



**United States House of Representatives
Committee on Oversight and Reform**

**“Facial Recognition Technology (Part III):
Ensuring Commercial Transparency & Accuracy”
January 15, 2020**

Written Testimony of Meredith Whittaker
Cofounder, AI Now Institute, New York University¹

Chairwoman Maloney, Ranking Member Jordan, and members of the Committee, thank you for inviting me to speak today. My name is Meredith Whittaker and I am the cofounder of the AI Now Institute at New York University. AI Now is the first university research institute dedicated to studying the social implications of artificial intelligence and algorithmic technologies (AI). Our work examines the rapid proliferation of AI systems through social domains such as criminal justice, healthcare, employment, and education. In particular, we focus on concerns in the areas of bias and inclusion, safety and critical infrastructure, rights and liberties, and labor. As we identify problems in each of these spaces, we work to address them through robust research, community engagement, and key policy interventions.

Until recently, I was also a longtime tech worker. I founded Google’s Open Research Group, and left the company in July 2019 after 13 years. While at Google, I led efforts to examine the ethics and fairness of AI systems. I also helped lead worker organizing, pushing back against unethical uses of Google’s technology, as well as against workplace bias and discrimination.

ai
The widespread deployment of facial recognition technologies raises important concerns that require urgent attention from lawmakers and regulators. **It is imperative that lawmakers act to protect fundamental rights and liberties and to ensure that these powerful technologies do not exacerbate inequality and enable social control as they transform core institutions and infrastructures.**

In this testimony, I will make six key points:

¹ This testimony was prepared in collaboration with staff across the AI Now Institute. It incorporates insights and research from many disciplinary perspectives, in service of offering a nuanced set of arguments and observations about a significant and complex topic. In particular, Roel Dobbe, Theodora Dryer, Genevieve Fried, Ben Green, Amba Kak, Joy Lisi Rankin, Varoon Mathur, Andrea Nill Sanchez, Deborah Raji, Rashida Richardson, and Jason Schultz all offered insights, sources, and editorial contributions.

1. **Facial recognition reflects and amplifies historical and present-day discrimination.** Even if it were possible to make facial recognition accurate for everyone, ensuring accuracy does not address the social context in which it will be deployed, and will not reduce harms like abuse and discriminatory deployment. Facial recognition allows businesses and governments to intrude into people’s lives without detection, and currently there are few guardrails in place to curtail biased and oppressive uses. Facial recognition is usually deployed by those who already have power—like employers, landlords, and police—to surveil and control those who have less power. Therefore, problems like racial profiling are likely to worsen with tools like facial recognition, especially as these technologies are disproportionately deployed to surveil Black, Latinx, and immigrant communities who already face systemic oppression and overpolicing.
2. **Most facial recognition systems in use are developed and sold by private companies.** This is true regardless of whether governments or private actors are the end users. The commercial nature of these technologies means they are shielded from accountability and oversight by claims of corporate secrecy. This structural secrecy makes it difficult for the public and regulators to understand how and where facial recognition is being used, and to detect and redress harms.
3. **There is a blurry line between public and private facial recognition.** This means we need to examine commercial systems and the incentive structures driving their development even in discussions that focus on government use. Data-sharing agreements between private facial recognition vendors and government also require scrutiny. Further, many private deployments of facial recognition assume government intervention, as when retail establishments use it to “detect” shoplifters and turn them over to police.
4. **Affect recognition and facial analysis pose particular dangers.** In addition to problems with basic facial detection and identification, attempts to “recognize” emotions or “types” of people on the basis of facial expression lack sound scientific support and further embed bias and discrimination within our society.
5. **Standards and technical fixes are not enough to solve the problems with facial recognition.** Standards for facial recognition assessment and auditing are a step in the right direction; however, such technical standards will never be sufficient to ensure that facial recognition is just or ethical. Further, narrow or weak standards run the risk of providing “checkbox certification,” allowing vendors and companies to assert that their technology is safe and fair without accounting for how it will be used, or its fitness for a given context. If such standards are positioned as the sole check on facial recognition systems, they could function to obfuscate harm instead of mitigate it.
6. **It is time to halt the use of facial recognition in sensitive social and political contexts, by both government and private actors.** Facial recognition poses an

existential threat to democracy and liberty, and fundamentally shifts the balance of power between those using facial recognition and the populations on whom it's applied. This is true both in government and commercial contexts. While auditing standards and transparency are necessary to answer fundamental questions, they will not address these harms. It is urgent that lawmakers act to halt the use of facial recognition in sensitive social and political domains until the risks are fully studied and adequate regulations that center the communities most affected are in place.

In this testimony I use the broad term “facial recognition” to include a range of technical capabilities, including face detection (recognizing a face in an image),² facial identification and verification (recognizing a single face, and distinguishing it from others), and facial analysis (inferring demographics, identity, and interior traits based on face data). While these constitute discrete capabilities that are often treated separately within the AI research field, the deployment of these tasks raises shared concerns. These functions are also often linked or packaged together, as when facial analysis is sold as an “add-on” to facial recognition products. Furthermore, many systems for facial analysis are trained on the same datasets used to develop facial recognition and face-detection systems, meaning that bias and limitations from those datasets can affect performance on all tasks.³

Facial recognition is inaccurate and reflects and amplifies historical and present-day discrimination

In weighing the use of facial recognition, it is important to recognize that in most cases, the “user” of a facial recognition tool is not the general public. The user is the business or government that licenses facial recognition systems from technology companies, applying it across a range of diverse applications—from choosing which job candidate to hire, to detecting shoplifters, to identifying criminals. Facial recognition provides businesses and governments with powerful biometric surveillance tools that increase their reach into people's lives, facilitating monitoring and control without clear guardrails.

² The majority of facial recognition systems in deployment use 2-D images as training data. Some facial recognition systems also use thermal data and other face data collected by sensors, sketches, video, or 3-D images.

³ See Letter from Concerned Researchers, *On Recent Research Auditing Commercial Facial Analysis Technology* (Mar 26, 2019), <https://medium.com/@bu64dcjrytwitb8/on-recent-research-auditing-commercial-facial-analysis-technology-19148bda1832>; see also Dina Bass, *Amazon Schooled on AI Facial Technology By Turing Award Winner*, Bloomberg (Apr. 3, 2019), <https://www.bloomberg.com/news/articles/2019-04-03/amazon-schooled-on-ai-facial-technology-by-turing-award-winner>.

Members of the general public have very little say over the application of the technology, although they are frequently the targets of its use. Those surveilled and targeted by facial recognition are very often unaware it is being used to surveil them and to shape decisions about their lives. Nor are they informed of the systems and companies capturing and processing their biometric data. Whether it's applied by governments or private actors, facial recognition is usually deployed by those who already have power—like employers, landlords, and police—to surveil and control those who have less power.

Many current uses undermine the constitutional rights of free association, free expression, and due process, while also enabling suspicionless surveillance and social control.^{4,5,6} A clear example of this tendency is the partnership between IBM and the New York Police Department, in which IBM trained a facial recognition model to classify people by skin tone, using NYC surveillance footage of the public collected without consent. IBM's system was intended to allow police to search the database by ethnicity, providing a tool for racial profiling alongside mass surveillance.⁷ Woodrow Hartzog, a law professor and privacy scholar, put it bluntly: "Facial recognition can be incredibly harmful when it's inaccurate and incredibly oppressive the more accurate it gets."⁸ The use of facial recognition for police surveillance also disproportionately affects Black communities, immigrant communities, and other communities of color who already face overpolicing and are most vulnerable to targeting and discrimination.

There is already evidence of the harm that can result from inaccurate facial recognition systems. In the United Kingdom (UK), documents uncovered through Freedom of Information (FOI) requests revealed that eight trials of a facial recognition system used by the police in London had an average 96 percent error rate, persistently misidentifying residents as criminals and leading to detention and harassment.⁹ A trial of facial recognition to identify drivers in New York failed completely, with 100 percent error rates, meaning that the technology correctly identified

⁴ An Act Establishing a Moratorium on Face Recognition and Other Remote Biometric Surveillance Systems: Hearing Before the Massachusetts Joint Committee on the Judiciary (Mass. Oct. 21, 2019) (statement of Marc Rotenberg, Exec. Dir., Elec. Privacy Info. Ctr.), <https://epic.org/testimony/congress/EPIC-FacialRecognitionMoratorium-MA-Oct2019.pdf>.

⁵ Andrew Guthrie Ferguson, *Facial Recognition and the Fourth Amendment* (October 21, 2019), <https://ssrn.com/abstract=3473423>.

⁶ Woodrow Hartzog and Evan Selinger, *Surveillance as Loss of Obscurity*, 72 Wash. & Lee L. Rev. 1343 (2015), <http://scholarlycommons.law.wlu.edu/wlulr/vol72/iss3/10>.

⁷ George Joseph & Kenneth Lipp, *IBM Used NYPD Surveillance Footage to Develop Technology That Lets Police Search by Skin Color*, *The Intercept* (Sept. 6, 2018), <https://theintercept.com/2018/09/06/nypd-surveillance-camera-skin-tone-search>.

⁸ Olivia Solon, *Facial Recognition's 'Dirty Little Secret': Millions of Online Photos Scraped Without Consent*, *NBC News* (March 12, 2019), <https://www.nbcnews.com/tech/internet/facial-recognition-s-dirty-little-secret-millions-online-photos-scrape-d-n981921>

⁹ Westminster Hall Debate: Facial Recognition and the Biometrics Strategy, sponsored by Darren Jones MP (May 1, 2019) (Briefing of Big Brother Watch), <https://bigbrotherwatch.org.uk/wp-content/uploads/2019/05/Big-Brother-Watch-briefing-on-Facial-recognition-and-the-biometric-strategy-for-Westminster-Hall-debate-1-May-2019.pdf>.

no one.¹⁰ Apple’s facial recognition incorrectly identified a student as a thief, leading to a false arrest, while in Brazil a similar case led to a woman being identified as a criminal.^{11,12} A student at Brown University was falsely identified as a bombing suspect, leading to death threats and online abuse.¹³ Uber’s face recognition authentication system locked transgender Uber drivers out of their accounts, failing to recognize them and leaving them unable to work.¹⁴ And in Florida, police used the FACES facial recognition system to identify Willie Allen Lynch as a suspect based on a cell phone picture.¹⁵ The system came back with a very low confidence match, which was used to prosecute Lynch. However, at trial, the fact that a facial recognition system was used to identify Lynch, along with the low confidence match results—evidence that could prove Lynch’s innocence or at least establish reasonable doubt—were withheld from the defense.

Research underscores the significant problems of bias and inaccuracy in commercial facial recognition systems. Although many of these systems may boast high overall accuracy rates, they perform considerably less well when their accuracy is measured against different demographic subgroups. Their failures are particularly profound for Black women, Native Americans, gender minorities, young and old people, and other underrepresented groups. The National Institute of Standards and Technology (NIST),^{16,17} researchers Joy Buolamwini, Timnit

¹⁰ Paul Berger, *MTA’s Initial Foray Into Facial Recognition at High Speed Is a Bust*, Wall St. J. (Apr. 7, 2019), <https://www.wsj.com/articles/mtas-initial-foray-into-facial-recognition-at-high-speed-is-a-bust-1155464200>.

¹¹ Bob Van Voris, *Apple Face-Recognition Blamed by N.Y. Teen for False Arrest*, Bloomberg (Apr. 22, 2019), <https://www.bloomberg.com/news/articles/2019-04-22/apple-face-recognition-blamed-by-new-york-teen-for-false-arrest>.

¹² Pedro Maia, *The Usage and Dangers of Facial Recognition Technology*, Impakter (Sept. 12, 2019), <https://impakter.com/the-usage-and-dangers-of-facial-recognition-technology>.

¹³ Jeremy C. Fox, *Brown University Student Mistakenly Identified as Sri Lanka Bombing Suspect*, Boston Globe (Apr. 28, 2019), <https://www.bostonglobe.com/metro/2019/04/28/brown-student-mistaken-identified-sri-lanka-bombings-suspect/0hP2YwyYi4qrCEdxKZCpZM/story.html>.

¹⁴ Jaden Urbi, *Some Transgender Drivers Are Being Kicked off Uber’s App*, CNBC (Aug. 8, 2018), <https://www.cnbc.com/2018/08/08/transgender-uber-driver-suspended-tech-oversight-facial-recognition.html>; Steven Melendez, *Uber Driver Troubles Raise Concerns About Transgender Face Recognition*, Fast Company (Aug. 9, 2018), <https://www.fastcompany.com/90216258/uber-face-recognition-tool-has-locked-out-some-transgender-drivers>.

¹⁵ Rashida Richardson, Jason M. Schultz, & Vincent M. Southerland, *Litigating Algorithms 2019 US Report: New Challenges to Government Use of Algorithmic Decision Systems*, AI Now Inst. (Sept. 2019), <https://ainowinstitute.org/litigatingalgorithms-2019-us.pdf>.

¹⁶ Patrick Grother, Mei Ngan, & Kayee Hanaoka, NIST Interagency Report 8280, *Face Recognition Vendor Test (FRVT), Part 3: Demographic Effects* (Dec. 2019), <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf>.

¹⁷ Mei Ngan & Patrick Grother, NIST Interagency Report 8052, *Face Recognition Vendor Test (FRVT): Performance of Automated Gender Classification Algorithms* (April 2015), <https://nvlpubs.nist.gov/nistpubs/ir/2015/NIST.IR.8052.pdf>.

Gebru and Inioluwa Deborah Raji,^{18,19} the ACLU and UC Berkeley,²⁰ and many others^{21,22,23,24} have performed audits and other tests that all confirm: **facial recognition does not work as advertised, and its inaccuracies and errors are generally worst for populations that are already facing societal discrimination.**

Despite pressing civil rights and liberties concerns, significant research detailing facial recognition's bias, and the lack of affirmative evidence validating its accuracy and utility, facial recognition's use is accelerating across sensitive domains affecting hundreds of millions of people. The technology is supplanting time clocks at job sites,²⁵ airline boarding passes at airports,²⁶ keys or other entry mechanisms for housing units,²⁷ safety systems or protocols at schools,²⁸ security at sport stadiums and event locations,²⁹ and it's being used to monitor children at summer camp and to authenticate gig workers when they log in to work, to name

¹⁸ Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, Proceedings of Machine Learning Research 81:1-15 (2018), <http://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>.

¹⁹ Inioluwa Deborah Raji & Joy Buolamwini, *Actionable Auditing: Investigating the Impact of Publicly Naming Biased Performance Results of Commercial AI Products*, Proceedings of the Conf. on Artificial Intelligence, Ethics, and Society (2019), https://www.aies-conference.com/2019/wp-content/uploads/2019/01/AIES-19_paper_223.pdf.

²⁰ Jacob Snow, *Amazon's Face Recognition Falsely Matched 28 Members of Congress With Mugshots*, ACLU (July 26, 2018), <https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28>.

²¹ KS Krishnapriya, Kushal Vangara, Michael C. King, Vitor Albiero, & Kevin Bowyer, *Characterizing the Variability in Face Recognition Accuracy Relative to Race*, Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition Workshops (2019), <https://arxiv.org/abs/1904.07325>.

²² Cynthia M. Cook, et al., *Demographic Effects in Facial Recognition and Their Dependence on Image Acquisition: An Evaluation of Eleven Commercial Systems*, IEEE Transactions on Biometrics, Behavior, and Identity Science 1:1 (Jan. 2019), <https://ieeexplore.ieee.org/document/8636231>.

²³ Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, & Emily Denton, *Saving Face: Investigating the Ethical Concerns of Facial Recognition Auditing*, Proceedings of the AAAI/ACM Conf. on AI, Ethics, and Society (2020), <https://arxiv.org/pdf/2001.00964v1.pdf>

²⁴ Morgan Klaus Scheuerman, Jacob M Paul, & Jed R. Brubaker, *How Computers See Gender: An Evaluation of Gender Classification in Commercial Facial Analysis Services*, Proceedings of the ACM on Human-Computer Interaction 3:1-33 (Nov. 2019), <https://dl.acm.org/doi/10.1145/3359246>.

²⁵ *4 Reasons to Use Time Clocks With Facial Recognition*, Buddy Punch (Jun. 19, 2018), <https://buddypunch.com/blog/time-clocks-facial-recognition>.

²⁶ Francesca Street, *How Facial Recognition Is Taking Over Airports*, CNN (Oct. 8, 2019), <https://www.cnn.com/travel/article/airports-facial-recognition/index.html>

²⁷ Ginia Bellafante, *The Landlord Wants Facial Recognition in Its Rent-Stabilized Buildings. Why?*, N.Y. Times (Mar. 28, 2019), <https://www.nytimes.com/2019/03/28/nyregion/rent-stabilized-buildings-facial-recognition.html>

²⁸ Sarah St. Vincent, *Facial Recognition Technology in US Schools Threatens Rights: Children of Color at Greatest Risk*, Human Rights Watch (June 21, 2019), <https://www.hrw.org/news/2019/06/21/facial-recognition-technology-us-schools-threatens-rights>

²⁹ Kevin Draper, *Madison Square Garden Has Used Face-Scanning Technology on Customers*, N.Y. Times (Mar. 13, 2018), <https://www.nytimes.com/2018/03/13/sports/facial-recognition-madison-square-garden.html>.

only a small number of examples.^{30,31} These are all cases in which bias and error can have significant life-altering consequences, denying people access to resources, shelter, and liberty.

In late 2018, the landlord of Atlantic Plaza Towers in Brooklyn, New York, shared his intention to install Stonelock biometric access technology and replace key fobs with facial recognition. Since the building was rent stabilized, major modifications like changing mechanisms for entry required approval from a state agency, and the landlord applied for such approval in March 2019. While the use of facial recognition for authentication and entry in housing is increasing, in this case tenants were alerted to the deployment before it happened, something that is not legally required in most residences. On learning of the switch, tenants rapidly organized and filed a challenge to the state's Homes and Community Renewal department, asking the agency to block the facial recognition system on privacy and ethical grounds.³²

In their defense against the use of this technology, these tenants brought forth a number of specific concerns that lawmakers across the country should pay attention to. First, who would own their biometric data once it was collected? How would it be stored, and what were the rules around sharing and reuse? They also raised concerns about bias and discrimination, referencing research that showed persistent errors and inaccuracies in commercial facial recognition systems that were most pronounced for demographic groups that lived in the building and surrounding neighborhood—women and Black and Latinx people. The residents also raised concerns about how the technology could be abused, noting that the landlord had previously used video footage to harass and monitor tenants, a practice they feared facial recognition would only exacerbate.³³

Proponents of facial recognition rarely account for the fundamental power imbalance built into the way facial recognition is developed and deployed.³⁴ These technologies work to increase existing power asymmetries in ways that benefit those already in positions of privilege. Any responsible assessment of facial recognition and its risks needs to be carried out with a sober understanding of the history of racial and gender-based discrimination, and should recognize the potential of this technology not only to enable forms of mass surveillance and social control

³⁰ Elizabeth Weise & Molly Horak, *Hey Mom, Did You See This? Camps Are Using Facial Recognition, Latest Use of Controversial Tech*, USA Today (Jul. 17, 2018), <https://www.usatoday.com/story/tech/news/2018/07/17/facial-recognition-helps-mom-and-dad-see-kids-camp-photos-raises-privacy-concerns-some/780725002>.

³¹ See Jaden Urbi, Steven Melendez, *supra* note 13.

³² Opposition to Owner's Application for Modification of Services to Install a Facial Recognition Entry System, *In the Matter of the Owners' Application for Modification of Services v. Tenants of Atlantic Plaza Towers*, Docket Nos. GS2100050D, GS2100080D, (NYS Housing & Community Renewal Office of Rent Administration/MCI Unit, Apr. 30, 2019), <https://www.legalservicesnyc.org/storage/PDFs/%20opposition%20to%20facial%20recognition%20entry%20system%20app.pdf>.

³³ Erin McElroy, *Disruption at the Doorstep*, Urban Omnibus (Nov. 6, 2019), <https://urbanomnibus.net/2019/11/disruption-at-the-doorstep>.

³⁴ Khari Johnson, *AI Ethics Is All About Power*, VentureBeat (Nov. 11, 2019), <https://venturebeat.com/2019/11/11/ai-ethics-is-all-about-power>.

that harm people who are already suffering from social discrimination, but also to endanger our collective freedoms.

From aviation to healthcare, there are few—if any—contexts in which American society permits companies to treat the public as experimental subjects, deploying untested, unverified, and faulty technology that has been proven to amplify bias and discrimination. With consequences that extend from threatening people's livelihoods to putting them in mortal danger due to misidentifying them as criminal suspects, it is clear that this technology leaves the public even more vulnerable than in the past—empowering institutions that may manipulate and discriminate against certain members, rather than truly protecting all of our interests or desire for agency and privacy.

Facial recognition is a commercial technology. And corporate secrecy prevents oversight and accountability of both government and private use

The majority of facial recognition systems in use are developed and sold by private companies. This means that even in discussions that focus on government use, we need to examine commercial systems and the incentive structures driving their development.³⁵

It also means that we need to challenge claims of corporate secrecy that prevent scrutiny and accountability. The ability to access facial recognition systems in order to audit and examine them is regularly blocked, guarded behind veils of corporate secrecy. This prevents researchers, journalists, lawmakers, and the public from fully understanding where, how, and with what consequences this technology is being used. It also means that access to use facial recognition is effectively only available to institutions (such as law enforcement or large corporations) that can afford to develop or license costly systems.

The contracts between facial recognition companies and the customers who license and use facial recognition systems are also generally shrouded in secrecy. Within large companies, these contracts are very closely guarded; in some cases, the contract itself requires that neither party—company or customer—disclose the existence of a contract, let alone how a given system will be used, and for what purpose. This is one example of the structural obscurity that protects corporate interests, conceals harm and misuse, and prevents lawmakers and the public from determining where, how, and whether such systems are appropriate. This is true of both government use of corporate systems and private use.

³⁵ Clare Garvie & Laura M. Moy, *America Under Watch: Face Surveillance in the United States*, Georgetown Law Center on Privacy & Technology (May 16, 2019), <https://www.americaunderwatch.com>.

The lack of diversity in tech shapes how AI companies work. It influences what kinds of products are built, who they are designed for, and who benefits from their deployment—but even that data is shrouded in secrecy. Substantial evidence shows that the companies developing facial recognition technologies are not reflective of society at large: women comprise only 15 percent of AI research staff at Facebook and 10 percent at Google; only 2.5 percent of Google’s workforce is Black, while Facebook and Microsoft are each at 4 percent.³⁶ Yet some firms attempt to extend trade secrecy even to their diversity data: both Oracle and Palantir made such claims in an attempt to block the Center for Investigative Reporting from accessing the equal employment opportunity data it files with the Department of Labor.^{37,38}

This culture of secrecy and lack of oversight allows facial recognition companies the freedom to make unvalidated claims about accuracy and efficacy. Stonelock, the company selling facial recognition to the landlord of Atlantic Plaza Towers, claimed that its system did not exhibit the racial, gender, and other biases found in similar systems. However, the company never submitted any evidence to substantiate this claim, nor did it open its system for validation and testing, effectively asking lawmakers, tenants, and the public to take them at their word.³⁹ Amazon made similar claims to accuracy that were not supported by research findings.⁴⁰ The company also refused to submit its facial recognition system to NIST for auditing, claiming that they were unable to modify it to comply with test specifications.⁴¹ In such cases, truth-in-advertising laws applied to AI companies would be helpful, holding companies liable for misrepresentations made in marketing, and giving the Federal Trade Commission or other designated agencies leverage for enforcement.

In the context of military or law enforcement use of facial recognition and other technical systems, there’s often double obscurity: corporate secrecy on one side, and classification or law enforcement transparency exemptions on the other. This is particularly troubling given that these are domains where some of the most serious risks of harm are present. It is worrying that

³⁶ Sarah Myers West, Meredith Whittaker, Kate Crawford, *Discriminating Systems: Gender, Race and Power in AI*. AI Now Inst. (Apr. 2019), <https://ainowinstitute.org/discriminatingsystems.html>.

³⁷ Jamillah Bowman Williams, *Diversity As A Trade Secret*, 107 Geo. LJ 1684 (2019), <http://scholarship.law.georgetown.edu/facpub/2097>.

³⁸ Will Evans & Sinduja Rangarajan, *Oracle and Palantir Said Diversity Figures Were Trade Secrets. The Real Secret: Embarrassing Numbers*, The Center for Investigative Reporting (Jan. 7, 2019), <https://www.revealnews.org/article/oracle-and-palantir-said-diversity-figures-were-trade-secrets-the-real-secret-embarrassing-numbers>.

³⁹ Amicus Letter of Rashida Richardson, AI Now Institute, in Support of Opposition to Owner’s Application for Modification of Services to Install a Facial Recognition Entry System, *In the Matter of the Owners’ Application for Modification of Services v. Tenants of Atlantic Plaza Towers*, Docket Nos. GS2100050D, GS2100080D (NYS Housing & Community Renewal Office of Rent Administration/MCI Unit, Apr. 30, 2019), <https://ainowinstitute.org/dhcr-amicus-letter-043019.pdf> (Owner’s Letter on file with author).

⁴⁰ See Letter from Concerned Researchers, Dina Bass, *supra* note 3.

⁴¹ Drew Harwell, *Federal Study Confirms Racial Bias of Many Facial-Recognition Systems, Casts Doubt on Their Expanding Use*, Washington Post (Dec. 19, 2019), <https://www.washingtonpost.com/technology/2019/12/19/federal-study-confirms-racial-bias-many-facial-recognition-systems-casts-doubt-their-expanding-use>.

new legislation intended to regulate and oversee facial recognition and other AI systems often includes exemptions for law enforcement.⁴²

Large tech companies are among the few organizations with the resources needed to develop and deploy machine-learning-based facial recognition and other AI systems at scale. While there are many facial recognition and AI startups, most of them license computational infrastructure from Amazon, Microsoft, or Google. Some also license their core technology—the facial recognition model itself—from these or other vendors, repackaging it for one or another domain-specific use case and selling this to customers.

Developing and deploying facial recognition and other AI systems at scale requires a combination of powerful computational infrastructure, massive amounts of biometric data, and the capital to recruit and retain rare and highly paid AI engineers. This combination of resources is both extremely expensive and very difficult to procure, even for those with the capital, since data collection is generally predicated on existing market reach. This combination of resources is not available to law enforcement and government agencies. Google, Amazon, Microsoft, and Facebook are leaders in this domain. Their position as dominant internet companies gave them access to vast amounts of consumer data, and spurred their investment in large-scale computational infrastructure. Over the last decade, these companies helped shape the field of AI. It is not surprising that DeepFace, a deep-learning facial recognition model that was the first to demonstrate the effectiveness of training facial recognition models using massive face datasets, was developed by Facebook, relying on the company's access to vast amounts of face data gathered from consumer profiles.⁴³

Currently, the companies building and selling facial recognition systems and other AI technologies are not subject to regulation and oversight capable of holding them accountable for the harms and errors their technology might inflict. This also applies to government use of commercial AI systems. If a company builds and deploys harmful technology, and misinforms a state or private actor of its capabilities, there are few remedies to hold the company accountable.

There is a blurry line between public and private facial recognition systems

Amazon's Ring, a surveillance doorbell system installed by individuals and businesses, provides a significant and troubling example of the complex interconnections between government and

⁴² See, e.g., Commercial Facial Recognition Privacy Act of 2019, S.847, 116th Congress (2019-2020), <https://www.congress.gov/bill/116th-congress/senate-bill/847>.

⁴³ Yaniv Taigman, et al., Facebook AI Research, *DeepFace: Closing the Gap to Human-Level Performance in Face Verification*, Proceedings of the 2014 IEEE Conf. on Computer Vision and Pattern Recognition (2014), <https://ieeexplore.ieee.org/document/6909616>.

private use of AI-enabled surveillance systems. Ring enables persistent surveillance of homes and neighborhoods, and while it does not currently include facial recognition, Amazon has filed a facial recognition patent in this space,⁴⁴ and appears to be planning to connect facial recognition capabilities to a “neighborhood watch list” database of people deemed suspect.⁴⁵ This raises serious concerns, given the documented racial bias in Amazon’s current facial recognition systems,⁴⁶ coupled with evidence that Ring’s use has led to a number of instances of racial targeting, in which people of color are reported as suspicious based on Ring footage.⁴⁷

But Ring is not simply a problematic consumer-facing service. It is also a pipeline to law enforcement. Amazon has partnered with at least 400 local police departments, enlisting officers as Amazon spokespeople to convince residents to install Ring systems. In exchange, police get access to a dashboard of Ring surveillance footage, either directly from users who opt in to share, or by submitting a request to Amazon.⁴⁸

Amazon Ring offers a clear example of the way private deployments of surveillance technology, including facial recognition, enable a backdoor to police and government surveillance. This is particularly troubling when it extends law enforcement monitoring into spaces previously inaccessible to them without a warrant, such as commercial properties or personal residences.⁴⁹ This example also shows how commercial facial recognition companies leverage government interest in data to expand their reach and acquire more data, which benefits their financial interests but does not ensure they are beholden to the needs of either the government or residents. Government enforcement reliance on Ring data or other privately run infrastructure could also pose serious issues if Amazon were to discontinue Ring’s development, institute a steep subscription fee, or make other changes well within the rights of a private company. The problems faced by the New York Police Department in attempting to retrieve data from Palantir

⁴⁴ U.S. Patent Application No. 15/984,298, Publication No. US 2018/0341835 A1 (published Nov. 29, 2018)(Amazon Technologies, Inc., applicant), https://www.aclunc.org/docs/Amazon_Patent.pdf.

⁴⁵ Sam Biddle, *Amazon’s Ring Planned Neighborhood “Watch Lists” Built on Facial Recognition*, *The Intercept* (Nov. 26, 2019), <https://theintercept.com/2019/11/26/amazon-ring-home-security-facial-recognition>.

⁴⁶ Inioluwa Deborah Raji & Joy Buolamwini, *Actionable Auditing*, *supra* note 18.

⁴⁷ Caroline Haskins, *Amazon’s Home Security Company Is Turning Everyone Into Cops*, *Motherboard* (Feb. 7, 2019), https://www.vice.com/en_us/article/qvyvzd/amazons-home-security-company-is-turning-everyone-into-cops.

⁴⁸ Drew Harwell, *Doorbell-Camera Firm Ring Has Partnered With 400 Police Forces, Extending Surveillance Concerns*, *Washington Post* (Aug. 28, 2019), <https://www.washingtonpost.com/technology/2019/08/28/doorbell-camera-firm-ring-has-partnered-with-police-forces-extending-surveillance-reach>.

⁴⁹ Evan Selinger, *Why You Can’t Really Consent to Facebook’s Facial Recognition*, *OneZero* (Sep. 30, 2019), <https://onezero.medium.com/why-you-cant-really-consent-to-facebook-s-facial-recognition-6bb94ea1dc8f>.

presents a cautionary example.⁵⁰

Data sharing between private companies and governments is a problem that extends beyond Ring. In many places, there is total lack of statutory, case law, or agency rules governing the sharing of biometric data with governmental agencies, third parties, or law enforcement. There are many cases of behind-the-scenes data-sharing arrangements that allow data collected by the private sector to be transferred and used by law enforcement, and, due to a lack of transparency, it is likely there are many more instances the public is not yet aware of. It is notable that Amazon Ring's relationship with law enforcement, including significant data-sharing agreements, was disclosed to the public by journalists following Freedom of Information Act (FOIA) requests, and not acknowledged by the company or police departments beforehand.

In other examples, a pretext of civic good is used to justify government use of private-sector technologies that ultimately serve to aid law enforcement. San Diego has installed thousands of microphones and cameras on street lamps in recent years. Marketed as an effort to study traffic and parking conditions, the data has ultimately proven to be of little use in improving traffic. Instead, the police took advantage of the infrastructure, using video footage from these traffic lights in more than 140 cases without any oversight or accountability.⁵¹ Similarly, the City of Miami is actively considering a 30-year contract with Illumination Technologies, providing the company with free access to set up light poles containing cameras and license-plate readers, collecting information that will filter through the Miami Police Department (and that the company can use in unchecked ways).⁵²

Commercial facial recognition systems deployed by private actors raise many of the same concerns as government use, especially since such systems are frequently used to inform meaningful decisions about people. Facial recognition systems are often applied in ways that presume government and law enforcement intervention. For example, a facial recognition system used by a private business to identify shoplifters assumes that a suspect will be turned over to the criminal justice system. Similarly, facial recognition used by a landlord to monitor tenants and enforce building rules, if marshalled as evidence supporting eviction, also presumes government intervention. In these cases, those using facial recognition are corporations or private actors, not a government agency. However, the harm of such use is no less real, and is likely to be prejudiced against traditionally disadvantaged populations.

⁵⁰ Michael Price & Emily Hockett, *Palantir Contract Dispute Exposes NYPD's Lack of Transparency*, Brennan Center for Justice (Jul. 20, 2017), <https://www.brennancenter.org/our-work/analysis-opinion/palantir-contract-dispute-exposes-nypds-lack-transparency>.

⁵¹ Joshua Emerson Smith, *As San Diego Increases Use of Streetlamp Cameras, ACLU Raises Surveillance Concerns*, Los Angeles Times, (Aug. 5, 2019), <https://www.latimes.com/california/story/2019-08-05/san-diego-police-ramp-up-use-of-streetlamp-cameras-to-crack-cases-privacy-groups-raise-concerns>.

⁵² Daniel Rivero, *Miami Could Let Company Put Surveillance Poles on Public Property for Free*, WLRN, October 9, 2019, <https://www.wlrn.org/post/miami-could-let-company-put-surveillance-poles-public-property-free>.

Affect recognition and facial analysis pose particular dangers

Many facial recognition systems also offer analysis capabilities, claiming to be able to detect gender, age, ethnicity, and other characteristics. Affect recognition is one type of facial analysis (which also extends beyond the face).⁵³ It claims to automatically detect a person's emotional state or inner qualities—from their personality, to their mental health, to whether or not they are competent, based on their physical appearance and mannerisms. Such systems are already being deployed widely, often alongside or as a component of facial recognition systems that identify and track individuals. These systems inform sensitive decisions that shape people's lives and access to resources, and they deserve particular scrutiny and rapid regulatory action.⁵⁴

The assertion that it's possible to determine a person's interior characteristics based on their facial expression through affect recognition is not backed by scientific consensus, and the technology reflects discredited pseudoscientific practices from the past, including physiognomy, phrenology, and race science, which interpreted physical differences between people as signs of their inner worth and used this to justify social inequality.⁵⁵ A comprehensive survey of over one thousand papers led by psychologist Lisa Feldman Barrett and a team of psychologists and engineers found that the claims made by affect recognition companies are not supported by the scientific literature on emotional expression. The authors conclude decisively that “no matter how sophisticated the computational algorithms . . . it is premature to use this technology to reach conclusions about what people feel on the basis of their facial movements.”⁵⁶

Beyond the lack of scientific foundation, affect recognition also encodes racial bias. Researcher Dr. Lauren Rhue found systematic racial biases in two well-known affect recognition programs: when she ran Face++ and Microsoft's Face API on a dataset of 400 NBA player photos, she found that both systems assigned Black players more negative emotional scores on average, no matter how much they smiled.⁵⁷

⁵³ Affect recognition can also include systems that analyze more than just facial expressions – for example, tone of voice and gait are also included in some affect recognition systems. For the purposes of this testimony, we focus on systems that draw on facial expression.

⁵⁴ Kate Crawford, Roel Dobbe, Theodora Dryer, Genevieve Fried, Ben Green, Elizabeth Kaziunas, Amba Kak, Varoon Mathur, Erin McElroy, Andrea Nill Sánchez, Deborah Raji, Joy Lisi Rankin, Rashida Richardson, Jason Schultz, Sarah Myers West, & Meredith Whittaker, *AI Now 2019 Report*, AI Now Inst. (2019), https://ainowinstitute.org/AI_Now_2019_Report.html.

⁵⁵ Stephen Jay Gould, *The Mismeasure of Man* (1981).

⁵⁶ Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M. Martinez, & Seth D. Pollak, *Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements*, *Psychological Science in the Public Interest*, 20(1), 1–68. <https://doi.org/10.1177/1529100619832930>.

⁵⁷ Lauren Rhue, *Racial Influence on Automated Perceptions of Emotions* (November 9, 2018), <https://ssrn.com/abstract=3281765>.

However, the evidence of biased inaccuracy and the lack of scientific foundation have not stalled the commercial deployment of affect recognition systems, and the industry is predicted to grow to over \$90 billion by 2024.⁵⁸ This technology is already being used to make sensitive determinations that are shaping people's lives, from deciding whether a job candidate will be a good worker,⁵⁹ to assessing whether a patient in medical care is in pain,⁶⁰ to detecting shoplifters before they steal,⁶¹ to tracking whether students in the classroom are attentive⁶² (ignoring studies that showed significant risks associated with the deployment of emotional AI in the classroom).⁶³ Tech companies—including Amazon,⁶⁴ Microsoft,⁶⁵ Affectiva,⁶⁶ Noldus,⁶⁷ Kairos,⁶⁸ and Sightcorp,⁶⁹ to name a handful—continue to sell affect recognition as part of their facial recognition offerings. Many third parties license these features from these companies and apply them in ways that aren't transparent to the public.

The example of the AI company HireVue is instructive. The company licenses affect recognition technology from Affectiva⁷⁰ and sells AI video-interviewing systems to large firms like Goldman Sachs and Unilever, marketing its system as capable of determining which job candidates will be successful workers and which won't based on a remote video interview. HireVue uses affect

⁵⁸ Paul Sawers, *Realeyes Raises \$12.4 Million to Help Brands Detect Emotion Using AI on Facial Expressions*, VentureBeat (June 6, 2019), <https://venturebeat.com/2019/06/06/realeyes-raises-12-4-million-to-help-brands-detect-emotion-using-ai-on-facial-expressions>.

⁵⁹ Drew Harwell, *Rights Group Files Federal Complaint against AI-Hiring Firm HireVue, Citing 'Unfair and Deceptive' Practices*, Washington Post (Nov. 6, 2019), <https://www.washingtonpost.com/technology/2019/11/06/prominent-rights-group-files-federal-complaint-against-ai-hiring-firm-hirevue-citing-unfair-deceptive-practices>.

⁶⁰ Clarice Smith, *Facial Recognition Enters into Healthcare*, Journal of AHIMA (Sept. 4, 2018), <https://journal.ahima.org/2018/09/04/facial-recognition-enters-into-healthcare>.

⁶¹ Lisa Du & Ayaka Maki, *These Cameras Can Spot Shoplifters Even Before They Steal*, Bloomberg (Mar. 4, 2019), <https://www.bloomberg.com/news/articles/2019-03-04/the-ai-cameras-that-can-spot-shoplifters-even-before-they-steal>.

⁶² Mark Lieberman, *I Know How You Felt This Semester*, Inside Higher Ed (Feb. 20, 2018), <https://www.insidehighered.com/digital-learning/article/2018/02/20/sentiment-analysis-allows-instructors-to-hape-course-content>.

⁶³ Andrew McStay, *Emotional AI and EdTech: Serving the Public Good?*, Learning, Media and Technology (Nov. 5, 2019), <https://doi.org/10.1080/17439884.2020.1686016>.

⁶⁴ Tom Simonite, *Amazon Says It Can Detect Fear on Your Face. Are You Scared?*, Wired (Aug. 18, 2019), <https://www.wired.com/story/amazon-detect-fear-face-you-scared>.

⁶⁵ Microsoft Azure, *Cognitive Services: Face*, <https://azure.microsoft.com/en-us/services/cognitive-services/face> (last visited: Jan. 13, 2020).

⁶⁶ Affectiva, *Emotion AI Overview*, <https://www.affectiva.com/emotion-ai-overview> (last visited Jan. 13, 2020).

⁶⁷ Noldus, *Emotion Analysis: FaceReader*, <https://www.noldus.com/facereader> (last visited Jan. 13, 2020).

⁶⁸ Luana Pascu, *New Kairos Facial Recognition Camera Offers Customer Insights*, Biometric Update, (Sept. 11, 2019), <https://www.biometricupdate.com/201909/new-kairos-facial-recognition-camera-offers-customer-insights>.

⁶⁹ F.A.C.E. API by Sightcorp, <https://face-api.sightcorp.com> (last visited Jan. 13, 2020).

⁷⁰ Ria Lupton, *Affectiva CEO Rana El Kaliouby Shares Applications for Emotion AI at True North*, BetaKit (Jun. 7, 2018), <https://betakit.com/affectiva-ceo-rana-el-kaliouby-shares-applications-for-emotion-ai>.

recognition to analyze these videos, examining facial movements, speech patterns, tone of voice, and other indicators.⁷¹ Based on these factors, in combination with other assessments, the system makes recommendations about who should be scheduled for a follow-up interview, and who should not get the job. HireVue's training data is selected from video of existing workers who have been deemed successful at a given firm.⁷² This implies that people who look and behave like those already hired and promoted are more likely to be selected. The potential for encoding and automating existing biases is clear. In a report examining HireVue and similar tools, authors Jim Fruchterman and Joan Mellea are blunt about the implications of such bias for disabled people. "[HireVue's] method massively discriminates against many people with disabilities that significantly affect facial expression and voice: disabilities such as deafness, blindness, speech disorders, and surviving a stroke," they write.⁷³

In addition to affect recognition, facial recognition systems are using facial analysis to catalog and determine peoples' identities and attributes based on their faces, including estimating age, ethnicity, gender, and more. Such methods can also be harmful. Microsoft, Amazon, and (until recently) IBM all offer facial recognition services that include the option to classify people's gender as either male or female based on an image of their face. Such features not only disregard the fluid nature of gender identity, but potentially endanger people who don't "fit" one or another binary gender category.^{74,75} Research has shown that gender classification systems persistently misclassify transgender people, and fail to identify non-binary people.⁷⁶

In the same vein, a much-maligned 2016 paper claimed to be able to determine sexual orientation based on a facial image.⁷⁷ While the claims made by the paper were roundly rebuked, the publication of a model making such claims still posed significant danger, especially given that being gay is illegal in at least 71 countries.⁷⁸ Other researchers applied the same

⁷¹ HireVue, <https://www.hirevue.com> (last visited Jan. 13, 2020).

⁷² Richard Feloni, *I Tried the Software That Uses AI to Scan Job Applicants for Companies Like Goldman Sachs and Unilever Before Meeting Them*, Business Insider (Aug. 23, 2017), <https://www.businessinsider.com/hirevue-ai-powered-job-interview-platform-2017-8#larsen-showed-me-what-a-recruiter-would-see-when-analyzing-my-answers-8>.

⁷³ Jim Fruchterman and Joan Mellea, *Expanding Employment Success for People with Disabilities*, Benetech (Nov. 2018), <https://benetech.org/about/resources/expanding-employment-success-for-people-with-disabilities-2>.

⁷⁴ Foad Hamidi, Morgan Klaus Scheuerman, & Stacy M Branham, *Gender Recognition or Gender Reductionism?: The Social Implications of Embedded Gender Recognition Systems*, Proceedings of the 2018 CHI Conf. on Human Factors in Computing Systems (Apr. 2018), <https://dl.acm.org/doi/10.1145/3173574.3173582>.

⁷⁵ Rachel Metz, *AI Software Defines People as Male or Female. That's a Problem*, CNN Business (Nov. 21, 2019) <https://www.cnn.com/2019/11/21/tech/ai-gender-recognition-problem/index.html>.

⁷⁶ Morgan Klaus Scheuerman et al., *How Computers See Gender*, *supra* note 23.

⁷⁷ Michal Kosinski & Yilun Wang, *Deep Neural Networks Are More Accurate Than Humans at Detecting Sexual Orientation From Facial Images*, *Journal of Personality and Social Psychology*, 114:2, 246 (Feb. 2018), <https://osf.io/zn79k>.

⁷⁸ Daniel Avery, *71 Countries Where Homosexuality Is Illegal*, Newsweek (Apr. 4, 2019) <https://www.newsweek.com/73-countries-where-its-illegal-be-gay-1385974>.

flawed logic, claiming to have developed AI models that could detect criminality based on a person's face.⁷⁹

Affect recognition and similar facial-analysis technologies function to classify and catalog people in ways that have significant consequences. They place the authority to determine a person's interior characteristics and identity in the hands of technology that is not only scientifically unfounded, but often used by those with power to inform significant judgements about people in more vulnerable positions. How someone might contest an automated assessment about their feelings, their worth, or their character remains an open question.

Any regulation of facial recognition must be sure to address affect recognition and similar systems that claim to be able to catalog and read people's identities and interior states based on automated detection of physical features.

Standards and technical fixes aren't enough to solve the problems with facial recognition

With mounting evidence of facial recognition's inaccuracy and failure, researchers and companies have worked to "debias" facial recognition, focusing on technical fixes and standards for testing and validation in an attempt to ensure accuracy and fairness.^{80,81} Recent legislation has also called for standardized auditing and assessment criteria for facial recognition and other AI technologies.⁸² Such standards can be helpful in setting criteria, giving the government, industry, and the AI field systematic approaches to determine whether or not a given system can be developed, sold, applied to one or another use case, or procured for government contracts.

While this is a step in the right direction, these approaches are not enough on their own. If they are implemented without care, they could do more harm than good.

AI systems, including facial recognition, model the world based on the data they're trained on during their development. Training data is at the core of how AI systems, including facial recognition, recognize and understand the world.⁸³ If a population is omitted from the data used to develop a model—by excluding images of people with darker skin, for example—then these people will be missing from the AI model's representation of the world. The excluded group therefore won't be recognized in the resulting system.

⁷⁹ Xiaolin Wu & Xi Zhang, Shanghai Jiao Tong University, *Automated Inference on Criminality using Face Images* (Nov. 13, 2016), <https://arxiv.org/pdf/1611.04135v1.pdf>.

⁸⁰ Inioluwa Deborah Raji & Joy Buolamwini, *Actionable Auditing*, *supra* note 19.

⁸¹ Facial Identification Scientific Working Group, <https://fiswg.org/index.htm> (last visited Jan. 13, 2020).

⁸² Commercial Facial Recognition Privacy Act of 2019, S.847, *supra* note 42.

⁸³ Kate Crawford & Trevor Paglen, *Excavating AI: The Politics of Images in Machine Learning Training Sets* (September 19, 2019), <https://www.excavating.ai>.

Standards for measuring performance and accuracy generally work by running an AI model against a standardized test dataset, called a *benchmarking dataset*, and measuring its performance for a given task. For instance, a facial recognition system could be tested on the task of one-to-one facial matching, or one-to-many matching, or on an analysis task such as gender identification. The system's performance on a given task is measured against a designated test dataset in order to understand how well the system works for that task. If a test dataset does not reflect the conditions, demographics, and environment where a facial recognition system will be deployed, measurements of performance using this dataset become meaningless, failing to account for real-world conditions in any informative way.

In assessing the advantages and limits of assessment standards, it is critical to examine the test benchmarking datasets they rely on. Benchmarking datasets act as ground truth against which researchers and developers compare systems, and thus they work to define—and misdefine—criteria like fairness, accuracy, and performance.

For example, Labeled Faces in the Wild (LFW) is a canonical facial recognition benchmarking dataset that helped shape the field of machine vision, and facial recognition in particular, by setting the standard against which researchers measured the accuracy of their systems. It consists of over 13,000 labeled images scraped from Yahoo News between 2002 and 2004, picturing celebrities, power players, and the newsworthy.⁸⁴ Many developers have worked to “improve” the performance of their systems on LFW, citing the successful performance of their systems on the dataset to buttress claims about their accuracy. Yet according to researchers Hu Han and Anil K. Jain, the diversity of Labeled Faces in the Wild is limited: 77 percent of the images feature male faces, 81 percent of the images show light-skinned people, and very few images contain children or elderly people.⁸⁵ Thus, during a critical period in its recent development, the AI field's understanding of facial recognition performance was largely based on whether it accurately recognized mainly white men, and this was the goal researchers and companies optimized for. With this in mind, the persistent racial and gender bias across facial recognition systems should come as no surprise.

LFW is not the only benchmarking dataset whose contents and history require attention. Current benchmarks also fall short of capturing the representation required for a reliable assessment of a model's performance upon release. Researchers Inioluwa Deborah Raji and Genevieve Fried surveyed over 100 facial recognition benchmarking datasets and found “dissonance between the perceived functionality of these systems under current evaluation norms and the reality of

⁸⁴ Gary B. Huang, et al., *Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments* (2008), <http://vis-www.cs.umass.edu/papers/lfw.pdf>.

⁸⁵ Hu Han & Anil K. Jain, *Age, Gender and Race Estimation from Unconstrained Face Images*, Michigan State University Technical Report (2014), http://biometrics.cse.msu.edu/Publications/Face/HanJain_UnconstrainedAgeGenderRaceEstimation_MSUTechReport2014.pdf

their performance when deployed.”⁸⁶ In other words, many of the systems that “pass” current benchmark evaluations continue to underperform in real-life contexts. Additionally, there is currently no standard practice to document and communicate the histories and limits of benchmarking datasets, and thus no way to determine their applicability to a particular system or suitability for a given context.⁸⁷

While limited and non-diverse benchmarking datasets fail to accurately measure facial recognition performance and harm, the practice of creating more diverse face datasets raises significant ethical and privacy questions. Creating such datasets requires the collection of additional face data, often from populations who have historical reasons to be wary of such efforts and who may not want their images used to develop surveillance technology.⁸⁸ Such efforts can violate privacy and lead to the tokenization of those included in these datasets, amplifying stereotypes and serving to make people visible to technical systems that work to harm their communities.⁸⁹

To obtain data, companies and researchers have a history of bypassing meaningful consent, scraping data from Google Image Search,⁹⁰ YouTube,⁹¹ Flickr,^{92,93} Wikipedia,⁹⁴ and even mug-shot databases.⁹⁵ Some data collection methods border on exploitation. For example, last year, journalists revealed that Google was offering “darker skinned” unhoused people five dollars in exchange for their face data. According to one staffer working on this project, the team gathering the data was instructed to target the unhoused “because they’re the least likely to say anything to the media.”⁹⁶ A number of datasets also use surveillance footage without consent.⁹⁷

⁸⁶ Inioluwa Deborah Raji & Genevieve Fried, *About Face: A Survey of Facial Recognition Evaluation*, Meta-Evaluation workshop at AAAI Conf. on Artificial Intelligence (Forthcoming 2020).

⁸⁷ Efforts like Datasheets, model cards, and fact sheets represent attempts to develop such standards, but they are currently prototypes, and have not been adopted widely within the AI field.

⁸⁸ Inioluwa Deborah Raji, et al., *Saving Face*, *supra* note 23.

⁸⁹ Anna Lauren Hoffmann, *Where Fairness Fails: Data, Algorithms, and the Limits of Antidiscrimination Discourse*, *Information, Communication & Society*, 22:7, 900-915 (2019), <https://www.tandfonline.com/doi/full/10.1080/1369118X.2019.1573912>.

⁹⁰ Wilma A. Bainbridge, Phillip Isola, & Aude Oliva, *The Intrinsic Memorability of Face Photographs*, *Journal of Experimental Psychology: General*, 142(4), 1323–1334 (2013) <https://doi.org/10.1037/a0033872>.

⁹¹ Lior Wolf, Tal Hassner, & Itay Maoz, *Face Recognition in Unconstrained Videos With Matched Background Similarity*, IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (July 2011), <https://ieeexplore.ieee.org/document/5995566>.

⁹² Inioluwa Deborah Raji, et al., *Saving Face*, *supra* note 23.

⁹³ Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, & Evan Brossard. *The MegaFace Benchmark: 1 Million Faces for Recognition at Scale*, Intl. Conf. on Computer Vision and Pattern Recognition (2016), <https://arxiv.org/abs/1512.00596>.

⁹⁴ Rasmus Rothe, Radu Timofte, & Luc Van Gool, *Deep Expectation of Real and Apparent Age From a Single Image Without Facial Landmarks*, *International Journal of Computer Vision*, 126(2-4): 144-57 (2018), <https://data.vision.ee.ethz.ch/cvl/rrothe/imdb-wiki>.

⁹⁵ Olivia Solon, *Facial Recognition's 'Dirty Little Secret'*, *supra* note 8.

⁹⁶ Ginger Adams Otis & Nancy Dillon, *Google Using Dubious Tactics to Target People With 'Darker Skin' in Facial Recognition Project: Sources*, N.Y. Daily News (Oct. 2, 2019),

Given the significant improvement in camera technology optimized to enable video tracking and data capture, such privacy violations are even more likely in the future.⁹⁸ Additionally, certain unregulated partnerships can lead to biometric data collected for one purpose being repurposed in exploitative ways by corporations and governments. For instance, the Chinese facial recognition company CloudWalk Technology will provide the Zimbabwe government with a massive facial recognition program in exchange for the face data of Zimbabweans,⁹⁹ and the FBI and ICE were discovered to have made use of face data from the DMV as well as local databases in order to target and identify individuals.¹⁰⁰

The way in which identity and ethnicity is categorized within “diverse” datasets can also raise problems. Such datasets usually treat race and other attributes as fixed and visually recognizable. The people whose data is included in these datasets rarely have the opportunity to self-identify, and assumptions, stereotypes, and even facial measurements are used to assign people to identity categories that don’t usually account for multifaceted identities (Black women or Latinx transgender women, for example). When these datasets are used as benchmarks against which bias and accuracy are assessed, they will inevitably provide an incomplete measure, excluding people whose identities are not represented.¹⁰¹ Even if it were desirable from a privacy and ethics standpoint, there are serious questions about whether it is even possible to “include” a comprehensive set of interlocking identities such that a given dataset could truly ensure fair and accurate results for everyone. This is most evident in the case of disability, which includes a wide array of physical and mental-health conditions that may come and go within the course of a lifetime, or even a day, meaning that “simply expanding a dataset’s parameters to include new categories, in an attempt to account for ‘disability,’ won’t work to ensure disabled people are represented.”¹⁰²

<https://www.nydailynews.com/news/national/ny-google-darker-skin-tones-facial-recognition-pixel-20191002-5vxpgowknffnfbmy5eg7epsf34-story.html>.

⁹⁷ Cade Metz, *Facial Recognition Tech Is Growing Stronger, Thanks to Your Face*, N.Y. Times (Jul. 13, 2019),

<https://www.nytimes.com/2019/07/13/technology/databases-faces-facial-recognition-technology.html>; see also Brainwash Dataset, <https://megapixels.cc/brainwash> (last visited Jan. 13, 2020).

⁹⁸ Jay Stanley, *The Dawn of Robot Surveillance: AI, Video Analytics, and Privacy*, ACLU (2019), https://www.aclu.org/sites/default/files/field_document/061119-robot_surveillance.pdf.

⁹⁹ Amy Hawkins, *Beijing’s Big Brother Tech Needs African Faces*, Foreign Policy (Jul. 24, 2018), <https://foreignpolicy.com/2018/07/24/beijings-big-brother-tech-needs-african-faces>.

¹⁰⁰ Drew Harwell, *FBI, ICE Find State Driver’s License Photos Are a Gold Mine for Facial-Recognition Searches*, Washington Post (Jul. 7, 2019), <https://www.washingtonpost.com/technology/2019/07/07/fbi-ice-find-state-drivers-license-photos-are-gold-mine-facial-recognition-searches>.

¹⁰¹ Alex Hanna, Emily Denton, Andrew Smart, & Jamila Smith-Loud, *Towards a Critical Race Methodology in Algorithmic Fairness*, Conference on Fairness, Accountability, and Transparency (2020), <https://doi.org/10.1145/3351095.3372826>.

¹⁰² Meredith Whittaker, Meryl Alper, Cynthia L. Bennett, Sara Hendren, Liz Kaziunas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, & Sarah Myers West, *Disability, Bias, and AI*, AI Now Inst. (Nov. 2019), <https://ainowinstitute.org/disabilitybiasai-2019.pdf>.

Standardized assessment (or audit) protocols are also limited in scope. Only the systems, populations, and tasks that are explicitly tested will be scrutinized, which will inevitably fail to account for many important questions and potential harms that require attention.¹⁰³ For example, an assessment focused on whether a facial recognition system performs equally well across population subgroups defined by gender will not identify disparities in classification across race, ability, or age, which may be more relevant depending on the context in which the system will be applied.

The focus on addressing bias and justice concerns through technical standards and testing may also distract from other issues. While such methods can provide researchers, regulators, and the public with important information, they are insufficient to ensure the safe deployment of a facial recognition system. Current auditing standards rarely include the qualitative considerations necessary to properly evaluate the technology.¹⁰⁴ Furthermore, the communities who will bear the risk of deployment and the civil-society groups fighting for their interests are often not consulted and included in defining an assessment process that addresses their concerns. Such gaps led researchers from Google, MIT, and the University of Toronto to conclude recently that, while such standards may improve the visibility of certain failure modes of these systems, “well intentioned attempts at algorithmic auditing can have effects that may harm the very populations these measures are meant to protect.”¹⁰⁵

More research is needed to develop better ways to evaluate these systems, taking into account the need to look beyond accuracy metrics and toward a more holistic view of the technology’s risks. Funding such efforts should be a priority, and barriers to such work, from trade secrecy to law enforcement transparency exemptions, need to be lifted to ensure democratic oversight.

Standardized approaches to measuring and assessing AI systems including facial recognition represent a step in the right direction, but they only speak to a limited set of concerns, and we cannot rely on them to steer important decisions on facial recognition’s use. If we depend too much on narrow or weak standards, we run the risk of providing “checkbox certification,” allowing vendors and companies to assert that their technology is safe and fair without accounting for how it will be used, or its fitness for a given context. If such standards are positioned as the sole check on such systems, they could function to obfuscate harm instead of mitigate it.

¹⁰³ Inioluwa Deborah Raji, et al., *Saving Face*, *supra* note 23.

¹⁰⁴ Mei Ngan & Patrick Grother, NIST Interagency Report 8052, *supra* note 17.

¹⁰⁵ Inioluwa Deborah Raji, et al., *Saving Face*, *supra* note 23.

It is time to halt the use of facial recognition in sensitive social and political contexts, by both government and private actors

Facial recognition is a technology that, once deployed, is very difficult to dismantle. Therefore, we must be extremely cautious about allowing its use in any context. Given that we are still in the early days of research on its impacts, the general lack of transparency and accountability for its use, and the significant risks it poses, the best approach to protecting the public is to put a halt to facial recognition, by both government and private actors, in sensitive social and political contexts such as criminal justice, healthcare, education, employment, and use of public space. Harms in these contexts are nearly impossible to remedy, especially when the harm is community-wide. Various cities, counties, and states across the US have already demonstrated strong leadership by taking these steps. It is now time for the federal government to follow suit.

For example, one of the few mechanisms currently in place to protect the public is exemplified by laws like Illinois' Biometric Information Privacy Act (BIPA), which currently allows individuals to sue companies for nonconsensual commercial facial recognition. Such approaches should be adopted more widely and expanded to include the right to sue government misuse. But litigation alone cannot address the problems with technologies like facial recognition. First, bringing a case requires not only evidence of misconduct, but also proof that the technology was used in the first place. Both of these are often difficult due to the obscurity in which such systems are deployed and the corporate secrecy that prevents public research and scrutiny. Litigation also requires the resources to pursue a case, which many of those likely to be harmed or targeted don't have.

Notice and consent meant to ensure that those subject to facial recognition are aware, and agree to its use, is also not feasible. Not only are the typical online notices rarely legible, but most users lack the power to decline, especially when few alternatives exist. Companies like Facebook also routinely ignore their own policies and break their promises.¹⁰⁶ And increasingly, facial recognition is being applied in contexts where non-consent would bar people from access to public space and opportunity, such as in airports, concert venues, and schools. Beyond this, it's unclear how such consent could work in practice, given the current applications and infrastructures underlying these technologies. How would someone opt out of, or opt in to, facial recognition used in retail establishments, airports, smart-city infrastructure, and other so-called smart environments, and what fundamental changes would enabling meaningful opt-out require?¹⁰⁷ How can someone be sure that biometric face data has not been processed or

¹⁰⁶ Evan Selinger, *Why You Can't Really Consent to Facebook's Facial Recognition*, *supra* note 49.

¹⁰⁷ Wojciech Wiewiórowski, European Data Protection Supervisor, *Facial Recognition: A Solution in Search of a Problem?* (Oct. 28, 2019),

https://edps.europa.eu/press-publications/press-news/blog/facial-recognition-solution-search-problem_en.

collected by such a system, and thus that their right to opt out has been respected? We don't have viable answers to these questions.

AI companies have also turned to voluntary AI principles and ethical statements that commit them to develop and deploy AI, including facial recognition, in beneficial ways. Microsoft's facial recognition principles assert that the company will "not deploy facial recognition technology in scenarios that we believe will put these freedoms at risk."¹⁰⁸ Google's AI Principles promise that Google will not develop "technologies whose purpose contravenes widely accepted principles of international law and human rights."¹⁰⁹ In 2018, Google cited these principles when it announced that it would be putting a pause on its plans to launch facial recognition products.¹¹⁰ Axon, the largest manufacturer of police body cameras, also made the choice to halt deployment of facial recognition, stating that "face recognition technology is not currently reliable enough to ethically justify its use."¹¹¹

Voluntary ethical principles and statements are a positive step. They acknowledge the problem and provide a rough standard by which to assess an organization's conduct, and in some cases they guide decision-making. But we cannot count on the AI industry to self-regulate. Ethical principles are not enough to address the serious risks of facial recognition. They fail to ensure accountability, and they allow companies to announce their commitment to beneficial conduct without submitting to regulation, oversight, or accountability to the communities who are harmed by their technologies. A pattern of decision-making at these companies, which includes pursuing investments and projects that contradict their own principles,^{112,113} also reveals that many AI firms choose revenue and growth over accountability.

Significant research is necessary to answer questions on whether this technology can be used in a way that is safe and fair, and we need to leave room for "no" as an answer to these critical questions. Such research requires access to private infrastructures, data, and

¹⁰⁸ Rich Sauer, Microsoft, *Six Principles to Guide Microsoft's Facial Recognition Work*, Microsoft On the Issues (Dec. 17, 2018), <https://blogs.microsoft.com/on-the-issues/2018/12/17/six-principles-to-guide-microsofts-facial-recognition-work>.

¹⁰⁹ Sundar Pichai, Google, *AI at Google: Our Principles*, The Keyword, June 7, 2018, <https://www.blog.google/technology/ai/ai-principles>.

¹¹⁰ Will Knight, *Google Says It Won't Sell Face Recognition for Now—But It Will Be Hard to Slow Its Use*, MIT Technology Review (Dec. 14, 2018), <https://www.technologyreview.com/f/612606/google-will-stop-providing-face-recognition-but-it-will-be-hard-to-curb-its-use>.

¹¹¹ First Report of the Axon AI & Policing Technology Ethics Board (June 2019), https://www.policingproject.org/s/Axon_Ethics_Report_vfinal-English.pdf.

¹¹² Olivia Solon, *Microsoft Hires Eric Holder to Audit AnyVision Over Use of Facial Recognition on Palestinians*, NBC News (Nov. 15, 2019), <https://www.nbcnews.com/tech/security/microsoft-hires-eric-holder-audit-anyvision-over-use-facial-recognition-n1083911>.

¹¹³ Alexia Fernández Campbell, *The Employee Backlash Over Google's Censored Search Engine for China, Explained*, Vox (Aug. 17, 2018), <https://www.vox.com/2018/8/17/17704526/google-dragonfly-censored-search-engine-china>.

documentation that is currently unavailable to all but the people employed by companies that produce these systems. Similarly, well-resourced enforcement regimes would need to be constructed across state and federal agencies, in ways that ensure the communities on whom facial recognition is used have meaningful opportunities to review and reject its use.

Over the past year, there has been growing pushback against facial recognition, much of it organized by community groups resisting the deployment of these technologies in their everyday lives.¹¹⁴ This grassroots work has led to a number of bans and moratoriums.

San Francisco was the first to pass a ban on government use of facial recognition. It is significant that in Silicon Valley's backyard, the people who build these systems, and who understand their capabilities and limitations, didn't feel comfortable having them used in their own communities. Indeed, tech workers and Amazon shareholders, among others close to this technology, have joined the call to halt the sale of facial recognition for government surveillance.^{115,116,117} The cities of Somerville,¹¹⁸ Oakland,¹¹⁹ Berkeley,¹²⁰ Brookline,¹²¹ and most recently San Diego¹²² joined San Francisco and passed their own bans and moratoriums on government use.

¹¹⁴ Kate Crawford & Meredith Whittaker, *AI in 2019: A Year in Review*, AI Now Inst. (Oct. 9, 2019), <https://medium.com/@AINowInstitute/ai-in-2019-a-year-in-review-c1eba5107127>.

¹¹⁵ Alexa Lardieri, *Amazon Employees Protesting Sale of Facial Recognition Software*, U.S. News and World Report (Oct. 18, 2018), <https://www.usnews.com/news/politics/articles/2018-10-18/amazon-employees-protesting-sale-of-facial-recognition-software>.

¹¹⁶ ACLU, et al., *Letter From Nationwide Coalition to Amazon CEO Jeff Bezos Regarding Rekognition*, ACLU (Jun. 18, 2020), <https://www.aclu.org/letter-nationwide-coalition-amazon-ceo-jeff-bezos-regarding-rekognition>.

¹¹⁷ Brian Brackeen, *Facial Recognition Software Is Not Ready for Use by Law Enforcement*, TechCrunch (Jun. 25, 2018), <https://techcrunch.com/2018/06/25/facial-recognition-software-is-not-ready-for-use-by-law-enforcement>.

¹¹⁸ Sarah Wu, *Somerville City Council Passes Facial Recognition Ban*, Boston Globe (Jun. 27, 2019), <https://www.bostonglobe.com/metro/2019/06/27/somerville-city-council-passes-facial-recognition-ban/SfaqQ7mG3DGulXonBHSCYK/story.html>.

¹¹⁹ Sarah Ravani, *Oakland Bans Use of Facial Recognition Technology, Citing Bias Concerns*, San Francisco Chronicle (Jul. 17, 2019), <https://www.sfchronicle.com/bayarea/article/Oakland-bans-use-of-facial-recognition-14101253.php>.

¹²⁰ Sara Merken, *Berkeley Bans Government Face Recognition Use, Joining Other Cities*, Bloomberg Law (Oct. 16, 2019), <https://news.bloomberglaw.com/privacy-and-data-security/hold-berkeley-bans-government-face-recognition-use-joining-other-cities>.

¹²¹ Tori Bedford, *Brookline Votes To Ban Face Surveillance Technology*, WGBH News (Dec. 11, 2019), <https://www.wgbh.org/news/local-news/2019/12/11/brookline-votes-to-ban-face-surveillance-technology>; see also *Article 25: Explanation Filed by Petitioners*, Special Town Meeting of Brookline, Massachusetts, pp. 44-47, (Nov. 19, 2019), available at <https://www.brooklinema.gov/DocumentCenter/View/20115/Article-Explanations-2019-STM>.

¹²² Katy Stegall, *3-Year Ban on Police Use of Facial Recognition Technology in California to Start in the New Year*, San Diego Union Trib. (Dec. 20, 2019), <https://www.sandiegouniontribune.com/news/public-safety/story/2019-12-20/3-year-ban-on-police-use-of-facial-recognition-technology-in-california-to-start-in-the-new-year>.

And Portland, Oregon, is considering the strongest ban yet, which would limit both commercial and governmental deployment.¹²³

Evidence shows that when communities are informed about the flaws and risks of facial recognition, they move to reject its use. Lawmakers should protect the public interest and heed communities' wishes, putting a halt to deployment by both government and the private sector until the risks are fully studied and adequate regulations are in place.

Recommendations for the path forward

- **Halt both governmental and commercial use of facial recognition in sensitive social and political contexts until the risks are fully studied and adequate regulations are in place.** In 2019, there has been a rapid expansion of facial recognition in many domains. Yet there is mounting evidence that this technology causes serious harm, most often to people of color and the poor, and none of the current technical mitigation methods adequately addresses these concerns. There should be a moratorium on all uses of facial recognition in sensitive social and political domains—including surveillance, policing, education, and employment—where facial recognition poses risks and consequences that cannot be adequately remedied retroactively.
- **Ban the use of affect recognition in important decisions that impact people's lives and access to opportunities.** Until then, AI companies should stop deploying it. Given the contested scientific foundations of affect recognition technology—a related class of systems that claim to detect things such as personality, emotions, mental health, and other interior states—it should not be allowed to play a role in important decisions about human lives, such as who is interviewed or hired for a job, the price of insurance, patient pain assessments, or student performance in school. This ban should be accompanied with federally funded research on the adequacy of existing laws and regulations to address these concerns.
- **Apply “truth-in-advertising” laws to AI products and services, including facial recognition.** The hype around AI is only growing, leading to widening gaps between marketing promises and actual product performance. Researchers and lawmakers struggle to measure and understand these gaps due to trade secrecy and other barriers that prevent access to vital information about these

¹²³ Sean Captain, *Portland Plans to Propose the Strictest Facial Recognition Ban in the Country*, Fast Company (Dec. 2, 2019), <https://www.fastcompany.com/90436355/portlands-proposed-facial-recognition-ban-could-be-the-strictest-yet>.

systems. With these gaps come increasing risks to both individuals and commercial customers, often with grave consequences. Much like other products and services that have the potential to seriously impact or exploit populations, AI companies should be held to high standards for what they can promise, especially when the scientific evidence to back these promises is inadequate and the longer-term consequences are unknown.

- **Craft expanded biometric privacy laws that regulate both public and private actors.** Biometric data, from DNA to face data, is at the core of many harmful AI systems, including facial recognition. Over a decade ago, Illinois adopted the Biometric Information Privacy Act (BIPA), which has now become one of the strongest and most effective privacy protections in the United States. BIPA allows individuals to sue for almost any unauthorized collection and use of their biometric data by a private actor, including for surveillance, tracking, and profiling via facial recognition. BIPA also shuts down the gray and black markets that sell data, making it vulnerable to breaches and exploitation. States that adopt BIPA should expand it to include government use, which will mitigate many of biometric AI's harms, especially in parallel with other approaches like moratoriums and prohibitions.
- **Require technology companies to waive trade secrecy and other legal claims that hinder oversight and accountability mechanisms.** Corporate secrecy laws are a barrier to oversight, accountability, and due process when they are relied on to obscure technologies used in ways that affect the public. They can inhibit necessary government oversight and enforcement of consumer protection laws, thus contributing to the "black box effect" that makes it hard to assess bias, contest decisions, or remedy errors. Anyone procuring these technologies for use in the public sector should have the right to demand that vendors waive these claims before entering into any agreements. Additionally, limiting the use of these legal claims across the board will help facilitate better oversight by state and federal consumer-protection agencies and enforcement of false and deceptive practice laws.
- **Require algorithmic impact assessments in both public and private sectors, and establish frameworks that ensure the communities on whom facial recognition and other AI technologies are used have decision-making power about how, and whether, these technologies are applied.** When communities have information about the use of facial recognition and similar technologies, they often act to stop it, showing that the interests of those applying these systems are not always aligned with the desires of those on whom the systems are being used. Such frameworks should also give communities the ability to audit and interrogate the systems.