

Comments Regarding the Office of Management and Budget's
Request for Information on Responsible Procurement of Artificial Intelligence in Government
by the

Electronic Privacy Information Center
Brennan Center for Justice at New York University School of Law
Electronic Frontier Foundation
Fight for the Future
TechTonic Justice
UnidosUS

April 29, 2024

I. Introduction

The Electronic Privacy Information Center (EPIC) and the Brennan Center for Justice at New York University School of Law (Brennan Center for Justice), alongside the undersigned civil society organizations, submit these comments in response to the Office of Management and Budget (OMB)'s Request for Information on Responsible Procurement of Artificial Intelligence in Government, published March 29, 2024.¹

EPIC is a public interest research center in Washington, D.C., established in 1994 to secure the fundamental right to privacy in the digital age for all people through advocacy, research, and litigation.² EPIC has a long history of promoting a human-rights based approach to AI policy³ and has developed specific expertise over AI procurement, authoring a comprehensive state AI procurement report⁴ and co-leading development of the forthcoming IEEE AI procurement standard.⁵ The Brennan Center for Justice is a nonpartisan law and policy institute that works to reform, revitalize, and defend our country's systems of democracy and justice.

This comment highlights two key gaps in current federal procurement processes for AI—**lack of vendor transparency** and **procedures for adapting to new AI risks**—and suggests additional, AI-specific procurement practices to facilitate compliance with OMB Memo M-24-10, *Advancing*

¹ 89 Fed. Reg. 22196 (Mar. 29, 2024), <https://www.federalregister.gov/documents/2024/03/29/2024-06547/request-for-information-responsible-procurement-of-artificial-intelligence-in-government>.

² EPIC, "About Us," accessed April 24, 2024, <https://epic.org/about/>.

³ See, e.g., EPIC, "AI and Human Rights," accessed April 24, 2024, <https://epic.org/issues/ai/>; EPIC, "Comments on the OMB's Draft Memorandum on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," December 5, 2023, <https://epic.org/wp-content/uploads/2023/12/EPIC-OMB-AI-Guidance-Comments-120523-L.pdf> [hereinafter "EPIC OMB AI Memo Comment"].

⁴ Grant Fergusson, *Outsourced and Automated: How AI Companies Have Taken Over Government Decision-Making*, EPIC, September 2023, 5–25, <https://epic.org/wp-content/uploads/2023/09/FINAL-EPIC-Outsourced-Automated-Report-w-Appendix-Updated-9.26.23.pdf> [hereinafter "Outsourced & Automated Report"].

⁵ IEEE SA, "Standard for the Procurement of Artificial Intelligence and Automated Decision Systems," accessed April 24, 2024, <https://standards.ieee.org/ieee/3119/10729/>.

*Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence.*⁶ This comment uses government acquisition of social media monitoring tools—an AI use case that OMB Memo M-24-10 presumes to be rights-impacting—as a case study of how these recommendations should apply.

II. Current Federal Procurement Processes Lack the Transparency and Adaptability Required for Agencies to Comply with OMB AI Guidance

Responsive to Questions 1–3.

AI systems are not comparable to other procured products and services or even non-AI software. AI and automated decision-making systems are *opaque* and *unexplainable* processes, characteristics that displace agency oversight, understanding, and control of critical agency processes in ways that other procured products and services do not.⁷ As Deirdre Mulligan and Kenneth Bamberger put it, “government agencies purchasing and using these [automated] systems most often have no input into—or even knowledge about—their design or how well that design aligns with public goals and values.”⁸ The traditional values that procurement promotes—price, competition, and innovation—do not adequately address the transparency, expertise, and oversight needed to manage AI risk, hindering agencies’ ability to ensure that procured systems protect rights and safety.⁹

To effectively manage AI vendors and their systems—and comply with federal AI guidelines—agencies need (1) meaningful transparency from vendors about their processes and AI functionality and (2) mechanisms to adjust procurement processes and contracts in response to new and changing AI risks. Currently, however, few vendors provide the information agencies need to conduct meaningful AI testing, audits, or impact assessments.¹⁰ Nor do most agencies retain sufficient records about vendor AI systems to adequately monitor system performance.¹¹ When high-risk AI incidents do occur, agencies must frequently rely on vendor-initiated reporting and system updates to identify and remedy harms.¹² Contractual limitations on the government’s rights to disclose information and data pertaining to procured

⁶ Shalanda D. Young (director, Office of management and Budget (OMB)), M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence, OMB, <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf> [hereinafter “OMB AI Memo”].

⁷ Deirdre K. Mulligan and Kenneth A. Bamberger, “Procurement as Policy: Administrative Process for Machine Learning,” *Berkeley Technology Law Journal*, vol. 34 (October 4, 2019): 781, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3464203; see also Outsourced & Automated Report, 19–25.

⁸ Mulligan and Bamberger, 786.

⁹ Mulligan and Bamberger, 787; see also Outsourced & Automated Report at 5–25; see generally EPIC OMB AI Memo Comment.

¹⁰ See Danielle Keats Citron, “Open Code Governance,” *University of Chicago Legal Forum* vol. 2008, no. 1 (2008): 357, <https://chicagounbound.uchicago.edu/cgi/viewcontent.cgi?article=1430&context=ucf>; Frank Pasquale, “Restoring Transparency to Automated Authority,” *Journal on Telecommunications and High Technology Law* 9, no. 235 (February 17, 2011): 235–36, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1762766 (exploring solutions to trade secrecy); Katherine Fink, “Opening the Government’s Black Boxes: Freedom of Information and Algorithmic Accountability,” *Information, Communication and Society* 21, no. 10 (May 30, 2017): 1–19, <https://www.tandfonline.com/doi/full/10.1080/1369118X.2017.1330418?scroll=top&needAccess=true>.

¹¹ Outsourced & Automated Report at 19–20; Mulligan and Bamberger, 786; Robert Brauneis and Ellen P. Goodman, “Algorithmic Transparency for the Smart City,” *Yale Journal of Law and Technology* 20, no. 103 (2018): 152, https://yjolt.org/sites/default/files/20_yale_j_l_tech_103.pdf.

¹² Outsourced & Automated Report at 8–9; Mulligan and Bamberger, 786.

systems, as well as limitations on how and when agencies can directly address errors, bugs, or other issues in vendor AI systems, also hinder agency oversight of these systems.

III. Additional AI Procurement Practices Are Necessary Across the Procurement Lifecycle

Procurement should not be a vehicle for skirting agencies' AI-related obligations, nor should it be a vehicle for undue vendor influence over government decisions. To ensure responsible and rights-respecting AI procurement, OMB should require federal agencies to supplement Federal Acquisition Regulation (FAR) requirements with additional risk identification and management practices before, during, and after procurement. The recommendations listed in this comment align with the Blueprint for an AI Bill of Rights,¹³ the National Institute of Standards and Technology (NIST)'s AI Risk Management Framework (AI RMF 1.0),¹⁴ the OMB's recent AI guidance (OMB Memo M-24-10),¹⁵ and the Risk Management Framework for the Procurement of Artificial Intelligence (RMF PAIS 1.0) authored by the AI Procurement Lab and the Center for Inclusive Change.¹⁶

OMB should also reinforce that its procurement guidance applies to all AI systems acquired for law enforcement purposes. EO 14110 requires a separate National Security Memorandum only for "AI used as a component of a national security system or for military and intelligence purpose."¹⁷ All other systems, including law enforcement systems, are within the scope of OMB's guidance. For dual use systems, such as systems procured to collect domestic intelligence and initiate criminal investigations, OMB should clarify how its guidance would apply. Both OMB and agencies that rely on dual use systems should make every effort to ensure that procurement of these systems aligns with the procurement standards OMB establishes to protect privacy, civil rights, and civil liberties.

A. Pre-Solicitation

Responsive to Questions 4, 9, and 10.

Procured AI systems bring with them not only the risks inherent to AI, but also risks arising from vendor relationships: informational disadvantages and accountability hurdles. An agency may be in a worse position to further its mission if it procures an inappropriately high-risk, harmful, or unexplainable AI system. Therefore, agencies should conduct a *needs assessment* at the pre-solicitation stage encompassing the following:

¹³ Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People*, White House, October 2022, <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

¹⁴ National Institute of Standards and Technology (hereinafter "NIST"), *Artificial Intelligence Risk Management Framework*, Department of Commerce, January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

¹⁵ OMB AI Memo.

¹⁶ Cari L. Miller and Gisele Waters, Risk Management Framework for the Procurement of AI Systems (RMF PAIS 1.0), AI Procurement Lab and Center for Inclusive Change, 2024, https://inclusivechange.hubspotpagebuilder.com/rmf_pais_1_download?submissionGuid=56365dd4-4f68-4f04-abf8-c675ab357b1e.

¹⁷ Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, Exec. Order No. 14110, 88 Fed. Reg. 75191 (2023), § 4.8, <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>.

1. Conducting a preliminary **assessment of potential rights or safety impacts** of available AI tools as part of FAR-required market research,¹⁸ which should involve **external stakeholder engagement** with potentially impacted communities and civil society.
2. Evaluating the risks and benefits of procuring AI for a specified use case *compared to* (a) **agency-led AI development** and (b) **non-AI alternatives**.

Too often, agencies are marketed shiny AI products and services without knowing their accompanying risks, impacts, and limitations. A pre-solicitation needs assessment serves as necessary friction to ensure agencies prioritize AI procurement that aligns with federal AI risk management standards—and centers the rights and safety impacts of procured AI *before* expending significant resources. Agencies should publicly disclose both pre-solicitation needs assessments and any impact assessments conducted further along in the procurement process.

Social Media Monitoring: Agencies seeking to acquire AI-based social media monitoring tools for broad-scale detection of security threats and risks should take stock of existing social media surveillance programs: have they yielded useful and reliable information, and at what cost to civil liberties, civil rights, or human safety?¹⁹ Not only is evidence of their utility lacking, they have also led to wrongful and discriminatory arrests, interrogations and investigations, and chilling effects on online expression.²⁰ This analysis alone should lead agencies to halt procurement.

A risk-benefit analysis of introducing AI in this context should examine whether the technology will lead to more reliable and less biased inferences of threats and risks—or just inject more complexity, errors, and biases into law enforcement operations. For example, agencies should assess how racial and ethnic biases embedded in natural language processing, the machine-learning process underlying many social media monitoring tools, would amplify harms to marginalized populations that are already disproportionately scrutinized, such as Black, Latino and Muslim communities.²¹

EXAMPLE

B. Solicitation and Evaluation

Responsive to Questions 5, 6, 9, and 10.

Agencies must ensure that AI vendors can meet the requirements of OMB Memo M-24-10 and other federal AI guidelines *before* awarding them the contract. The solicitation process is typically the first—and in the case of sealed bidding, potentially the only—chance for agencies to evaluate whether vendor practices and systems conform to these guidelines.²²

¹⁸ 48 CFR § 10.

¹⁹ Rachel Levinson-Waldman, Harsha Panduranga, and Faiza Patel, “Social Media Surveillance by the U.S. Government,” Brennan Center for Justice, January 7, 2022, <https://www.brennancenter.org/our-work/research-reports/social-media-surveillance-us-government>; EPIC, *Social Media Monitoring: Government Surveillance of Public Space*, June 2018, <https://archive.epic.org/privacy/surveillance/spotlight/0518/Social-Media-Monitoring.pdf>.

²⁰ *Ibid.*, Complaint, Doc Society v. Blinken, No. 19CV03632 (D.D.C. December 5, 2019) (case originally filed as Doc Society v. Pompeo), <https://knightcolumbia.org/documents/399e32ad77>.

²¹ Anjalie Field et al., “A Survey of Race, Racism, and Anti-Racism in NLP,” *Association for Computational Linguistics* vol. 1 (2021): 1905–1925, <https://aclanthology.org/2021.acl-long.149.pdf>; Abubakar Abid, Maheen Farooqi, and James Zou, “Persistent Anti-Muslim Bias in Large Language Models,” *AAAI/ACM Conference on AI, Ethics, and Society* (2021): 298–306, <https://dl.acm.org/doi/10.1145/3461702.3462624>.

²² Sealed bids are evaluated without discussions. See, 48 CFR § 14.101.

Ideally, FAR should be amended to incorporate the OMB guidelines.²³ In the meantime, agency Contracting Officers (COs) should designate these guidelines as a special responsibility standard under the Regulation. Section 9.104-2 authorizes COs to develop “special standards of responsibility” with the “assistance of appropriate specialists” when these are necessary for a “particular acquisition or class of acquisitions.”²⁴ Vendors must show that they can fulfill these standards before a contract is awarded.²⁵ For example, the Federal Transit Administration imposes a special quality assurance requirement for prospective bus manufacturers imposing special testing and manufacturing controls on vendors.²⁶

AI systems are a “class of acquisitions” that should require special reporting and documentation standards. These standards should enable agencies to lead testing, evaluation, and assessment that is: (1) socio-technical in nature, and (2) tailored to how agency officials might use, abuse, or be hindered by these systems, and how their operation would affect people and communities.²⁷ Information provided by the vendor (e.g., results of the vendor’s own error and bias testing of their system) should be assessed alongside other sources of data and analysis, such as relevant academic literature and consultations with affected individuals and communities. These standards should also give agencies the option of commissioning independent testing and audits of vendor-provided claims and information (e.g., testing vulnerabilities in the vendor’s systems that requires technical expertise beyond the agency’s purview).

At a minimum, these standards should require vendors to provide the following information:

1. Proof of vendors’ **AI development, testing, and governance policies**, including incident response policies.
2. All information necessary for the procuring agency to conduct an **AI impact assessment**, including a list of an AI system’s intended uses and limitations; the source, quality, appropriateness and limitations of data used in the AI’s design, development, training, testing, or operation; and any documented risks and risk mitigation protocols.
3. All information necessary to conduct an **independent AI audit**, including the results of any vendor AI testing, evaluation, validation, or red-teaming exercises; information regarding AI model architecture, explainability, cybersecurity, and human oversight; and any relevant AI system accessibility measures to manage language barriers and disabilities. Vendors that reach the contract negotiation stage should expect to provide more comprehensive technical access to their systems for independent testing (e.g., by depositing the source code, model weights and training parameters in escrow).²⁸

²³ Hannah Quay-de la Vallee, Ridhi Shetty, and Elizabeth Laird, Federal Government's Power of the Purse: Enacting Procurement Policies and Practices to Support Responsible AI Use, Center for Democracy & Technology, April 2024, 35, <https://cdt.org/wp-content/uploads/2024/04/2024-04-27-AI-Federal-Procurement-report-1.pdf>.

²⁴ 48 CFR § 9.104-2(a).

²⁵ 48 CFR § 9.104-2(b).

²⁶ Federal Transit Administration, “Determining Contractor Responsibility,” Department of Transportation, updated December 28, 2017, <https://www.transit.dot.gov/funding/procurement/third-party-procurement/determining-contractor-responsibility>.

²⁷ Reva Schwartz et al., *NIST Special Publication 1270: Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*, NIST, March 2022, 10, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>.

²⁸ Outsourced & Automated Report, 48.

Vendors should also be required to disclose this information to the public, or versions thereof that omit only those details that are strictly necessary to protect legitimate trade secrets.

If the vendor is providing a variety of AI tools within a single system (such as sentiment analysis and predictive analytics), the risks, biases and limitations of each tool and relevant mitigations should be separately documented and discussed in their proposal. Creating standard forms and practices for vendors to explain AI testing, evaluation, validation, and red-teaming results both during solicitations and throughout the life cycle of a contract will enable agencies to more effectively review vendor practices even if they lack comprehensive AI expertise.

Reporting Templates and Best Practices: OMB should evaluate and recommend templates for providing information on AI systems and vendor testing. Some promising emerging templates include model cards (which explain how the AI system works),²⁹ dataset cards (which explain how data is collected, generated, or used by the system),³⁰ and transparency reports (which detail a company’s policies and processes).³¹ In the context of social media monitoring, minimum disclosures include: *error rates and thresholds*, such as the rate at which sentiment analysis and predictive functions inaccurately attribute negative attitudes to social media posts;³² *model limitations*, such as the inability to detect sarcasm, satire, or irony;³³ and *data sources and safeguards*, including how publicly available social media data is scraped and updated, safeguards to prevent the collection of private data, and the types of data used to train the system to analyze sentiment and make threat predictions.

EXAMPLE

Risk management should not just be a compliance measure. AI solicitations should incentivize competition among vendors to provide the strongest possible protections for civil liberties and civil rights – not just faster, bigger, or more sophisticated AI systems. When evaluating proposals, agencies can prioritize non-price factors such as past performance and how well they can satisfy technical requirements.³⁴ Solicitations and related documents should specify that the vendor’s past performance addressing rights and safety-related risks, and their ability to provide model features that mitigate these risks, will be evaluated. Vendor-provided risk mitigation features should supplement – but not displace – agencies’ assessments of the risks and benefits of procuring a particular AI system.

Incentivizing Risk Management: The use of sentiment analysis tools on social media is unproven and risky. Sentiment analysis tools that are programmed to always associate neutral sentiment with a specified list of identity terms representing protected characteristics (e.g., “gay,” “Muslim,” and “feminist”) may lower the risk of discrimination compared to tools that lack such a feature.³⁵ Ultimately, however, vendors’

EXAMPLE

²⁹ Margaret Mitchell et al., “Model Cards for Model Reporting,” *Proceedings of the Conference on Fairness, accountability, and Transparency* (January 29, 2019): 220–229, <https://dl.acm.org/doi/10.1145/3287560.3287596>; Hugging Face, “Model Cards,” accessed April 24, 2024, <https://huggingface.co/docs/hub/en/model-cards>.

³⁰ Hugging Face, “Dataset Cards,” accessed April 24, 2024, <https://huggingface.co/docs/hub/en/datasets-cards>.

³¹ See, e.g., Digital Services Act, “Article 13 – Transparency reporting obligations for providers of intermediary services,” December 15, 2020, <https://digitalservicesact.cc/dsa/art13.html> (mandating standardized reporting on content moderation activities by internet platforms).

³² See, e.g., Salesforce, “Detect Sentiment Model Card,” accessed April 24, 2024, https://help.salesforce.com/s/articleView?id=sf.bi_integrate_transformation_detectSentimentModelCard.htm&type=5.

³³ Mayur Wankhade, Annavarapu Chandra Sekhara Rao, and Chaitanya Kulkarni, “A survey on sentiment analysis methods, applications, and challenges,” *Artificial Intelligence Review* vol. 55 (2022): 5731–5780, <https://link.springer.com/article/10.1007/s10462-022-10144-1>.

³⁴ 48 CFR § 15.304.

³⁵ See, e.g., Salesforce, “Detect Sentiment Model Card.”

provision of technical safeguards should not displace agency judgment that this technology is far too inaccurate for deployment in law enforcement and other high-risk contexts. In particular, sentiment analysis conducted on non-Latin derived languages, such as Arabic and Urdu, is likely unreliable and error prone given the lack of high-quality web data on such languages (the most common source of training data for large language models).³⁶

The solicitation and proposal evaluation process may also surface risks to civil liberties and rights that agencies did not foresee in previous needs or impact assessments. OMB should designate this as a trigger point for agencies to reevaluate the risks of procuring AI, and whether pursuing non-AI strategies would lower the risks. AI solicitations should be withdrawn if the additional risks cannot be meaningfully prevented or mitigated.

Reevaluating Risks: During the Federal Bureau of Investigation’s (FBI) solicitation of proposals for a “Social Media Tactical Exploitation Tool,” a Q&A with prospective vendors surfaced concern from vendors that the Bureau may require them to run searches on First Amendment activity, such as keyword searches covering “protected groups like activist organizations.” Vendors also raised concern that other technical requirements, such as creating generic social media accounts for surveillance purposes, and location-based analysis of social media posts, would violate platforms’ Terms of Service (TOS) designed to protect user privacy.³⁷

The possibility of a constitutional violation is a core rights impact that must trigger additional agency risk management processes. A needs assessment would have almost certainly flagged the privacy and constitutional impacts of social media monitoring, requiring the Bureau to engage with impacted communities and examine non-AI alternatives. If these risks did not emerge before solicitations, however, vendor concerns should have triggered: (1) additional review by the agency’s Chief Artificial Intelligence Officer (CAIO) and privacy, civil liberties and civil rights offices or staff, (2) reconsideration of the necessity of procuring AI-based monitoring capabilities, and (3) amendments to contract solicitations and related documents to address any identified risks or impacts should the procurement move forward. For example, an agency could amend its Statement of Work to only allow techniques that do not violate constitutional rights or platforms’ TOS and require additional independent audits and the implementation of anonymous whistleblower processes. If an agency is unable to develop sufficient risk mitigation (e.g., the capabilities it desires cannot meaningfully abide by constitutional requirements), it should not procure these capabilities or move forward with procurement at all.

C. Negotiation and Contract Award

Responsive to Questions 7–10.

Needs and impact assessments conducted pre-solicitation or during solicitation should be updated to capture the risks specific to the vendor that was ultimately awarded the contract. Competing AI systems

³⁶ Gabriel Nicholas and Aliya Bhatia, *Lost in Translation: Large Language Models in Non-English Content Analysis*, Center for Democracy & Technology, May 2023, 17, <https://cdt.org/wp-content/uploads/2023/05/non-en-content-analysis-primer-051223-1203.pdf>.

³⁷ Federal Bureau of Investigation (FBI), “Request for Proposals – Social Media Exploitation – Amendment 2,” Department of Justice, updated January 21, 2022, <https://sam.gov/opp/63867a4d178d4717ac246d61c955cc05/view>; FBI, “QA_RFP 15F06722R0000005,” accessed April 24, 2024, row 53, https://sam.gov/api/prod/opps/v3/opportunities/resources/files/174b38e572f64a40be24ff6317bb8d80/download?&status=arc_hived&token= (Excel spreadsheet compiling FBI answers to industry questions, see row 53).

within the same use context can incorporate vastly different data, training processes, model architectures, and testing protocols, leading to variations in risk profiles.

Risks identified throughout the procurement process and associated risk mitigation protocols should be consolidated into a single repository that is accessible to the agency and vendor, such as an **AI risk register** developed and maintained by the agency. This register should, at a minimum, identify:

1. All risks associated with the general use case, the type of AI system acquired, and the awardee’s AI governance policies, including risks that may not be explicitly covered by OMB Memo M-24-10;
2. The party responsible for mitigating these risks (with a strong presumption that agencies will conduct all socio-technical risk mitigation measures while overseeing and auditing any vendor testing protocols);
3. The information required to monitor and mitigate these risks; and
4. Any other AI risk management practices required to manage these risks.

The development of an AI risk register before contract award not only facilitates more effective risk monitoring at the start of a contract period, but also enables agency COs to raise new concerns and add new, vendor- and system-specific contract provisions during AI contract negotiations. The register should be made publicly accessible to the fullest extent possible.

Social Media Monitoring Risk Register: The risks associated with social media monitoring tools may vary depending on the types of languages analyzed, the sources of training data, or the platforms covered. For example, sentiment analysis tools may be too error-prone for law enforcement use. A risk register should capture these risks, corresponding mitigations, as well as agency and vendor responsibilities in conducting and overseeing these mitigations (e.g., agency consults with independent linguistics and AI experts to assess reliability and usefulness of sentiment analysis; vendor to restrict sentiment analysis function to non-law enforcement uses).

EXAMPLE

D. Contract Monitoring

Responsive to Questions 9–10.

Under both new and existing AI contracts, contract monitoring will serve as the primary method for ensuring an agency’s procured AI complies with federal AI use requirements. For future acquisitions, COs should negotiate for contractual clauses that enable agencies to continually test, monitor, and mitigate risks consistently with OMB Memo M-24-10 and other applicable guidelines. In particular, COs should negotiate for data rights that enable meaningful access to the system’s inner workings for auditing and other risk monitoring purposes (and strictly for those purposes only). This data may include descriptions of training datasets, model weights, and the results of vendor testing and evaluations.³⁸ To ensure compliance with use case inventory and other transparency requirements, COs should negotiate for contract clauses that enable agencies to publicly disclose information that is critical to understanding the

³⁸ 48 CFR § 52.227–14; Stephen Casper et al., “Black-Box Access is Insufficient for Rigorous AI Audits,” *ArXiv*: 2401.14446 (2024), <https://arxiv.org/pdf/2401.14446.pdf>.

system’s impacts on rights and safety, such as accessible visualizations or summaries of model behavior, as well as records of related incidents and incident responses.³⁹

For AI systems acquired before OMB Memo M-24-10, agencies should exercise their data rights under the contract to seek as much of the aforementioned information as possible. If reporting and disclosure gaps remain, COs should, in coordination with agency Chief AI Officers (CAIOs), negotiate modifications to the contract supported by appropriate financial consideration.⁴⁰ Agencies should refrain from exercising their option to renew contracts if they are unable to determine that the AI systems at issue comply with OMB Memo M-24-10 and other federal AI use requirements.

Social Media Monitoring Contract Review: CAIOs should work with COs to review existing agency contracts for social media monitoring systems to ensure that they comply with OMB Memo M-24-10. This review should cover both social media monitoring systems as well as components of agency systems that incorporate such monitoring (e.g., predictive policing systems that rely on social media analysis). CAIOs should also commission testing for error, discrimination, and bias, and assess how system outputs (e.g., sentiment scores or threat predictions) facilitate decisions or actions that impact rights (e.g., investigative activities) and develop metrics for measuring the reliability of these outputs.

EXAMPLE

IV. Conclusion

The undersigned organizations welcome OMB’s interest in advancing responsible AI procurement throughout the federal government. Many rights and safety impacts of federal AI come from procured AI systems, and OMB guidelines on AI procurement procedures will help agencies align their procurement practices with OMB Memo M-24-10, Executive Order 14110, and other federal AI guidance. For any questions, please contact Grant Fergusson (EPIC) at fergusson@epic.org and Amos Toh (Brennan Center for Justice) at toha@brennan.law.nyu.edu.

Respectfully submitted,

Electronic Privacy Information Center
Brennan Center for Justice
Electronic Frontier Foundation
Fight for the Future
TechTonic Justice
UnidosUS

³⁹ For an analysis of some existing AI contract clauses, see Outsourced & Automated Report at 41–48.

⁴⁰ Christopher G. Griesedieck Jr. and Michael T. Francel, “New Rules, Pre-Existing Contract: A FAR-Reaching Issue?,” American Bar Association, March 4, 2024, https://www.americanbar.org/groups/public_contract_law/publications/procurement-lawyer/2024/winter/new-rules-pre-existing-contract-far-reaching-issue/.