

Full Statement of Miranda Bogen
Director, AI Governance Lab at the Center for Democracy & Technology
For the Privacy and Civil Liberties Oversight Board Public Forum on the Role of Artificial
Intelligence (AI) in Counterterrorism and Related National Security Programs
July 11, 2024

Thank you to the Privacy and Civil Liberties Oversight Board (“PCLOB”) for the opportunity to provide these comments about the privacy and civil liberties issues associated with the use of artificial intelligence (AI) in counterterrorism and national security. My name is Miranda Bogen, and I am the Director of the AI Governance Lab at the Center for Democracy & Technology (“CDT”), a nonprofit, nonpartisan organization that defends civil rights, civil liberties, and democratic values in the digital age. For nearly three decades, CDT has worked to ensure that rapid technological advances promote our core values as a democratic society. Prior to taking this role, I worked with developers and deployers of advanced AI and machine learning models and systems at Meta, where I was directly involved in defining processes for managing risks presented by these technologies and building approaches and guidance to encourage the adoption of more responsible AI development practices.

Though technologies commonly associated with the concept of artificial intelligence have rapidly evolved in recent years, many of the challenges this type of technology presents and the questions it poses have been surfacing for some time. The newest AI-powered methods and tools may provide new capabilities to promote our national security. But we urge caution — especially when considering uses in high-stakes context such as national security — given the many well-known but unresolved risks that AI systems pose to people’s rights and safety.

Below, we address some themes responsive to questions PCLOB posed in their invitation to provide insight in the context of this panel: how AI will impact privacy and civil liberties, whether human operators will be able to understand and provide sufficient oversight to AI-powered systems, and how such oversight could be most effective. We also offer reflections on the role PCLOB or analogous independent agencies could play in providing oversight to the uses and implementation of AI so that national security and intelligence agencies are not left to “grade their own homework.”

The use of AI will continue to exacerbate privacy and civil liberties risk.

One way in which intelligence agencies may seek to use AI is to help analyze and act on huge swaths of text, audio, image, and video intelligence. We are deeply concerned, however, that without appropriate safeguards and oversight, this technology will be deployed to facilitate and dramatically expand indiscriminate surveillance and increase reliance on automated tools to inform national security activities, despite the many limitations of this technology and the risks to civil rights and civil liberties it poses.

A helpful mental model can be to divide AI applications into *predictive* AI systems that analyze or predict likely actions or outcomes, and *generative* AI systems which generate new content based on prompts. While this can be an imprecise distinction (for example, these categories can

converge when systems powered by generative AI models, such as chatbots, are leveraged to analyze large amounts of data or produce analysis or predictions), it can be useful in understanding relevant issues that AI developers, deployers, and oversight entities must attend to.

Both types of AI systems pose acute concerns related to privacy and civil liberties. For example, incomplete, unrepresentative or biased training data can lead to erroneous or discriminatory outcomes, which can both cause direct harm to people and divert attention and resources away from other areas of need. Imagine a predictive AI system that was designed to generate leads, but was built on selective or biased training data; such a system will cause investigators to waste time and resources chasing bad leads, leaving genuine security dangers unattended to. Or consider a facial recognition system trained and implemented so poorly that it frequently triggers false alarms, leading investigators to subject innocent people to undue law enforcement attention while failing to register the needed level of concern when a real threat appears amid the noise.¹ Use of AI without appropriate civil liberties and civil rights protections can also lead to the suppression of dissent, oppression of marginalized groups, and the supercharging of surveillance.²

The problems of bias in predictive AI systems are well-known, and unfortunately these issues persist — perhaps even more perniciously — in the context of generative AI.³ Large language models, for instance, have been shown to embed implicit biases⁴ with potentially dramatic effects: one study found that LLMs were more likely to suggest that speakers of African American Vernacular English be convicted of crimes and sentenced to death than speakers of standard American English.⁵ Concerningly, many such biases are unlikely to be detected using

¹ See, e.g. Lizzie Dearden, “Facial Recognition Wrongly Identifies Public as Potential Criminals 96% of Time, Figures Reveal”, The Independent, May 7, 2019, <https://perma.cc/YZ36-RC6A>.

² Paul Mozur, “In Hong Kong Protests, Faces Become Weapons”, N.Y. Times, July 26, 2019, <https://www.nytimes.com/2019/07/26/technology/hong-kong-protests-facial-recognition-surveillance.html>; Paul Mozur, “One Month, 500,000 Face Scans: How China Is Using A.I. to Profile a Minority”, N.Y. Times, Apr. 14, 2019, <https://www.nytimes.com/2019/04/14/technology/chinasurveillance-artificial-intelligence-racial-profiling.html>; Paul Mozur, “Inside China’s Dystopian Dreams: A.I., Shame and Lots of Cameras,” N.Y. Times, July 8, 2018, <https://perma.cc/27U7-S365>; Lena Masri, “Facial Recognition is Helping Putin Curb Dissent With the Aid of U.S. Tech”, Reuters, Mar. 28, 2023, <https://www.reuters.com/investigates/special-report/ukraine-crisis-russia-detentions>;

Khari Johnson, “Iran to Use Facial Recognition to Identify Women Without Hijabs”, Ars Technica, Jan. 11, 2023, <https://arstechnica.com/tech-policy/2023/01/iran-to-use-facialrecognition-to-identify-women-without-hijabs>.

³ Leonardo Nicoletti and Dina Bass, “Humans are Biased. Generative AI is Even Worse,” Bloomberg, June 9, 2023, <https://www.bloomberg.com/graphics/2023-generative-ai-bias>.

⁴ Xuechunzi Bai, Angelina Wang, Ilia Sucholutsky, and Thomas L. Griffiths, Measuring Implicit Bias in Explicitly Unbiased Large Language Models, *arXiv*, February 2024, <https://arxiv.org/abs/2402.04105>; Noel Ayoub, Karthik Balakrishnan, Marc S. Ayoub, Thomas F. Barrett, Abel P. David, and Stacey T. Gray, MD, *Inherent Bias in Large Language Models: A Random Sampling Analysis*, Mayo Clinic Proceedings: Digital Health, April 11, 2024, [https://www.mcpdigitalhealth.org/article/S2949-7612\(24\)00020-8/fulltext](https://www.mcpdigitalhealth.org/article/S2949-7612(24)00020-8/fulltext).

⁵ Valentin Hofmann, Pratyusha Ria Kalluri, Dan Jurafsky, and Sharese King, *Dialect prejudice predicts AI decisions about people’s character, employability, and criminality*, *arXiv*, March 2024, <https://arxiv.org/pdf/2403.00742>.

common performance and safety benchmarks. Using such biased AI systems in intelligence activities may therefore undermine national security more than protect it.

In many cases, AI outputs can be highly arbitrary, because the process of training machine learning and AI models unavoidably involves a significant amount of randomness.⁶ In predictive AI systems, research has shown that prediction errors exhibit so much variance for some members of a population that results are effectively random.⁷ Generative AI systems, meanwhile, are designed purposely to include random noise through a system parameter called “temperature” — this is what makes text and image generation systems appear to engage in more human-like or “creative” manner.⁸ This sort of uncertainty in the outcomes of AI systems suggests that AI unreliability is not a question of *if*, but *when*, so deployers or users of AI systems in high stakes contexts must proactively account for this uncertainty and plan for the erroneous outcomes that will result.

Privacy and civil liberties protections that govern the use of AI systems will be critical to ensuring systems are deployed responsibly and that people are treated fairly. PCLOB is the only independent government agency that is in a position to provide the necessary oversight with respect to national security systems. It has the necessary access to classified information and to technical expertise on both the Board and within the staff to the Board. While PCLOB may not be able to itself “look under the hood” of AI systems to determine whether they treat people fairly, were trained on properly collected data, or have reliable outputs, it should be able to ensure that the agencies employing AI have taken these steps and to interrogate the outcomes and resulting action of AI applications.

To illustrate, consider the role PCLOB could play in assessing the use of an AI-powered system to help determine who could be designated as a surveillance target, placed on watchlists, designated for counterterrorism and intelligence assessments, or made the subject of a preliminary investigation. In the context of targeting non-U.S. persons abroad for surveillance under Section 702 of the Foreign Intelligence Surveillance Act (FISA 702), AI might be employed to essentially “nominate” potential targets based on the characteristics of the hundreds of thousands of persons targeted for FISA 702 surveillance in the past, including their locations, communications patterns and even their statements made in public and in private settings. As PCLOB has emphasized, the standards for conducting this surveillance are quite low⁹ — there is no requirement that the target be a suspected agent of a foreign power or any

⁶ A. Feder Cooper, Jonathan Frankle, and Christopher De Sa, *Non-Determinism and the Lawlessness of Machine Learning Code*, Proceedings of the 2022 Symposium on Computer Science and Law (CSLAW '22), November 2022, <https://arxiv.org/pdf/2206.11834>.

⁷ A. Feder Cooper, Katherine Lee, Madiha Zahrah Choksi, Solon Barocas, Christopher De Sa, James Grimmelmann, Jon Kleinberg, Siddhartha Sen, and Baobao Zhang, *Arbitrariness and Social Prediction: The Confounding Role of Variance in Fair Classification*, The Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24).

⁸ See e.g. “Experiment with parameter values,” Google Cloud, accessed July 3, 2024, <https://cloud.google.com/vertex-ai/generative-ai/docs/learn/prompts/adjust-parameter-values>.

⁹ See Privacy and Civil Liberties Oversight Board, *Report On The Surveillance Program Operated Pursuant To Section 702 of the Foreign Intelligence Surveillance Act*, September 28, 2023,

suspicion that the target is involved in criminal or threatening activity — which could leave AI with a free hand in designating targets. Employing AI for intelligence surveillance targeting in an environment with such loose standards could significantly amplify intrusive surveillance that poses extreme risks to individual privacy.

Use of AI in designating individuals for assessments and to be the subjects of preliminary investigations is similarly fraught. Indeed, preliminary investigations may be initiated “on the basis of *any allegation or information* indicative of possible criminal or national security-threatening activity,”¹⁰ a standard so low it appears an AI recommendation could satisfy it entirely. Assessments and preliminary investigations themselves can result in invasive activities against individuals, as well as spur on more invasive surveillance and threatening law enforcement activities.

We recommend PCLOB review and publicly report on whether AI recommendations impact designation of individuals as surveillance targets, inclusion on watchlists, designation for assessments, designation as subjects of investigations, or similar actions. If such use is occurring or being contemplated, PCLOB should assess intelligence agencies’ treatment of these AI-powered systems in a number of areas: Are final decisions made by humans, and what level of deference is given to AI recommendations? Are the personnel interacting with and responding to AI systems specially trained, such as in the area of how to address risk of automation bias? What measures are taken to review the efficacy of AI systems in terms of training data they were built on, and the input data that their recommendations are based on? What testing measures were conducted to assess the effectiveness of systems, and what ongoing auditing processes exist to evaluate AI systems’ recommendations? And finally, is the risk to privacy and civil liberties of this AI so disproportionate to the potential benefit that the use ought to be foregone?

AI-powered systems remain unreliable and difficult to scrutinize, making oversight critically important.

In contexts like national security, where stakes are high both for public safety and for people’s rights, ensuring that sources of analysis are accurate and robust is critical. Simply put, the intelligence community should not assume that AI-augmented analysis is by default more accurate than human analysis, as it may often be *less* accurate.

While AI systems can appear to confidently produce predictions or analysis, they are at their foundation statistical tools that are inherently limited by the data they are trained and tested on and by the uncertainty inherent in predictive tasks. For instance, if an AI system is designed to

[https://documents.pcllob.gov/prod/Documents/OversightReport/054417e4-9d20-427a-9850-862a6f29ac42/2023%20PCLOB%20702%20Report%20\(002\).pdf](https://documents.pcllob.gov/prod/Documents/OversightReport/054417e4-9d20-427a-9850-862a6f29ac42/2023%20PCLOB%20702%20Report%20(002).pdf)

¹⁰ See, Department of Justice, *The Attorney General’s Guidelines For Domestic FBI Operations*, <https://www.justice.gov/archive/opa/docs/guidelines.pdf>; see also Michael German and Emily Hockett, “Standards for Opening an FBI Investigation So Low They Make the Statistic Meaningless,” Brennan Center for Justice, May 2, 2017, <https://www.brennancenter.org/our-work/analysis-opinion/standards-opening-fbi-investigation-so-low-they-make-statistic>.

make predictions based on an inherently subjective assessment, such as whether an individual has engaged in “suspicious activity,” the systems will reflect such subjectivity. Systems oriented toward analyzing or predicting mechanistic, or clearly defined, conditions will be less vulnerable to this pitfall — but may still perform poorly or unevenly across populations due to low quality or unrepresentative data. Additionally, even where substantial data is available, research has shown that prediction efforts are plagued by inherent limitations in that data, leading to errors that even increasingly advanced machine learning techniques cannot necessarily overcome.¹¹ Research into how AI systems can better indicate levels of uncertainty in their output is in progress, but remains nascent.

Meanwhile, generative AI systems continue to face the issue of hallucination, even when these systems are prompted to provide citations or chain-of-thought reasoning (in which a system is told to describe how it has arrived at its output in a step by step manner).¹² If an AI system can’t accurately reveal factors that informed its conclusions, human analysts won’t be able to assess whether or to what extent the system has provided information of unique value, or whether it is relying on unauthorized signals or spurious correlations to generate analysis or recommendations. The use of AI systems to triage and prioritize information for human analysis and review can, theoretically, reduce the likelihood that system errors and biases will lead to flawed actions or decision-making — but attention must be paid to whether human reviewers have lowered system thresholds in a manner that may be leading to increased error or whether they are overly deferring to the outputs of a systems without conducting reasonable oversight.¹³

To maintain some degree of confidence in the performance of an AI system, PCLOB should ensure agencies rely on training data to develop AI systems that was lawfully and ethically gathered and is relevant to the system’s intended uses; that they provide transparency into how systems were customized, fine-tuned, and validated for national security purposes to spot faulty assumptions or risks that such specification may have introduced; and that they maintain visibility into how these systems are integrated into operational work and how their outputs are acted on to ensure intended safeguards against errors and biases remain intact. Robust, independent, and ongoing evaluation of systems’ overall accuracy, robustness, and other characteristics can gauge the extent to which these errors are likely to occur in a given system (and spot whether performance is degrading over time). At the same time, technical evaluations

¹¹ Certain features are simply unmeasurable, certain relevant information goes unmeasured and is therefore absent from datasets, and certain information is imperfectly measured. Ian Lundberg, Rachel Brown-Weinstock, Susan Clampet-Lundquist, Sarah Pachman, Timothy J. Nelson, Vicki Yang, Kathryn Edin, and Matthew J. Salganik, *The origins of unpredictability in life outcome prediction tasks*, PNAS, June 4, 2024, <https://www.pnas.org/doi/full/10.1073/pnas.2322973121>.

¹² See e.g. Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli, *Hallucination is Inevitable: An Innate Limitation of Large Language Models*, arXiv, January 2024, <https://arxiv.org/pdf/2401.11817>; Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman, *Language Models Don’t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting*, 37th Conference on Neural Information Processing Systems (NeurIPS 2023), December 2023, <https://arxiv.org/abs/2305.04388>.

¹³ For example, a recent investigation found that when the Israeli military deployed an AI-powered target selection system, the threshold for authorizing military strikes was lowered and that human investigation of system recommendations were deprioritized in favor of generating broader target lists. Yuval Abraham, “Lavender’: The AI machine directing Israel’s bombing spree in Gaza,” +972 Magazine, April 3, 2024, <https://www.972mag.com/lavender-ai-israeli-army-gaza>.

of a system can make no guarantees about the veracity or reliability of specific outputs, so agencies should make efforts to ensure that automatically generated analysis or recommendations remain subject to meaningful human oversight. And finally, technical oversight related to the performance of a system is irrelevant if the purpose of the system is inappropriate to begin with, such as unauthorized mass surveillance systems or pseudoscientific emotion recognition systems.¹⁴

Human decision makers with subject matter and domain expertise can and should maintain meaningful oversight over the use of AI systems, but this will require proactive effort.

Some have hypothesized that increasingly advanced AI systems will prove challenging to understand or control, but complex organizations of all kinds, including government agencies, exhibit similar characteristics — making oversight all the more important.¹⁵ National security institutions must put in place both internal as well as independent governance mechanisms to promote responsible use of AI, such as clearly assigning decision-making and internal oversight responsibilities, requiring review and approval by high-level officials for procurement of systems and use cases that present particularly high risks, and ensuring legal, civil rights, and privacy officials have comprehensive visibility into how departments and agencies are using AI and are included as part of the decision making process through the AI development, procurement, and deployment lifecycle.

In addition to these organizational and procedural interventions, human oversight can be embedded in AI system design, monitoring, and oversight over AI-recommended actions. In the context of a system's design, interdisciplinary teams and internal and external governance professionals can review a system's key parameters and training data, and efforts can be made to mitigate the risk that people tasked with labeling data (either for model training in supervised learning contexts or for safety efforts such as reinforcement learning from human feedback¹⁶) will inject biases, subjective opinions, or avoidable errors into the AI systems informed by that data.¹⁷ Human users can and should be trained in how to use and interpret the output of AI systems, as well as be given the ability to flag suspicious or apparently anomalous or biased outputs so that these signals can be fed back to system developers. Efforts should be made to prevent automation bias, or the tendency for people to naturally assume automated systems are correct even in the face of conflicting evidence. National security authorities should also ensure

¹⁴ Article 19, *Emotional Entanglement: China's emotion recognition market and its implications for human rights*, January 2021, <https://www.article19.org/emotion-recognition-technology-report>.

¹⁵ Jay Stankey, "How Can Smart, Ethical Individuals Form Dumb, Amoral Government Agencies?" ACLU, September 6, 2013, <https://www.aclu.org/news/national-security/how-can-smart-ethical-individuals-form-dumb-amoral-government>.

¹⁶ Nathan Lambert, Louis Castricato, Leandro von Werra, and Alex Havrilla, "Illustrating Reinforcement Learning from Human Feedback (RLHF)," Hugging Face, December 9, 2022, <https://huggingface.co/blog/rlhf>.

¹⁷ Stephen Casper and Xander Davies et al, *Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback*, arXiv, July 2023, <https://arxiv.org/pdf/2307.15217>.

that humans are not placed in the position of being a “moral crumple zone,” where they are given nominal oversight over automated systems and thus held accountable for the systems’ behavior, but offered insufficient training or opportunity to truly exercise that oversight.¹⁸ And in cases where AI systems are used to recommend or take specific actions that may affect individuals’ rights (e.g., targeting a person for surveillance), human approval should be required before such actions are triggered.

To translate this concept into the intelligence surveillance targeting arena, consider that intelligence analysts are sometimes given a drop down menu to prompt them to justify their targeting decisions. That menu should require the identification of actual evidence that statutory or agency-driven criteria were met. “AI recommendation” shouldn’t be on the menu at all — or if present, should never be enough to justify the targeting decision in the absence of other evidence that the analyst is prompted to identify. This is how human approval, with evidence-based justification, might be required in the intelligence surveillance targeting context.

PCLOB could play an important role in ensuring that human reviewers are in place to assess the results of AI use, that these reviewers are not inappropriately deferring to the outputs of the AI system instead of conducting proper oversight, and that they have the necessary transparency into the operation of an AI system, and the training to spot things like faulty assumptions on which an application of AI is based.

While the recommendations above largely focus on ensuring AI is as accurate and efficacious as possible, it is also essential that AI technologies are used in a manner that upholds constitutional values, civil rights, and civil liberties. Agencies should conduct impact assessments to determine whether an AI system risks being biased or otherwise violating constitutional and human rights. Certain AI technologies or uses should be prohibited because they pose an unacceptable risk to rights (e.g., AI profiling or risk scoring systems that attempt to predict an individual’s future criminality or foreign intelligence-related activities).

The PCLOB already plays an important role in ensuring that the use of intelligence and national security authorities to protect against terrorism are used consistently with the right to privacy and civil liberties. National security use of AI will open up more opportunities and responsibilities for oversight that may tax PCLOB’s limited resources and could distract it from other necessary oversight activities. For these reasons, we suggest that you focus your AI-related oversight to specific, publicly-announced “deep dives,” that you invite Congress to provide additional resources to enable AI oversight activities, and that you consider supporting the creation of a different body, modeled after the PCLOB, to focus specifically on AI.

PCLOB was established to protect privacy and civil liberties in the fight against terrorism. When fighting terrorism is one of the purposes of a national security program, PCLOB has properly asserted a role in overseeing such a program. FISA 702 surveillance is a good example of a program with such mixed purposes. However independent oversight is also critical for

¹⁸ Madeleine Clare Elish, *Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction*, *Engaging Science, Technology, and Society* 5 (2019), <https://estsjournal.org/index.php/ests/article/view/260>.

AI-powered national security activities wholly unrelated to terrorism, or that have only a marginal or incidental anti-terrorism purpose. The use of AI on the battlefield, to support covert activity, to screen intelligence analysts and many other national security purposes may be out-of-scope for PCLOB, but still require independent oversight. That is why CDT is calling for the creation of a “PCLOB for AI.” Should Congress opt to create such a body, careful lines between the responsibilities of PCLOB and of the new independent body to oversee use of AI in the national security arena would need to be drawn.

AI should not be permitted to circumvent rules and safeguards established for intelligence agencies and personnel.

In addition to risks posed in terms of how AI systems operate and the type of recommendations they provide, PCLOB should also recognize and account for how AI might be used to circumvent rules and safeguards built upon the assumption of human activity. For example, in last year’s PCLOB report on FISA 702, the Board recommended requiring that personnel obtain court approval before reviewing the results of a US person query of 702-collected communications.¹⁹ If such a requirement became law in the future, personnel might seek to use AI to circumvent it: Rather than seeking court approval to personally read communications returned from a US person query, FBI agents might simply task an AI system with reviewing the 702-collected communications of certain US persons, provide assessments of contact with foreign targets, and take the position that because there was no human review of the information returned from the query, court approval would not be required. The statutory rule requiring court approval will have been circumvented. Instead, the AI system would offer recommendations for which individuals should be subject to monitoring or investigation.

Potential for AI to circumvent constitutional rules is also a danger given the intelligence community’s stance that communications content has not been searched if it is only subject to automated scanning and not available for human review, a position fundamental to FISA 702’s Upstream collection system.²⁰ Based on this narrow interpretation, agencies might press for AI to scan communications and provide recommendations based on content, all while claiming that the lack of human review or collection into government databases means a Fourth Amendment search has not occurred.

PCLOB should examine and report on whether AI is being used in this or any other manner to circumvent rules and limits established for intelligence agencies and their personnel.

¹⁹ See Privacy and Civil Liberties Oversight Board, *Report On The Surveillance Program Operated Pursuant To Section 702 of the Foreign Intelligence Surveillance Act*, September 28, 2023, [https://documents.pcllob.gov/prod/Documents/OversightReport/054417e4-9d20-427a-9850-862a6f29ac42/2023%20PCLOB%20702%20Report%20\(002\).pdf](https://documents.pcllob.gov/prod/Documents/OversightReport/054417e4-9d20-427a-9850-862a6f29ac42/2023%20PCLOB%20702%20Report%20(002).pdf)

²⁰ Robert S. Litt, *The Fourth Amendment in the Information Age*, Yale Law Journal, April 27, 2016, <https://www.yalelawjournal.org/forum/fourth-amendment-information-age>

PCLOB should assess compliance with and sufficiency of existing executive policies on agencies' use of AI.

As an independent oversight agency with access to classified programs, PCLOB is uniquely poised to assess the effectiveness of administration policy on agencies' use of AI.

Executive-wide policy on AI is primarily based on two items: the Office of Management and Budget AI governance memorandum²¹ and the forthcoming memorandum on national security uses of AI, set to be published later this month. PCLOB should review agency responses to these items in a number of respects. First, agencies that oversee counterterrorism operations such as the FBI and DHS could fall into the regulatory rubric of the OMB memorandum or under the authority of the national security memo. Dual purpose agencies should seek to follow the OMB rules by default, and limit application of the national security memo to specific programs centered on national security. PCLOB is well positioned to vet and report on whether agencies are following this approach, or if national security mandates as a component of agencies is being used as a pretext to avoid applying the requirements of OMB memorandum more broadly. Second, PCLOB should examine how effectively both the OMB and national security memorandum are spurring the adoption of rules that effectively safeguard privacy and civil liberties. If the current structure contains gaps or fails to anticipate risks, PCLOB should directly recommend improved agency practices, as well as publicly report on the need to solicit input from outside stakeholders and experts.

As intelligence and national security agencies deepen their pursuit of and investment in new technologies like artificial intelligence, the careful consideration of the privacy and civil liberties implications of AI-powered systems is both necessary and urgent. Many issues presented by more advanced AI systems resemble risks that are well-known but remain unaddressed, and many interventions that could help to address these longstanding impacts will lay a strong foundation for more responsible deployment and use of systems of increasing complexity. Independent oversight and expertise will play a critical role in ensuring that decisions around the appropriate use of AI-powered tools remain grounded in human rights and core democratic values.

²¹ Office of Management and Budget, M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (Mar. 28, 2024), <https://perma.cc/RKK7-SMYJ>. See also, Center for Democracy & Technology, "CDT Welcomes Final OMB Guidance on Federal Agencies' Use of AI, and Now Looks Toward Earnest Implementation", March 28, 2024, <https://perma.cc/MQ34-VUWG>