The California Sunday Magazine

The long history of — and recent backlash to — facial recognition

Adrien Chen

January 22, 2020

The sudden explosion of facial recognition into public consciousness can make it seem a force as unstoppable as the weather, but the technology didn't come out of nowhere. The first facial-recognition technology was arguably created by the French police officer Alphonse Bertillon more than a century ago. In the late 1800s, Bertillon devised a method for identifying criminals based on their physical features. The index cards assigned to each person included 11 physical measurements plus standardized photographic portraits and a "verbal portrait." This early biometric system allowed the body to be abstracted into information; it was "a transformation of the body's signs into text," wrote the late critic Allan Sekula. In theory, this text made the process of identifying people less prone to human error. Bertillon's system became widely used in France and soon spread to the United States, where it gained brief popularity until it was replaced by a much quicker and more reliable ruler: fingerprinting.

Now, the human face is cataloged on a scale that Bertillon could not have imagined. A report by Georgetown Law's Center on Privacy and Technology estimates that law-enforcement facial-recognition technology affects more than 117 million American adults. Within four years, the Department of Homeland Security aims to scan 97 percent of all passengers on outbound international flights. Since 2017, Facebook has used facial recognition to tag people in photos, and every month, it seems, another goofy photo-filter app that may or may not be a honey pot for a Russian data-mining operation goes viral.

Why has the use of facial recognition become such a hot-button issue now? The most obvious answer is that the technology has been improved, streamlined, and commercialized to the point that it has become widely accessible, available for purchase for as low as 40 cents an image if you opt for Amazon's facial-recognition software plan. The earliest automated facial-recognition systems developed in the 1960s required human operators to manually enter facial features for a computer to learn, not unlike Bertillon's first system. Computer scientists can now teach computers to teach themselves to recognize faces.

The 9/11 attacks were a sort of big-bang moment for automated facial recognition. The infamous surveillance-camera images of two of the hijackers, Mohamed Atta and Abdulaziz al-Omari, passing through airport security in Portland, Maine, suggested to many that facial recognition could have identified them and prevented the attack. In November 2001, the Senate held a hearing on biometrics where Dianne Feinstein argued that "because these cameras didn't use

facial biometric systems, security was not alerted, and the hijackers remained free to carry out their bloody plans." By December, stock for Visionics, an early face-recognition pioneer, soared by more than 300 percent.

Like other efforts to secure the United States from foreign terrorists, facial-recognition technology had unexpected consequences. It was Visionics that developed the first real-time law-enforcement application of facial recognition, in 2001, on civilians in Tampa, Florida, placing cameras throughout the city's downtown in the hope of finding criminals with outstanding warrants. As agencies like the FBI and Customs and Border Protection adopt the technology, it is becoming clear that there is no guarantee it will reduce human bias in identifying people. Joy Buolamwini, who studies the social implications of technology, has been raising awareness of the tendency for facial algorithms to misidentify black faces, which she suggests is the result of a stark lack of diversity in artificial-intelligence research. Then there is the issue of who is subjected to these technologies in the first place. When Detroit, which is about 80 percent black, installed a system of hundreds of cameras — some of which had facial-recognition technology — many residents saw it as a new chapter in the long history of black Americans being subjected to unique surveillance in the name of public safety.

Three cities have banned the use of facial-recognition technology by authorities, with many more considering legislation to do the same. Detroit's system was met with vociferous protests. The routine searching of mug shots and driver's licenses by authorities has been decried by civil-liberties groups as amounting to a "perpetual lineup." The scrutiny of China's use of facial-recognition technology to suppress its Uighur minority is in no small part fueled by a fear that something similar might happen in the U.S. Perhaps facial recognition has become a symbolic last straw in our techno-skeptic moment. Even still, the metrics tell us a growing number of people are willing to opt in, or perhaps just unwilling to resist the lure of technological convenience.

How a computer learns to identify a face

In 1994, Dr. Joseph Atick was plugging away in his office at Rockefeller University in New York City. He had theorized that computers could one day process biological information — a face, specifically — the way the human brain processes visual information, and was spending long nights with a team writing code to make it happen. "It was something like 4 a.m., after a long day of failures. We were just correcting things, then submitting to the computer to compile," says Atick. He left his office with his teammates to stretch his legs and use the bathroom. When they came back into the room, the computer greeted them by name in its tinny voice. "That was a defining moment," Atick says. Though the technology has developed some 25 years since Atick's software recognized him, the basic principles remain the same. Atick and Benji Hutchinson of NEC Corporation of America, one of the world's largest providers of biometric technology, walk us through what it takes to teach a computer facial recognition.

1. "A face is the most important pattern we as humans recognize," says Atick. "Our mothers, our brothers and sisters." But a software doesn't see this image as a person — yet.

2. To untrained software, a facial image is just a bunch of pixels, which are then converted into a series of numerical values that represent the intensity and direction of light and shadow at each point. From there, it can begin to recognize patterns that correspond to facial features: an eye socket, a jawline, or a nose.

3. To learn how to identify a unique face, the software is trained to look at a certain set of landmarks, usually 68 of them. Where these data points fall is unique for every person, like a fingerprint — researchers and scientists call it a "faceprint." Other methods of identification, like Surface Texture Analysis, which maps and catalogs skin texture, follow a similar process to identify lines, pores, or wrinkles.

4. So far, the software can recognize a face, but only when the person is looking directly at the camera. What if the person turns his head? What if he's looking down? By taking the "landmark" data of the off-angle face and then scaling and rotating it, the software can learn to adjust images of faces that are oriented differently.

5. Atick and his colleagues used to manually write facial-recognition algorithms. Today, the process is driven by machine learning, but "humans are still involved," Hutchinson says. "Lab technicians make the technology more accurate by feeding it more images." Private companies can now create their own facial-recognition systems without much understanding for how or why their accuracy may skew one way or another.

6. The software is only as strong as its training data. "If all they're seeing is a particular class of people, they'll make errors on the 'out-of-sample' class," Atick says. Many facial-recognition systems have trained on datasets that are largely white and male. "We see dramatically lower error rates with certain demographics than others for whom we haven't invested as much in R&D," says Hutchinson.

Illustrations by Tim Peacock

Why law enforcement is growing more dependent on the technology

By Chris Outcalt

Two years ago, Arapahoe County Sheriff's Office investigator Jim Hills changed his morning work routine. He used to boot up his computer at 6 a.m. and find email blasts from big-box stores in the Denver, Colorado, metro area — crime bulletins detailing high-volume thefts along with a screen grab of the suspect from the store's security-camera footage. Hills would take a quick look, but unless he recognized the individual in the frame, there wasn't much he could do. Delete. But in early 2018, the department piloted a new facial-recognition software. Now, managers, many in charge of stores beyond the boundaries of Arapahoe County, often reach out to Hills directly.

Facial-recognition software is widely used by local and federal law-enforcement agencies across the country, including the New York Police Department and the FBI. But concerns about the technology's accuracy and questions about privacy and its susceptibility to abuse have led more than a dozen state and city legislatures to ban the software's use to varying degrees. No such laws currently exist in Colorado, and the Arapahoe County Sheriff's Office, a 790-person agency covering a predominantly white suburb southeast of Denver, has approved the technology for department-wide use. The office is applying it to anything from auto thefts to more violent crimes such as assaults.

The origins of the Arapahoe County program stretch back more than a decade. At the time, the sheriff's office was a founding member of an information-sharing consortium among Colorado law-enforcement agencies — a group that, at first, had nothing to do with facial recognition. The

idea was to create a centralized database of items such as arrest records, mug shots, and other investigative documents to make it easier for detectives to find useful information from other departments. “Folks out there committing crimes don’t care about jurisdictions and borders,” says Captain Jared Rowlison, who oversees the investigative unit. Last year, the company that maintains the consortium database, Lumen, approached the sheriff’s office about testing its new facial-recognition software. Investigator Hills attended a training to learn how to use the program, which has a drag-and-drop photo-search feature and produces a list of results based on the quality of the match. That first day, Hills had a case in mind to test it out — a stolen car from the parking lot of a Planet Fitness gym. The gym had captured good-quality security footage of a possible suspect; Hills plugged a screen grab of the guy into the Lumen program. The list of high-probability hits included the booking photo of someone who’d recently been arrested elsewhere in Colorado. Hills called up the department; turned out, not only did this agency still have the guy in custody, but they’d found him in the stolen car Hills was looking for. Hey, this is pretty legit, Hills thought.

Since then, Hills has become one of the go-to users of the program at the sheriff’s office. He quickly realized it worked well for generating leads on so-called “push-out thefts” — when shoplifters roll through an exit with a shopping cart full of unpaid merchandise, often in clear view of security cameras — and began making a point to check his morning email for crime bulletins outlining those cases. Hills estimated he’s helped solve close to 50 since he started using facial recognition. Another case involved a man and woman who’d cashed between $5,000 and $20,000 in fraudulent traveler’s checks at Costco stores across the Denver area. The pair had opened memberships at the chain using fake names, but the ID photos were real, and Hills plugged the images into the facial-recognition database. A manager identified one of the hits, and charges are pending.

Investigator Tara Young, a colleague of Hills, approached him last year about whether the program might be useful in a case she’d been working that had gone cold. It involved a woman who went on a date with a guy she barely knew who stole her wallet when she got up to go to the bathroom. The police didn’t have the option of running his plate because the parking-lot security footage was garbage. All the woman knew was his first name. But she did have a pretty good picture of his face on her cellphone. Young had Hills run the image in the facial-recognition database and, sure enough, got a hit. She put together a photo lineup, and the woman picked the guy out right away. “There’s nothing worse than having a good case but not having enough information to pursue it,” Hills says.

Rowlison is well-aware of the concerns surrounding facial-recognition technology and its biases. The department only compares photos gathered during an investigation to a database of mug shots, unlike some sheriff’s offices and federal agencies, like the FBI and ICE, that incorporate driver’s license photos into their databases. The sheriff’s office also does not write arrest warrants based on a single hit, no matter the level of confidence the system assigns the results. More identification is needed, Rowlison says, to meet the threshold of probable cause. “We firmly believe that this is simply a tool to help add a piece to a puzzle that is a criminal investigation.”

A CEO, camp counselor, high school administrator, and radiologist on why they’re embracing facial recognition

Jon Miller, CEO of PopID and Cali Group, which operates an international chain of burger restaurants called CaliBurger

About four years ago, we started putting kiosks in our burger restaurants. The goal was to reduce labor costs and increase revenue — there's a lot of data showing that people spend more money on kiosks than when a human takes their order. One theory is that it's not odd for people to order extra bacon or a double burger at a kiosk, because they're not talking to a person, where there's guilt associated with ordering too much food. But when we put in kiosks, the ordering times went way up, from 30 seconds to a minute, which created operational issues. The best idea we had was: We'll put a camera on the kiosk and use facial recognition, so when a person walks up to the kiosk, they can see their past orders and order the same thing again. They don't have to go through the process of customizing with no tomatoes and lettuce. Then we added a payment system, so that people could pay using their faces.

Basically, the first time you come to the kiosk, you place your order and swipe your credit card. Then the system asks, "Would you like us to remember you?" If you accept, we take your picture, ask you for your name and phone number, and that's it. We weren't sure how people would respond to facial recognition, but consumer acceptance of faces for payment was surprisingly high, and I think Apple's Face ID has a lot to do with that. About half of the people who interact with the kiosk want to use facial recognition, and 80 percent of the people who use facial recognition to log in also use their face to pay. The whole thing's been engineered to be in compliance with all these laws that say you can't do biometric analysis on someone without their consent. So when you walk up to the kiosk, we crop your face so that nobody behind you is subject to any facial-recognition analysis — there's a formula to select the closest face to the camera. Then we do the biometrics on your face and pull up your past order. And when you pay, we do the biometrics again to make sure it's still you standing there, so we can authenticate your credit card. Our kiosks actually don't reduce labor costs — we end up redeploying our staff so that, instead of taking people's orders, they can walk around interacting with guests.

Carly Crowley, outdoor education coordinator, Prescott Pines Camp in Prescott, Arizona

Things have definitely changed from when I began working as a camp counselor. Over the years, I've seen parents wanting a more personalized experience, where instead of searching through a photo album to catch a glimpse of their kid, they have programs or a system that sends photos of their child directly to them. Waldo Photos, the facial-recognition app, does that for us. We have campers here who just don't like to be in front of the cameras. So the photographer has to be like, "Your parents signed up for this program," and then they're like, "Oh, OK." With the older kids, they're kind of like, "Agh, my mom signed up for that?!" But the younger kids are just like, "Hey! Take my picture!" They love to be in front of the camera. With some of the older ones, we negotiate and say, "OK, I need at least three pictures of you today smiling and having fun."

Once we started using Waldo, parents didn't call or email as much asking about their kids. The only time they contacted us was when their kid looked unhappy in the photo. They call to make sure their kid is having fun. When that happened, we would go out purposefully and get some pictures of their kid smiling. It's hard to explain the context of some photos: Maybe they just happened to be not smiling, but if you talked to them, they were having a great time.

Mike Matranga, executive director of security and school safety, Texas City Independent School District in Texas City, Texas

After the shooting at Santa Fe High School in 2018, I was the guy who received a phone call from our superintendent. He said, “We as educators should not be making decisions that guys like you are making.” It’s 2020, and we’re still trying to deploy the same tactics and techniques that we used 20 or 30 years ago, and that’s just the wrong method.

When students register for school, we enter their information into a database, which includes things like their parents’ contact information, their address, any health considerations, as well as their photos. It’s important to note that we only put students in the database if there’s a reason to be in the database, like if a student has been assigned to an alternative school or if they’ve been suspended, or if local law enforcement alerts us. We also include sex offenders within a certain range of the school and the county’s most wanted. The information comes from open sources, like the county website.

A few months ago, we had a tip come through our anonymous reporting app. There was a Snapchat photo of a student holding a gun in one of our bathrooms on campus. We identified the student and the next day met him at the front door as he was entering. The student didn’t have a weapon with him but admitted he had a gun and said it wasn’t his. The student who actually owned the weapon was absent. We knew he might be coming back to school the next day, so we moved the facial-recognition software to the exteriors to beef up our cameras. As the student approached the school, it sent a message to the entire administrative staff and our deputies on campus, who identified the student right away. He was arrested on the spot.

The first time we used the technology was at graduation last year. This might be a bold statement, but I’m certain we had the safest graduation ceremony in the nation. We had a camera system deployed on three different gates. We had a student who was assigned to an alternative school — and therefore wasn’t allowed to attend extracurricular activities — who attempted to come inside, and we were able escort him out with no incident.

Amy Horner, director of Radiation Oncology for Centura Health and Parker Adventist Hospital in Colorado

With health care, medical errors, especially in a hospital setting, are a big focus. Let’s say we have a room full of patients waiting for treatment, and there are two Mrs. Smiths, and they’re both breast-cancer patients. The two-step verification is very routine for our staff, so let’s say they were distracted that day and didn’t get the right birthday. It’s those kind of human errors we are looking to eliminate. Once radiation is delivered to a patient, it’s pretty hard to take it back.

We beta-tested this technology before it was FDA-approved, and it’s been more than a year since we’ve been using it. Now, the very first day they come for their CT scan, we take down all of their treatment information and take their photo on an iPad. From that day on, everything is linked with that photo. When they arrive for their radiation treatment, they stand in front of the camera and get their face scanned, then the program verifies their identity and even tells the technician which accessories they’ll need during the treatment.

The patients love it. They are able to carry on conversations as they get verified — “Oh hey, how was your weekend?” “It was great!” — and never miss a beat.

Illustrations by Anuj Shrestha

There are few laws regulating facial recognition. That might change soon.

In May, city Supervisor Aaron Peskin introduced a bill that made San Francisco the first city in the United States to ban the use of facial recognition by law enforcement and other agencies. In June, the Boston suburb of Somerville became the second to enact a similar ban, then Oakland the third, each citing the technology's propensity to endanger civil rights as a major concern. Three states, including California, prevent police from using it in body cameras. Now, there are more than a dozen bans being considered in cities and states across the country.

In 2018, to prove the fallibility of the tech, the American Civil Liberties Union put every sitting member of the House and Senate through Amazon's Rekognition software, comparing their faces with 25,000 publicly available mug shots. Twenty-eight of them came back with false positives, springing some lawmakers into action. Though no federal regulations have yet been put in place, ten bills that address facial recognition's use are being considered — by both sides of the aisle.

Representative Jimmy Gomez of California's 34th Congressional District

I was one of the 28 congressmen who was misidentified by the Amazon Rekognition software. I wasn't surprised. I know for a fact that I have to be more careful when I get pulled over by a police officer. But it concerns me more for people who are working two or three jobs, driving down the street when they get pulled over because they matched some mug shot. They miss their job, their car gets towed, they don't have money to pay the impound fee.

Part of the reason the government hasn't implemented policies and legislation is the technology has been advancing so quickly. Republicans and Democrats, we oftentimes disagree where we want government involved in people's lives. Republicans are concerned about how this technology is being used against conservative demonstrators or the mass surveillance of American citizens. Democrats are concerned about that, too, but more about the fact that it's being disproportionately used against people of color.

Matthew Feeney, director of the Project on Emerging Technologies, Cato Institute

A lot of people will say, "Well, look, if you're not doing anything wrong, what's the problem?" I worry there will be some kind of observer effect, where people feel as if they shouldn't engage in legal activity because they know they can be searched via some kind of facial-recognition system. It's a serious concern, and it doesn't just apply to protests or religious minorities. It applies to people who go to abortion clinics, people who go to gun clubs.

Facial recognition is a technology, and technology is not good or bad by definition. I'm a civil libertarian, but even I can see a world in which I might be comfortable with facial recognition: Law-abiding citizens should be pushed from databases, the only searchable people should be those with outstanding warrants for violent crime — but not a single law-enforcement agency in the country comes close to satisfying those conditions.

Illustration by Joan Wong

Where is this all headed? A glimpse into our potential future

In 2017, researchers in China filed more than 900 patents for facial-recognition technology, nearly ten times the number in the U.S. In a nation blanketed by a network of some 200 million surveillance cameras — and where the government has fueled the private sector's development of artificial intelligence — technologists have proposed ways to incorporate facial recognition

into every aspect of society. As China's technology sector continues to spread its ideas and investment dollars around the globe, these tools come with risks. Here, three experts on technology and policy imagine how these applications might impact how we work, shop, and even learn.

PATENT CN109255739A

A technology that collects and analyzes the facial expressions and movements of students inside a classroom and categorizes the behavior it detects as "good" or "bad." After weighting these various behaviors, the algorithm will calculate the total "learning efficiency" of each classroom.

In the past 20 years, there's been a bunch of research in computer science that tries to look for particular facial expressions that might indicate that learning is taking place. One study argues that wrinkling the mouth or furrowing the brow is correlated with learning. So perhaps these systems track students' eyeballs and gazes, and look at whether their mouths are wrinkled or their eyebrows are furrowed. It will also look at students' movements — either just the top half of their bodies, or how they're moving around a classroom. Maybe if the student is sitting still and looking toward the front, the system takes it as a sign that they're engaged.

Education is keen on generating feedback in real time; I could see this data being fed back to students and teachers, perhaps in the form of a dashboard in the classroom, with a "learning indicator" that goes from zero to 100. And if we're being Orwellian, you could also see the school principal with a huge, real-time dashboard in their office, indicating which classrooms are doing well and which are not. And, in an educational system where administrators try to gather as many data points on students as possible, this data on classroom performance will likely be added to student records.

In an old-fashioned school, students who aren't paying attention know to sit down, face the front, and not look too bored. Students would learn to game this system. They aren't stupid — they'll just figure out how to move in the right way: If I know not to move, then I won't move. If I'm supposed to furrow my brow, I will furrow my brow.

A problem we talk about in education is educators teaching to the test, and this is a similar phenomenon: If the machine punishes me for getting students to move around, then we won't move around. If I were a teacher, I'd be making sure my students are moving in the right way, with the right expressions on their faces. I might start coordinating all of it myself: telling students, "Move now. Stay still. Make this face."

— NEIL SELWYN, PROFESSOR OF EDUCATION AT MONASH UNIVERSITY, AUSTRALIA

PATENT CN109299973A

A platform that uses facial recognition to identify consumers and push targeted advertisements to them

It could be anything from a kiosk or a billboard or one of those 7-foot-tall boxes with ads in a mall. You'd walk past it, and it would collect your face and compare it to a database, and then the billboard will give you advertising tailored to you. It could be based on who you are, or your age and income and gender— whether you're more likely to be interested in a mountain bike or a cane. Or where you are, or who you're with. Maybe it's the color you're wearing, or the time of

day, or the ten past locations where another billboard spotted you. Or simply your face's appearance or disposition: You look sad, so you want to watch a sad movie. If you're euphoric, maybe you're shown a big-ticket item, and you make a purchase you might find unwise if you were more sober. Any data point could be fair game.

Some might say there's nothing sinister here. They're taking information you put out every day — it's your face, it's not a secret — and guiding you to products you like, in the exact moment you want them. What might make people less comfortable is the asymmetry of power. We like to think we're in command of our choices when we purchase things, but that might not remain the case. Let's say I'm in a rush — I'm more likely to eat a quick, unhealthy meal. And if that pattern gets exploited by a machine, that means I'm not as in command of my own eating pattern as the private company observing that pattern. It's not hard to imagine the software targeting someone for an opportunistic deal: showing different people different products based on their mental state, or even different prices for the same product.

In many states, if you buy a timeshare, you have — by law — a number of days to reconsider the purchase. Lawmakers recognize that certain sales tactics can get people to behave in ways that don't ultimately serve their own interests. But in this case, we're talking about a person trying to sell you something, who has all this granular data on how you've behaved in the past and how you might behave in the future — information that might not even be accessible to you.

If this system tracks the places people go over a long period of time, we might also think about how the hypothetical billboard company retains its data. Most ad companies aren't interested in collecting embarrassing or explosive data about people they're trying to market to, but as soon as you start assembling this kind of information, it could have that effect — maybe you're spotted with someone you're having an affair with, or going to a clinic for an illness you don't want a future insurance company to know about. All of this stuff could be accessible to a third party who wants to buy the data, or a government.

— BEN SOBEL, RESEARCHER ON LAW, TECHNOLOGY, AND DIGITAL MEDIA AT THE BERKMAN KLEIN CENTER FOR INTERNET AND SOCIETY

PATENT CN109829691A

A punch-card system in which employees clock in and out of work using face and voice recognition

We've seen these kinds of tools being used in high-security places like hospitals and nursing homes, as well as in agriculture. We've also seen it being used by companies like Uber to verify remote workers. A concern is that these tools could make workers' jobs more difficult: If for some reason the tech isn't perfect, employees would be locked out. Studies have shown that face recognition doesn't work as well with people with darker skin, who make up a large portion of workers in low-wage industries. And in the U.S., particularly in agriculture, there's concern about immigration status. If someone is undocumented and these systems require facial recognition to sign into work, workers would wonder, How do I know this information isn't being turned over to law or immigration enforcement? Law enforcement can be present in workplaces — in retail settings, whenever theft happens, the police are called to the scene — so that biometric data could potentially be exposed.

— AIHA NGUYEN, PROJECT LEAD OF THE LABOR FUTURES INITIATIVE AT DATA & SOCIETY

When photographer Chien-Chi Chang was covering the ongoing pro-democracy protests in Hong Kong, he began to notice how surveillance cameras — the CCTV cameras that pepper cities around the world — were being destroyed and vandalized by protesters. Police in Hong Kong have been known to use facial recognition in crime-solving, but as the protests continue, many suspect that authorities are applying the technology to identify and target activists, too. Protesters damage the cameras in creative ways, sometimes smashing the lenses or covering them with plastic wrap. Often, as one protester destroys a surveillance camera, another will shield him or her with an umbrella, one more physical barrier against a technology many in Hong Kong have grown to fear.