MIT Case Studies in Social and Ethical Responsibilities of Computing

The Bias in the Machine: Facial Recognition Technology and Racial Disparities

Sidney Perkowitz¹

¹Charles Howard Candler Emeritus Professor of Physics, Emory University

Published on: Feb 05, 2021

License: [Creative Commons Attribution-NonCommercial 4.0 International License \(CC-BY-NC 4.0\)](https://creativecommons.org/licenses/by-nc/4.0/)

ABSTRACT

Facial recognition technology (FRT) appears in uses from providing secure access to smartphones, to identifying criminal suspects from surveillance images as a tool of the justice system. Citizen's rights and social justice groups, alongside the research community, have identified undesirable societal consequences arising from the uncritical use of FRT algorithms, including false arrest and excessive government surveillance. Within the United States, these consequences have disproportionately affected people of color, both because algorithms have typically been less accurate when applied to nonwhite people, and because—like any new forensic technology—FRT systems are being incorporated into systems and institutions with their own histories of disparities. As research continues to address racial biases in the performance of FRTs, examples from older forensic technologies, such as fingerprint identification, might offer insights into improving real-world deployment.

Sidney Perkowitz

Charles Howard Candler Emeritus Professor of Physics, Emory University

Keywords: *facial recognition, justice system, racial equity, false arrest*

Author Disclosure(s): Portions of this case study originally appeared in Sidney Perkowitz, “The Bias in the Machine: Why Facial Recognition Has Led to False Arrests,” *Nautilus*, August 19, 2020, and are reproduced with permission.

In January 2020, Robert Williams, a Black man, was wrongfully arrested due to an inaccurate facial recognition algorithm, a computerized approach that analyzes human faces and identifies them by comparison to database images of known people. Williams was handcuffed and arrested in front of his family by Detroit police without being told why, then jailed overnight after the police took mugshots, fingerprints, and a DNA sample.

The next day, detectives showed Williams a surveillance video image of a Black man standing in Shinola, a store that sells expensive watches. It immediately became clear that Williams was not the person shown in the video. Detailing his arrest in the [Washington Post](#), Williams wrote, “The cops looked at each other. I heard one say that ‘the computer must have gotten it wrong.’”¹ Williams learned

that in investigating a theft from the store, a facial recognition system had tagged his driver's license photo as matching an image from the surveillance video. The error from the algorithm was then compounded when police relied upon the assertion of a private investigator that the algorithm had correctly matched the photos. Rather than following appropriate procedures—in which police must first confirm a machine match, then seek additional evidence sufficient for an arrest—they simply arrested Williams. Only after detectives saw Williams in person and compared him to the surveillance image did they see the differences, and the case against him collapsed. But those steps came *after* Williams had spent thirty hours in jail and posted a \$1,000 bond.²

The Williams arrest was briefly unique because it received public attention. As the American Civil Liberties Union (ACLU) observed, over four thousand police departments across the United States use facial recognition systems, making it highly likely that other people have been wrongly implicated in crimes.³ Indeed, soon several similar cases came to light. In February 2019, another Black man, Nijeer Parks, was accused of shoplifting and attacking a police officer in Woodbridge, New Jersey. Facial recognition software had identified him as a suspect from a driver's license photo. After he spent ten days in jail and paid \$5,000 in legal fees to defend himself, unassailable evidence showed that he had been miles away at the time the crime occurred.⁴ Meanwhile, in July 2019, a facial recognition match sent Detroit police to arrest a third Black man, Michael Oliver. The charge, a felony count of larceny, was dropped when a judge agreed that he had been misidentified, as shown by mismatched tattoos.⁵

These cases highlight the significant problem with facial recognition technology in criminal investigations: if the algorithm wrongly selects someone as a suspect, that can harm an innocent person. Moreover, new technologies, such as facial recognition algorithms, are being incorporated within systems of policing and surveillance (within the United States and elsewhere) that already have histories of racial disparity and inequality. These preexisting patterns heighten the stakes surrounding the adoption of any new technology, especially for people of color, as emphasized by scholars such as Ruha Benjamin in her ground-breaking study *Race After Technology* and Brian Jefferson in *Digitize and Punish*.⁶ Consider a concrete example: after Nijeer Parks was arrested, he seriously considered admitting guilt and accepting a plea bargain—even though he had been arrested based on a faulty image match and he had been thirty miles away when the crime was committed—because he had prior criminal convictions (from about a decade earlier), which could have led to a severe sentence if he had been convicted at trial.⁷

As recently as 2018 and 2019, detailed studies by researchers at MIT and Microsoft Research and at the US National Institute of Standards and Technology (NIST) identified persistent inaccuracies in algorithms that were designed to detect and/or identify faces when applied to people of color.⁸ The algorithms were also less accurate for women than for men, with the largest errors, up to 35 percent, arising for female faces of color, according to the MIT/Microsoft study.⁹ This report and its authors, Joy

Buolamwini and Timnit Gebru, have been instrumental in raising the issue of bias in facial analysis software. Buolamwini, who founded the Algorithmic Justice League to support her efforts, explains the origins of her work in the documentary film *Coded Bias* (2020).^{[10](#)}

These recent efforts highlight at least two major challenges: constructing algorithms that have comparably high accuracy for identifying individuals across racial, ethnic, gender, and age groups; and incorporating such technologies into real-world systems—such as policing and government surveillance—that already have long histories of systematic inequalities.

Early Challenges

One reason for the racial disparities in the performance of facial recognition technologies is the relative lack of nonwhite faces in the sample data sets that have been used to develop the algorithms. The poor representation of people of color from around the world, and their range of facial features and skin shades, creates what researchers call a “demographic bias” built into the technology. On the other hand, the criminal databases used by police in the United States oversample people of color, leading to further racial disparities. In *Digitize and Punish*, Jefferson notes, “In some cities, 80 percent of the black male population is listed in [criminal justice] databases.”^{[11](#)} Likewise, in *Race After Technology*, Benjamin describes a gang database in California in which people of color constitute 87 percent of the entries; some entrants in the database were infants, less than a year old.^{[12](#)}

Disparities in the data with which facial recognition technologies have been developed stretch back to the origins of the field. Consider, for example, the efforts of American mathematician and artificial intelligence (AI) pioneer Woodrow Wilson “Woody” Bledsoe. Beginning in 1959, he produced several innovations in pattern recognition and computer vision that made him a major figure in the development of AI.^{[13](#)} Although his pioneering efforts in facial recognition have since been surpassed, they underscore the long-standing challenges of addressing bias in data collection, sampling, and analysis.

Working at the Sandia Corporation in Albuquerque, New Mexico, Bledsoe and a colleague first invented a way for a machine to recognize alphanumeric characters. They projected a character, say the letter “A,” onto a rectangular 10 x 15 array of photocells. Each photocell represented a discrete picture element or pixel and received a binary 1 or 0 depending on whether or not it contained part of the image. The pixels were sampled in adjacent groups of two or more called “n-tuples” to account for the spatial relations among them. Further manipulation produced a set of binary digits embodying the letter “A.” This process found and stored the bits and a resulting unique score for every character; then an unknown character was identified by comparing its score to the values in memory. The method worked well, correctly recognizing up to 95 percent of handwritten or printed numerals.

But when Bledsoe and colleagues later tried to extend the method to identify faces, they ran into difficulties. N-tuple groups worked poorly for the intricate pattern of a face and the large amount of data it takes to represent it in detail. In addition, whereas a letter or number is a static two-dimensional object, a face is a variable three-dimensional one. The photographic appearance of a face can vary greatly with lighting, the tilt of the head, and the facial expression. It also changes with the subject's age.

To cope with these complexities, instead of n-tuples Bledsoe's team turned to human operators who manually measured characteristic parameters from photographs of faces, such as the distance between the pupils of the eyes or from top to bottom of an ear. By the 1960s, Bledsoe could show that with this hybrid method, a computer using stored facial measurements reduced by 99 percent the number of images a person would have to sift through to match a new photo—a huge improvement when examining large photographic datasets. Then, in 1973, the entire process was automated with software that measured eyes, ears, and so on in a photo of a face without human intervention.

Bledsoe's seminal work at Sandia was classified and his later efforts in facial recognition were funded and classified by the Department of Defense or the Central Intelligence Agency, so details of his research have not been widely known. But in early 2020, writer Shaun Raviv described in *Wired* what he learned from examining Bledsoe's life and an archive of his work given to the University of Texas after Bledsoe's death in 1995.¹⁴ The facial recognition study, Raviv reports, began with a data set of photos of four hundred white males. In the archive, Raviv saw no references to women or people of color, or their images in dozens of marked-up photos that represent Bledsoe's facial measurements.

Machine-Learning Approaches

Bledsoe's early approach relied upon examining specific facial features. More recent efforts aim to mimic humans' extraordinary ability to recognize faces, by considering an entire face at once. As explained by Sofia Landi, a neuroscientist at Rockefeller University who studies human facial recognition, "If they ask you to describe the face of a loved one it's hard actually to think about all the details...in our memory I think that we store them as a holistic percept."¹⁵

Translating such tasks to computers requires examining all the data needed to describe a face at high resolution, a daunting task when sifting through huge data sets, even with fast computers.

Researchers, however, have found ways to reduce the data requirements. One method, called "eigenfaces," was developed in the 1980s and 1990s. It uses mathematical and statistical analysis to develop a set of so-called standard faces. These can be combined in different proportions to represent any given face reasonably accurately, using much less data than a straightforward pixel-by-pixel analysis. More recently, machine learning, a key part of AI, is bringing the newest changes. Yet even

with these recent efforts, the biases that arise from a lack of diversity still appear, for much the same reason that Bledsoe's data sets were formative in developing his algorithms.

The impact of machine learning on the field is shown in data from NIST, which for years has invited producers of facial recognition algorithms to submit them for testing. In 2019, NIST presented its analysis of 189 algorithms from 99 mostly commercial developers. These were checked against federal databases with 18 million images of 8.5 million people for general accuracy and across different demographic groups. The NIST team assessed the algorithms' performance in two applications: 1:1 matching, in which a face is compared to a stored image for verification, as in confirming the validity of a passport; and 1:n matching, in which a face is compared to a whole data set, typically to find a criminal suspect. For each algorithm, the researchers determined the number of false negatives, in which a face that should be matched to one in the database is not, and false positives, in which a face is matched to the wrong one.¹⁶

By many measures, the data showed significant improvement in recent years. For example, the rate at which algorithms failed to find any match between a submitted face and one in the database dropped from 4 percent in 2014 to only 0.2 percent by 2018. Newer algorithms were also less sensitive to the variations in facial appearance that plagued early efforts. The best of these algorithms, the NIST researchers found, came out of an "industrial revolution" in facial recognition, the adoption of deep convolutional neural networks (CNN) to analyze faces, another holistic approach.¹⁷

A neural network is a computing system that can be taught to carry out certain tasks, somewhat like the connected neurons in a biological brain. Applied to facial recognition, a CNN mimics human visual perception (although not all the specifics of human facial recognition, which are still being uncovered). In general, neurons in specialized regions of our brains register certain general elements in what the eyes see, such as the edges of objects, lines tilted at particular angles, color, and motion. The brain assembles these results into a meaningful whole that allows a person, for example, to quickly recognize a friend even under obscured or varied conditions.

As in the n-tuple method, in a CNN the pixels forming an image are analyzed in spatially adjacent clumps, but succeeding stages go deeper. Like the regions in the brain, each stage seeks different types of general pictorial elements rather than seeking specific elements such as eyes, nose, and so on. The results are mathematically manipulated, passed on through the stages, and accreted to finally produce an integrated representation of a face.

Crucially, this process is specifically designed to recognize faces by first exposing the CNN to a large data set of varied facial images. The network analyzes and identifies them, and receives feedback about whether the results are correct. Then its interior parameters are adjusted to improve the quality of the identifications. With enough faces and iterations, this trains the system to develop a

comprehensive approach to recognizing faces. But to make the approach truly equitable across population groups, it must be trained with an equitable data set. A majority-white data set, for example, does not produce accurate results for dark faces, and training a system using only images of men does not guarantee valid results for women. Lack of diversity characterized the data sets that Bledsoe used to develop his early methods, and diversity is now recognized as an important requirement within the CNN method.

Within NIST's testing, CNN-based algorithms performed best; but overall, different algorithms differed in how well they identified people of different races, genders, and ages.¹⁸ These results echo earlier studies of 1:1 matching and are the first to explore demographic effects in 1:n matching. Errors in each application yield different undesirable outcomes. A false positive in a 1:1 search can allow unauthorized access to a physical site or device that requires a match to an approved face in a database; a false positive in a 1:n search for a criminal suspect puts the subject at risk for unwarranted accusations.

In 1:n matching, the 2019 NIST data show that the most accurate algorithms are also the most reliable across demographic groups. Less proficient ones gave higher rates of false positives for Black females compared to Black males, white males, and white females in a database of 1.6 million mugshots maintained by the US Federal Bureau of Investigations (FBI). For 1:1 matching, some algorithms falsely matched Black and Asian faces ten to one hundred times more often than white ones. Notably, however, some algorithms from Asian countries yielded fewer false positives for Asians than for whites. This, the report notes, shows how the degree of diversity in a training data set can affect the racial performance of a CNN.¹⁹

Other research has more explicitly explored how lack of diversity affects the training of a neural network. In 2012, B. F. Klare and A. K. Jain at Michigan State University, with colleagues, tested 1:1 facial matching against police mugshots.²⁰ Different types of algorithms they examined were all less accurate for Black faces than white or Hispanic ones. One of these biased algorithms was a neural network that used a training data set. The researchers found that its error rate for Black faces decreased when this data set was limited to Black faces, and also improved when the training data set included equal numbers of Black and white faces.

This suggests how to make biased training databases more equitable. In one recent demonstration, researchers at the biometrics company Onfido developed an algorithm trained on a demographically unbalanced data set that could nonetheless yield less biased results.²¹ Its facial training images came from different continents in varying proportions, such as 0.5 percent from Africa compared to 61 percent from Europe. This yielded a false positive rate sixty-three times higher for African faces than for European ones. But when the researchers used statistical methods to train with more African faces

than their small numbers alone would provide, the discrepancy was reduced to a factor of 2.5. A team of academic researchers in Germany has developed other approaches to mitigate demographic bias without needing to retrain the algorithm.²² Such efforts continue: a review published in 2020 includes references to more than seventy research papers published since 2009 that assess the degree of bias in facial recognition algorithms and various ways to try to mitigate it.²³

Beyond the Laboratory

As researchers continue to develop methods with which to increase the accuracy of facial recognition technologies within controlled, laboratory settings, difficult questions remain about applications of the technology in real-world settings.

In an email interview, Patrick Grother, lead biometric scientist for the NIST report, explained that a facial recognition algorithm actually returns a list of likely candidates. In the context of a police action, the next step is supposed to be that a human investigator confirms that there is a good match within the list and then seeks additional evidence, such as eyewitness testimony or forensic evidence from a crime scene, to justify arresting a particular subject. Several recent examples, however, indicate that these steps have not always been followed properly, leading to the arrest of individuals who had been (mis)identified by facial recognition technologies.²⁴

Neither humans nor algorithms are infallible; for example, recent studies have documented subjective flaws in eyewitness identification of suspects.²⁵ Similar weaknesses of human judgment can affect the use of facial recognition technologies. Computer scientist Timnit Gebru (formerly at Google, where she co-led its ethical artificial intelligence team, and coauthor of the MIT/Microsoft paper cited earlier) has noted that individuals might trust the output of an algorithm—thinking that it must be “objective”—more than their own judgment about whether two images match, despite underlying biases or inaccuracies in the technology.²⁶

Such concerns are more than idle speculation, as other examples from forensic technologies have shown, such as the long-established form of identification, fingerprint analysis. Much like computerized scanning of facial images in a criminal investigation, a suspect fingerprint is checked against an FBI database of millions of fingerprints by the Integrated Automated Fingerprint Identification System (IAFIS). A trained expert makes a final determination from the resulting short list of partial matches, based on existing standards for how many matching features are required for a valid identification. Even after decades of experience, calibration, and scrutiny of such methods, however, high-profile mistakes can still occur—not only by novices but even by experts.

In March 2004, a terrorist bombing of commuter trains in Madrid, Spain, killed 191 people and injured 2,000.²⁷ Spanish police sent fingerprints from the crime scene to the FBI. An experienced lead FBI

examiner followed by two others mistakenly identified one of the IAFIS matches as coming from Brandon Mayfield, an American-born lawyer and convert to Islam living in Oregon. His fingerprints were on file due to his service in the US Army, rather than any prior criminal activity. He and his family were secretly monitored and had their home searched, and he was wrongly imprisoned for two weeks. Mayfield later sued the US Justice Department and obtained a public apology and an unusually large settlement of \$2 million.^{[28](#)}

In a later study of what went wrong, an international panel of experts concluded that under pressure to make a crucial high-profile decision, the lead examiner overly trusted that the IAFIS list must surely include a fully valid match. This influenced both the lead examiner's judgment and that of the other examiners, who were told of the lead's decision—an example of “confirmation bias,” the tendency to interpret evidence to support a person's existing beliefs or choices. Although the FBI initially dismissed objections by the Spanish police to the proposed match, the FBI eventually conceded the mistake, and has since changed its fingerprinting protocols.^{[29](#)}

What came next is also instructive. In 2009 and 2016, high-level federal scientific panels considered the FBI's Madrid fingerprint failure, firearms analysis, and other problematic areas in forensic science. The panels recommended changes such as confirming the scientific validity of the methods, establishing rigorous standards for successful matches, and providing better training for practitioners and users.^{[30](#)}

In the fingerprinting case, close scrutiny led to reforms and oversight of required procedures at federal agencies. The contrast with present uses of facial recognition software is significant. Local policing in the United States is dispersed across thousands of separate agencies and jurisdictions. Although the recent NIST study indicates that different commercial products yield varying levels of accuracy, including different levels and types of errors across racial, ethnic, and gender categories, at present local police departments can choose among scores of commercially available algorithms without clear guidance about their proper uses or reliability.^{[31](#)}

For these reasons, some policymakers have begun to call for legislation to establish standards and regulation for facial recognition technologies. In a series of hearings held in 2019, the US House Oversight Committee achieved bipartisan support in favor of regulation, and such a bill might soon become US law.^{[32](#)} In the meantime, the local governments of Somerville, Massachusetts, and of Oakland and San Francisco, California, banned any city department from using the technology, including the police.^{[33](#)} Responding to recent cases of unjust treatment of people of color, Timnit Gebru has likewise called for the technology to be “banned at the moment. I don't know about the future.”^{[34](#)} An editorial in the *Washington Post* also recently advocated a moratorium on unregulated facial recognition technologies. The editors argued that such a moratorium would protect those most at risk

while giving Congress time to develop rules to support “positive applications of facial recognition technology used correctly.”³⁵

The challenges—and the stakes—are only going to grow. Ubiquitous in homes, businesses, and public spaces, a billion surveillance cameras are projected to be placed in over fifty countries by 2021, one for every eight people on Earth.³⁶ Within the United States alone, as of 2016, facial images of half of American adults—117 million people—appeared in databases accessible to police.³⁷

Discussion Questions

1. Technologies such as facial recognition algorithms are deployed within existing structures and institutions. Ruha Benjamin observes that when one conceptualizes racial bias as a technical glitch, “the idea becomes ‘let’s fix it with better data’—rather than stepping back and thinking: is it really possible to create fair and just algorithms in an unfair and unjust society?”³⁸ If you were a local official, what specific steps would you suggest before deciding whether to implement facial recognition technology for official uses within your community?
2. Researchers at the Georgetown University Center on Privacy and Technology have recently observed that the “benefits of face recognition are real” and that the technology “has been used to catch violent criminals and fugitives.” However, they also note that with so much facial data being collected all the time, even from millions of people who have not had any direct interactions with the police or the legal system—and who may not realize that their images have been incorporated into databases used for criminal investigations—the situation is “unprecedented and highly problematic.”³⁹ Do you agree? What are some new questions that facial recognition technologies raise regarding public safety, individual privacy, and inequitable impacts on minority communities?

Bibliography

Albright, T. “[Why Eyewitnesses Fail](#).” *PNAS* 114, no. 30 (2017): 7758–64.

Ballantyne, M., R. Boyer, and L. Hines. “[Woody Bledsoe: His Life and Legacy](#),” *AI Magazine* 17, no. 1 (1996): 7–20.

Benjamin, R. [Race After Technology: Abolitionist Tools for the New Jim Code](#). Medford, MA: Polity, 2019.

Boyer, A., and R. Boyer. “A Biographical Sketch of W. W. Bledsoe,” in [Automated Reasoning: Essays in Honor of Woody Bledsoe](#), edited by R. Boyer, 1–30. Dordrecht, Netherlands: Kluwer Academic Publishers, 1991.

Bruveris, M., J. Gietema, P. Mortazavian, and M. Mahadevan. [Reducing Geographic Performance Differentials for Face Recognition](#). Preprint, submitted February 27, 2020,

<https://arxiv.org/pdf/2002.12093.pdf>; also presented in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)* (March 2020), 98–106.

Buolamwini, J., and T. Gebru. “[Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification](#).” *Proceedings of Machine Learning Research* 81 (2018): 1–15.

Cosgrove, E. “[One Billion Surveillance Cameras Will Be Watching around the World in 2021, a New Study Says](#).” CNBC, December 6, 2019.

Drozdzowski, P., C. Rathge, A. Dantcheva, N. Damer, and C. Busch. [Demographic Bias in Biometrics: A Survey on an Emerging Challenge](#), *IEEE Transactions on Technology and Society* 1, no. 2 (June 2020): 89–103.

Editorial Board, *Washington Post*. “[Unregulated Facial Recognition Must Stop Before More Black Men are Wrongfully Arrested](#).” *Washington Post*, December 31, 2020.

Ferguson, A. [The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement](#). New York: NYU Press, 2017.

Garvie, C. “[The Untold Number of People Implicated in Crimes They Didn’t Commit Because of Face Recognition](#).” American Civil Liberties Union, June 24, 2020.

Garvie, C., A. Bedoya, and J. Frankle. [The Perpetual Line-up: Unregulated Police Face Recognition in America](#). Washington, DC: Georgetown Law Center on Privacy and Technology, 2016.

Grother, P., M. Ngan, and K. Hanaoka. [Face Recognition Vendor Test. Part 3: Demographic Effects](#). NISTIR 8280. Gaithersburg, MD: National Institute of Standards and Technology, December 2019.

Hill, K. “[Wrongfully Accused by an Algorithm](#).” *New York Times*, June 24, 2020; updated August 3, 2020.

Hill, K. “[Another Arrest, and Jail Time, Due to a Bad Facial Recognition Match](#).” *New York Times*, December 29, 2020; updated January 6, 2021.

Jefferson, B. *Digitize and Punish: Racial Criminalization in the Digital Age*. Minneapolis: University of Minnesota Press, 2020.

Kantayya, S. [Coded Bias](#), documentary film. 7th Empire Media, 2020. Film.
<https://www.codedbias.com/>

Klare, B., M. J. Burge, J. C. Klontz, R. W. Vorder Bruegge, and A. K. Jain. “[Face Recognition Performance: Role of Demographic Information](#).” *IEEE Transactions on Information Forensics and Security* 7, no. 6 (2012): 1789–1801.

Landi, S., and Freiwald, W. A. “[Two Areas for Familiar Face Recognition in the Primate Brain](#).” *Science* 357, no. 6351 (2017): 591–95.

Lichtblau, E. “[U. S. Will Pay \\$2 Million to Lawyer Wrongly Jailed](#).” *New York Times*, November 30, 2006.

Melton, M. “[Government Watchdog Questions FBI On Its 640-Million-Photo Facial Recognition Database](#).” *Forbes*, June 4, 2019.

Metz, R. “[Beyond San Francisco, More Cities Are Saying No to Facial Recognition](#).” *CNN Business*, July 17, 2019.

National Research Council of the National Academies. [Strengthening Forensic Science in the United States: A Path Forward](#). Washington, DC, The National Academies Press, 2009.

National Research Council of the National Academies. [Identifying the Culprit: Assessing Eyewitness Identification](#). Washington, DC: The National Academies Press, 2014.

Newman, L. H. “[Facial Recognition Has Already Reached Its Breaking Point](#).” *Wired*, May 22, 2019.

O’Brien, H. “[The New Jim Code](#)”: Ruha Benjamin on Racial Discrimination by Algorithm, *New Statesman*, September 26, 2019.

Ovide, S. “[A Case for Banning Facial Recognition](#).” *New York Times*, June 9, 2020.

President’s Council of Advisors on Science and Technology. [Forensic Science in Criminal Courts: Ensuring Scientific Validity of Feature-Comparison Methods](#). Washington, DC: Executive Office of the President, President’s Council of Advisors on Science and Technology, September 2016.

Raviv, S. “[The Secret History of Facial Recognition](#).” *Wired*, January 1, 2020.

Sciolino, E. “[Bombings in Madrid: The Attack; 10 Bombs Shatter Trains in Madrid, Killing 192](#).” *New York Times*, March 12, 2004.

Simonite, T. [Congress Is Eyeing Face Recognition, and Companies Want a Say](#). *Wired*, November 23, 2020.

Slutsky, M., and Landi, S., dirs. [What Is My face?](#), documentary film. Labocine, 2019. Film, 4 min, <https://www.labocine.com/films/what-is-my-face>.

Stacey, R. “[Report on the Erroneous Fingerprint Individualization in the Madrid Train Bombing Case](#).” *Forensic Science Communications* 7, no. 1 (January 2005).

Stokes, E. “[Wrongful Arrest Exposes Racial Bias in Facial Recognition Technology](#).” CBS News, November 19, 2020.

Strittmatter, K. [We Have Been Harmonised: Life in China's Surveillance State](#). London: Old Street, 2019.

Terhörst, P., M. L. Tran, N. Damer, F. Kirchbuchner, and A. Kuijper. “[Comparison-Level Mitigation of Ethnic Bias in Face Recognition](#).” *Proceedings of the IEEE International Workshop on Biometrics and Forensics (IWBF)*, 1–6 (April 2020).

Terhörst, P., J. N. Kolf, N. Damer, F. Kirchbuchner, and A. Kuijper. “[Post-Comparison Mitigation of Demographic Bias in Face Recognition Using Fair Score Normalization](#).” Preprint, submitted July 7, 2020. <https://arxiv.org/abs/2002.03592>.

Van Noorden, R. “[The Ethical Questions That Haunt Facial-Recognition Research](#),” *Nature* 587 (November 19, 2020): 354–58.

Williams, R. “[I Was Wrongfully Arrested Because of Facial Recognition. Why Are Police Allowed to Use It?](#)” *Washington Post*, June 24, 2020.

Copyright © the Author(s) 2021

Footnotes

1. R. Williams, “[I Was Wrongfully Arrested Because of Facial Recognition. Why Are Police Allowed to Use It?](#),” *Washington Post*, June 24, 2020. ↵
2. K. Hill, “[Wrongfully Accused by an Algorithm](#),” *New York Times*, June 24, 2020; updated August 3, 2020. ↵
3. C. Garvie, “[The Untold Number of People Implicated in Crimes They Didn’t Commit Because of Face Recognition](#),” American Civil Liberties Union, June 24, 2020. ↵
4. K. Hill, “[Another Arrest, and Jail Time, Due to a Bad Facial Recognition Match](#),” *New York Times*, December 29, 2020; updated January 6, 2021. ↵
5. E. Stokes, “[Wrongful Arrest Exposes Racial Bias in Facial Recognition Technology](#),” CBS News, November 19, 2020. ↵
6. R. Benjamin, [Race After Technology: Abolitionist Tools for the New Jim Code](#) (Medford, MA: Polity, 2019); B. Jefferson, *Digitize and Punish: Racial Criminalization in the Digital Age* (Minneapolis: University of Minnesota Press, 2020). See also A. G. Ferguson, [The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement](#) (New York: NYU Press, 2017). ↵

7. Hill, "[Another Arrest, and Jail Time.](#)" ↵
8. J. Buolamwini and T. Gebru, "[Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification.](#)" *Proceedings of Machine Learning Research* **81 (2018)**: 1–15; P. Grother, M. Ngan, and K. Hanaoka, [Face Recognition Vendor Test. Part 3: Demographic Effects](#), NISTIR 8280 (Gaithersburg, MD: National Institute of Standards and Technology, December 2019). ↵
9. Buolamwini and Gebru, "[Gender Shades.](#)" ↵
10. S. Kantayya, [Coded Bias](#), directed by S. Kantayya (7th Empire Media, 2020), documentary film. <https://www.codedbias.com/> ↵
11. Jefferson, *Digitize and Punish*, 1. ↵
12. Benjamin, *Race After Technology*, 6. ↵
13. A. Boyer and R. Boyer, "A Biographical Sketch of W. W. Bledsoe," in [Automated Reasoning: Essays in Honor of Woody Bledsoe](#), ed. R. Boyer (Dordrecht: Kluwer Academic Publishers, 1991), 1–30; M. Ballantyne, R. Boyer, and L. Hines, "[Woody Bledsoe: His Life and Legacy](#)," *AI Magazine* **17**, no. 1 (1996): 7–20; S. Raviv, "[The Secret History of Facial Recognition](#)," *Wired*, January 1, 2020. ↵
14. Raviv, "[The Secret History of Facial Recognition.](#)" ↵
15. S. Landi and W. A. Freiwald, "[Two Areas for Familiar Face Recognition in the Primate Brain](#)," *Science* **357**, no. **6351** (2017): 591–95. See also the 2019 documentary film [What Is My face?](#), directed by M. Slutsky and S. Landi ([Labocine](#), 2019), film, 4 min, <https://www.labocine.com/films/what-is-my-face>. ↵
16. Grother, Ngan, and Hanaoka, [Face Recognition Vendor Test. Part 3.](#) ↵
17. Grother, Ngan, and Hanaoka, [Face Recognition Vendor Test. Part 3.](#) ↵
18. Grother, Ngan, and Hanaoka, [Face Recognition Vendor Test. Part 3.](#) ↵
19. Grother, Ngan, and Hanaoka, [Face Recognition Vendor Test. Part 3.](#) ↵
20. B. Klare et al., "[Face Recognition Performance: Role of Demographic Information](#)," *IEEE Transactions on Information Forensics and Security* **7 (2012)**: 1789–1801. ↵
21. M. Bruveris et al., "[Reducing Geographic Performance Differentials for Face Recognition](#)," <https://arxiv.org/abs/2002.12093v1> (submitted February 27, 2020); also presented at *Institute of Electrical and Electronic Engineers Winter Applications of Computer Vision Workshops (WACVW)* (March 2020), 98–106. ↵

22. P. Terhörst et al., "[Comparison-Level Mitigation of Ethnic Bias in Face Recognition](#)," *Proceedings of the IEEE International Workshop on Biometrics and Forensics (IWBF)*, 1–6 (April 2020); P. Terhörst et al., "[Post-Comparison Mitigation of Demographic Bias in Face Recognition Using Fair Score Normalization](#)," <https://arxiv.org/abs/2002.03592> (submitted July 7, 2020). [↵](#)
23. P. Drozdowski et al., "[Demographic Bias in Biometrics: A Survey on an Emerging Challenge](#)," *IEEE Transactions on Technology and Society* **1** no. 2 (June 2020): 89–103. [↵](#)
24. Hill, "[Wrongfully Accused by an Algorithm](#)"; Hill, "[Another Arrest, and Jail Time](#)"; Stokes, "[Wrongful Arrest Exposes Racial Bias](#)." [↵](#)
25. National Research Council of the National Academies, [Identifying the Culprit: Assessing Eyewitness Identification](#) (Washington, DC: The National Academies Press, 2014); T. Albright, "[Why Eyewitnesses Fail](#)," *Proceedings of the National Academy of Sciences* **114** no. 30 (2017): 7758–64. [↵](#)
26. Gebru quoted in S. Ovide, "[A Case for Banning Facial Recognition](#)," *New York Times*, June 9, 2020. [↵](#)
27. E. Sciolino, "[Bombings in Madrid: The Attack; 10 Bombs Shatter Trains in Madrid, Killing 192](#)," *New York Times*, March 12, 2004. [↵](#)
28. E. Lichtblau, "[U. S. Will Pay \\$2 Million to Lawyer Wrongly Jailed](#)," *New York Times*, November 30, 2006. [↵](#)
29. R. Stacey, "[Report on the Erroneous Fingerprint Individualization in the Madrid Train Bombing Case](#)," *Forensic Science Communications* **7**, no. 1 (January 2005). [↵](#)
30. National Research Council of the National Academies, [Strengthening Forensic Science in the United States: A Path Forward](#) (Washington, DC: The National Academies Press, 2009); President's Council of Advisors on Science and Technology, [Forensic Science in Criminal Courts: Ensuring Scientific Validity of Feature-Comparison Methods](#) (Washington, DC: Executive Office of the President, President's Council of Advisors on Science and Technology, September 2016). [↵](#)
31. Grother, Ngan, and Hanaoka, [Face Recognition Vendor Test. Part 3](#). [↵](#)
32. M. Melton, "[Government Watchdog Questions FBI on Its 640-Million-Photo Facial Recognition Database](#)," *Forbes*, June 4, 2019; L. H. Newman, "[Facial Recognition Has Already Reached Its Breaking Point](#)," *Wired*, May 22, 2019; T. Simonite, "[Congress Is Eyeing Face Recognition, and Companies Want a Say](#)," *Wired*, November 23, 2020. [↵](#)
33. R. Metz, "[Beyond San Francisco, More Cities Are Saying No to Facial Recognition](#)," *CNN Business*, July 17, 2019. [↵](#)

34. Gebru quoted in S. Ovide, "[A Case for Banning Facial Recognition](#)." ↵
35. Editorial Board, "[Unregulated Facial Recognition Must Stop Before More Black Men are Wrongfully Arrested](#)," *Washington Post*, December 31, 2020. ↵
36. E. Cosgrove, "[One Billion Surveillance Cameras Will Be Warching around the World in 2021, a New Study Says](#)," *CNBC*, December 6, 2019. See also R. Van Noorden, "[The Ethical Questions That Haunt Facial-Recognition Research](#)," *Nature* 587 (November 19, 2020): 354–58; and K. Strittmatter, [We Have Been Harmonised: Life in China's Surveillance State](#) (London: Old Street, 2019). ↵
37. C. Garvie, A. Bedoya, and J. Frankle, [The Perpetual Line-up: Unregulated Police Face Recognition in America](#) (Washington, DC: Georgetown Law Center on Privacy and Technology, 2016). ↵
38. Benjamin quoted in H. O'Brien, "'[The New Jim Code](#)': Ruha Benjamin on Racial Discrimination by Algorithm," *New Statesman*, September 26, 2019. ↵
39. C. Garvie, A. Bedoya, and J. Frankle, [The Perpetual Line-up](#). ↵