



Officers
Chair
Judith L. Lichtman
National Partnership for
Women & Families
Vice Chairs
Margaret Huang
Southern Poverty Law Center
Derrick Johnson
NAACP
Thomas A. Saenz
Mexican American Legal
Defense and Educational Fund
Secretary
Fatima Goss Graves
National Women's Law Center
Treasurer
Lee A. Saunders
American Federation of State,
County and Municipal Employees

Board of Directors
Abed Ayoub
American-Arab
Anti-Discrimination Committee
Gloria L. Blackwell
AAUW
Ray Curry
International Union, UAW
Jocelyn Frye
National Partnership for
Women & Families
Jonathan Greenblatt
Anti-Defamation League
Mary Kay Henry
Service Employees International Union
Damon Hewitt
Lawyers' Committee for
Civil Rights Under Law
David H. Inoue
Japanese American Citizens League
Virginia Kase Solomon
League of Women Voters of the
United States
Marc Morial
National Urban League
Janet Murguía
UnidosUS
Svante Myrick
People For the American Way
Janai Nelson
NAACP Legal Defense and
Educational Fund, Inc.
Christian F. Nunes
National Organization for Women
Rabbi Jonah Pesner
Religious Action Center
of Reform Judaism
Rebecca Pringle
National Education Association
Lisa Rice
National Fair Housing Alliance
Kelley Robinson
Human Rights Campaign
Anthony Romero
American Civil Liberties Union
Liz Shuler
AFL-CIO
Fawn Sharp
National Congress of American Indians
Maria Town
American Association of
People with Disabilities
Randi Weingarten
American Federation of Teachers
John C. Yang
Asian Americans Advancing Justice |
AAJC

April 29, 2024

David A. Myklegard
Deputy Federal Chief Information Officer
Office of Management and Budget
725 17th St., N.W.
Washington, D.C. 20503

Christine J. Harada
Senior Advisor
Office of Federal Procurement Policy
Office of Management and Budget
725 17th St., N.W.
Washington, D.C. 20503

Submitted electronically via www.regulations.gov

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

Dear Mr. Myklegard and Ms. Harada,

On behalf of The Leadership Conference on Civil and Human Rights (The Leadership Conference)¹, its Center for Civil Rights and Technology, and the undersigned organizations, we write in response to the Office of Management and Budget's (OMB) *Request for Information: Responsible Procurement of Artificial Intelligence in Government, FR Doc. 2024-06547, filed: 3/28/2024*. The Request for Information (RFI) seeks comments to help inform OMB's development of "an initial means to ensure the responsible procurement of AI by Federal Agencies." The RFI poses questions for response, including on the protection of privacy, civil rights, and civil liberties.

People who face discrimination based on race, ethnicity, national origin, religion, sex, sexual orientation, gender identity, income, immigration status, or disability are more likely to be harmed by automated systems and often lack the resources to respond to harms when they

President and CEO
Maya Wiley

¹ The Leadership Conference is a coalition charged by its diverse membership of more than 250 national organizations to promote and protect the civil and human rights of all persons in the United States. Through its membership and its Media/Telecommunications Task Force, The Leadership Conference works to ensure that civil and human rights, equal opportunity, and democratic participation are at the center of communication, public education, and technology policy debates.

occur. These harms are well documented and span numerous sectors, including housing,² employment,³ financial services and credit,⁴ insurance,⁵ public health and health care,⁶ education,⁷ public accommodations,⁸ government benefits and services,⁹ and policing.¹⁰

The Leadership Conference, as a coalition and through the Center on Civil Rights and Technology, is actively engaged in policy development to ensure civil rights remains at the center of the development and use of new technologies, especially where those technologies are rights and safety impacting. The work of our federal agencies affects the lives of people every day. Ensuring that technologies procured by

² Valerie Schneider, *Locked Out by Big Data: How Big Data, Algorithms and Machine Learning May Undermine Housing Justice*, 52.1 Colum. Hum. Rts. L. Rev. 251, 254 (2020), https://blogs.law.columbia.edu/hrlr/files/2020/11/251_Schneider.pdf; Lauren Kirchner & Matthew Goldstein, *Access Denied: Faulty Automated Background Checks Freeze Out Renters*, The Markup & N.Y. Times (May 28, 2020), <https://themarkup.org/locked-out/2020/05/28/access-denied-faulty-automated-background-checks-freeze-out-renters>; Brief of the American Civil Liberties Union Foundation, The Lawyers' Committee for Civil Rights Under Law, The National Fair Housing Alliance, and The Washington Lawyers' Committee for Civil Rights and Urban Affairs, as Amici Curiae Supporting Appellant and Reversal, *Opiotennione v. Bozzuto Mgmt. Co.*, No. 21-1919, (4th Cir. 2021), ECF No. 49-2, <https://www.lawyerscommittee.org/wp-content/uploads/2022/08/3.-Opiotennione-v.-Bozzuto-Mgmt-Corp-amicus-brief.pdf>.

³ Miranda Bogen & Aaron Rieke, *Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias*, Upturn, (Dec. 2018), <https://www.upturn.org/static/reports/2018/hiring-algorithms/files/Upturn%20-%20Help%20Wanted%20-%20An%20Exploration%20of%20Hiring%20Algorithms,%20Equity%20and%20Bias.pdf>; Alina Köchling & Marius Claus Wehner, *Discriminated by an algorithm: a systematic review of discrimination and fairness by algorithmic decisionmaking in the context of HR recruitment and HR development*, 13 Bus. Res. 795 (2020), <https://doi.org/10.1007/s40685-020-00134-w>.

⁴ Julia Angwin et al., *Minority Neighborhoods Pay Higher Car Insurance Premiums Than White Areas With the Same Risk*, ProPublica (Apr. 5, 2017), <https://www.propublica.org/article/minority-neighborhoods-higher-car-insurance-premiums-white-areas-same-risk>; Maddy Varner & Aaron Sankin, *Suckers List: How Allstate's Secret Auto Insurance Algorithm Squeezes Big Spenders*, The Markup (Feb. 25, 2020), <https://themarkup.org/allstates-algorithm/2020/02/25/car-insurance-suckers-list>.

⁵ Robert Bartlett et al., *Consumer-Lending Discrimination in the FinTech Era* (Nat'l Bureau Econ. Res. Working Paper No. 25943, 2019), <https://www.nber.org/papers/w25943>; Bertrand K. Hassani, *Societal Bias reinforcement through machine learning: a credit scoring perspective*, 1 AI & Ethics 239 (2020), <https://link.springer.com/article/10.1007/s43681-020-00026-z>.

⁶ Ziad Obermeyer et al., *Dissecting racial bias in an algorithm used to manage the health of populations*, 366 Sci. 447 (2019), <https://www.science.org/doi/10.1126/science.aax2342>; Trishan Panch et al., *Artificial intelligence and algorithmic bias: implications for health systems*, 9 J. Glob. Health (2019) (offering definitions of algorithmic bias in health systems), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6875681/>; Natalia Norori et al., *Addressing bias in big data and AI for health care: A call for open science*, 2 Patterns (2021), <https://doi.org/10.1016/j.patter.2021.100347>.

⁷ Todd Feathers, *Major Universities Are Using Race as a "High Impact Predictor" of Student Success*, The Markup (Mar. 2, 2021), <https://themarkup.org/machine-learning/2021/03/02/major-universities-are-using-race-as-a-high-impact-predictor-of-student-success>; Maureen Guarcello et al., *Discrimination in a Sea of Data: Exploring the Ethical Implications of Student Success Analytics*, Educause Rev. (Aug. 24, 2021), <https://er.educause.edu/articles/2021/8/discrimination-in-a-sea-of-data-exploring-the-ethical-implications-of-student-success-analytics>.

⁸ Alex P. Miller & Kartik Hosanagar, *How Targeted Ads and Dynamic Pricing Can Perpetuate Bias*, Harv. Bus. Review (Nov. 8, 2019), <https://hbr.org/2019/11/how-targeted-ads-and-dynamic-pricing-can-perpetuate-bias>; Jennifer Valentino-DeVries et al., *Websites Vary Prices, Deals Based on Users' Information*, Wall St. J. (Dec. 24, 2012), <https://www.wsj.com/articles/SB1000142412788732377204578189391813881534>.

⁹ Rashida Richardson et al., *Litigating Algorithms 2019 Report: New Challenges to Government Use of Algorithmic Decision Systems*, AI Now Inst. (Sept. 2019), <https://ainowinstitute.org/litigatingalgorithms-2019-us.pdf>; Hadi Elzayn et al., *Measuring and Mitigating Racial Disparities in Tax Audits 3-4* (Stan. Inst. for Econ. Pol'y Rsch., Working Paper, Jan. 30, 2023), https://dho.stanford.edu/wp-content/uploads/IRS_Disparities.pdf.

¹⁰ Aaron Sankin et al., *Crime Prediction Software Promised to Be Free of Biases. New Data Shows It Perpetuates Them*, The Markup (Dec. 2, 2021), <https://themarkup.org/prediction-bias/2021/12/02/crime-prediction-software-promised-to-be-free-of-biases-new-data-shows-it-perpetuates-them>; Todd Feathers, *Gunshot-Detecting Tech is Summoning Armed Police to Black Neighborhoods*, Vice: Motherboard (July 19, 2021), <https://www.vice.com/en/article/88nd3z/gunshot-detecting-tech-is-summoning-armed-police-to-black-neighborhoods>.

agencies, including AI systems, procured by agencies work, is critical. Together with its AI Memo, OMB's rules for responsible procurement of AI will be a significant step in ensuring AI is effectively governed across the federal government. Responsible AI procurement practices, including practices that focus on civil rights, are essential for maintaining public trust and ensuring efficient use of taxpayer funds. Accordingly, we agree with OMB's focus on ensuring technology tools, including artificial intelligence (AI) systems, are consistent with our democratic values centered on civil rights and liberties and the protection of privacy. AI procured by federal agencies should behave safely and equitably. AI systems should not be biased or discriminate against individuals or traditionally marginalized groups and communities. AI systems should benefit, not harm, the public. We support OMB placing requirements to ensure that these core AI principles, as well as those embodied in the recent Memorandum, *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence* (AI M-Memo) and in Executive Order 14110, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence* (AI Executive Order), are followed.

It is critical that agencies and vendors understand the risks associated with the use of AI tools before the technology is acquired and put into use. To that end, OMB's policy for the responsible procurement of AI in government must require that any AI system procured by a federal agency adhere to the requirements of OMB's M-Memo and AI EO. In addition, the policy should cover the following areas so that the public is assured an AI system is trustworthy:

- Procurement assessment for federal contracting must have clear requirements of transparency about testing and performance with equity metrics to ensure safety.
- Companies providing AI systems to the federal government should be obligated to ensure clear accountability on the safe and equitable performance of the technologies being procured. They must be held accountable for how their systems function, including ensuring that their systems are not biased and are fit for their intended purpose.
- Vendors should be required to provide information to the procuring agency on the system's capabilities and limitations, including information on the appropriate, as well as inappropriate, uses of the system.
- OMB should provide draft contract language, including data privacy protections, to federal agencies to use in procurement with suppliers of AI.
- OMB's draft contract language should require vendors to provide appropriate access to data, models, and other relevant information to enable testing and evaluation of those systems.
- OMB's procurement guidance should ensure that federal agencies procure AI systems that incorporate relevant standards, like the NIST Risk Management Framework, and other standards that NIST or others may develop to assess and evaluate AI for issues like bias.

Procurement Guidance Needs to Mitigate Harms

As we stated in our coalition’s comments on OMB’s M-Memo, agencies using an AI system must show those systems are safe, trustworthy, and enable rights-protecting outcomes before it is put into use. Consequently, those obligations should also apply to the vendors providing AI to an agency. The procurement guidance that OMB puts in place must build on the M-Memo and continue to center AI policy on equity and rights. As we noted:

“With the Memo, OMB is taking a significant step in ensuring AI is effectively managed throughout the Federal Government. New agency requirements and guidance for AI governance, innovation, and risk management, including through specific risk management practices for uses of AI that impact the rights and safety of the public are very necessary. While AI has the potential to improve operations and efficiency across the Federal Government, those outcomes will only be achieved if people impacted by those systems trust the decisions being made. The Memo provides actionable guidance to agencies setting out how to ensure the use of AI upholds our democratic values and earn that trust.”¹¹

OMB’s procurement guidance should protect the American public against existing and potential harms from AI, including threats to people’s rights, opportunities, jobs, economic well-being, and access to critical resources and services. Marginalized communities served by the different agencies across the federal government are those that bear the most risk from the use of untested and unsafe AI systems. People expect that risks associated with other regulated products will be identified, mitigated, and made known. Likewise, we expect that AI technology is safe and effective—i.e. that it works for everyone.

Building on Prior Policies

OMB’s procurement guidance must build on the measures this administration has already taken to begin the process of mitigating the risks of AI, including Executive Order 14091 (“Further Advancing Racial Equity and Support for Underserved Communities Through the Federal Government”), the AI Bill of Rights, NIST’s AI Risk Management Framework, the 2023 National AI Research and Development Strategic Plan, and the plan for the National AI Research Resource. Executive Order 14091, for example, instructs federal agencies to use all their available authorities to combat algorithmic discrimination. The National AI R&D Strategic Plan calls for investments in technical research to develop frameworks for accountability, fairness, privacy, and bias, as well as research to understand and mitigate the social and ethical risks. These efforts provide a strong basis for a national AI strategy that is centered on equity and civil rights.

In an August 4, 2023 letter to the White House, leading civil rights and civil society organizations said, “(F)ederal agencies funding, acquiring, or using an AI system have a responsibility to ensure that the system works and is fit for purpose.” The groups further stated that the federal government should not use AI systems unless they are shown to be effective and safe. No definition of safe and effective is meaningful unless it is explicit and clear that it includes being non-discriminatory and non-violative of

¹¹ Leadership Conference OMB AI Guidance Memo Comments (Dec. 5, 2023), <https://civilrights.org/wp-content/uploads/2023/12/Leadership-Conference-OMB-AI-Comments.pdf>.

civil and human rights. Simply put, AI should work, and work for everyone.¹² Therefore, protecting civil liberties and advancing equity should be at the center of OMB’s procurement policy in this area. The public should be protected against existing and potential harms from AI—including threats to people’s rights, opportunities, jobs, economic well-being, and access to critical resources and services.

General Considerations

In addressing the “lifecycle” of an AI system through its design, development, and deployment, there are key aspects to consider that can contribute to ensuring those systems are fair and equitable:

Anti-bias expertise: The vendor team should have the right skill sets, such as civil rights expertise, to mitigate bias in the AI system. This includes having a diverse workforce to develop and deploy the technology.

Data governance and assessments: Availability of relevant data is necessary for any AI system. Agencies should ensure that data governance mechanisms are in place at the start of the procurement process. Flaws and potential bias must be addressed before deploying the system in real life uses. Agencies should also determine whether and how data are shared with vendors..

AI impact assessments to determine public benefit: Agencies should define the public benefit goal use that an AI system is intended to achieve. That public benefit should be included in procurement documentation. The public benefit should be compliant with the principles of non-discrimination and equal treatment. Alternative solutions should also be considered during the process of evaluating the AI system. An initial impact assessment addressing the following areas should also be conducted:

- How the AI system meets needs and benefits the public.
- Human and socio-economic impacts.
- Data quality and the potential for inaccuracy or bias.
- Potential unintended consequences.

Community engagement: Early engagement with potentially impacted communities can help inform the process. Impacted communities include impacted work groups and their representative unions. Community engagement should be ongoing and iterative throughout the lifecycle of an AI system. Procurement guidance should also include the requirement to consider the languages of potentially impacted communities.

Avoiding “black box: AI: Explainability and interpretability of AI is important so agencies can best understand the system, identify risks, and determine whether the outcomes comply with civil rights.

Technical and ethical limitations of AI deployment: During the evaluations, vendors should identify areas of potential bias within their training data and measures taken to address that issue.

¹² Letter from the Center for American Progress, The Leadership Conference on Civil and Human Rights, and The Center for Democracy and Technology to the White House (Aug. 4, 2023).

AI lifecycle management: Implementation plans, sustainable and ongoing evaluation methods, and feedback mechanisms are critical to ensure ethical use. Real-world deployment may reveal issues not apparent in the procurement process. Agencies should establish appropriate oversight mechanisms to allow scrutiny of AI systems throughout their lifecycle.

Liability and risk: Risks should be allocated to the parties best able to manage them. That may mean liability for certain areas rests with the agency, particularly around the use and application of the AI solution and in relation to data access. Liability can also rest with a vendor, including areas focused on technical aspects, security, and quality assurance.

In addition, agencies should require vendors to demonstrate the following key practices when providing AI systems to evaluate proposals:

- An internal AI ethics approach, with examples of how it has been applied to design, develop, and deploy AI solutions.
- Processes to ensure accountability of outcome of algorithms.
- Processes that will avoid unfair discriminatory outcomes and/or disparities for historically marginalized communities and individuals.
- Model testing under a range of conditions.
- Definitions of acceptable model performance.

Agencies must conduct due diligence to ensure that vendors have the appropriate capacity to provide information about the AI they are seeking to provide. Due diligence should help establish the accuracy, veracity, and credibility of a vendor's proposal and will be useful in determining what other supporting documentation is needed. Agencies need to ensure that risks are identified, and that the vendor takes steps to mitigate those risks.

We appreciate that procurement policy needs to address which requirements should be placed on vendors and which requirements should be the responsibility of agencies. For example, it may be more suitable to require a vendor to conduct detailed testing and assessments of an AI system, rather than the agency that is seeking to procure the system. Regardless of who is responsible for specific requirements, when it comes to the use of AI systems, no system that is put in use should be biased or discriminatory. In addition, there should be incentives for the development of systems that are equitable or developed with a "benefits-driven" framework.

Auditability: Enable auditability by implementing process logs to gather data across modeling, training, testing, verifying, and implementing the system across the lifecycle.

Model testing: Testing a model on an ongoing basis is needed to show continued accuracy and avoid erroneous decisions that could negatively impact individuals. Contracts should establish how a model will be monitored once deployed.

Consider alternatives: Initial impact assessments, including potential risks, can help determine how to proceed. It is important for an agency to acquire AI systems that are narrowly tailored to the purpose that the agency seeks to address or to a specific articulable challenge an agency forecasts.

End-of-Life: Agencies should consider the end-of-life processes for an AI system.

Contractual Obligations

Agencies should incorporate clear and specific contractual obligations to mitigate the risks related to the deployment and use of AI. Agencies should:

- Define both the public benefit sought, as well as outline potential harm to be prevented and mitigated if they occur. Those harms should include consideration of bias, discrimination, privacy violations, safety concerns, and other adverse impacts.
- Define expected testing requirements for AI systems. Agencies should be involved in testing to ensure usability, fairness, and safety outcomes are consistent with the desired outcomes and the administration's AI policies.
- Use third-party audits to verify vendor compliance.
- Require vendors to promptly report any identified AI harms related to the system provided.
- Establish procedures for reporting and escalating AI harm incidents.
- Ensure that government personnel receive adequate training on using AI systems so they can understand the safety aspects and potential risks.
- Set performance benchmarks related to AI harm mitigation, including penalties, termination of the contract, or legal action.

To minimize risks, vendors should be required to:

- Conduct a risk assessment specific to the AI being provided. Based on that assessment, the vendor should develop a risk mitigation plan to address identified risks, including bias, privacy, and unintended consequences. Risk assessments should be conducted prior to deploying AI systems.
- Ensure that the AI system's decision-making process is transparent and explainable. "Black box" algorithms, including proprietary models, should be accompanied by clear documentation that explains how their AI models work. That documentation should include a description of the system's capabilities and limitations.
- Securely protect data, including storage, transmission, and access.
- Address bias in the AI, including conducting regular audits and fairness assessments to identify and rectify bias.
- Validate testing models, including against diverse datasets.
- Be accountable for the performance and outcomes of the AI systems. If harm is found, the vendor should assist the agency in investigating and taking corrective action.
- Establish monitoring mechanisms to detect issues, including unintended consequences. Vendors should also establish a process to provide regular reports on system performance, including any incidents.

- Provide user manuals that cover AI system operation, risks, and mitigation strategies.

Conclusion

We appreciate this opportunity to comment on procurement guidance, which, building on OMB’s M-Memo, should provide a roadmap for a practical and actionable approach to identify, measure, and mitigate harms before AI is put into use, as well as to evaluate existing systems. As we have stated before:

“For the use of AI to be successful, agencies must ensure that the benefits and risks of AI are considered early on and throughout the AI lifecycle, through design, development, and deployment. Before procuring or using an AI system, an agency should understand its limitations, recognize its intended uses as well as potential misuses, consider how to ensure the AI works for the people, and prevent harm.”

Without clear guidance on how to ensure accountability, transparency, and explainability, agencies will fail to ensure oversight of algorithmic decision-making and facilitate risk of harm. Moreover, given that federal government procurement rules and purchasing practices often have a strong influence on markets, well-crafted procurement guidance can help set a baseline for responsible AI use. OMB guidance will set expectations for government use of AI and is necessary to align use of AI systems with the stated policy goals of fairness and equity.

Thank you for your consideration of our concerns and views. Please direct any questions about these comments to Koustubh “K.J.” Bagchi, vice president of the Center for Civil Rights and Technology at The Leadership Conference, at bagchi@civilrights.org or Frank Torres, privacy and AI fellow at The Leadership Conference, at torres@civilrights.org.

Sincerely,

The Leadership Conference on Civil and Human Rights
Access Now
Asian Americans Advancing Justice – AAJC
Communications Workers of America
National Consumer Law Center (on behalf of its low-income clients)
NETWORK Lobby for Catholic Social Justice
The Trevor Project
United Church of Christ Media Justice Ministry