP is a Federal security authorization program that establishes a standard assessment and approval process for the use of cloud services, resulting in an authorization package that agencies use to validate a vendor's security claims. In contrast, OMB's guidance is not predominantly focused on broadly applicable security controls, but rather the context-specific nature of how AI is used in the wide array of agency missions and the impact of using AI to make or inform agencies' decisions and actions. For example, the requirement to assess the impact of use of a text-summarization system and test it in a real-world context will differ significantly depending on whether it is used to summarize news articles or to review applications for veteran's benefits. Accordingly, OMB has determined that the FedRAMP process is not a suitable vehicle for validating compliance with the guidance's context-specific risk management requirements, due to the fundamentally different purposes they serve.

Through the updates made to the memorandum based on public comment and agencies' forthcoming implementation of OMB's requirements, OMB is striving to ensure consistent implementation of the new risk-management requirements, while also respecting the unavoidable differences in how AI systems may be used in different contexts and for different purposes.

L. Balancing Recommendations to Improve External Workforce Hiring with Recommendations to Upskill and Protect Existing Federal Workers

Some commenters requested that OMB include considerations for Federal employee job quality and recommended that agencies provide opportunities for upskilling employees on AI concepts, which would allow incumbent workers to grow their skillsets and advance their public-sector careers. Commenters suggested that investments in AI training opportunities for frontline Federal workers who interact with the technology would also help OMB achieve more effective and efficient implementation of its guidance. OMB is supportive of efforts to upskill incumbent workers and maintains language in the memorandum that encourages agencies to pursue such opportunities.

IV. Risk Management

M. Overpromoting AI as a Solution

A few commenters expressed concern that the draft guidance pushed agencies to pursue wholesale adoption of AI, rather than a balanced approach to assessing an array of potential technical and non-technical solutions. Commenters proposed that agencies document the expected benefit of AI relative to non-AI options and suggested that the memorandum avoid language directing agencies

to work towards "AI maturity." Other commenters requested that OMB provide guidance to help agencies focus on project requirements and relevant mission areas, instead of the "AI hype cycle." The draft memorandum, however, contained language—which remains in the final text—that directed agencies to make evidence-based determinations about whether AI was suited for a given task. The final memorandum also includes language to emphasize that agencies are encouraged to prioritize AI development and adoption for the public good and where the technology can be helpful in understanding and tackling large societal challenges, such as using AI to improve the accessibility of government services, reduce food insecurity, address the climate crisis, improve public health, and grow the nation's economic competitiveness in a way that benefits people across the nation.

Relatedly, some commenters were concerned that efforts to elevate AI innovation would result in a wave of outsourced AI contracts. Commenters suggested that the memorandum should establish a presumption that work related to the government's continued design, development, testing, and operation of AI will be insourced whenever doing so is feasible. OMB did not incorporate this edit, as it would be impractical and intrusive to set broad resourcing requirement given the wide variation in agency use cases. However, the memorandum includes text that encourages agencies to consider a comprehensive strategy for filling their workforce gaps, which includes a recommendation on upskilling existing staff to assist in AI and AI-enabling roles. This builds on language in Executive Order 14110 directing a Government-wide AI talent surge to accelerate the placement of key AI and AI-enabling talent in high-priority areas and to advance agencies' data and technology strategies.

N. Further Considerations for Facial Recognition Technology

A few public commenters expressed concern about the reliability and fairness of facial recognition technology, citing prominent press coverage of controversial outcomes from law enforcement use. In addition, some commenters had particular reservations about the process for training facial recognition and other biometric technologies, which may involve obtaining images from online sources without the originator's consent.

In response, OMB added a new recommendation in the memorandum on responsibly procuring AI for biometric identification. For example, agencies are encouraged to validate the accuracy, reliability, and validity of the AI's ability to match identities. OMB also expanded the list of presumed rights-impacting AI uses to include "Conducting biometric identification of a suspected criminal, known criminal, missing person, or victim, or for one-to-many identification in publicly accessible spaces."

O. Risk Management Practices

Based on several comments, the final memorandum provides additional clarifications on a number of the guidance's risk management practices. For example:

i. OMB has decoupled the requirements for continuous monitoring and periodic human review. The continuous monitoring requirement predominantly focuses on technical performance monitoring, where agencies must detect changes to a system's functionality. By contrast, periodic human review is a more holistic assessment where agencies must regularly evaluate risks from the use of AI to determine whether the deployment context, risks, benefits, and agency needs have evolved over time.

ii. OMB has provided additional clarification on the requirement to consult and incorporate feedback from affected groups in important contexts such as fraud prevention and law enforcement investigations, where consulting with the targeted individual is impractical but consulting with a representative group may be appropriate. Additionally, OMB clarified that an agency is not required to terminate use of AI in the event of negative feedback from affected groups or the public, but rather, if an agency assesses the referenced feedback and has determined that the use of AI in a given context would cause more harm than good, the agency should not use the AI.

iii. OMB has offered greater clarity on the opt-out requirement for AI-enabled decisions. Where practicable and consistent with applicable law and governmentwide guidance, agencies must provide a mechanism for individuals to opt-out of the AI functionality in favor of a human alternative. However, the guidance identifies limited circumstances where such opt-out requirements may not be appropriate—when, for example, the AI functionality is solely used for the prevention, detection, and investigation of fraud[13] or cybersecurity incidents, or the conduct of a criminal investigation.

iv. OMB has added language to clarify that where agencies already maintain an appeals or secondary human review process (for example, a statutory appeal mechanism for a denial of benefits), agencies can make use of and expand such processes, as appropriate, to meet the memorandum's related AI-specific requirements.

P. Reporting on an Agency's Data Outputs Over Third-Party Data Inputs

A number of commenters suggested that OMB's risk management practices that direct actions on data governance and assessment should focus on data "outputs" rather than data "inputs." The commenters asserted that requirements to exhaustively document an AI's training data inputs, provenance, and reliability would be infeasible for certain types of AI, particularly ones that were trained on datasets that measured hundreds of gigabytes. Commenters also expressed the concern that the draft disclosure requirements would implicate a vendor's proprietary information and could require revealing trade secrets.

These commenters recommended a few solutions. First, they suggested that OMB should revise the data requirements to focus on summary information sufficient to document robustness (i.e., the ability of a system to maintain its level of performance under a variety of circumstances, such as

[13] As noted in the memorandum, some uses of AI in these categories, such as the use of biometrics for identity verification, may be subject to requirements in other guidance that would necessitate an option to opt-out, and the AI memorandum does not replace, supersede, or otherwise interfere with any such requirements.

if new data were introduced) and how the data is fit for purpose. Alternatively, commenters stated that the memorandum should be re-scoped and only apply to the agency's use of data, for example when an agency fine-tunes a vendor's base model. This would focus the data requirements on the data an agency is contributing for training and use, rather than the underlying data a vendor used to produce and maintain an AI.

OMB acknowledges the commenter feedback on the need to protect intellectual property and proprietary information. The memorandum's data-focused requirements call for "sufficient descriptive information," which OMB believes captures the suggestion to focus on detailed summaries for input data. However, OMB did not remove all considerations related to vendor-provided data, as research has shown that flawed training data significantly influences the AI's operational performance downstream.

Alternatively, some commenters provided recommendations on expanding data reporting, to include data from third parties. This included comments suggesting that the characterization of AI risks include considerations for copyright and other intellectual property generally, and that provenance documentation should include information about any licenses obtained to use training materials. Other commenters noted that agencies using safety-impacting and rights-impacting AI systems should be required to keep a record of data provenance, procurement, preprocessing, storage, and security to facilitate auditing. The final memorandum includes reporting requirements for data provenance and preparation, among other relevant data reporting measures. Recordkeeping for topics such as procurement or system security measures is already governed by an array of existing laws and policies, and therefore is not addressed specifically in the memorandum.

Q. Agency's Ability to Collect and Apply Demographic Data for Disparity Assessments

Some commenters expressed concern that agencies would face barriers in obtaining the necessary demographic data to conduct the memorandum's required disparity assessments, given the multi-step processes required to ensure that data collection complies with the Paperwork Reduction Act (PRA). These commenters supported the goal of having secure, privacy-preserving standards. The commenters also wanted to ensure that collecting demographic data for the purpose of assessing AI systems for bias is not itself a form of discrimination. One commenter called on OMB to issue further guidance to agencies to help agencies navigate risks of litigation and other legal risks when collecting data in regulated or sensitive contexts. In relation to the comments on the PRA, commenters also recommended the amendment of OMB's guidance on the Privacy Act of 1974 to eliminate barriers to linking demographic information to records of individuals for the purpose of reducing bias. OMB recognizes the relevance of considering the relationship between the PRA and Privacy Act on this subject; however, the recommended approaches are beyond the scope of this memorandum.

In contrast, some commenters advocated against an expanded requirement for agencies to collect and store demographic data, citing concerns about data minimization requirements and whether information about an individual would be used for a purpose other than what it was originally collected for. Drawing from lessons learned from OMB's work on implementing the President's equity executive orders and Equitable Data Working Group, OMB adjusted the AI memorandum to focus on actions agencies may take using data to which they have ready access—actions such as documenting when they are using data that contains information about a class protected by Federal antidiscrimination laws and putting controls in place to mitigate disparities that lead to unlawful discrimination or harmful bias.

R. Revisiting Deadline for Implementation of Risk Management Practices

Many commenters expressed concern that the draft memorandum's proposed August 1, 2024 deadline was too aggressive for agencies to successfully implement the minimum risk management practices in Section 5(c). Commenters noted that this timeline does not account for OMB's forthcoming action on acquisition of AI systems and services, which OMB has been directed to take by September 2024 (180 days after EO 14110's signing). In response, OMB has moved the deadline to achieve compliance with the guidance's minimum practices to December 1, 2024 in the final guidance, providing agencies with an additional 120 days.

S. Prohibiting Select Artificial Intelligence Use Cases and Limiting the Use of Artificial Intelligence Until Certain Conditions are Met

A number of commenters requested that the memorandum identify prohibited uses of AI. Some commenters suggested that AI be prohibited if, for example, it does not demonstrate scientific validity or undermines human rights or the rule of law. Commenters also recommended specific use cases that should be prohibited, such as predictive policing systems, emotion recognition systems, biometric categorization systems, and predictive risk systems in the law enforcement and immigration context.

Executive Order 13960 directs agencies to follow a range of relevant principles when designing, developing, acquiring, and using AI, including the proposition that AI must be accurate, reliable, effective, purposeful, lawful, and respectful of our Nation's values. Elaborating on those principles, the OMB memorandum provides practical guidance and requirements to ensure that agencies use AI only when its risks are manageable, and that the agencies do, in fact, manage those risks. The memorandum directs agencies not to use AI when its benefits do not meaningfully outweigh the risks. OMB did not incorporate the recommendation to prohibit select AI use cases, as the memorandum addresses the concerns cited while balancing the need to allow the use of AI when it is beneficial.

Relatedly, some commenters requested that the memorandum instruct agencies not to begin removing barriers to AI implementation, or to move forward with AI implementation, until after

the preliminary actions set forth in Executive Order 14110 have been completed—especially those related to the use of AI in health care. OMB did not incorporate this recommendation, as promoting AI innovation is a core objective of the memorandum, as directed both by statute and by Executive Order 14110, and the memorandum's provisions are intended to strike the appropriate balance between benefitting from the use of AI and managing its risks.

T. Environmental Risk

Some commenters recommended that OMB expand the memorandum's risk assessment requirements to include considerations for environmental impact, particularly when agencies use or benefit from using large amounts of compute power to train complex AI models. These commenters offered examples where populations or groups could be negatively impacted by environmental degradation. In response, OMB has introduced language into the memorandum's procurement recommendation section encouraging agencies to consider the environmental impact of computationally intensive AI services. In addition, OMB included text that elevates the importance of prioritizing AI for the public good and encourages agencies to look to AI that can help tackle large societal challenges, such as the climate crisis.

U. Leveraging Existing Best Practices

A few commenters suggested that the memorandum should incorporate the Organisation for Economic Co-Operation and Development's Principles on Artificial Intelligence and the NIST Artificial Intelligence Risk Management Framework. Section 5(c) of the memorandum encourages agencies to supplement the initial baseline set by OMB's guidance with additional best practices such as those developed by NIST and other reputable organizations.

V. Organizational Transparency, External Engagement, and Diversity

V. Advancing Diversity and Equity

Many commenters emphasized the need to further elevate equity and inclusivity as outcomes for public sector AI. In response, the final memorandum provides additional language to promote the institutionalizing of diverse perspectives—including from underserved communities, such as people with disabilities, low-income neighborhoods, and LGBTQ individuals—in a range of areas, such as in the AI governance process and in the development of enterprise AI strategies. In particular, OMB has clarified that the role of the Chief Artificial Intelligence Officer is critical to promoting equity and inclusion within agencies' AI governance structures and is responsible for the incorporation of diverse perspectives into decision-making processes. Additionally, OMB is encouraging agency AI Governance Boards to consult external experts that may offer a range of technical, civil rights, labor relations, and sector-specific expertise. OMB has also added a requirement that agency AI strategies publicly articulate plans for encouraging diverse

perspectives throughout the AI development or procurement lifecycle, including how to determine whether an AI use case is meeting the agency's equity goals and civil rights commitments.

W. Consultation with Federal Labor Organizations and Collective Bargaining

A few commenters suggested that the guidance require agencies to consult with Federal workers and their representatives at certain milestones. This would include consultation before an agency makes a decision to procure or implement AI and prior to the selection of an AI vendor contract. These commenters also requested that the memorandum require agencies to follow the direction given in Executive Order 14003, *Protecting the Federal Workforce*, to negotiate over the subjects identified in 5 U.S.C. § 7106(b)(1). Commenters requested that engagements on collective bargaining occur at multiple points across the AI's design, development, procurement, or use phases.

Consistent with Executive Order 14003 and Executive Order 14025, *Worker Organizing and Empowerment*, OMB strongly supports the Federal Government's role in protecting, empowering, and rebuilding the career Federal workforce. Executive Order 14025 affirms that "as AI creates new jobs and industries, all workers need a seat at the table, including through collective bargaining, to ensure that they benefit from these opportunities." OMB has added language to the memorandum to help clarify to agencies select scenarios on AI use and management where they should consider applicable collective bargaining obligations and engagement with Federal workers and Federal labor unions. For example, the final memorandum: encourages Agency AI Governance Boards to consult with external experts, as appropriate, on methods for engaging the Federal workforce; explains that Federal employees (to include those represented by Federal labor organizations) may be among the groups affected by a particular agency use of AI, and if so must be consulted in the agency's design, development, and use of that AI; and affirmatively recognizes that agencies should be meeting collective bargaining obligations. Some of the commenters' labor-specific recommendations can be more appropriately addressed in the forthcoming principles and best practices for employers to be issued by the Secretary of Labor pursuant to Section 6 (b)(i) of Executive Order 14110, and the OMB memorandum encourages agencies to promote and incorporate these principles and practices when they are complete. The Administration continues to work diligently through a range of avenues to address specific comments regarding the intersection of labor relations and AI.

X. Sharing Agency Documentation with the Public

Some commenters expressed that the memorandum should require agencies to publish and seek public comment on AI documentation. In particular, commenters suggested that agencies should be required to publish the AI system's underlying source code and training data, and provide both a technical and plain language explanatory analysis of the AI system and its components. Relatedly, commenters asserted that agencies should not be given any discretion regarding providing AI documentation, and some requested particular updates to the memorandum's AI use

case inventory requirements to make them stricter and less flexible. OMB did not incorporate these suggestions, which could create conflict or confusion given the number of existing laws and governmentwide policies that already limit the extent of documentation that agencies are permitted or encouraged to share publicly. This includes limits on the publication of sensitive law enforcement information, classified information, personally identifiable information, or an agency's deliberative materials.

VI. Third-Party Requirements

Y. Procurement

Many comments that addressed procurement stressed the need for standardized technical and acquisition requirements. Particularly, a number of commenters highlighted the benefits of providing agencies with example contract clauses in order to ensure that new and existing contracts would reflect the memorandum's call for adequate data protection safeguards, particularly in terms of protecting private sector trade secrets. Some commentors further suggested requiring use of discrete pre-award and post-award benchmarks as a means of ensuring compliance without compromising efficiency. Other commenters raised related points on the importance of transparency, highlighting the requirement under the OPEN Government Data Act[14] to release a wide range of government data to the public. Generally, commenters requested application of the memorandum's risk management requirements to agency procurement of AI and emphasized the importance of striking a balance between relevant statutes and related guidance on data protection, privacy, and transparency.

Consistent with section 104 of the AI in Government Act of 2020 and section 7224(d) of the Advancing American AI Act, this memorandum does not require specific language or terms in contracts for the acquisition of AI systems or services; it instead provides recommendations for responsible procurement. As noted in the memorandum, however, OMB will be positioned to address procurement matters via a separate, forthcoming action focused on that topic, as required by section 7224(d) of the Advancing American AI Act and Section 10.1(d)(ii) of Executive Order 14110.

Z. Applicability to Federal Financial Assistance

A few commenters requested that the memorandum's risk management requirements be extended to state and local governments' administration of Federally funded programs, including by adding conditions to the award of Federal funding. Federal financial assistance is outside the scope of this memorandum. Consistent with the provisions of the AI in Government Act of 2020, the Advancing American AI Act, and Executive Order 14110 (which directs the publication of this guidance), the focus is on agencies' own use of AI. However, the final memorandum encourages agencies,

[14] Pub. L. No. 115-435, 132 Stat. 5529 (2019), https://www.congress.gov/115/bills/hr1770/BILLS-115hr1770ih.pdf.

consistent with applicable law, to consider the memorandum's minimum risk management practices when choosing criteria for the award of Federal financial assistance.

VII. Content Authentication

AA. Pursue Content Authentication for U.S. Government Information

Commenters recommended that the memorandum should require a mechanism for distinguishing genuine U.S. Government information, digital content, and documents from AI-generated fakes. Section 4.5 of Executive Order 14110 directs a series of actions for "identifying and labeling synthetic content produced by AI systems, and to establish the authenticity and provenance of digital content, both synthetic and not synthetic, produced by the Federal Government or on its behalf." This includes a separate direction for OMB to issue guidance to agencies for "for labeling and authenticating such content that they produce or publish." Therefore, these recommendations are best addressed in other OMB or agency documents.