



Center for American Progress
1333 H Street NW, Suite 100E
Washington, DC 20005

202.682.1611

americanprogress.org

December 5, 2023

Clare Martorana
U.S. Federal Chief Information Officer
Office of the Federal Chief Information Officer Office of Management and Budget
725 17th Street, N.W.
Ste. 50001
Washington, D.C. 20503

Submitted electronically via www.regulations.gov

Re: Request for Comments: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum; FR Doc. 2023-24269, 23 Nov. 2023.

Dear Ms. Martorana,

Artificial Intelligence (AI) has the potential to dramatically impact nearly all the federal government's operations and engagement with the American people. The opportunities and risks of AI, combined with the power of the state, places an even greater burden on the federal government to ensure its ethical use of AI, especially when AI is used for purposes that impact safety or rights. The federal government, its contractors, and its grantees' use of AI should only be allowed if those AI systems are effective, safe, and nondiscriminatory.¹

The Center for American Progress (CAP) has called for the Executive Branch to take immediate action on AI² and has joined the civil rights community and other allies in advocating for centering the White House Blueprint for an AI Bill of Rights (AI Bill of Rights)³ in any AI executive order (EO) and Office of Management and Budget (OMB) guidance.⁴ The President's Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence⁵ (AI EO) represents a strong start on the guidance needed to ensure the development of safe and effective AI.⁶

The accompanying draft OMB guidance "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence"⁷ is essential guidance for federal agencies and contractors that are already using AI systems.⁸ Particularly worthy of praise is the draft guidance's definition of safety-impacting and rights-impacting AI and purposes that are presumed to be safety-impacting and rights-impacting. These thoughtful definitions are an important step forward in outlining the conditions in which the use of AI requires extra attention and caution.

Below CAP provides the following response to the Request for Comments: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum; FR Doc. 2023-24269. Please contact Adam Conner (aconner@americanprogress.org: 202-669-5671) with any questions.

Sincerely,

The Center for American Progress

Adam Conner
Vice President, Technology Policy
The Center for American Progress

The draft OMB AI guidance puts principles from the 2022 White House Blueprint for an AI Bill of Rights into effect for the federal government’s use of AI

The White House AI Bill of Rights “identified five principles that should guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence”⁹ along with examples of how to implement those principles in practice. The draft OMB AI guidance integrates aspects of the AI Bill of Rights in its guidance to federal agencies within the limits of executive authority.

The integration of the AI Bill of Rights principles of “Safe and Effective Systems” and “Algorithmic and Discrimination Protections” into federal agency guidance is reflected in the minimum requirements for safety-impacting and rights-impacting AI, including: 1) AI impact assessments that include examining purpose and risks, 2) Testing and independent evaluation, 3) Mitigating emerging risks to rights and safety, 4) Tasks to ensure AI advances equity, dignity, and fairness.¹⁰

The AI Bill of Rights principle of “Human Alternatives, Consideration, and Fallback”¹¹ can be seen in the requirements for periodic human review, human training and assessment, and appropriate human consideration.¹² This direct application of the AI Bill of Rights practice for agencies using AI systems is essential to ensuring their trustworthiness and is how the government can lead by example by putting the AI Bill of Rights in practice.

Conversely, the draft OMB AI Guidance lacks needed specifics for agencies to implement the AI Bill of Rights Principles of “Notice and Explanation” and “Data Privacy.”¹³ Recommendations regarding the addition of those two principles can be found below.

Agencies must be instructed to include “Notice and Explanation” to users who encounter any AI systems

The ongoing minimum requirements for a safety-impacting or rights-impacting AI are outlined in Section 5.c.iv.D through Section 5.c.iv.H of the draft OMB guidance. In particular, Section 5.c.iv.H, “Provide public notice and plain-language documentation through the AI use case inventory,” states that

“Agencies must ensure...that the AI’s entry in the use case inventory serves as adequately detailed and generally accessible documentation of the system’s functionality that provides public notice of the AI to its users and the general public. Where practicable, agencies should include this documentation or link to it in contexts where people will interact with or be impacted by the AI.”¹⁴

This language suggests that the AI’s listing in the use case inventory is sufficient public notice of the use of the AI to the users, which seems an unlikely proposition as it would require users to know about and seek out an agency’s AI use case inventories. The recommendation for agencies to “include this documentation or link to it in contexts where people will interact with or be impacted by the AI” is a kind of notice and explanation. But the suggestion to include or link to AI use case inventories as a way to alert users of interaction with a specific AI system and to learn more information about that AI system is a poor user experience. The existing agency AI use case inventories listed on AI.gov¹⁵ tend to list all AI use cases in one large PDF or .csv file,¹⁶ making it difficult to easily digest and forcing a user to search for whatever AI system they are interacting with in a PDF or spreadsheet. A more viable alternative to this approach is detailed below.

The section on practices that must be followed on an “ongoing basis” for safety-impacting or rights-impacting AI in Section 5.c.iv.D through Section 5.c.iv.H of the draft OMB AI guidance should add a separate requirement to provide notice and transparency to a public end user that is interacting with or being affected by a specific AI system. This notice and explanation requirement should, in real-time, provide information about the specific system involved and if linking to an AI use case, should link to a specific individual use case explanation, not a list of all AI use cases. This notice and transparency requirement is in line with one of the principles of the AI Bill of Rights, “Notice and Explanation”, which states “You should know that an automated system is being used and understand how and why it contributes to outcomes that impact you.”¹⁷ As the AI Bill of Rights notes in its detailed discussion of the “Notice and Explanation” principle under “What should be expected in automated systems” these notices should be “Provide clear, timely, understandable, and accessible notice of use and explanations,” “Provide explanations as to how and why a decision was made or an action was taken by an automated system,” and “Demonstrate protections for notice and explanation.”¹⁸ Accessible notice and explanation is a fundamental expectation people should have when interacting with AI and the US Government should model its development here.

The draft OMB AI guidance does not provide guidance on “Data Privacy,” one of the principles in the AI Bill of Rights

The draft OMB AI guidance does not provide guidance on implementing data privacy practices for AI systems used by the government, contractors, or grantees from the “Data Privacy” principle of the AI Bill of Rights. The guidance’s primary mention of data in the minimum practices for safety-impacting and rights-impacting Artificial Intelligence is in Section 5.iv.A.3 “The quality and appropriateness of the relevant data” and Section 5.v.A.3 “using representative data.”¹⁹ While federal agencies are governed by existing laws, including the Privacy Act of

1974,²⁰ and understanding the need for new federal privacy legislation and to address government's use of data, there is no reason for OMB to not outline additional requirements around data privacy such as a directive to collect only strictly required information, stipulate data retention guidelines, and ensure basic data privacy protections. For example, OMB should require federal agencies to uphold the Data Privacy principle of the AI Bill of Rights by having agencies and contractors adopt new data minimization standards, such as those proposed in the bipartisan American Data Privacy and Protection Act (ADPPA) introduced in the 117th Congress.²¹

The purposes that are presumed to be Safety-Impacting or Rights-Impacting AI are among the most important contributions in the draft OMB guidance to the broader development of responsible AI

The federal government wields enormous power to potentially impact lives for the better or worse. In outlining the specific areas where AI is presumed to be safety-impacting or rights-impacting, the draft OMB AI guidance is setting out the first U.S. specific roadmap for how to think about actual scenarios where AI is poised to do the most potential harm and requires additional responsibility. This is particularly true when the entity is the federal government, which operates with tremendous power and where the public has more limited recourse than when dealing with the private sector.

As the draft OMB AI guidance only binds the federal government, contractors, and grantees in cases of AI for potential safety-impacting and rights-impacting purposes, any complaints that these minimum requirements are overly burdensome or will impact innovation should be ignored. In fact, these requirements will ensure that when the government uses AI for these purposes, innovation will explicitly not trade off against protecting safety or rights. If the approach of safety-forward innovation is adopted by those whose AI systems are used by the government, contractors, or grantees who fall under the OMB AI guidance, then that is a net innovation benefit for all.

Additional clarification and transparency are required for any determination that an AI does not meet the definition of “safety-impacting AI” or “rights-impacting AI”

Section 5b “Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights-Impacting” of the draft OMB guidance²² states:

“Where an agency currently uses or plans to use AI for a purpose described below, the CAIO, in coordination with other relevant officials as specified by the agency, may make a determination (or reverse a prior determination) that the AI application or component does not match the definitions of “safety-impacting AI”²³ or “rights-impacting AI”²⁴ and is therefore not subject to the minimum practices. The agency CAIO may make or reverse this determination only with a documented context-specific and system-specific risk assessment. Any such determination or reversal must be reported to OMB within 30 days.”²⁵

The determination that AI used for purposes that are presumed to be safety-impacting or rights-impacting, but that does not meet the definition of “safety-impacting AI” or “rights impacting AI,” is a confusing one to imagine in execution. Additional clarity is needed to specify what can only be anticipated as the limited circumstances in which AI technology being used for safety-impacting or right-impacting purposes would not meet the definition of “safety-impacting AI” or “rights-impacting AI.” Otherwise, the ability to make this determination looms large as a potential loophole that could exempt vast swathes of AI that should otherwise be subject to the minimum practices for safety-impacting and rights-impacting AI. In the draft guidance the CAIO retains sole authority to determine if an AI meets the definition of “safety-impacting AI” or “rights-impacting AI” without any potential recourse for disagreement.²⁶ OMB should direct agencies to create an additional level of review that allows for a decision by the CAIO to be revisited through a process that raises the definition issue to the Deputy Secretary or equivalent that the CAIO reports to.²⁷ Finally, these definition determinations by the CAIO must only be reported to OMB, but agencies and OMB should be required to jointly make these determinations public in order to provide transparency and prevent abuse.

Clarity and transparency are required for any waivers from minimum practices for safety-impacting and rights-impacting AI

Additional clarity is needed to outline the circumstances under which agencies can seek waivers from having to meet minimum practices for either safety-impacting or rights-impacting AI.²⁸ Delineating these minimum requirements will be essential to providing critical assurance to the public and those in charge of oversight that the use of AI that is safety-impacting or rights-impacting or for safety-impacting or rights-impacting purposes will be monitored. Clarity is needed to ensure a very high bar for issuance and uniformity in that approach. Similar to the note above for definitions, the draft guidance gives an agency CAIO sole authority to issue waivers from minimum practices without any potential recourse for disagreement.²⁹ OMB should direct agencies to create an additional level of review that allows for a waiver decision by the CAIO to be revisited through a process that raises the issue to the Deputy Secretary or equivalent that the CAIO reports to.³⁰ Agencies must be required to report such waivers to OMB within 30 days, make any waiver determinations public, and publish a master list of such waivers.

OMB should establish a process to add new presumed safety-impacting or rights-impacting AI purposes

As noted above, the OMB AI guidance purposes that are presumed to be safety-impacting or rights-impacting³¹ are a critical step in ensuring a balanced assessment around the purpose and risk of agencies using AI. However, the guidance does not outline any process to add new presumed safety-impacting or rights-impacting AI purposes in the future. As AI technology continues to advance, flexibility should be provided to ensure these purposes can keep pace with new developments. OMB should establish a process to add new purposes to the initial list of safety-impacting or rights-impacting AI, possibly under the section “Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights- Impacting.”³² OMB should also create a mechanism by which the public can submit new safety-impacting or rights-impacting purposes for consideration. The responsibility for adding new purposes presumed to

be safety-impacting or rights-impacting should be driven by the interagency council to coordinate AI that will consist of the agencies CAIO's once designated in Section 10.1 of the AI EO.³³

The OMB AI Guidance does not clearly state that not adopting AI that is safety-impacting or rights-impacting should always be a strong option and that failing any part of the minimum requirements should result in an agency not using or to stop using AI

The draft OMB AI guidance allows the adoption of AI by federal agencies, so long as the process requirements in the minimum practices for AI that are automatically presumed to be safety-impacting or rights-impacting are completed. The draft guidance does not and should clearly state that not adopting AI that is safety-impacting or rights-impacting should always be a strong option for an agency. OMB should issue clearer general criteria to make clear when an AI system should not be considered as an option.

OMB should also add to Section 5.iv. "Minimum Practices for Either Safety-Impacting or Rights-Impacting AI"³⁴ to specifically state that AI that fails to meet individual components of the minimum practices should result in the agency not using or to stop using that AI.

Failing any part of the AI impact assessment should also result in an agency not using or stopping using AI. The AI impact assessment minimum practice requirement for safety-impacting or rights-impacting AI should clearly state and require that failing any of the three parts of the AI impact assessment (purpose and benefit, risk, and quality of appropriateness of the data) should result in the agency not using the AI systems. In the current draft of the AI impact assessment requirements only the second requirement "The potential risks of using AI" section explicitly states "if the benefits do not meaningfully outweigh the risks, agencies should not use the AI."³⁵ The overall requirements for "Complete AI impact assessment"³⁶ should explicitly state that failing any three parts of the assessment below should disqualify an AI system from use by the agency. Additionally, the first requirement "The intended purpose for the AI and its expected benefit"³⁷ should explicitly state that if the AI cannot be shown to meet its expected purpose or benefit, it should not be used. Similarly, the third requirement "The quality and appropriateness of the relevant data"³⁸ should make explicit that if data cannot be shown to meet the fitness and quality to meet the AI's intended purpose, the AI should not be used.

Section 5.c.iv.C "Independently evaluate the AI"³⁹ should be strengthened to explicitly state that if the independent reviewing authority does not find that the AI "works appropriately and as intended, and that its expected benefits outweigh its potential risks," the AI system should not be used.

Section 5.c.iv.D "Conduct ongoing monitoring and establish thresholds for periodic human review" states that ongoing monitoring and periodic human reviews are "to determine whether the existing implementation of the minimum practices in this section adequately mitigates any new risk."⁴⁰ This should be clarified to explicitly state that if a new, inadequately mitigated risk comes to light during monitoring or periodic human review, then a process must be initiated to mitigate the risk or to stop using the AI.

Labor and Workforce

In Section 5.b.ii, we recommend the addition of “collective bargaining, workplace organizing, union membership, or concerted activity” to part G to ensure the further protection of workers from AI systems and to continue the Biden-Harris administration's historic support for organized labor and the working people of America.

Concerns Around Disability and Civil Rights

The disability community is impacted by AI in numerous safety-impacting and rights-impacting purposes and must be top of mind as a consideration in any AI usage by agencies, the OMB guidance should reflect this. While AI has great promise for increasing accessibility in both the online and real world,⁴¹ accessibility cannot be the only issue in which the disability community is engaged. Members of the disability community have the potential to be harmed across multiple areas of presumed safety-impacting or rights-impacting purposes of AI through , and accessibility is a distinct subset of many AI issues. The draft guidance should specifically include a focus on the disability community by incorporating the following changes to the guidance.

The definitions of “Rights-Impacting AI”⁴² and “Safety-Impacting AI”⁴³ should be modified to specifically reference and include those with disabilities (additions below in **bold and underlined**):

Rights-Impacting AI: AI whose output serves as a basis for decision or action that has a legal, material, or similarly significant effect on an individual’s or community’s:

1. Civil rights, civil liberties, or privacy, including but not limited to freedom of speech, voting, human autonomy, **disability rights**, and protections from discrimination, excessive punishment, and unlawful surveillance;
2. Equal opportunities, including equitable access to education, housing, credit, employment, **disability rights, reasonable accommodations**, and other programs where civil rights and equal opportunity protections apply; or
3. Access to critical resources or services, including healthcare, financial services, social services, transportation, non-deceptive information about goods and services, and government benefits or privileges.

Safety-Impacting AI: AI that has the potential to meaningfully impact the safety of:

1. Human life or well-being, including loss of life, serious injury, bodily harm, biological or chemical harms, occupational hazards, harassment or abuse, **language access, disability**, or mental health, including both individual and community aspects of these harms;
2. Climate or environment, including irreversible or significant environmental damage;
3. Critical infrastructure, including the critical infrastructure sectors defined in Presidential Policy Directive 2143 and the infrastructure for voting and protecting the integrity of elections; or,

4. Strategic assets or resources, including high-value property, information marked as sensitive or classified by the Federal Government, and intellectual property.

Additionally, Section 5.b.ii “Purposes That Are Presumed to be Rights-Impacting”⁴⁴ should add a section addressing the impact on the disability community:

L. Impacting those with disabilities or medical conditions including but not limited to accommodations, accessibility, protections including the Americans with Disabilities Act. Agencies should be instructed to consider the specific impact on the disabled community in all safety-impacting and rights-impacting purposes. The safety-impacting and rights-impacting purposes that are likely to impact the disability community include but are not limited to Section 5.b.ii B (Law enforcement or surveillance-risk), C (immigration, asylum, or detention status), D (Detecting or measuring emotions or attention), E (in education), F (Tenant screening), G (Determining the terms and conditions of employment), H (decisions regarding medical devices), J (government benefits or services), and K (Recommendations or about child welfare, child custody).

Finally, the OMB Guidance should recommend the appointment of a Federal Chief Accessibility Officer to coordinate the implementation of AI for greater accessibility in government.⁴⁵

Generative AI as a new feature in existing approved and procured office productivity software

The leading office software suites, Microsoft 365 and Google Workplace,⁴⁶ have already announced they will be bringing generative AI to their millions of enterprise customers, which include almost every federal agency.⁴⁷ While most agencies use the government specific versions of office productivity software, they have very few if any functional differences with the commercial versions. Microsoft has already announced its generative AI roadmap for the government.⁴⁸

Section 5.v.d.v “Responsibility Procuring Generative AI”⁴⁹ appears to refer to generative AI and contracts to procure it only in the future tense. It ignores the likely widespread rollout of generative AI to the federal government within existing contracts and the addition of generative AI as a feature to office productivity software systems that are already procured by nearly every federal agency. Additional procurement guidance should be provided to clarify generally the addition of AI features to existing procured software. Specifically, given the near-term widespread availability of generative AI as a feature in every federal agency’s already procured office productivity software, OMB should require federal guidance before enabling widespread deployment of generative AI in existing office productivity software. OMB should task the Office of the Federal Chief Information Officer and the CIO Council to craft such guidance in coordination with the new interagency council of agency CAIOs.⁵⁰

Procurement

OMB should clarify in the guidance that the new AI EO and accompanying draft OMB AI guidance applies to federal agencies and their contractors. OMB should also provide clarity to ensure that federal grants, especially to state and local governments, are also covered by the Memo. AI is a new technology and vendors may be prone to making promises that are unrealistic or impossible.⁵¹ Procurement guidelines should emphasize that vendors claims should be scrutinized to pre-procurement and after-procurement and information on systems, especially false claims, should be shared across the federal government to provide maximum information for future procurements.

Elevate prior AI EO and OMB AI guidance on reporting agency statutory regulatory authorities for AI

The new draft OMB guidance “Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence” notes “The principles of this memorandum do not, by contrast, govern agencies’ regulatory actions designed to prescribe law or policy regarding non-agency uses of AI.”

While understanding regulation is not the focus of this OMB guidance, the proposed document provides an opportunity to elevate and prioritize one strong aspect of the previous Executive Order 13859 of “Maintaining American Leadership in Artificial Intelligence”⁵² from February 11, 2019 and OMB guidance M-21-06 “Guidance for Regulation of Artificial Intelligence Applications”⁵³ from November 17, 2020. EO 13859 required “the heads of implementing agencies that also have regulatory authorities shall review their authorities relevant to applications of AI and shall submit to OMB plans to achieve consistency with the memorandum” which OMB M-21-06 noted “Executive Order 13859 requires that implementing agencies with regulatory authorities review their authorities relevant to AI applications and submit plans to OMB on achieving consistency with this Memorandum. The agency plan must identify any statutory authorities specifically governing agency regulation of AI applications, as well as collections of AI-related information from regulated entities” with a template requiring “Statutory Authorities Directing or Authorizing Agency Regulation of AI Applications. List and describe any statutes that direct or authorize your agency to issue regulations specifically on the development and use of AI applications.” The Department of Health and Human Services was the only agency to publicly release their collection of AI authorities.⁵⁴ This inventory of agency regulatory authorities for AI is a strong exercise for agencies to undergo to better understand their authorities in relation to AI. This regulatory undertaking would be of immense interest to Congress, civil society, and the public as we work to understand if new AI laws are needed or if existing AI authorities are sufficient.

In order to aid all agencies in achieving the goals outlined in the President’s new Executive Order 14110 “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,”⁵⁵ the new draft OMB AI guidance should resurface the requirement from OMB M-21-06⁵⁶ for agencies to inventory their full list of potential AI regulatory authorities, add the requirement to inventory for AI regulatory authorities for the agencies and CAIOs in Section 3.a and 3.b.,⁵⁷ require that list to be submitted to OMB and the new White House AI Council

established in EO 14110 within 90 days, and the list of such agency regulatory authorities should also be made public by the agencies and the White House on AI.gov.

Questions from the Requests for Comment

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?

As noted above, the disability community is impacted by numerous purposes in the list of purposes that are presumed to be safety-impacting or rights-impacting. Too often when it comes to AI, the disability community's concerns are reduced to accessibility when the impact of AI on the disability community is much broader. The disability community's specific needs should be added as a specific consideration in Section 5(b) with language like this:

Impacting those with disabilities or medical conditions including but not limited to accommodations, accessibility, protections including the Americans with Disabilities Act. Agencies should be instructed to consider the specific impact on the disabled community in all safety-impacting and rights-impacting purposes. The safety-impacting and rights-impacting purposes that are likely to impact the disability community include but are not limited to Section 5.b.ii B (Law enforcement or surveillance-risk), C (immigration, asylum, or detention status), D (Detecting or measuring emotions or attention), E (in education), F (Tenant screening), G (Determining the terms and conditions of employment), H (decisions regarding medical devices), J (government benefits or services), and K (Recommendations or about child welfare, child custody).

In Section 5.b.ii, we recommend the addition of "collective bargaining, workplace organizing, union membership, or concerted activity" to part G to ensure the further protection of workers from AI systems and to continue the Biden-Harris administration's historic support for organized labor and the working people of America.

Finally, we recommend a category for the administration of democratic processes, including but not limited to election administration, voter registration, public comments submitted to the government, and other aspects of the democratic process overseen by federal agencies.

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

The minimum practices identified for safety-impacting and rights-impacting AI are a commendable step by the administration to set higher level safeguards for certain usages of AI in the federal government, including with contractors and grantees. However, it is critical that these determinations are made in a way that can be iterative, changed, and added to. Given the nascent yet rapidly advancing nature of the technology, there will undoubtedly be additional

uses and expanded current uses that can and should be deemed safety- or rights-impacting in the future. OMB should set out processes for evaluating new and updating existing cases in a thorough, efficient, and timely manner prior to any new or updated uses of AI in government to ensure this approach is future-proof. Specifically for existing applications of AI for new uses or with enhanced models, OMB should stipulate exact thresholds for when these are reviewed and outline the process by which agencies must do so and when.

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

The annual use case inventory should be in an easily readable format and also available in a machine-readable format that does not require significant technical expertise to access, understand, or analyze. NIST should work with the interagency council to coordinate AI in government to help prototype standards to be used for this use case inventory. Additionally, it should include the below specific information for each unique use case, even when shared among the same models or system including:

- Name, description, and developer information of AI model/system
- Exact ways the model is utilized and for which use cases
- Model inputs and outputs by data type, anonymized for maximum privacy protection
- Who the intended recipients of the model's decisions are and for what purpose
- Information regarding human in the loop on AI-generated decisions
- Last time model was assessed for rights- and safety-impacting determinations and what those outcomes were

¹ Center for American Progress and others, "Letter to the White House on Forthcoming AI Executive Order," Center for American Progress, August 3, 2023, available at <https://www.americanprogress.org/wp-content/uploads/sites/2/2023/08/Letter-to-WH-on-AI-EO.pdf>

² Adam Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence," Center for American Progress, April 25, 2023, available at <https://www.americanprogress.org/article/the-needed-executive-actions-to-address-the-challenges-of-artificial-intelligence/>.

³ The White House, "Blueprint for an AI Bill of Rights," available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights/> (last accessed December 2023).

⁴ Megan Shahi and Adam Conner, "Priorities for a National AI Strategy," August 10, 2023, available at <https://www.americanprogress.org/article/priorities-for-a-national-ai-strategy/>; Center for American Progress and others, "Letter to the White House on Forthcoming AI Executive Order;" and Center for American Progress and others, "Letter to the White House on Artificial Intelligence Executive Order," September 5, 2023, available at https://www.americanprogress.org/wp-content/uploads/sites/2/2023/09/Letter-to-WH-on-AI-Exec-Order_9-6.pdf.

⁵ Executive Office of the President, "Executive Order 14110 of October 30, 2023: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," November 1, 2023, available at <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence>.

⁶ Center for American Progress, "STATEMENT: CAP's Alondra Nelson Praises Executive Order on Artificial Intelligence," Press release, October 30, 2023, available at <https://www.americanprogress.org/press/statement-caps-alondra-nelson-praises-executive-order-on-artificial-intelligence/>.

-
- ⁷ Shalanda D. Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” available at <https://www.whitehouse.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-draft-for-public-review.pdf> (last accessed December 2023).
- ⁸ The White House, “The Government is Using AI to Better Serve the Public,” available at <https://ai.gov/ai-use-cases/> (last accessed December 2023).
- ⁹ The White House, “Blueprint for an AI Bill of Rights.”
- ¹⁰ Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Page 15, Section 5.iv.A and Page 17, Section 5.iv.cB and Section 5.iv.C, Page 17, Section 5.iv.E and Page Section 5.v.A.
- ¹¹ The White House, “Blueprint for an AI Bill of Rights.”
- ¹² Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Page 17, Section 5.iv.D and Page 17, Section 5.iv.F and Page 18, Section 5.iv.G.
- ¹³ The White House, “Notice and Explanation,” available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights/notice-and-explanation/> (last accessed December 2023); and The White House, “Data Privacy,” available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights/data-privacy-2/> (last accessed December 2023).
- ¹⁴ Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Page 18, Section 5.c.iv.H.
- ¹⁵ The White House, “The Government is Using AI to Better Serve the Public.”
- ¹⁶ Department of Health and Human Services, “Artificial Intelligence Use Cases - FY2022,” available at <https://www.hhs.gov/sites/default/files/hhs-ai-use-cases-inventory.pdf> (last accessed December 2023); and USDA, “2023 Public Inventory of AI Use Cases,” available at <https://www.usda.gov/sites/default/files/documents/2023%20Public%20Inventory%20of%20AI%20Use%20Cases%20-%20USDA.csv> (last accessed December 2023).
- ¹⁷ The White House, “Notice and Explanation,” available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights/notice-and-explanation/> (last accessed December 2023).
- ¹⁸ The White House, “Notice and Explanation.”
- ¹⁹ Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies,” Page 15-16, Section 5.iv.A.3 and Page 19, Section 5.v.A.3.
- ²⁰ U.S. Department of Justice, “Privacy Act of 1974,” available at <https://www.justice.gov/opcl/privacy-act-1974>.
- ²¹ Center for American Progress and others, “Letter to the White House on Forthcoming AI Executive Order;” and American Data Privacy and Protection Act of 2022, H.R.8152, 117th Cong, 2nd sess., (June 21, 2022), available at <https://www.congress.gov/bill/117th-congress/house-bill/8152/text>.
- ²² Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies,” Page 10-11, Section 5.b.
- ²³ Ibid, Page 25.
- ²⁴ Ibid, Page 26.
- ²⁵ Ibid, Page 10-11, Section 5.b.
- ²⁶ Ibid, Page 11, Section 5.b.
- ²⁷ Ibid, Page 7, Section 3.b.iv.
- ²⁸ Ibid, Page 14, Section 5.c.iii.
- ²⁹ Ibid, Page 14, Section 5.c.iii.
- ³⁰ Ibid, Page 7, Section 3.b.iv.
- ³¹ Ibid, Page 10, Section 5.b.
- ³² Ibid, Page 10, Section 5.b.

-
- ³³ Executive Office of the President, “Executive Order 14110 of October 30, 2023: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.”
- ³⁴ Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Page 15, Section 5.c.iv.
- ³⁵ Ibid, Page 15, Section 5.c.iv.A.2.
- ³⁶ Ibid, Page 15, Section 5.c.iv.A.
- ³⁷ Ibid, Page 15, Section 5.c.iv.A.1.
- ³⁸ Ibid, Page 15, Section 5.c.iv.A.3.
- ³⁹ Ibid, Page 16, Section 5.c.iv.C.
- ⁴⁰ Ibid, Page 17, Section 5.c.iv.D.
- ⁴¹ OpenAI, “Be My Eyes,” March 14, 2023, available at <https://openai.com/customer-stories/be-my-eyes>.
- ⁴² Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies,” Page 24
- ⁴³ Ibid, Page 25.
- ⁴⁴ Ibid, Page 12, Section 5.b.ii.
- ⁴⁵ Shahi and Conner, “Priorities for a National AI Strategy.”
- ⁴⁶ Nihal Krishan, “Microsoft rolls out generative AI roadmap for government services” *FedScoop*, October 31, 2023, available at <https://fedscoop.com/microsoft-rolls-out-generative-ai-roadmap-for-government-services/>; and Karen Dahut, “Accelerating the mission with generative AI,” *Google*, July 31, 2023, available at <https://cloud.google.com/blog/topics/public-sector/accelerating-mission-generative-ai>.
- ⁴⁷ “Google Workspace,” Fedramp Marketplace, available at <https://marketplace.fedramp.gov/products/F1206081364> (last accessed December 2023); and “Office 365 Multi-tenant & Supporting Services,” Fedramp Marketplace, available at <https://marketplace.fedramp.gov/products/MSO365MT> (last accessed December 2023).
- ⁴⁸ Krishan, “Microsoft rolls out generative AI roadmap for government services.”
- ⁴⁹ Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Page 21, Section 5.v.d.v.
- ⁵⁰ Shahi and Conner, “Priorities for a National AI Strategy.”
- ⁵¹ Sharon Goldman, “Princeton University’s ‘AI Snake Oil’ authors say generative AI hype has ‘spiraled out of control,’” *VentureBeat*, Aug. 23, 2023, available at <https://venturebeat.com/ai/princeton-university-ai-snake-oil-professors-say-generative-ai-hype-has-spiraled-out-of-control/>.
- ⁵² Executive Office of the President, “Executive Order 13859 of February 11, 2019: Maintaining American Leadership in Artificial Intelligence,” February 14, 2019, available at <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>.
- ⁵³ Russell T. Vought, “Memorandum For The Heads Of Executive Departments And Agencies,” November 17, 2020, available at <https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-06.pdf>.
- ⁵⁴ Department of Health and Human Services, “OMB M-21-06 (Guidance for Regulation of Artificial Intelligence Applications),” available at <https://www.hhs.gov/sites/default/files/department-of-health-and-human-services-omb-m-21-06.pdf> (last accessed December 2023).
- ⁵⁵ Executive Office of the President, “Executive Order 14110 of October 30, 2023: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.”
- ⁵⁶ Vought, “Memorandum For The Heads Of Executive Departments And Agencies.”
- ⁵⁷ Young, “Proposed Memorandum For The Heads Of Executive Departments And Agencies: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Page 4, Section 3.a; Page 5, Section 3.b.