Automated fingerprint analysis is one step closer to reality

August 14 2017

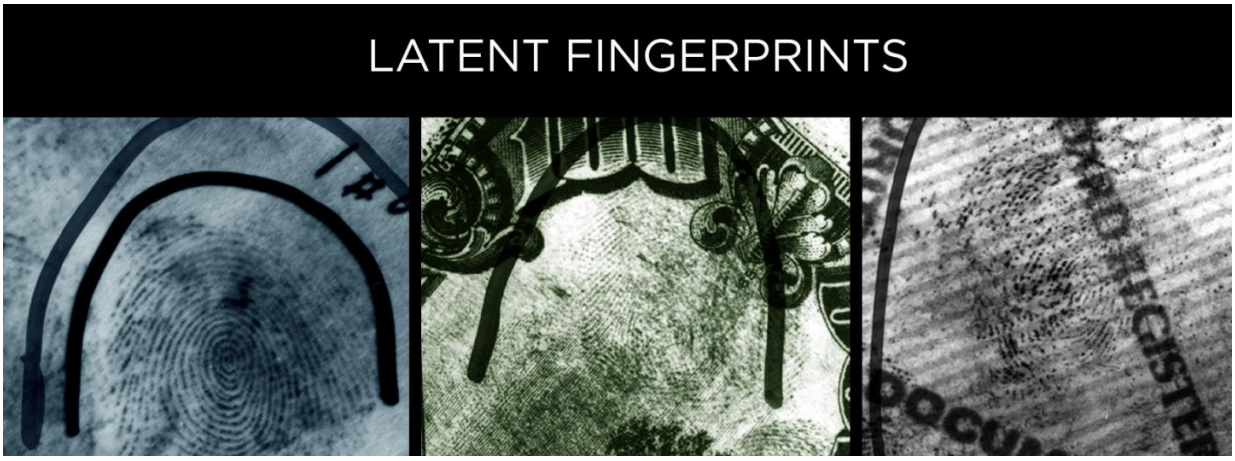

Fingerprints left at a crime scene -- so-called latent prints -- are often partial, distorted and smudged. Credit: Chugh et al., Hancek/NIST

The first big case involving fingerprint evidence in the United States was the murder trial of Thomas Jennings in Chicago in 1911. Jennings had broken into a home in the middle of the night and, when discovered by the homeowner, shot the man dead. He was convicted based on fingerprints left at the crime scene, and for most of the next century, fingerprints were considered, both in the courts and in the public imagination, to be all but infallible as a method of identification.

More recently, however, research has shown that fingerprint examination

can produce erroneous results. For instance, a 2009 report from the National Academy of Sciences found that results, "are not necessarily repeatable from examiner to examiner," and that even experienced examiners might disagree with their own past conclusions when they re-examine the same prints at a later date. These situations can lead to innocent people being wrongly accused and criminals remaining free to commit more crimes.

But scientists have been working to reduce the opportunities for human error. This week, scientists from the National Institute of Standards and Technology (NIST) and Michigan State University report that they have developed an algorithm that automates a key step in the fingerprint analysis process. Their research has been published in *IEEE Transactions on Information Forensics and Security*.

"We know that when humans analyze a crime scene fingerprint, the process is inherently subjective," said Elham Tabassi, a computer engineer at NIST and a co-author of the study. "By reducing the human subjectivity, we can make fingerprint analysis more reliable and more efficient."

A Key Decision Point

If all fingerprints were high-quality, matching them would be a breeze. For instance, computers can easily match two sets of rolled prints—those that are collected under controlled conditions, as when you roll all 10 fingers onto a fingerprint card or scanner.

"But at a crime scene, there's no one directing the perpetrator on how to leave good prints," said Anil Jain, a computer scientist at Michigan State University and a co-author of the study. As a result, fingerprints left at a crime scene—so-called latent prints—are often partial, distorted and smudged. Also, if the print is left on something with a confusing

background pattern such as a twenty-dollar bill, it may be difficult to separate the print from the background.

That's why, when an examiner receives latent prints from a crime scene, their first step is to judge how much useful information they contain.

"This first step is standard practice in the forensic community," said Jain. "This is the step we automated."

Following that step, if the print contains sufficient usable information, it can be submitted to an Automated Fingerprint Identification System. The AFIS (pronounced AY-fiss) then searches its database and returns a list of potential matches, which the examiner evaluates to look for a conclusive match.

But the initial decision on fingerprint quality is critical.

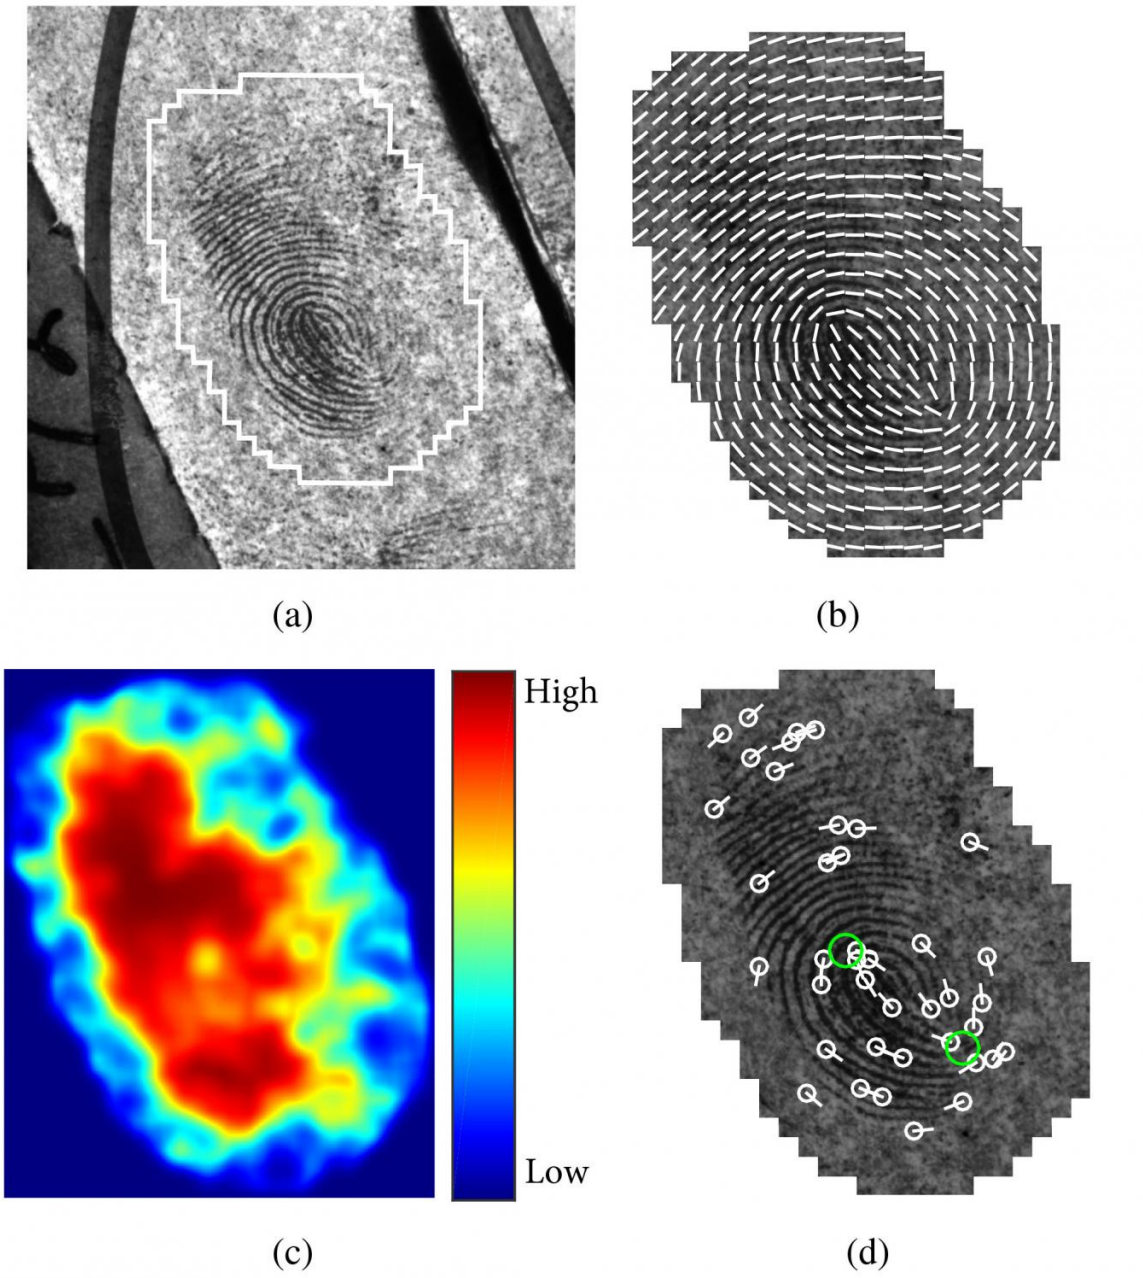

Automatically extracted features of a latent fingerprint: (a) Input latent with manually marked region of interest, (b) ridge flow overlaid on the cropped latent, (c) ridge quality map, and (d) features that can be used as points of comparison, including minutiae (white circles) and core points (green circles). Credit: Chugh et al.

"If you submit a print to AFIS that does not have sufficient information, you're more likely to get erroneous matches," Tabassi said. On the other hand, "If you don't submit a print that actually does have sufficient information, the perpetrator gets off the hook."

Currently, the process of judging print quality is subjective, and different examiners come to different conclusions. Automating that step makes the results consistent. "That means we will be able to study the errors and find ways to fix them over time," Tabassi said.

Automating this step also will allow fingerprint examiners to process evidence more efficiently. That will allow them to reduce backlogs, solve crimes more quickly, and spend more time on challenging prints that require more work.

Training the Algorithm

The researchers used machine learning to build their algorithm. Unlike traditional programming in which you write out explicit instructions for a computer to follow, in machine learning, you train the computer to recognize patterns by showing it examples.

To get training examples, the researchers had 31 fingerprint experts analyze 100 latent prints each, scoring the quality of each on a scale of 1 to 5. Those prints and their scores were used to train the algorithm to determine how much information a latent print contains.

After training was complete, researchers tested the performance of the algorithm by having it score a new series of latent prints. They then submitted those scored prints to AFIS software connected to a database of over 250,000 rolled prints. All the latent prints had a match in that database, and they asked AFIS to find it.

This testing scenario was different from real casework, because in this test, the researchers knew the correct match for each latent print. If the scoring algorithm worked correctly, then the ability of AFIS to find that correct match should correlate with the quality score. In other words, prints scored as low-quality should be more likely to produce erroneous results—that's why it's so important to not inadvertently submit low-quality prints to AFIS in real casework—and prints scored as high-quality should be more likely to produce the correct match.

Based on this metric, the scoring algorithm performed slightly better than the average of the human examiners involved in the study.

What made this breakthrough possible, beside recent advances in machine learning and computer vision, was the availability of a large dataset of latent prints. Machine learning algorithms need large datasets for training and testing, and until now, large datasets of latent fingerprints have not been available to researchers, largely due to privacy concerns. In this case, the Michigan State Police provided the researchers with the testing dataset, after having first stripped the data of all identifying information.

The next step for the researchers is to use an even larger dataset. This will allow them to improve the algorithm's performance and more accurately measure its error rate.

"We've run our algorithm against a database of 250,000 prints, but we need to run it against millions," Tabassi said. "An algorithm like this has to be extremely reliable, because lives and liberty are at stake."

More information: Tarang Chugh et al, Latent Fingerprint Value Prediction: Crowd-based Learning, *IEEE Transactions on Information Forensics and Security* (2017). DOI: 10.1109/TIFS.2017.2721099