Facial Recognition for Counter-Terrorism: Neither a Ban Nor a Free-for-All

Scott Robbins

1 Introduction

This chapter starts from the fact that new technology has given new power to the state to automate the identification of previously known terrorists who are organizing attacks on the citizens that the state is supposed to protect. The power to do this (and associated powers), if it works effectively, would help in countering terrorism. Facial recognition technologies (FRTs) promise to give the state precisely that power.

Using FRTs, it is claimed, the state could *verify* that people are who they say they are, *identify* people appearing in images or video feeds, *characterize* their behavior and emotions, and check that they are not a suspected terrorist. For example, FRTs are deployed to verify that a person going through border control is indeed the person pictured on an identification document (e.g., a passport).[1] Interpol has deployed its Project FIRST system to help state authorities identify foreign terrorist fighters (FTFs).[2] In a truly horrifying example, the company Faception claims to be able to detect terrorists and pedophiles based on the characteristics of their face [16].[3] This power, however, has been challenged. This challenge, for some, should result in a complete ban on the use of this technology.

S. Robbins (✉)
Center for Advanced Security, Strategic and Innovation Studies (CASSIS), University of Bonn, Bonn, Germany
e-mail: srobbins@uni-bonn.de

[1] The company Veridas, for example, advertises that their FRT system can be used at border controls. See https://veridas.com/government-institutions/.

[2] https://www.interpol.int/en/Crimes/Terrorism/Identifying-terrorist-suspects.

[3] It would take another paper to discuss the many failings of even proposing this. While this is probably being used in earnest by both the state and private corporations— it should not be. The science behind applications like these are pseudo-science and have many of the same characteristics as the now disreputable practice of phrenology [2].

© The Author(s) 2021

A. Henschke et al. (eds.), *Counter-Terrorism, Ethics and Technology*,
Advanced Sciences and Technologies for Security Applications,
https://doi.org/10.1007/978-3-030-90221-6_6

The reasons put forward for such a ban are that FRTs suffer from pervasive bias resulting in the benefits and harms being unequally distributed amongst groups, the state will inevitably use these technologies for illegitimate purposes, and that the existence of FRTs chill our behavior (i.e., causes people to censor themselves for fear of surveillance).

At the moment, the state faces little restriction over how they use FRTs. There are plenty of examples of the state's use of FRTs for purposes that give people pause. For example, police departments in the U.S. have used FRTs to identify and monitor activists and protestors of color [1]. Setting up a surveillance network powered by FRTs will, it is argued, significantly increase the risk of the state abusing its power. This risk will be associated with an increase in citizens chilling their behavior.

It is paramount that if the state can use this technology to increase their power to counter terrorism, this power is constrained such that the abuses and chilled behavior do not occur. I argue below for five conditions on the use of FRTs. First, the state must create institutional constraints that only allow FRTs to be used in places where people do not (and should not) enjoy a reasonable expectation of privacy (e.g., airports, border crossings). Second, the cameras equipped with FRT must be marked to assure the public that they are not being surveilled in places that they should have a reasonable expectation of privacy. Third, FRTs should be restricted to finding serious criminals (e.g., terrorists). Fourth, the state should not use third-party companies that violate the first three conditions during the creation or use of its service. And fifth, third-party companies should not be able to access or read the sensitive data collected by the state. With these conditions satisfied, given the effectiveness of FRT, the state can harness FRT's power to counter-terrorism.

2 The Basics of Facial Recognition

The goal of FRT is to verify, identify, characterize, or check someone against a watch list based on an image of a particular person. Most people will have some experience with this because Facebook, Google, and Apple all use F.R. in commercial applications. Apple's FaceID lets users into their phones using the camera on their phone to verify their identity. Google and Facebook have long used F.R. to identify and auto-tag photos with the people in them.

The four goals of FRTs (verify, identify, characterize, and watch list)[4] should be distinguished as they carry different ethical concerns. Verification is merely matching two images to check if they are the same person. This is a one-to-one comparison. An organization may want to verify that the person wearing a security badge is the same as the person's picture on the badge. This can be difficult for humans to do—but relatively easy for FRTs.

[4] There is also a fifth goal that is simply *detection* which would merely involve detecting that there is a face in an image. This is necessary in order for the success of the other goals but needn't concern us here.

Using FRTs to identify is a one-to-many relationship. An image of a face has to be checked against a 'faces' database to determine who, exactly, they are.[5] This goal of FRTs has had its spotlight in the media recently because the FBI used FRT to identify those who participated in the January 6, 2021 insurrection at the U.S. capitol [6]. High-quality images taken from that day are fed into an FRT that can compare the faces on those images to a database of faces that have identities attached to them.

FRTs for the characterization of a particular person aim to label people as having a particular emotional state or as, for example, terrorists based solely on an image of their face. In China, for example, FRTs have been used in classrooms to detect the level of engagement of the students [12]. The aforementioned company Faception claims to be able to detect everything from 'professional poker player' to a person with a 'high IQ' to a 'terrorist'. This has been dismissed as modern day phrenology by some academics—as there does not seem to be any evidence that facial features have a relationship to personality, profession, or criminal behavior (see e.g., [28]).

3 Arguments for an FRT Ban

3.1 Disparate Impact

FRTs suffer from pervasive bias. This means that FRTs perform exceptionally well for some groups, while it performs terribly for other groups. With this in mind, we can use 'bias' how we use it in everyday language: FRTs are biased against dark-skinned people. FRTs, in one study, performed 5–10% worse for African Americans compared with Caucasians. Buolamwini and Gebru found error rates as high as 34% for African females compared with a low 0.8% error rate for Caucasian males. This could be due to their being a lack of images of dark-skinned people used to train the algorithm. Or it could be that dark-skinned faces are harder for current algorithms to translate into computer language and extract useful patterns out of. Whatever the reason, current algorithms have a huge problem recognizing dark-skinned faces.

Problems like these mean that the benefits and harms brought by FRTs are unequally distributed amongst groups of people. Those for whom the technology does not work as well with, will not be able to be verified by FRTs—causing suspicion and further intrusive surveillance. Furthermore, they will be misidentified more frequently. This may cause them to be suspected as a terrorist or other serious criminal. For example, on January 9, 2020, Robert Williams was arrested in front of his wife and two daughters. The reason for this arrest was that an FRT misidentified him as a person who stole watches from a store in a robbery that took place 18 months earlier [15]. If Robert Williams was just an unfortunate misidentification due to the FRT not being 100% accurate, we could accept this—misidentification also happens

[5] I am speaking loosely here. The 'faces' in the database are actually computer generated representations of faces. Depending upon the specific methodology used these can be more or less robust. For an overview of some of the specific technical methodologies see [11].

when done by human beings. However, FRTs consistently misidentify (or fail to identify) people of color significantly more often than whites. Then, it follows that people of color will disproportionately experience the harms caused by these technologies. In a liberal democracy, the principle of everyone being equal under the law is violated by technologies with this problem.

Meanwhile, the benefits of FRTs will be disproportionately received by middle-aged white males. Not only will they be identified more reliably—meaning that they will get through security lines without further intrusive surveillance, but they will disproportionately feel the benefits of convenience that these technologies promised in the first place. Joy Buolamwini (mentioned above) started to analyze FRTs precisely because she couldn't get FRTs to recognize her face. At one point, she put on a white mask triggering the program to recognize hers as a face [13]. The point is that the convenience promised by FRTs is distributed unfairly. This compounds the problem above because the same group that disproportionately experiences the harms of FRTs also disproportionately fail to experience its benefits. This problem must be overcome if FRTs are to be used anywhere. The main point is that many FRTs don't work. If a particular technology doesn't work, then we shouldn't use it. However, this does not mean that the technology will not work in the future. In this paper, I assume that we will only be using FRTs that work with an appropriate level of effectiveness for everyone.

3.2 Chills Behavior

> *Surveillance conducted with facial recognition systems is intrinsically oppressive. The mere existence of facial recognition systems, which are often invisible, harms civil liberties, because people will act differently if they suspect they're being surveilled* [7].

Many institutions and scholars echo this sentiment about FRT. Evan Selinger and Brenda Leong, channeling philosopher Benjamin Hale, argue that pervasive effective FRTs would undermine our free will—and would prevent ethical behavior caused by that will—replacing it with "I acted ethically because someone was watching" [22]. The freedom to choose to do the right thing whether or not someone is watching is central to the liberal democratic ideal of autonomy. In everyday life we encounter many scenarios that require ethical reasoning and action. Coffee, for example, might be for sale based on the honor system. Customers are supposed to leave a Euro after they take a coffee. People should have the right to be honorable. When someone (or something) is watching, then we don't get the chance to be honorable. Our actions are evaluated in light of someone watching—which, when you leave the Euro for the coffee isn't as honorable as if you were to leave the Euro without someone watching. FRTs, therefore, should be banned (or so concludes their argument).

Furthermore, the freedom to gather in large groups to protest injustice should not be hampered by the knowledge that you will be identified by FRTs and be labeled as a subversive. The freedom of assembly is enshrined in liberal democratic

constitutions and declarations of human rights. The U.N. Declaration of Human Rights, in article 20, states that "Everyone has the right to freedom of peaceful assembly and association" [25], and the United States Constitution gives citizens the "right of the people peaceably to assemble, and to petition the Government for a redress of grievances ("U.S. Senate: Constitution of the United States", n.d. 26)." The right to assemble and air grievances can be the last resort to create necessary change. A 2020 study of the U.S. civil rights movement, for example, showed that it was activism and protests which "drove media coverage, framing, congressional speech, and public opinion on civil rights" [27].

A person who was horrified by the murder of George Floyd and wants to voice their support for systemic change in the policing system in the U.S. should be able to do so. However, they may fear that FRTs will identify them as taking part in a protest (and may further document what exactly the protest was)—which may cause them to lose their job or harm their chances for jobs in the future. If that protest were to turn violent then it may be that all attendees get labeled as violent protestors—regardless of their intentions and actions at the protest. A 2013 report showed that:

> *surveillance of Muslims'* quotidian activities has created a pervasive climate of fear and suspicion, encroaching upon every aspect of individual and community life. Surveillance has chilled constitutionally protected rights—curtailing religious practice, censoring speech and stunting political organizing [23].

Ordinary citizens' rights to practice religion, speak their minds, and politically organize have been shown to be hindered by surveillance. FRTs increase this risk dramatically—as their chances of being identified with FRTs is far greater.

The recent events of January 6, 2021, in which pro-Trump groups converged on the capital and staged a violent insurrection, may cause one pause here. Don't we want these people to have their behavior 'chilled'? Many liberals cheered the use of FRTs to identify people to have them arrested.

There are two essential things to note here. First, the right to assembly, speech, and political organizing does not include the right to violently overthrow the government. Chanting and holding up signs in front of the Capitol building should not be chilled—whether or not we agree with the assemblers. Carrying weapons, engaging with police, and threatening Congress members are not included in the right to peaceful assembly. Second, while this protest did turn violent and illegal, that does not mean that each individual who attended this protest deserves to be stigmatized without participating in the actual insurrection. While those that stormed the capitol should fear consequences brought on by the state, those that simply protested the election results should not.

If people who merely intended to voice their grievances did not attend this protest simply because they feared being identified and face the consequences, their right to peaceful assembly was violated. This causes harm even if that person is wrong about what might happen to them. They are unsure—and therefore change their behavior. This is why there must be both institutional barriers to technology being used this

way *and* transparency in law enforcement and government to assure the public that this is so [20].[6]

One might think that with CCTV we already face this problem. CCTV captures images of people all the time—and sometimes that footage is distributed in order to identify someone that has committed a crime. Think of an armed robbery at a gas station. The suspect might be captured on CCTV footage in front of the gas station—and could be captured because someone recognized them on that footage. If, two days later, someone else walks into that same gas station they are also captured on CCTV footage. However, because normal CCTVs are not equipped with FRTs, they will never be identified as there was no crime committed (the footage has no reason to be 'looked' at—by a computer or a human). FRTs have the capability to continuously identify and store the information related to people that come across its view. This affords the state the ability to easily identify anyone who attended a particular protest—whether or not they committed a crime. This amounts to intrusive surveillance without cause. Of course, the state can claim that they do not store the information unless crimes are committed; however without clear and transparent institutional (and possibly technological) barriers to such use, it will be difficult for people to act as if they are not being surveilled using FRTs.

3.3 Scope Creep

> *Facial recognition enables surveillance that is oppressive in its own right; it's also the key to perpetuating other harms, civil rights violations, and dubious practices. These include rampant, nontransparent, targeted drone strikes; overreaching social credit systems that exercise power through blacklisting; and relentless enforcement* of even the most trivial of laws, like jaywalking and failing to properly sort your garbage cans [21].

The arguments for a ban rest on the premise that FRTs will creep pervasively into society and be used for all kinds of things they weren't initially used for. In the above quote, FRTs are envisioned for Chinese-style social credit systems and enforcement of things like jaywalking. The idea is that once this technology is out there, it will be normalized. We will come to expect it—and then it will be used everywhere.

To highlight this, I offer the following example. Let's say that FRTs are extremely effective. The government has intelligence that five New York City individuals are planning on carrying out a terrorist attack. It is decided to upgrade the CCTV network to include FRT. If any one of those individuals is captured by the smart CCTV cameras, then the authorities will be notified. It is agreed that this will be their best chance to stop the terrorist attack. Unfortunately, that upgrade cost a lot of money. In an attempt to raise money, the mayor decides that the FRT can simply start automatically ticketing J walkers. J walking is illegal, and many people do it—so using FRT to ticket them will raise a lot of money.

[6] More on this in Sect. 4.1.

Jeroen van den Hoven calls this 'information injustice.' He argues that people may not object to their data being used for a particular purpose; however, when that same data is used for another purpose, an injustice has occurred. If your library search data is collected to provide better services by the library, this may be something you agree to. However, if that same data is used to collect information on your tastes and pass them to others for advertising purposes, then informational injustice has occurred [10]. In this case, the use of FRTs to catch terrorists is now repurposed for catching J walkers. While an argument may justify the use of FRTs to catch terrorists, it cannot be used to catch J walkers without a new justification.

The problem is that once the surveillance apparatus includes pervasive FRTs, the barriers to using it for things not originally intended are very low. This is not the case for regular CCTV cameras. The cost of employing people to pour through that video and attempt to identify individual J walkers wouldn't be worth the money raised by ticketing them. CCTV's technological limitations naturally restrict law enforcement's ability to use them for anything—protecting people's reasonable expectation of privacy.

3.4 An Outright Ban

For some, the concerns above, taken together, creates a case for an outright ban of the technology. Like San Francisco, some cities have enacted such a ban [3].

Selinger and Werner believe that FRTs are "so inherently toxic that it deserves to be completely rejected, banned, and stigmatized" [21]. In another post, they conclude that "The future of human flourishing depends upon facial recognition technology being banned before the systems become too entrenched in our lives" [7].

In what follows, I argue that an outright ban may not be justified. First, there are contexts in which our expectation of privacy is simply non-existent. Second, the chilling effects are not necessarily going to happen, nor are they necessarily bad things. Finally, the scope creep that critics are concerned about is not inevitable. If the technology works as advertised, then there are some restricted contexts where these harms do not materialize.

4 Conditions for the Use of Facial Recognition

Given the argument for bans on FRT and the privacy and free speech rights enshrined in liberal democratic constitutions and human rights declarations, it is clear that the state must justify the use of FRTs before they can be used to capture terrorists. This is not a technology that simply improves upon a power that the state already had; instead, it is an entirely novel power. That is the power to identify anyone that comes into view of an FRT equipped camera without a human being watching the video feed.

Here I will outline the conditions that FRT should be subject to operate in a liberal democracy justifiably. I expand on each in the sections below. The context in which FRT is being used must be one in which the public does not have a reasonable expectation of privacy. Second, the only goal should be to prevent serious crimes like terrorism from taking place. Finally, FRTs to store and capture biometric facial data in a database, the individual in question must be suspected of committing a serious crime.

4.1 Reasonable Expectation of Privacy

In a famous case in the United States, the supreme court ruled that Charles Katz had a reasonable expectation of privacy when he closed the phone booth door [4, Chap. 1]. This meant that the evidence collected by the state who was listening in on his conversations in that phone booth had to be thrown out. This notion of a 'reasonable expectation of privacy' is fundamental to how the value of privacy is interpreted in liberal democracies. It is not just a legal notion but a notion which grounds how we act. In our bedrooms, we have a reasonable expectation of privacy, so we can change clothes without fear of someone watching. When Charles Katz closed the door to the phone booth he was using, he enjoys a reasonable expectation of privacy—he believes that no one should listen to his conversation.

Facial data captured by FRTs should be at least as protected as voice data. CCTVs in the public sphere should not be collecting information on individuals—something that happens when CCTVs are equipped with FRT. When I walk down my street, I have a reasonable expectation that my comings and goings are not being recorded—whether it be a police officer following me around or by a smart CCTV camera recognizing my face. Regular CCTVs do not record individuals' comings and goings; rather, they record what happens at a particular location.

The difference is that a CCTV camera does not record a line in a database that includes my identity and the location that I was 'seen' at. CCTV equipped with FRT *can* record such a line in a database—significantly empowering the state to perform searches that tell them much about my comings and goings. Not only should these searches be linked to clear justifications; but there should be clear justifications for collecting such intimate data (their comings and goings) on individuals.

This reasonable expectation can be overridden if I have committed a serious crime or plan on committing a serious crime. This is because my right to privacy would be overridden by the "rights of other individuals…to be protected by the law enforcement agencies from rights violations, including murder, rape, and terrorist attack" [17, 110]. If one were to be in the process of planning a terrorist attack, it would not be a surprise to them that they were being surveilled. Terrorists take active measures to prevent surveillance that they expect to occur. This may seem to justify the placing of smart CCTVs in public spaces to identify terrorists.

CCTV cameras are currently placed in many public spaces. If something happens, the authorities can review the CCTV footage to see who was responsible. In this case,

the place itself is being surveilled. Data on individuals is not 'captured' in any sense. There is no way to search a database of CCTV footage for a particular name. One must look at the footage. However, if this CCTV camera were to be "smart" and capture biometric facial data along with video footage, then each individual who is captured by this camera is being surveilled. The authorities now know each person that comes into this camera's view and what time they were there. This, even though an overwhelming majority of people coming into any CCTV camera's view has not, and does not plan to, commit a serious crime. Their privacy has been invaded.

This has ethical implications regarding scope creep and chilling behavior discussed in Sect. 3. If FRT enabled CCTV cameras are in operation, then it is easy for the state to add new uses for the technology. A simple database search could reveal everyone who goes into an area with many gay bars. A gay man in a country where homosexuality is considered unacceptable but not illegal may chill their behavior—that is, not go to gay bars to fear those visits being documented. While the FRT enabled CCTV cameras were initially installed to counter terrorism, the ability to easily search for anyone that has come across it makes it easy to use it for other, illegitimate purposes.

The state could simply state that they will only use FRTs with a warrant targeted against an individual suspect of a serious crime. For example, the authorities may have good information regarding the planning of a terrorist attack by a particular person. It is imperative that they find this person before they are able to execute the attack. They obtain a warrant and then use the city's network of FRT-enabled CCTV cameras to 'look' for this person. If this person's face is captured by one of these cameras, then the authorities are immediately notified.

If we bracket issues of efficacy and disparate impact, it appears that this would be a useful power to the state—and subject to restrictions that protect privacy. The issue is not whether or not to use FRTs, but *how* they can and should be used. However, these would be merely institutional and perhaps legal barriers that are subject to interpretation. The scope of national security is little understood. Donald Trump used the concept to justify the use of collecting cell-phone location data to track suspected illegal immigrants [14]. The power enabled by FRTs is so great, and the justifications to use them will be so little understood, that it will be near impossible for regular citizens to feel and act as if they have privacy—even if they do, in principle, have it. Your partner may promise to never read your journal unless you are either dead or in a coma; however, the fact that she has a key and knows where it is will probably cause you do self sensor what you write down—just in case. With a journal, and with your general comings and goings, you should enjoy a reasonable expectation of privacy.

However, there are some public spaces where individuals do not enjoy a reasonable expectation of privacy. Airports and border crossings are two such examples. For better or worse, we now expect little privacy in these contexts. Authorities are permitted to question us, search our bags, search our bodies, submit us to millimeter scans, etc. It would be rather odd to think that our privacy was invaded more by our faces being scanned and checked against a criminal database. On regular public sidewalks, I would be horrified to find out that the state recorded my comings and

goings; however, I would be shocked to find out the state did not record each time I crossed into and out of the country. This points to the idea that there may be places where we *should* have a reasonable expectation of privacy—whether we do or not.

A recent U.S. supreme court case illustrates this nicely. Timothy Carpenter was arrested for armed robbery of Radio Shacks and T-Mobile stores. The police used a court order (which is subject to less standards than a warrant) to obtain GPS data gathered by his cell phone and collected by the telecommunications companies MetroPCS and Sprint. In an opinion written by chief justice John Roberts, the supreme court ruled that Timothy Carpenter should have a reasonable expectation of privacy concerning his constant whereabouts. The government cannot simply, out of curiosity, obtain this data [24]. This prevents the widespread use of smart CCTV cameras in plain sight to undermine our 'reasonable expectation of privacy.' The state should not use conspicuous surveillance as a way to claim that no one has a reasonable expectation of privacy where these cameras exist. The critical point is that there are public spaces where citizens of a liberal democracy *should* have a reasonable expectation of privacy.

Therefore, *if* there are places where citizens *should not* have a reasonable expectation of privacy *and* FRTs are effective (they do not cause unequally distribute false positives and false negatives across different groups), it may be justifiable to use FRTs in those places. People expect the state to protect them from terrorism. If FRTs contribute to keeping citizens safe from terrorists, then there is a good reason to use them. However, based on the analysis above, they cannot simply be used anywhere as there are places where citizens *should* have a reasonable expectation of privacy.

The above points to the allowable use of regular CCTV cameras in public spaces but prevents FRTs from operating in those same public spaces.[7] The problem now is: How will the public know the difference? This is a serious problem. After all, the right to free expression may be 'chilled' because people believe that the state is surveilling their actions. I may worry that because my friend lives above a sex shop, the state's surveillance may cause them to believe I frequent the sex shop rather than visit my friend. I may, therefore, not visit my friend very often. Or I may not join a Black Lives Matter protest because I believe the state is using FRTs to record that I was there. This is the "chilling effect" mentioned in Sect. 3.2. This can occur even if the state is *not* engaging in such surveillance. The only thing that matters is that I believe it to be occurring.

The 'chilling effect' puts the burden on the state to assure the public that such unjustified surveillance is not happening. Where it is justified, there are appropriate safeguards and oversight to prevent misuse, etc. This requires institutional constraints, laws, and effective messaging. As [20] argue, institutional constraints and laws alone will not *assure* the public that the state is not practicing unjustified

[7] It is not for this paper to evaluate the ubiquitous use of regular CCTV cameras in public spaces. I only claim that regular CCTV does not violate our reasonable expectation of privacy if it is the place that is being surveilled and not individual people (e.g. when our identities, time, and location are stored in a searchable database).

intrusive surveillance. And vice versa, effective messaging alone will not *ensure* that the state is not practicing unjustified intrusive surveillance.

For example, if the state creates laws that prevent the use of FRT on regular city streets but the cameras that are used look the same as the smart CCTV cameras that have FRT in airports, then the public will not be assured that facial recognition is not taking place. This sets up the conditions for the chilling effect to occur. However, if the state uses cameras that are clearly marked for facial recognition in places like airports, and cameras that are clearly marked 'no facial recognition' on city streets but no laws are preventing them from using FRT on city streets, then the public has a greater chance of being assured. However, nothing is preventing the state from using the footage of those cameras and running facial recognition on them after the video has been captured. Therefore, it takes both institutional constraints (bound by law) *and* effective messaging to meet the standards which support liberal democratic values like free expression.

This creates two conditions for the state's use of FRT. First, the state must create institutional constraints that only allow FRTs to be used in places where people do not (and should not) enjoy a reasonable expectation of privacy (e.g., airports, border crossings). Second, the cameras equipped with FRT must be marked to assure the public that they are not being surveilled in places that they should have a reasonable expectation of privacy.

4.2 Cause for the State's Use of FRTs

The state should not simply use new technology because it exists. There must be a purpose for using technology that is greater than the harms and privacy infringements that occur due to that technology. It would be odd to use wiretaps to surveil a serial jaywalker. Wiretaps are used in highly restrictive situations involving serious criminals. FRTs should be no different. The point is, that "justifications matter." Collecting facial data by using FRTs for countering terrorism does not mean that the data is now fair game for any other use. Each use must have its moral justification—and if that justification no longer obtains, then that data should be destroyed [8, 257].

Terrorism is a serious enough risk (in terms of possible harm—not necessarily in terms of likelihood) that it features as a justification employed by those advocating the use of FRTs. In these cases, one does not feel as if the privacy rights of terrorists are so strong that they should not be surveilled. We expect the government to do what they can to find people like this. Their privacy rights are overridden by others' rights not to be injured or killed in a terrorist attack.

The problem is that FRTs must also surveil everyone that comes into view of one of its cameras. That is, each face is used as an input to an algorithm that attempts to match that face to an identity and/or simply check whether that face matches one of the identities of suspected terrorists. In a technical sense, this technology could only be used for the legitimate purpose of finding terrorists. However, as argued above—the difficulty in assuring the public that this is the case will have a chilling

effect. Furthermore, the real possibility of scope creep makes placing these cameras, in places where people should have a reasonable expectation of privacy, dangerous.

This means that no matter the cause, FRTs should not be employed in places where innocent people have a reasonable expectation of privacy (as argued above). However, once we restrict its use to those places where there is no reasonable expectation of privacy, then finding serious criminals using FRTs poses no ethical problem (providing that it reaches a threshold of effectiveness). The third condition for the use of FRTs is that FRTs should be restricted to finding serious criminals (e.g., terrorists).

4.3 Reliance on Third-Party Technology

The state's reliance on third-party technology companies to facilitate surveillance is perhaps the area where the most violations of liberal democratic values occur. For example, the government cannot simply scrape the entire internet of pictures of people, match the faces to names, create a detailed record of things you have done, places you have gone, people you have spent time with, etc. Especially without a just cause. This amounts to intrusive surveillance of every individual. In liberal democracies, there must be a justification (resulting in a warrant approved by a judge) to engage in such surveillance of an individual. Surveilling a million people should not be considered more acceptable than the surveillance of one person. However, Clearview A.I. has been scraping images from the web and creating digital identities for years. Many police departments and government agencies are now using this third-party company to aid in using FRTs [9].

This causes significant ethical concern for three reasons: first, some third-party companies do not follow the constraints already mentioned above; second, sensitive data is being stored and processed by third-party companies that have institutional aims that could incentivize the misuse or abuse of this data; and third, the role that these companies play in surveillance may reduce the public's trust in them.

4.3.1 Contracting out the Bad Stuff

When I first encountered FRT at an airport, I was a bit squeamish. It took me some time to understand why. Indeed, I am not against using such technology to prevent terrorists from entering the country or detecting people who are wanted in connection with a serious crime or find children on missing person lists.[8] I also did not feel that I had a reasonable expectation of privacy. I expect to be questioned by a border guard and have my passport checked. I expect that my bag or my body could be searched.

[8] Although I was concerned that my face could be checked against those captured at, for example, Black Lives Matters protests around the world—and that I could face scrutiny due to my participation. This concerns the just causes for FRTs discussed earlier.

And I expect to be captured on camera continuously throughout the airport. So why did I have this immediate adverse reaction towards the use of FRT by the state?

The answer lies in my knowledge regarding the contracting out of such work to third-party technology companies. I am expected to trust the state and the third party technology company that is behind the technology. Are they capturing my biometric face data and storing it on their third-party servers? Are there institutional barriers preventing them from reusing or selling that data for their benefit? Is the data captured, sent, stored, and processed in line with best security practices? In short, I fear that even if the proper laws and constraints regarding the state's use of FRTs are in place, that third-party technology company is not bound by them or does not respect them.[9]

This is wrong. There are laws in place that prevent the United States, for example, contracting out intrusive surveillance on their citizens to other countries. So the U.S.—not being able to collect data on its citizens—cannot ask the U.K. to collect data on a U.S. citizen. The same should be true for FRTs. Suppose the U.S. cannot gather facial data on the entire U.S. population (practicing bulk surveillance). In that case, the U.S. should also not contract such work out to a third-party company—or use a third party company that has engaged in this practice. If I contract the work of killing an enemy to somebody else, that does not absolve me of all responsibility regarding the murder of that enemy.

It is not, in principle, unacceptable to use tools created by third-party companies. Third-party companies often have the resources and incentives to create far better tools than the government could create. Silicon valley technology companies attract many creative and motivated thinkers—and pay them a salary that the government could not afford. It would be detrimental to say that the government cannot use tools created by these companies. However, big data and artificial intelligence have made this relationship much more complicated.

Rather than merely purchasing equipment, the government is now purchasing services and data. A.I. algorithms created by third-party companies are driven by the collection of vast amounts of data. If this algorithm is to be used by the state, the state must ensure that the data driving it was collected according to laws governing the state's data collecting capabilities. Furthermore, the hosting of the data that the government collects is increasingly being contracted out to cloud services like Amazon Web Services. This is so because this data processing is extremely resource-intensive and something that third-party companies are more efficient at. This creates a situation where our biometric facial data may have to be sent to a third-party company for storage and/or processing. The company in question must have no ability to see/use this data. This is so for two reasons. First, these companies have institutional aims[10] that have nothing to do with the security of the state. This creates incentives for companies to use this data for their aims—creating an informational injustice [10]. Furthermore, this blurring of institutional aims (e.g.,

[9] It should be noted that strong data protection laws like Europe's GDPR can prevent some of this from taking place.

[10] See Miller [18] for an excellent discussion on the blurring of institutional purposes.

maximizing profits *and* countering terrorism) could be detrimental to the company. As a result of NSA programs like PRISM, which purportedly allows the state to gain access to the company servers of Google and Facebook [5], rival companies are now advertising that they are outside of U.S. jurisdiction and can therefore be used without fear of surveillance.[11]

Second, this data is now being entrusted to companies that may not have the same security standards or oversight expected for the storage and processing of sensitive surveillance data. Recently the Customs and Border Patrol contracted out facial recognition to a third-party company which was breached in a cyber-attack causing the photos of nearly 100,000 people to be stolen. Customs and Border Patrol claimed no responsibility—saying it was the third-party company's fault. The state should be responsible for the security of surveillance data [19, 35].

This discussion should cause constraints on how the state uses third-party companies to facilitate surveillance. Condition number four for the state's use of FRTs is that the state should not use third-party companies that violate the first three conditions during the creation or use of its service. This means that the state should know about the services they are using. Furthermore, a fifth condition is that the third-party company should not be able to access or read the sensitive data collected by the state. This keeps the state in control of this sensitive surveillance data.

5 Conclusion

What has been written above agrees with much of what proponents of a ban argue. The large difference is that I do not believe that FRTs will necessarily creep into society in a pervasive way. The five conditions I argue for above prevents this type of creep. Furthermore, the chilling effect so feared by proponents of a ban will not necessarily occur. This only happens when there is pervasive use of FRTs in places where people should have a reasonable expectation of privacy. By restricting FRTs use to those places where people should not have a reasonable expectation of privacy, this concern can be alleviated.

However, the concern that FRTs suffer from pervasive bias is serious. There may not be FRTs that are effective at all. This should prevent their use by the state. Until it can be shown that these technologies work in a way that won't disproportionately distribute the harms and benefits amongst groups, FRTs should not be used. What is called for, then, is a moratorium rather than a ban. Once it has been shown that FRTs are effective, the state should use them within the limits outlined above.

[11] ProtonMail, for example, claims that "ProtonMail is incorporated in Switzerland and all our servers are located in Switzerland. This means all user data is protected by strict Swiss privacy laws." https://protonmail.com/.

References

1. Cagle M (2016) Facebook, Instagram, and Twitter provided data access for a surveillance product marketed to target activists of colorlACLU of Northern CA. October 11, 2016. https://www.aclunc.org/blog/facebook-instagram-and-twitter-provided-data-access-surveillance-product-marketed-target
2. Chinoy S (2019) Opinionlthe racist history behind facial recognition. The New York Times, July 10, 2019, sec. Opinion. https://www.nytimes.com/2019/07/10/opinion/facial-recognition-race.html
3. Conger K, Fausset F, Kovaleski SF (2019) San Francisco bans facial recognition technology. The New York Times, May 14, 2019, sec. U.S. https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html
4. Farivar C (2018) Habeas data: privacy vs. the rise of surveillance tech. Brooklyn
5. Greenwald G (2015) No place to hide: Edward Snowden, the NSA, and the U.S. surveillance state, reprint. Picador, New York
6. Harris M (2021) How facial recognition technology is helping identify the U.S. capitol attackers - IEEE Spectrum. IEEE Spectrum: Technology, Engineering, and Science News. January 11, 2021. https://spectrum.ieee.org/tech-talk/artificial-intelligence/machine-learning/facial-recognition-and-the-us-capitol-insurrection
7. Hartzog W, Selinger E (2018) Facial recognition is the perfect tool for oppression. Medium. August 2, 2018. https://medium.com/s/story/facial-recognition-is-the-perfect-tool-for-oppression-bc2a08f0fe66
8. Henschke A (2017) Ethics in an age of surveillance: personal information and virtual identities. Cambridge University Press, New York
9. Hill K (2020) The secretive company that might end privacy as we know it. The New York Times, January 18, 2020, sec. Technology. https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html
10. van den Hoven J (1999) Privacy of information injustice? In: Mendina GT (ed) Ethics and electronic information in the twenty-first century. Purdue University Press, West Lafayette, Indiana, USA, pp 139–150
11. Introna LD, Nissenbaum H (2009) Facial recognition technology: a survey of policy and implementation issues. The center for catastrophe preparedness and response. https://nissenbaum.tech.cornell.edu/papers/facial_recognition_report.pdf
12. Kuo L (2019) China brings in mandatory facial recognition for mobile phone users. The Guardian, December 2, 2019, sec. World news. http://www.theguardian.com/world/2019/dec/02/china-brings-in-mandatory-facial-recognition-for-mobile-phone-users
13. Lohr S (2018) Facial recognition is accurate, if you're a white guy. The New York Times, February 9, 2018, sec. Technology. https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html.
14. Lutz E (2020) Trump's immigration crackdown has taken a dystopian turn. Vanity Fair. February 7, 2020. https://www.vanityfair.com/news/2020/02/trump-immigration-crackdown-has-taken-a-dystopian-turn-cell-phone-data
15. Mayor P (2020) ACLU letter to detroit public safety headquarters, June 24, 2020. https://cdn.arstechnica.net/wp-content/uploads/2020/06/dpd_complaint_v_final.pdf
16. McFarland M (2016) Terrorist or pedophile? This start-up says it can out secrets by analyzing faces. Washington Post, May 24, 2016. https://www.washingtonpost.com/news/innovations/wp/2016/05/24/terrorist-or-pedophile-this-start-up-says-it-can-out-secrets-by-analyzing-faces/
17. Miller S (2008) Terrorism and counter-terrorism: ethics and liberal democracy. Wiley. https://www.wiley.com/en-us/Terrorism+and+Counter+Terrorism%3A+Ethics+and+Liberal+Democracy-p-9781405139434
18. Miller S (2019) Whither the University? Universities of technology and the problem of institutional purpose. Sci Eng Ethics 25(6):1679–1698. https://doi.org/10.1007/s11948-019-00147-7

19. Robbins S (2021) Machine learning & counter-terrorism: ethics, efficacy, and meaningful human control. Doctoral thesis, Delft. Technical University of Delft, The Netherlands. https://repository.tudelft.nl/islandora/object/uuid:ad561ffb-3b28-47b3-b645-448771eddaff
20. Robbins S, Henschke A (2017) The value of transparency: bulk data and authoritarianism. Surveill Soc 15 (3/4): 582–589. https://doi.org/10.24908/ss.v15i3/4.6606
21. Selinger E, Hartzog W (2018) Amazon needs to stop providing facial recognition Tech for the Government. Medium. June 21, 2018. https://medium.com/s/story/amazon-needs-to-stop-providing-facial-recognition-tech-for-the-government-795741a016a6
22. Selinger E, Leong B (2021) The ethics of facial recognition technology. SSRN scholarly paper ID 3762185. Rochester, Social Science Research Network, NY. https://papers.ssrn.com/abstract=3762185
23. Shamas D, Arastu N (2013) Mapping muslims: NYPD spying and its impacts on American Muslims. Long Island, New York, CUNY School of Law, USA. https://www.law.cuny.edu/wp-content/uploads/page-assets/academics/clinics/immigration/clear/Mapping-Muslims.pdf
24. Sorkin AD (2018) In carpenter, the supreme court rules, narrowly, for privacy. The New Yorker. June 22, 2018. https://www.newyorker.com/news/daily-comment/in-carpenter-the-supreme-court-rules-narrowly-for-privacy
25. Universal Declaration of Human Rights (2015) October 6, 2015. https://www.un.org/en/universal-declaration-human-rights/
26. U.S. Senate: Constitution of the United States. n.d. Accessed February 4, 2021. https://www.senate.gov/civics/constitution_item/constitution.htm
27. Wasow O (2020) Agenda seeding: How 1960s Black protests moved Elites, public opinion and voting. Am Polit Sci Rev 114(3):638–659. https://doi.org/10.1017/S000305542000009X
28. Whittaker M, Crawford K, Dobbe R, Fried G, Kaziunas E, Mathur V, West SM, Richardson R, Shultz J, Schwartz O (2018) AI Now 2018. AI Now Institute. December 2018. https://ainowinstitute.org/AI_Now_2018_Report.html

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.