ACLU

December 5, 2023

Via regulations.gov.
Office of Management and Budget
Docket No. OMB-2023-0020

RE: Request for Comments: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum

The American Civil Liberties Union writes in response to the Request for Comment published by the Office of Management and Budget (OMB) on its draft memorandum on "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence"[1] (Draft Memorandum). These comments highlight the many provisions of the Draft Memorandum that would meaningfully preserve and advance civil rights and civil liberties and presents recommendations on where the Draft Memorandum may be strengthened. These comments urge OMB to:

- Clarify and broaden the scope of the Draft Memorandum by reemphasizing its application to law enforcement and immigration agencies, reconsidering the exclusion of national security systems and intelligence agencies, clarifying its application to less-advanced systems, and extending its applicability to federally funded programs;
- Maintain or expand the scope of AI uses cases that are presumed to be rights-impacting;
- Maintain the strength of key minimum practices while bolstering others, including by cabining agency discretion to waive or otherwise avoid the Draft Memorandum's requirements, bolstering notice provisions, and ensuring that community feedback mechanisms are equitable and approachable.

I.   OMB Should Clarify and Broaden the Scope of the Draft Memorandum to Encompass Critical AI Use Cases

A.   OMB Should Adopt the Proposed Definition of "Covered Agencies" and Clarify that It Includes Law Enforcement and Immigration

OMB should retain the Draft Memorandum's broad definition of "covered agencies." The Draft Memorandum proposes to apply to any "agency" as defined in 44 U.S.C. § 3502(1). That definition is appropriately broad. The AI revolution is already affecting all aspects of life, including commerce, finance, education, employment, governmental benefits, healthcare, and more. The definition of "agency" in § 3502 is commensurate to the broad sweep of AI—and the challenges it brings, encompassing "any executive department, military department, Government

---

corporation, Government controlled corporation, or other establishment in the executive branch of the Government."[2]

In the final draft of the memorandum, OMB should emphasize that this definition encompasses law enforcement, immigration, and similar agencies. Uses of AI in those agencies have a pronounced impact on rights and safety. Law enforcement agencies across the country have deployed algorithmic systems such as facial recognition technology and predictive policing systems, often with harmful results. For example, facial recognition algorithms have been shown to have higher false match rates when used to try to identify people of color.[3] To date, there are at least six known cases of Black people being wrongfully arrested and jailed based on police reliance on faulty facial recognition "matches."[4] A recent report by the U.S. Government Accountability Office found that of seven federal law enforcement agencies within the U.S. Department of Justice and Department of Homeland Security investigated, all seven initially used facial recognition technology "without requiring staff take facial recognition training," and four lacked "policies or guidance specific to facial recognition technology that address civil rights and civil liberties."[5] Likewise, predictive policing algorithms rely on data—such as arrest rates—that has baked into it the over-surveillance and disparate policing of communities of color.[6]

---

[2] 44 U.S.C. § 3502(1).

[3] National Institute of Standards & Technology, Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects (NISTIR 8280) (2019), https://www.nist.gov/news-events/news/2019/12/nist-study-evaluates-effects-race-age-sex-face-recognition-software.

[4] *See* Eyal Press, *Does A.I. Lead Police to Ignore Contrary Evidence:?*, New Yorker (Nov. 13, 2023), https://www.newyorker.com/magazine/2023/11/20/does-a-i-lead-police-to-ignore-contradictory-evidence; Kashmir Hill, *Eight Months Pregnant and Arrested After False Facial Recognition Match*, N.Y. Times (Aug. 6, 2023), https://www.nytimes.com/2023/08/06/business/facial-recognition-false-arrest.html; Kashmir Hill & Ryan Mac, *'Thousands of Dollars for Something I Didn't Do'*, N.Y. Times (Mar. 31, 2023), https://www.nytimes.com/2023/03/31/technology/facial-recognition-false-arrests.html; Kashmir Hill, *Another Arrest, and Jail Time, Due to a Bad Facial Recognition Match*, N.Y. Times (Dec. 29, 2020), https://www.nytimes.com/2020/12/29/technology/facial-recognition-misidentify-jail.html; Kashmir Hill, *Wrongfully Accused by an Algorithm*, N.Y. Times (Aug. 3, 2020), https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html; Elisha Anderson, *Controversial Detroit Facial Recognition Got Him Arrested for a Crime He Didn't Commit*, Detroit Free Press (July 10, 2020), https://www.freep.com/story/news/local/michigan/detroit/2020/07/10/facial-recognition-detroit-michael-oliver-robert-williams/5392166002.

[5] Government Accountability Office, Facial Recognition Services: Federal Law Enforcement Agencies Should Take Actions to Implement Training, and Policies for Civil Liberties (2023), https://www.gao.gov/products/gao-23-105607.

[6] *See e.g.*, Will Douglas Heaven, *Predictive Policing Algorithms Are Racist. They Need to be Dismantled*, MIT Technology Review (2020), https://www.technologyreview.com/2020/07/17/1005396/predictive-policing-algorithms-racist-dismantled-machine-learning-bias-criminal-justice/; Rashida Richardson et al., *Dirty Data, Bad*

Similarly, AI and algorithmic systems have outsized impacts on individuals' rights and safety in the immigration system. For example, since 2012, Immigration and Customs Enforcement has used an algorithmic "risk classification assessment" to help to determine whether individuals should be held in detention during the pendency of removal proceedings. The algorithmic assessment, however, was tilted in favor of detention, and "over time, the [assessment] became a means of imposing detention in nearly all cases, instead of driving ICE to use alternatives to detention as originally intended."[7] More recently, immigration agencies have expanded their use of AI-powered surveillance, such as automated license plate readers,[8] social media monitoring,[9] facial recognition and other biometric surveillance,[10] automated surveillance towers "capable of identifying and capturing human faces,"[11] and data mining a vast array of public and private databases.[12] Major data brokers provide algorithmically driven data services to

Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice, 94 N.Y.U. L. Rev. 15 (2019), https://www.nyulawreview.org/online-features/dirty-data-bad-predictions-how-civil-rights-violations-impact-police-data-predictive-policing-systems-and-justice.

[7] Robert Koulish and Kate Evan, *Punishing with Impunity: The Legacy of Risk Classification in Immigration Detention*, 36 Geo. Immigr. L.J. 1, 36 (2021), https://www.law.georgetown.edu/immigration-law-journal/in-print/volume-36-number-1-fall-2021/punishing-with-impunity-the-legacy-of-risk-classification-in-immigration-detention-2; *accord* Robert Koulish, *Immigration Detention in the Risk Classification Assessment Era*, 16 Conn. Pub. Int. L.J. 3, 27 (2016), https://cpilj.law.uconn.edu/wp-content/uploads/sites/2515/2018/10/16.1-Immigration-Detention-in-the-Risk-Classification-Assessment-Era-by-Robert-Koulish.pdf ("[W]here ICE had broad discretion to detain or to release subjects, RCA almost always (94.6 percent of the time) considered individuals to be high or medium flight risks and usually (62.1 percent of the time) considered them to be high or medium public safety risks.")' Adi Robertson, *ICE Rigged Its Algorithms to keep Immigrants in Jail, Claims Lawsuit*, The Verge (Mar. 3, 2020), https://www.theverge.com/2020/3/3/21163013/ice-new-york-risk-assessment-algorithm-rigged-lawsuit-nyclu-jose-velesaca.

[8] Benjamin Hayes, *U.S. Immigration and Customs Enforcement Use of Automated License Plate Reader Databases*, 33 Geo. Immigr. L.J. 145 (2019), https://www.law.georgetown.edu/immigration-law-journal/in-print/u-s-immigration-and-customs-enforcement-use-of-automated-license-plate-reader-databases; Vasudha Talla, *Documents Reveal ICE Using Driver Location Data From Local Police for Deportations*, ACLU NorCal (Mar. 13, 2019), https://www.aclunc.org/blog/documents-reveal-ice-using-driver-location-data-local-police-deportations.

[9] Charlie Savage, Visa Applicants' Social Media Data Doesn't Help Screen for Terrorism, Documents Show, N.Y. Times (Oct. 5, 2023), https://www.nytimes.com/2023/10/05/us/social-media-screening-visa-terrorism.html (despite lack of efficacy, Biden administration plans to maintain program of surveilling social media profiles); *Timeline of Social Media Monitoring for Vetting by the Department of Homeland Security and the State Department*, Brennan Center for Justice (May 21, 2021), https://www.brennancenter.org/our-work/research-reports/timeline-social-media-monitoring-vetting-department-homeland-security-and (describing both manual and automatic screening of the social media accounts).

[10] Mijente et al., Who's Behind ICE 58-63 (2023), https://mijente.net/wp-content/uploads/2023/02/Who-is-Behind-ICE-The-Tech-and-Data-Companies-Fueling-Deportations_v4.pdf.

[11] Mijente et al., The Deadly Digital Border Wall 12 (2021), https://notechforice.com/wp-content/uploads/2021/10/Deadly.Digital.Border.Wall_.pdf.

[12] Alvaro Bedoya, *The Cruel New Era of Data-Driven Deportation*, Slate (Sept. 22, 2020), https://slate.com/technology/2020/09/palantir-ice-deportation-immigrant-surveillance-big-data.html; Drew Harwell, *ICE Investigators Used a Private Utility Database Covering Millions to Pursue Immigration Violations*, Washington Post (Feb. 26, 2021), https://www.washingtonpost.com/technology/2021/02/26/ice-private-utility-data/;

immigration agencies such as "social media surveillance, access to jail booking data, face recognition and 'geolocation analysis & geographic mapping' of cellphones," including "large volume online batching" and analysis of the data.[13]

Emphasizing that law enforcement, immigration, and related agencies fall within the Draft Memorandum's ambit is necessary to foreclose abuse of accommodations made for law enforcement in the Draft Memorandum. For example, the Draft Memorandum advises that agencies' inventories of AI use cases may not have to include "sensitive law enforcement"[14] information and that immediate notice may not be practicable for individuals involved in law enforcement proceedings.[15] It is critical the final memorandum underscore its applicability to law enforcement, immigration, and similar agencies to preclude these exceptions from swallowing the rule.

B. OMB Should Reconsider Exceptions for National Security Systems and Intelligence Agencies

OMB should reconsider its wholesale exclusion of national security systems and its broad carve-outs for intelligence and defense agencies in the Draft Memorandum.[16] These AI systems may involve some of the federal government's most consequential decisions, but OMB's approach would allow them to continue delaying the adoption of critical safeguards. These agencies should not be categorically excluded from the basic protections contemplated by OMB. Instead, OMB's memorandum should encompass such agencies wherever possible and permit calibration of AI requirements only for specific use cases where a modification is strictly justified by national security needs and subject to robust internal and external oversight.

The Draft Memorandum largely parallels President Biden's recent Executive Order on "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence"[17] in its treatment of these systems. The Executive Order excludes "national security systems" from its requirements,[18] and subjects those systems to a separate, future process, culminating in submission of a "National Security Memorandum on AI" to the President.[19] Intelligence agencies—including the Office of the Director of National Intelligence, the Central Intelligence

---

Douglas MacMillan and Elizabeth Dwoskin, *The War Inside Palantir: Data-Mining Firm's Ties to ICE Under Attack by Employees*, Washington Post (Aug. 22, 2019), https://www.washingtonpost.com/business/2019/08/22/war-inside-palantir-data-mining-firms-ties-ice-under-attack-by-employees.

[13] Sam Biddle, *LexisNexis Sold Powerful Spy Tools to U.S. Customs and Border Protection*, The Intercept (Nov. 16, 2023), https://theintercept.com/2023/11/16/lexisnexis-cbp-surveillance-border.

[14] Draft Memorandum at 4 n.8.

[15] Draft Memorandum at 20 n.35.

[16] These comments use "intelligence agencies" to refer to the agencies defined in 50 U.S.C. § 3003.

[17] Executive Order 14110 of October 30, 2023, sec. 4.8, 88 Fed. Reg. 75191, 75204 (Nov. 1, 2023).

[18] *E.g.*, Executive Order 14110 of October 30, 2023, secs. 4.1(a)(ii), 10.1(a), 10.1(i) (excluding national security systems from various AI governance protocols).

[19] Executive Order 14110 of October 30, 2023, sec. 4.8.

Agency, the National Security Agency, and the intelligence elements of the Federal Bureau of Investigation, the Department of Homeland Security, and the Drug Enforcement Administration[20]—would also be subject to that future National Security Memorandum.[21] OMB's Draft Memorandum adopts a similar approach[22] by exempting the intelligence agencies from AI use-case inventories and from adhering to the Draft Memorandum's minimum risk management practices.[23]

OMB should change its approach and avoid "national security" exceptionalism wherever possible. U.S. intelligence agencies and the military have raced to integrate AI into some of the government's most profound decisions: who it surveils, who it adds to government watchlists, who it labels a "risk" to national security, and even who it targets using lethal weapons.[24] In many of these areas, the deployment of AI already appears to be well underway. But the public knows almost nothing about the systems that agencies like the FBI, DHS, CIA, and National Security Agency are developing or deploying, and even less about what concrete safeguards exist to ensure fairness, equal treatment, privacy, and due process—if any.[25]

The Draft Memorandum suggests that these agencies' uses of AI are "governed through other policy,"[26] but that is a significant overstatement. The policy sources it identifies are largely general statements of principles without meaningful accountability mechanisms or binding rules. For example, ODNI's *Principles for Artificial Intelligence Ethics for the Intelligence Community* describes six high-level guidelines—including a commitment to be "transparent and accountable," but the public to date has seen little evidence of either.[27] The Defense Department recently released a toolkit "to help DoD personnel design, develop, deploy, and use AI systems responsibly," but using the toolkit is voluntary.[28] Overall, the intelligence and defense agencies lack specific rules and safeguards for their AI systems, as well as clear processes to implement and enforce those rules.

In addition, the elements of the National Security Memorandum mandated in the Executive Order lack the specificity of OMB's Draft Memorandum, and there is no guarantee the

---

[20] 50 U.S.C. § 3003; Executive Order 12333 of December 4, 1981, sec. 3.5(h).
[21] Executive Order 14110 of October 30, 2023, sec. 4.8(a)-(b).
[22] Draft Memorandum at 3 (exempting "national security systems" from the scope of the Draft Memorandum).
[23] Draft Memorandum at 4, 10 (excluding the Intelligence Community from use case inventories and minimum risk management practices).
[24] *See, e.g.*, National Security Commission on Artificial Intelligence, Final Report at 81, 109–10 (2021), https://assets.foleon.com/eu-west-2/uploads-7e3kk3/48187/nscai_full_report_digital.04d6b124173c.pdf.
[25] *See Id.* at 141–54, 395–410.
[26] Draft Memorandum at 3 n.6.
[27] Office of the Director of National Intelligence, Intelligence Community Principles of Artificials (2020), https://www.dni.gov/index.php/newsroom/reports-publications/reports-publications-2020/3634-principles-of-artificial-intelligence-ethics-for-the-intelligence-community-1692377385.
[28] Department of Defense, *CDAO Releases Responsible AI (RAI) Toolkit for Ensuring Alignment With RAI Best Practices* (Nov. 14, 2023), here.

eventual memorandum will include meaningful protections. To fill this critical gap, OMB should—to the extent it is able—extend the Draft Memorandum's provisions to national security systems that affect people in the United States and subject the Intelligence Community to the requirements to provide use-case inventories, subject to a strict declassification review process, and adhere to minimum risk-management practices.

A careful reading of the Executive Order and the Advancing American AI Act reveals that OMB is not precluded, at the very least, from extending the minimum risk-management practices to intelligence and defense agency uses of AI that are not "national security systems." Although the Advancing American AI Act does not mention "national security systems," the Executive Order exempts those systems from the scope of OMB's final memorandum.[29] The Executive Order does not, however, exclude intelligence and defense agencies wholesale.

Moreover, although the Advancing American AI Act does exclude some agencies from its provisions, those exclusions apply only to the Act's use case inventories. The Act excludes intelligence agencies from its provisions generally,[30] but nothing in that Act expressly exempts intelligence agencies from the minimum risk-management practices. In addition, the Act's exclusion of defense agencies is expressly limited to the use-case inventories,[31] and the minimum practices are still applicable, which the Draft memorandum acknowledges.[32] It may also be permissible to extend other provisions in the Draft Memorandum to intelligence and defense agencies.

To the extent that OMB is precluded by either the Executive Order or statutory law from mandating protections for national security systems or intelligence and defense agencies, we urge the President and Congress to revisit those limitations. Regardless, to ensure accountability, OMB should require *all* agencies to identify which AI systems they believe qualify as "national security systems" and the basis for that determination, given the potential misuse of this label to shield AI systems deployed for purposes such as ordinary law enforcement and immigration.

C. OMB Should Clarify that "Covered AI" Includes Less Advanced Algorithmic Systems

Although the definition of "artificial intelligence" adopted by the Draft Memorandum[33] is capacious enough to accommodate many automated, computerized processes already in place,

---

[29] Executive Order 14110 of October 30, 2023, secs. 10.1(b)(iv), 10.1(i).
[30] Pub. L. No. 117-263, div. G, title LXXII, subtitle B, § 7228, https://www.congress.gov/117/plaws/publ263/PLAW-117publ263.pdf.
[31] *Id.* § 7225(d).
[32] Draft Memorandum at 5.
[33] Draft Memorandum at 22 (artificial intelligence means an "artificial system" that meets one of five prongs, including "perform[ing] tasks under varying and unpredictable circumstances without significant human oversight," "solv[ing] tasks requiring human-like perception [or] cognition," "think[ing] or act[ing] like a human," "approximat[ing] a cognitive task," or "act[ing] rationally."

OMB should clarify in the final draft of the memorandum that less sophisticated algorithmic decision-making is within the scope of the guidance. In particular, OMB should clarify that systems that rely on or are derived from statistical modeling are covered by the memorandum.

That clarification is crucial because many critical, rights-impacting governmental functions are already being carried out by algorithmic systems that are derived from statistical modeling and may fall short of popular depictions of "artificial intelligence."[34] For example, a litigator with the ACLU's Idaho affiliate recently testified before a U.S. Senate Committee about his experience challenging black-box algorithms used by Idaho to determine individuals' Medicaid benefits.[35] The algorithmic system reduced or denied benefits—sometimes by more than 30 percent—without explanation.[36] The algorithmic system was implemented without notice, and the State of Idaho and its private vendor attempted to hide its functioning behind trade secrets claims.[37] The ACLU of Idaho eventually prevailed in court and learned that Idaho's system was "a set of formulas in a fairly basic Microsoft Excel spreadsheet," which computed each person's benefits in "hidden cells," leaving state officials unable to explain how or why it reached its benefits determinations.[38] Despite its outsized impact on individuals' rights, Idaho's algorithmic system lacked critical safeguards, based on underlying models that "Department staff had just brainstormed," but "never validated, standardized, or audited the instrument."[39]

Similarly, the PATTERN risk assessment developed by the U.S. Department of Justice is used to inform programming and release decisions for individuals incarcerated in federal facilities. PATTERN scores can be calculated by adding up whole numbers based on roughly a dozen pieces of information about a person, and these scores may be calculated using paper-based forms or processes.[40] While a tool like PATTERN may appear to be simple, the tool was developed using statistical modeling techniques, including "machine learning boosted regression

---

[34] *See* James Vincent, *Top AI Researchers and CEOs Warn Against "Risk of Extinction" in 22-Word Statement*, The Verge (May 30, 2023), https://www.theverge.com/2023/5/30/23742005/ai-risk-warning-22-word-statement-google-deepmind-openai.

[35] Testimony of Ritchie Eppink, Hearing AI in Government Before the S. Comm. On Homeland Security & Government Affairs (May 16, 2023), https://www.hsgac.senate.gov/hearings/artificial-intelligence-in-government.

[36] *Id.* at 2.

[37] *Id.* at 3.

[38] *Id.*

[39] *Id.*

[40] See Federal Bureau of Prisons, PATTERN Risk Assessment, https://www.bop.gov/inmates/fsa/pattern.jsp (last visited November 27, 2023).

procedures,"[41] and it is used in ways that, like seemingly more complex AI systems, raise serious concerns about transparency, accuracy, and fairness.[42]

   Algorithmic systems like PATTERN and the tools used in Idaho are likely pervasive,[43] and their harms are exactly the harms the Draft Memorandum is meant to address. Excluding such rudimentary systems would significantly undermine the memorandum's efficacy. Importantly, the Draft Memorandum's definition of "artificial intelligence" can be read to encompass such algorithmic systems—for example, as a "set of techniques . . . that is designed to approximate a cognitive task."[44] OMB should clarify that this language should be interpreted to include such systems that rely on statistical modeling, especially given that many of the definition's prongs seem to focus on advanced and emerging AI technologies such as those that "learn from experience and improve performance," that "solve tasks requiring human-like perception [or] cognition," or that are "designed to think or act like a human."[45]

   More generally, the definition hinges on terms—like "human-like perception, "act rationally," "cognitive task," and "embodied robot," among others—that lack clear definitions in AI research or policy communities.[46] In addition, the structure of the definition, where AI systems are described as designed to "act like...human[s]," "act rationally," "achieve goals," or "perform tasks," employs a type of AI anthropomorphism that may obscure the roles and responsibilities of human decision-makers who make critical decisions shaping the design and deployment of AI systems, including deciding whether the systems are built or used in the first place.[47] Efforts to regulate government uses of AI should steer clear of definitions that may

---

[41] National Institute of Justice, 2021 Review and Revalidation of the First Step Act Risk Assessment Tool at 16 (2021), https://nij.ojp.gov/library/publications/2021-review-and-revalidation-first-step-act-risk-assessment-tool.

[42] See Formal Statement of the American Civil Liberties Union For a Stakeholder Engagement Session on First Step Act Implementation, ACLU (September 27, 2022), https://www.aclu.org/wp-content/uploads/document/ACLU_PATTERN_Public_Comment.pdf; Coalition Letter on the Use of PATTERN Risk Assessment in Prioritizing Release in Response to the COVID-19 Pandemic, ACLU (April 3, 2020), https://www.aclu.org/letter/coalition-letter-use-pattern-risk-assessmentprioritizing-release-response-covid-19-pandemic; ACLU, Comment Letter to Department of Justice on PATTERN First Step Act (Sept. 3, 2019), https://civilrights.org/resource/comment-letter-to-department-of-justice-on-pattern-first-step-act/.

[43] Julia Angwin, The Seven-Year Struggle to Hold an Out-of-Control Algorithm to Account, The Markup (Oct. 8, 2022), https://themarkup.org/newsletter/hello-world/the-seven-year-struggle-to-hold-an-out-of-control-algorithm-to-account; Dillon Reisman, *How the Government Relies on Algorithms to Allocate Healthcare Benefits—and Why These Secret Formulas Threaten Patients' Fundamental Rights*, ACLU NJ (Aug. 9, 2022), https://www.aclu-nj.org/en/news/howgovernment-relies-algorithms-allocate-healthcare-benefits-and-why-these-secret-formulas.

[44] Draft Memorandum at 22.

[45] *Id.*

[46] Murray Shanahan, *Talking About Large Language Models*, arXiv (2022), https://arxiv.org/abs/2212.03551.

[47] David Watson, *The Rhetoric and Reality of Anthropomorphism in Artificial Intelligence*, 29 Minds and Machines 417-440 (2019), https://link.springer.com/article/10.1007/s11023-019-09506-6.

miscommunicate the capabilities of AI systems, contribute to unwarranted AI "hype,"[48] or diminish the role of human decision-makers, and should instead focus on the concrete inputs and outputs of AI systems and the tasks such systems are used to support or complete. Absent adjustments to the definition itself, appropriate revisions and clarifications to the "technical context" for interpreting the definition, as described below, could help mitigate these issues.

OMB could significantly clarify and strengthen the scope of "covered AI" in the memorandum with appropriate edits to the "technical context" included in the Draft Memorandum for interpreting the definition of "artificial intelligence."[49] In addition to clarifying that systems relying on or derived from statistical modeling are considered AI for the purposes of the memo, OMB should clarify part 2 of the "technical context," which states that "[t]his definition of AI does not include robotic process automation or other systems whose behavior is defined only by human-defined rules or that learn solely by repeating an observed practice exactly as it was conducted."[50] "Robotic process automation" is a term often used to describe a broad class of technologies, including systems that use or may be integrated with artificial intelligence or machine learning methods.[51] As highlighted in a recent report about government uses of AI systems, process automation of varying complexity is often closely intertwined—and at times inextricable from—government uses of statistical modeling and AI systems.[52] OMB should define or clarify what processes are considered "robotics process automation," and to the extent that the language in part 2 excludes algorithmic systems such as PATTERN or those deployed in Idaho's Medicaid benefits determinations, OMB should remove such language altogether.

OMB should also clarify part 3 of the "technical context," which states in part that "the technical complexity of a system (e.g., the number of parameters in a model, the type of model, or the amount of data used for training purposes) is not a relevant consideration for determining whether it constitutes AI."[53] Information about a model's complexity can be useful for understanding a system's functions and impacts. If the goal of part 3 is to clarify that no system should be considered too simple to be a covered AI system under the memo due to its relative

---

[48] Melissa Heikkilä, *Unpacking the Hype Around OpenAI's Rumored New Q* Model*, MIT Technology Review (Nov. 27, 2023) ("And while it might be great PR, these hype cycles do more harm than good for the entire field by distracting people from the real, tangible problems around AI.").

[49] Draft Memorandum at 23.

[50] Draft Memorandum at 23.

[51] For example, several large robotic process automation (RPA) vendors with public sector clients describe how RPA and AI can be leveraged together. *See, e.g.*, *AI and RPA*, UiPath, https://www.uipath.com/automation/ai-and-rpa (last visited November 27, 2023);*Combine the Power of RPA and AI to Empower Rapid End-to-End Business Process Automation*, Automation Anywhere, https://www.automationanywhere.com/rpa/intelligent-automation (last visited November 27, 2023).

[52] Grant Fergusson, Outsourced and Automated: How AI Companies Have Taken Over Government Decision-Making, Electronic Privacy Information Center (2023), https://epic.org/wp-content/uploads/2023/09/FINAL-EPIC-Outsourced-Automated-Report-Appendix-Included.pdf

[53] Draft Memorandum at 23.

![ACLU logo]

lack of technical complexity as measured through parameter count, training data size, or related metrics, that point should be made explicitly.

  D. OMB Should Extend the Draft Memorandum's Provisions to Federally Funded Programs

The Draft Memorandum currently applies only to federal agencies—namely, any "executive department, military department, Government corporation, Government controlled corporation, or other establishment in the executive branch," with a few enumerated exceptions.[54] The scope of Draft Memorandum at the federal level is appropriately broad, and OMB should not narrow the definition of "agency" in the final memorandum. However, the Draft Memorandum appears to exclude state and local programs receiving federal assistance; this exclusion will leave many particularly dangerous uses of AI unregulated, and OMB should take steps to expand the scope of the memorandum's applicability to federally funded programs. In addition, it is crucial that OMB ensure that the Draft Memorandum reach federal contractors and AI procured by federal agencies, which the Draft Memorandum already contemplates.[55]

The exclusion of federally funded programs is particularly pernicious because federal funds may help support uses of AI with significant impacts on rights and safety. For example:

- As described, above, Idaho's Department of Health and Welfare was employing algorithmic systems to determine benefits for federally funded Medicaid programs.[56] Although the system cut some individuals' benefits by as much as 30 percent, officials were unable to explain why determinations were reached, and litigation by ACLU of Idaho revealed that the system was implemented without meaningful safeguards.
- An ACLU and Human Rights Data Analysis Group audit of an algorithmic risk-scoring system used to inform child welfare decision-making in Allegheny County, Pennsylvania highlighted several ways in which the algorithm's design and deployment could enable algorithmic bias.[57] The risk-scoring system could potentially disproportionately flag Black families and families with disabilities for investigation. The audit highlighted the

---

[54] *See* 44 U.S.C. § 3502(1).

[55] Draft Memorandum at 21.

[56] Testimony of Ritchie Eppink, Hearing AI in Government Before the S. Comm. On Homeland Security & Government Affairs (May 16, 2023), https://www.hsgac.senate.gov/hearings/artificial-intelligence-in-government.

[57] Marissa Gerchick et al., *How Policy Hidden in an Algorithm is Threatening Families in This Pennsylvania County*, ACLU (Mar. 14, 2023), https://www.aclu.org/news/womens-rights/how-policy-hidden-in-an-algorithm-is-threatening-families-in-this-pennsylvania-county; Marissa Gerchick et al., ACLU, The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool (2023), https://www.aclu.org/the-devil-is-in-the-details-interrogating-values-embedded-in-the-allegheny-family-screening-tool. Allegheny County and its Department of Human Services receive federal funds. *DHS Funding*, Allegheny County, https://www.alleghenycounty.us/Human-Services/About/Funding-Sources.aspx (last visited Nov. 22, 2023); *Allegheny, County Of*, TAGGS, https://taggs.hhs.gov/Detail/RecipDetail?arg_EntityId=swAAHUn5jiXXGX5RfqF%2Fmg%3D%3D (last visited Nov. 22, 2023).

system's use of existing government databases, including county child welfare, juvenile probation, and behavioral health records. Problematically, those databases reflect the lives of those who have more contact with government agencies and systems shaped by historical and ongoing discrimination—not necessarily those who pose greater "risk" to their children. Additionally, the outcome the tool predicts is the risk of child removal by the County, based on its historical practices. Because government databases, including those regarding child removal statistics, reflect systems shaped by historical and ongoing discrimination, using them to identify the characteristics of households more likely to have a child removed means selecting from a pool of factors that over-represents some groups of people and underrepresent others.

- "[C]rime-fighting grants" provided through the U.S. Department of Housing and Urban Development have been used by local housing authorities to deploy AI-powered surveillance.[58] For example, in "rural Scott County, Va., cameras equipped with facial recognition [technology] scan everyone who walks past them, looking for people barred from public housing."[59] Numerous other uses of facial recognition and similar technology in federally funded housing have been well documented.[60] Likewise, public housing authorities may rely on algorithmically driven tenant screening, including criminal background checks used as a prerequisite for public housing, often with discriminatory effects on over-policed populations.[61]

- Various federal programs such as the Department of Justice's Community Oriented Policing Services program and the Edward Byrne Memorial Justice Assistance Grant Program, among others,[62] fund local law enforcement deployment of AI-powered surveillance.[63] These can include facial recognition technology, automated license plate

---

[58] Douglas MacMillan, *Eyes on the Poor: Cameras, Facial Recognition Watch Over Public Housing*, Washington Post (May 16, 2023), https://www.washingtonpost.com/business/2023/05/16/surveillance-cameras-public-housing.
[59] *Id.*
[60] *Id.*; Dan Bateyko, *Taken for Granted: Where's the Oversight of AI and Federal Funding?*, CDT (Aug. 7, 2023), https://cdt.org/insights/taken-for-granted-wheres-the-oversight-of-ai-and-federal-funding.
[61] DeMetria McCain, Principal Deputy Assistant Secretary for Fair Housing and Equal Opportunity, U.S. Department of Housing and Urban Development, Memorandum on Implementation of the Office of General Counsel's Guidance on Application of Fair Housing Act Standards to the Use of Criminal Records by Providers of Housing and Real Estate-Related Transactions 2 (June 10, 2022), https://www.hud.gov/program_offices/fair_housing_equal_opp/fheo_guidance ("[H]ousing providers sometimes utilize third-party companies to independently screen and reject applicants using algorithms that may contain racial or other prohibited bias in their design."); *see* Comments of the ACLU, Tenant Screening Request for Information, Docket No. FTC-2023-0024 (May 30, 2023), https://www.aclu.org/wp-content/uploads/2023/07/2023.05.30-ACLU-Comment-to-FTC-CFPB-Tenant-Screening-RFI.pdf (describing private uses of algorithmic tenant screening).
[62] Brian Naylor, *How Federal Dollars Fund Local Police*, NPR (June 9, 2020), https://www.npr.org/2020/06/09/872387351/how-federal-dollars-fund-local-police.
[63] Matthew Guariglia & Dave Maass, *How Police Fund Surveillance Is Part of the Problem*, EFF (Sept. 23, 2020), https://www.eff.org/deeplinks/2020/09/how-police-fund-surveillance-technology-part-problem.

![ACLU](ACLU logo)

readers, predictive policing, gunshot detection and social media surveillance.[64] Even COVID-19 relief funds have been used for such AI-driven law enforcement technologies.[65]

Excluding these use cases—and countless others—would drastically undermine the Draft Memorandum's efficacy. President Biden's Executive Order already recognizes the importance of reaching federally funded state and local level programs in part. For example, it requires the federal Secretary of Health and Human Services to issue a plan "addressing the use of automated or algorithmic systems in the implementation by States and localities of public benefits and services administered by the Secretary."[66]

OMB should follow the lead established by President Biden and ensure the protections in the Draft Memorandum extend to federally funded state and local programs. Recognizing that such programs can have complex administrative regimes, OMB has several options to ensure that the extension is effective. Those options might include empowering CAIOs to survey state and local use cases of AI in federally funded programs, to issue corresponding guidance, and to review rulemaking authority related to the use of AI in federally funded programs. In the final memorandum, OMB may also identify specific federally funded programs, including as illustrative examples, that should adhere to the requirements of the memorandum.

II.   OMB Should Maintain or Expand the Scope of Rights-Impacting AI

Under the Draft Memorandum, covered agencies must implement minimum risk management practices for AI that impacts rights and safety, and OMB should either expand the scope of "rights-impacting AI" or—at minimum—refuse to narrow it.

In the Draft Memorandum, "rights-impacting AI" is defined as AI that has a "legal, material, or similarly significant effect" on individuals or communities across three critical groups of rights: (1) civil rights and civil liberties, including privacy, freedom of speech, voting, autonomy, and the right to be free from discrimination, excessive punishment, and unlawful surveillance; (2) equitable access to critical life opportunities that have traditionally been protected by the law—such as education, housing, credit, and employment; and, (3) access to critical resources or services—such as "healthcare, financial services, social services,

---

[64] Tate Ryan-Mosely, *How US Police Use Counterterrorism Money to Buy Spy Tech*, MIT Technology Review (Dec. 7, 2022), https://www.technologyreview.com/2022/12/07/1064354/police-counterterrorism-money-buy-spy-tech-surveillance-fema; Brandon Block, *Federal Aid Is Supercharging Local WA Police Surveillance Tech*, Crosscut Cascade PBS (July 26, 2023), https://crosscut.com/investigations/2023/07/federal-aid-supercharging-local-wa-police-surveillance-tech.

[65] Chris Baumohl, *Two Years In, COVID-19 Relief Money Fueling Rise of Police Surveillance*, EPIC (Mar. 9, 2023), https://epic.org/two-years-in-covid-19-relief-money-fueling-rise-of-police-surveillance; Anastasia Valeeva, Wihua Li & Susie Cagle, *Rifles, Tasers and Jails: How Cities and States Spent Billions of COVID-19 Relief*, The Marshall Project (Sept. 7, 2022), https://www.themarshallproject.org/2022/09/07/how-federal-covid-relief-flows-to-the-criminal-justice-system.

[66] Executive Order 14110 of October 30, 2023, sec. 7.2(b)(i).

transportation, non-deceptive information about goods and services, and government benefits or privileges."[67]

The definition's broad scope is necessary to meet existing and imminent challenges as AI continues to develop at an incredible pace across governmental agencies and economic sectors. As described throughout these comments, artificial intelligence impacts individuals in numerous ways, including through various governmental uses:

- the use of AI-powered surveillance such as facial recognition technology and automated license plate readers by law enforcement;
- sophisticated AI-driven surveillance used by immigration agencies, including social media monitoring, facial recognition and other biometric surveillance, and automated surveillance towers "capable of identifying and capturing human faces";
- the use of AI to identify "risks" to national security and for placement on government watchlists;
- automated determinations of eligibility for governmental benefits or benefits amounts;
- scoring families for investigation by child welfare agencies; and,
- algorithmic screening of tenants in public housing.

In addition, well-documented AI usage by private entities may be reflected in governmental uses as well:

- AI tools are used to recruit and source job applicants, including through targeted advertising and specialized recruitment platforms, often with discriminatory effects on women, people of color, and older individuals;[68]

---

[67] *Id.*

[68] See Sheridan Wall & Hilke Schellmann, *LinkedIn's Job-Matching AI was Biased. The Company's Solution? More AI.*, MIT Tech. Rev. (June 23, 2021), https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence; *Facebook EEOC Complaints*, ACLU (Sept. 25, 2019), https://www.aclu.org/cases/facebook-eeoc-complaints; Ariana Tobin & Jeremy Merrill, *Facebook is Letting Job Advertisers Target Only Men*, ProPublica (Sept. 18, 2018), https://www.propublica.org/article/facebook-is-letting-job-advertisers-target-only-men; Ava Kofman & Ariana Tobin, *Facebook Ads Can Still Discriminate Against Women and Older Workers, Despite a Civil Rights Settlement*, ProPublica (Dec. 13, 2019), https://www.propublica.org/article/facebook-ads-can-still-discriminate-against-women-and-older-workers-despite-a-civil-rights-settlement; *see* Pauline Kim & Sharion Scott, *Discrimination in Online Employment Recruiting,* 63 St. Louis University Law Journal 1, 8 (2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3214898 (describing study that found that simulated users who were identified as male or female and who engaged in identical web browsing activities related to job searches on Google were shown very different ads, with an ad for coaching on higher paying executive jobs shown significantly more often to men).

- automated screening tools are used to "to filter out or rank applicants with automated resume screening based on knockout questions, keyword requirements, or specific qualifications or characteristics";[69]
- targeted advertising for housing and tenant screening scores can excluded people of color, with adverse housing history, or with criminal records unrelated to their qualification for housing from being able to find housing;[70]
- the algorithms underlying credit scores, including emerging, alternative models, rely on historic data and consequently "bake in" existing systemic biases, resulting in in disproportionately lower scores for already marginalized populations;[71]
- algorithmic systems and similar technologies are used in evaluating mortgage applications[72] and in insurance underwriting,[73] which have been found to disproportionately deny loans to people of color;
- some clinical decision tools used to recommend patient care have been found to have a racial bias, despite not considering race as an express input;[74] and,
- private and public educational institutions are deploying AI acquired from private vendors to detect student behavior captured in video surveillance or to analyze student sentiment in social media posts and private messages.[75]

---

[69] ReNika Moore, ACLU, Testimony Before the Equal Employment Opportunity Commission (Jan. 31, 2023), https://www.eeoc.gov/meetings/meeting-january-31-2023-navigating-employment-discrimination-ai-and-automated-systems-new/moore (citing Aaron Rieke & Miranda Bogen, *Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias*, Upturn (Dec. 10, 2018), https://www.upturn.org/work/help-wanted).

[70] Comments of the ACLU, Tenant Screening Request for Information, Docket No. FTC-2023-0024 (May 30, 2023), https://www.aclu.org/wp-content/uploads/2023/07/2023.05.30-ACLU-Comment-to-FTC-CFPB-Tenant-Screening-RFI.pdf (describing private uses of algorithmic tenant screening); Charge of Discrimination, HUD v. Facebook (Mar. 28, 2019), https://www.hud.gov/sites/dfiles/Main/documents/HUD_v_Facebook.pdf; Muhammad Ali et al., Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead to Skewed Outcomes (2019), https://arxiv.org/abs/1904.02095; Kaveh Waddell, *How Tenant Screening Reports Make it Hard for People to Bounce Back from Tough Times*, Consumer Reports (Mar. 11, 2021), https://www.consumerreports.org/electronics/algorithmic-bias/tenant-screening-reports-make-it-hard-to-bounce-back-from-tough-times-a2331058426.

[71] National Consumer Law Center, Past Imperfect: How Credit Scores and Other Analytics "Bake In" and Perpetuate Past Discrimination at 5-7 (2016), https://www.nclc.org/resources/past-imperfect-how-credit-scores-and-other-analyticsbake-in-and-perpetuate-past-discrimination; Student Borrower Protection Center, Educational Redlining at 15 (2020), https://protectborrowers.org/wp-content/uploads/2020/02/Education-Redlining-Report.pdf.

[72] Emmanuel Martinez & Lauren Kirchner, *The Secret Bias Hidden in Mortgage-Approval Algorithms*, The Markup (Aug. 25, 2021), https://themarkup.org/denied/2021/08/25/the-secret-bias-hidden-in-mortgage-approval-algorithms.

[73] Ronda Lee, *AI Can Perpetuate Racial Bias in Insurance Underwriting*, Yahoo! Money (Nov. 1, 2022), https://money.yahoo.com/ai-perpetuates-bias-insurance-132122338.html.

[74] Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 Sci. 447 (Oct. 25, 2019), https://www.science.org/doi/10.1126/science.aax2342.

[75] Letter from Rep. Lori Trahan et al. to Hon. Miguel Cardona, Secretary, U.S. Department of Education (Oct. 19, 2023), https://trahan.house.gov/uploadedfiles/final_edtech_letter_10.12.23.pdf; Chad Marlow et al., ACLU, Digital

The critical impacts of AI extend beyond those examples, and the broad definition of "rights-impacting AI" is necessary to cover the broad scope of AI's impact. To ensure the final memorandum provides meaningful, comprehensive protections, OMB should resist calls to narrow the scope of "right-impacting AI."

Moreover, the final memorandum should expand on—or at least maintain—the Draft Memorandum's list of presumptively rights-impacting uses of AI. Uses that are presumed to be rights impacting include "[d]ecisions to block, remove, hide, or limit the reach of protected speech," law enforcement use of risk assessments and facial recognition, school discipline, tenant screening, determining the terms of employment, and recommendations or decisions about child welfare or child custody.[76] Presumptions help reduce the cost to individuals if agencies incorrectly apply the definition of "right-impacting AI." The list of use cases that are presumed to be rights-impacting are too important and too fundamental to leave to agencies' discretion in the first instance. Instead, the presumptions appropriately place the burden on agencies to demonstrate the *lack* of impact on rights in the most critical instances.

However, many of the presumptive rights-impacting use cases are deployed by state and local governmental entities receiving federal funding—use cases that are currently excluded from the scope of the Draft Memorandum. As urged above, the final memorandum should extend its requirements to state and local programs that receive federal assistance. In addition, if the Draft Memorandum's scope is expanded to include national security systems and intelligence agencies, the list of presumptively rights-impacting uses cases should be adjusted accordingly to include national security decisions about surveillance, governmental watchlisting, searches at the border or related to immigration, and national security-related risk assessments.

III.   **The Final Memorandum Should Maintain the Strength of Key Minimum Practices While Bolstering Others**

A. The Draft Memorandum Appropriately Requires Independent Assessments of AI Harms and Discontinuing Harmful AI

The Draft Memorandum underscores that an AI system should be deployed for a particular use case only if its harms are outweighed by its benefits and the harms have been

---

Dystopia: The Danger in Buying What the EdTech Surveillance Industry Is Selling at48-50 (2023), https://www.aclu.org/report/digital-dystopia-the-danger-in-buying-what-the-edtech-surveillance-industry-is-selling; Elizabeth Laird et al., Center for Democracy & Technology, Off Task: EdTech Threats to Student Privacy and Equity in the Age of AI (2023), https://cdt.org/insights/report-off-task-edtech-threats-to-student-privacy-and-equity-in-the-age-of-ai; Ari Sen & Dereka K. Bennett, *Tracked: How Colleges Use AI to Monitor Student Protests*, Dallas Morning News (Sept. 20, 2022), https://interactives.dallasnews.com/2022/social-sentinel; Mark Keierleber, *Meet the Gatekeepers of Students' Private Lives*, The 74 (May 2, 2022), https://www.the74million.org/article/meet-the-gatekeepers-of-students-private-lives.
[76] Draft Memorandum at 12-13.

appropriately mitigated.[77] Several key components of the minimum risk management practices underscore this commitment to discontinuing—or not deploying—harmful AI:

- The Draft Memorandum instructs that AI impact assessments should assess the "potential risks of using AI, as well as what, if any, additional mitigation measures, beyond these minimum practices, the agency will take to help reduce these risks." The Draft Memorandum emphasizes, "The expected benefits of the AI functionality should be considered against its potential risks, and if the benefits do not meaningfully outweigh the risks, agencies should not use the AI."[78]

- Agencies are also required to test AI in its "real-world context," which must "mirror as closely as possible the conditions in which the AI will be deployed" and consider both the technology itself and feedback from likely operators. Through these results, agencies should demonstrate that the AI will "achieve its expected benefits while sufficiently mitigating risks associated with the AI, or else the agency should not use the AI."[79] Real-world testing must be accompanied by an independent review.[80]

- Ongoing monitoring should continue to assess risks to rights and safety, including in a real-world context, and develop new mitigation techniques to reduce those risks. Again: "Where the AI's risks to rights or safety exceed an acceptable level and where mitigation is not practicable, agencies must stop using the affected AI as soon as is practicable."[81]

The Draft Memorandum's emphasis on real-world, independent evaluation of risks and benefits of AI—and discontinuing, decommissioning, or not deploying AI with excessive risks—is appropriate and should be maintained in the final memorandum. Experience has demonstrated that real-world testing environments are necessary to identify harms from human-AI interaction, which may arise from the data used to train the model, the model's use in unintended or unanticipated environments, incorrect interpretation of its output by operators, or other factors.

For example, as highlighted by an independent analysis of the system, one lending model developed as an alternative to traditional credit scores "charged higher interest rates and origination fees" for borrowers who had attended Historically Black Colleges and Universities (HBCUs) and Hispanic-Serving Institutions (HSIs).[82] Subsequent investigations revealed that although the lending model did not rely on specific institutions, it did categorize those institutions based on "average incoming standardized test scores," which led to the

---

[77] Draft Memorandum at 20.
[78] Draft Memorandum at 15.
[79] Draft Memorandum at 16.
[80] *Id.*
[81] Draft Memorandum at 17.
[82] Student Borrower Protection Center, Educational Redlining at 15 (2020), https://protectborrowers.org/wp-content/uploads/2020/02/Education-Redlining-Report.pdf.

disproportionate treatment of HBCUs and HSIs.[83] Testing in a real-world context by an independent entity was critical to identifying those disparities, and it is possible that further internal real-world testing could have not only identified those disparities before the system was deployed but also led to mitigations, which are still ongoing.[84]

Further, real-world, independent evaluation of the risks and benefits of AI should include feedback from likely operators. For example, the Allegheny Family Screening Tool ("AFST") was developed and implemented by a county child welfare agency to help make screening decisions about whether or not to investigate reports of alleged child maltreatment.[85] The AFST's risk score was designed to reflect the likelihood that the agency would remove a child from the family within two years of the report—not the likelihood that a child will be abused or neglected,[86] a fact acknowledged by the AFST's creators.[87] A later study observed that call screeners who used the tool found that screeners and their supervisors disagreed with the use of child removal as a proxy for child abuse or neglect. Some explained that, based on their experience, "children were often placed in foster care without any concerns of child abuse or neglect."[88] Conversely, another caseworker indicated that a child might never enter the foster system even where legitimate concerns of child maltreatment do exist. On balance, then, the primary tool users found the risk being predicted by the tool to be both over- and under-inclusive. The dissonance between the tool developers and those who would eventually use the tool points to the importance of requiring truly independent evaluations of the pros and cons of using an AI tool and ensuring that pre-deployment real world testing includes solicitation and meaningful consideration of the views of, among others, those who would ultimately use the tool.

---

[83] Relman Colfax PLLC, Initial Report of the Independent Monitor 22 (2021), https://www.relmanlaw.com/cases-406.

[84] Relman Colfax PLLC, Third Report of the Independent Monitor 8 (2021), https://www.relmanlaw.com/cases-406.

[85] Marissa Gerchick et al., *How Policy Hidden in an Algorithm is Threatening Families in This Pennsylvania County*, ACLU (Mar. 14, 2023), https://www.aclu.org/news/womens-rights/how-policy-hidden-in-an-algorithm-is-threatening-families-in-this-pennsylvania-county; Marissa Gerchick et al., ACLU, The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool (2023), https://www.aclu.org/the-devil-is-in-the-details-interrogating-values-embedded-in-the-allegheny-family-screening-tool.

[86] Allegheny Cnty. Dep't of Hum. Servs., *Frequently Asked Questions* 6 (updated Apr. 2019), https://www.alleghenycountyanalytics.us/wp-content/uploads/2019/05/FAQs-from-16-ACDHS-26_PredictiveRisk_Package_050119_FINAL-8.pdf.

[87] Rhema Vaithianathan et al., Allegheny Family Screening Tool: Methodology, Version 2, at 7 (2019), https://www.alleghenycountyanalytics.us/wp-content/uploads/2019/05/Methodology-V2-from-16-ACDHS-26_PredictiveRisk_Package_050119_FINAL-7.pdf ("[T]here are valid concerns that the AFST model, and other models trained to predict system outcomes like out-of-home placement, may be predicting the risk of institutionalized or system response rather than the true underlying risk of adverse events."). Ethicists evaluating the AFST concluded that the mismatch between the tool's measurement and use case would be mitigated through subsequent investigation. Rhema Vaithianathan et al., *Developing Predictive Models to Support Child Maltreatment Hotline Screening Decisions: Allegheny County Methodology and Implementation,* Sec. 2, at 4 (2017), https://www.alleghenycountyanalytics.us/wp-content/uploads/2019/05/Methodology-V1-from-16-ACDHS-26_PredictiveRisk_Package_050119_FINAL.pdf.

[88] Hao-Fei Cheng et al., *How Child Welfare Workers Reduce Racial Disparities in Algorithmic Decisions*, 2022 Proc. CHI Conf. on Human Factors in Computing Sys. 13 (2022), https://dl.acm.org/doi/10.1145/3491102.3501831.

The final memorandum should embody this commitment to assessing AI harms and discontinuing harmful AI throughout its provisions. For example, Section 4 includes numerous references to removing barriers to the development and deployment of AI, but almost none to managing the risks posed by AI—and none to discontinuing AI where harms outweigh the risks. It is critical that agencies' AI strategies and other documentation center the need to discontinue harmful uses of AI. Further, the Draft Memorandum contemplates that documentation of minimum risk management practices[89]—including independent assessments of the AI—will be made available to OMB; OMB should further require that documentation to be available in an accessible, approachable format to the public.

     B.   The Draft Memorandum Correctly Recognizes that Minimum Practices Should Be Applicable Even if AI only Influenced the Decision or Outcome

The Draft Memorandum appropriately recognizes that although AI may not independently make decisions or fully automate a task, it may nonetheless carry risks to rights and safety. For example, the definition of "risks from the use of AI" expressly applies "regardless of whether . . . the AI merely informs the decision or action, partially automates it, or fully automates it."[90] Risks from the use of AI also include where "the humans involved . . . are not aware of how or to what extent the AI influenced or automated the decision or action."[91] Several other provisions in the Draft Memorandum, such as purposes that are presumed to be rights- and safety-impacting, apply where the AI "is used to control or meaningfully influence" an outcome.[92]

The Draft Memorandum's treatment of AI used to "influence" decisions or outcomes corresponds to how AI is actually used in practice, and OMB should maintain this approach. For example, one predictive model used in colleges and universities evaluates individual students' likelihood of academic success and assigns them a corresponding "risk score." One investigation found the model's risk scores correlated with students' race, and in some cases, expressly incorporated it as a "high-impact predictor."[93] Academic advisors often review students' risk scores, and although the model did not independently make decisions about students, its scores might nonetheless "leave advisers with an immediate and potentially life-changing impression of students and their prospects within a given major."[94] Although AI did not make the final

---

[89] Draft Memorandum at 13.
[90] Draft Memorandum at 24.
[91] *Id.*
[92] *Id.* at 11-12.
[93] 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.
[94] *Id.*

determination, its influence was significant, and OMB should ensure the final memorandum covers such scenarios.

      C. The Discretion to Waive Minimum Practices or Otherwise Avoid the Memorandum's Obligations Should Be Further Cabined

In the final memorandum, OMB should ensure that waivers of the Draft Memorandum's provisions and other limitations on its scope are not abused. The Draft Memorandum permits agencies to waive the minimum risk management practices.[95] Under the Draft Memorandum, a waiver would be permitted for "a specific covered AI application or component" after the Chief AI Officer makes "a written determination, based upon a system-specific risk assessment, that fulfilling the requirement would increase risks to safety or rights overall or would create an unacceptable impediment to critical agency operations."[96]

OMB should cabin agencies' discretion to make this determination. In particular, the final memorandum should include descriptions or definitions of what it means to increase risks to safety or rights or to unacceptably impede "critical agency operations." Examples of use cases that meet these requirements—or do not meet these requirements—would also be helpful. At minimum, OMB should clarify that ordinary administrative burdens of discontinuing an AI system due to its harms does not meet these thresholds.

In addition, the Draft Memorandum recognizes other exercises of agency discretion that limit the Memorandum's applicability:

- Agencies retain significant discretion in the first instance to determine which systems qualify as "covered AI" or rise to the level of rights- and safety-impacting AI[97] and do not appear to be required to report the former to OMB.
- "National security systems" are exempt from the Draft Memorandum's requirements altogether, and although the Memorandum incorporates an existing definition for the term,[98] it does not provide procedures for determining what systems meet that definition or reviewing those determinations.[99]
- Covered AI is also subject to the use case inventory only "to the extent practicable" and as permitted for "the protection of privacy and of sensitive law enforcement, national

---

[95] Draft Memorandum at 14-15.

[96] Draft Memorandum at 14.

[97] Draft Memorandum 4-5, 10 ("Agencies must review each use of AI that they are developing or using to determine whether it matches the definition of safety-impacting or rights impacting.").

[98] 44 U.S.C. § 3552(b)(6).

[99] The term "national security system" is applicable to federal cybersecurity rules, and the law generally defers to agency heads in evaluating "information security programs" for national security systems. *See* 44 U.S.C. §§ 3555(c), 3557. Such deference is inappropriate when the evaluation is not regarding the security of an agency's own systems, but the system's impact on individuals' and communities' rights and safety.

security, and other protected information,"[100] but the Draft Memorandum provides no guidelines in assessing practicability or when covered AI implicates "sensitive" or "protected information."

These limitations create opportunities for agencies to avoid the Draft Memorandum's obligations.

OMB may avoid abuses of the waiver process and other limitations by ensuring robust reporting and review by OMB of agencies' use of those authorities. The Draft Memorandum already envisions reporting in some instances. For example, beginning "with the use case inventory for 2024, agencies will be required, as applicable, to identify and report . . . any related extensions and waivers granted under" the Draft Memorandum.[101] Agencies are similarly expected to document and report to OMB determinations that a use of AI is not, in fact, safety- or rights-impacting.[102] Other determinations, however, are not reported to OMB, including the determination that a system does not qualify as "covered AI" or is deemed to be a "national security system" or that a use case involves "sensitive law enforcement" and "other protected information."

OMB should consequently expand the Draft Memorandum's reporting requirements to include each of these scenarios. In each instance, the agency should detail the scope, justification, and evidence supporting the determination. The threshold determination that a particular system constitutes "covered AI" is particularly crucial, as the entirety of the Draft Memorandum's obligations depend on that determination. However, overextending reporting for that threshold determination could overburden both agencies and OMB. Consequently, OMB should ensure that its forthcoming "detailed instructions for the [AI use case] inventory"[103] contemplate when agencies should report liminal determinations that particular systems did *not* constitute "covered AI."

In addition to robust reporting requirements for determinations regarding "covered AI," "national security systems," and "protected information," OMB should consider other safeguards to cabin agency discretion. These safeguards might include review by the agency AI oversight board, agency civil rights authorities, or OMB. OMB should also consider time-limiting agency determinations that exempt systems from the Draft Memorandum's obligations.

D. Notice Should Be Further Bolstered and Should Be Provided Before Use of Rights- or Safety-Impacting AI

---

[100] Draft Memorandum at 4 n.8.
[101] Draft Memorandum at 4-5.
[102] Draft Memorandum at 11.
[103] Draft Memorandum at 4.

The Draft Memorandum envisions providing notice of AI's use through two means: "plain language documentation" in the AI use case inventory and notice to "negatively affected individuals."[104] The Draft Memorandum establishes important parameters around both sets of notices to ensure that they are meaningful for impacted individuals and communities. The inventory, for example, must be in plain language and "generally accessible" while still being "adequately detailed."[105] The inventory must also be located "in contexts where people will interact with or be impacted by the AI."[106] Similarly, the individual notice for negative actions must be timely, consistent with the Plain Writing Act, and available "in multiple languages and through alternative formats and channels," as appropriate.[107] Agencies are encouraged, but not required to provide explanations for decisions and actions by AI.[108]

In the final memorandum, OMB should provide further parameters around the notice requirements for plain language, explainability, timeliness, and accessibility. It should also consider mandating individual notice *prior* to the use of rights- or safety-impacting AI. Each of these five principles are central to the Administration's *Blueprint for an AI Bill of Rights*,[109] and OMB should integrate them into the final memorandum.

First, OMB should bolster the Draft Memorandum's plain language requirements. The Draft Memorandum contemplates agencies' adherence to the Plain Writing Act "if applicable."[110] The Plain Writing Act, however, is applicable only to communications about benefits, taxes, or legal requirements,[111] and by some measures, agency compliance with the Act has been uneven.[112] The final memorandum should supplement that baseline requirement by requiring that both the use case inventory and, especially, notices of negative actions be provided in a short form. The short form should, in clear and plain language, state the decision taken, explain how the decision was reached—as described below—and provide clear instructions for

---

[104] Draft Memorandum at 18, 20.
[105] *Id.*
[106] *Id.*
[107] *Id.* at 20.
[108] *Id.* at 20.
[109] White House Office of Science and Technology Policy, Blueprint for an AI Bill of Rights (2022), https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf [hereinafter "AI Bill of Rights"].
[110] Draft Memorandum at 20; Pub. L. No. 111-274 (codified at 5 U.S.C. § 301 note), https://www.congress.gov/111/plaws/publ274/PLAW111publ274.pdf.
[111] Pub. L. No. 111-274, sec. 3(2) (codified at 5 U.S.C. § 301 note), https://www.congress.gov/111/plaws/publ274/PLAW111publ274.pdf.
[112] *See* Center for Plain Language, 2022 Federal Plain Language Report Card (2023), https://centerforplainlanguage.org/2022-federal-plain-language-report-card.

recourse.[113] The *AI Bill of Rights* similarly underscored that notices should be "brief and clear" and able to be understood quickly.[114]

Second, OMB should make explanations of AI's decisions mandatory. Under the Draft Memorandum, agencies are only "strongly encouraged to provide explanations for such decisions and actions."[115] Such explanations "might include . . . how and why the AI-driven decision or action was taken" or at least "the general nature of such AI decisions through context such as the data that the decision relied upon, the design of the AI, and the broader decision-making context in which the system operates."[116] The explanations should be mandatory. The basic tenets of due process require that a governmental "decisionmaker should state the reasons for his determination and indicate the evidence he relied on,"[117] and decisions by AI should be no different. Explainability is a cornerstone of the Administration's *AI Bill of Rights*, which stated, "Automated systems should provide explanations that are technically valid, meaningful and useful to you and to any operators or others who need to understand the system, and calibrated to the level of risk based on the context."[118]

Third, the final memorandum should include more concrete requirements for "timely" notice. Timely notice is crucial in ensuring that individuals may appeal or contest an AI's negative impacts, as contemplated by the Draft Memorandum.[119] The final memorandum should specify that the appeals process should give "people ample time and opportunities to respond . . . and opportunities to extend" deadlines,[120] while reducing "burdens such as time spent gathering records and documentation needed to prove eligibility, travel time associated with developing and submitting the collection, or even time waiting to speak with agency personnel."[121] After appeals are filed, agencies should speedily resolve them, as "lost time [waiting for a government response] operates as a kind of tax—a 'time tax'—and it imposes a serious burden on our people as they interact with the Government."[122] The *AI Bill of Rights* also recognizes the importance of notice in seeking redress, stating, "Human consideration and fallback are only useful if they are

---

[113] *See* Center for Democracy & Technology et al., Civil Rights Standards for 21st Century Employment Selection Procedures at 25 (2022), https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures.

[114] AI Bill of Rights at 43.

[115] Draft Memorandum at 20.

[116] *Id.* at 20 n.37.

[117] Goldberg v. Kelly, 397 U.S. 254, 271 (1970).

[118] AI Bill of Rights at 40.

[119] Draft Memorandum at 20.

[120] Office of Management and Budget, Improving Access to Public Benefits Through the Paperwork Reduction Act (M-22-10) at 17 (2022), https://www.whitehouse.gov/wp-content/uploads/2022/04/M-22-10.pdf.

[121] *Id.* at 2; *accord* Office of Management and Budget, Strategies for Reducing Administrative Burden in Public Benefit and Service Programs at 1 (2022), https://www.whitehouse.gov/wp-content/uploads/2022/12/BurdenReductionStrategies.pdf (describing informational and learning, compliance, psychological, and redemption costs).

[122] Executive Order 14058 of December 13, 2021, sec. 1, 86 Fed. Reg. 71357, 71357 (Dec. 16, 2021).

conducted and concluded in a timely manner. . . . In time-critical systems [such as in healthcare or employment], this mechanism should be immediately available or, where possible, available before the harm occurs."

Fourth, the final memorandum should more robustly address accessibility, not only in notice provided to users but also in agencies' assessment and evaluation of AI. Federal law already requires agencies to provide accessible experiences to people with disabilities, and the final memorandum should expressly incorporate those requirements. President Biden's Executive Order on "Transforming Federal Customer Experience and Service Delivery To Rebuild Trust in Government" directed agencies to "ensur[e] the accessibility of services for customers with disabilities"[123] and Section 508 of the Rehabilitation Act requires that agencies "ensure" that any "electronic and information technology allows . . . individuals with disabilities who are members of the public . . . to have access to and use of information and data" from federal agencies that is comparable to individuals without disabilities.[124] The *AI Bill of Rights* recommends that AI systems "should be designed, developed, and deployed by organizations in ways that ensure accessibility to people with disabilities" and that notices and explanations be accessible.[125] Unfortunately, in the two decades since Section 508 was passed, not all agencies have met its requirements,[126] and OMB should ensure that accessibility is a requirement in agencies' use of AI, assessing its harms, establishing appropriate risk management practices, and providing notice and explanation of AI's use.

Finally, OMB should consider expanding its notice requirement to provide meaningful notice *prior* to the use of rights- and safety-impacting AI. Timely, individual notice prior to the use of AI is essential for protecting critical rights, as recognized by the *AI Bill of Rights*,[127] such as seeking accommodations for disabilities.[128] For example, providing notice to an individual with a disability *after* AI has been used to make a negative decision about their job application gives the individual no opportunity to seek meaningful accommodations. The Equal Employment Opportunity Commission has stated that prior notice is a "promising practice" for compliance with the disabilities rights laws, including "[i]nforming all job applicants and employees who are

---

[123] Executive Order 14058 of December 13, 2021, sec. 6(d), 86 Fed. Reg. 71357, 71364 (Dec. 16, 2021).

[124] 29 U.S.C. § 794d(1).

[125] AI Bill of Rights at 27, 43.

[126] U.S. Department of Justice & U.S. General Services Administration, Section 508 Report to Congress and the President at 9 (2023), https://www.justice.gov/crt/section-508-home-page-0 (finding that 55% of agency accessibility statements require "some remediation to meet requirements).

[127] AI Bill of Rights at 43 ("Users should receive notice of the use of automated systems in advance of using or while being impacted by the technology.").

[128] *See* Center for Democracy & Technology et al., Civil Rights Standards for 21st Century Employment Selection Procedures at 23 (2022), https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures; Center for Democracy & Technology, Comments on Revised Proposed Rules to Implement Local Law 144 of 2021 on Automated Employment Decision Tools at 7 (2023), https://cdt.org/insights/cdt-comments-scrutinize-nycs-revised-rules-that-leave-even-more-workers-unprotected-from-algorithmic-bias/ (ten days' notice likely insufficient for requesting accommodations).

being rated that reasonable accommodations are available" and "[d]escribing, in plain language and in accessible formats, the traits that the algorithm is designed to assess."[129] OMB should consequently consider requiring notice to be provided to individuals *prior* to any use of rights- and safety-impacting AI, including for purportedly positive outcomes, with the burden on agencies to demonstrate that that prior notice would be impracticable.

E. Feedback Mechanisms Should be Equitable and Approachable, Consistent with Prior Federal Guidance

The Draft Memorandum appropriately requires agencies to "consult affected groups, including underserved communities, in the design, development, and use of the AI" and in "the event of negative feedback, agencies must consider not deploying the AI or removing the AI from use."[130] We suggest stating expressly that the consultation and consideration of any feedback must be meaningful, not just *pro forma*. This means the feedback should be solicited both prior to the adoption of AI and, if implemented, during deployment. In addition, the feedback should be used in the pre-adoption evaluation of AI's benefits versus its harms as well as in impact evaluations once deployed. Finally, to promote accountability for following these guidelines, agencies should be required to disclose what the feedback was and how they acted on it, much as agencies do as part of the notice and comment process for proposed regulations.

One 2019 paper underscores the importance of *meaningfully* soliciting feedback from impacted communities. The paper reported the results of workshops where individuals from "affected communities" were asked about an algorithmic risk-scoring tool used in child welfare screening decisions.[131] Aside from the limited utility of soliciting feedback only after the tool was put in use, only 18 of the 95 workshop participants were people whose families had been investigated by child welfare authorities. The remainder worked for or with the child welfare agency using the risk-scoring tool, with varying degrees of direct contact with impacted families.[132] During the workshop, participants expressed concerns with the risk-scoring tool, including system-level concerns such as "low expectations of the benefits that the system could provide"; bias among caseworkers and built into data or tool itself; lack of information about how the tool weighed data; and, families' inability to dispute a score.[133] Rather than recommending that these concerns be considered in evaluation of the tool, the authors of the paper essentially advised better messaging or tweaks to the tool that would change perceptions

---

[129] U.S. Equal Employment Opportunity Commission, The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees (2022), https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence.
[130] Draft Memorandum at 19.
[131] Anna Brown, et al., Toward Algorithmic Accountability in Public Services, at 1 (2019), https://dl.acm.org/doi/pdf/10.1145/3290605.3300271.
[132] *Id.* at 3-4.
[133] *Id.* at 7-9.

but not its functioning[134]—a response that may have affected underrepresentation of impacted groups.

Moreover, agencies must ensure their methods for soliciting feedback are equitable, approachable, and accessible. The avenues for soliciting feedback envisioned by the Draft Memorandum may inadvertently be exclusionary. For example, the Draft Memorandum contemplates soliciting feedback through the Federal Register, post-transaction forms, or public hearings.[135] Although those are traditional means for agency interaction with the public—and should indeed be employed—they may require individuals to have sufficient time and resources to meaningfully respond. As OMB has observed, "these efforts are too often perceived by stakeholders and agencies alike as being siloed, inaccessible, or irrelevant compliance exercises with unclear purpose or benefit to communities, demonstrating insufficient consideration of the needs, interests, and priorities of diverse populations."[136] OMB has concluded that "there is no one-size-fts-all approach to stakeholder engagement" but identified a number of promising practices for ensuring engagement meets stakeholders where they are:[137]

- "collect[ing] feedback from communities not through a specific agency scope, but rather organized around a life event: surviving a natural disaster, a child with a disability transitioning to adulthood, and a Service member seeking civilian employment for the first time";
- deploying "well-designed mobile sites [to] drive better digital experiences";
- "directly increas[ing] community capacity by providing resources (grants, financial assistance, and procurement opportunities) directly into underserved communities"; and,
- simplifying forms through usability testing, and avoiding "legalistic, dense form[s] [that] could deter eligible [individuals] out of receiving these protections or add unnecessary stress for already stressed" communities.

---

[134] For instance, since changing the variables used by the tool to calculate risk scores to include family strengths and not just weaknesses would increase costs, the authors suggested "modifications [that] would produce a model that is mathematically equivalent to the original, but which may be perceived and responded to very differently by users and affected individuals." To the extent participants questioned whether use of the tool resulted in better outcomes for families, the authors recommended "[c]onveying how making use of data and algorithms leads to improved family and process outcomes," as if that was a foregone conclusion and not a metric to incorporate into future impact evaluations. And to the extent participants raised concerns about the "opacity of algorithmic decisionmaking," the recommendation was not to disclose more information, but a concern that such disclosure would not have the (desired) impact of improving caseworker-family relationships so further research was warranted about the usefulness of being transparent. *Id.* at 9-10.

[135] *Id.*

[136] Office of Management and Budget, Study to Identify Methods to Assess Equity: Report to the President at 30 (2021), https://www.whitehouse.gov/wp-content/uploads/2021/08/OMB-Report-on-E013985-Implementation_508-Compliant-Secure-v1.1.pdf.

[137] *Id.* at 32-35.

The final memorandum should direct agencies to assess burdens on impacted communities and develop strategies to reduce barriers when gathering feedback from impacted communities.

      F.   Non-Discrimination Considerations Should Extend Beyond Protections Currently Provided by Federal Civil Rights Law

The Draft Memorandum requires agencies to assess the discriminatory effects of AI, including disparate impact.[138] However, the Draft Memorandum generally limits those provisions to protected classes under federal civil rights laws.[139] The protections provided by federal law, however, represent an absolute minimum and may exclude characteristics worthy of protection. For example, federal protections do not extend to sources of income, education level, or history with the criminal legal system. Consequently, the final memorandum should require agencies to also assess impacts on historically marginalized groups not protected by federal law and mitigate those impacts if they are not relevant to the AI's intended use.

Conclusion

The ACLU applauds OMB's significant step forward in bolstering civil rights protections as the rapid advance of AI in government and throughout society continues. OMB's Draft Memorandum is significant, but more work is needed to build off this next step. To ensure that the final memorandum provides robust protections for civil rights and civil liberties, OMB should:

- To the extent possible, include national security systems and intelligence agencies within the Draft Memorandum's scope, clarify its application to less-advanced systems, and extend its applicability to federally funded programs;
- Maintain or expand the scope of AI use cases that are presumed to be rights-impacting;
- Strengthen key minimum practices, including by cabining agency discretion to waive or otherwise avoid the Draft Memorandum's requirements, bolstering notice provisions, and ensuring that community feedback mechanisms are equitable and approachable.

If you have any questions about these comments, please do not hesitate to contact Cody Venzke, Senior Policy Counsel, at cvenzke@aclu.org.

Sincerely,

American Civil Liberties Union

---

[138] Draft Memorandum at 18, 20.
[139] *Id.*