BROOKINGS

Report

10 actions that will protect people from facial recognition software

Darrell M. WestThursday, October 31, 2019

Editor's Note:

This report from The Brookings Institution's Artificial Intelligence and Emerging Technology (AIET) Initiative is part of "AI Governance," a series that identifies key governance and norm issues related to AI and proposes policy remedies to address the complex challenges associated with emerging technologies.

Facial recognition (FR) software inspires intense reactions from many people. On the one hand, a number of individuals worry that FR will usher in an Orwellian nightmare of mass surveillance and privacy intrusions. They see FR combined with ubiquitous video cameras, artificial intelligence (AI), and data analytics as a formula for harming humanity and restricting individual freedom.

Yet at the same time, FR has been used to locate missing people, improve the security of schools and airports, help those who are visually impaired, and counter terrorism. It is the ultimate dual-use technology, with a wide range of beneficial and dangerous uses. The very same tools that horrify individuals also can help people and save lives.

Given uncertainties over its usage, some advocates have proposed an outright ban on FR deployment until we have a better understanding of its ramifications. Others suggest a short-term moratorium while regulations are enacted. Still others believe FR should be deployed, but with appropriate guardrails designed to protect people from nefarious uses.

"[Facial recognition] is the ultimate dual-use technology. ... The very same tools that horrify individuals also can help people and save lives."

In this paper, I propose 10 actions that will protect people from the greatest risks associated with FR software; these include limiting data storage and sharing, mandating accuracy standards, instituting third-party assessments, and more.

1. Limit data storage time

One of the greatest FR fears involves the lengthy digital storage of pictures and video, which can then be used in a variety of potentially pernicious ways. Individuals worry that information will be misused and impinge on their personal freedom. Rather than helping people, many are concerned about massive visual databases with few limits on FR applications.

A reform that would be useful in assuaging these fears and limiting misuse is restrictions on the amount of time for which video footage and images can be stored. If video is useful in cases of imminent danger, there is no need to retain data once that danger has passed. Indeed, for many applications, time limits would enable people to gain certain FR benefits while minimizing its risks.

Reasonable people can debate what an appropriate length of time would be for data storage. The period undoubtedly varies by usage types, with some information needing to be held for longer times and others being able to be discarded after short durations. Some images compiled for imminent danger threats have high value for short periods of time, while others may need to be

held longer. A machine learning approach known as "federated learning" ensures that data never leaves the camera itself (i.e., it is never sent to central data hub), thereby improving data security.

2. Restrict data sharing

People are bothered when information compiled for purpose A is used for purposes B, C, or D. We have seen examples of this when state motor vehicles departments sell license pictures to third-party organizations for FR in other areas. Those kinds of transfers concern people because they are not generally aware of the sale of their personal information, and they may not approve of their image being used by commercial firms for other purposes. Transferring images from organization A to B or C needs to be subject to clear standards regarding the justified need for the information in B or C. Companies should not be able to transmit or sell FR images with no limits. Rather, there has to be a clear rationale for the transmission of FR images.

In looking at public opinion, individuals vary in their FR attitudes depending on the context of the application. For example, a Brookings Institution survey found that FR favorability differed in various settings. As shown below, people were most supportive to FR applications designed to protect students in schools (41% favorable) but less favorable (27% favorable) to FR usage in retail stores to prevent theft. Opinions about FR in airports (31% favorable) and stadiums (33% favorable) were in the middle of this opinion spectrum.

Views about facial recognition	**Unfavorable**	**Favorable**
In retail stores to prevent theft	50%	27%
In airports to establish identity	44	31

Views about facial recognition	Unfavorable	Favorable
In stadiums to protect people	44	33
In schools to protect students	38	41

Source: Brookings Institution survey, Oct. 8, 2018

3. Provide clear notification in public areas

Companies or government agencies that photograph, film videos of, or otherwise collect information on people for use in FR should post clear notifications in public areas that they are engaging in these activities. That alerts the public that their face is being recorded so that those who object can avoid those areas. This step would heighten consumer awareness about FR use and enable people to take whatever personal steps they desire to protect themselves from recording.

4. Mandate accuracy standards

One of the major FR problems is varying degrees of accuracy across racial groups. Since much of the training data are incomplete or unrepresentative of the general population, the FR based on that information is biased. In the FR area, software has a much higher accuracy level for white as opposed to nonwhite faces. In particular, FR has problems with darker skin tones and therefore is less accurate with those people. With FR recognition being deployed in law enforcement, border security, retail, and airports, differential results across groups can serve as the basis for bias, discrimination, or disparate results of various sorts.

"With FR recognition being deployed in law enforcement, border security, retail, and airports, differential results across groups can serve as the basis for bias, discrimination, or disparate results of various sorts."

It is challenging, of course, to determine how high accuracy levels should be before FR is deployed on a widespread basis. Clearly, the degree depends on the magnitude of the impact on people's lives. In law enforcement, where misidentification can lead to arrest or incarceration, the accuracy level should be very high. As a warning sign, an in-depth Cardiff University analysis found thousands of false-positive matches in Australia, as well as FR systems that did not operate well in low-light conditions.

Employment applications also have major consequences, so high standards need to be set there as well. Companies that discriminate or whose FR generates biased results should be held accountable for those issues. They should be subject to meaningful penalties for racial biases or other abuses associated with facial recognition products.

5. Deploy third-party assessments

Third-party assessments represent a way for organizations to have greater confidence in their FR products and services. Consumers want to make sure products do what they are supposed to do and avoid problematic applications. Having "gold star" ratings or the equivalent of the government-backed Energy

Star ratings would help others know what the capabilities are and what their worries should be. Third-party organizations can help build confidence in the way in which FR is being deployed and the accuracy of its assessments.

6. Reduce collateral information collection

Some FR applications compile a considerable amount of collateral information unrelated to the primary purpose of the recording. For example, when police officers wearing body cameras go to the scene of an incident, they pick up film not just of possible perpetrators but innocent bystanders who just happen to cross their line of sight. Unless the evidence is clearly related to the incident under investigation, there is no need for unrelated information to be retained by law enforcement beyond the investigation. Images can be blurred or simply destroyed when they no longer have investigative value.

7. Require opt-in for marketing applications

Opting-in should be required for marketing applications linked to FR. This includes FR that attaches your name to a commercial profile and delivers ads straight to you. Since consumers are very sensitive about personal privacy, FR for marketing purposes should be limited with opt-in provisions. That will allow consumers to restrict their exposure to the level with which they are comfortable.

There are other situations when opt-out would be desirable. In cases of low threat where there is little demonstrable need for long-term data storage, giving people the means to remove their images would reassure the public that pictures are not being used for purposes they do not approve would be valuable. A right to be forgotten would build public confidence in FR applications.

8. Develop technical standards

Technical standards are a common way for companies to safeguard products. As an illustration, when mobile technology was under development, experts gathered to determine common standards for communications, security, and interoperability. All smartphones had to meet these specifications in order to be sold. International bodies extended standards overseas, which made it easy for smartphones to work across the globe.

"FR should be held accountable in the same way other dual-use products are governed."

The same logic applies to FR in that there should be technical standards that make sure applications are secure and privacy is protected. Common rules would enable people to address widespread fears and work to limit problems that emanate from them. FR should be held accountable in the same way other dual-use products are governed. The Institute of Electrical and Electronics Engineers is working on FR standards right now—and it is time for the National Institute of Standards and Technology (NIST) to do the same thing.

9. Certify certain usages

The security aspects of enterprise systems are certified by organizations such as the International Organisation for Standardization (ISO). That group specifies whether particular products meet government rules and regulations, with third-party organizations testing products for compliance. That helps consumers and businesses to know what applications do and whether they meet specified rules. It

is a way to guarantee a high degree of uniformity, which can boost confidence in sectors as a whole.

In the United States, NIST certifies products. That agency uses public databases to test FR and certify particular applications. However, there have been complaints that NIST testing relies upon scraped data from private websites and that its tests don't translate to everyday scenarios. Rather than testing with a broad range of data, it focuses too much on applications linked to law enforcement and that test passage depends on image quality and operational features. Certification should combine automated testing with human checks in order to have the most reliable testing and certification.

10. Make sure FR testing is based on representative data and in actual field conditions

It is important that FR testing for certification, technical standards, and government compliance be based on representative and non-idiosyncratic data, taking place in the field under real-life conditions. In a commercial system with lots of proprietary information, it is crucial that representative databases be used for baseline testing and product certification. Idiosyncratic information sources such as images collected for mug shots are not representative of the overall population and therefore have limited utility for testing purposes. People need to have confidence that FR testing relies upon a broad range of images, field conditions, and population groupings. With FR performance levels varying by light conditions and image resolution, it is crucial to have reliable field testing.

---

policymakers and the public. The conclusions and recommendations of any Brookings publication are solely those of its author(s), and do not reflect the views of the Institution, its management, or its other scholars.

Microsoft provides support to The Brookings Institution's Artificial Intelligence and Emerging Technology (AIET) Initiative. The findings, interpretations, and conclusions in this report are not influenced by any donation. Brookings recognizes that the value it provides is in its absolute commitment to quality, independence, and impact. Activities supported by its donors reflect this commitment.

Report Produced by **Center for Technology Innovation**