The Ethics of Facial Recognition Technology

Evan Selinger and Brenda Leong

(Forthcoming in *The Oxford Handbook of Digital Ethics ed. Carissa Véliz.)*

1. Facial Recognition Technology Primer: What Is It and How Is It Used?

Since the face is a unique part of the human body that is deeply linked to personal, social, and institutional identities, whoever controls facial recognition technology wields immense power. That power is the subject of intense debate—debate that has legal implications for privacy and civil liberties, political consequences for democracy, and a range of underlying ethical issues. While we have our own views on the subject, the primary goal of this chapter is to clarify what some of the most fundamental ethical issues are as well as specify the key conceptual distinctions that have to be grasped to fully understand them.

We begin by defining basic terms. The media, the public, and even the designers and producers of various image-based systems are prone to using the category 'facial recognition technology' inconsistently. Sometimes, they stretch the term too far and apply it to image-based technologies that analyse faces without identifying individuals.

Facial scanning systems come in four main varieties and they can each be understood as having different use cases, benefits, and risks (FPFa, 2018). The most basic applications use facial *detection*, such as what you might see through your camera—the small, square overlay that moves around to frame the face(s) of the people in your field of vision. This technology does not collect personally identifiable information (PII) and finds human faces to allow the camera to do things like focus or apply a playful filter on them and count people passing a certain spot.

The next level on the continuum is called facial *characterization*, sometimes also referred to as facial analysis and emotion detection. In this case, more detailed information is collected by analyzing a single image. Marketers might use an interactive billboard at a bus stop, or a screen

mounted above a product display, to collect information such as gender, approximate age range, and potential emotional indicators (e.g. 'smiling' or 'sad'), that can be combined with other data, such as how long the person looked at the screen, or where else they went within a store. This technology also can benefit visually disabled individuals by describing on-screen images to them (e.g. a man and a woman seated on a towel on the beach, laughing) (Newton, 2016). While facial characterization programs do not routinely create or retain personably identifiable facial templates, the process, particularly in systems purporting to detect emotions and dispositions, engender concern for two reasons: (1) scientifically questionable presumptions and technological inaccuracies underlie classifying and interpreting faces (Buolamwini and Gebru, 2018) and (2) unproven facial characterization systems have been applied towards controversial ends, such as inferring if someone is gay or straight (Economist, 2017).

The term 'facial recognition' most precisely applies to two variations of biometric systems that create an identifiable template of a unique person: *verification* and *identification*. Biometrics are any measure of a personal characteristic that is unique to an individual and can be used to distinguish one human being from another (IBIA, 2018). Like many biometric systems, facial recognition systems create templates, a point-based design that proprietary software derives from a person's facial structure; every company's system performs this function differently and images can only be matched using the same system as the initial enrollment. To enroll someone in a facial recognition database, the system scans their face (live or from an image), creates a template, and then stores that information as the baseline for future matches (IBIA, 2018).

Verification is facial recognition in a one-to-one matching system in which software answers the question, 'Is this person who they are claiming to be?' The output is a simple 'yes'

or ‘no’ to validate the claimed identity. An example of verification is accessing your phone by having a screen scan your face to match a saved template. By contrast, *identification* is a one-to-many matching process in which software answers the question, ‘Can an algorithm determine who this unknown person is?’ Law enforcement uses identification systems when running a collected image against an existing database, such as one containing mugshots or driver’s license holders. The system scans the new image—possibly from a video tape at a public venue, or an image from a camera on-scene—creates a template, and then attempts to match it to a previously enrolled individual.

Currently, many services—some of which we have mentioned already—rely on these functions. For example, convenience services include logging onto phones, sorting and organizing family photos, authenticating identities on platforms for goods and services, and profiling consumers to provide personalized recommendations. Hotels, conferences, and concerts are exploring facial identification to create VIP experiences for their members and registrants, enabling the transition from taxi, to lobby, to room, or checking into a performance or event, with minimal delays, lines, or other points of friction along their path, (Revfine undated).

Additionally, companies are using facial identification to assist the blind and low vision communities through audio or braille interfaces. Other programs are using facial characterization functions to help people on the autistic spectrum interpret emotional expressions (Gay, et al, 2013). Educators are deploying facial characterization in personalized learning; and many schools have installed security systems for campus access. Medical researchers are finding new ways to use image studies for diagnosis and treatment (Hallowell, et al, 2018). Finally, governments are using facial recognition with concentrated efforts in law enforcement (e.g. suspect identification and tracking missing persons), international security (e.g. border controls

and terrorist watch lists), and domestic security (e.g. safety protections at large events, tracking hate groups, and searching for persons of interest). Each of these uses engenders ethical, legal, and privacy concerns.

2. Standards, Measures, and Disproportionately Distributed Harms

Since facial recognition technology is used across the globe, it is beyond the scope of this chapter to comprehensively review all of the standards (Welinder and Palmer, 2018). A lack of consensus exists around whether quality testing should be mandated, who should be authorized to conduct testing, and what standards should be adopted to ensure facial recognition technology is used responsibly. Even just the matter of deciding what images in datasets should be used to train facial recognition programs has generated extensive controversy, particularly around the issue of image source (i.e. whether people consented to that use of their images) and diversity (Hill, 2020).

Setting operational standards is complicated because there are many variables to consider. For example, since photos collected from video or other sources for identification purposes may be of low quality and resolution, facial recognition technology operators should never accept outputs without strong controls for review and oversight. Overall, they need to be able to clearly demonstrate the sufficiency of their training data, set thresholds appropriately, and analyze results correctly across a range of situations.

For technical standard purposes, it is illustrative to note that the National Institute of Standards and Technology (NIST) within the U.S. Department of Commerce tests and ranks many commercially available systems, including those used by U.S. government agencies. Manufacturers *voluntarily submit* their systems for testing; if a system is not ranked, its producer has chosen not to submit it. NIST testing offers a thorough review of each system's capabilities,

including false positives and false negatives, broken down across demographic groups when applicable. The best systems—the 10 to 15 manufacturers at the top of the rankings—show excellent accuracy, above 99 per cent in almost all contexts. These companies represent the vast majority of market share, particularly by government customers at all levels (e.g. federal, state, and local/municipal).

However, it would be grave error to underestimate the fact that the NIST list also includes the remaining ~100 companies that offer some form of facial recognition service, with decreasing levels of accuracy. Particularly bad outcomes are concentrated across gender and racial categories, especially for people of colour, people who are transgender, and women, with black women experiencing the highest percentage of demographically sortable mistakes (Gorther, Ngan and Hanaoka, 2019).

These outcomes have led many critics to emphasize that everyone is not equally vulnerable to being harmed by mistaken identifications and everyone does not experience the same sense of unease when using, or being subject to a facial recognition system. This is the ethical and legal problem of *disproportionately distributed harm*. Since the context of using facial recognition technology varies, the harms can range from delays and inconvenience, to embarrassment, harassment, false accusations, and imprisonment. Unfortunately, when the demographic majority of society find themselves less at risk of error or harm than minorities, it can incline them to be less risk-averse than if they had an equal share in the peril (Guardian, 2020). This is especially problematic in the case of facial recognition technology because 'once deployed, [it] is very difficult to dismantle' due to lock-in effects and path-dependency (Whittaker 2020, p. 21). Related problems exist in situations in which uses of facial characterization generate 'socially toxic effects' by reinforcing discredited stereotypes about race

and gender (Browne, 2009 and Stark, 2019) and bolstering discredited forms of inquiry, such as phrenology, which is now re-emerging in digital form (Chinoy, 2019).

Ultimately, what should count as the sufficient standard of accuracy for any system will vary based on application and context. On an iPhone, Apple's system verifies an image that is stored locally on the device. Apple's system, called FaceID, uses an infrared camera, a depth sensor and a dot projector to map 30,000 points on a face and create an 3D scan. (The 3D technology is one of the ways to prevent access by someone simply holding up a picture of the phone's owner to gain access). The detail and level of certainty for this match yield roughly a false positive rate of 1 in 10 million. It is an entirely acceptable standard for phone access but is far below the standards that would be required for terrorist watchlists or criminal prosecution, if such use cases should ever be approved at all (Angwin, et al, 2016). Likewise, since the mismatch rate for identification is higher for some demographics than others, it would be inexcusable to identify suspects for law enforcement purposes without an expertly trained human in the loop (Lynch, 2018).

3. Erosions of Trust

In order for a society to function, there needs to be a *viable level of trust* among a range of people and institutions (Waldman, 2018). In this section, we outline some of the harms that facial recognition technologies can cause when their use erodes people's trust in social institutions. We begin by considering how the use of facial recognition is affecting the level of trustworthiness between citizens and law enforcement. In a society in which law enforcement is perceived to be untrustworthy, important ethical goals, such as providing justice, become difficult, if not impossible.

In a series of reports, the Georgetown Law Center on Privacy and Technology reached several alarming conclusions about both the accuracy and the social impacts of facial recognition technology in the context of law enforcement. While this analysis only covers U.S. cases and therefore is not globally representative, it still remains useful for explanatory purposes. Some of the main findings of the reports include the following observations: law enforcement have purchased citywide face surveillance networks that can scan the faces of city residents in real time as they walk down the street; law enforcement has not always been transparent with the public about how they are using facial recognition technology; law enforcement agents, who are not bound by federally standardized procedures for using facial recognition technology, have engaged in shoddy practices (e.g. inappropriately submitting forensic sketches and celebrity photos, and doctoring low-quality photos, including copying and pasting facial features from someone else's face onto a photo); law enforcement are searching databases containing name-face links of over half of American adults without acquiring explicit consent from citizens to do so; and, government agents have employed legally contested practices when using facial recognition systems to enforce surveillance for international departures from airports (Garvie, Bedoya, and Frankle, 2016; Rudolph, Moy, and Bedoya, 2017; Garvie and Moy, 2019; Garvie, 2019)

What these research findings suggest is that police in the U.S. are integrating facial recognition into their institution in ways that *should diminish the trust of a well-informed public* and inspire political will for more robust restrictions on mass surveillance technologies (Ferguson, forthcoming). Indeed, recent surveys about police use of facial recognition already show that trust among younger people and people of colour is substantially lower than older, white people (Smith, 2019). Unsurprisingly, theorists like Zoé Samudzi who are attuned to

problems with the police acting in biased and untrustworthy ways are arguing that the debates over making facial recognition systems more accurate for minorities like black people are failing to grapple with a more important and fundamental question: 'In a country where crime prevention already associates blackness with inherent criminality, why would we fight to make our faces more legible to a system designed to police us?' (Samudzi, 2019).

Trust is also central to many legal and ethical issues concerning the use of sensitive biometric data in contexts such as employment, benefits determinations, and marketing. Misuse and abuse in these domains can *erode consumer and stakeholder trust in the fairness, equality, and reliability of these processes*, and, depending on the context and circumstances, harms can result that impact individuals, groups, or society as a whole (Barocas and Selbst, 2016). Individual harms may include 'loss of opportunity,' 'economic loss,' 'loss of liberty,' and 'societal detriment.' Examples of loss of opportunity harms include informational injuries related to employment, insurance and social benefits, housing, and education. For example, if an employer uses a biased facial scanning system during an interview to evaluate the applicant for characteristics of friendliness or other aspects that would make her a 'good fit' for company culture, and ultimately treats this analysis as the deciding factor over her resume, performance, or other qualifications. Economic loss harms relate to credit, differential pricing, and narrowing of choice, such as have been demonstrated based on gender (Vigdor, 2019). Loss of liberty harms include the negative effects of surveillance, such as suspicion, incarceration, and others that we will discuss in more detail below (Gillard, 2019). Lastly, 'social detriment' harms arise from the development of filter bubbles and confirmation bias, and the stigmatization of groups leading to dignitary harms and stereotype reinforcement. (FPF, 2017).

Collectively, these harms can *adversely impact society as a whole by eroding the trust people have in their ability to be treated fairly, succeed on their own merits, and receive equal justice* (Berle, 2020). Given how high these stakes are, strong counter measures are being proposed. For example, the organization *AI Now* is calling for affect recognition to be banned in situations in which 'important decisions' are made 'that impact people's lives and access to opportunities,' including 'who is interviewed or hired for a job, the price of insurance, patient pain assessments, or student performance in school' (Crawford and Whittaker, 2019).

In addition to the direct harms to individuals and groups whose data is being collected or used, there are harms to those individuals who have chosen to opt out of particular platforms or services (e.g. social media), attempted to avoid a particular technology (e.g. not buying 'smart' home appliances), or requested to have their information excluded from data sets (e.g. exercising unsubscribe or do not sell options, among others). Such individuals might believe they are doing enough to protect themselves from unwanted analysis, but machine learning-based algorithms—including facial characterization systems—might still be able to infer information about them, align with similarly situated individuals, and use those inferences for marketing, pricing, employment, housing, or educational recommendations that directly impact them (Deane, 2018). If individuals believe they have no escape from the ubiquity of surveillance, they will be even more likely to *lose trust in assurances that their data and privacy choices are being respected or enforced* (Brandom, 2018).

While it is difficult to resolve the ethical issues around facial recognition technology in legal and regulatory practices, governments and policy organizations are trying to do so. Numerous papers and studies are emphasizing the importance of establishing transparent and accountable governance, and the emerging paradigm almost unanimously calls for ethical ideals

centred around ‘reasonable’ and ‘trustworthy’ practices. For example, the World Economic Forum proposed a framework that ‘seeks to address the need for a set of concrete guidelines to ensure the trustworthy and safe use of this technology’ (WEF, 2020). And the Ada Lovelace Institute commissioned an independent review of the governance of biometric data to ‘make recommendations for reform that will ensure biometric data is governed consistently with human rights, the public interest and public trust’ (Ada Lovelace, 2020). These are merely illustrative examples of regulatory bodies seeking practical tools for ensuring that an appropriate amount of trust exists between authorities and citizens. Assessing the extent to which such approaches can positively impact the social contract would take us beyond the scope of this chapter.

4. Ethical Harms Associated With Perfect Facial Surveillance

Although facial recognition technology systems can be inaccurate, it is important to keep in mind that humans tend to be worse. Human accuracy at identification is both highly evolved (e.g. we can remember faces with remarkably short initial exposure) but also highly unreliable. Even people with the best identification skills have high error rates, and when we rely on our biological capabilities to make matches out of a large group of possibilities, the process is slow, frequently biased, and subject to degrading accuracy over time. Despite these human limitations, the alternative of highly accurate, fast, and efficient automated identification system is not inherently preferable.

Crucially, facial recognition technology can pose significant risks when they operate correctly as well as the obvious concerns when they are inaccurate or make biased recommendations. If every facial recognition technology system worked perfectly every single time, on every demographic and population, concerns about surveillance, criminal forensics, and

long-term tracking and profiling would remain. In particular, Orwellian worries about living under conditions of ubiquitous surveillance might intensify.

Awareness of being watched affects individuals' behaviour regardless of whether they intend any wrongdoing. Surveillance can affect individuals' perceptions of themselves and others, and is said to have a 'chilling effect', which means that people will be inclined to self-censor their public statements and activities and even unintentionally conform their behaviour to acceptable group norms (Kaminiski and Witnov, 2015). Because the chilling effect limits self-expression, creativity, and growth, it harms democratic societies by depriving the marketplace of ideas from receiving input from all its members. It also leads to 'othering' people whose personalities, conditions, or behaviour deviate from average or mainstream expectations, including minorities, neuro-atypical individuals, and activists of all stripes who have ethical reasons for wanting to challenge aspects of the status quo (Kaminiski and Witnov, 2015).

To identify another harm that perfect surveillance can yield, philosopher Benjamin Hale proposes a futuristic thought experiment: imagine a society that gets closer to the ideal of perfect policing and uses ubiquitous facial surveillance as a tool to deter people from criminal activity (Gershgorn, 2020). According to Hale, this approach to governance risks eroding a normative ideal that, in the Kantian system of ethics, is central to self-determined decision-making: 'freedom of the will' (Hale, 2005).

To illustrate the problem, Hale asks us to consider a person who is considering breaking a law against adultery and cheating on a spouse but, solely out of fear of getting caught through facial surveillance and subsequently punished, feels compelled to be law-abiding (Hale, 2005, p.145). While sneaky people might get away with cheating on their spouses in a society without ubiquitous facial surveillance, in one in which ubiquitous facial surveillance monitors public

movements, it will be difficult, if not impossible, to have a discreet rendezvous or purchase a surreptitious gift. In such a society—one in which Hale presupposes facial recognition technology generates information that can be widely accessed—someone who commits infidelity is almost guaranteed to get caught and suffer reputational and legal consequences.

Hale's main point is that people who lived under the threat of constant facial surveillance will be disincentivized to consider such matters as whether it truly is ethical for people to freely choose to commit to a monogamous sexual relationship, and whether one should be allowed to choose to be the type of person who would intentionally deceive a spouse, absent risk of discovery. From Hale's perspective, there is a terrible cost to increasing legal compliance this way. It erodes the motivation for people to engage in ethical deliberation about how they should act and who they should be (Hale, 2005, p. 150). In such a world, ethical intentions as well as moral character and personal virtues that require the exercise of free choice, such as sincerity or integrity, would all be compromised because they would be hard to develop. 'Taking responsibility for one's actions by claiming that the action accords with who one is—"I stayed true because I love you, because I am your companion and I am honest"—can never be uttered without the attendant: "I did it because someone else was watching"' (Hale, 2005, p. 151).

5. Alienation, Dehumanization, and Loss of Control

Philosopher Philip Brey clarifies how facial recognition systems can be used to cause the harms of alienation, dehumanization, and loss of control (Brey, 2004). To narrow our focus to these problems, he pinpoints two adverse results that can follow from biometrics, including facial templates, digitally encoding highly personal aspects of our bodies. First, reducing an important part of our being to digital information can lead us to view a piece of our physical selves as having a new and essential function: the face becomes a medium for transferring the

information that automated systems need to make an identification. People can experience this shift in phenomenological perspective as dehumanizing because an intrinsic aspect of their person, such as their unique faces that have deep connections to their life experiences, is translated into things that only have instrumental value, such as passwords, PINs, and barcodes.

Brey further contends that externalizing a body part by turning it into an informational equivalent can have profound implications for how power is deployed. Because the process separates an aspect of the self from its owner, others can seize control of it. To use terms associated with loss of property and ownership, this is a process of alienation—a setting aside of what is yours that grants someone else access, use, and authority. Once this happens, Brey maintains, your face is no longer exclusively 'yours'; it exists in a form that you might not understand or even recognize. In such a scenario in which your body has new purposes and new meanings, Brey contends it is questionable whether you can retain control of 'your' data or 'yourself.'

6. The Facial Recognition Slippery Slope Debate

How facial recognition technology might be used in the future is having a profound impact on current ethical, legal, and political debates. While not always explicitly stated, the following two questions underlie these debates:

- Is it reasonable to believe that facial recognition technology has such distinctive and powerful affordances that, over time, slippery slope conditions will incline societies, even ones committed to democratic principles, to look for ever-expansive ways to invasively use it?

- If so, will the slippery slope influence behaviour so strongly that *ethically unacceptable* erosions of freedom, dignity, and democracy occur—or, at a minimum, will be more likely to occur than often is acknowledged?

Woodrow Hartzog and Evan Selinger answer both questions affirmatively and claim 'facial recognition technology is the most uniquely dangerous surveillance mechanism ever invented' (Hartzog and Selinger, 2018). Believing the technology poses unique threats to basic liberties, they recommend *enacting a ban across public and private sectors* to prevent slippery slope drivers from expanding the scale and scope of facial surveillance (Selinger and Hartzog, 2018).

By contrast, Adam Theirer claims slippery slope allegations lack credibility because the slippery slope advocates illegitimately take 'a kernel of truth…that a new technology could pose risks if used improperly…and extrapolate from it hyper-dystopian predictions ripped from the plots of sci-fi books and shows' (Thierer, 2019). From his perspective, slippery slope arguments about facial recognition technology are 'technopanics' that resemble prior exaggerated concerns over technologies such as 'instant photography, sensors, CCTV, and RFID' that well-functioning societies have successfully adapted to (Thierer, 2019). Given the putatively logical flaws of slippery slope arguments, Thierer recommends rejecting them, including the allegation that democratic countries will become more inclined to embrace authoritarianism if they adopt widespread facial recognition technology (Thierer, 2019).

7. Slippery Slope Arguments

To critically evaluate a slippery slope argument about facial recognition technology, one needs a sound framework for assessing slippery slope arguments in general. The most basic question, therefore, is: what is a slippery slope argument and how should it be assessed?

A slippery slope argument has an if-then structure and presents a causal sequence linking a seemingly unobjectionable commencing action to an *ethically objectionable final outcome* (Enoch, 2001, p. 631). Specifically, slippery slope arguments attempt to identify tragic situations that will arise primarily due to lack of foresight about medium or long-term consequences. While 'slippery slopes appear everywhere' (Devine, 2018, p. 392), the argument is especially prominent in debates about technology and technology policy, especially in controversies over biotechnology (Holm and Takala, 2007) and privacy (Scalia, 2013; Morales, Ram, and Roberts, 2020; and ACLU of Illinois, 2017).

In the case of facial recognition, the slippery slope argument entails the assertion that all parties agree a society of ubiquitous, public surveillance is inimical to the ideals of a free, democratic order, and that this outcome is so tragically harmful and exceptionally difficult to undo once achieved, that it should be actively prevented with intentional strategic measures in all policy development for technology implementation. Then, the question becomes: is there any level of use case of facial recognition that can be 'safely' implemented that will not have an unacceptably high likelihood of leading to this final, undesired outcome?

For example, is it reasonable to expect that something as seemingly low-risk as widespread photo tagging on Facebook (occurring in conjunction with similar features within other consumer activities) will normalize facial recognition technology to such a degree that too many citizens uncritically accept such practices, which then expand to more intrusive use cases, and ultimately permit the government to engage in ubiquitous facial surveillance? If the answer to this question is 'yes', it will impact current practices as well as future policy and legal determinations, including the social choice to limit individual rights in the light of broader social good.

Many people are already divided over *whether consent is a sufficient ethical requirement* for individuals who want to use particular consumer applications. Continuing with the Facebook example, Facebook's basic facial recognition technology policy is on solid legal ground because in complying with local laws, it does not make facial recognition services available everywhere; and when the services are available, the company follows standard privacy-by-design practice by only allowing tagging to be available after users give affirmative expressed consent. This policy appears to respect user autonomy: it allows them to choose or refrain from selecting the option without financial penalty; and, it provides users with a default design (i.e., the functions initially are turned off) that protects their privacy.

It might be, however, that the conditions under which Facebook legally permits users to give consent fail to meet the normative requirements for offering consent in an ethically legitimate manner, in light of a slippery slope conclusion. From this perspective, the conditions for opting-in fall short of the standard Nancy Kim calls 'consentability,' particularly regarding informed consent (Kim, 2019). In Kim's interpretation, informed consent does not require knowing every conceivable risk, but it *does* necessitate that when parties make an offer, they include a reasonable presentation of the significant risks that accepting it entails. Slippery slope proponents maintain that because Facebook does not tell users that the common practice of tagging photos can cause them to become receptive to more invasive applications of facial recognition technology, it leaves them in the dark about a fundamental danger (Selinger and Hartzog, 2019c).

There are other aspects of 'consentability' that are impacted as well. Since the more invasive applications of facial recognition technology, including by law enforcement, are known to create disproportionate harms for minorities, it also violates the 'collective autonomy'

condition of consentability (Selinger and Hartzog, 2019c). Collective autonomy is the ethical requirement that a democracy safeguard *fundamental liberties for all* even when doing so requires preventing some groups (e.g. the majority of citizens) from exercising low stakes expression of choice (e.g. tagging photos) in situations in which such choices would prevent others (e.g. minorities) from being able to exercise high-level autonomy interests, like freedom of movement and association (Selinger and Hartzog, 2019c).

Does the Facebook example, then, encapsulate the first step on a slippery slope to an Orwellian society?

8. Two Types of Slippery Slope Arguments: Fallacious and Reasonable

Older critical reasoning textbooks characterized the slippery slope argument as an informal fallacy (Hurley, 1982). Slippery slope proposals were deemed fallacious because proponents can always present a series of wildly speculative predictions leading to ultimate doom. The possibility of being causally disingenuous, however, does not mean all slippery slope arguments are, in principle, poorly formulated.

Contemporary philosophers acknowledge that there are, in fact, reasonable slippery slope arguments, and philosopher Douglas Walton created a framework for making them. Walton's account has ten characteristics that collectively cover the following temporal sequence: it begins with a debate over whether choosing a certain course of action is wise; it continues by identifying specific causal drivers that will propel behaviour to a grey area; in the grey area, people lose control over what to do next; once momentum takes events further, an uncontrollable slide commences, that results in an inevitable catastrophe (Walton, 2017).

While we applaud Walton's insights, we believe his framework only identifies what could be called a *reasonable, guaranteed version of the slippery slope argument*. By modifying

its structure, one can formulate something else: a *reasonable, achievable version of the slippery slope argument*. The main difference between the direct and achievable versions is that the achievable version does not posit catastrophe as inescapable after people lose control in a grey area. Instead, a more modest claim is made—namely, that the chances of a catastrophe occurring are more likely than advocates pushing for the first step are willing to acknowledge.

To make an achievable version of a slippery slope argument, one must do the following: 'explicitly specify plausible mechanisms that could drive slippage from one step to another " and 'rigorously explain why the mechanisms deserve due consideration' (Frischmann and Selinger, 2019, 39 and 41). A parallel between achievable and guaranteed slippery slopes thus can be drawn to the 'direct' and 'weak' versions of Langdon Winner's thesis that 'inherently political technologies' exist (Winner, 1986). The strong version is that adopting 'a given technical system unavoidably brings with it conditions for human relationships that have a distinctive political cast—for example, centralized or de-centralized, egalitarian or inegalitarian, repressive or liberating' (Winner, 1986). The weak version only 'holds that a given kind of technology is strongly compatible with, but does not strictly require, social and political relationships of a particular stripe' (Winner, 1986).

To better appreciate this distinction, consider the following hypothetical example that Eugene Volokh and David Newman offer while presenting their 'defense of the slippery slope.'

> [Imagine] a proposal to put video cameras on street lamps to catch or deter street criminals. On its own, the plan may not seem that susceptible to police abuse, as long as the tapes are viewed only when someone reports a crime and otherwise recycled every day or two. Many people may be inclined to support installing the cameras, even if they would oppose a more intrusive extension of the policy, such as linking the cameras to face-recognition software or permanently archiving the

tapes.

> But once the government implements the policy and invests money in buying, installing, and wiring thousands of cameras, the costs of implementing the next step plummet. Comprehensive surveillance becomes much cheaper and thus politically easier. The money already invested may persuade a bloc of swing voters to endorse a broader surveillance operation, even if they originally opposed the camera program on cost grounds. Faced with this prospect, then, those who support the cameras but reject the archiving must decide: Should we implement the limited camera policy now and risk that it will lead to permanent surveillance records in the future? Or should we reject the limited camera policy we want for fear of the more intrusive policy that we oppose? (Volokh and Newman, 2003).

Based on our classification, Volokh and Newman present an achievable version of a slippery slope because they are not making a case that if, at one moment in time, video cameras are installed on street lamps, at a subsequent moment in time it is inevitable that the infrastructure will include new facial recognition capabilities. Instead, they highlight a basic principle of transaction cost theory. It is easy to influence behaviour at scale by lowering transaction costs, e.g. minimizing how much money, time, effort, or other resources is required to do something. More specifically, Volokh and Newman are pointing out that when society invests resources like money, labour, and infrastructure, the buy-in can have a pronounced influence on attitudes toward future choices about what to do with it. In the case of surveillance systems, once considerable investment is made, subsequently refraining from extracting the most surveillance value can seem wasteful.

As Volokoh and Newman's hypothetical case emphasizes, the temptation to maximize surveillance potential is especially strong when proposals for bolstering assets can be framed as inexpensive efficiency enhancements. In the real world, this argument has already been made for

adding facial recognition capabilities to existing cameras for the purpose of more easily identifying criminals. New York Governor Andrew Cuomo stated that the use of cameras at toll booths to scan vehicle license plates 'is almost the least significant contribution that this electronic equipment can actually perform,' and so it makes more sense to take surveillance to 'a whole new level' (Furfaro, Bain, and Brown, 2018). Since it is not expensive to add facial recognition technology capabilities to the cameras and link the system to a variety of databases, Cuomo recommended this course of action (Furfaro, Bain, and Brown, 2018).

9. Facial Recognition Technology as Unique and Slippery

What is it about facial recognition technology that leads some people to be concerned that its widespread use will incentivize causal drivers that lead to catastrophic outcomes? The answer is that, given how much information facial recognition technology systems powered by artificial intelligence systems can quickly process, they deserve to be classified as *unique* compared to any other technology that presently can be used for surveillance purposes. If this categorization is correct, that is, if facial surveillance is like a 'phase transition' in physics (Stanley, 2019) or a highly toxic substance like plutonium (Stark, 2019) and lead paint (Read, 2020), then there are plausible reasons for believing that its features present uniquely dangerous slippery slope risks.

Is facial recognition technology truly unique? Some say 'no' and insist that it is spurious to demarcate it as a singular entity (Thierer, 2019 and Castro, 2019). In response, others base their arguments for uniqueness on four claims that we present below. *Taken together, these claims along with related points about powerful causal drivers (normalization and function creep), justify that proponents can advance a valid, achievable version of the slippery slope argument.*

First: faces play such a special, existential role in human lives that it is almost inconceivable to imagine large-scale human societies that do not place an extremely high value on unconcealed faces. Faces are the 'primary means by which humans recognize... each other' (Rifkin, et. Al, 2018, p. 310). Even if it turns out that, contrary to popular wisdom, the appearance of our individual faces is not definitively unique, the likelihood of finding duplicates remains debatable (Meester, Preneel, and Wenmackers, 2019). While facial features can change, it nevertheless is difficult for adult faces to dramatically alter in the absence of extreme circumstances. This is why, unfair as it might be, people are shocked when celebrities get plastic surgery and stop looking like classic images of themselves, and why cutting-edge, partial medical face transplants have triggered 'intense' and 'heated' reactions (Pearl, 2017).

Faces are also the 'primary means by which humans...interact with each other' and express themselves (Rifkin, et. al 2018, p. 310). Faces have this preeminent status because facial expressions and even micro-expressions are powerful forms of body language, and the source of speech. Because faces are the *foremost intermediary between our private interior lives and the ways we make ourselves publicly available*, the experience of face-to-face interaction traditionally has been associated with immediacy and intimacy, and some philosophers have even made the case that the face (understood broadly as the living presence of another person) is the basis from which ethical relationships and ethical responsibilities arise (Levinas, 1969). Our faces come with us wherever we go, and the central roles they play in identification, communication, and social interaction go a long way towards explaining why, with exceptions (e.g. burqas), most contemporary societies *expect us to keep our faces visible*. Indeed, hiding your face in public can provide a good justification for others to infer that you are behaving suspiciously.

Second, more information can readily be extrapolated from faces than other biometrics. When it comes to identification, faces are the most reliable conduits for linking our on- and offline lives—much more so than iris patterns, fingerprints, gait, *et cetera*. Additionally, faces are the basis from which a host of inferences and predictions related to facial characterization are possible—that is, inferences and predictions about everything from mood, to likelihood of telling the truth, sexual preferences, and the propensity for criminal behaviour. Rampant criticism that junk science underlies many uses of facial characterization does not appear to be deterring these ambitions (Chinoy, 2019 and Varghese, 2019). Neither are assessments that illegitimate scientific claims are further embedding 'bias and discrimination within our society' (Whittaker, 2020, p. 2). The incomparably high value associated with extracting information and parsing meaning from faces thus makes them the ideal target for extensive analysis from both the private and public sectors—from law enforcement, to educators and advertisers, and many more groups.

Third, under current conditions, there are lower transaction costs for extracting diverse forms of information from faces than other biometrics (Hartzog and Selinger, 2018). Unlike DNA and fingerprints, for example, no physical interaction is needed. Faces can be scanned from sensors remotely, and passively (Dormehl, 2020), with software that can quickly infer information or seek matches. While in the future, such easy scanning might be applied at scale to gait recognition or other related biometrics, nothing comparable exists today. This asymmetry is also aided by widespread activities that link personal information directly to faces, such as registering for driver's licenses and passports and diverse online activities (e.g. social media, job locating services, and employee rosters). In short, the massive quantity of existing photos, the extraordinary amount of information linked to faces in existing databases, and the ease of

expanding databases through activities like scraping make facial recognition systems the ideal plug-and-play technology.

Fourth, deep legal gaps exist, placing few limits on the use of this technology. While laws differ around the world, the legal lacuna regulating facial recognition technology are deep and the lack of legal guardrails enables permissive, if not promiscuous uses. In the European Union, there is growing concern about consolidation of resources among governments (FindBiometrics, 2020). In the United States, Congress does not impose any restrictions on how the government uses facial recognition technology and the courts have not mandated any meaningful limitations. Furthermore, the law presumes citizens lose some reasonable expectation of privacy from face scanning and analysis as soon they are in public or share name-face connections with third parties (Wehle, 2014), and remains unclear on the legitimacy of scraping the web for supposedly 'public' images (Vermont Attorney General, 2020). There is a thin line dividing technology companies and government agencies, and the contracts between facial recognition companies and the customers who license and use facial recognition systems are also generally 'shrouded in secrecy' (Whittaker 2020, p. 8).

These four reasons for considering facial recognition technology as unique are then combined with two additional factors to justify an achievable version of the slippery slope argument: normalization and the creeping expansion of surveillance powers and surveillance technologies (Frischmann and Selinger, 2018).

While civil rights and privacy advocates have identified the harms that facial recognition technology can cause, consumer applications of it are accompanied by positive representations in a variety of media, including advertisements, that single out the technology as fun and rewarding to use in daily life. Whether to open a phone, move quickly through an otherwise frustratingly

long line, pay for a purchase, tag photos, or have your visage turned into an amusing avatar or matched with a famous painting, faces and facial recognition technologies are continually represented as the ideal currencies for doing things joyfully and efficiently. Since these experiences can drive expectations and re-engineer desires, they raise the possibility that 'if citizens expect to be immersed in facial recognition technology wherever they go, they might become open to allowing law enforcement to behave like everyone else' (Selinger and Hartzog, 2019a).

From this perspective, it is a mistake to believe that the risks associated with different applications of facial recognition technology can be neatly cabined. When using facial recognition technology to unlock an iPhone, Apple deserves the highest grade for the privacy-by-design approach it takes to verification. Faceprints are encrypted and stored locally on the device, so there is little chance of them being compromised or abused. Nevertheless, as people become used to unlocking phones with faces, they might be inclined to accept other uses of facial recognition technology.

Finally, the various mechanisms of surveillance creep as discussed throughout have led some scholars and municipalities to contend that facial recognition is 'truly one-of-a-kind technology' that currently has more potential than any other to be used to weaken the 'obscurity' protections that are essential to the following ends and liberties: pursuing self-development, intimate relationships, freedom from chilling effects, freedom from the pressure to be conventional, and freedom to participate fully in democratic and civic life (Hartzog and Selinger, 2019; Selinger and Hartzog, 2019b).

How likely is it that the factors listed will propel us further down the slippery slope? This is not the sort of measurement where either side can justify their position with an appeal to

uncontestable, mathematically formulated probabilities; current analysis cannot provide anything like mathematically precise articulation of these risks. Again, slippery slope outcomes are defeasible—the prognostics do not designate inevitabilities. How events unfold will depend on contextual factors, such as the actions of social, political, and legal institutions, the voices who make arguments about facial recognition technology and how much authority and resources they have, how legal and policy reforms are crafted, and how legal frameworks are interpreted and applied. While its daunting to weigh all these features, one thing is certain: the longer it takes to determine appropriate regulations, the harder it will be to rollback an infrastructure that deploys society-shaping power.

Bibliography

ACLU of Illinois (2017) 'Surveillance Cameras Are A Slippery Slope' in Cunningham, A. (ed) *Privacy and Security in the Digital Age.* New: Greenhaven Publishing, pp 156-164.

Ada Lovelace Institute (2020) Independent Review of the Governance of Biometric Data. *Ada Lovelace Institute.* 24 January at https://www.adalovelaceinstitute.org/ada-lovelace-institute-announces-independent-review-of-the-governance-of-biometric-data/

Angwin, J., Larson, J., Mattu, S. and Kirchner, L. (2016) 'Machine Bias.' *ProPublica*. 23 May, at https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

Barocas, S. and Selbst, A. (2016) 'Big Data's Disparate Impact.' *California Law Review* vol. 104: 671-98 at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2477899.

Berle, I. (2020) 'Face Recognition Technology: Compulsory Visibility and Its Impact on Privacy and the Confidentiality of Personal Identifiable Images.' *Law, Governance and Technology Series* (41) 1st ed. Sutton: Springer at https://www.amazon.com/Face-Recognition-Technology-Confidentiality-Identifiable/dp/3030368866/ref=sr_1_1?dchild=1&keywords=9783030368876&linkCode=qs&qid=1584314952&s=books&sr=1-1

Bitzionis, T. (2020) 'EU Planning Shared Network of Face Biometrics Databases.' *FindBiometrics Global Identity Management*. 24 February, at https://findbiometrics.com/biometrics-news-eu-planning-shared-network-facial-recognition-databases-022401/

Brandom, R. (2018) 'Shadow Profiles Are the Biggest Flaw in Facebook's Privacy Defense.' *The Verge*. 11 April at https://www.theverge.com/2018/4/11/17225482/facebook-shadow-profiles-zuckerberg-congress-data-privacy

Brey, P. (2004) 'Ethical Aspects of Facial Recognition Systems in Public Places.' *Journal of Information, Communication, and Ethics in Society* vol. 2: 97-109.

Browne, S. (2009) 'Digital Epidermalization: Race, Identity, and Biometrics.' *Critical Sociology* vol. 36(1): 131-150.

Buolamwini, J. and Gebru, T. (2018) 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification.' *Proceedings of Machine Learning Research, Conference on Fairness, Accountability, and Transparency,* at http://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf

Castro, D. (2019) 'In Attempt to Ban Facial Recognition Technology, Massachusetts Could Inadvertently Ban Facebook, iPhones, and More.' *Information Technology & Innovation Foundation*. 21 October at https://itif.org/publications/2019/10/21/attempt-ban-facial-recognition-technology-massachusetts-could-inadvertently

Chinoy, S. (2019) 'The Racist History Behind Facial Recognition.' *The New York Times*. 11 July at https://www.nytimes.com/2019/07/10/opinion/facial-recognition-race.html

Crawford, K, and Whittaker, M, et al (2019) 'AI Now 2019 Report.' *AI Now Institute*. December at https://ainowinstitute.org/AI_Now_2019_Report.pdf

Deane, M. (2018) 'AI and the Future of Privacy.' *Medium, Towards Data Science.* September 5 at https://towardsdatascience.com/ai-and-the-future-of-privacy-3d5f6552a7c4

Devine, P. (2018) 'On Slippery Slopes' *Philosophy* vol. 93: 375-393.

Devlin, H (2020) 'AI Systems Claiming to 'read' Emotions Pose Discrimination Risks.' *The Guardian*. 16 February at https://www.theguardian.com/technology/2020/feb/16/ai-systems-claiming-to-read-emotions-pose-discrimination-risks

Dormehl, L. (2020) 'U.S. Military Facial Recognition Could Identify People from 1 km Away.' *Digital Trends*. 18 February at https://www.digitaltrends.com/cool-tech/military-facial-recognition-tech-kilometer/

Economist, The (2017) 'Advances in AI are used to spot signs of sexuality.' *The Economist* September 9th edition, at https://www.economist.com/science-and-technology/2017/09/09/advances-in-ai-are-used-to-spot-signs-of-sexuality

Enoch, David. (2001) 'Once You Start Using Slippery Slope Arguments, You're on a Very Slippery Slope.' *Oxford Journal of Legal Studies* vol. 21(4): 629–647.

Ferguson, A. (forthcoming). 'Facial Recognition and the Fourth Amendment' *Minnesota Law Review* vol. 105.

Frischmann, B. and Selinger, E. (2018) *Re-Engineering Humanity*. New York: Cambridge University Press.

Furfaro, D., Bain, J., and Brown, R. (2018) 'Inside Cuomo's Plan to Have Your Face Scanned at NYC Toll Plazas.' *New York Post*. 20 July at https://nypost.com/2018/07/20/inside-cuomos-plan-to-have-your-face-scanned-at-nyc-toll-plazas

Future of Privacy Forum (2018) 'Privacy Principles for Facial Recognition Technology in Commercial Applications.' *Future of Privacy Forum*. September, at https://fpf.org/wp-content/uploads/2019/03/Final-Privacy-Principles-Edits-1.pdf

Future of Privacy Forum (2018a) 'Understanding Facial Detection, Characterization, and Recognition Technologies.' *Future of Privacy Forum*, March, at https://fpf.org/wp-content/uploads/2018/09/FPF_FaceRecognitionPoster_R5.pdf.

Future of Privacy Forum (2017) 'Unfairness by Algorithm: Distilling the Harms of Automated Decision-Making.' *Future of Privacy Forum.* 11 December, at https://fpf.org/2017/12/11/unfairness-by-algorithm-distilling-the-harms-of-automated-decision-making/

Garvie, C. (2019) 'Garbage In, Garbage Out: Face Recognition on Flawed Data.' *Georgetown Law Center on Privacy and Technology*. 16 May, at https://www.flawedfacedata.com/

Garvie, C. and Moy, L. (2019) 'America Under Watch: Face Surveillance in the United States.' *Georgetown Law Center on Privacy and Technology*. 16 May, at https://www.americaunderwatch.com/

Garvie, C., Bedoya A., and Frankle J. (2016) 'The Perpetual Line-Up: Unregulated Police Face Recognition in America.' *Georgetown Law Center on Privacy and Technology*. October 18, at https://www.perpetuallineup.org/

Gay, V., Leijdekkers, P., and Wong, F. (2013) 'Using Sensors and Facial Expression Recognition to Personalize Emotional Learning for Autistic Children.' *PHealth 2013 Conference Proceedings*. 71-76, at https://books.google.com/books?hl=en&lr=&id=FTKVAQAAQBAJ&oi=fnd&pg=PA71&dq=personalize+learning+facial+recognition&ots=4ul6xuvxJL&sig=kgVPfx9yJT2H1J_IEVhjiwLhGjs#v=onepage&q=personalize%20learning%20facial%20recognition&f=false

Gershgorn, D. (2020) 'Exclusive: Live Facial Recognition Is Coming to U.S. Police Body Cameras.' *Medium One Zero.* 5 March, at https://onezero.medium.com/exclusive-live-facial-recognition-is-coming-to-u-s-police-body-cameras-bc9036918ae0

Gillard, C. (2019) 'Privacy's Not an Abstraction.' *Fast Company*. 25 March, at https://www.fastcompany.com/90323529/privacy-is-not-an-abstraction

Grother, P., Ngan, M., and Hanaoka, K. (2019) 'Facial Recognition Vendor Test, Part 3: Demographic Effects.' *National Institute of Standards and Technology*, *U.S. Department of Commerce*. December, at https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf

Hallowell, N., Parker, M. and Nellåker, C. (2018) 'Big data phenotyping in rare diseases: some ethical issues.' *Genetics in Medicine* vol. 21(2): 272-274. 15 June.

Hale, B. (2005) 'Identity Crisis: Face Recognition Technology and Freedom of the Will' *Ethics, Place & Environment* vol. 8(2): 141-158.

Hartzog, W. and Selinger, E., (2019) 'Just a Face in the Crowd? Not Anymore.' *The New York Times*. April 18: A 25.

Hartzog, W. and Selinger, E. (2018) 'Facial Recognition is the Perfect Tool for Oppression.' *Medium Artificial Intelligence*. 2 August, at https://medium.com/s/story/facial-recognition-is-the-perfect-tool-for-oppression-bc2a08f0fe66

Hill, K. (2020) 'The Secretive Company That Might End Privacy As We Know It.' *The New York Times*. 18 January, at https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html

Holm S. and Takala, T. (2007) 'High Hopes and Automatic Escalators: A Critique of Some New Arguments in Bioethics.' *Journal of Medical Ethics* vol. 33(1): 1-4.

Hurley, P. (1982) *A Concise Introduction to Logic*. Belmont: Wadsworth.

International Biometrics and Identity Association (2018) 'Biometrics Explained: Answers to 13 Basic Biometrics Questions.*' IBIA*, at https://www.ibia.org/download/datasets/4346/IBIA-Biometrics-Explained-final-final-web.pdf.

Kaminiski, M. and Witnov, S. (2015) 'The Conforming Effect: First Amendment Implications of Surveillance, Beyond Chilling Speech.' *University of Richmond Law Review* vol. 49: 465.

Kim, N. (2019) *Consentability: Consent and Its Limits*. New York: Cambridge University Press.

Lévinas, E. (1969) *Totality and Infinity: An essay on Exteriority*. Pittsburgh: Duquesne University Press.

Lynch, J. (2018) 'Face Off: Law Enforcement Use of Facial Recognition Technology.' *Electronic Frontier Foundation publication*. 12 February. (*since updated* 28 May 2019), at https://www.eff.org/wp/law-enforcement-use-face-recognition

Maryland v. King, 469 U.S. 435 (2013) (Scalia, J., dissenting).

Meester, R., Preneel, B., and Wenmackers, S. (2019) 'Reply to Lucas & Henneberg: Are Human Faces Unique?' *Forensic Science International* vol. 297: 217-220.

Morales, I. Ram, N., and Roberts, J. (2020) 'DNA Collection at the Border Threatens All Americans.' *The New York Times*. 23 January, at https://www.nytimes.com/2020/01/23/opinion/dna-collection-border-privacy.html

Newton, C. (2016) 'Facebook begins using artificial intelligence to describe photos to blind users.' *The Verge*. 5 April at https://www.theverge.com/2016/4/5/11364914/facebook-automatic-alt-tags-blind-visually-impared

Pearl, S. (2017) 'Changing Faces.' *Aeon*. 15 November, at https://aeon.co/essays/what-do-face-transplants-say-about-identity-and-wellbeing

Read, M. (2020) 'Why We Should Ban Facial Recognition Technology.' *New York Magazine*. 30 January at https://nymag.com/intelligencer/2020/01/why-we-should-ban-facial-recognition-technology.html

Revfine (no date). '4 Use Cases of Facial Recognition in The Hospitality Industry.' *Revfine*, *Technology Tips* at https://www.revfine.com/facial-recognition-hospitality-industry/

Rifkin, W., Kantar, R., Ali-Khan, S., Plana, N., Diaz-Sisco, R., Tsakiris, M., and Rodriguez, E. (2018) 'Facial Disfigurement and Identity: A Review of the Literature and Implications for Facial Transplants.' *AMA Journal of Ethics* vol. 20(4): 309-323.

Rudolph, H., Moy, L., and Bedoya, A. (2017) 'Not Ready for Takeoff: Face Scans at Airport Departure Gates.' *Georgetown Law Center on Privacy and Technology*. 21 December at https://www.airportfacescans.com/

Samudzi, Z. (2019) 'Bots Are Terrible at Recognizing Black Faces: Let's Keep It That Way.' *Daily Beast*. 8 February at https://www.thedailybeast.com/bots-are-terrible-at-recognizing-black-faces-lets-keep-it-that-way

Selinger, E. and Hartzog, W. (2019a) 'Our Government Should Not Be Conducting Facial Surveillance.' *Medium*, *One Zero*. 5 March at https://onezero.medium.com/our-government-should-not-be-conducting-facial-surveillance-54cc13f1ea61

Selinger, E. and Hartzog, W. (2019b) 'When Happens When Employers Can Read Your Facial Expression.' *The New York Times*. 21 October: A23, at https://www.nytimes.com/2019/10/17/opinion/facial-recognition-ban.html

Selinger, E. and Hartzog, W. (2019c) 'Why You Can't Really Consent to Facebook's Facial Recognition.' *Medium*, *One Zero*. 30 September at https://onezero.medium.com/why-you-cant-really-consent-to-facebook-s-facial-recognition-6bb94ea1dc8f

Selinger, E. and Hartzog, W. (2018) 'Amazon Needs to Stop Providing Facial Recognition Tech for the Government.' *Medium*. 21 July at https://medium.com/s/story/amazon-needs-to-stop-providing-facial-recognition-tech-for-the-government-795741a016a6

Smith, A. (2019) 'More Than Half of U.S. Adults Trust Law Enforcement to Use Facial Recognition Responsibly.' *Pew Research Center, Internet and Technology*. 5 September at https://www.pewresearch.org/internet/2019/09/05/more-than-half-of-u-s-adults-trust-law-enforcement-to-use-facial-recognition-responsibly/

Stark, L. (2019) 'Facial Recognition Technology is the Plutonium of AI.' *XRDS*, vol. 25;: 50-55, at https://xrds.acm.org/article.cfm?aid=3313129.

Stanley, J. (2019) 'The Dawn of Robot Surveillance: AI, Video Analytics, and Privacy.' *American Civil Liberties Union*, 17 June at https://www.aclu.org/sites/default/files/field_document/061819-robot_surveillance.pdf

Thierer, A. (2019) 'The Great Facial Recognition Technopanic of 2019.' *The Bridge*, 17 May at https://www.mercatus.org/bridge/commentary/great-facial-recognition-technopanic-2019

Varghese, S. (2019) 'The Junk Science of Emotion-Recognition Technology.' *The Outline*, 21 October at https://theoutline.com/post/8118/junk-emotion-recognition-technology?zd=1&zi=jy2zxjml

Vermont Attorney General (2020) 'Attorney General Donovan Sues Clearview AI for Violations of Consumer Protection Act and Data Broker Law,' *Office of the Vermont Attorney General*, 10 March at https://ago.vermont.gov/blog/2020/03/10/attorney-general-donovan-sues-clearview-ai-for-violations-of-consumer-protection-act-and-data-broker-law/

Vigdor, N. (2019) 'Apple Card Investigated After Gender Discrimination Complaints.' *The New York Times*, 10 November at https://www.nytimes.com/2019/11/10/business/Apple-credit-card-investigation.html

Volokh, E., and Newman, D. (2003) 'In Defense of the Slippery Slope.' *Legal Affairs* March/April, at https://www.legalaffairs.org/issues/March-April-2003/scene_marapr03_volokh.msp

Walton, D. (2017) 'The Slippery Slope Argument in the Ethical Debate on Genetic Engineering Of Humans.' *Science and Engineering Ethics* vol. 23(6): 1507-1528 at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3107283.

Waldman, A. (2018) *Privacy as Trust: Information Privacy for an Information Age*. New York City: Cambridge University Press.

Welinder, Y., and Palmer, A. (2018) 'Face Recognition, Real-Time Identification, and Beyond.' In *The Cambridge Handbook of Consumer Privacy*, edited by. Evan Selinger, Jules Polonetsky, and Omer Tene, 102-124. Rochester, NY. Cambridge University Press.

Wehle (née Brown), K. (2014) 'Anonymity, Faceprints, and the Constitution.' *George Mason Law Review* vol. 21: 409-466 at http://www.georgemasonlawreview.org/wp-content/uploads/2014/03/Brown-Website.pdf.

Winner, L. (1986) *The Whale and the Reactor: A Search for Limits in an Age of High Technology*. Chicago: University of Chicago Press.

Wittaker, M. (2020) 'Facial Recognition Technology (Part III): Ensuring Commercial Transparency & Accuracy.' Written Testimony to the United States Representatives Committee on Oversight and Reform. 15 January at https://ainowinstitute.org/oversight-committee-testimony-whittaker.pdf

World Economic Forum (2020) 'A Framework for Responsible Limits on Facial Recognition.' *WEF White Paper*. February at http://www3.weforum.org/docs/WEF_Framework_for_action_Facial_recognition_2020.pdf