BIOMETRICS

FACE RECOGNITION

POLICING PROJECT FIVE-MINUTE PRIMERS:

FACE RECOGNITION

Written by Allen Slater, Policing Project Student Fellow

One of the cornerstones of front-end accountability is the principle that issues of public concern—everything from law enforcement's use of surveillance technologies and data sharing to its use of force and arrest practices—should be the subject of robust engagement between police departments and the communities they serve. But productive public debate requires access to accurate information so that the policies and priorities can be carefully considered.

The Policing Project's biometrics series aims to aid in this purpose by providing basic information on one of the more complex—and rapidly changing—areas of policing: the use of biometric technologies. In each entry in this series we will explore four questions about a particular type of biometric technology:

- **How are police using this technology?**
- **How does the technology work?**
- **How accurate is the technology?**
- **Are there reasons why the public should be concerned?**

We begin with **face recognition**, as area the Policing Project is already steeped in as we staff Axon's independent Artificial Intelligence (AI) and Policing Technology Ethics Board, which produced a report that came out strongly against adding face recognition to body-worn cameras.

Though still considered an emerging technology, face recognition has already garnered debate (and been prominently featured in recent headlines) – both for its widespread use and ethical concerns.

HOW ARE POLICE USING FACE RECOGNITION?

At its core, face recognition is an identification tool. Because a person's face is fairly unique, police can use face recognition to identify (or confirm the identity) of people in a variety of contexts, such as during stops, after an arrest, from an image obtained during in the course of an investigation, or even during real time video surveillance.

Often this identification serves an investigative purpose. For example, police might use face recognition to identify someone suspected of a crime or to determine if someone has an outstanding warrant. The technology can also be used for security purposes, including to determine if someone should be given access to a secure facility.

Face recognition technology is currently being used by a variety of law enforcement agencies at the local, state, and federal levels. Although we know some information about how law enforcement is using face recognition, the truth is that because police use of face recognition technology is largely unregulated, most jurisdictions can operate this technology with little internal or external oversight, which means some might be using the technology in secret.

We do know that multiple federal agencies use face recognition technologies in law enforcement investigations. The FBI, in cooperation with local and state agencies, oversees and uses a face recognition database containing more than 600 million images of people's faces. U.S. Customs and Immigration Enforcement (ICE) is also using face recognition, including by accessing multiple states' DMV databases. The U.S. Department of Homeland Security is already using face recognition at a number of airports and is pushing for broader implementation of the practice.

We do not know much about precisely how federal agencies are using this technology, only that they are. Congressional lawmakers recently acknowledged the need to examine how law enforcement's use of facial recognition technology affects citizens. In addition to requesting testimony from FBI and DHS officials, Congress has tasked the U.S. Government Accountability Office (GAO) with evaluating the FBI's facial recognition program.

Face recognition is also being used widely on a state and local level. Although there is no exhaustive list, reports of face recognition use have surfaced in Oregon, Florida, Ohio, Pennsylvania, Michigan, Maryland, California, and New York, just to name a few.

HOW DOES FACE RECOGNITION WORK?

There are a wide variety of private companies that market face recognition technology to law enforcement: Amazon Web Services is one of the highest profile companies in this space, but, so far, has a relatively small law enforcement footprint. Some of the major face recognition vendors already working with law enforcement across the country include DataWorks Plus, Vigilant Solutions, and NEC.

The precise details of how a face recognition algorithm works will vary depending on which company's software is being used. Regardless, the programs are typically powered by AI or machine learning, and involve four basic steps:

1. **Face detection**: A user submits a probe image (the photograph or still from a video that contains the face the user wants to identify), and a computer algorithm searches for and identifies the part of the image containing a face;

2. **Feature normalization**: The software prepares that face image for processing by "normalizing" it – standardizing its features, such as orientation, levels of zoom, resolution, or brightness;

3. **Feature extraction**: The software reads and maps the face image, creating a digital "faceprint" that the software can compare against other images to find a match;

4. **Face matching**: The software then actively compares faceprints with images from a database to find images that are a "match"; the user is then informed of the results.

It is important to note that face recognition programs do not provide simple results, such as, "Yes, there is a match." Instead, the programs often provide users with a probabilistic estimate (a measure of how confident the algorithm is) of whether the probe image matches an image (or images) from the database.

These databases containing photos accessed by law enforcement during a face recognition search, called target databases, vary between jurisdictions. Some agencies maintain databases built exclusively from local mugshots, while others utilize a combination of mugshots and state driver's license databases or the FBI's Next Generation Identification Interstate Photo System. (The use of driver's license databases in this manner is certainly not without its controversy and some members of Congress have expressed concern over the FBI's use of state DMV databases.)

IS FACE RECOGNITION ACCURATE?

The performance of a face recognition algorithm (its rate of false negatives and false positives) will depend on many factors, the most significant factor being the quality of the probe image. A good probe image must be of high enough resolution (not grainy), and also needs to satisfy a variety of other conditions, such as adequate lighting in the photograph and the correct orientation of the subject's head.

Under ideal conditions—a probe photograph taken from the perfect angle with the right lighting—almost every face recognition algorithm will be capable of 99% accuracy. But accuracy limited to these conditions is not particularly helpful when it comes to investigative policing, where probe images will almost never be ideal.

What's more, even with the highest quality images, the potential for error in face recognition still exists across all demographics. According to some studies, women and people with darker skin tones face the greatest risk of misidentification. These differences are often the result of differences in training data—in other words,

the AI/machine-learning that powers the face recognition algorithms learned from data sets that were predominantly white and predominantly male so that it performs much better on images from those groups.

These two factors – overall unreliability in the field and inequitable performance – led the Axon AI Ethics Board report on face recognition to call on Axon to not develop face recognition technology for its body-worn cameras. Axon accepted that recommendation. Additionally, the city of Orlando, which conducted two separate trials of Amazon's Rekognition software, declined to implement a face recognition program for similar reasons.

ARE THERE REASONS WHY THE PUBLIC SHOULD BE CONCERNED?

As it stands, there are a few reasons why one might be concerned with law enforcement's use of face recognition technology, particularly if left unregulated:

First, as discussed above, at the moment, the technology does not work particularly well in real-world scenarios. This means sometimes the technology can (and likely will) produce false positives. Misidentifications can have all sorts of negative potential consequences. Imagine the potential impact on an officer's decision to use force if, for example, a face recognition algorithm misidentifies someone as a wanted murderer. In other cases, false positives have led to wrongful arrests. These negative consequences are particularly troubling if they fall disproportionately on certain races or genders, as discussed above.

Second, if and when the technology begins to work as advertised, some critics are particularly concerned that face recognition has the potential to fundamentally alter daily life and day-to-day policing. For example, unlike identification tools like fingerprints or DNA, face recognition paired with pervasive CCTV or body-worn cameras might allow police to track an individual's movements, both in real-time and historically. With body-worn cameras, in particular, this risks turning a transparency tool into a surveillance tool. As a result, some fear that this technology might even have a chilling effect on free speech if people feel they're constantly being watched.

Although a few jurisdictions have banned face recognition and others have adopted transparent policies, those agencies are currently the exception, not the rule. The potential concerns mentioned above cannot be overstated given how unregulated law enforcement use of face recognition is at the moment.