JILIS

How Should Facial Recognition Be Regulated?

Woodrow Hartzog
Professor of Law and Computer Science
Northeastern University

Facial recognition systems are being rapidly adopted around the world. Governments are seeking to use them for real-time surveillance, identity authentication, and crime prevention. Industry is seeking to develop facial recognition tools for a wide array of uses in schools[1], shops, airports, places of worship, and virtually every other space in public life. Facial recognition is being pitched as the answer for truancy, shoplifting, violence[2], incivility, and a host of other societal problems. To government and industry, facial recognition can make life safe, civil, and easy.

But facial recognition is also the most dangerous and oppressive technology ever invented. It enables intrinsically oppressive surveillance because people will act differently if they know that everything they do is being watched and they are being identified everywhere they go. Yet, oppressive surveillance is just one of the many dangers of facial recognition. The technology also has a disproportionate impact on people of color and other vulnerable populations. Not only do marginalized populations bear the brunt of facial recognition surveillance first and most intensely, but these systems are also biased[3] towards them because they are typically under-represented in the training data and the design process. Facial recognition can also fuel harassment and violence because it makes stalking in real time and online easy, cheap, and compatible with mob behavior.

Facial recognition also is being implemented in systems that are being used to deny[4] people fundamental rights and opportunities based upon arbitrary tracking of people's movements, habits, relationships, interests and thoughts. Governments are deploying facial recognition as a tool to help them relentlessly and perfectly enforce[5] even minor laws such as jaywalking and petty theft, which can be suffocating to a society. Facial recognition systems can completely eliminate the practical obscurity[6] we rely upon to move about in our daily lives and amplifies surveillance capitalism[7], that is, "the unilateral claiming of private human experience as free raw

1) Schools Can Now Get Facial Recognition Tech for Free. Should They? (WIRED)
https://www.wired.com/story/realnetworks-facial-recognition-technology-schools/

2) In the Face of Danger, We're Turning to Surveillance (WIRED)
https://www.wired.com/story/surveillance-safety/

3) Federal study confirms racial bias of many facial-recognition systems, casts doubt on their expanding use (The Washington Post)
https://www.washingtonpost.com/technology/2019/12/19/federal-study-confirms-racial-bias-many-facial-recognition-systems-casts-doubt-their-expanding-use/

4) Anonymity, Faceprints, and the Constitution (Kimberly L. Wehle)
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2394838

5) 22 eerie photos show how China uses facial recognition to track its citizens as they travel, shop — and even use toilet paper (BUSINESS INSIDER)
https://www.businessinsider.com/how-china-uses-facial-recognition-technology-surveillance-2018-2

6) Obscurity and Privacy (Evan Selinger, Woodrow Hartzog)
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2439866

7) High tech is watching you (The Harvard Gazette)
https://news.harvard.edu/gazette/story/2019/03/harvard-professor-says-surveillance-capitalism-is-undermining-democracy/

material for translation into behavioral data."

As facial recognition systems proliferate, a large debate is being waged over how best to regulate them. In this short piece, I argue that, given the high likelihood of abuse, the procedural approaches to regulating facial recognition, such as consent regimes, notice and choice regimes, FIPs-based regimes, and warrant requirements, are doomed to fail. The only way to meaningfully limit the abuses of this technology is to ban it.

One of the most popular proposals to regulate facial recognition is to require that companies get peoples' consent before these systems can be used. Consent is one of the most common and intuitive concepts in all of privacy and data protection law. Unfortunately, it will never protect us. The problem is that meaningful, informed consent for facial surveillance is impossible to achieve at scale. Here's why:

First, consent is an illusion. The consent people give to facial surveillance is usually mediated by technology, which means it cannot help but be engineered to produce particular results. When it comes to consent, design is everything. The realities of technology at scale mean that the services we use must necessarily be built in a way that constraints our choices. Imagine a world where every user got to dictate their own terms in an open text box instead of a boilerplate terms of use. Companies would never get off the ground. Instead, we get boxes to check, buttons to press, switches to activate and deactivate, and other settings to fiddle with. Companies have an incentive for people to give their consent, so it is to their advantage to make users believe the consent they are giving is meaningful.

Interface design also nudges people by sending them signals and making tasks easier or harder which encourage them to act in predictable ways. Companies deploy manipulative 'dark patterns'[8] in user interfaces to exploit our built-in tendencies to prefer shiny, colorful buttons and ignore dull, grey ones. They may also shame us into feeling bad about withholding data or declining options[9]. They might simply make exercising consent possible but costly through forced work, subtle misdirection, and incentive tethering.

Requiring companies to ask for peoples' consent as the main way to regulate facial recognition paves the way for abuse and self-dealing at the margins. At scale, these margins matter. Even among those acting in good faith, we are left with the problem of relying on the notion of consent and choice to do more work for us than it's capable of. We risk looking around at the robust new frameworks for data protection, the rules built to encourage meaningful consent over personal information, patting ourselves on the back and saying 'mission accomplished,' when that isn't true. It wasn't even the right mission.

Second, consent requests are overwhelming. To hear people tell it, consent requests are something we can never get enough of. There seems to be no problem in privacy that cannot be remedied by chucking a few more switches, delete buttons, and privacy settings at people. Companies promise more and better consent requests, and then, when privacy harms happen, we collectively decide they should have asked for even more permissions.

Consent requirements are attractive in isolation. Who wouldn't want more power over things that affect our lives? But with this power often comes a practical obligation. If you do not actively deny consent, you are at risk. Companies can take your inaction as acquiescence. As I've written elsewhere,while you might remember to adjust your privacy settings on Facebook, what about Instagram, Twitter, Google, Amazon, Netflix, Snapchat, Siri, Cortana, Fitbit, Candy Crush, your smart TV, your robot vacuum cleaner, your WiFi-connected car, and your child's Hello Barbie?

Mobile apps can ask users for over 200 permissions and even the average app asks for about five. Imagine if even a fraction of them wanted your permission to use facial recognition technology on top of all that. Many can relate to the experience of a child asking for candy, over and over, until the requests become too much to ignore and we give in, simply to quiet

8) Dark Patterns
https://www.darkpatterns.org/

9) Are you sure? — how user interfaces undermine consent (Medium)
https://uxdesign.cc/how-user-interfaces-undermine-consent-81551cf48777

them. Willpower can feel like a finite, vulnerable, and subjective resource, and systems are designed to deplete and erode it. Once our willpower and ability to make choices has been compromised, the consent users have been given is meaningless.

People only have twenty four hours in a day (fewer if you sleep) and every company wants you to make choices. Even if we consolidated all of our choices, the tension between simplicity and nuance inherent in one of the most complex and fraught environments imaginable would seem irresolvable. This is because nuance gets glossed over when companies try to simplify and shorten information. Risk is either hidden through abstraction or made so explicit and voluminous we don't even know where to begin.

This is to say nothing of how ineffectual "notice" and "disclosure" requirements are regarding facial recognition. People often do not see such disclosures and even if they are made prominent have little ability to avoid them if facial recognition becomes ubiquitous. It is difficult for people to avoid stores entirely or even assert their right not to be surveilled at every store they visit. In public places, it might not even be possible to ask people for their consent, even when people want privacy in public. This means sometimes the only actual option for people who do not want to be surveilled will simply be to stay at home. This makes choosing privacy, a fundamental right and precondition for human flourishing, a costly option.

Finally, consent is too narrowly focused on individuals instead of society. Notions of individual consent don't fit well with privacy as a collective value. Just because lots of people agree to something does not mean it is good for society. When privacy is thought about in such individualistic, transactional terms, peoples' sense of privacy is always being negotiated against what others value.

What makes us think that the collective result of atomized decisions about facial recognition will be best for our overall privacy, anyway? Scholars have noted that a large body of research shows that peoples' privacy preferences are uncertain, contextually dependent, and malleable. The availability of knowledge doesn't necessarily translate into meaningfully informed decisions about being watched.

Lawmakers are often attracted to consent requirements and mandated disclosures, like the warnings in public places that facial recognition is being used, because they are cheap and counterbalance 'information disparity' —that is, the reality that companies and governments often know much more than those being watched regarding the wisdom of the decisions they make. People are being asked to consider the many different risks of facial recognition for every single encounter or context. This is an impossibly complex calculation to make about future risks and consequences.

If facial recognition is so dangerous that it requires formal permission, and choices can only meaningfully be made in elusive, demanding, and bounded environments with preconditions such as 'freely given, specific, informed, retractable, and unambiguous,' then why would we allow companies and governments to engage in what feels like a fiction, even under optimal conditions? If companies and governments ask for consent for face surveillance, they will almost always get it. We need more robust protections against systems as dangerous as facial recognition technologies.

Lawmakers have also proposed looking at the "fair information practices" to regulate facial recognition technology. The FIPs are the set of aspirational principles developed over the past fifty years used to model rules for responsible data practices. Thanks to the FIPs, data protection regimes around the world, including Japan's new data protection law, require those collecting and using personal information to be accountable, prudent, and transparent. They purportedly give data subjects control over their information by bestowing rights of correction and deletion.

FIPs-based regimes were relatively well-equipped for the first wave of personal computing. The FIPs provide a common set of values, which is necessary as data flows from one country to another at the speed of light. But facial recognition systems push FIPs principles like data minimization, transparency, choice, and access to the limit. They are procedural rules that actually help to authorize and enshrine facial surveillance through permission-granting protocols that slow, but do not stop abuse and

widespread surveillance. Key data protection concepts like requiring "legitimate interests" to process data, purpose limitations for the data collected, and data minimization cannot adequately mitigate the harms of facial recognition systems. Even legitimized use of facial recognition for fraud prevention, such as legitimate interests which are not outweighed by the rights of individuals, and processing necessary for performance of a contract or necessary to comply with a legal obligation, is capable of great societal harm. What might result in a net benefit to a particular individual in a specific, limited circumstance will not necessarily benefit society as a whole.

The same procedural critique also applies to proposed search warrant requirements for law enforcement use of facial recognition technologies. While regimes like this often require government actors to seek permission before conducting a limited search based upon probable cause, they do not protect against all the harms of surveillance. Most warrant requests are in fact granted and facial recognition searches on people in public[10] places are regularly considered to not encroach on peoples' privacy, despite evidence[11] to the contrary. Warrants are important safeguards for due process and privacy, but like consent requirements, they will end up only slowing, then further entrenching surveillance systems. This might be acceptable for all other surveillance systems, but it will not be enough for facial recognition technologies, which are uniquely dangerous. Our faces are central to our identity. They are difficult and sometimes illegal to hide. These tools can draw from an existing legacy of photos and videos that link our names and faces. The result is that we will all be watched more closely and more regularly than ever before.

Given the ineffectiveness of consent regimes, fair information practices, and warrant requirements, there seems to be only one remaining strategy to mitigate facial recognition technologies – they must be selectively or outright prohibited. This could be done several different ways, such as prohibiting face prints from being linked to names or being stored in databases, prohibiting real-time face surveillance, prohibiting government procurement of facial recognition technologies, or prohibiting the use of facial recognition and face & affect characterization in particular contexts such as in schools, in an interview or in an employment setting. Governments might issue a blanket prohibition on all types of facial recognition (or at least face surveillance) then grant exceptions for very limited and specific circumstances, such as for use to aid people with disabilities or emergency circumstances.

Whatever approach lawmakers take, it seems clear that the appetite for this technology is quite strong and facial recognition will continue to be adopted and inevitably abused until lawmakers take the technology seriously and treat it as unique. The world has never seen anything like facial recognition. Our privacy and ability to flourish as people and communities depends upon comprehensive and robust rules so that we can simply remain a face in the crowd.

10) The Public Information Fallacy（Woodrow Hartzog）
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3084102
11) The Public Information Fallacy（Woodrow Hartzog）

For more information

Privacy's Blueprint: The Battle to Control the Design of New Technologies (2018)
(Japanese translation by Tatsuhiko Yamamoto, Satoshi Narihara, Takayuki Matsuo, and Mayu Terada through Keiso Shobo Press forthcoming).

Why You Can No Longer Get Lost in the Crowd, THE NEW YORK TIMES(April 17, 2019).

What Happens When Employers Can Read Your Facial Expressions?, THE NEW YORK TIMES(Oct. 17, 2019).

Professor of Law and Computer Science
Northeastern University

Woodrow Hartzog

Woodrow Hartzog is a Professor of Law and Computer Science at Northeastern University School of Law and the Khoury College of Computer Sciences. He is also a JILIS Senior Researcher, a Resident Fellow at the Center for Law, Innovation and Creativity (CLIC) at Northeastern University, a Faculty Associate at the Berkman Klein Center for Internet & Society at Harvard University, a Non-resident Fellow at The Cordell Institute for Policy in Medicine & Law at Washington University, and an Affiliate Scholar at the Center for Internet and Society at Stanford Law School.

He is the author of *Privacy's Blueprint: The Battle to Control the Design of New Technologies*, published in 2018 by Harvard University Press (Japanese translation forthcoming with Keiso Shobo Publishing). His book with Daniel Solove, *Breached!: Why Data Security Law Fails and How to Improve It*, is under contract with Oxford University Press.