Microsoft

The Official Microsoft Blog

Microsoft On the Issues The AI Blog
Transform

All Microsoft ⌄

Facial recognition technology: The need for public regulation and corporate responsibility

Jul 13, 2018 | Brad Smith - President

[f] [twitter] [in]

All tools can be used for good or ill. Even a broom can be used to sweep the floor or hit someone over the head. The more powerful the tool, the greater the benefit or damage it can cause. The last few months have brought this into stark relief when it comes to computer-assisted facial recognition – the ability of a computer to recognize people's faces from a photo or through a camera. This technology can catalog your photos, help reunite families or potentially be misused and abused by private companies and public authorities alike.

Facial recognition technology raises issues that go to the heart of fundamental human rights protections like privacy and freedom of expression. These issues heighten responsibility for tech companies that create these products. In our view, they

Related Stories

Dec 6, 2018 | Brad Smith

Facial recognition: It's time for action >

Apr 29, 2019 | Julie Brill

Our support for meaningful privacy protection through the Washington Privacy Act >

Dec 17, 2018 | Rich Sauer

Six principles to guide Microsoft's facial recognition work >

Feb 11, 2019 | Brad Smith

Next Generation Washington: Our priorities for 2019 >

Stay Connected

also call for thoughtful government regulation and for the development of norms around acceptable uses. In a democratic republic, there is no substitute for decision making by our elected representatives regarding the issues that require the balancing of public safety with the essence of our democratic freedoms. Facial recognition will require the public and private sectors alike to step up – and to act.

We've set out below steps that we are taking, and recommendations we have for government regulation.

First, some context

Facial recognition technology has been advancing rapidly over the past decade. If you've ever seen a suggestion on Facebook or another social media platform to tag a face with a suggested name, you've seen facial recognition at work. A wide variety of tech companies, Microsoft included, have utilized this technology the past several years to turn time-consuming work to catalog photos into something both instantaneous and useful.

So, what is changing now? In part it's the ability of computer vision to get better and faster in recognizing people's faces. In part this improvement reflects better cameras, sensors and machine learning capabilities. It also reflects the advent of larger and larger datasets as more images of people are stored online. This improvement also reflects the ability to use the cloud to connect all this data and facial recognition technology with live cameras that capture images of people's faces and seek to identify them – in more places and in real time.

Advanced technology no longer stands apart from society; it is becoming deeply infused in our personal and professional lives. This means the potential uses of facial recognition are myriad. At an elementary level, you might use it to catalog and search your photos, but that's just the beginning. Some uses are already improving security for computer users, like recognizing your face instead of requiring a password to access many Windows laptops or iPhones, and in the future a device like an automated teller machine.

Some emerging uses are both positive and potentially even profound. Imagine finding a young missing child by recognizing her as she is being walked down the street. Imagine helping the police to identify a terrorist bent on destruction as he walks into the arena where you're attending a

sporting event. Imagine a smartphone camera and app that tells a person who is blind the name of the individual who has just walked into a room to join a meeting.

But other potential applications are more sobering. Imagine a government tracking everywhere you walked over the past month without your permission or knowledge. Imagine a database of everyone who attended a political rally that constitutes the very essence of free speech. Imagine the stores of a shopping mall using facial recognition to share information with each other about each shelf that you browse and product you buy, without asking you first. This has long been the stuff of science fiction and popular movies – like "Minority Report," "Enemy of the State" and even "1984" – but now it's on the verge of becoming possible.

Perhaps as much as any advance, facial recognition raises a critical question: what role do we want this type of technology to play in everyday society?

The issues become even more complicated when we add the fact that facial recognition is advancing quickly but remains far from perfect. As reported widely in recent months, biases have been found in the performance of several fielded face recognition technologies. The technologies worked more accurately for white men than for white women and were more accurate in identifying persons with lighter complexions than people of color. Researchers across the tech sector are working overtime to address these challenges and significant progress is being made. But as important research has demonstrated, deficiencies remain. The relative immaturity of the technology is making the broader public questions even more pressing.

Even if biases are addressed and facial recognition systems operate in a manner deemed fair for all people, we will still face challenges with potential failures. Facial recognition, like many AI technologies, typically have some rate of error even when they operate in an unbiased way. And the issues relating to facial recognition go well beyond questions of bias themselves, raising critical questions about our fundamental freedoms.

Politics meets Silicon Valley

In recent weeks, the politics of the United States have become more intertwined with these technology developments on the West Coast. One week in the middle of June put the issues raised by facial recognition technology in bold relief for me and

other company leaders at Microsoft. As the country was transfixed by the controversy surrounding the separation of immigrant children from their families at the southern border, a tweet about a marketing blog Microsoft published in January quickly blew up on social media and sparked vigorous debate. The blog had discussed a contract with the U.S. Immigration and Customs Enforcement, or ICE, and said that Microsoft had passed a high security threshold; it included a sentence about the potential for ICE to use facial recognition.

We've since confirmed that the contract in question isn't being used for facial recognition at all. Nor has Microsoft worked with the U.S. government on any projects related to separating children from their families at the border, a practice to which we've strongly objected. The work under the contract instead is supporting legacy email, calendar, messaging and document management workloads. This type of IT work goes on in every government agency in the United States, and for that matter virtually every government, business and nonprofit institution in the world. Some nonetheless suggested that Microsoft cancel the contract and cease all work with ICE.

The ensuing discussion has illuminated broader questions that are rippling across the tech sector. These questions are not unique to Microsoft. They surfaced earlier this year at Google and other tech companies. In recent weeks, a group of Amazon employees has objected to its contract with ICE, while reiterating concerns raised by the American Civil Liberties Union (ACLU) about law enforcement use of facial recognition technology. And Salesforce employees have raised the same issues related to immigration authorities and these agencies' use of their products. Demands increasingly are surfacing for tech companies to limit the way government agencies use facial recognition and other technology.

These issues are not going to go away. They reflect the rapidly expanding capabilities of new technologies that increasingly will define the decade ahead. Facial recognition is the technology of the moment, but it's apparent that other new technologies will raise similar issues in the future. This makes it even more important that we use this moment to get the direction right.

The need for government regulation

The only effective way to manage the use of technology by a government is for the government proactively to manage this use itself. And if there are concerns about how a technology

will be deployed more broadly across society, the only way to regulate this broad use is for the government to do so. This in fact is what we believe is needed today – a government initiative to regulate the proper use of facial recognition technology, informed first by a bipartisan and expert commission.

While we appreciate that some people today are calling for tech companies to make these decisions – and we recognize a clear need for our own exercise of responsibility, as discussed further below – we believe this is an inadequate substitute for decision making by the public and its representatives in a democratic republic. We live in a nation of laws, and the government needs to play an important role in regulating facial recognition technology. As a general principle, it seems more sensible to ask an elected government to regulate companies than to ask unelected companies to regulate such a government.

Such an approach is also likely to be far more effective in meeting public goals. After all, even if one or several tech companies alter their practices, problems will remain if others do not. The competitive dynamics between American tech companies – let alone between companies from different countries – will likely enable governments to keep purchasing and using new technology in ways the public may find unacceptable in the absence of a common regulatory framework.

It may seem unusual for a company to ask for government regulation of its products, but there are many markets where thoughtful regulation contributes to a healthier dynamic for consumers and producers alike. The auto industry spent decades in the 20th century resisting calls for regulation, but today there is broad appreciation of the essential role that regulations have played in ensuring ubiquitous seat belts and air bags and greater fuel efficiency. The same is true for air safety, foods and pharmaceutical products. There will always be debates about the details, and the details matter greatly. But a world with vigorous regulation of products that are useful but potentially troubling is better than a world devoid of legal standards.

That's why Microsoft called for national privacy legislation for the United States in 2005 and why we've supported the General Data Protection Regulation in the European Union. Consumers will have more confidence in the way companies use their sensitive personal information if there are clear rules of the

road for everyone to follow. While the new issues relating to facial recognition go beyond privacy, we believe the analogy is apt.

It seems especially important to pursue thoughtful government regulation of facial recognition technology, given its broad societal ramifications and potential for abuse. Without a thoughtful approach, public authorities may rely on flawed or biased technological approaches to decide who to track, investigate or even arrest for a crime. Governments may monitor the exercise of political and other public activities in ways that conflict with longstanding expectations in democratic societies, chilling citizens' willingness to turn out for political events and undermining our core freedoms of assembly and expression. Similarly, companies may use facial recognition to make decisions without human intervention that affect our eligibility for credit, jobs or purchases. All these scenarios raise important questions of privacy, free speech, freedom of association and even life and liberty.

So what issues should be addressed through government regulation? That's one of the most important initial questions to address. As a starting point, we believe governments should consider the following issues, among others:

- Should law enforcement use of facial recognition be subject to human oversight and controls, including restrictions on the use of unaided facial recognition technology as evidence of an individual's guilt or innocence of a crime?

- Similarly, should we ensure there is civilian oversight and accountability for the use of facial recognition as part of governmental national security technology practices?

- What types of legal measures can prevent use of facial recognition for racial profiling and other violations of rights while still permitting the beneficial uses of the technology?

- Should use of facial recognition by public authorities or others be subject to minimum performance levels on accuracy?

- Should the law require that retailers post visible notice of their use of facial recognition technology in public spaces?

- Should the law require that companies obtain prior consent before collecting individuals' images for facial recognition? If so, in what situations and places should this apply? And

what is the appropriate way to ask for and obtain such consent?

- Should we ensure that individuals have the right to know what photos have been collected and stored that have been identified with their names and faces?

- Should we create processes that afford legal rights to individuals who believe they have been misidentified by a facial recognition system?

This list, which is by no means exhaustive, illustrates the breadth and importance of the issues involved.

Another important initial question is how governments should go about addressing these questions. In the United States, this is a national issue that requires national leadership by our elected representatives. This means leadership by Congress. While some question whether members of Congress have sufficient expertise on technology issues, at Microsoft we believe Congress can address these issues effectively. The key is for lawmakers to use the right mechanisms to gather expert advice to inform their decision making.

On numerous occasions, Congress has appointed bipartisan expert commissions to assess complicated issues and submit recommendations for potential legislative action. As the Congressional Research Service (CRS) noted last year, these commissions are "formal groups established to provide independent advice; make recommendations for changes in public policy; study or investigate a particular problem, issue, or event; or perform a duty." Congress' use of the bipartisan "9/11 Commission" played a critical role in assessing that national tragedy. Congress has created 28 such commissions over the past decade, assessing issues ranging from protecting children in disasters to the future of the army.

We believe Congress should create a bipartisan expert commission to assess the best way to regulate the use of facial recognition technology in the United States. This should build on recent work by academics and in the public and private sectors to assess these issues and to develop clearer ethical principles for this technology. The purpose of such a commission should include advice to Congress on what types of new laws and regulations are needed, as well as stronger practices to ensure proper congressional oversight of this technology across the executive branch.

Issues relating to facial recognition go well beyond the borders of the United States. The questions listed above – and no doubt others – will become important public policy issues around the world, requiring active engagement by governments, academics, tech companies and civil society internationally. Given the global nature of the technology itself, there likely will also be a growing need for interaction and even coordination between national regulators across borders.

Tech sector responsibilities

The need for government leadership does not absolve technology companies of our own ethical responsibilities. Given the importance and breadth of facial recognition issues, we at Microsoft and throughout the tech sector have a responsibility to ensure that this technology is human-centered and developed in a manner consistent with broadly held societal values. We need to recognize that many of these issues are new and no one has all the answers. We still have work to do to identify all the questions. In short, we all have a lot to learn. Nonetheless, some initial conclusions are clear.

First, it's incumbent upon those of us in the tech sector to continue the important work needed to reduce the risk of bias in facial recognition technology. No one benefits from the deployment of immature facial recognition technology that has greater error rates for women and people of color. That's why our researchers and developers are working to accelerate progress in this area, and why this is one of the priorities for Microsoft's Aether Committee, which provides advice on several AI ethics issues inside the company.

As we pursue this work, we recognize the importance of collaborating with the academic community and other companies, including in groups such as the Partnership for AI. And we appreciate the importance not only of creating data sets that reflect the diversity of the world, but also of ensuring that we have a diverse and well-trained workforce with the capabilities needed to be effective in reducing the risk of bias. This requires ongoing and urgent work by Microsoft and other tech companies to promote greater diversity and inclusion in our workforce and to invest in a broader and more diverse pipeline of talent for the future. We're focused on making progress in these areas, but we recognize that we have much more work to do.

Second, and more broadly, we recognize the need to take a principled and transparent approach in the development and application of facial recognition technology. We are undertaking work to assess and develop additional principles to govern our facial recognition work. We've used a similar approach in other instances, including trust principles we adopted in 2015 for our cloud services, supported in part by transparency centers and other facilities around the world to enable the inspection of our source code and other data. Similarly, earlier this year we published an overall set of ethical principles we are using in the development of all our AI capabilities.

As we move forward, we're committed to establishing a transparent set of principles for facial recognition technology that we will share with the public. In part this will build on our broader commitment to design our products and operate our services consistent with the UN's Guiding Principles on Business and Human Rights. These were adopted in 2011 and have emerged as the global standard for ensuring corporate respect for human rights. We periodically conduct Human Rights Impact Assessments (HRIAs) of our products and services, and we're currently pursuing this work with respect to our AI technologies.

We'll pursue this work in part based on the expertise and input of our employees, but we also recognize the importance of active external listening and engagement. We'll therefore also sit down with and listen to a variety of external stakeholders, including customers, academics and human rights and privacy groups that are focusing on the specific issues involved in facial recognition. This work will take up to a few months, but we're committed to completing it expeditiously .

We recognize that one of the difficult issues we'll need to address is the distinction between the development of our facial recognition services and the use of our broader IT infrastructure by third parties that build and deploy their own facial recognition technology. The use of infrastructure and off-the-shelf capabilities by third parties are more difficult for a company to regulate, compared to the use of a complete service or the work of a firm's own consultants, which readily can be managed more tightly. While nuanced, these distinctions will need consideration.

Third, in the meantime we recognize the importance of going more slowly when it comes to the deployment of the full range

of facial recognition technology. Many information technologies, unlike something like pharmaceutical products, are distributed quickly and broadly to accelerate the pace of innovation and usage. "Move fast and break things" became something of a mantra in Silicon Valley earlier this decade. But if we move too fast with facial recognition, we may find that people's fundamental rights are being broken.

For this reason, based in part on input from the Aether Committee, we're moving more deliberately with our facial recognition consulting and contracting work. This has led us to turn down some customer requests for deployments of this service where we've concluded that there are greater human rights risks. As we're developing more permanent principles, we will continue to monitor the potential uses of our facial recognition technologies with a view to assessing and avoiding human rights abuses.

In a similar vein, we're committed to sharing more information with customers who are contemplating the potential deployment of facial recognition technology. We will continue work to provide customers and others with information that will help them understand more deeply both the current capabilities and limitations of facial recognition technology, how these features can and should be used, and the risks of improper uses.

Fourth, we're committed to participating in a full and responsible manner in public policy deliberations relating to facial recognition. Government officials, civil liberties organizations and the broader public can only appreciate the full implications of new technical trends if those of us who create this technology do a good job of sharing information with them. Especially given our urging of governments to act, it's incumbent on us to step forward to share this information. As we do so, we're committed to serving as a voice for the ethical use of facial recognition and other new technologies, both in the United States and around the world.

We recognize that there may be additional responsibilities that companies in the tech sector ought to assume. We provide the foregoing list not with the sense that it is necessarily complete, but in the hope that it can provide a good start in helping to move forward.

Some concluding thoughts

Finally, as we think about the evolving range of technology uses, we think it's important to acknowledge that the future is not simple. A government agency that is doing something objectionable today may do something that is laudable tomorrow. We therefore need a principled approach for facial recognition technology, embodied in law, that outlasts a single administration or the important political issues of a moment.

Even at a time of increasingly polarized politics, we have faith in our fundamental democratic institutions and values. We have elected representatives in Congress that have the tools needed to assess this new technology, with all its ramifications. We benefit from the checks and balances of a Constitution that has seen us from the age of candles to an era of artificial intelligence. As in so many times in the past, we need to ensure that new inventions serve our democratic freedoms pursuant to the rule of law. Given the global sweep of this technology, we'll need to address these issues internationally, in no small part by working with and relying upon many other respected voices. We will all need to work together, and we look forward to doing our part.

Tags: [AI](#), [Brad Smith](#), [facial recognition technology](#), [policy](#)

What's new	Microsoft Store	Education	Enterprise	Developer	Company
NEW Surface Pro 6	Account profile	Microsoft in education	Microsoft Azure	Microsoft Visual Studio	Careers
NEW Surface Laptop 2	Download Center	Office for students	Microsoft Industry	Windows Dev Center	About Microsoft
NEW Surface Go	Microsoft Store support	Office 365 for schools	Data platform	Developer Network	Company news
Xbox One X	Returns		Find a solution provider		Privacy at Microsoft
Xbox One S	Order tracking	Deals for students & parents	Microsoft partner resources	TechNet	Investors
VR & mixed reality	Store locations			Microsoft developer program	Diversity and inclusion
Windows 10 apps	Buy online, pick up in store	Microsoft Azure in education	Microsoft AppSource	Channel 9	Accessibility

Office apps

Health

Financial services

Security

Office Dev Center

Microsoft Garage

🌐 English (United States)