DOI:10.1145/3474096

Toby Walsh

Viewpoint
The Troubling Future for Facial Recognition Software

Considering the myriad perspectives of facial recognition technology.

GEORGE ORWELL'S NOVEL *1984* got one thing wrong. A surveillance state will not have people watching people, as the Stasi did in East Germany. Computers will be the ones watching people. Technology lets you perform surveillance at an industrial scale.

This is already happening in China, where facial recognition software is being used by law enforcement for catching relatively minor offenders such as jaywalkers to enabling much more disturbing activities such as tracking Uyghurs. The West has also seen a rise in the use of such software. For example, the controversial company Clearview AI has scraped approximately three billion photographs from the Web, which the company uses to sell facial recognition services to agencies including the U.S. Federal Bureau of Investigation.

Fortunately, pushback is starting to happen against these developments. In June 2020, IBM announced it would no longer sell, research, or develop facial recognition software. Amazon and Microsoft quickly followed suit, announcing moratoria on selling such services to the police pending federal regulation.

Local and national governments in the U.S. are hitting the pause button. San Francisco, Boston, and several other cities have introduced bans.

IMAGE BY VARUNA

And the Facial Recognition and Biometric Technology Moratorium Act introduced by Democratic lawmakers in June 2020 attempts, as the name suggests, to impose a moratorium on the use of facial recognition software. Professional societies such as ACM, along with organizations including Human Rights Watch and the UN, have also called for regulation.

A major ethical concern behind many of these calls is bias. Researchers including MIT's Joy Buolamwini have demonstrated the technology often works better on men than women, better on white people than Black people, and worst of all on Black women. And while some facial recognition software has been improved in response, significant biases remain. In June 2020, in the first known case of its type, a man in Detroit was arrested in front of his family for burglary because he was mistakenly identified by facial recognition software. It may come as no surprise the man was Black.

Society should, however, be very careful of calls for a moratorium on

the use of facial recognition software because of such biases. It is, of course, entirely unacceptable to see Black people incarcerated due to a biased algorithm. But we risk shooting ourselves in the foot if we use bias as reason to call for regulation.

One day, facial recognition software may well be less biased than humans. Our ability to recognize faces is highly variable, has a significant hereditary component, and is often biased toward people of one's own race. Therefore, beating humans at face recognition is not a tall order.

There is promising, if somewhat slow, progress on making facial recognition software less biased. This ranges from using more representative datasets (such as the 10K U.S. Adult Face database that consists of a large demographically balanced set of faces) to debiasing methods (such as selectively resampling biased data to make it less so).

One day, just as with playing chess, reading X-rays, or translating spoken Mandarin into written English, computers might therefore easily outperform humans at facial recognition and do so in a much less biased way than humans. And at this point, government agencies will be morally obliged to use facial recognition software since it will make fewer mistakes than humans do.

Banning facial recognition because of its biases is therefore problematic. The problem is it overlooks the many other harms automated facial recognition may bring into people's lives. The technology will, for example, challenge many fundamental rights such as the right to privacy and the right to protest.

Previously, if you were in a large crowd protesting about, say, the climate emergency or Black Lives Matter, you were anonymous. Now, software can identify you in real time. For instance, in 2018, Chinese police found a criminal at a music concert attended by 60,000 people using some facial recognition software. Human eyes cannot do that. Only computer eyes can.

Local and national governments in the U.S. are hitting the pause button.

There is promising, if somewhat slow, progress on making facial recognition software less biased.

There are, of course, benefits from the use of facial recognition software. In 2018, for example, police took such software into orphanages in New Delhi and were able to reunite nearly 3,000 children with their parents. This was a great good.

How then do we proceed? Many people must come together to decide how to navigate this future: ethicists, technologists, politicians, and sociologists, to name just a few. Regulation is likely to be key and we have already seen example of bans being successful implemented in public settings.

The ethics of facial recognition software is complex. It begins with the collection of datasets often without explicit consent from the people concerned. And it then ranges over many issues, from its use on vulnerable populations like the Uyghurs in China, to the dilemma of a technology with both good and bad uses.

Face recognition may be one of the first uses of AI to trouble us greatly. But it will not be the last. Ultimately, this is about the world we will invent. And all of society must be engaged in this debate. C

Toby Walsh (tw@cse.unsw.edu.au) is Professor of Artificial Intelligence at the University of New South Wales in Sydney Australia and at CSIRO Data61. He is a fellow of both the ACM and the Australian Academy of Science, and a strong advocate for limits to ensure AI is used to improve our lives. He has authored three books on AI for general audiences, the most recent is *Machines Behaving Badly: The Morality of AI*.