The Data



A landmark facial-recognition study in 2018 found that error rates were 40 times as high for dark-skinned women as they were for light-skinned men. Accuracy may have improved since, but real-world uses are creating new problems.

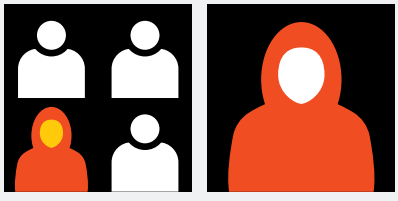
Where Everybody Knows Your Face...

Facial-recognition technology (FRT) dates back 60 years. Just over a decade ago, deep-learning methods tipped the technology into more useful—and menacing—territory. Now, retailers, your neighbors, and law enforcement are all storing your face and building up a fragmentary photo album of your life.

Yet the story those photos can tell inevitably has errors. FRT makers, like those of any diagnostic technology, must balance two types of errors: false positives and false negatives.

In best-case scenarios—such as comparing someone's passport photo to a photo taken by a border agent—false-negative rates are around two in 1,000 and false positives are less than one in 1 million.

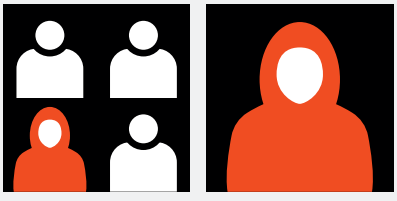
In the rare event you're one of those false negatives, a border agent might ask you to show your passport and take a second look at your face. But as people ask more of the technology, more ambitious applications could lead to more catastrophic errors. Let's say that police are searching for a suspect and they're comparing an image taken with a security camera with a previous "mug shot" of the suspect. The three possible outcomes are shown below.



a) The FRT identifies the suspect, since the two images are of the same person, according to the software. Success!



b) The FRT matches another person in the footage with the suspect's probe image. A false positive, coupled with sloppy verification, could put the wrong person behind bars and lets the real criminal escape justice.



c) The FRT fails to find a match at all. The suspect may be evading cameras, but if the cameras just have low-light or bad-angle images, this creates a false negative. This type of error might let a suspect off and raise the cost of the manhunt.

Authorized licensed use limited to: IEEE Editors-in-Chief. Downloaded on November 26,2026 at 00:11:25 UTC from IEEE Xplore. Restrictions apply.

>99%

Accuracy that FRT makers claim for best-case scenarios



EASY TO RECOGNIZE → HARD TO RECOGNIZE

Training-data composition, differences in how sensors detect faces, and intrinsic differences between groups, such as age, all affect an algorithm's performance. The United Kingdom estimated that its FRT exposed some groups, such as women and darker-skinned people, to risks of misidentification as much as two orders of magnitude greater than it did to others.

What happens with photos of people who aren't cooperating, or vendors that train algorithms on biased datasets, or field agents who demand a swift match from a huge dataset? Here, things get murky.

Consider a busy trade fair using FRT to check attendees against a database, or gallery, of images of the 10,000 registrants, for example. Even at 99.9 percent accuracy you'll get about a dozen false positives or negatives, which may be worth the trade-off to the fair organizers. But if police start using something like that across a city of 1 million people, the number of potential victims of mistaken identity rises, as do the stakes.

What if we ask FRT to tell us if the government has ever recorded and stored an image of a given person? That's what U.S. Immigration and Customs Enforcement agents have done since June 2025, using the Mobile Fortify app. The agency conducted more than 100,000 FRT searches in the first six months. The size of the potential gallery is at least 1.2 billion images.

At that size, assuming even best-case images, the system is likely to return around 1 million false matches, but at a rate at least 10 times as high for darker-skinned people, depending on the subgroup.

Responsible use of this powerful technology would involve independent identity checks, multiple sources of data, and a clear understanding of the error thresholds, says computer scientist Erik Learned-Miller of the University of Massachusetts Amherst: "The care we take in deploying such systems should be proportional to the stakes." ■



THE NEGATIVES OF FALSE POSITIVES

2020: Robert Williams's wrongful arrest cost him detention. The ensuing settlement requires Detroit police to enact policies that recognize FRT's limits.



ALGORITHMIC BIAS

2023: Court bans Rite Aid from using facial recognition for five years over its use of a racially biased algorithm.



TOO FAST, TOO FURIOUS?

2026: U.S. immigration agents misidentify a woman they'd detained as two different women.

CLOCKWISE FROM BOTTOM RIGHT: VICTOR J. BLUE/BLOOMBERG/GETTY IMAGES; ISTOCK (7)

Authorized licensed use limited to: IEEE Editors-in-Chief. Downloaded on November 26,2026 at 00:11:25 UTC from IEEE Xplore. Restrictions apply.