UNIDOSUS
STRONGER COMMUNITIES. STRONGER AMERICA.

SUPPLEMENTAL WRITTEN TESTIMONY ON CIVIL RIGHTS IMPLICATIONS OF FEDERAL USE OF FACIAL RECOGNITION TECHNOLOGY

Submitted to

The U.S. Commission on Civil Rights

Submitted by

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

May 14, 2023

On behalf of UnidosUS,* we respectfully submit this supplemental testimony to respond, as requested by Commissioner Magpantay, to testimony presented by the Department of Homeland Security (DHS) at the March 8, 2024, briefing: "Civil Rights Implications of the Federal Use of Facial Recognition Technology: Guidance for Meaningful Federal Oversight."

The testimony by Peter Mina, Deputy Officer for Programs and Compliance with the Office of Civil Rights and Civil Liberties (CRCL) at the DHS (referred to as the "CRCL testimony"), outlines commitments by the DHS to using biometric systems according to overarching principles of "civil rights, civil liberties, and privacy." DHS Directive 026-11, issued Sept. 2023, on the Use of Face Recognition and Face Capture Technologies (hereinafter referred to as the "Directive"), describes "how DHS will ensure that its use of face recognition and face capture technologies is subject to extensive testing and oversight." Press reports also provide a glimpse of the many expanding uses of Artificial Intelligence (AI) across DHS functions.

These statements were made in a context in which the DHS has been the subject of specific investigations by the Office of Inspector General, the Government Accountability Office (GAO), and research by advocacy groups, including:

- A September 2023 GAO report found that law enforcement at the DHS and the Department of Justice (DOJ) lacked basic protocols or training around the use of facial recognition technologies (FRT).
- Immigration and Customs Enforcement (ICE) used facial recognition technology 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 Customs and Border Protection (CBP), ICE, and the Secret Service purchased and used commercial geolocation data in violation of their privacy policies and that the 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 with DHS privacy policies and because the DHS Privacy Office did not follow or enforce 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 to stop purchasing location data by the end of Sept 2023.

* UnidosUS is a nonprofit, nonpartisan organization serving 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.

- 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 Department of Homeland Security are buying access to and using vast volumes of people's cell phone location information extracted from smartphone apps.

Unfortunately, the record suggests only a partial or delayed response to these serious concerns. Oral testimony at the hearing highlighted that DHS's training efforts were only getting underway in April. Thus, subsequent progress on training of staff is a legitimate and worthy topic for further inquiry by the Commission. The CRCL's written testimony merely describes broad themes as it considers DHS's use of FRT in light of civil rights and civil liberties, including such critical subjects as bias, accuracy, and validation.

We support work by DHS to improve its trainings and programs, but note as follows:

- Despite decades of deploying these technologies, the Directive generally fails to provide adequate specific and context, outlining only broad and poorly defined commitments to oversight processes in principle, without providing specific activities that would be needed to ensure that a system's use of technologies, including artificial intelligence (AI) or machine learning models as well as data surveillance, are protective of civil rights and civil liberties in practice. Vaguely worded ideals are not a substitute for an enforceable governance framework anchored by constitutional norms and developed through an inclusive and transparent democratic process.
- A democratic, human-centered, and rights-protecting governance model for surveillance systems and the technologies that drive and scale them requires formalizing agency engagement with impacted groups and their use of these channels for continuous input, feedback, and assessment.
- Assurances about the "technical" accuracy and validity of systems are fine, but fail to correct the human factors and systemic incentives that will drive outcomes in practice. To understand how systems operate in the world, collecting and evaluating empirical data on that is reviewed by impacted groups with the power to influence these systems is the only way to reconcile uses with fairness and other values.

DHS Statements Reveal Considerable Gaps in Governance of FRT and Biometric Surveillance

We cannot expect surveillance technologies to advance democratic governance without actionable standards based on impacts in the real world. As a law enforcement agency combining criminal and civil responsibilities, AI uses by DHS and its sub-agencies are among the most high-risk and rights-impacting. The steep imbalance of power between immigrants and communities of color and law and immigration enforcement heightens risks for abuse, information gaps, and bias.

The ground truth for any framework governing the use of surveillance models in the law and immigration enforcement setting is the experience of individuals—whether approaching a TSA checkpoint at the airport, peacefully assembling in protest, or being tracked and targeted

through a combination of scraped social media profiles and stored biometric data—not in a testing laboratory detached from these contexts.

Accuracy is a core component of assessing system validity overall, as rightfully noted throughout the CRCL testimony and DHS Directive. Yet statistical benchmarks alone are insufficient metrics by which to measure whether a system, and the practices and policies that govern it, are aligned with constitutional values and sound democratic practices. We can "solve" the accuracy problem, and still get the uses of the technology very, very wrong.

Progress is also likely to be hampered by the OMB policy to allow waivers for "law enforcement and national security" related uses of AI without the safeguards that should exist for rights- or privacy-impacting uses of AI. Such waivers erode any incentive to do the hard work of aligning the design of systems with rights and liberties. Unfortunately, the final policy permits agency officials to decide whether an AI use can be excluded from even the minimum safeguards it offers. It gives officials broad authority to exempt a system if they believe doing so would impede operations or increase overall risks to safety or rights. It also allows officials to sidestep them altogether if AI merely informed an action and was not the "principal basis" for it. Given the profound internal pressures AI officials will face, these loopholes could easily swallow the policy.

But regardless of whether a policy requires it, federal agencies must ensure that AI uses align with our nation's core civil rights, constitutional principles, and democratic values. Communities of color, including Latinos, are among the first to be targets of tech surveillance and the last to benefit from its advances. Given its reach and power, deploying AI responsibly and ethically requires new and innovative forms of governance. Checking a box—or escaping review altogether—will fail to address profound power imbalances at the heart of our shared technological future.

Rather than providing exemptions or waivers for law and immigration enforcement uses, we need enforceable and binding standards for all surveillance technologies and the AI models that scale them. Therefore, DHS should develop a more tailored approach to highly sensitive use cases and ensure the Department's uses include necessary safeguards.

The failure to align FRT and other biometric systems with an appropriate governance framework perpetuates vulnerable communities' exposure to unaccountable and harmful surveillance practices. People in immigrant and mixed-status communities are far too often test cases for policies that roll back protected rights and liberties, such as DHS's decades of dragnet surveillance. These same communities have, historically, been left behind and left out of both technological advances and tech governance.

In light of our Fourth Amendment, we always must balance individual liberties and protections against government overreach, with the need to govern and prevent harm. For this reason, among many, we should expect law enforcement institutions to develop both specific and tangible safeguards fashioned for specific factual circumstances, as well as monitoring and

training that reconciles our values with concrete practices at key decision points. This course of action creates supervisory incentives to follow the rule of law and the policies at moments that may be exigent.

Without actionable, enforceable, and democratic mechanisms to animate these principles in real-world applications, use of these systems will side-step accountability and harm groups of people who are disproportionately subjected to them. Deputy Officer Mina's testimony highlights several aspects of the DHS Directive on FRT, including:

- Requirements for testing in accordance with "national standards."*
- Periodic testing of current uses of biometric surveillance tools, and
- A process for core oversight by offices within DHS to review all new uses of biometric surveillance technologies before implementation.

While important, these broad prescriptions lack even a semblance of sufficient specificity to achieve accountability or transparent uses. They also fail to situate appropriate assessments of these tools (and their outputs) in the context of real-world use cases or solicit feedback from impacted communities on their flaws. Additionally, in terms of testing, there is no assertion the evidence and results of these tests would be publicly shared, nor is there a detailed public plan to address the findings of the tests and make improvements.

Having deployed these systems for decades, DHS is now undertaking what the CRCL testimony describes as "the difficult task of retrofitting complex systems" with civil rights, civil liberties, and privacy considerations. However, the legacy of unchecked deployment, plus the absence of a more developed and detailed approach to address flaws and lessons learned, fail to inspire confidence that the DHS will align its deployment of biometric surveillance systems with sound policy and constitutional norms.

DHS's guidance and the CRCL testimony also fail to establish the minimum benchmarks that are needed, beyond vague references to technical "standards" that do not yet really exist. The truth is that we are building the plane while flying it.

A democratic, human-centered, and rights-protecting governance model for AI-driven surveillance systems requires novel mechanisms for agency engagement with impacted groups for continuous input, feedback, and assessment.

The DOJ and DHS should evaluate current uses in light of each of these 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

* The phrase "national standards" in Mr. Mina's testimony appears to reference DHS Directive 026-11 instruction for key DHS offices to "[d]evelop accuracy and performance metrics, and procedures for testing and evaluating FR and FC technologies in accordance with International Organization for Standardization/ International Electrotechnical Commission (ISO/IEC) standards and technical guidance issued by National Institute of Standards and Technology (NIST)."

indicates, to risks and 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.

There are sound reasons for both departments to take a far closer look at current uses and the lessons those offer before rushing to adopt new ones. Successful systems are more difficult than appears at first glance to build and execute. For example, the National Institute of Standards and Technology's (NIST) AI Risk Management Framework (RMF) directs entities to make systems "Fair—with Bias Managed." Goals like achieving a "fair" model can be in tension with the purely technical or statistical accuracy of a model, as Brian Christian explains in his excellent book, The Alignment Problem, in which he provides specific examples of researchers' efforts to grapple with algorithmic bias in parole decision-making.

For this and other reasons, the NIST RMF further notes that fairness can be a contested, situational, and difficult factor to satisfy—yet fairness along multiple dimensions is a critical factor to get right in law and immigration enforcement settings, both legally and morally. Such decisions cannot, therefore, be a function of math alone—human judgment and democratic input, as well as transparency about the tradeoffs to the extent they exist, is crucial. Defining model fairness, for example, and the inputs and mechanisms needed to satisfy that definition is itself a civil rights policy matter.

Only by creating a means to incorporate impacted perspectives could DHS—or any agency—establish a set of applicable and responsive constitutional and accountable practices for biometric surveillance or other powerful forms of AI models. Notably, CRCL's testimony makes only a single passing reference to public engagement despite asserting that the CRCL considers themes of discrimination, accuracy, scale, use, perception, redress, and unintended consequences when reviewing and supporting the DHS surveillance programs.

The CRCL testimony also references a commitment to "*providing information* on the DHS Directive through its public engagement venues" (*emphasis added).* Yet nothing in the written testimony or DHS Directive provides a plan for agency consultation and engagement with impacted groups to inform model testing, policy-setting, or even to inventory and understand the real-world contexts and consequences of the agency's current and potential use cases.

CRCL's testimony acknowledges that "designing [biometric systems] with civil rights considerations from the beginning avoids the difficult task of retrofitting complex systems." Yet it lacks a commitment to meaningfully include and consult impacted and community groups to inform system design, assessment, and pre-deployment plans. In other words, even in the event we could achieve perfect technical accuracy for model outputs, the agency would still have to address open questions about the applications of the CRCL's nominal guiding principles of discrimination, accuracy, scale, use, perception, redress, and unintended consequences.

Consistent and comprehensive impact assessments that focus on real-world settings and the impact of technology are a vital source of feedback and learning to strengthen the defensibility, effectiveness, and fairness of specific uses of technology. Evaluations in partnership with impacted groups and other stakeholders can surface overlooked issues, generate empirical insights on how systems perform, and develop data on impacts to—and the experiences of—affected people. Findings could directly inform iterative improvements to policies, model training, and other policies and training as systems evolve.

Creating a system of democratized governance to shape standards and systems is also imperative to earn public trust. An insightful recent paper notes that participatory design methods are increasingly at the forefront of AI research and exploration:

> Community-based participatory design is an approach to designing computing technologies with and for different publics, with the aim of forming more equitable relationships between algorithmic systems and often-marginalized publics. [] Computing systems are rarely developed entirely by the publics they serve; and in this way, participatory design is a situated practice of future-making through which heterogeneous communities collaboratively imagine new sociotechnical futures. While participatory design has a long tradition in shaping the design of computing systems, it has more recently become a means to co-create artificial intelligence (AI) transparency and accountability artifacts, such as model cards, design workbooks, and user agreements. [*Citations omitted.*]

The authors envision five dimensions for the participatory design of user agreements. Applying a comparable vision to the present context, we could translate these as a call for:

1) Participatory development of performance standards for models;
2) Structures within model designs that anticipate and defend against potential harms;
3) Opportunities to provide and revoke informed consent;
4) Complaint mechanisms for harms when they occur and a means of redress;
5) Disclosures and labeling of limitations and performance; and
6) External information gathering about potential and actual harms to drive iteration on standards.

"Fulsome consideration" of these systems by DHS, as CRCL's testimony promises, could not possibly be fulsome without including such feedback. When governance of technology, as it does here, overemphasizes the purely technical aspects of model standardization testing and output validation, it too often problematically ignores the implications of—and flaws intrinsic to—rights-impacting uses.

Yet another problematic and intrinsic limitation of most data-driven and AI–powered surveillance systems are the pervasive power imbalances between deployers of technologies and the people they are used to target. Because minorities are, by definition, less represented in datasets, data-driven models have less information about these groups than they do about a given majority. Further, models can draw subtle inferences from data in myriad ways that the way things are now is the way they should be in the future—in effect, mistaking what is distributionally true for what is morally or ethically true about human difference and potential.

Sometimes, failure modes are more obscure. Models have been caught gathering clues related to race or gender in hiring decisions, for example, from word choice or specific activities listed on a resume, even when information on race and gender has been omitted, leading companies to discard the models as too inherently biased to be used.

Relatively simple automated decision-making models have also been shown to be deeply biased and to lack predictive value in areas such as mortgage lending, given the number of factors that function as proxies for race, even when protected class is omitted. Moreover, the black box nature of many models means that subtle forms of bias may remain undetected without specific steps, including interrogation of the model for bias, impact assessments, and other forms of actual empirical evaluation.

These differences can be deadly if unnoticed. For example, as an article on the underdeveloped science of data quality for AI in *Stanford Medicine* notes:

> Scientists at Duke University Hospital, for instance, designed an AI program to identify children at risk of sepsis, a dangerous response to an infection. But the program took longer to flag Latino kids than white kids, possibly delaying the identification and treatment of Latino children with sepsis. The bias, it turned out, existed because doctors themselves took longer to diagnose sepsis in Latino kids. This taught the AI program that these children might develop sepsis more slowly or less often than white children.

Three years into the effort, researchers learned that doctors took longer to diagnose the sepsis in Latino children, possibly because, among other reasons, Latino families were awaiting the arrival of hospital translators. Without external reviewers to gut-check and validate the conditions on the ground that produce data, many AI builders, or users, may not know what they do not know. A September 2020 report on uses of AI in sepsis monitoring programs, *Repairing Innovation*, notes that:

> …all too often, potential solutions remain just that—potential solutions, which may work in theory, given pre-set conditions. Rarely are these solutions tested, verified, or even used "in the wild." For this reason, we need fewer studies proposing how AI technologies could be used to address existing problems in the abstract and more studies exploring how and in what ways could AI technologies be integrated into existing social processes such that they actually address those problems.

In other words, unless designed well, data-driven models have blind spots related to serious but sometimes intangible considerations essential to human values and decision-making–such as respect for civil rights, civil liberties, fairness, redressability, transparency, and privacy. The testimony here shows that the DHS is, unfortunately, far from a comprehensive policy on FRT or AI models.

In particular, we call on HUD, DOJ, and DHS (and all federal agencies) to give impacted communities a voice in governance through practical mechanisms that provide a means of feedback for agencies about the uses and impacts of technologies in real time. We outline below 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.

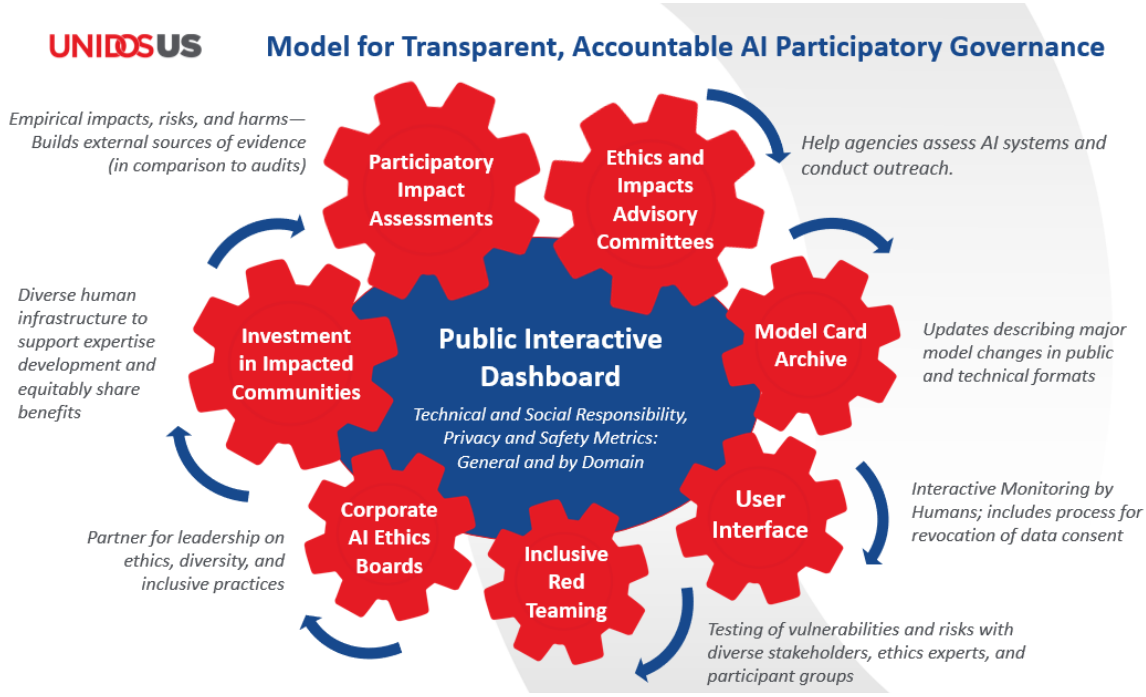

In addition, the Departments' Use Inventories and proposed risk management approaches should be organized according to the "AI Risks and Trustworthiness" issues described by NIST, which highlight that AI systems should meet baselines for each of the following factors: 1) Valid

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.

The DOJ and DHS 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 safety- and rights-impacting uses.

For its part, the Commission should ask DHS as follows:

1) What stakeholder engagement is planned around specific use cases?
2) What empirical testing has been conducted with real-world users (*e.g.*, has TSA measured whether or not airline passengers fully understand their ability to withhold consent at checkpoints to FRT uses? What evidence suggests that they do or do not feel at liberty to make a choice?)
3) What changes in biometric surveillance policy or other steps are DHS considering, if any, in light of the OMB EO?
4) What programs at DHS will likely receive waivers in December and what is the timeline for aligning those uses with the requirements of the EO?
5) What are the implications for FRT of Secretary Mayorkas's new AI initiative?
6) What evidence of consent for data collection does DHS regarding its well-publicized collections of DMV and utility data that include US citizens?
7) Why does DHS feel a need to hold onto non-citizen data for an astonishing 75 years?
8) What training has been completed at this stage in terms of the timeline and what has yet to be done?
9) Does the planned training adequately capture the human factors and limitations of the tech use cases? Will the training be shared with the Commission? Has it been independently evaluated?
10) What are safety evaluations that have been conducted by DHS regarding risks to migrants and new, more hazardous corridors that may be created in light of the plan to expand DHS surveillance technologies and towers at the border?

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.