Inquiring Minds– 27 September 2019

Melissa Butler, Moderator

Should Facial Recognition Technology Be Regulated?

It's popping up everywhere—tiny cameras linked to facial recognition software allowing users to gather and analyze data about passers-by as well as to identify particular individuals. But the use of this technology presents practical, legal, and ethical issues.

What are the costs and benefits of this technology? Do the costs outweigh the benefits or vice versa?

Are the problems inherent in the technology itself or merely flaws that will be eliminated as it gets better?

Should it be regulated? How? By government? By self-regulation among tech companies?

Should states and localities follow San Francisco's example and ban its use by government? Should tech companies be prohibited from providing these tools to repressive regimes?

Does facial recognition technology render privacy obsolete? Should we have a "right to be forgotten?" a "right to be unnoticed"? Or are these Luddite fantasies?

Jessica Davis, "The Problem with AI Facial Recognition," *InformationWeek*, 2/11/19

AI facial recognition algorithms are far from perfect. . . .Shelf-mounted cameras. . . give physical stores demographic information that could guide how they market to individual customers. It's something that could give them a competitive edge against online retailers such as Amazon, that have been leveraging customer data all along.

But using cameras to capture photos of your customers in a way they may not even notice seems like it could be crossing that line between cool technology and creepy technology. Is it too invasive? Beyond that, there could be other problems, too. What if the software misidentifies a man as a woman and offers him a discount on feminine hygiene products? What are the consequences?

The consequences may not be hugely significant in the retail setting. A customer could get miffed, talk about it on social media, and not go back to that store for a while. But could the consequences be higher in other applications of machine vision and AI-driven facial recognition software?

Turns out, there is a great deal of concern about AI facial recognition software, which is commercially available from a number of big vendors, including Microsoft, IBM, and Amazon.

Most recently the focus is on a study about how some commercial algorithms are not as accurate at identifying darker-skinned people and women as they are at identifying lighter-skinned men. It's a topic that's been covered before. For instance, in July 2018 the American Civil Liberties Union (ACLU) applied an Amazon algorithm to photos of members of the US Congress and the algorithm identified 28 of them as people who have been arrested for a crime.

[A study]co-authored by MIT graduate student Joy Buolamwini, who also is founder of The Algorithmic Justice League, an organization that describes itself as dedicated to fighting bias in algorithms. . . says that these algorithms are best at identifying lighter-skinned men. Their performance isn't as good when identifying women or darker-skinned people. . . .In machine learning, the results can be biased or inaccurate based on the volume and type of training data used. For instance, Amazon used machine learning to screen resumes of job applicants and ended up with a pool of mostly male candidates. That's probably because the historic pool of data used to train the algorithms included more men than women.

By adding more data or data sources to the pool used to train the algorithm, vendors may have improved the accuracy of their AI facial recognition systems.. . . [S]ystems allow organizational customers to set a threshold or confidence level. This may be set based on the type of action the organization plans for the results.. . .For instance, in terms of gender, these systems may decide that someone is male. But they will also provide a confidence score, essentially saying that they are 67% (or some other percentage) certain that the person is male. Retailers have already set the level of confidence score they are willing to accept. So if someone is deduced to be male with a 67% confidence score and the retailer has set the threshold level at 60%, the customer will see the offer customized for a man. If the retailer sets the threshold score at 70%, the customer's 67% score would not meet that threshold and the customer would see a generic offer that could be made to any customer, male or female.

If the stakes are high, for example, in a law enforcement application where someone's life trajectory may change, the organization may set the threshold at 99%. If the stakes are not as high they may set the threshold at a much lower level.

Heightened privacy awareness

Still, are there privacy issues with collecting customer images, particularly in the age of the EU's General Data Protection Regulation (GDPR) and other new data privacy laws? [In some applications] the images of customers are not retained. However, aggregated data about the demographics of the customers who visit a particular display is retained and analyzed to help retailers gain insights into their customers.

Should enterprises be experimenting with AI facial recognition software? That probably depends on the application and the level of risk that is entailed. For physical store retailers looking to gain an edge against their digital competitors, these applications could open up a world of data and insights that have not before been available.

Other machine vision technology that looks for matches of images of people against a database of known images has been used to fight child trafficking in the case of Thorn. Likely matches are surfaced by the AI and the human-in-the-loop makes the final determination of whether a match has been found. The benefit of using AI in this type of application, whether it is identifying missing children or identifying a suspected terrorist at a crowded sporting event in real time, is that the algorithm can analyze and make a match in seconds. But in these kinds of high-risk applications, having a human in the loop to make the final call is probably an important safeguard against mistakes in this nascent technology.

That's something that enterprises should keep in mind. This technology is still new, so obviously it's not perfect, and it should be handled with care. Also, like many emerging technologies, it lacks many regulations to govern its use. So far. Those regulations will likely be coming in the years ahead.

Nicole Lindsey. “San Francisco Bans Facial Recognition Technology” May 24, 2019

Amidst growing concerns about the privacy issues raised by facial recognition technology, San Francisco became the first major U.S. city to prohibit the technology. As part of a broader anti-surveillance ordinance (“Stop Secret Surveillance Ordinance”) introduced by Supervisor Aaron Peskin and approved overwhelmingly by the city’s Board of Supervisors, the facial recognition technology ban will apply to all of the city’s 53 different departments, including the San Francisco Police Department.

Details of the facial recognition technology ban

The new ordinance is expected to go into effect in less than 30 days, meaning that it will have a direct and tangible impact on how San Francisco deals with surveillance technology. The city – known as one of the most tech-savvy and tech-friendly in the world – is now at the forefront of issues like privacy and the responsible use of technology. As the Board of Supervisors noted, there is a fine line between “good policing” and becoming a “police state”, and the recent rapid improvements in facial recognition technology (which now enable near real-time tracking of individuals, even in large crowds) have raised many concerns that the technology was advancing too far, too fast. By the city slapping a total ban on the technology, it would help to prevent the proverbial genie from escaping out of the bottle.

It is important to note, however, that the ban on facial recognition technology for city agencies does not apply to private individuals or private businesses. Individuals, for example, are still able to use their home security cameras. And they will still be

---

able to use the facial recognition technology built into digital mobile devices like the iPhone. Private businesses in the city – such as the local supermarket – would still be able to use surveillance technologies in order to prevent or deter crime from happening on their premises.. . . And, even though the San Francisco Police Department is no longer able to authorize the use of new facial recognition technology without first getting the approval of the city government's Board of Supervisors, it will still be able to use existing surveillance technology that is already in place, such as police body cameras and license plate readers. Moreover, if needed in a criminal investigation, the San Francisco Police Department would be able to request that footage from surveillance cameras be used to help solve the case. Finally, federally controlled facilities at both the San Francisco International Airport and the Port of San Francisco – two entry points into California from possibly dangerous foreign destinations – would still be able to use face recognition technology.

There are several major issues that still need to be worked out for facial recognition technology in order to make it more palatable for city leaders. One major issue, for example, is the perceived notion that facial recognition technology comes with its share of biases and inaccuracies that could lead to a disproportionately negative effect on certain demographic groups. A recent MIT and Georgetown study found that, since facial recognition technology is almost always "trained" using datasets featuring faces of white males, that it is not as accurate in finding matches for women and people of color. If used by a police department, for example, it might lead to a lot of false positives and the wrongful arrest of people of color. On a related note, another fear is that facial recognition technology will be deployed in a way that will disproportionately impact certain neighborhoods in the city. In the name of cracking down on crime in notoriously high-crime neighborhoods, for example, police officers might deploy facial recognition technology primarily in parts of the city that are known to have large African-American or Hispanic-American populations.

Given these issues and the overall controversy swirling around facial recognition technology in general, it is perhaps no surprise that civil liberties and privacy advocates – such as the ACLU and the Electronic Frontier Foundation (EFF) – have wholeheartedly supported the new ordinance. As Matt Cagle, a technology and civil liberties attorney with the ACLU of Northern California, points out, a total ban on facial recognition technology is needed for the smooth functioning of a "healthy democracy." There is a reasonable expectation that, if you are walking around the city, that someone is not tracking your every move.

However, not everyone is supportive of the surveillance tech ban. Some tech lobbying groups, for example, have said that a more practical step would have been a "moratorium" on the use of the technology, rather than an outright ban. And one anti-crime group known as Stop Crime SF has said that the ban would remove a possible deterrent to specific sorts of crime, such as property crime. Tech companies

that make the facial recognition systems, though, have been notably absent from the debate. In general, they support "safeguards" on the use of the technology, but as might be expected, do not support a total and outright ban on facial recognition software.

Possible implications of the San Francisco facial recognition technology ban
So will other cities follow suit, now that San Francisco has decided to move ahead with its facial recognition tech ban? Two other cities – Oakland, California and Somerville, Massachusetts – could be next up on the list of municipalities to ban facial recognition technology. And, overall, there seems to be growing momentum for cutting back on surveillance technology and protecting consumer privacy. At least seven different California municipalities are working on surveillance ordinances of their own, and there is even mounting speculation that some type of action might be taken at the federal level. In March, for example, the U.S. Senate introduced a new bill that would have placed new safeguards around the use of facial recognition technology.

San Francisco is the first major U.S. city to prohibit facial recognition technology to protect citizen privacy but the ban does not apply to private sector.

Going forward, it will be interesting to see if legislation at both the local level and the federal level can keep pace with the remarkable amount of innovation happening in the technology space. As various forms of artificial intelligence (AI) become more and more powerful, it is perhaps only a matter of time before more cities and states take a step back to consider what the impact might be on personal privacy.

An unusual consensus emerged recently between artificial intelligence researchers, activists, lawmakers and many of the largest technology companies: Facial recognition software breeds bias, risks fueling mass surveillance and should be regulated. Deciding on effective controls and acting on them will be a lot harder.

On Tuesday, the Algorithmic Justice League and the Center of Privacy & Technology at Georgetown University Law Center unveiled the Safe Face Pledge, which asks companies not to provide facial AI for autonomous weapons or sell to law enforcement unless explicit laws are debated and passed to allow it. Microsoft Corp. last week said the software carries significant risks and proposed rules to combat the threat. Research group AI Now, which includes AI researchers from Google and other companies, issued a similar call…..None of the biggest makers of the software – companies like Microsoft, Google, Amazon.com Inc., Facebook Inc. and IBM – has signed the Safe Face Pledge yet.

Large tech companies may be reluctant to commit to a pledge like this, even if they're concerned about negative consequences of the software. That's because it could mean walking away from lucrative contracts for the emerging technology. The market for video surveillance gear is worth $18.5 billion a year, and AI-powered

equipment for new forms of video analysis is an important emerging category, according to researcher IHS Markit. Microsoft and Facebook said they're reviewing the pledge. Google declined to comment.

"There are going to be some large vendors who refuse to sign or are reluctant to sign because they want these government contracts," said Laura Moy, executive director of the Center on Privacy & Technology.

Microsoft is still selling facial recognition software to governments, a fact that the American Civil Liberties Union took the company to task for this week. It asked Microsoft to halt the sales and join the organization's call for a federal moratorium on government use of the technology.

The use of facial recognition for surveillance, policing and immigration is being questioned because researchers, including Buolamwini, have shown the technology isn't accurate enough for critical decisions and performs worse on darker-skinned people.

Providers have responded differently to the scrutiny. Microsoft is defending government contracts generally, while asking for laws to regulate the space. Amazon took issue with research by the ACLU into the Rekognition program it sells to police departments, but the company has also said it's working to better educate police on how to use the software. Companies including Microsoft, Facebook and Axon, a maker of police body cameras, have formed AI ethics boards and Google published a set of more-general AI principles in June.

The Safe Face Pledge asks companies to "show value for human life, dignity and rights, address harmful bias, facilitate transparency" and make these commitments part of their business practices. This includes not selling facial recognition software to identify targets where lethal force may be used. The pledge also commits companies to halt sales of face AI products that are not "subject to public scrutiny, inspection, and oversight."

…While employees and customers can pressure companies to act ethically with regard to AI, more attention needs to be focused on laws and government oversight, said Ryan Calo, a law professor at the University of Washington, who is on the board of AI Now and gets funding from Microsoft. Without broad regulation, if some companies refuse to sell the software, others will step in.

"We have been attempting to get companies to cease providing tools to the government, rather than trying to ensure the government doesn't do things we don't agree with," Calo said. "They are government agencies – we ought to be able to police them. We can't ask technology companies to make it all go away.

Check out the supplementary article by Andy Kessler called "Have No Fear of Facial Recognition" on Shellpoint.info.
