UNIDOSUS
STRONGER COMMUNITIES. STRONGER AMERICA.

WRITTEN TESTIMONY ON CIVIL RIGHTS IMPLICATIONS OF FEDERAL USE OF FACIAL RECOGNITION TECHONOLOGY

Presented at

"Civil Rights Implications of the Federal Use of Facial Recognition Technology: Guidance for Meaningful Federal Oversight"

Submitted to

The U.S. Commission on Civil Rights

Submitted by

Laura MacCleery, Senior Policy Director
Claudia Ruiz, Senior Civil Rights Analyst
Policy & Advocacy
UnidosUS
Raul Yzaguirre Building
1126 16th Street NW, Suite 600
Washington, DC 20036-4845

March 8, 2024

On behalf of UnidosUS, we respectfully submit this testimony on the pressing issue of achieving appropriate governance for the federal use of facial recognition technology (FRT). UnidosUS is a nonprofit, nonpartisan organization that serves as the nation's largest Hispanic civil rights and advocacy organization. Since 1968, we have challenged the social, economic, and political barriers that affect Latinos through our unique combination of expert research, advocacy, programs, and an Affiliate Network[*] of more than 300 community-based organizations across the United States and Puerto Rico.

There is a clear and urgent need for updated regulations to address ongoing infringements by uses of these technologies to constitutional principles such as due process, equal protection, and privacy. How governments set standards for technology acquisition by the federal government, including by law enforcement and immigration, could be a substantial lever to drive more responsible and democratic processes and design in areas where it may matter most for our society.

Unfortunately, such decisions have historically been characterized by a lack of principled approaches, transparency, or accountability. Because communities of color and immigrants are the first to be targeted with powerful new forms of surveillance and are the last to benefit from technological changes, doing better by our communities in these critical areas is one of the most important civil rights issues of our time.

The question of how to balance state power with the preservation of appropriate zones for the privacy of individuals and groups within a democracy is not new. They go back to our founding and our Constitution. As our Founders knew, any healthy democracy must have effective ways to address threats to law and order and the rule of law at the same time that it preserves space for non-violent protest, the exercise of free speech and the ability to travel, and protects a zone of privacy around the individual that is both intellectual and physical. While the specific boundaries around longstanding doctrines in these areas may shift over time, grappling fully and sensitively with these questions remains a central task.

We simply have not done this work to integrate powerful new technologies that allow cheap and routine surveillance of movements and biometrics, and that are aided by data collection on every aspect and moment of our lives. Such information continues to be collected largely without notice to—or consent from—individuals, and without the ability to request deletion or even know what data is known about them.

[*] UnidosUS Affiliate Network, https://www.unidosus.org/about/affiliates/.

The absence of fundamental privacy and rights-preserving norms and laws makes FRT, and other forms of biometric surveillance, a Wild West. Our significant and specific concerns about FRT use by law and immigration enforcement agencies on Latinos, immigrants, and communities of color are grounded in three observations:

- First, current uses of FRT undermine democratic norms and principles and threaten immigrant communities and communities of color. We cannot allow an infrastructure of invasive surveillance and unchecked data-sharing to undermine cherished constitutional freedoms.
- Second, rather than providing exemptions or waivers for law and immigration enforcement uses, as suggested by the Office of Management and Budget 's (OMB) draft AI Memorandum (in its "Request for Comments on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum," hereinafter called the "OMB Memo"), we need enforceable and binding standards for all surveillance technologies and AI models used to scale them.
- Third, our failure to align these systems with any appropriate governance framework perpetuates communities' exposure to unaccountable and opaque uses of technologies, including FRT and other forms of biometric surveillance. We must formalize mechanisms that elevate the voices of impacted communities in setting policy and set baselines for privacy with better laws that drive improved systems and designs.

Many of my colleagues from the civil rights community rightly raise the issue of inaccuracies in the data that specifically and disproportionately impact communities of color. We share their deep concerns.

We saw them in action in the context of UnidosUS's work in Puerto Rico on the expanded Child Tax Credit, when the government's identification systems routinely failed to recognize darker-skinned images of tax filers. As Dr. Joy Buolamwini's excellent book, *Unmasking AI*, explains, addressing these sources of bias will not be simple. Machine learning models are also highly prone to biased inferences generally, and discriminatory outcomes are pervasive. Nevertheless, these technologies would pose risks even if we could comprehensively address such profound issues of bias and inaccuracy. We focus here primarily on risks that infect these uses across other dimensions and pertain to biometric surveillance.

I. Creating a system compatible with democratic norms Is essential.

Wherever they work and live, and whomever they are, everyone deserves access to basic democratic rights, including the right to privacy, the right to travel, and the right to due process of law. Yet Latinos and immigrants, like other historically marginalized communities, have endured a legacy of surveillance by the U.S. government, resulting in disproportionate racial profiling, targeting, and tracking by law and immigration enforcement.

We must not shy away from hard cases balancing trade-offs between effective law enforcement and maintaining public safety with protecting and preserving individual rights and liberties. Democracy's lifeblood is the accommodation of diversity and dissent, enabled by associational freedom and zones of autonomy for individuals to develop away from state interference or coercion. Mass surveillance enabled by unfettered automated technologies can upend the foundations on which democratic self-rule relies. We must identify uses by both governments and private actors that clearly are steps towards the abusive or anti-democratic uses of surveillance, and bar them outright, such as social credit scoring and behavioral or emotional monitoring.

While high-tech surveillance capabilities are frequently touted as serving public safety, their use also presents a clear source of systematic civil rights and liberties violations. In other countries, we have seen that such technologies offer chilling possibilities for oppressing freedoms by authoritarian regimes. These clear and imminent dangers demand oversight by the government to balance effective law enforcement with constitutional norms inherent to a free society and to prevent abuse in specific cases.

We already are living under networks of rapidly proliferating and intrusive surveillance systems. Data purchased by governments include personal and consumer information from utility companies and third-party data brokers, as well as private contracts for tools with facial and biometric recognition capacities. These dragnet forms of surveillance by the Department of Homeland Security (DHS) and law enforcement infringe on the rights of millions, including millions of U.S. citizens. Consider that:

- A September 2023 GAO report found that law enforcement at DHS and the DOJ lacked basic protocols or training around the use of facial recognition technologies. In response, DHS Sec. Mayorkas published a memo articulating a policy commitment to constitutional principles.

- Immigration and Customs Enforcement (ICE) has used FRT to search the driver's license photographs of around 1 in 3 (32%) of all adults in the U.S. The agency has access to the driver's license data of 3 in 4 (74%) adults and tracks the movements of cars in cities home to nearly 3 in 4 (70%) adults. When 3 in 4 (74%) adults in the U.S. connected the gas, electricity, phone, or internet in a new home, ICE was able to automatically learn their new address.

- DHS has expanded the use of facial recognition technology on travelers, including U.S. citizens, at airports and land borders without obtaining consent.

- A report by the Department of Homeland Security Office of the Inspector General (DHS OIG) revealed that U.S. Customs and Border Protection (CBP), ICE, and the Secret Service purchased and used commercial geolocation data in violation of their privacy policies and that DHS components have failed to develop policies governing the purchase and use of location data. According to DHS OIG these failures "occurred because the components did

not have sufficient internal controls to ensure compliance its own privacy policies." The report recommended that CBP and ICE discontinue the use of commercial geolocation until they have developed and implemented sufficient policies, including conducting a privacy impact assessment. CBP promised Sen. Ron Wyden (D.-OR) to stop purchasing location data by the end of Sept 2023.

- In July 2022, the American Civil Liberties Union (ACLU) published thousands of pages of previously unreleased records about how CBP, ICE, and other parts of the DHS are buying access to and using vast volumes of cellphone location information extracted from smartphone apps.

- According to the Electronic Frontier Foundation, CBP has installed about 300 different types of surveillance towers from the California coast to the tip of Texas. CBP's 2023 and 2024 budgets also cover the deployment of approximately 174 additional towers along the U.S.–Mexico border. Studies by Sam Chambers of the University of Arizona document how the presence of the towers cause migrants to take more dangerous and remote pathways.

As this makes clear, use by governments of AI tools even in cases involving core matters of civil liberties, are a cause for concern for vulnerable populations and impose on the privacy rights of both immigrants and U.S. citizens. This does not inspire confidence that the right guardrails are in place to use FRT or AI in ways that are consistent with democratic principles, fairness, due process, and other important constitutional values.

The lack of specific requirements is deeply problematic. For the 62.1 million Latinos living in this country, the risks of overreach from intrusive surveillance are pervasive. In addition to the nearly 20 million immigrants who identify as Latino in this country and more than 10.6 million U.S. citizens of any racial or ethnic identification who live in mixed-status households, they also face unique risks to the infringement of basic rights from oversurveillance.

Perceived efficiencies from current and planned uses in criminal justice, immigration enforcement and related uses will likely lead agencies to continue to gloss over deeply concerning data security, stewardship, privacy, and civil liberties concerns. We have also seen fearmongering used to justify billions of dollars of investment in biometric surveillance infrastructure at the border and automated license plate reading technologies across the nation, along with other forms of routinized mass surveillance.

It is important for the Administration to act decisively to address these risks. The consequences of now-routine forms of data collection on communities and individuals could, if left unchecked, provide tools for overreach on immigration that is unmistakably authoritarian. In particular, immigrant and mixed-status communities are canaries in the coal mines on civil liberties because their lives are used as test cases for policies that roll back protected rights and liberties. Tellingly, these same communities have historically been left behind and left out of both technological advances and governance of new technologies.

The combination of these factors produces surveillance infrastructure that creates digital suspect classes, places entire communities under heightened scrutiny, and alters the capacities and concentration of law enforcement resources at the community level. These biased forms of over-policing are unrelated to actual risk or probable cause, and thus contrary to basic principles of law.

II. We need enforceable and binding standards for all surveillance technologies and AI models.

The choice often posed between rights-protecting model design and individual privacy, on the one hand, and effective law and immigration enforcement, on the other, is a false one, as we outlined in our comments on the OMB's AI Memo. Permitting agencies to duck accountability and oversight of these tools through waivers, as the OMB Memo proposes, would erode any incentive to do the hard work of aligning the design of systems with rights.

Such alignment is exactly the work before us. Once appropriate incentives and protections are in place, sound design and intentional guardrails and limitations can make both a reality. The failure to use privacy by design and other rights-protecting principles and mechanisms should not be characterized as a function of the technology when it is, instead, a human choice to sanction unaccountable, untransparent and dangerous practices.

We owe communities who will be impacted first and, possibly, worst a system that embeds core personal rights and democratic freedoms, assures transparency, protects our personal autonomy, and creates a baseline for fair rules of the road in tech design. Doing so is essential to stop abuse of the tools, but it will also provide incentives for the right kinds of innovation so that business models support, rather than undermine, human dignity.

The task of the OMB Memo for the agencies is to establish "proper controls" over government uses of AI for current and near-future models and uses. We believe the Memo is a solid start, but its approach is incomplete or lacks important clarity in a number of areas that could benefit from substantially more operational structure for agencies, and that OMB should more fully leverage the work of NIST.

For example, the agencies' assignment under the Memo to achieve "maturity" for AI systems begs the question of how—and who—defines that success and on what grounds. Agencies will need constructive guidance on common technical issues arising from current uses and mitigations for AI systems, as well as to be informed about helpful developments and technical and sociotechnical challenges that arise in particular contexts and use cases.

For FRT, as for other forms of powerful technologies at scale, including the AI models used to make them more widely available, basic benchmarks and governance systems are sorely needed. In the absence of clear regulations for deploying these technologies, current practices related to government use of FRT fail on each of the government's own benchmarks outlined in

the National Institute for Standards and Technology's (NIST) AI Risk Management Framework (AI RMF).

Specifically, the AI RMF provides a set of characteristics to assess the trustworthiness of AI models, including FRT systems, which highlight that AI systems should meet baselines for each of the following factors: 1) Valid and Reliable; 2) Safe; 3) Secure and Resilient; 4) Accountable and Transparent; 5) Explainable and Interpretable; 6) Privacy-Enhanced; and 7) Fair—with Harmful Bias Managed.

Agencies—including law enforcement and immigration agencies—should evaluate current uses in light of each of these values across their entire portfolio of AI uses, in consultation with NIST and other experts familiar with the evolving science for each of these measures, paying concentrated attention, as the OMB Memo indicates, to risks, and to safety- and rights-impacting uses. In mapping current uses of AI and algorithmic tools, agencies should also:

- Detail and explain the technological limitations of a tool given its use cases and relevant human factors,
- Identify the adequacy of any current evaluations of the training data, model design, and impacts, and any mitigations for known and potential risks,
- Describe the extent of involvement or consultation with impacted communities (more on this below) on design, risks, impacts, or other aspects of the model or system,
- Explain the adequacy and conclusions of external audits and impact assessments that are underway or have been done, and
- Fully characterize the socio-technical context at the agency related to human interactions with the technology, evidence on experiences of internal and external users, and other factors.

Additional clarity and detailed instruction on how to align success for the "maturity" of AI systems with the requirements and standards from the NIST RMF, alongside new requirements for consultation with impacted communities, as we outline below, would provide a much more developed set of parameters for "success" and could help to generate much more transparent and aligned processes across the government, including for FRT systems, which are deployed in highly sensitive use cases.

Technology also brings a rising risk of so-called "automation bias"–as the OMB Memorandum calls our propensity to place undue faith in outputs generated by automated tools (sometimes even in the face of contradictory information derived from empirical evidence). Such concerns are heightened when power imbalances are pervasive and there is every incentive, given the exigencies of urgent situational judgments—such as those that occur in policing and at our borders—to disregard a tool's limitations.

Although the NIST RMF framework calls for AI to be "privacy-enhancing," OMB's approach fails to ensure that this will matter where it is needed most. Instead, the Memo's proposed waivers are likely to allow some of the most problematic and rights-infringing deployments to continue to avoid even basic forms of public accountability. For example, as a law enforcement agency combining criminal and civil responsibilities, DHS or its sub-agencies may claim that law enforcement and national security exemptions apply or that an activity is "mission critical."

Given the need to create government policies that protect us from the specter of abuse, OMB, with the advice of this Commission, should instead lead a process of taking full account of current practices and fix them in short order. At a minimum, the OMB should create additional clarity regarding when agencies can seek waivers or exemptions from having to meet risk management requirements. We agree with the Leadership Conference, which stated in their comments to the OMB Memo that the following is needed:

- When a waiver or exception is granted, there should be a mechanism to seek reconsideration of such a decision.
- The Memo should be clear that waivers and exceptions sunset annually and should be reevaluated in light of these documented harms and risks.
- The Memo should require that agencies consider less rights-impacting alternatives before they are eligible for consideration for a waiver or exception.
- The Memo should require that agencies publicly report seeking waivers or exemptions, and the grounds for this request and its resolution and timing be reported.

Law and immigration enforcement uses in fact require more—not fewer—procedural safeguards. In lieu of providing waivers, we should require privacy by design, emphasizing those that are compatible with effectiveness, such as data minimization, access controls, federated learning, and other privacy-enhancing techniques for government AI uses, as described in our recommendations, below. Collection by agencies of biometric data, including faceprints, DNA, and other biometrics, should also receive substantially heightened scrutiny—and attendant safeguards—given its power, scale, and indelibility.

In keeping with the above, we fully support the OMB Memo's provisions on procurement policies that underscore that AI contracts should align with national values and law, including "those addressing privacy, confidentiality, copyright, human and civil rights, and civil liberties." Since waivers for law enforcement or mission-critical functions could undermine progress in assuring that federal tax dollars are not spent on systems incompatible with this requirement, consistency across federal procurement policy provides another reason to substantially narrow or eliminate them.

IV. We must envision a regulatory ecosystem with transparency and accountability for impacted communities.

Our failure to align these systems with any appropriate governance framework means that communities are subject to unaccountable and opaque tools and models. To ensure that they are based on more than the specific incentives of the enforcement agencies, it is essential to create diverse and inclusive advisory bodies to help craft and review policies, channel and facilitate community input, and evaluate the evidence on uses and outcomes. This means formalizing governance mechanisms that elevate and incorporate the voices of impacted communities in setting policy.

As described in our Written Testimony on Governance of Artificial Intelligence and Comments on OMB Draft AI Memorandum, a responsible, accountable, and transparent approach to AI governance that includes privacy-enhancing techniques and use limitations complemented by community-informed governance, rigorous oversight, and public transparency can safeguard against anti-democratic misuse of these technologies. In particular, we call on the federal government, including the DOJ and the DHS, to give impacted communities a voice in the governance process through practical mechanisms that provide a means of feedback for agencies about the uses and impacts of technologies in real time.

Democracies learn in public, using a deliberative process that assesses harms and trade-offs, looks at technical capacities and implications for shared values, and allows stakeholders to weigh in. AI governance, including for FRT and biometric tools, to be democratic in nature, should anticipate potential harms and include mechanisms for accountability to the people they impact. Too often, the bias or flaws in models are understood too late—so we must get better at both predicting and preventing foreseeable harms through good design: Impacted groups are ideally positioned to tell technologists what they may not know.

In the current work being done even at NIST, there is an emphasis on metrics that were created without civil rights scrutiny or the involvement of impacted communities. These fail to account for the socio-contextual dynamics, and real-world application of constitutional norms like equal protection or due process in light of limited metrics and limitations in tools. Because targeted communities cannot meaningfully redress surveillance harms from outside of the decision-making and standard-setting process, it is crucial to create and require means for public accountability and community input. Our traditional forms of notice-and-comment process, which takes years, will be insufficient to keep pace with the rapid change in capacities and uses of AI and surveillance tools.

Among other reasons, this is why fairness or bias "audits," which many institutions generally measure against statistical outputs, can and do still result in inequitable, discriminatory real-world outcomes. Notions of fairness and equity do not exist in a vacuum. What we define as "fair" or "equitable" must also include a qualitative and comprehensive assessment of outcomes, including collection of the evidence on how these systems affect vulnerable groups.

True equity requires their experiences, needs, and perspectives to shape governance frameworks and decisions around acceptable applications for these systems and the delicate and situational judgments and tradeoffs they entail. Only by balancing technical audits with impacted communities' lived experiences, developed through inclusive impact assessments, can societal effects be understood and addressed.

We call for a multifaceted and comprehensive governance model that includes inclusive red teaming, impact assessments, and consumer complaint collection, alongside a public leaderboard for metrics and a requirement for community advisory committees for each agency, sub-agency, or department, as depicted below.

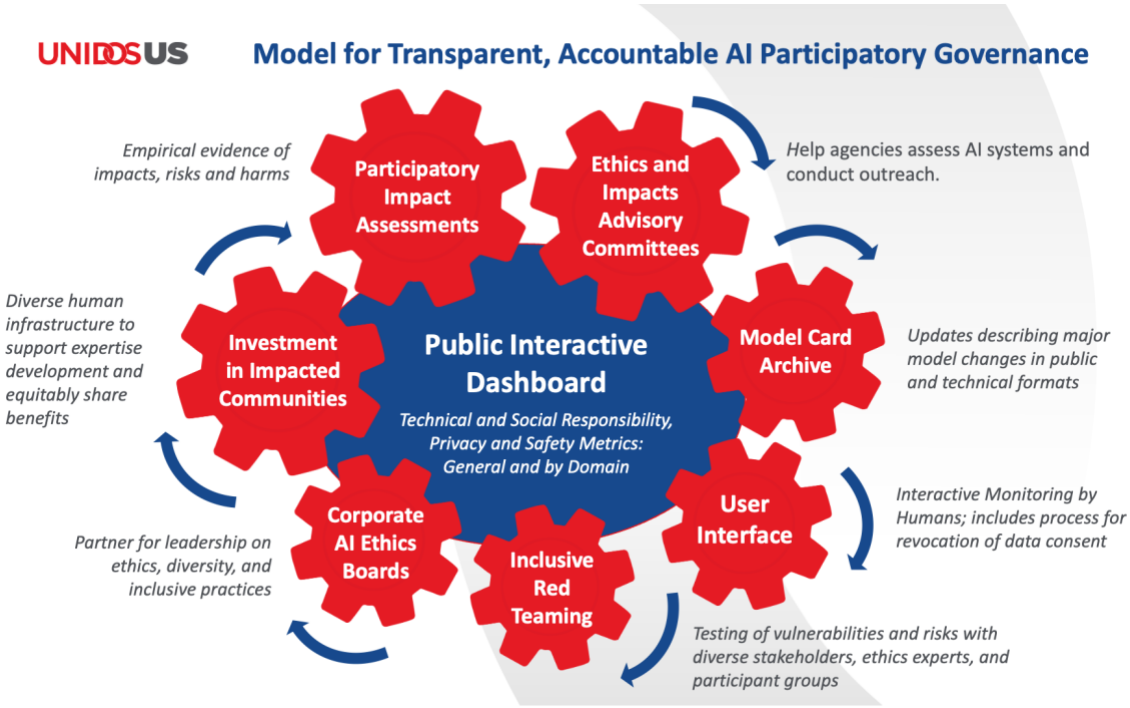

V. Additional recommendations that are essential to privacy, democracy, and civil rights.

Generations of disenfranchisement and exclusion have rendered civil rights interests and marginalized groups systemically unlikely to be able to drive major changes in the absence of social upheaval, such as we saw with the murder of George Floyd. The most critical step to secure the rights of impacted communities is to secure a new baseline of privacy rights for all, embedded in the technology. Because data-driven technologies have come to underpin nearly every facet of our society, we must secure our fragile democracy with Constitutional safeguards that effectuate better data stewardship and respect for personal freedoms and secures fairness in law and immigration enforcement.

To ensure that privacy approaches in the U.S. keep pace with technological design advances that can protect personal freedoms, we should regard the following national safeguards as a floor for FRT, AI, and related technologies. Uses and models should:

- Require consent and understandable disclosures about personal data collection, transfer, retention, and use.
- Guarantee individuals' right to access, correct, delete, and port personal data held by companies.
- Limit data collection, sharing, and retention to what is reasonably necessary for providing services.
- Mandate security safeguards proportional to data sensitivity, like encryption and access controls.
- Provide heightened protections for sensitive data like financial, health, or location information.
- Enable opt-outs and accessible consumer controls of data sales, targeted advertising, and granular profiling, including extending permission to authorized third parties to opt out on their behalf.
- Bar unlawful discrimination from algorithmic assessments.
- Establish accountability through external audits and civil fines for violations and empower state Attorneys General and individuals with the right to pursue legal damages for violations.
- Require bias testing, model documentation, and human checks on consequential algorithmic determinations.
- Mandate consideration and uses where appropriate of specific privacy-enhancing technologies and techniques that can be implemented for lawful government AI uses that enhance security and do not decrease effectiveness, including but not limited to, the following tools:
    - Branching and segmentation—branching divides data by sensitivity, while segmentation divides datasets by purpose, access needs or other criteria.
    - Differential privacy—adds controlled statistical noise to datasets to mask individual identities while enabling useful aggregated insights.
    - Federated learning—trains AI models without centralizing sensitive data.
    - Homomorphic encryption—allows computing on encrypted data without decrypting it first.

- Secure multiparty computation—enables multiple entities to jointly analyze data while keeping their separate inputs concealed.

- Data anonymization—this includes removing or obscuring personal identifiers like names and social security numbers from records used to train AI systems.

- Data minimization—limiting collected data strictly to the minimum necessary for a given authorized purpose.

- Encryption—encrypting stored and transmitted data to secure it against unauthorized access.

- Access controls—restricting and auditing data and system access to essential personnel.

Consistent with due process and the concerns articulated above, we should also address specific measures to prohibit autonomous AI determinations of any immigration or criminal justice outcomes, such as consequential predictive decisions about criminal justice and immigration outcomes like risk assessments, bail, parole, and sentencing terms. As the National Academies of Sciences report on FRT indicates, we should require human validation and a host of procedural and transparency protections for any FRT- or AI-informed recommendations, resisting the temptation to make automation bias the default in systems pertaining to human life and liberty. More specifically, we should:

- Prohibit AI tools from determining criminal risk scores and sentencing guidance absent human validation and require statistical bias testing and public reporting of aggregated outcomes.

- Include clear and public descriptions of how AI is used as well as assurances of anti-bias processes such as: diverse training data, bias testing with impacted groups, expert oversight committees, regular algorithmic auditing and impact assessments by independent third parties, use of AI models as advisories rather than sole determinants, outcomes monitoring, and public reporting.

- Laws should establish clear accountability for harms with transparency, explanations of AI-informed decisions to individuals, and appeals processes to contest discriminatory, unfair, or incorrect results affecting their case outcomes.

- Provide public reporting on aggregate metrics like accuracy, fairness indicators across disaggregated demographics, and compliance with anti-bias standards to facilitate oversight.

- Conduct civil rights impact assessments with community representation prior to developing or deploying such AI tools. Incorporate diverse perspectives into design requirements and require explanations for any that are not incorporated.

- Institute ongoing review by oversight boards or independent bodies to monitor data practices and ensure accountability for justified needs. These should be multi-stakeholder committees with decisional power that evaluate whether continued use of such AI tools remains appropriate or should be terminated based on the public interest.
- Other basic requirements for law enforcement uses consistent with Constitutional rights include the need to:
    - Require warrants—compel investigators to demonstrate probable cause and receive judicial approval to gather data on specific targets.
    - Mandate specificity—warrants and data requests must precisely define information sought related to a case or investigation.
    - Minimize data collection—collect only the data required to serve the investigative purpose.
    - Set time limits—data warrants should expire after a defined reasonable period.
    - Ensure safety with encryption—enable secure transmission and storage of information to prevent unauthorized access to data.
    - Require training—law enforcement and immigration personnel should be trained on limitations, civil liberties considerations, and proportionality.

We deeply appreciate your interest in this topic and stand ready to assist the Commission. For additional information, please contact Laura MacCleery, Senior Director of Policy, at lmaccleery@unidosus.org, and Claudia Ruiz, Senior Civil Rights Analyst, at cruiz@unidosus.org.