BRENNAN CENTER FOR JUSTICE

December 5, 2023

Office of Management and Budget
Submitted Via Federal eRulemaking Portal at www.regulations.gov

RE: Brennan Center for Justice Response to "Request for Comments on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Memorandum," OMB-2023-0020

The Brennan Center for Justice at NYU School of Law is a nonpartisan law and policy institute that works to reform, revitalize, and defend our country's systems of democracy and justice. We submit the following comments in response to the Office of Management and Budget's (OMB) Request for Comments on its "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence" draft memorandum (OMB Memorandum).

The Brennan Center commends OMB for its efforts to develop transparency measures and risk management practices for safety- and rights-impacting use of artificial intelligence (AI). We submit these comments and recommendations to further strengthen these critical safeguards.

1. **The IC and national security systems should not be excluded wholesale from the minimum practices outlined in the OMB Memorandum.**

The OMB Memorandum contains broad carveouts for the Intelligence Community (IC) and national security systems. The Advancing American AI Act of 2022 exempts the IC and Department of Defense (DoD) from the requirement to publish AI use case inventories.[1] The OMB Memorandum, however, goes further by entirely exempting the IC from implementing minimum risk management practices.[2] Similarly, "national security systems" are not mentioned anywhere in the Advancing American AI Act, but are also exempted from the ambit of the OMB Memorandum.[3] We recognize that Executive Order 14110 provides that the IC, DoD, and national security systems will be covered by a separate national security memorandum,[4] but are of the view that consigning rights-impacting AI use by significant parts of the federal government to an opaque national security track seriously jeopardizes the privacy, rights, and liberties of millions of people in the United States. As has become clear in the context of policing and criminal justice, too often the brunt of algorithmic discrimination is borne by communities of color, heightening the risks of allowing a large segment of AI to remain outside the framework that is meant to ensure that these systems are effective, accurate, fair, and rights-respecting.

[1] Advancing American AI Act, Pub. L. 117-263, title LXXII, subtitle B, §§ 7225(d), 7228 (2022) (exempting the Department of Defense and the intelligence community), https://www.congress.gov/117/plaws/publ263/PLAW-117publ263.pdf.
[2] Shalanda D. Young (director, Office of Management and Budget (OMB)), memorandum, Re: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence, 88 Fed. Reg. 75625 (November 3, 2023), sec. 5 ("Managing Risks From the Use of Artificial Intelligence") (hereinafter "OMB Memorandum").
[3] OMB Memorandum, sec. 2(a).
[4] 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.

National security systems are defined by statute to include information systems whose deployment "involves intelligence activities."[5] Under the provisions of Executive Order 12333, this means all IC systems used for data collection, retention, analysis, and dissemination.[6] Since the IC includes federal agencies that are primarily domestically focused, such as the Federal Bureau of Investigation (FBI)[7] and certain parts of the Department of Homeland Security (DHS), a broad swath of systems that directly affect persons living in the U.S. are excluded from basic testing and risk management practices. This is particularly concerning because the FBI and DHS have deployed AI systems that exhibit the types of problems and risks that the OMB Memorandum is meant to mitigate.

The FBI's massive facial recognition program through which it has access to billions of face prints is used for both law enforcement and intelligence purposes.[8] It has been repeatedly criticized by the Government Accountability Office for: inadequate accuracy testing, including a failure to test performance in common use scenarios and to test for false positives;[9] lack of training, including no training requirements for agents using external facial recognition technology like Clearview AI, leaving 95 percent of agents not trained;[10] lack of systems to track agents' use of non-FBI facial recognition tools;[11] and lack of specific policies to mitigate risks to civil rights and liberties from this technology.[12] The Bureau's lax attitude toward the reliability and fairness of its systems is particularly alarming given the existence of several research studies documenting facial recognition's inaccuracy in identifying non-white faces,[13] and multiple reports of erroneous matches resulting in wrongful arrests and incarceration of Black Americans.[14] Moreover, the FBI has also developed, in partnership with DoD, technology to identify individuals from video footage of public spaces, which could pose serious risks to First Amendment protected activities if it is used to, for example, identify protestors for surveillance and other measures.[15]

[5] 44 USC § 3552(b)(6)(A)(i)(I).
[6] United States Intelligence Activities, Exec. Order No. 12333, 46 Fed. Reg. 59941 (1981), part 2.3, https://www.archives.gov/federal-register/codification/executive-order/12333.html.
[7] 50 U.S.C. § 3003(4)(H). While the statute includes only the "intelligence elements" of the Federal Bureau of Investigation (FBI), in practice, the FBI and the Director of National Intelligence consider the entire Bureau to fall under the intelligence community (IC) definition. FBI, "About: Leadership & Structure," last accessed December 5, 2023, https://www.fbi.gov/about/leadership-and-structure/intelligence, and Office of the Director of National Intelligence, "Members of the IC," last accessed December 5, 2023, https://www.dni.gov/index.php/what-we-do/members-of-the-ic.
[8] See Ernest J. Babcock, "Privacy Impact Assessment for the Facial Analysis, Comparison, and Evaluation (FACE) Services Unit," FBI, May 1, 2015, sec. 3, https://www.fbi.gov/how-we-can-help-you/more-fbi-services-and-information/freedom-of-information-privacy-act/department-of-justice-fbi-privacy-impact-assessments/facial-analysis-comparison-and-evaluation-face-services-unit, and Erin M. Prest, *Privacy Impact Assessment for the Facial Analysis, Comparison, and Evaluation (FACE) Phase II System,* Department of Justice, July 9, 2018, sec. 3, https://www.fbi.gov/file-repository/pia-face-phase-2-system.pdf/view.
[9] See Government Accountability Office (GAO), *Face Recognition Technology: FBI Should Better Ensure Privacy and Accuracy*, May 16, 2016, 25–32, https://www.gao.gov/assets/gao-16-267.pdf.
[10] See GAO, *Facial Recognition Services: Federal Law Enforcement Agencies Should Take Actions to Implement Training, and Policies for Civil Liberties*, September 5, 2023, 26–28, https://www.gao.gov/assets/gao-23-105607.pdf.
[11] See GAO, *Facial Recognition Technology: Federal Law Enforcement Agencies Should Better Assess Privacy and Other Risks*, June 29, 2021, 20–26, https://www.gao.gov/assets/gao-21-518.pdf.
[12] See generally GAO, *Facial Recognition Services,* https://www.gao.gov/assets/gao-23-105607.pdf.
[13] Patrick Grother, Mei Ngan, and Kayee Hanaoka, *Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects,* National Institute of Standards and Technology (NIST), December 19, 2019, https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf, and Steve Lohr, "Facial Recognition Is Accurate, if You're a White Guy," *New York Times*, February 9, 2018, https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html.
[14] See Khari Johnson, "Face Recognition Software Led to His Arrest. It was Dead Wrong," *Wired*, February 28, 2023, https://www.wired.com/story/face-recognition-software-led-to-his-arrest-it-was-dead-wrong/.
[15] Drew Harwell, "FBI, Pentagon helped research facial recognition for street cameras, drones," *Washington Post*, March 7, 2023, https://www.washingtonpost.com/technology/2023/03/07/facial-recognition-fbi-dod-research-aclu/.

Other FBI systems are also deployed for both domestic intelligence and criminal investigation purposes. For example, the Bureau's various social media monitoring programs, for which it has contracted with companies such as Dataminr, ZeroFox, and Babel Street, use some version of AI to trawl the web.[16]

DHS's Office of Intelligence and Analysis (I&A), which is also part of the IC, undertakes broad social media monitoring for domestic intelligence purposes. It uses an AI tool called Tangles that constantly monitors and scrapes the web to build detailed profiles of individuals.[17] During the 2020 racial justice protests, I&A built dossiers on protestors using social media and other sources, generating a chorus of criticism from Congress and the department's oversight offices.[18] The DHS Office of the General Counsel found that I&A employees did not have an adequate understanding of what constitutes protected speech,[19] while its inspector general found that I&A "does not have comprehensive policies and procedures to ensure [that] its employees effectively collect [open source intelligence] and adhere to privacy protections."[20] More recently, I&A has surveilled Americans discussing abortion online,[21] and DHS has launched a program to trawl the internet for dangerous "narratives and grievances," raising First Amendment concerns.[22] I&A is now developing AI data analysis tools for use by the law enforcement parts of the department.[23]

Given the impact of programs like these on people in the U.S., their wholesale exemption leaves in place risky AI systems without safeguards and should be replaced by a system that institutes protections, testing, and oversight comparable to those for agencies and systems covered under the OMB Memorandum.

2. OMB should clarify provisions exempting agencies from the obligation to produce information about AI use cases in public AI use case inventories and exercise oversight over such exemptions.

According to Section 3(a)(iv) footnote 8 of the OMB Memorandum, agencies must publicly report use cases "to the extent practicable and consistent with applicable law and governmentwide guidance, including those concerning the protection of privacy and of sensitive law enforcement, national security, and other protected information."[24] We appreciate that the OMB Memorandum references only governmentwide—rather than agency-specific—policies in this provision. At the same time, we believe that the standards articulated in footnote 8 are too broad and vague. We recommend that they be clarified as set out below.

---

[16] See Ken Klippenstein, "FBI Hired Social Media Surveillance Firm that Labeled Black Lives Matter Organizers 'Threat Actors,'" *Intercept*, July 6, 2023, https://theintercept.com/2023/07/06/fbi-social-media-surveillance-zerofox/, and *Select Comm. To Investigate the January 6th Attack on the U.S. Capitol,* 117th Congress (July 26, 2022) (interview transcript of Jennifer Moore, executive assistant director, FBI), 10–11, 98, https://www.govinfo.gov/content/pkg/GPO-J6-TRANSCRIPT-CTRL0000916069/pdf/GPO-J6-TRANSCRIPT-CTRL0000916069.pdf (explaining use of social media monitoring for predicated investigations and intelligence gathering).

[17] Office of Intelligence and Analysis (I&A), *Office of Intelligence and Analysis Operations in Portland*, DHS, April 20, 2021, 55, https://www.wyden.senate.gov/imo/media/doc/I&A%20and%20OGC%20Portland%20Reports.pdf; see also Joey Scott, "LAPD Is Using Israeli Surveillance Software That Can Track Your Phone and Social Media," *Knock LA*, November 27, 2023, https://knock-la.com/lapd-is-using-israeli-surveillance-software-that-can-track-your-phone-and-social-media/.

[18] I&A, *Office of Intelligence and Analysis Operations in Portland,* Department of Homeland Security (DHS), https://www.wyden.senate.gov/imo/media/doc/I&A and OGC Portland Reports.pdf.

[19] I&A, *Office of Intelligence and Analysis Operations*, 20.

[20] Office of the Inspector General, *The Office of Intelligence and Analysis Needs to Improve Its Open Source Intelligence Reporting,* DHS, July 6, 2022, 11, https://www.oig.dhs.gov/sites/default/files/assets/2022-07/OIG-22-50-July22.pdf.

[21] Jana Winter, "DHS monitored 'social media reactions' to Roe, collected legally protected speech, bulletin shows," Yahoo! News, updated November 16, 2022, https://news.yahoo.com/dhs-monitored-social-media-reactions-to-roe-collected-legally-protected-speech-bulletin-shows-001254616.html.

[22] Ken Dilanian, "DHS launches warning system to find domestic terrorism threats on public social media," NBC News, May 10, 2021, https://www.nbcnews.com/politics/national-security/dhs-launches-warning-system-find-domestic-terrorism-threats-public-social-n1266707.

[23] Privacy Office, *Privacy Impact Assessment for the DHS Data Analysis Tools,* DHS/ALL/PIA-055(a), DHS, June, 13, 2023, 19–20, https://www.dhs.gov/sites/default/files/2023-05/privacy-pia-cbp068-cbpmobileapplication-may2023.pdf.

[24] OMB Memorandum, 4 n8. Section 5(c)(iv)(H) of the OMB Memorandum includes a similar exemption.

First, OMB should make clear that agencies cannot use this provision to avoid providing information regarding systems about which information has already been publicly released—for example through press reports, FOIA documents, Congressional testimony, privacy impact assessments, systems of records notices, and other sources. Public release undermines any argument for wholesale withholding information about a system based on practicability or protection of information.

Second, the decision to exempt use cases on these grounds should not be left solely to the agency. For each invocation of an exemption per footnote 8, OMB should require agencies to submit a detailed rationale and relevant documentation, as well as to specify the system-specific policies they have in place to mitigate any risks to privacy, civil rights, and civil liberties. To increase transparency, OMB should annually publish information on the number of exemptions by agency.

Third, OMB should issue public guidelines making clear that practicability would justify an exemption from transparency only in the most exceptional circumstances and specify factors that are relevant to such a determination.

Fourth, OMB should similarly issue public guidelines on the situations in which laws and guidance concerning the protection of privacy, sensitive law enforcement, national security, and other protected information could justify not publishing information about AI use cases in a use case inventory, with transparency as the default principle.

In particular, OMB should clarify that—absent a specific provision or provisions in a law or governmentwide guidance that clearly require that no information about a system be disclosed—footnote 8 does not allow agencies to avoid disclosing the very existence of an AI system, a description of the capabilities of that AI system, the purpose of the agencies' use of that AI system, and the rules and restrictions that govern the use of the system. Although OMB has yet to publish guidance on what information will be included in use case inventories, it seems highly unlikely that such disclosures would reveal information about individuals, and the fact that a system collects or uses information about individuals does not by itself trigger privacy concerns about disclosing *system*-level information.[25] Nor would such disclosures reveal information about particular investigations that could be regarded as sensitive law enforcement information. Local laws already require police departments to disclose information about their surveillance technologies, providing evidence that system level transparency is achievable.[26] Indeed, public reporting of AI systems may help improve such systems, allowing for input and advice from outside experts and encouraging agencies to deliberate carefully about their use of AI.

Lastly, allowing agencies to avoid publicly reporting use cases on the grounds of "sensitive national security information" appears redundant because of the broad carveouts for national security systems, the IC, and DoD discussed above. We recommend that OMB strike this phrase or, at minimum, clarify for the public what types of systems would be covered in this category. We also request that OMB specify what type of information is covered by the phrase "other protected information," noting that classified information would be covered by other laws and regulations and the carveouts for national security systems, the IC, and DoD. Any use of these exceptions should, in any event, be narrowly construed and subject to OMB oversight.

[25] In the unlikely circumstance that use case inventories would include information about individuals, agencies should be required to report that information in a privacy-preserving manner—rather than choose not to report that information at all.
[26] See Stevie Degroff and Albert Fox Cahn, *New CCOPS On the Beat: An Early Assessment of Community Control of Police Surveillance Laws*, Surveillance Technology Oversight Project, February 10, 2021, 6, https://static1.squarespace.com/static/5c1bfc7eee175995a4ceb638/t/602430a5ef89df2ce6894ce1/1612984485653/New+CCOPS+On+The+Beat.pdf (showing that, as of December 2020, fourteen local governments required annual reports of law enforcement's surveillance technology acquisition and use).

3. OMB should require agencies' evaluation of data to be based on prevailing academic norms rather than historical data.

We appreciate that OMB has highlighted the importance of ensuring the "quality and appropriateness" of training data when developing safety- and rights-impacting AI,[27] and note that in some cases this may require investing in better data collection, consistent with privacy concerns, and abandoning AI use cases until more suitable data exists.[28] Data on the criminal justice system especially suffers from quality issues that would undermine the validity and accuracy of any AI trained on it. To take one example, the First Step Act of 2018 requires the Bureau of Prisons to develop an algorithm to predict the "recidivism risk" of people in federal custody. As currently deployed, that algorithm, "PATTERN," focuses on calculating a narrow type of recidivism: the risk that someone will be re-arrested (or, much more rarely, returned to federal prison). But arrest and criminality are not the same thing, and conflating the two has led PATTERN's predictions to suffer from unwarranted racial disparities.[29] A better and more accurate tool would abandon the focus on re-arrest entirely and seek to predict solely the risk of re-conviction or re-imprisonment.[30] Remarkably, though, the data to train such an algorithm does not exist—at least not in one place. Instead, it is scattered across databases maintained by state courts and law enforcement agencies.[31] These problems—atomized and incomplete data that speak to only one narrow aspect of a social problem—are unfortunately common in the criminal justice context.[32]

Agencies must, as the OMB Memorandum advises in Section 5(c), exercise special care when designing systems operating in this space. OMB should further clarify that agencies' evaluation of the "quality and appropriateness" of data to be used in safety- and rights-impacting AI should be based on prevailing academic norms rather than historical use cases. It is not enough, for example, that federal agencies have historically measured recidivism based on the risk of re-arrest. Agencies should determine tabula rasa whether a dataset measures what it purports to and, if not, whether that renders it inappropriate for use in machine learning. And OMB should encourage agencies to consult with a range of civil society voices in making these determinations.

---

[27] OMB Memorandum, sec. 5(c)(iv)(3).

[28] See Reva Schwartz et al., "Towards a Standard for Identifying and Managing Bias in Artificial Intelligence," NIST, March 2022, 15–19, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf (warning against a bias toward focusing "on which datasets are available or accessible, rather than what dataset might be most suitable").

[29] Ames Grawert, "Brennan Center's Public Comment on the First Step Act's Risk and Needs Assessment Tool," Brennan Center for Justice, September 4, 2019, https://www.brennancenter.org/our-work/research-reports/brennan-centers-public-comment-first-step-acts-risk-and-needs-assessment, and Ames Grawert and Patricia Richman, The First Step Act's Prison Reforms: Uneven Implementation and the Path Forward, September 23, 2022, 3–4, https://www.brennancenter.org/our-work/research-reports/first-step-acts-prison-reforms.

[30] See Jessica M. Eaglin, "Constructing Recidivism Risk," *Articles By Maurer Faculty* (2017), 94, https://www.repository.law.indiana.edu/cgi/viewcontent.cgi?article=3647&context=facpub (noting that arrest is one of "the least procedurally protected instances of contact with the criminal justice system").

[31] To its credit, the Department of Justice (DOJ) continues to investigate ways to improve PATTERN, including potentially refocusing the tool to predict re-conviction risk. See Office of the Attorney General, *First Step Act Annual Report*, DOJ, April 2023, 10–11, https://www.ojp.gov/first-step-act-annual-report-april-2023.

[32] Until very recently, government crime data largely focused on a handful of offenses—hardly a complete picture of crime in the United States. See generally Ames Grawert and Noah Kim, "Understanding the FBI's 2021 Crime Data," Brennan Center for Justice, October 4, 2022, https://www.brennancenter.org/our-work/research-reports/understanding-fbis-2021-crime-data, and "U.S. Crime Rates and Trends — Analysis of FBI Crime Statistics," Brennan Center for Justice, October 16, 2022, https://www.brennancenter.org/our-work/research-reports/us-crime-rates-and-trends-analysis-fbi-crime-statistics.

4. OMB should tighten standards for waivers and exercise oversight over them.

Section 5(c)(iii) of the OMB Memorandum allows an agency's Chief AI Officer (CAIO) to grant a waiver from risk management practices for safety- and rights-impacting AI uses based on an assessment that fulfilling the requirements would either "increase risks to safety or rights overall" or constitute an "unacceptable impediment to critical agency operations."[33] This standard is overbroad and vague and leaves too much discretion in the hands of CAIOs. We recommend the following:

First, OMB should only allow waivers for a set period and subject to annual review. The minimum standards are meant to ensure that AI works (through testing), that it doesn't harm people (through safeguards against bias and civil liberties' protections), and that its use is understood by those who are subject to it (through transparency). Agencies' use of faulty or biased technology should not be allowed to continue indefinitely through waivers.

For example, DHS's Customs and Border Protection (CBP) uses CBP One, an app that requires asylum seekers at the U.S.-Mexico border to use facial recognition in order to schedule appointments at ports of entry and requires non-citizens to use facial recognition to request travel authorization for parole processes.[34] Not only is there abundant evidence about facial recognition programs' inaccuracy in identifying darker skinned faces, but there are also reports that CBP One in particular "is failing to register many people with darker skin tones"—such as those who have made their way to the southern border from Haiti and African countries—making it near impossible for them to apply for asylum.[35] Managing the flow of asylum seekers at the border is clearly a high priority for CBP, but it should not get an open-ended pass to continue using technology that may be biased rather than take steps to improve its systems.

Second, it is unclear what would constitute an "unacceptable impediment to critical operations" and an increased "risk to safety or rights overall." OMB should clarify what is meant by these terms; clarify how adherence to these standards would be assessed; provide examples of scenarios in which a waiver is or is not appropriate; and give detailed instructions or a template for what information must be shared with OMB.

Finally, the OMB Memorandum leaves too much discretion in the hands of the CAIO, inviting operational pressures to prevail over the need for testing and safeguards. We have seen time and again that internal offices charged with protecting civil rights and civil liberties have given way in the face of operational imperatives.[36] There is no reason to believe that CAIOs would behave differently, particularly considering the vague standard put forward in the OMB Memorandum. Any waivers should be subject to review by OMB, and OMB should make publicly available information on systems for which waivers are granted, including: the scope of the waiver (e.g., which risk management requirements the agency will not need to implement with respect to that use case); the period for which the waiver applies; and the date on which OMB will review the waiver.

5. OMB should maximize publicly available information about AI uses.

In addition to public reporting, the OMB Memorandum in several places requires agencies to report, solely to OMB, information such as agency lists of safety- and rights-impacting AI[37] and how agencies are

[33] OMB Memorandum, sec. 5(c)(iii).
[34] Privacy Office, *Privacy Impact Assessment for CBP One*, DHS/CBP/PIA-068, DHS, May 12, 2023, 18, 24, https://www.dhs.gov/sites/default/files/2023-05/privacy-pia-cbp068-cbpmobileapplication-may2023.pdf.
[35] Melissa del Bosque, "Facial Recognition Bias Frustrates Black Asylum Applicants to US, Advocates Say," *Guardian*, February 8, 2023, https://www.theguardian.com/us-news/2023/feb/08/us-immigration-cbp-one-app-facial-recognition-bias.
[36] Spencer Reynolds and Alia Shahzad, *Holding Homeland Security Accountable: How to Strengthen Departmental and Congressional Oversight,* Brennan Center for Justice, October 26, 2023, 2–8, https://www.brennancenter.org/media/11860/download.
[37] OMB Memorandum, sec. 5(b).

implementing risk management practices.[38] In the interest of transparency, OMB should make this information public as part of use case inventories.

We note that while DoD and the IC are both exempt from maintaining AI use case inventories, the OMB Memorandum provides a mechanism for DoD to report to OMB AI use cases that are not used as part of a national security system.[39] The IC should have this same requirement.

6. Internal agency mechanisms are not sufficient to address AI risks.

We appreciate that OMB will be exercising some degree of oversight over agencies' interpretations of the requirements it has articulated. However, as set out in this submission, the OMB Memorandum does not apply to national security systems and IC uses of rights-impacting AI, leaving these critical systems without risk management practices, antidiscrimination testing, and impact assessments. We have already encouraged OMB to expand its oversight and review responsibilities to the greatest extent possible. Additional independent oversight would contribute significantly to building trust in the use of AI, and we therefore recommend that an appropriately resourced outside body modeled on the Privacy and Civil Liberties Oversight Board be constituted to serve as an additional check on agency use of AI.

***

We appreciate OMB's commitment to ensuring that people's rights and safety are protected from potentially harmful AI systems and urge you to take the steps outlined above to better accomplish this goal. We also appreciate OMB's engagement with civil society and encourage you to continue to seek input from stakeholders. For any questions, please contact Faiza Patel at patelf@brennan.law.nyu.edu.

Sincerely,

Faiza Patel
Senior Director
Brennan Center for Justice
Liberty and National Security Program

[38] OMB Memorandum, sec. 5(c).
[39] OMB Memorandum, sec. 3(a)(v).