# COGNITIVE TIMES

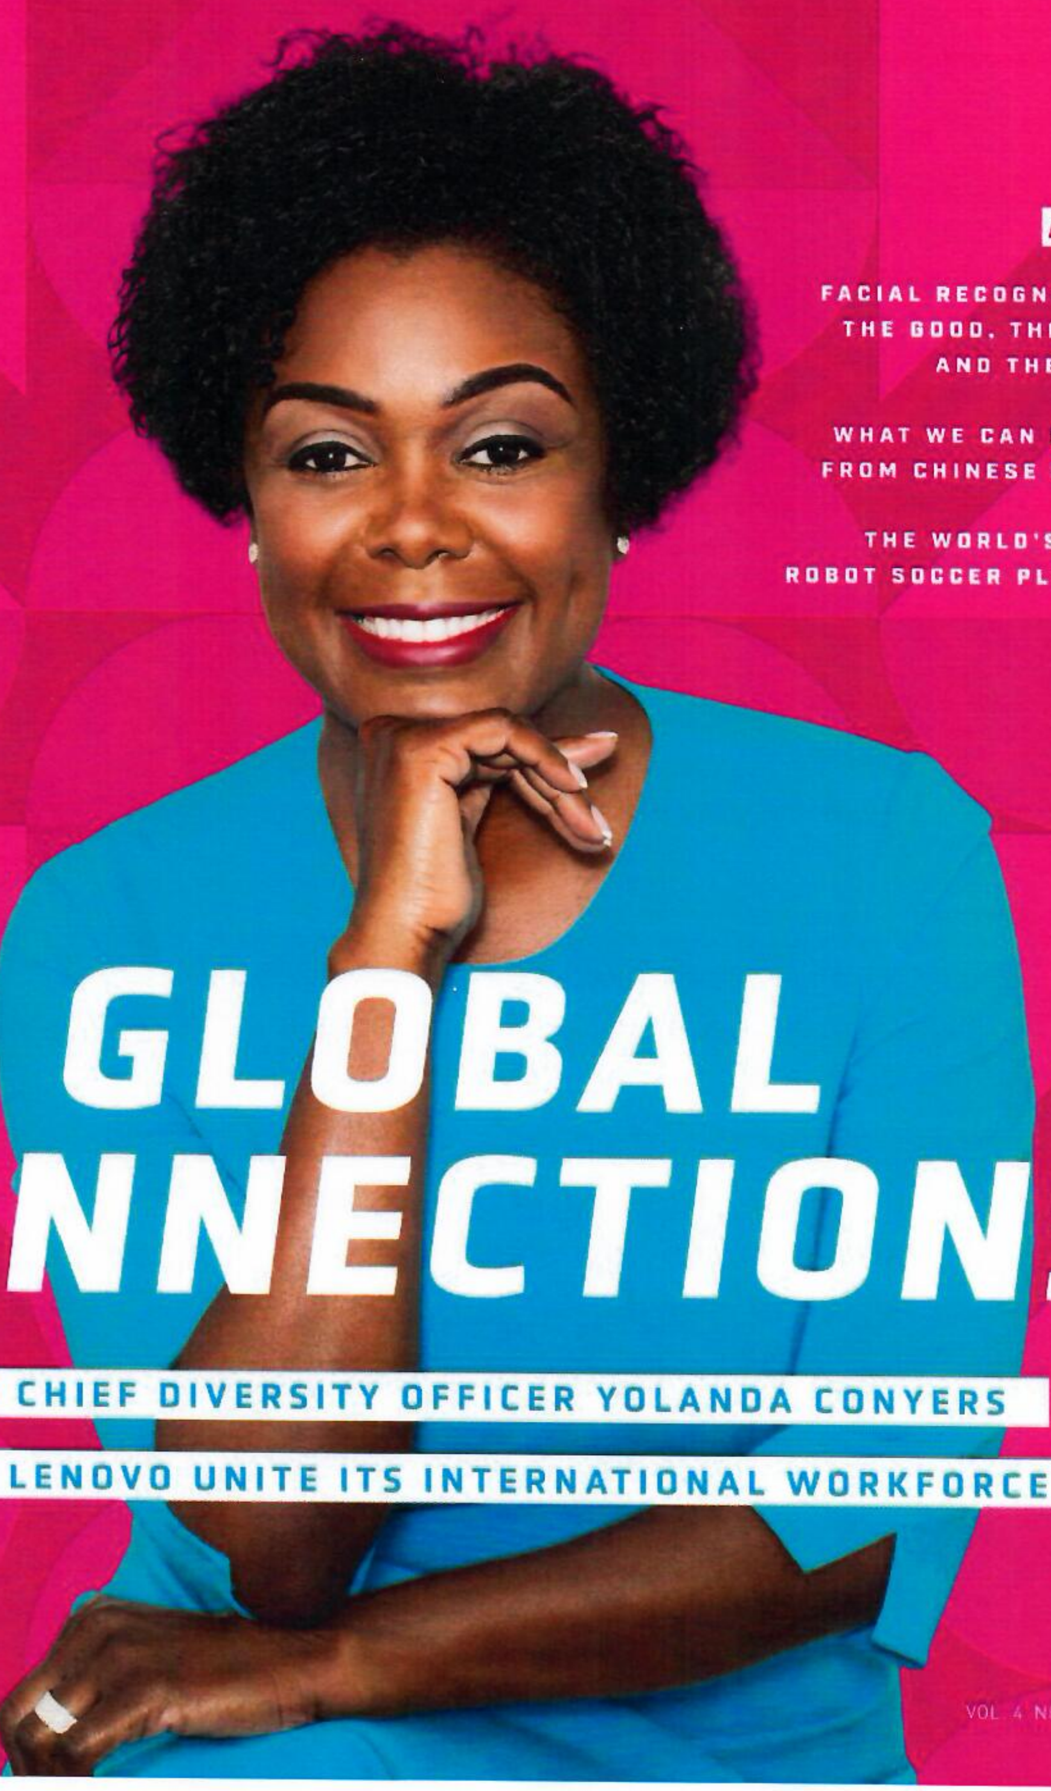

ALSO

FACIAL RECOGNITION:
THE GOOD, THE BAD,
AND THE UGLY

WHAT WE CAN LEARN
FROM CHINESE SCI-FI

THE WORLD'S BEST
ROBOT SOCCER PLAYERS

## GLOBAL CONNECTIONS

HOW CHIEF DIVERSITY OFFICER YOLANDA CONYERS
HELPS LENOVO UNITE ITS INTERNATIONAL WORKFORCE

VOL. 4 NO. 3 // 2019

# FEATURES

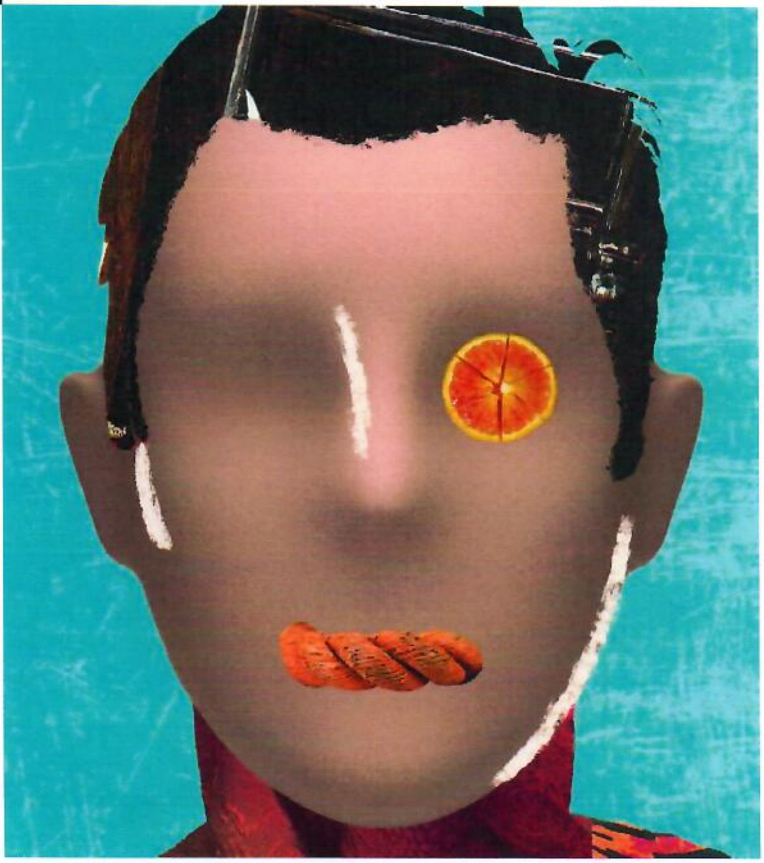

# tHe match game

**Facial recognition technology is revolutionizing law enforcement and simplifying identification, but is it ready for prime time?**

It's early on a Sunday morning when you hear a loud banging on your front door. It's someone with the FBI, shouting, "Open up. We have a search warrant!" Frightened and confused, you get out of bed, pull on some clothes, and answer the door.

The federal agents explain that they've obtained their warrant based on a witness's identification of you as one of the crooks who pulled off a recent late-night bank heist in Austin, Texas. They don't seem to care when you insist you were at home—1,300 miles away in Los Angeles—at the time that the crime occurred.

So who is this witness falsely implicating you? Facial recognition technology built into a camera at the scene of the robbery. Even though you're confident in your alibi ultimately proving your innocence, it's likely going to cost you loads of time and money to do so.

Think this scenario is fiction? Think again. Similar technology—deployed by companies like ResolutionView and Amazon, as well as Apple—is now available throughout much of the world. Along with the growing prevalence of terrorist attacks on public spaces has come greater demand for security tools enabling the identification of potential threats before innocent lives are lost. That's helped drive the market for AI-powered facial recognition to a value of $4.51 billion in 2016, with a forecasted continued growth to $9.06 billion by 2024.

Just as these dollar figures have grown, so too have complaints about the technology's use, including well-founded concerns about privacy and fears about racial profiling. Yet all is not bleak in the realm of facial recognition. As with so much of technology, it comes down to what we as a society determine are its acceptable and unacceptable uses. There's some good, some bad, and some ugly that deserves our consideration.

## How It Works

Here's an extremely basic breakdown of the steps involved in facial recognition: (1) image capture, (2) the distance between the eyes and other prominent facial features of the subject are mapped, (3) the image is converted to grayscale and cropped, (4) the image is converted to a template used by a search engine for facial comparison results, and (5) an algorithm searches for a match to the image by comparing the template to others on file.

Algorithms can be non-trainable or trainable. Non-trainable algorithms use fixed common feature representations to characterize face images. Similarities between faces are measured within these set parameters. The trainable ones, like the one used to customize user shopping by online retailer Zappos, aren't fixed. They learn over time about the preferences of particular customers and ideally improve in their anticipation of a customers' specific needs and wants, bit by bit modifying themselves with each new lesson in order to improve accuracy. So when converting an image of a face to a searchable template, a trainable algorithm would make changes to itself so as to, in theory, avoid repeating mistakes. Regardless, proprietary algorithms and the training data used to create them are generally secret. As such, they are not subject to public scrutiny. This makes it tougher for any underlying shortcomings the algorithms to be corrected, as we will see below.

Regardless of which algorithm is used, facial recognition systems generally compare the image taken in step (1) with a database in step (5). For example, according to *Governing* magazine, at least 39 states use facial recognition software linked to their Department of Motor Vehicles.

## The Good

In 1996, Lynn Cozart disappeared just days before he was to be sentenced by a Pennsylvania court for molesting three children. Investigators searched for him for years, but the trail went cold. Then, in 2015, Pennsylvania State Police sent Cozart's mug shot to the FBI's Next Generation Identification database, which contains more than 30 million facial records. The FBI's team responsible for the search—called Facial Analysis, Comparison and Evaluation Services—matched the mug shot to the face of one "David Stone" of Muskogee, Oklahoma, who worked at a local Wal-Mart. After 19 years, Cozart was brought to justice.

Then there's the case of Wu Xieyu, who was suspected of murder in 2016 and went on the run. On April 26 of this year, he was arrested just minutes after appearing in an airport in the Chinese city of Chongqing. Only six months prior to the arrest, the airport had upgraded its surveillance system to include facial recognition technology. Xieyu had kept his whereabouts unknown with the help of more than 30 identification cards, but the new tech alerted authorities when it made a 98 percent match with pictures of him in a fugitive database. Another eyebrow-raising example is the case of a fugitive caught in a crowd of about 50,000 people attending a Jacky Cheung concert in the Chinese city of Nanchang. That's why, as more fully set forth by various FBI presentations, including one by Richard W. Vorder Bruegge, facial recognition technology can be used to identify fugitives, as in the cases above, missing persons, and persons of interest.

Positive applications of AI facial recognition technology also exist in the commercial context. Apple recently reported that its new iPhone will allow you secure access via facial recognition technology, meaning no more pesky passwords. This technology makes your devices less suscep-

tible to hackers. And Apple is not alone in embracing this technology. Other merchants, including Mastercard, have started to use it in lieu of passwords.

Another example is the California eatery CaliBurger, which has linked facial recognition to its loyalty program. This isn't thought of as being an impingement on privacy since loyalty members have already agreed to share personal data with the brand. The software, which is installed in CaliBurger ordering kiosks, recognizes registered members when they approach, activates their loyalty accounts and, using previous searches, displays the customer's favorite meals. While some might find such AI learning about your preferences creepy—which is understandable—others enjoy the time-saving convenience that such learning brings to their shopping experience. Finally, some retailers have used the technology to screen for known shoplifters entering their stores via FaceFirst, a California-based facial recognition company.

### The Bad

Still, this technology remains far from perfect and is prone to the misidentification of subjects. There has been plentiful evidence of shortcomings in the algorithmic training of facial recognition. For example, a 2012 study titled "Face Recognition Performance: Role of Demographic Information," co-authored by the FBI, found that if a system was mostly trained on the faces of white people and then operated on the faces of black people, it might discard information useful in discerning features of black faces. The study also determined that the technology found females more difficult to recognize than males and had low accuracy rates for people ages 18-30. In fact, when the London Metropolitan Police recently tested facial recognition technology, it was found to have an error rate as high as 81%.

A "false negative" occurs when the technology fails to match a photo that you take with one that is, in fact, in the system. What's much worse is a "false positive"—when the system incorrectly matches a person's face to one in the database. As reported by the Electronic Frontier Foundation, the FBI's facial recognition system "may not be sufficiently reliable to accurately locate other photos of the same identity, resulting in an increasing percent of misidentification." According to EFF, increasing the number of people in the database doesn't improve accuracy. This idea is buttressed by Brendan Klare, CEO of Rank One Computing and co-author with the FBI of the 2012 study. As he puts it, "garbage in, garbage out." In other words, a database of poor-quality mug shots won't improve accuracy as it grows in quantity.

These errors have significant effects. Facial recognition equipment at the Notting Hill Carnival in London resulted in about 35 false matches and the erroneous arrest of an individual who was flagged as being wanted on a warrant. The ACLU scanned the faces of all 535 members of Congress against 25,000 public mug shots, using Amazon's open Rekognition API. While none of the members of Congress were in the mug shots, the system nonetheless generated 28 false positives.

Still another concern is the invasion of privacy. Numerous reports explain how facial technology is being used by companies like IBM that scrape about one million photos from websites like Flickr in contravention of its privacy policies. In that case, IBM released a collection of the photos, coded to describe the subjects' appearance, without anybody consenting to the procedure. "None of the people I photographed had any idea their images were being used [by IBM] in this way," says Greg Peverill-Conti, a public relations executive who found that more than 700 photos wound up in IBM's facial-recognition-research "training dataset."

In a related example, the Federal Trade Commission complaint that led to this year's $5 billion settlement with Facebook alleged that the social media company deceived users about its ability to turn off a facial recognition tool that offers photo tag suggestions—the implication being that Facebook collected a growing library of images without user consent. Some companies, like Microsoft, have removed photos used for facial recognition purposes to address such concerns.

### The Ugly

Not long ago, two undercover police officers bought $50 worth of crack cocaine from a man in Florida. They didn't arrest him on the spot. Instead, one of the officers took a photo of the suspect while pretending to make a phone call, though the image quality he captured was poor. The officers used the statewide facial recognition system to compare the iPhone photos to those in the mug shot database. A match was made with Willie Lynch, who was subsequently arrested for the crime. However, the state's algorithm only expressed a "one-star" vote of confidence in the match—the lowest rating. At trial, the officers testified that they recognized Lynch not based on their eyewitness account but because of the results of the facial recognition system. But even the state's own expert expressed uncertainty about how the matching algorithm worked. While Lynch's defense lawyers demanded to see other photos of potential suspects identified by the system as evidence of its inaccuracy, their request was denied.

## A BASIC BREAKDOWN OF THE STEPS INVOLVED IN FACIAL RECOGNITION

**1. Image capture**

**2. The distance between the eyes as well as other prominent facial features of the subject are mapped.**

**3. The image is converted to grayscale and cropped.**

**4. The image is converted to a template used by a search engine for facial comparison results.**

**5. An algorithm searches for a match to the image by comparing the template to others on file.**

Such a denial is particularly troubling. According to the EFF, "criminal databases," like the one used by the Florida authorities, "include a disproportionate number of African Americans, Latinos, and immigrants, due in part to racially biased police practices." This partially explains why the 2012 study, in EFF's view, "showed that accuracy rates for African Americans were lower than for other demographics." This bias may be explained by what the study called the "other race effect," in that "humans have consistently demonstrated a decreased ability to recognize subjects from races different from their own."

Consequently, the Congressional Black Caucus has voiced concerns about "algorithmic bias" in software like Amazon's Rekognition API. In a May 24, 2018, letter to Jeff Bezos, the caucus chairman Cedric L. Richman wrote, "We are troubled by the profound negative unintended consequences this form of artificial intelligence could have for African Americans, undocumented immigrants, and protestors." Nearly 40% of the false matches of members of Congress in the ACLU study were people of color.

That being said, the high error rate in that particular study can, according to Klare, be attributed to the threshold setting used: a 80% setting to indicate a match instead of the 90% recommended for law enforcement. Matt Wood of Amazon Web Services likewise pushed back at the ACLU findings in a blog post: "When we set the confidence threshold at 99% (as we recommend in our documentation), our misidentification rate dropped to zero." In light of these sorts of errors in implementing software, he suggests concerns about false positives may be overblown.

But algorithmic bias isn't the only concern of the CBC and other critics. Another is the increasing presence of surveillance of everyday, perfectly legitimate activity. This "will only further erode the public's trust in law enforcement," the CBC's Richman wrote in his letter to Bezos. A report by the Georgetown Center on Privacy and Technology reports that only one agency, the Ohio Bureau of Criminal Investigation, has a face recognition policy expressly prohibiting the use of such technology to track those engaged in protected free speech in public. The Fourth Amendment would not likely provide such speakers any protections, since it only protects those with a reasonable expectation of privacy. However, the First Amendment does protect public speakers against retribution by law enforcement for voicing unpopular views during lawful protests not amounting to vandalism.

### What's Next?

San Francisco recently banned the use of AI-powered facial recognition technology. Other cities, such as Somerville, Massachusetts, have followed suit. But Cozart would not have been captured by the FBI years after disappearing without the use of such technology. As a result, a middle ground is likely the best approach toward the use and application of AI-powered facial recognition technology. This is especially true when studies show that the technology, when used in conjunction with human input, leads to more accurate results than when AI is used alone.

As far as abuses go, such as with social networks, they no doubt should be policed and punished. So, too, should abusive uses of the technology by state and federal government that chill free speech. Overreliance on the technology by law enforcement out to convict suspects has no quick fix, although executive orders providing more nuanced instructions for use of the technology could help. In the meantime, trust will need to be placed in the courts to ferret out, and sanction, prosecutions that rely solely on "one star" photos for felony convictions that are otherwise unreliable. It is likely that many of these issues will be resolved by the next generation of the technology, which will be more accurate.

By appreciating the good, bad, and ugly of facial recognition technology, the immense benefits will not be thrown out merely because of the associated negatives.

BY RYAN LONG

ILLUSTRATIONS BY ALEX BERGER