HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

44th PARLIAMENT, 1st SESSION

Standing Committee on Access to Information, Privacy and Ethics

EVIDENCE

NUMBER 015

Monday, April 4, 2022

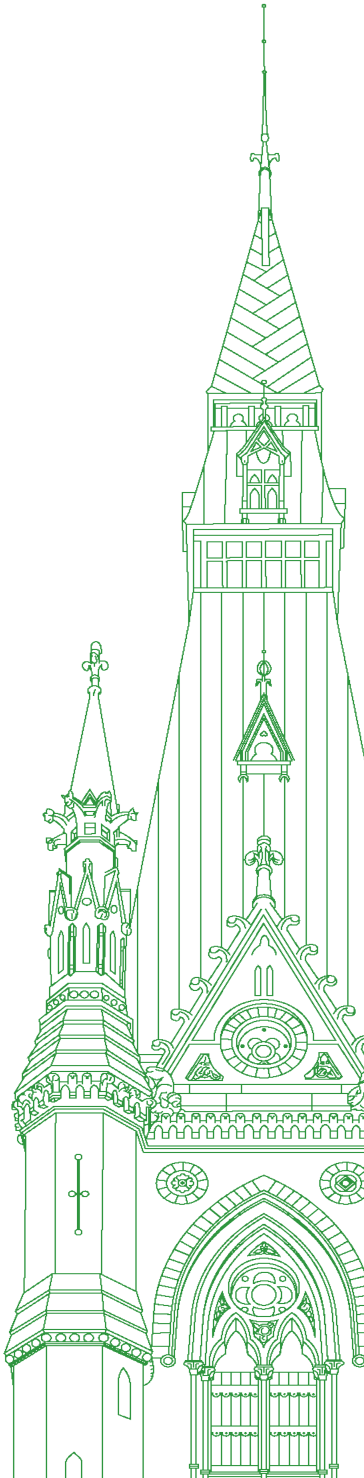

Chair: Mr. Pat Kelly

Standing Committee on Access to Information, Privacy and Ethics

Monday, April 4, 2022

● (1100)

[*English*]

The Chair (Mr. Pat Kelly (Calgary Rocky Ridge, CPC)): I call this meeting to order. Welcome to meeting number 15 of the House of Commons Standing Committee on Access to Information, Privacy and Ethics.

Pursuant to Standing Order 108(3)(h) and the motion adopted by the committee on Monday, December 13, 2021, the committee is resuming its study of the use and impact of facial recognition technology.

Today's meeting is taking place in a hybrid format, pursuant to the House order of November 25, 2021. Members are attending in person in the room and remotely using the Zoom application. Per the directive of the Board of Internal Economy on March 10, 2022, all those attending the meeting in person must wear a mask, except for members who are at their place during proceedings.

For those participating by video conference, click on the microphone icon to activate your mike. Please mute your mike when you are not speaking.

For witnesses participating for the fist time, in this type of meeting you have the option for interpretation. At the bottom of your screen, you can select floor, which is in either language, or French or English for translation. For those in the room, you can use the earpiece and select the desired channel.

I would remind everyone that all comments should be addressed through the chair.

Members in the room should raise their hand to speak. For members on Zoom, please use the "raise hand" function. The clerk and I will manage the speaking order as best we can. We appreciate your patience and understanding.

I welcome all of our witnesses. We have four witnesses this morning: Dr. Rob Jenkins, professor, University of York; Mr. Sanjay Khanna, strategic adviser and foresight expert; Ms. Angelina Wang, computer science graduate researcher, Princeton University; and Dr. Elizabeth Anne Watkins, post-doctoral research associate, Princeton University.

We will begin with Dr. Jenkins.

You have five minutes for your opening statements.

Professor Rob Jenkins (Professor, University of York, As an Individual): Good morning.

Thank you, Mr. Chair and members of the committee.

My name is Rob Jenkins. I'm a professor of psychology at the University of York in the U.K., and I speak to the issue of face recognition from the perspective of cognitive science.

I'd like to begin by talking about expectations of face recognition accuracy and how actual performance measures up to these expectations.

Our expectations are mainly informed by our experience of face recognition in everyday life, and that experience can be highly misleading when it comes to security and forensic settings.

Most of the time we spend looking at faces, we're looking at familiar faces, and by that I mean the faces of people we know and have seen many times before, including friends, family and colleagues. Humans are extremely good at identifying familiar faces. We recognize them effortlessly and accurately, even under poor viewing conditions and in poor quality images. The everyday success of face recognition in our social lives can lead us to overgeneralize and to assume that humans are good at recognizing faces generally. We are not.

Applied face recognition, including witness testimony, security and surveillance, and forensic face matching, almost always involves unfamiliar faces, and by that I mean the faces of people we do not know and have never seen before.

Humans are surprisingly bad at identifying unfamiliar faces. This is a difficult task that generates many errors, even under excellent viewing conditions and with high quality images. That is the finding not only for randomly sampled members of the public but also for trained professionals with many years of experience in the role, including passport officials and police staff.

It is essential that we evaluate face recognition technology, or FRT, in the context of unfamiliar face recognition by humans. This is partly because the current face recognition infrastructure relies on unfamiliar face recognition by humans, making human performance a relative comparison, and partly because, in practice, FRT is embedded in face recognition workflows that include human operators.

Unfamiliar face recognition by humans, a process that is known to be error prone, remains integral to automatic face recognition systems. To give one example, in many security and forensic applications of FRT, an automated database search delivers a candidate list of potential matches, but the final face identity decisions are made by human operators who select faces from the candidate list and compare them to the search target.

The U.K. “Surveillance Camera Code of Practice” states that the use of FRT “...should always involve human intervention before decisions are taken that affect an individual adversely”. A similar principle of human oversight has been publicly adopted by the Australian federal government: “decisions that serve to identify a person will never be made by technology alone”.

Human oversight provides important safeguards and a mechanism for accountability; however, it also imposes an upper limit on the accuracy that face recognition systems could achieve in principle. Face recognition technologies are not 100% accurate, but even if they were, human oversight bakes human error into the system. Human error is prevalent in these tasks, but there are ways to mitigate it. Deliberate efforts, either by targeted recruitment or by evidence-based training, must be made to ensure that the humans involved in face recognition decisions are highly skilled.

Use of FRT in legal systems should be accompanied by transparent disclosure of the strengths, limitations and operation of this technology.

If FRT is to be adopted in forensic practice, new types of expert practitioners and researchers are needed to design, evaluate, oversee and explain the resultant systems. Because these systems will incorporate human and AI decision-making, a range of expertise is required.

Thank you.

● (1105)

The Chair: Thank you very much, Dr. Jenkins.

Now we have Mr. Khanna.

You have five minutes.

Mr. Sanjay Khanna (Strategic Advisor and Foresight Expert, As an Individual): Mr. Chair, thank you very much for the opportunity to speak to you and members. I will be speaking about facial recognition technology in terms of the individual, digital society and government.

I am a consultant in the areas of strategic foresight, scenario planning and global change, and I am an adjunct professor in the Master of Public Policy in Digital Society program at McMaster University.

A key foresight method that I use for planning for the future is scenario planning. As Canada navigates the most uncertainty it has faced since the start of the post-war period, scenario planning can play a role in helping legislators to inform resilient strategy and public policy. I see the following as important issues to address with facial recognition, which I will refer to as FRT.

One, people are being targeted by FRT without meaningful consent and/or in ways they do not understand. Two, societies that are increasingly unequal include populations of people who cannot advocate for their interests related to FRT's current or possible use. Three, legislators will always be behind the curve if they do not take the time to explore the plausible futures of digital society and the role of novel technologies such as FRT within them.

I will speak to these concerns from the perspectives of the individual, of society and of government.

In terms of the individual, our faces open doors for us and can lead to doors being closed on us. We experience biases across the spectrum from negative to positive and implicit to explicit based on how our faces are perceived and on other factors related to our appearance. This fundamental reality shapes our lives.

With an FRT-enabled world, what might it mean to be recognized by technical systems in which FRT is embedded?

What might it mean for FRT to be combined with sentiment analysis to quickly identify feelings at vulnerable moments when a person might be swayed or impacted by commercial, social or political manipulation?

What might it mean for a person to be identified as a potential social, political or public safety threat by FRT embedded into security robots?

What might it mean for a person to be targeted as a transactional opportunity or liability by FRT embedded into gambling or commercial services?

Technologies associated with FRTs, such as big data, machine learning and artificial intelligence, amplify these potential risks and opportunities of FRT and other biometric technologies. While some individuals may welcome FRT, many are concerned about being targeted and monitored. In cases in which rights are infringed, individuals may never know how or why; companies may choose not to reveal the answers, and there may not be meaningful consent.

In such cases, there will be no accessible remedies for individuals impacted by commercial, legal or human rights breaches.

In terms of digital society, Canadian society faces unprecedented challenges. Rising social and racial inequalities in our country have been worsened greatly by the pandemic. Canadians are experiencing chronic stress and declining physical and mental health. Social resilience is undermined by disinformation and misinformation. Canada is addressing new and threatening challenges to the post-war order. The climate crisis is a co-occurring threat multiplier.

Despite these challenges, major technology companies are profiting from opportunities amidst the unprecedented risk and so have gained additional leverage in relation to government and our digital society. In the process, a few companies have accrued considerable power with trillion-dollar-plus valuations, large economic influence and a lock on machine learning and artificial intelligence expertise.

As I speak, technology leaders are imagining the next FRT use cases, including how FRT might be used more widely in business, government and industry. Some tech companies are exploring threats and opportunities that would justify use cases that may be unlawful today but could be viable in new circumstances, from a change in government to a shocking security event to changes in labour laws.

In terms of government, a society facing constant disruption has not proved to be a universally safe one for Canadians. The realities of harms and potential harms to individuals and of the risks and opportunities for business and government puts effective governance in the spotlight. At a time of unprecedented risk, parliamentarians have a responsibility to make sense of societal change and to comprehend plausible futures for FRT amidst the use of sophisticated surveillance systems in “smarter” cities, growing wealth and income inequality, threatened rights of children and marginalized communities.

Creating effective law and policy related to FRT should involve due contemplation of plausible futures.

● (1110)

I respect that for you, as legislators, this is a challenging task, given the often short-term horizons of elected individuals and parties. However, prospective thinking can complement the development of legislation to deal with novel and often unanticipated consequences of technologies as potent as FRT, which is inextricably linked with advances in computer vision, big data, human computer interaction, machine learning, artificial intelligence and robotics.

The Chair: Mr. Khanna, I'm sorry. I'm going to have to interrupt you.

Mr. Sanjay Khanna: I have one more paragraph

The Chair: You're a little bit over time. Thank you for your opening statement.

Mr. Fergus.

Hon. Greg Fergus (Hull—Aylmer, Lib.): While you were speaking, I heard Mr. Khanna say it was his last paragraph. I am wondering if we could make the exception to hear that.

The Chair: If you can spit it out in about 15 seconds or less, then I'll do that.

Mr. Sanjay Khanna: A government responding to the “now” in this space will always remain behind the curve. Some technology companies and start-ups are betting that governments won't catch up. Legislators should take steps to correct this impression by instituting guardrails over longer horizons that strengthen Canadians' resilience.

The Chair: Thank you very much.

I do apologize to witnesses when periodically I have to cut them off, but we are somewhat governed by the clock.

My apologies. It will probably not be the last time I have to do that in this meeting.

We will move along now to Ms. Wang.

Please go ahead with your opening statement. You have up to five minutes.

Ms. Angelina Wang (Computer Science Graduate Researcher, Princeton University, As an Individual): Hi, I'm Angelina Wang, a graduate researcher in the computer science department at Princeton University. Thank you for inviting me to speak today.

I will give a brief overview of the technology behind facial recognition, as well as highlight some of what are, in my view, the most pertinent technical problems with this technology that should prevent it from being deployed.

These days, different kinds of facial recognition tasks are generally accomplished by a model that has been trained using machine learning. What this means is that rather than any sort of hand-coded rules, such as that two people are more likely to be the same if they have the same coloured eyes, the model is simply given a very large dataset of faces with annotations, and instructed to learn from it. These annotations include things like labels for which images are the same person, and the location of the face in each image. These are typically collected through crowdsourcing on platforms like Amazon Mechanical Turk, which has been known to have homogeneous worker populations and unfavourable working conditions. The order of magnitude of these datasets is very large, with the minimum being around 10,000 images, and the maximum going up to millions. These datasets of faces are frequently collected just by scraping images off the Internet, from places like Flickr. The individuals whose faces are included in this dataset generally do not know their images were used for such a purpose, and may consider this to be a privacy violation. The model uses these massive datasets to automatically learn how to perform facial recognition tasks.

It’s worth noting here that there is also lots of pseudoscience on other kinds of facial recognition tasks, such as gender prediction, emotion prediction, and even sexual orientation prediction and criminality prediction. There has been warranted backlash and criticism of this work, because it's all about predicting attributes that are not visually discernible.

In terms of what some might consider to be more legitimate use cases of facial recognition, these models have been shown over and over to have racial and gender biases. The most prominent work that brought this to light was by Joy Buolamwini and Timnit Gebru called “Gender Shades”. While it investigated gender prediction from faces, a task that should generally not be performed, it highlighted a vitally important flaw in these systems. What it did was showcase that hiding behind the high accuracies of the model were very different performance metrics across different demographic groups. In fact, the largest gap was a 34.4% accuracy difference between darker skin-toned female people and lighter skin-toned male people. Many different deployed facial recognition models have been shown to perform worse on people of darker skin tones, such as multiple misidentifications of Black men in America, which have led to false arrests.

There are solutions to these kinds of bias problems, such as collecting more diverse and inclusive datasets, and performing disaggregated analyses to look at the accuracy rates across different demographic groups rather than looking at one overall accuracy metric. However, the collection of these diverse datasets is itself exploitative of marginalized groups by violating their privacy to collect their biometric data.

While these kinds of biases are theoretically surmountable with current technology, there are two big problems that the current science does not yet know how to address. These are the two problems of brittleness and interpretability. By brittleness, I mean that there are known ways that these facial recognition models can break down and allow bad actors to circumvent and trick the model. Adversarial attacks are one such method, where someone can manipulate the face presented to a model in a particular way such that the model is no longer able to identify them, or even misidentify them as someone completely different. One body of work has shown how simply putting a pair of glasses that have been painted a specific way on a face can trick the model into thinking one person is someone entirely different.

The next problem is one of interpretability. As I previously mentioned, these models learn their own sets of patterns and rules from the large dataset they are given. Discovering the precise set of rules the model is using to make these decisions is extremely difficult, and even the engineer or researcher who built the model frequently cannot understand why it might perform certain classifications. This means that if someone is misclassified by a facial recognition model, there is no good way to contest this decision and inquire about why such a decision was made in order to get clarity. Models frequently rely on something called “spurious correlations,” which is when a model uses an unrelated correlation in the data to perform a classification. For example, medical diagnosis models may be relying on an image artifact of a particular X-ray machine to perform classification, rather than the actual contents in the image. I believe it is dangerous to deploy models for which we have such a low understanding of their inner workings in such high-stakes settings as facial recognition.

Some final considerations I think are worth noting include that facial recognition technologies are an incredibly cheap surveillance device to deploy, and that makes it very dangerous because of how quickly it can proliferate. Our faces are such a central part of our identities, and generally do not change over time, so this kind of surveillance is very concerning. I have only presented a few technical objections to facial recognition technology today, and taken as a whole with the many other criticisms, I believe the enormous risks of facial recognition technology far outweigh any benefits that can be gained.

Thank you.

● (1115)

The Chair: Thank you.

Dr. Watkins, you have up to five minutes.

Dr. Elizabeth Anne Watkins (Postdoctoral Research Associate, Princeton University, As an Individual): Thank you for the chance to speak today.

My name is Elizabeth Anne Watkins and I am a post-doctoral research fellow at the Center for Information Technology as well as the human-computer interaction group at Princeton University, and an affiliate with the Data & Society research institute in New York.

I'm here today in a personal capacity to express my concerns with the private industry use of facial verification on workers. These concerns have been informed by my research as a social scientist studying the consequences of AI in labour contexts.

My key concerns today are twofold: one, to raise awareness of a technology related to facial recognition yet distinct in function, which is facial verification; and two, to urge this committee to consider how these technologies are integrated into sociotechnical contexts, that is, the real-world humans and scenarios forced to comply with these tools and to consider how these integrations hold significant consequences for the privacy, security and safety of people.

First I'll give a definition and description of facial verification. Whereas facial recognition is a 1:n system, which means it both finds and identifies individuals from camera feeds typically viewing large numbers of faces, usually without the knowledge of those individuals, facial verification, on the other hand, while built on similar recognition technology, is distinct in how it's used. Facial verification is a 1:1 matching system, much more intimate and up close where a person's face, directly in front of the camera, is matched to the face already associated with the device or digital account they're logging in to. If the system can see your face and predict that it's a match to the face already associated with the device or account, then you're permitted to log in. If this match cannot be verified, then you'll remain locked out. If you use Face ID on an iPhone, for example, you've already used facial verification.

Next I'll focus on the sociotechnical context to talk about where this technology is being integrated, how and by whom. My focus is on work. Facial verification is increasingly being used in work contexts, in particular gig work or precarious labour. Amazon delivery drivers, Uber drivers and at-home health care workers are already being required in many states in the U.S., in addition to countries around the world, to comply with facial verification in order to prove their identities and be allowed to work. This means the person has to make sure their face can be seen and matched to the photo associated with the account. Workers are typically required to do this not just once, but over and over again.

The biases, failures and intrinsic injustices of facial recognition have already been expressed to this committee. I'm here to urge this committee to also consider the harms resulting from facial verification's use in work.

In my research, I've gathered data from workers describing a variety of harms. They're worried about how long their faces are being stored, where they're being stored and with whom they're being shared. In some cases, workers are forced to take photos of themselves over and over again for the system to recognize them as a match. In other cases, they're erroneously forbidden from logging into their account because the system can't match them. They have to spend time visiting customer service centres and then wait, sometimes hours, sometimes days, for human oversight to fix these errors. In other cases still, workers have described being forced to step out of their cars in dark parking lots and crouch in front of their headlights to get enough light for the system to see them. When facial verification breaks, workers are the ones who have to create and maintain the conditions for it to produce judgment.

While the use of facial recognition by state-based agencies like police departments has been the subject of growing oversight, the use of facial verification in private industry and on workers has gone on under-regulated. I implore this committee to allocate attention to these concerns and pursue methods to protect workers from the biases, failures and critical safety threats of these tools, whether it's through biometric regulation, AI regulation, labour law or some combination thereof.

I second a recent witness, Cynthia Khoo, in her statement that recognition technology cannot bear the legal and moral responsibility that humans are already abdicating to it over vulnerable people's lives. A moratorium is the only morally appropriate regulatory response.

Until that end can be reached, accountability and transparency measures must be brought to bear not only on these tools, but also on company claims that they help protect against fraud and malicious actors. Regulatory intervention could require that companies release data supporting these claims for public scrutiny and require companies to perform algorithmic impact assessments, including consultation with marginalized groups, to gain insight into how workers are being affected. Additional measures could require companies to provide workers with access to multiple forms of identity verification to ensure that people whose bodies or environments cannot be recognized by facial verification systems can still access their means of livelihood.

At heart, these technologies provoke large questions around who gets to be safe, what safety ought to look like, and who carries the burden and liability of achieving that end.

Thank you.

● (1120)

The Chair: Thank you very much for that opening statement.

Now we'll move to questions.

We'll begin with Mr. Williams for six minutes.

Mr. Ryan Williams (Bay of Quinte, CPC): Thank you very much, Mr. Chair, and thank you very much to our witnesses who are attending today. This is very interesting.

I'm going to start with Mr. Jenkins.

You've completed work regarding the accuracy of facial recognition by experts such as passport officers, and you've found a large amount of human error that exists. What are the error rates by humans versus machine learning software for facial recognition technology?

Prof. Rob Jenkins: It depends largely on the specifics of the task. In a task in which passport staff who have been trained and have many years' experience in the job are asked to compare live faces presented in front of them against photographed identity documents similar to passports, we typically see error rates of around 10%. That means for every 10 comparisons that are made, one of them is made erroneously. I'm talking about a decision on whether there's a match or a mismatch between the photo and the live person.

As for computer-based systems, we have very little understanding in how most of them operate in realistic conditions. Many of the test results that are reported by vendors are based on idealized situations in which image quality is reliably good and the conditions under which the match is being conducted are very good. That ignores the noise and complexity of the real world. So we just don't know enough about that, in my view.

● (1125)

Mr. Ryan Williams: Okay. Thank you.

When we compare it with other methods of identification, such as fingerprinting, do you have any data on that? What would the error rate be for fingerprints instead of using facial recognition?

Prof. Rob Jenkins: I can't quote a figure, but there are reasons that fingerprint matching can be more reliable in certain circumstances. One of the reasons is that facial appearance changes a lot according to lighting conditions and the distance from the face to the camera lens. Those particular problems are not present when it comes to matching fingerprints.

Mr. Ryan Williams: Are the errors made by humans versus computers the same errors, or are they completely different? You've mentioned a few of them.

Prof. Rob Jenkins: There are patterns of similarity, but there are also striking divergences between the errors that computers and humans make. Dr. Wang mentioned an example of where simply adding glasses to someone wouldn't affect a human perceiver's view of who is there, but seemingly superficial changes like that can really throw some computer systems and lead to incorrect answers that are unexpected.

Mr. Ryan Williams: You mentioned that FRT should always include human intervention. Will it ever be that with human and machine intervention we have 100% accuracy? What does that decrease that accuracy to?

Prof. Rob Jenkins: One of the benefits of having human oversight as a part of the system is that egregious errors of the type we were just discussing can be fished out and noticed for the errors that they are before being acted upon. For that reason, I think it's important have a human safeguard, but the fact that human face recognition is not infallible also means that we should expect humans to introduce errors into the system if they're given the final decision. That's the result of the cognitive biases we all carry with us. I'm talking about good-faith errors rather than prejudice or malicious intent.

Mr. Ryan Williams: In “Two Factors in Face Recognition”, you wrote, “Face recognition accuracy depends much more on whether you know the person's face than whether you share the same race.” How does this trend carry through into AI-based facial recognition software?

Prof. Rob Jenkins: Well, I think it's important to distinguish between differences in ability and prejudice. Both exist, but they're independent of each other.

Differences in ability to recognize faces reflect the viewer's social diet of faces—that is, the range of facial appearances they encounter. That's important for at least two reasons. First, we should expect demographic disparities in face recognition by humans even in the absence of prejudice. Second, the notion of a social diet of faces has a clear analogue in face recognition technology, specifically the composition of face databases that are used to train the algorithm.

Tackling prejudice is clearly important in its own right, but it would not eliminate demographic disparities in face recognition accuracy. That's a separate problem.

Mr. Ryan Williams: Okay.

Thank you, Mr. Chair. I'll let my 14 seconds go.

The Chair: All right. Thank you.

With that, I'll go to Mr. Fergus for six minutes.

Hon. Greg Fergus: Thank you very much, Mr. Chair.

I'd like to thank all the witnesses for being present here today. I appreciate it.

I have questions for several witnesses, so I'd appreciate it if the witnesses could be brief, yet pithy, in their comments.

Mr. Jenkins, in a question that you had from my colleague, Mr. Williams, you were asked about setting up fingerprinting versus facial verification. From what I heard from Dr. Wang, you and other witnesses, they're not quite the same thing.

Can you compare the two in terms of their accuracy and how facial recognition technology is used, as opposed to fingerprinting? I'm assuming it is really just a process of trying to match up a dataset to another dataset. Is that correct?

● (1130)

Prof. Rob Jenkins: There are some general similarities.

In both cases, the idea is to take a sample from the world—be it somebody's fingerprint line or their facial image alike—and compare it with some stored representation that you have and that you're expecting will provide a match.

The difficulty arises when the variability in the live capture from the person you're trying to identify...it can vary over time. You always have to account for that variability in attempting the match to the gallery of stored information.

Now—

Hon. Greg Fergus: In other words, the situation changes remarkably for the presentation of one's face, as opposed to the presentation of one's fingerprints. It might not be quite an apples-to-apples and oranges-to-oranges comparison.

Prof. Rob Jenkins: I think that's fair to say. We know for sure that different pictures of one person's face can be more varied than pictures of different people's faces. That's the nub of the problem.

Hon. Greg Fergus: Thank you very much for that.

Dr. Wang, thank you very much for your presentation. If I may suggest, I know that you only brought to our committee a couple of the problems that your research has identified. If there are others that you would like to share with this committee.... We have a common saying here that if we don't hear it or if we don't read it, we can't report on it. We would certainly appreciate it if you felt you had the time and could send us more examples of what you consider some of the limitations of facial verification.

I'd like to go back to the two big problems that you identified, which are brittleness and interpretability.

I was wondering if you could talk a bit more about the brittleness of it. Bad actors could circumvent the system, but there's also the vulnerability of people who have no intention of circumventing it, but are yet victims of the biases. I think you talked about machine learning and that all it does is extenuate the biases that would exist in society in general.

Am I correct?

Ms. Angelina Wang: Yes, you are.

For brittleness, because we don't really know what the model is picking up on in order to make certain identifications, we don't know what patterns it's relying on. Because humans might know that people are likely to wear makeup and put on glasses, they can control for these kinds of changes. If someone were to inadvertently do something a bit different with their face and how they're presenting themselves, this might not be tested for and the model might misidentify them.

Hon. Greg Fergus: I know this goes beyond what you testified today, but in some of the readings we've had, we've talked about the limitations of the technology, such as camera technology. There are clear biases in the faces that the technology will favour. It was created throughout, and has evolved since we started taking pictures. It favours white males, in particular. For every other category or group, there are varying levels of greater and greater inaccuracy.

Could you talk a bit more about that? Even if we were to try to correct for machine learning, we would still have a problem with the technology itself, and the biases that might be introduced by that technology.

Ms. Angelina Wang: Ever since cameras were invented, they have always worked a lot worse on people with darker skin tones. They haven't accounted for different lighting differences. The cameras have always been developed primarily on people with lighter skin tones. A lot of times in different lighting conditions, it just will not work as well on people with different skin tones. People's faces may blend into the background more, depending on what they look like.

● (1135)

Hon. Greg Fergus: Therefore, as a result, it perpetuates that bias that's already built into the system.

Ms. Angelina Wang: Exactly. The image quality will be different for different people.

Hon. Greg Fergus: Mr. Khanna and Dr. Watkins, I'm coming up close to the end of my time, but I'm going to see if I can get in a really quick question.

Mr. Khanna, you mentioned that politicians have to get ahead of the game.

Can you give us, very briefly, how we should get ahead of the game to try to put the right type of framework around FRT?

Mr. Sanjay Khanna: Yes. I think you should use a technique called scenario planning. I think for the purposes that you're using it, the Oxford scenario planning approach out of Oxford University is quite useful, because it involves multi-stakeholder engagements and—

The Chair: It was good that you got a clear answer.

If you have additional information that you'd like to provide to the committee, I welcome you to do so.

Mr. Fergus actually took his clock down to zero before he was finished asking his question.

Hon. Greg Fergus: I'm always pushing the envelope.

The Chair: Yes. Indeed, you are.

[*Translation*]

Mr. Villemure, you have the floor for six minutes.

Mr. René Villemure (Trois-Rivières, BQ): I want to say hello to all the witnesses. Thank you for your remarkable availability.

In this first round, my questions will be for Mr. Khanna and Mr. Jenkins.

Mr. Khanna and Mr. Jenkins, I have a very general question for you. I would ask that you respond in a few seconds and then we can dig deeper.

Does facial recognition mean the end of personal freedom?

I will turn the floor over to you, Mr. Khanna.

[*English*]

Mr. Sanjay Khanna: I'll take that.

Very briefly, it depends on the contextual environment of governance of the technologies. I also think that the nature of the government within which these technologies are being employed is very important. The legislative governance and other oversight mechanisms can change. In certain contexts and kinds of government, it could very well potentially mean that—

[*Translation*]

Mr. René Villemure: Thank you very much, Mr. Khanna.

Mr. Jenkins, yes or no.

Does it mean the end of personal freedom?

[*English*]

Prof. Rob Jenkins: Do you really want a yes or no?

[*Translation*]

Mr. René Villemure: If at all possible.

[*English*]

Prof. Rob Jenkins: Not on its own. No.

[*Translation*]

Mr. René Villemure: Thank you very much.

Mr. Jenkins, in your research, you talk about intrapersonal variability.

Could you elaborate on that?

[*English*]

Prof. Rob Jenkins: Yes. Each of us has one face, which has its own appearance. That appearance changes a lot of times, not only over the long term as we grow and age, but also from moment to moment, as viewpoints change, the lighting around us changes or as we change our facial expression or talk.

There's an awful lot of variation, and this is a problem. What you're trying to do, of course, in the context of facial recognition, is to establish which of the people you know or have stored in some database you are looking at right now. That variability is difficult to overcome. You're always in the position of not knowing whether the image you have before you could count as one of the people you know or it is somebody new.

I think the variability is fundamental to the problem that we're discussing. Different people vary in their appearance, but each person also varies in their appearance. Separating those two sources of variability to understand what you're looking at is computationally difficult.

[*Translation*]

Mr. René Villemure: Thank you very much.

Mr. Khanna, in past discussions, you have alluded to biometric terrorism.

Could you tell us more about that?

● (1140)

[*English*]

Mr. Sanjay Khanna: I'm trying to recall the particular conversation you're referring to, but certainly there are scenarios within which those kinds of questions are being explored, such as the extent to which someone's identity could be stolen to identify them as a terrorist actor.

There are many plausible scenarios. I'm not sure how facial recognition technology might play into that specifically, but this is where scenario planning and those sorts of techniques can be very useful to draw in the kinds of lines of inquiry that you are concerned about.

[*Translation*]

Mr. René Villemure: Could you tell us a little more about the type of government framework we should be thinking about?

[*English*]

Mr. Sanjay Khanna: I think the frameworks can emerge only from the kind of study that this committee is doing already. There may be studies that are occurring in parallel that you need to draw upon to look at these challenges more holistically. I think that's what I would ask.

Facial recognition technology is embedded in a whole bunch of other technologies, and to accelerate development requires machine learning, computer vision and a whole bunch of other sorts of areas. It needs to be looked at quite holistically in order for Parliament to develop that kind of holistic framework that's needed, I believe.

[*Translation*]

Mr. René Villemure: Thank you very much, Mr. Khanna.

Mr. Jenkins, considering the speed at which technology is evolving, is it too late to act?

[*English*]

Prof. Rob Jenkins: No, I don't think it's too late to act. I think it's important that we act now. We should proceed on the basis of evidence—what we know—and use that evidence to try to accomplish what we want.

[*Translation*]

Mr. René Villemure: Thank you very much, Mr. Jenkins.

I will leave my remaining 30 seconds to my colleagues.

[*English*]

The Chair: All right. Thank you. It's appreciated.

We'll move now to Mr. Green for six minutes.

Mr. Matthew Green (Hamilton Centre, NDP): Mr. Chair, I'll happily take those 30 seconds as offered.

Mr. Chair, I think we can all agree that the technical aspects of this committee. I'm not sure we're going to get as deep as we need to go in order to get the kind of report that is going to be required out of this in the time we have allotted, so I'm going to put some very concise questions to all of the witnesses, starting with Dr. Watkins.

Dr. Watkins, based on your subject matter expertise, what would be your top legislative recommendations to this committee? We're going to be putting together a report and hope to have some of these recommendations reflected back to the House for the government's consideration.

Dr. Elizabeth Anne Watkins: Thank you so much. I would say that I have three top recommendations.

The top one would be to establish a moratorium. It's simply too unreliable for the futures and the livelihoods to which we are allocating responsibility.

The second two recommendations would involve accountability and transparency.

We need better insight into how these tools are being used; where the data is being stored; how decisions are being made with them; whether or not humans are involved; and how these decisions are embedded within larger bureaucratic organizational structures around how decisions are being made. Some kind of documentation to give us insights into these processes, such as algorithmic impact assessments, would be very useful.

Further, we need some kinds of regulatory interventions to produce accountability and build the kinds of relationships between the government, private actors and the public interest so that the relationships can be built to ensure that the needs of the most vulnerable are addressed.

Mr. Matthew Green: Ms. Wang, what would be your top legislative recommendations to this committee for its consideration?

Ms. Angelina Wang: I don't think I have anything else to add to what Dr. Watkins has said.

Mr. Matthew Green: Okay.

Professor Khanna, what would be your recommendations to this committee?

Mr. Sanjay Khanna: I think the safeguards need to be increased, certainly for children, marginalized groups and first nations in particular. The COVID pandemic has made things worse for all of those populations, and it's important to consider what the trajectory is in order to figure out what kinds of harms could plausibly occur in the years to come, given the shocks we've already experienced.

Mr. Matthew Green: What kind of safeguards would you recommend? Do you have any specificity around that?

Mr. Sanjay Khanna: No. I would need to take some time to think about where, specifically, the strengthening could occur, but there are some reports—for instance, the UNICEF “Policy guidance on AI for children” of November 2021—that could be very valuable in this context.

● (1145)

Mr. Matthew Green: I would put to all witnesses that if, after this, you come up with some thoughts that you weren't able to articulate in our fast-fire rounds, to consider providing them to this committee for consideration in writing, and hopefully they will also be included in our report.

Professor Jenkins, what are your top legislative recommendations for this committee's consideration?

Prof. Rob Jenkins: I would say attention to human operators in the design and implementation of facial recognition systems, transparency and the development of an expert workforce in facial recognition.

Mr. Matthew Green: Thank you very much.

Professor Watkins, I noted that in a report called “Now you see me: Advancing data protection and privacy for Police Use of Facial Recognition in Canada” that “Danish liberal deputy Karen Melchior said during parliamentary debates that 'predictive profiling, AI risk assessment and automated decision-making systems are weapons of “math destruction”', because they are 'as dangerous to our democracy as nuclear bombs are for living creatures and life.'”

Given that kind of framing of “weapons of 'math destruction'”, you noted that there's going to be an important accountability in the private sector. I note that Amazon has just had its first unionization. Hopefully, there will be some discussions around this.

What safeguards should we be putting on the private sector to ensure that these “weapons of 'math destruction'” are not unleashed on the working class?

Dr. Elizabeth Anne Watkins: That's a fantastic question. The private sector often goes under-regulated when it comes to these sorts of technologies.

There's a really fascinating model available in the state of Illinois under their Biometric Information Privacy Act. They established that, rather than having a notice and consent form, whereby users have to opt out of having their information used, it's actually the reverse, so that users have to actually opt in. Users have to be consulted before any kind of biometric information is used.

Biometric information is defined quite widely in that legislation. As far as I can recall, it includes facial imprints as well as voice imprints. This legislation has been used to wage lawsuits against companies in the private sector—for example, Facebook—for using facial recognition in their photo-identification processes.

So looking at that kind of legislation, which places control over biometric information back into the hands of users from the get-go, would be very advantageous in terms of taking steps toward putting guardrails around the private sector.

Mr. Matthew Green: I will close by saying that in one of your papers, you and your colleagues wrote that, “Despite many promises that algorithmic systems can remove the old bigotries of biased human judgement, there is now ample evidence that algorithmic systems exert power precisely along those familiar vectors.” Can you comment on that statement?

Dr. Elizabeth Anne Watkins: Thank you.

While AI, machine learning and algorithmic technologies appear to be very futuristic, very innovative and brand new, they're based on data that has been gathered over years and decades, reflecting things like institutional biases, racism and sexism.

This data doesn't come from nowhere. It comes from these institutions that have engaged, for example, in over-policing certain communities. Processes like over-policing then produce datasets that make a criminal look a certain way, when we know that doesn't actually reflect reality. These are the institutional ways in which they see populations.

Those datasets are then the very datasets on which AI and machine learning learn and they learn what the world is. So rather than being innovative and futuristic, AI, machine learning and algorithmic processes are actually very conservative and very old-fashioned, and they are perpetuating the biases that we, as a society, ought to figure out how to step forward and get past.

Mr. Matthew Green: Thank you very much.

The Chair: We now go to Mr. Kurek for five minutes.

Mr. Damien Kurek (Battle River—Crowfoot, CPC): Thank you very much.

Thank you to the witnesses for providing your expertise to the committee. Let me first make a quick comment. As a number of my colleagues have said, the way these committee reports work is that only evidence presented can end up in the report. So if there is any further documentation, thoughts or evidence that you believe would be valuable for this committee to see, including your recommendations, please feel free to send it our way. It becomes incredibly helpful as we compile reports. Let me make that offer to all of you beyond simply answering the questions that are asked here today.

To follow up on a question Mr. Green asked, Dr. Khanna, do you support a moratorium on FRT until there is a framework in place?

● (1150)

Mr. Sanjay Khanna: I do on a personal level, absolutely.

Mr. Damien Kurek: Okay, I appreciate that.

I'll put the same question to Dr. Jenkins.

Would you support a moratorium until there's a framework in place?

Prof. Rob Jenkins: I'm not sure I have a strong view on the moratorium. I'm certainly attuned to the errors that can arise in these systems, and I tend to focus on those more so than the benefits. It may not be my place to speak for the good people of Canada.

Mr. Damien Kurek: Okay, I was just asking for your perspective on that, but thank you.

I think this committee, both in this study and others, has heard a lot about the concept of consent. Certainly, when you use facial recognition on an iPhone, an android or a computer, you're consenting for your picture to be used to log in and whatnot. That is very, very different from the widespread use of scraping the Internet for images and law enforcement making a determination. That's an important differentiation.

To Dr. Khanna, in 2016 it was reported that the federal government tested facial recognition technology on millions of travellers at Toronto Pearson International Airport. What type of negative ramifications could there be for those several million travellers who passed through border control at terminal 3 at Pearson between July and December 2016 when this pilot project was running? Could you outline what some of those concerns might be in a very real-world example?

Mr. Sanjay Khanna: I think part of the concern is that we don't know. There hasn't been transparency about what some of the implications and knock-on impacts may have been, and if there were, they may even be not clear to those who may have suffered harms that they are unaware of.

It's a very tricky and challenging space to get into, which is part of the reason why transparency is such a threat to people who sometimes circumvent the law in order to gather and test what can happen through that sort of surveillance.

I'll stop there before speculating further on that question.

I just want to add very briefly that there is this question of.... There's that song. I always feel like "somebody's watching me", and Canadians can now not feel paranoid that they might be feeling that way.

Mr. Damien Kurek: Sure, I think that's certainly one of the big challenges.

I'll go to Dr. Jenkins, if I could, on that similar vein of questioning. There are about 45 seconds here, I think.

On ethical concerns relating to a pilot project like I described at Pearson International Airport, would you have any comments that you could share with the committee?

Prof. Rob Jenkins: One of my concerns would be the possibility of misidentification that is then difficult to detect or undo. I think around 100,000 passengers per day travel through Heathrow Airport, so, if we had an accuracy of 99% in that context, we'd be talking about 100 misidentifications per day, which soon adds up. It just doesn't seem sustainable to me.

Mr. Damien Kurek: With that, I'll simply use the last few seconds of my time to say thank you and again extend the offer. Please feel free to send further information to the committee as you think further about these very important issues.

Thank you very much.

The Chair: Thank you, Mr. Kurek, for keeping us on schedule.

Now we have Ms. Saks for five minutes.

Ms. Ya'ara Saks (York Centre, Lib.): Thank you, Mr. Chair.

Thank you to our witnesses today.

I'm going to start with a pretty open-ended question, but I feel that there's reason to ask it.

We've heard a lot about what's wrong with this technology and why it's bad. Is there anything good about it?

Is anyone willing to take a stab to start?

Prof. Rob Jenkins: We use automatic face recognition as a blanket term, but it can be used in many different applications. Someone mentioned the convenience of unlocking a phone or accessing account details quickly using it privately in a way similar to a password. I think that is a very different situation than using it for ambient surveillance at the scale of an entire nation.

● (1155)

Ms. Ya'ara Saks: Okay.

Going on with that, Dr. Watkins mentioned the benefits of one-to-one facial verification versus general facial recognition, so there is some advantage use to the technologies. As Mr. Khanna mentioned, as legislators we have to think about how we're behind the ball here. The curve is trending further ahead of us. At the same time, is there a way in which we can set up basic fundamental legislative guardrails at this point, whether they're anchored in privacy or in preventing scraping from open-source platforms, that could create a safety net, to start? We're constantly going to be dealing with novel and emerging technologies, but are there key principles we can look at in guardrail legislation that we should be considering?

I'm wondering if Mr. Khanna or Dr. Watkins would have any suggestions here.

The Chair: Ms. Saks, I'm just pausing for a brief moment. I think there may have been other witnesses who wanted to answer your first open question.

Ms. Ya'ara Saks: Oh. I apologize. Thank you.

The Chair: You sort of addressed your second question to Mr. Khanna, so I'll let him answer that now. If Dr. Watkins wants to go after and answer either question, then let's do that.

Go ahead, Mr. Khanna.

Mr. Sanjay Khanna: Mr. Chair, my response is that I think there could be something akin to—this is not the right phrasing—a digital charter of rights for Canadians that allows them to own and have a portable and secure form of biometric data that is considered to be sacrosanct.

I know that's a bit ambitious as a thought, but it's something that comes to mind as we have this conversation.

Ms. Ya'ara Saks: Go ahead, Dr. Watkins.

Dr. Elizabeth Anne Watkins: Thank you so much for asking this question. This is such an important question that I've been having recently with colleagues. When I beat the drum about needing to get rid of facial verification, a lot of people will then say, “Well, then, what next? What instead?” It's because these systems are often in place to guarantee worker privacy, to prevent fraud and to protect security. Workers deserve to be safe and secure and to be protected from bad actors. But there need to be alternatives in place so that facial recognition and verification is not the only way and there are ways to give workers other options. They can opt out of the verification process and opt in with perhaps a password or fingerprints.

Again, I think algorithmic impact assessments would be a really great first step to start to shed light into some of these areas where we simply don't know the types of effects and impacts these technologies are having on communities across contexts. Some information-gathering missions in the form of impact assessments, in partnership between the private and public sectors to start to assess what these impacts and effects are, would go a long way.

Ms. Ya'ara Saks: Thank you.

Through you, Mr. Chair, I have one more open-ended question.

We hear a lot of talk about a moratorium. For me, my key question is about how to implement a moratorium. My key concern is actually about the relationship between private and public enforcement, that there are contracts set up in third party structures and currently there is a loophole.

To Dr. Watkins, Mr. Khanna or Dr. Jenkins, what would be key guardrails in a moratorium?

Dr. Elizabeth Anne Watkins: One thing that struck me in reading all the language around bans that have erupted in the past few years is that they're a great start. However, these bans typically only address the way in which these technologies are used by state-backed agencies—by police departments, for example. They don't curtail the way in which surveillance tools are used in retail stores, for example, or the ways in which these types of data can then be sold to law enforcement or the back doors, exactly as you're saying, between public and private sharing of data that's been collected without consent and without knowledge.

So some kind of regulations or guardrails around how data is transferred between public and private would be a good step.

The Chair: Thank you.

With that, we will move to Monsieur Villemure for two and a half minutes.

● (1200)

[*Translation*]

Mr. René Villemure: Thank you very much, Mr. Chair.

Ms. Watkins, I listened to your testimony and I feel your overall message could be summed up in two words: “Be careful”.

Would you agree with that?

[*English*]

Dr. Elizabeth Anne Watkins: It depends to whom you're addressing such care be taken. If the definition of care includes, for example, consultation with workers or consultation with labour interests or workers' advocates to.... I have not spoken to all workers in Canada and the U.S., and I can't speak for all of them. I know that there are some workers who do advocate for facial recognition because they say they want their accounts to be secure and they want to be safe and riders to be safe, which are all good goals. But taking care, to whom it is addressed, I don't know.

[*Translation*]

Mr. René Villemure: Thank you very much, Ms. Watkins.

Mr. Khanna, of the scenarios you spoke of earlier, which one would you choose to develop facial recognition technology as you know it right now?

[*English*]

Mr. Sanjay Khanna: Picking up on the comments of my fellow panellists, and I think they've covered very good ground, you have three geographies here, U.K., U.S. and Canada, where these technologies have been employed and where a large number of lessons have been learned, particularly in the academic community, which is teaching us a great deal about what we need to safeguard, and they're doing so independently, so I think drawing on those things is critical. In terms of scenarios, again, it's looking at these technologies like FRT in the broader context of AI, machine learning and other technologies that feed into, are part of and are embedded in FRT. Governments need to look at this holistic context. We're in a digital society and things are getting ahead of us. How do we create the safeguards as our society faces greater social and economic inequities in the years ahead in part just because of COVID and things leading up to it?

Thank you.

The Chair: Thank you.

You have two and a half minutes, Mr. Green.

Mr. Matthew Green: Thank you.

Ms. Wang, in your work, you examine the amplification of bias in machine learning systems. My fear is that this committee has spent a lot of time on facial recognition, but perhaps hasn't been able to fully grasp the impacts of AI and of machine learning. Could you briefly describe the concept of bias amplification in machine learning, and perhaps describe what some of the material consequences of bias amplification are, and who tends to be most impacted?

Ms. Angelina Wang: Bias amplification refers to a notion of bias that is often thought of as just a correlation in the data. This correlation could be between some particular demographic group and some concept that they are stereotypically related to. Because machine learning models are trying to pick up on any patterns that are available in the data to learn, they frequently amplify these biases and will overpredict them whenever they are deployed.

Mr. Matthew Green: Do you have any examples in law enforcement, for instance? We're hearing terms around predictive policing and a throwback to the *Minority Report* example. Would you care to comment on any research you may have found related to law enforcement's use of machine learning?

Ms. Angelina Wang: Sure, in predictive policing, if communities of colour and different neighbourhoods with higher proportions of Black citizens may have higher levels of crime, then predictive policing models may over-report those communities in the future to be more likely to have crime, even if that is not true, and will overamplify this compared to the base rate of what the correlation actually is.

Mr. Matthew Green: To address this problem you have a tool. What do you see as the main benefits of this tool and who do you envision using this as a way to enable pre-emptive data analysis?

Ms. Angelina Wang: I'm not sure what tool you're referring to, but I think measuring these correlations and being aware that even a model with very high accuracy may not be itself amplifying biases and might be creating the same biases that are in the dataset. Even if a model isn't adding additional biases, the existing dataset will already have these too.

Mr. Matthew Green: Thank you.

For the record, I thought I saw your work attached to a revised tool, but maybe I was mistaken.

Ms. Angelina Wang: That's referring to biases in visual datasets.

● (1205)

Mr. Matthew Green: Got it.

Thank you so much, I appreciate the insight into that.

The Chair: Thank you.

Mr. Bezan for five minutes.

Mr. James Bezan (Selkirk—Interlake—Eastman, CPC): Thank you, Mr. Chair.

I want to thank our witnesses for their time and expertise on this important study we're undertaking.

I want to go around to all four witnesses to ask them a question following on where Mr. Green was going.

When you take artificial intelligence and machine learning, tie that in with facial recognition and then the possible application of that in the criminal justice system, will this significantly impede constitutional rights, our charter freedoms that we have here in Canada, as potentially being used under the Criminal Code?

I will start with Ms. Wang.

Ms. Angelina Wang: I'm sorry. I don't think I'm familiar enough with that.

Mr. James Bezan: Essentially, if FRT and AI are used as part of evidence in the conviction of individuals, would that present problems under our Criminal Code and under the Charter of Rights and Freedoms? Can we rely on FRT as enough evidence to deal with our criminal justice system and protect the rights of individuals?

Ms. Angelina Wang: I think that because you can acquire facial images without any sort of consent, and that there are so many errors and you don't really know why a model would make a particular decision, then that would go against human rights.

Mr. James Bezan: Okay.

Professor Watkins.

Dr. Elizabeth Anne Watkins: Thank you. Forgive my ignorance with the Canadian criminal charter.

In the U.S., we have a right to freedom of movement. If facial recognition technology is collecting faces from people as they move through public space, then that means the decisions they make about which public spaces through which they move could be potentially chilled. The implementation of FRT into public surveillance would have a chilling effect on that particular right. That's just one of many examples.

Mr. James Bezan: Okay.

Mr. Khanna.

Mr. Sanjay Khanna: I believe this is going to be a test that works its way through the courts. Assuming FRT and machine learning algorithms are used to identify criminals and not just their social media postings and so on, as happened in Ottawa, then tests are going to need to happen against the charter, in my view, to develop some legal precedent around this. Certainly, harms are plausible.

Mr. James Bezan: Okay.

Mr. Jenkins.

Prof. Rob Jenkins: If facial recognition accuracy is low, then there are concerns about miscarriage of justice. If it's low for some people but high for others, there are concerns about equality. If it's high for everybody, there are concerns about privacy. Those are all of the options.

Mr. James Bezan: Okay.

As we're going through this and we're hearing loud and clear on the recommendations—accountability, transparency, putting in place a moratorium until we have actual legislation in place—how do we bring forward, as parliamentarians, the proper safeguards to ensure that facial recognition is being used correctly, that bias is removed, that discrimination is eliminated, or minimized at the very least, so that we can write into the Criminal Code, the Privacy Act, PIPEDA, the guardrails we need to make sure we're not relying overly heavily on facial recognition technology, keeping in mind that there are always going to be issues around public safety and national security?

I'll go to Mr. Khanna first.

Mr. Sanjay Khanna: I'll answer that I think this parliamentary committee is taking steps in that direction by drawing on such a wide group of interprofessional and interdisciplinary experts.

Another thing that's important is for there to be opportunities for employees of companies that have the largest datasets which might be used to be compelled to provide evidence on how they're using these technologies as well, in order to inform legislative approaches. They could be company employees who come out and are whistle-blowers, who are then able to report to these committees in some sort of way.

Drawing on what people know within industry, to equalize and create a proper symmetry between what you know as legislators and what companies know internally, is probably very important.

● (1210)

The Chair: Thank you.

Now we will we go to Ms. Hepfner for five minutes.

Ms. Lisa Hepfner (Hamilton Mountain, Lib.): Thank you very much.

Thank you to the witnesses for their time today. Through the chair, I want to take advantage of the fact that we have three different countries represented here.

Starting with Mr. Jenkins, maybe you can talk to us a little bit about whether the U.K. is looking at any sorts of rules or guardrails around AI. We've talked about how legislators should approach this before it gets too late. I'm wondering what other countries are doing.

Prof. Rob Jenkins: I'm sorry, that's probably not really in my area of expertise. I can speak to the cognitive science of face recognition, but I'm not an expert on the law or the policy.

Ms. Lisa Hepfner: So, you don't know whether other countries are looking into some sort of guardrails or moratoriums or at the legislation around AI or facial recognition technology.

Prof. Rob Jenkins: I know that they are, but I don't have a deep knowledge of those processes.

Ms. Lisa Hepfner: Maybe Ms. Wang or Ms. Watkins can weigh in from a U.S. perspective. Is there any legislation that's being looked at on the U.S. side in the same way as Canada?

Ms. Angelina Wang: I'm also not familiar with this.

Dr. Elizabeth Anne Watkins: This is not my area of expertise, but I will say that one area of legislation that's been particularly useful for workers in automated decision-making is in the GDPR, and its functional right to an explanation. While the GDPR does not actually have the words “right to an explanation”, a lot of the guardrails around ensuring that companies have to provide workers with insights into how decisions are being made about them by automated systems could be a really useful model.

Ms. Lisa Hepfner: Other than what we've heard, does anybody have any further advice on how, as legislators, we can help make this practice, if it comes, other than a moratorium? Maybe more specifically what guardrails could we put into place to make sure that the risks are mitigated somewhat?

Nobody wants to tackle that.

Mr. Sanjay Khanna: I'll just bring up something I said earlier on drawing on as much research and insight as you possibly can on racialized minorities, first nations, children, or anyone who is more vulnerable to this sort of exploitation, or could be made vulnerable by changing economic circumstances that the government of the day and members of the various parties are concerned about. Looking prospectively at this to figure out how to safeguard those individuals is probably very important in the mix.

Ms. Lisa Hepfner: We've also heard today a lot about how the biases in AI come from the human biases that we have in our society, because the machines are programmed by humans. I'm wondering if this is universal, because I did see briefly one study that...algorithms that were developed in Asia may not have the same discrimination problems that algorithms developed in North America have.

Perhaps, Ms. Wang, you can talk about that. Are there better ways to develop this technology so we can still get the benefit while mitigating some of the discrimination risks?

Ms. Angelina Wang: Thank you.

I think that each model is developed in the context of the different study that it's made by, and so models developed in Asia also have lots of biases. They are just a different set of biases than models that have been developed by Canadians or Americans.

For example, a lot of object recognition tools have shown that they are not as good at recognizing the same objects—for example, soap—from a different country than the country where the dataset came from.

There are ways to get around this, but this requires a lot of different people involved with different perspectives, because there really is just no universal viewpoint. I think there's never a way of getting rid of all the biases in the model, because biases themselves are very relative to a particular societal context.

● (1215)

The Chair: Thank you.

That takes us to the end of the second block. We're going to go into subsequent rounds now.

Just for the information of members, it does not appear that there will likely be a vote at this point, so we will be able to probably complete this meeting. There will be plenty of opportunity for members to get questions in.

With that, we go now to Mr. Williams.

Mr. Ryan Williams: Thank you, Mr. Chair.

I'm going to follow my colleague, Ms. Hepfner, on some of the questioning.

Mr. Jenkins, again, you've written about the other-race effect, which is a theory that own-race faces are better remembered than other-race faces. We know that facial recognition technology is very accurate with white faces, but its accuracy drops with other skin colours.

Could this be due to the other-race effect of the programmers, essentially a predominantly white programming team creating an AI that is better at recognizing white faces? Would the same bias apply to an FRT AI developed by a predominantly, let's say, Black programming team? What does your research show, and what are you seeing in your studies?

Prof. Rob Jenkins: Bias among programmers could be a factor, but I don't think we need to invoke that to understand the demographic group differences that we see in these automatic face recognition systems.

I think that can be explained by the distribution of images that are used to train the algorithms. If you feed the algorithms mostly, let's say, white faces, then it will be better at recognizing white faces than faces from other races. If you feed it mainly Black faces, it will be better at recognizing Black faces than white faces.

Maybe the analogy with language is helpful, here. It matters what's in your environment as you are developing as a human, and it also matters as you're being programmed as an artificial system.

Mr. Ryan Williams: Ms. Wang, we know that facial recognition technology is terribly inaccurate with correctly identifying non-white people. We've heard of error rates of up to 34% for darker-skinned females. This FRT-induced digital racism is unacceptable and further reinforces why this technology should not be used for law enforcement.

You've written about mitigating bias in machine learning. How do we end this digital racism?

Ms. Angelina Wang: It's very hard to think about, because none of these technologies are ever going to be used in a vacuum, and they're always situated in a particular social context. Even if you had some sort of facial recognition system that worked perfectly, or at least the same across different people with different skin tones, the way this is used, for example, for surveillance or policing, is itself still very racist. You can never really disentangle the technology from [*Technical difficulty—Editor*]

Mr. Ryan Williams: I want to follow up on one of my colleague's questions. Can this technology be used for good?

Something I've read about is having this technology used to help curb human trafficking, finding images using AI to identify, let's say, an individual who might have been 13 when they disappeared and is now older. Using that technology for good may be used in human trafficking or solving some of that.

To all of the panellists, are there ways to have that used as a positive aspect by law enforcement and not a negative? Are there ways you can see right now that it can be something that's protected when we're looking at legislation?

Prof. Rob Jenkins: I think you characterized facial recognition technology as a tool, and, in my view, that's exactly the correct characterization. You can use a tool to try to help other people, or you can use it to try to harm other people, so we need to understand the intent of people as well as understand the capabilities of the technology itself.

● (1220)

Mr. Ryan Williams: I have the same question for anyone else who can answer that in 40 seconds.

Mr. Sanjay Khanna: I might add that consumer companies, consumer brands and retailers are looking quite closely at the technology and are advancing how they think about sentiment analysis and perceiving how customers are feeling in a branded or transactional environment. Some people might not find that particularly threatening. They might find it a benefit in some way, but guardrails are still needed around that.

There are always going to be some economic arguments for traffic, for sales and for different kinds of marketing and sales engagement and transactional opportunities that probably need to be looked at, should these technologies be employed, from an oversight standpoint.

The Chair: Thank you.

Now we have Mr. Bains for up to five minutes.

Mr. Parm Bains (Steveston—Richmond East, Lib.): Thank you, Mr. Chair.

Thank you to all of our witnesses for taking the time today.

I want to leave an open question here for any of our witnesses.

Based on your responses to Mr. Green's earlier question, there seems to be a considerable amount of legislation needed before FRT is widely used.

My questions come from Richmond, British Columbia. It's home to a strong South Asian and Asian demographic. We learned from an earlier panel expert who joined us that the VPD is using FRT without a lot of oversight.

Are any of you aware of any British Columbia law enforcement agencies using FRT?

Mr. Khanna, are you aware of any of this?

Mr. Sanjay Khanna: No, I'm not aware of how the Vancouver Police Department is using FRT.

Mr. Parm Bains: Okay, and I'll stay with you, then.

In a paper, you and your colleagues acknowledge that machine learning systems perpetuate and amplify certain biases present in the data. As a result, you developed the revised tool to enable preemptive analysis of large-scale datasets. How does the revised tool mitigate these biases?

Mr. Sanjay Khanna: I think this could be another Sanjay Khanna who happens to be working in AI and machine learning. It's not me.

Mr. Parm Bains: Oh, okay.

Ms. Angelina Wang: I think that is for me.

What the revised tool mostly does is it tries to find different patterns and correlations present in datasets that are likely to propagate into models that are trained on the dataset. It is not guaranteed by any means to find all the possible correlations that could arise. It just surfaces potential ones to the users so they can be more aware of those dataset creations when they are using a model that has been trained on such a dataset.

Mr. Parm Bains: Thank you.

I would like to share the rest of my time with my colleague, Mr. Fergus.

Hon. Greg Fergus: Thank you very much, Mr. Bains. I appreciate it.

Moving on a little bit, Dr. Wang, you mentioned earlier in your testimony, and I want to make sure I got this right, that even if we were to solve for bias and discrimination, there are some concerns that have been brought into the use of machine learning in terms of identifying folks. Can you talk a little bit about that?

Ms. Angelina Wang: Sure, yes.

Two of the points that I brought up are interpretability and brittleness. For brittleness, back actors are able to just trick the model in different ways. In the specific study I'm referring to they print a particular pattern on a pair of glasses, and through this, they can actually trick a model into thinking they're somebody completely different.

The other part is transparency. Models right now are very uninterpretable, because they have been able to pick up on whatever patterns the model has figured out as able to help it best with its task. We don't necessarily, as people, know what patterns the models are relying on. They could be relying on—

Hon. Greg Fergus: I'm sorry to interrupt. It seems, in other words, the machines are not able to tell us what it is they're using to make that kind of evaluation.

Ms. Angelina Wang: Yes, exactly.

● (1225)

Hon. Greg Fergus: Mr. Khanna, you raised the possibility of a digital charter of rights for Canadians. This is a very intriguing idea. If you were to blue-sky a little bit, what would you expect would be some of the elements inside that kind of charter?

Mr. Sanjay Khanna: One would be sanctity of personal data, so protection of certain data, like facial data, that's very intimate to the individual. I think that's part of it, but I think it would also aim to ensure alignment with the Canadian Charter of Rights and Freedoms and also be potentially a secure repository of data that can be exchanged and verified and is much more cybersecure than might be out there.

Hon. Greg Fergus: Could I very quickly ask you, Mr. Khanna, to talk about the sanctity of personal data? Does that mean we would have the right, for example, to our images, that our facial images would be ours? It's our property. Has the horse left the barn on that? Can we pull that back in?

Mr. Sanjay Khanna: I think the horse has left the barn to a great extent, to the extent that you can't draw out of Clearview AI what has been already taken. When starting to think about this, particularly as children age, we aren't the only ones who have been exposed to facial recognition technology. There are current and multiple generations that are going to be affected by this. Thinking about those who haven't yet been exposed, for whom the horse hasn't left the barn, is super important.

The Chair: Thanks. We let you go quite a bit over time there, but the testimony was good and important, and for once we're not quite jammed up against a hard stop here.

Next is Monsieur Villemure.

Go ahead, please.

[*Translation*]

Mr. René Villemure: Mr. Jenkins, how could we inject a little ethics into all this facial recognition technology? Could radical transparency or the right to be forgotten be the way to go?

I will turn the floor over to you for two and a half minutes to talk to us about this.

[*English*]

Prof. Rob Jenkins: Yes, I certainly think transparency is important. We should aim for the situation where members of the public can understand how these technologies are being used; how they could be effective; how they could be affected by them; and how they may have been affected by them.

We know from studies of the use of these technologies in the U.S., for example, that there's very little in terms of an audit trail, and I think auditing the use of face recognition technologies is going to be an important part of using them more widely.

[*Translation*]

Mr. René Villemure: Is the concept of radical transparency that people usually refer to enough or not enough?

[*English*]

Prof. Rob Jenkins: I think it's probably not enough on its own. I think it's an important component of an ethical system.

[*Translation*]

Mr. René Villemure: My last question has to do with the notion of consent.

When our image is captured as we're walking in the street, it's pretty much impossible for us to give consent.

Because that would be next to impossible, what can we expect in terms of consent or protection?

[*English*]

Prof. Rob Jenkins: Yes, it's a very difficult question. I don't have a straight answer, but I'm slightly cautious about comparing face recognition technologies against a perfectly accurate and bias-free system, because that's not an option that's on the table.

We certainly know what we get from the kinds of decision systems that have been in use for decades. Current systems involve errors and involve bias, and we don't consent to being captured on CCTV in my country or by the eyes of other people. I think it's a complicated matter.

● (1230)

[*Translation*]

Mr. René Villemure: Thank you very much.

[*English*]

The Chair: Thank you.

Now we have Mr. Green.

Mr. Matthew Green: Just to clarify, it's five minutes?

The Chair: No. It's two and half, with a little generosity.

Mr. Matthew Green: Okay, there we go. I appreciate the generosity.

We will go to Professor Khanna.

Professor Khanna, I'm hoping to explore even deeper into the relationship between corporate use of this technology and the state. As I'm to understand, companies and organizations you have advised utilize FRT. What regulations and measures must these companies adhere to in order to protect data and privacy and Canadians currently?

Mr. Sanjay Khanna: Just to clarify, I haven't advised companies on their use of FRT. It would be a bit of an accident if they happened to be using it. The projects I've worked on haven't been FRT specific. A project I've worked on recently with the World Congress on Justice for Children on the future of child justice—

Mr. Matthew Green: If I could, with two minutes and 30 seconds, you referenced guardrails. Based on your experience, does Canada have an appropriate framework to regulate the use of facial recognition technology by private and state agencies?

Mr. Sanjay Khanna: Not yet, and that's what I'm hoping you and other legislators will get your fine minds around.

Mr. Matthew Green: Okay.

Professor Jenkins, coming out of this study, given the challenges of accurate facial recognition both by humans and by artificial intelligence, do you have any further recommendations to mitigate the negative consequences of relying on facial recognition for security purposes specifically?

Prof. Rob Jenkins: One of the main concerns is mistaken identity and just the idea that an innocent person could be apprehended, accused and even sentenced for a crime they didn't commit. That's clearly an error that we want to avoid, and we also want to avoid the opposite error of failing to apprehend someone who could be a great danger to other people.

That's not new. We've been trying to mitigate those problems ever since we've had eyewitness testimony, but it takes on a new form at the scale that face recognition technologies are being deployed. To my mind, that's the main difference.

Mr. Matthew Green: Ms. Wang brought up the point of interpretability and the idea that with humans, at least you can contest a decision. However, as it stands now, there's a difficulty in contesting decisions that are being made.

Do you have any input on ways in which we can mitigate interpretability and those various ways in which we can contest decisions?

Prof. Rob Jenkins: It's possible to ask a human how they reached a particular judgment, but we don't have a great deal of insight into why we make the decisions we make sometimes. Often, we're inventing justifications post hoc that sound plausible to others, and that's news to us as much as it is to them.

I'm not sure I can make recommendations that would transfer readily from that situation to decisions made by AI.

Mr. Matthew Green: You would agree that—oh. Thank you.

The Chair: Thank you. I gave you an extra minute, more or less.

Now we'll have Mr. Bezan for up to five minutes and/or a share, if you're going to do that.

Go ahead.

Mr. James Bezan: Thank you, Mr. Chair.

I want to direct my questions toward Professor Watkins.

You talked about the whole issue of putting in place a moratorium on the use of FRT until we have the proper guardrails in place through legislation and regulation. When is it appropriate to use FRT in the workplace, by government agencies and by individuals?

Dr. Elizabeth Anne Watkins: A step toward answering that question would be to use such legislative and regulatory tools as an algorithmic impact assessment, in tandem with consultation with marginalized groups.

I can't speak for the workers as to what kinds of safety and security technologies they would like to see in their workplaces. Consult with these groups to ask them what kinds of technology they are okay with and they would prefer to comply with. Provide them with alternatives, where they can opt out of technologies that they do not wish to comply with, yet still access their means of livelihood. Those would be good steps.

● (1235)

Mr. James Bezan: Do you believe that police agencies should be allowed to use FRT?

Dr. Elizabeth Anne Watkins: No.

Mr. James Bezan: It's not just a moratorium; you're talking about a complete ban on using FRT by police agencies, border service agencies and the government in general.

Dr. Elizabeth Anne Watkins: In high-risk scenarios, where lives and livelihood are on the line, not only are these technologies at present unreliable, but they also presume that social constructs, like race and gender, are machine-readable in a person's face. That is simply untrue.

Mr. James Bezan: When you start talking about companies like Clearview AI, which have a track record of mistakenly identifying people and having a prejudice in their AI technology with FRT, should those companies be banned?

Dr. Elizabeth Anne Watkins: I think their technologies should not be used in high-risk scenarios.

Mr. James Bezan: They would be still, in your mind, okay to be used by an employer in the workplace, even though they have a track record that definitely indicates a prejudice.

Dr. Elizabeth Anne Watkins: The workplace is a very high-risk scenario. They should not be used in the workplace. They should not be used in a public space. They should not be used by police.

Frankly, I think there ought to be a moratorium until we know more about how these tools are impacting communities.

Mr. James Bezan: Ms. Wang, would you like to weigh in on this? You've done extensive study on how FRT and Clearview, in particular, has been used to marginalized people.

Do you agree with what Professor Watkins has been saying here?

Ms. Angelina Wang: Yes, I do. We understand too little right now. We shouldn't deploy them yet, if ever.

Mr. James Bezan: Okay.

You've looked at the RCMP, I believe. Would the Canada Border Services Agency...?

We often have national security threats. It's probably best, in your opinion, then, that we should not be using FRT in any of our law enforcement agencies and border control agencies here in Canada.

Ms. Angelina Wang: Yes.

Mr. James Bezan: Do you want to take that last minute?

Mr. Damien Kurek: Thank you very much.

I'm going to throw it open to all the witnesses and I will go through one by one.

Could you list off as quickly as possible examples of FRT, either public or private, just for the committee's reference? That would be very helpful.

We will start with Dr. Watkins.

Dr. Elizabeth Anne Watkins: I'm sorry. Can you repeat the question? It's examples of already deployed FRT?

Mr. Damien Kurek: Yes. Just off the top of your head, do you have any examples for the committee to use as reference points of FRT that's in use?

Dr. Elizabeth Anne Watkins: As far as I know, FRT is currently being used on Uber drivers, on Amazon delivery drivers and on at-home health care workers who are required to log into their workplace using electronic visit verification.

In terms of FRT instead of FVT, as far as I know, many police departments across the U.S. are using FVT except in those cities where there have been bans and moratoriums of which there are a handful.

Mr. Damien Kurek: Thank you.

Dr. Wang.

Ms. Angelina Wang: The ones I can think of are HireVue and some of these interviewing platforms.

Mr. Damien Kurek: Thank you.

Dr. Jenkins.

Prof. Rob Jenkins: It's often used in border control in a number of countries and in processes related to border control, such as passport renewal, to verify that the person submitting the document is who they claim to be.

It's also used in retrospective review of crowd footage to try to identify suspects who may have been captured in CCTV footage, for example.

● (1240)

The Chair: Mr. Kurek asked a question that invited a long answer from four members. I am going to ask Mr. Khanna to very quickly respond to your question if he would like. Then we're going to go directly to Ms Khalid.

Mr. Sanjay Khanna: Okay. This would need to be confirmed, but there are stories of it being used in children's toys, children's applications and things like that. That needs to be verified, but I recall seeing that in a UN report.

Thanks.

The Chair: Thank you.

As what I believe will be our last questioner, Ms. Khalid, go right ahead.

Ms. Iqra Khalid (Mississauga—Erin Mills, Lib.): Thank you very much, Mr. Chair and, through you, thank you to the witnesses for your very compelling testimony today.

Just to add to the list that you have provided, I will say that in 2018 Taylor Swift used facial recognition technology to identify some of her stalkers. That was a very fascinating, interesting and, I think, complex use of technology.

I know we have been talking about moratoriums. Perhaps I will start by asking our witnesses what a moratorium would achieve in an environment in which technology and innovation occur at such a fast pace?

Perhaps I will start with Dr. Khanna.

Mr. Sanjay Khanna: A moratorium, as you know, is a pause. It's to gather information and insight both from within organizations and from outside organizations and to gather and assess and then determine what kinds of guardrails might be imposed should that moratorium be lifted.

If this is a question of “'math' destruction”, as Dr. Watkins has described, then it does make sense to employ a moratorium and create that pause for better decision-making.

Ms. Iqra Khalid: How long do you think that pause should last? Is it to find a perfect solution that we're looking for, or is it to find a workable balance between public safety, privacy and convenience of the general public?

Mr. Sanjay Khanna: I would defer to my colleagues who have studied the developments in artificial intelligence and facial recognition technology more closely.

Ms. Iqra Khalid: If any one of you wants to take that on, please go ahead.

Prof. Rob Jenkins: Twenty years ago I used to go around telling everyone that the trouble with these facial recognition systems was that they didn't work. These days I find myself spending more time saying the trouble with these facial recognition systems is that they do work.

Over the past five years there has been impressive progress in how well these systems can identify faces. That's not to say that errors are not made. Errors are made, and sometimes they are surprising and difficult to predict, but it's absolutely right that the landscape is changing very quickly and it would change through the duration of a moratorium.

Ms. Iqra Khalid: Thank you.

I'll change tracks a little bit, although I'm not sure who to address this question to. Do any of you know if there is current technology or a system that allows Canadians to take themselves off all facial recognition databases, or all artificial intelligence databases, to be completely anonymized?

Prof. Rob Jenkins: I suspect that there are people on the panel who know more about the technology than I do, but if algorithms are trained on a huge set of images, and one of those images, or more than one, is of you, then the cake is already baked. It's difficult to unbake the cake and remove the influence of any one individual from the database on the algorithm that emerges from it.

Ms. Iqra Khalid: Thanks for that. Basically, privacy laws and the protection of privacy laws in this instance are kind of that balance and not black or white, where you opt in or opt out. You're kind of there, baked into that cake, as you said, Dr. Jenkins.

In that case, then, we see that social media companies, for example, or other platforms build in these algorithms, the artificial intelligence that creates convenience in shopping. They purchase datasets for companies to buy their customers, basically, so that they can advertise to them. Are there any regulations that you think could be part of a potential bill of rights that would protect Canadians in the way in which their data is sold to these companies?

Does anybody want to take that on? I'm sorry. I just don't know who to address this to. It's a complex one.

● (1245)

The Chair: Do you want to rephrase it really quickly, or direct it to a specific witness?

Ms. Iqra Khalid: Dr. Khanna, if you want to take it, go ahead.

Mr. Sanjay Khanna: This is where we don't really know how to protect Canadians in that way on the commercial side. That's why there has been discussion of data portability, where Canadians would have the right to their own data but also to earn money from it should they consent to it being used transactionally.

There has been a lot of push-back against that. In Australia, News Corp was finally pushed by...or Google had to pay publications for the data they were using online. That could be done, at least conceptually, for citizens as well.

The Chair: Thank you very much to all the witnesses.

Actually, there are a couple of things I want to address from the chair here.

First of all, just at the very end, Mr. Khanna, in your response to Mr. Kurek's question, you referenced a report that talked about the use of FRT with respect to children's toys. I wonder if you could supply that report or give the information to our clerk so that the report might be available to the committee for its report to Parliament.

Mr. Sanjay Khanna: I'd be happy to do so.

The Chair: That would be very much appreciated.

I want to ask a question. In the examples that have come about today, I guess the most “benign” use of FRT, if that's the word for it, or one of the more benign uses spoken of, is the one that many of us are familiar with. That's the facial recognition to unlock an iPhone or a mobile device. An individual has consented to this use and has supplied a photo of themselves for their convenience and for the biometric security around their own phone. On a personal level, I find a fingerprint much more convenient and easier, if the device will allow that, than a photo, and more reliable.

If this is one for which there seems be, on this panel or around the table, one of the more easily supported uses of this, are there problems, even at that level, of where a consumer is readily, or at least relatively readily, consenting to this type of use?

I'll maybe ask each of our panellists to weigh in on this for a quick moment. Would this be an acceptable use of FRT? Would this be included in the moratoriums that some are asking for?

Let me start with you, Dr. Watkins, just for a quick answer.

Dr. Elizabeth Anne Watkins: Thank you so much. This is a great question.

I urge the committee to think about consent in a context in which consent takes place. Consent can often be much more complex than it looks from the outside. It's not always a yes or a no, or “no I don't want to do this, so I'm going to go to the next alternative”. Often, there are no alternatives. Often, there are financial pressures that people are facing that force them to comply with these kinds of protocols.

For example, the facial verification that's in place in many gig companies, there is no alternative. If they don't comply with facial verification, they're simply off the app.

The Chair: Go ahead, Dr. Jenkins.

Prof. Rob Jenkins: I agree with all of that. Informed consent goes a long way, but it has to be informed.

● (1250)

The Chair: Go ahead, Mr. Khanna.

Mr. Sanjay Khanna: The only thing I could add to my esteemed panellists here is another potentially benign use case. They may contradict me on this, but one of them is using FRT to prevent industrial accidents. If there are operators who are tired or sleepy—this could include long-haul truckers, others in nuclear or other kinds of industrial facilities, or in health, where they're falling asleep and not being alert to one's lack of attention—it could be potentially beneficial.

The Chair: Ms. Wang, it's over to you for the final word.

Ms. Angelina Wang: It's not always clear what your image is used for. I don't know that phones these days do this, but they could be collecting that data and using it to turn other models. Consent isn't always clear, and that applies to all of these cases.

The Chair: Indeed.

Thank you all.

Mr. Fergus, go ahead. You have your hand up.

Hon. Greg Fergus: Can I ask a follow-up question to yours? I thought it was a really interesting line of investigation that you were going on, and I want to—

The Chair: Yes. Go ahead, Mr. Fergus. We have a few moments left.

Hon. Greg Fergus: Following up on Mr. Kelly's question about the individual use and whether or not there's.... Let's say there's consent. When we use the facial recognition technologies to have access to our phones, are those images shared beyond the use of that phone and the owner of that phone? Are fingerprints used? Is that information shared beyond?

I thought the consent on those kinds of phones or security devices was between the end-user and the phone itself. Please correct me if I'm wrong, because I'd like to know.

Prof. Rob Jenkins: Different companies have taken different positions on that. For some companies, everything happens on the device. For other companies, that's not a part of the deal; it can go to the cloud and from there to other places.

The Chair: Would any of our other witnesses care to comment?

Go ahead, Dr. Watkins.

Dr. Elizabeth Anne Watkins: Thank you.

I agree with Dr. Jenkins. The variance that he described, the lack of certainty that we have about whether the data stays on the phone or if it stays on company servers, or if, in fact, it's used by a third party vendor and stored on their servers, shows the need for transparency that we ought to have on where this data is being stored and how it's being used.

The Chair: Go ahead, Ms. Wang.

Ms. Angelina Wang: There's a new version of training called “federated learning”, where you can keep the image on your device the entire time, but you still consent to an update. You tell them all how it should adjust its parameters such that it can better classify your own image. In this case, the image never leaves you or your phone, but the model is still able to use that information to improve itself. However, it is a bit ambiguous how consent would work there.

The Chair: All right.

Thank you so much to our witnesses. It's been a very informative panel.

With that, the meeting is adjourned.

Published under the authority of the Speaker of
the House of Commons

SPEAKER'S PERMISSION

Also available on the House of Commons website at the
following address: https://www.ourcommons.ca

Publié en conformité de l'autorité
du Président de la Chambre des communes

PERMISSION DU PRÉSIDENT

Aussi disponible sur le site Web de la Chambre des
communes à l'adresse suivante :
https://www.noscommunes.ca