ARTIFICIAL INTELLIGENCE | OPINION

Police Facial Recognition Technology Can’t Tell Black People Apart

AI-powered facial recognition will lead to increased racial profiling

By Thaddeus L. Johnson, Natasha N. Johnson on May 18, 2023

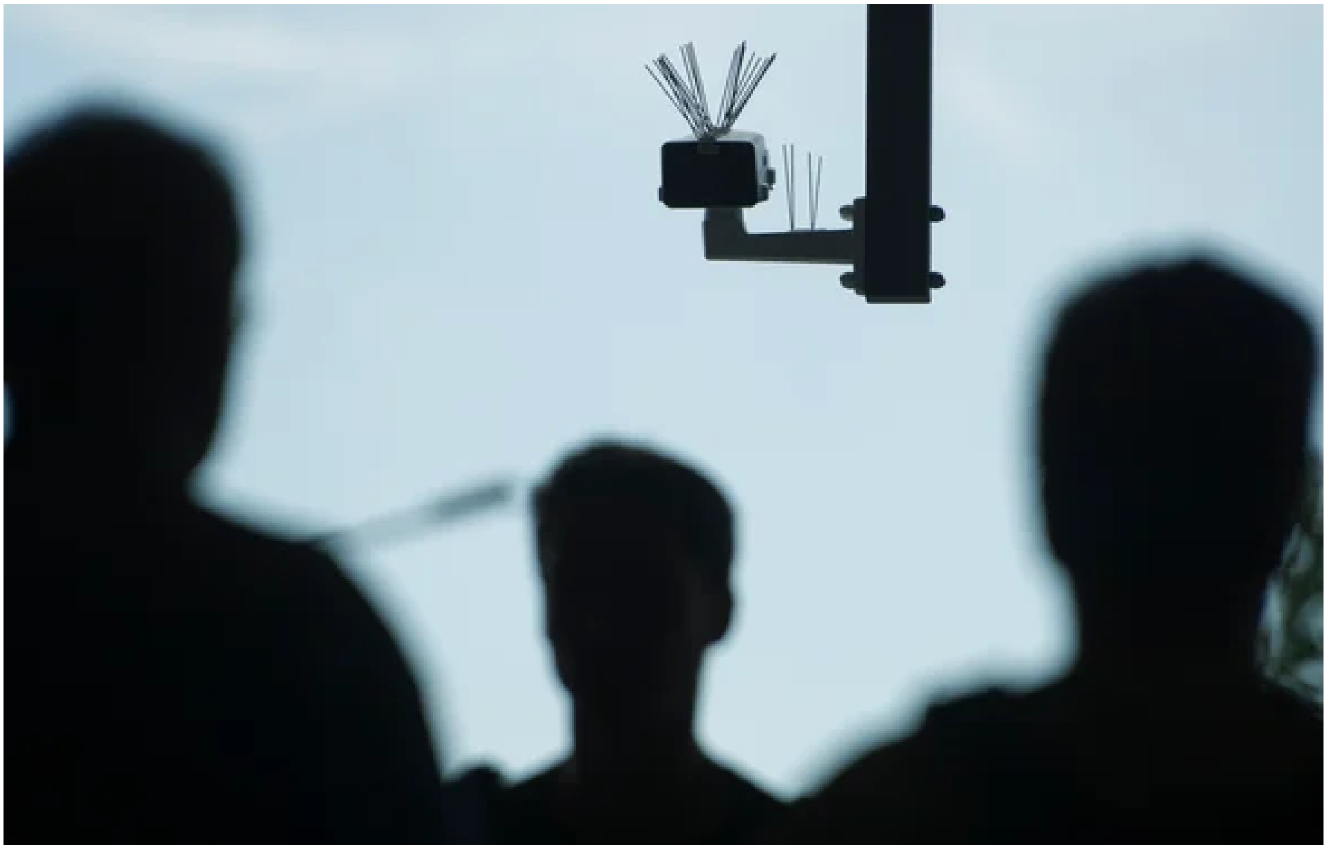

Credit: Steffi Loos/Getty Images

Imagine being handcuffed in front of your neighbors and family for stealing watches. After spending hours behind bars, you learn that the facial recognition software state police used on footage from the store identified you as the thief. But you didn’t steal anything; the software pointed cops to the wrong guy.

Unfortunately this is not a hypothetical. This happened three years ago to Robert Williams, a Black father in suburban Detroit. Sadly Williams' story is not a one-off. In a recent case of mistaken identity, facial recognition technology led to the wrongful arrest of a Black Georgian for purse thefts in Louisiana.

Our research supports fears that facial recognition technology (FRT) can worsen racial inequities in policing. We found that law enforcement agencies that use automated facial recognition disproportionately arrest Black people. We believe this results from factors that include the lack of Black faces in the algorithms' training data sets, a belief that these programs are infallible and a tendency of officers' own biases to magnify these issues.

While no amount of improvement will eliminate the possibility of racial profiling, we understand the value of automating the time-consuming, manual face-matching process. We also recognize the technology's potential to improve public safety. However, considering the potential harms of this technology, enforceable safeguards are needed to prevent unconstitutional overreaches.

FRT is an artificial intelligence–powered technology that tries to confirm the identity of a person from an image. The algorithms used by law enforcement are typically developed by companies like Amazon, Clearview AI and Microsoft, which build their systems for different environments. Despite massive improvements in deep-learning techniques, federal testing shows that most facial recognition algorithms perform poorly at identifying people besides white men.

Civil rights advocates warn that the technology struggles to distinguish darker faces, which will likely lead to more racial profiling and more false arrests. Further, inaccurate identification increases the likelihood of missed arrests.

Still some government leaders, including New Orleans Mayor LaToya Cantrell, tout this technology's ability to help solve crimes. Amid the growing staffing shortages facing police nationwide, some champion FRT as a much-needed police coverage amplifier that helps agencies do more with fewer officers. Such sentiments likely explain why more than one quarter of local and state police forces and almost half of federal law enforcement agencies regularly access facial recognition systems, despite their faults.

This widespread adoption poses a grave threat to our constitutional right against unlawful searches and seizures.

Recognizing the threat to our civil liberties, cities like San Francisco and Boston banned or restricted government use of this technology. At the federal level President Biden's administration released the “Blueprint for an AI Bill of Rights” in 2022. While intended to incorporate practices that protect our civil rights in the design and use of AI technologies, the blueprint’s principles are nonbinding. In addition, earlier this year congressional Democrats reintroduced the Facial Recognition and Biometric Technology Moratorium Act. This bill would pause law enforcement’s use of FRT until policy makers can create regulations and standards that balance constitutional concerns and public safety.

The proposed AI bill of rights and the moratorium are necessary first steps in protecting citizens from AI and FRT. However, both efforts fall short. The blueprint doesn’t cover law enforcement’s use of AI, and the moratorium only limits the use of automated facial recognition by federal authorities—not local and state governments.

Yet as the debate heats up over facial recognition’s role in public safety, our research and others’ show how even with mistake-free software, this technology will likely contribute to inequitable law enforcement practices unless safeguards are put in place for nonfederal use too.

First, the concentration of police resources in many Black neighborhoods already results in disproportionate contact between Black residents and officers. With this backdrop, communities served by FRT-assisted police are more vulnerable to enforcement disparities, as the trustworthiness of algorithm-aided decisions is jeopardized by the demands and time constraints of police work, combined with an almost blind faith in AI that minimizes user discretion in decision-making.

Police typically use this technology in three ways: in-field queries to identify stopped or arrested persons, searches of video footage or real-time scans of people passing surveillance cameras. The police upload an image, and in a matter of seconds the software compares the image to numerous photos to generate a lineup of potential suspects.

Enforcement decisions ultimately lie with officers. However, people often believe that AI is infallible and don’t question the results. On top of this using automated tools is much easier than making comparisons with the naked eye.

AI-powered law enforcement aids also psychologically distance police officers from citizens. This removal from the decision-making process allows officers to separate themselves from their actions. Users also sometimes selectively follow computer-generated guidance, favoring advice that matches stereotypes, including those about Black criminality.

There’s no solid evidence that FRT improves crime control. Nonetheless, officials appear willing to tolerate these racialized biases as cities struggle to curb crime. This leaves people vulnerable to encroachments on their rights.

The time for blind acceptance of this technology has passed. Software companies and law enforcement must take immediate steps towards reducing the harms of this technology.

For companies, creating reliable facial recognition software begins with balanced representation among designers. In the U.S. most software developers are white men. Research shows the software is much better at identifying members of the programmer's race. Experts attribute such findings largely to engineers' unconscious transmittal of "own-race bias" into algorithms.

Own-race bias creeps in as designers unconsciously focus on facial features familiar to them. The resulting algorithm is mainly tested on people of their race. As such many U.S.-made algorithms "learn" by looking at more white faces, which fails to help them recognize people of other races.

Using diverse training sets can help reduce bias in FRT performance. Algorithms learn to compare images by training with a set of photos. Disproportionate representation of white males in training images produces skewed algorithms because Black people are overrepresented in mugshot databases and other image repositories commonly used by law enforcement. Consequently AI is more likely to mark Black faces as criminal, leading to the targeting and arresting of innocent Black people.

We believe that the companies that make these products need to take staff and image diversity into account. However, this does not remove law enforcement's responsibility. Police forces must critically examine their methods if we want to keep this technology from worsening racial disparities and leading to rights violations.

For police leaders, uniform similarity score minimums must be applied to matches. After the facial recognition software generates a lineup of potential suspects, it ranks candidates based on how similar the algorithm believes the images are. Currently departments regularly decide their own similarity score criteria, which some experts contend raises the chances for wrongful and missed arrests.

FRT's adoption by law enforcement is inevitable, and we see its value. But if racial disparities already exist in enforcement outcomes, this technology will likely exacerbate inequities like those seen in traffic stops and arrests without adequate regulation and transparency.

Fundamentally police officers need more training on FRT's pitfalls, human biases and historical discrimination. Beyond guiding officers who use this technology, police and prosecutors should also disclose that they used automated facial recognition when seeking a warrant.

Although FRT isn't foolproof, following these guidelines will help defend against uses that drive unnecessary arrests.

This is an opinion and analysis article, and the views expressed by the author or authors are not necessarily those of Scientific American.

ABOUT THE AUTHOR(S)

Thaddeus L. Johnson, a former police officer, is a senior fellow at the Council on Criminal Justice and teaches criminology at Georgia State University.

Natasha N. Johnson is a faculty member at Georgia State University and director of its M.I.S. program in Criminal Justice Administration.