HOUSE OF COMMONS
CHAMBRE DES COMMUNES
CANADA

44th PARLIAMENT, 1st SESSION

Standing Committee on Access to Information, Privacy and Ethics

EVIDENCE

NUMBER 025

PUBLIC PART ONLY - PARTIE PUBLIQUE SEULEMENT

Thursday, June 9, 2022

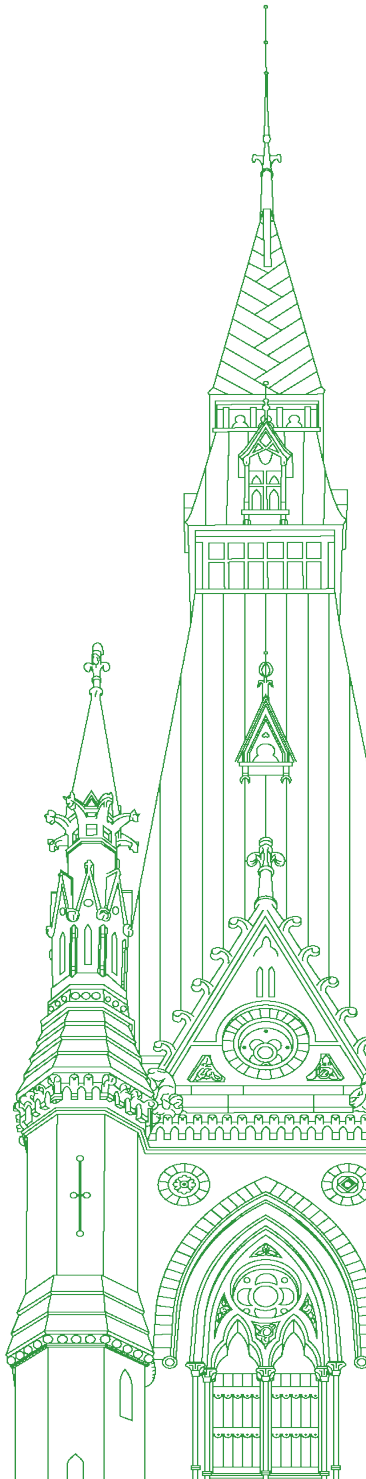

Chair: Mr. Pat Kelly

Standing Committee on Access to Information, Privacy and Ethics

Thursday, June 9, 2022

● (1625)

[*English*]

The Chair (Mr. Pat Kelly (Calgary Rocky Ridge, CPC)): I call this meeting to order.

[*Translation*]

Welcome to meeting number 25 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 on the use and impact of facial recognition technology.

[*English*]

I'd like to welcome our witness. We have today, as an individual, Nestor Maslej, research associate at the institute for human-centred artificial intelligence, Stanford University; and from the Privacy and Access Council of Canada, Sharon Polsky, president.

Mr. Maslej, you have up to five minutes for your opening statement.

Mr. Nestor Maslej (Research Associate, Institute for Human-Centered Artificial Intelligence, Stanford University, As an Individual): Good afternoon. I'd like to begin by thanking the chair and members of the committee for the invitation to speak today.

I'm Nestor Maslej, and currently I serve as a research associate for the Stanford Institute for Human-Centered AI. I am also a co-author and the lead researcher for the AI Index. Although my testimony today makes use of data from the AI Index, I am speaking as a private citizen, and my views are not representative of those of the Stanford Institute for Human-Centered AI.

The AI Index is an annual report, currently in its fifth edition, that aims to track, distill and visualize key trends in artificial intelligence. Our goal at the index is to be the best and most authoritative single source of information on trends in AI. The index aims to give policy-makers like you not only a deeper understanding of AI but also, crucially, an understanding that is grounded in empirical data.

It is this latter aim especially that informs my testimony today. I am here to answer the following question: What does data tell us about facial recognition technology? I will answer this question by tackling two sub-questions. First I will comment on capability. As of today, what can FRT do? Second I will examine usage. Who uses FRT—public and private actors—and how?

In terms of capability, there has been tremendous progress in the performance of facial recognition algorithms in the last five years. The index looked at data from the National Institute of Standards in Technology's face recognition vendor test, which comes from the U.S. Department of Commerce and measures how well FRT performs on a variety of homeland security and law enforcement tasks, such as facial recognition across photojournalism images, identification of child trafficking victims, deduplication of passports and cross-verification of visa images.

In 2017, some of the top-performing facial recognition algorithms had error rates anywhere from roughly 20% to 50% on certain FRVT datasets. As of 2021, none has posted an error rate greater than 3%, with the top-performing models registering an error rate of 0.1%, meaning that for every one thousand faces, these models correctly identify 999.

The index also shows that the performance of FRT deteriorates on masked faces but not by an overly significant degree. More specifically, performance is five to 16 percentage points worse depending on the FRT algorithm and dataset.

In terms of usage, FRTs are becoming increasingly deployed in both public and private settings. In 2021, 18 of 24 U.S. government agencies used these technologies: 16 departments for digital access or cybersecurity, six for creating leads in criminal investigations, and five for physical security. Moreover, 10 departments noted that they hoped to broaden its use. These figures are admittedly U.S.-centric, but they paint a picture of how widely governments use these tools and towards what end.

Since 2017, there has also been a total of $7.5 billion U.S. invested globally in funding start-ups dedicated to facial recognition. However, only $1.6 million of that investment has gone towards Canadian FRT start-ups. In the same time period, the amount invested in FRT technologies has increased 105%, which suggests that business interest in FRT is also growing. Our estimates also show that FRT is the 12th-most funded area out of 25 AI focus areas.

Lastly, a McKinsey survey of leading business executives, which we include in the index, shows that across all surveyed industries, only 11% of businesses had embedded facial recognition technology in their standard business processes, which trailed robotic process automation at 26% and natural speech understanding at 14% as the most embedded technologies.

In conclusion, I've presented some of the AI Index's key findings on the current capabilities and usage of FRT. It is my hope that the data I have shared usefully informs the committee's deliberation on the future regulation of facial recognition technologies in Canada. I'd be more than happy to answer any questions on the data I've presented and the implications that it may have.

Thank you.

● (1630)

The Chair: Thank you.

Now we'll go to Ms. Polsky for up to five minutes.

Ms. Sharon Polsky (President, Privacy and Access Council of Canada): Thank you so much, Chair, and good afternoon, members of the committee as well. Thank you for inviting me to appear before you today on behalf of the Privacy and Access Council of Canada.

My remarks today reflect round tables held by the council with members from across the public and private sectors, and with members of law enforcement, who agree that facial recognition is one of many digital tools that have great potential.

Like any technology, facial recognition is neither good nor bad, but it's easy to justify, especially when considered on its own. What people do with technology makes all the difference in reasonableness, proportionality and impact on lives.

Thirty-four years ago, our Supreme Court said that "privacy is at the heart of liberty in a modern state", that "privacy is essential for the well-being of the individual" and that privacy "is worthy of constitutional protection", and I dare say it still is, except that now we struggle to have any privacy, at home or away.

It's difficult now, if not impossible, to prevent our facial images being captured and analyzed and our movements and our associations being calculated and evaluated in real time. We are in view every time we go outside, and often inside as well, and our images are posted to the Internet, often without our knowledge. We haven't consented to those images being used, or to our gait, our keystrokes or other biometrics being analyzed and correlated with databases that have been amassed with information about each of us.

We haven't asked that the voice-activated devices or the messaging platforms that our children use at school and we use at work analyze our conversations or our emotions, or for our TVs to watch us watching them, yet that is now commonplace, thanks to governments and companies creating an unregulated global biometrics industry that's predicted to reach $59 billion U.S. by 2025, while the tech companies embedded in the public sector urge us to use our faces to pay for groceries and to get government services.

In the 40 years that computers have been part of our daily lives, though, there hasn't been any substantive education in Canada about privacy or access laws, or rights or responsibilities, so it's no surprise that Canadians trust that the laws themselves are enough to protect privacy or that just 14% rate their own knowledge of their privacy rights as "very good". In the meantime, there's been an onslaught of automated, privacy-invasive technologies and multi-million dollar investments in surveillance technologies to create safe communities across Canada purchased by the other 86% of people as well.

Certainly, facial recognition-enabled cameras in cars, body cams, doorbells and cellphones might help police identify a suspect or solve a crime, but even police admit that cameras and facial rec do not prevent crime, and there's little correlation between the number of public CCTV cameras and crime or safety, yet their unregulated sale and use are a self-fulfilling prophesy, because familiarity breeds consent.

Facebook, Cambridge Analytica, Cadillac Fairview and Tim Hortons are just the tip of the iceberg. Companies and governments can and do create or use technologies that violate our privacy laws because they can, because the current consent model is a fantasy, and because Mark Zuckerberg and others know that the risk of penalty is far less than the reward of collecting, manipulating and monetizing information about us.

We are at a moment, though, where three important changes are needed to help safeguard our democratic freedoms without impeding innovation and so that Canadians can regain trust in government, police and the public sector.

First, enshrine privacy as a fundamental human right for all Canadians, in person, online and at our borders.

Second, enact laws that require everyone who creates, purchases or uses technology to demonstrate that they actually have a clear and correct grasp of our privacy laws, rights and responsibilities.

Third, in the same way that vehicles and food must meet stringent government regulations before being allowed for sale or use in Canada, craft laws that put the onus on creators, requiring that technologies undergo comprehensive independent examination of their privacy access and algorithmic integrity, bias and impact before the product or platform may be acquired or used, directly or indirectly, and make sure the standards are set and the laws are written without the direct or indirect influence or input of industry.

● (1635)

Those are just a few highlights of a very complex issue that I am looking forward to discussing with you.

The Chair: Thank you.

With that, we'll proceed to questions.

For the first six minutes, we have Mr. Williams.

Mr. Ryan Williams (Bay of Quinte, CPC): Thank you to all our witnesses.

Through you, Mr. Chair, I'm going to start with Ms. Polsky.

I think you mentioned this in your recommendations. Do we need proper education campaigns for Canadians on digital consent and privacy in the digital age?

Ms. Sharon Polsky: We need substantive education that explains what privacy is. It doesn't exist yet. Really, it's no different from me tossing my car keys to the kid across the street and saying go have a good time but stay safe out there, without explaining what a stop sign is or what to do when they see one.

We need proper education, and we need the people who will be delivering the education to be educated first. We're lacking that right now.

Mr. Ryan Williams: We've talked before about consent fatigue: people who aren't reading the consent form that's about six pages long. People scroll through it, and with any kind of app that you download, it's the same kind of thing.

Do you see consent fatigue in your work, and what do we do about it?

Ms. Sharon Polsky: Well, consent fatigue is an interesting term. I think it's more a matter that people are resigned to the fact that no matter what they do or don't encounter in a so-called privacy policy, it's irrelevant, because the language that has been allowed—and frankly, embraced—by Canadian, European and other data protection regulators is allowed to be so vague as to be meaningless.

A perfect example is that you'll typically see the introductory fluff, “We respect your privacy,” and then it's, “We will collect your personal information only for business purposes,” and a list of other vague terms. Every for-profit organization's business purpose is to improve their bottom line and their net profit. Anything they can do to fulfill that obligation is a legitimate business purpose, as far as they're concerned. It's meaningless when it comes to protecting individuals. When we say yes to any of these, the companies essentially have carte blanche to share our information with their business partners, whoever they might be, wherever in the world. When it's outside of Canada, those companies do what they wish with it, for as long as they wish.

● (1640)

Mr. Ryan Williams: We know that the Privacy Act is outdated and needs to be updated, so how would you update the Privacy Act?

Ms. Sharon Polsky: On the Privacy Act, I'd say it's important to stop having so many fractured puzzle pieces of privacy legislation—federally and provincially and territorially. With each one, although they're very much alike—they're similar in most respects—they all have different exemptions, and it's almost impossible for anybody to know what law to comply with. If it's provincial, does it comply with this...or if it's health legislation, is it public sector? Then, when it crosses the line out of the country or to a different jurisdiction, it's a nightmare for compliance.

Have one overarching piece of legislation that covers the public sector, the private sector, the non-profit sector and political parties as well, please.

Mr. Ryan Williams: Thank you.

Almost every witness who has appeared before the committee—academics, lawyers and civil liberty experts—has called for a moratorium on the use of FRT by police forces.

I know that you've conducted round tables with law enforcement officers. What was their honest opinion on FRT use?

Ms. Sharon Polsky: The facial recognition that we use right now pulls a selection of mug shots from our database, and then a person actually has to look at the suspect picture and the database and compare them. That's fine. Not one of them could wrap their heads around the idea that there is such a thing as real-time, live facial recognition already in use in some jurisdictions.

They insisted that this is what we have today. They couldn't see beyond what they use today, or the implications for privacy and security of the new technology that they're not yet using.

Mr. Ryan Williams: In your opinion, would you say that the rank-and-file members understand the FRT that they're using?

Ms. Sharon Polsky: No. Very simply, no, because they're no different from most people across Canada, and I dare say elsewhere. Without education about the correct compliance requirements, what the legislation actually means, what the technology can actually do—not the sales pitch—all they can rely on is the sales pitch from a vendor whose interest is in their commission and their company's bottom line. They are not interested in our protection or our privacy, or, frankly, the police's problems.

Mr. Ryan Williams: Would the actual rank and file that you went through those round tables with support a moratorium on the use of FRT for police?

Ms. Sharon Polsky: When they can talk about themselves in their own lives, yes. I've spoken with many members of law enforcement from across Canada in different agencies, municipal, federal and military, and they basically say this: I'm not interested in being assumed to be a criminal. It's just a matter of time until I'm identified. I want to be able to go about my business anonymously. Just because I walk outside my door, I shouldn't always be under surveillance, with somebody—I don't know who or where—trying to figure out who I am, who I'm with and what I'm doing.

When the officers are in uniform, though, they have to toe the party line.

The Chair: Thank you for that.

Now we will go to Mr. Bains for up to six minutes.

Mr. Parm Bains (Steveston—Richmond East, Lib.): Thank you, Mr. Chair, and thank you to our witnesses for joining us today.

Mr. Maslej, the Institute for Human-Centered Artificial Intelligence index report for 2022 discusses diagnostic metrics that evaluate the model's impact or performance on, for example, population subgroups or minorities compared with the entire population. Can you comment on the research and investments being made to improve diagnostic metrics so that models do not misidentify persons from subgroups and minorities?

● (1645)

Mr. Nestor Maslej: Yes. That's an excellent question.

The index doesn't look in too much detail at how much investment is being put into that area at the moment, but interest is growing. One thing I will note on the data I cited from the NSIT FRT test is that the data I looked at—

The Chair: I'm going to interrupt you, Mr. Maslej. We're having a little trouble with your audio. We had tested it earlier, and I understand that it was all right, but it's not good right now. Can you ensure that you've selected the correct headset on your Zoom application?

Mr. Nestor Maslej: Let me try once again.

Is it better now?

The Chair: I will ask the translators....

No. I will suspend the meeting for a moment while we test this and get it straightened out.

● (1645)

(Pause)

● (1650)

The Chair: Thank you.

Mr. Bains, perhaps you can briefly repeat your question, and we'll get to Mr. Maslej's response.

Mr. Parm Bains: Yes.

We were talking about improving diagnostic metrics so that models do not misidentify persons from subgroups and minorities. I was getting your thoughts on that.

Mr. Nestor Maslej: My apologies for that trouble with the microphone.

The AI Index doesn't comment directly on the amount of investment that exists in that space, but it alludes to the fact that this is becoming an area of increasing concern across a lot of different spaces. I'll highlight two points here.

The first is that with the data I provided from the NIST FRVT test, this test looks at 1:1 verification. In a brief that I submitted to the committee, in figure 1.1, I show the success of different models on a variety of different datasets. Now, in this figure, one of the things that is very clear is that the best-performing model is the one that performs on visa photos, and that's the one where you had a correct identification 999 times out of 1,000, whereas the worst model is the one that is for the wild photos dataset.

The wild photos dataset is a dataset of individuals whose faces might be obscured partially by shadows, or perhaps they weren't looking directly into the camera, and the top-performing models identified correctly 970 out of 1,000 faces. It's still very high, but there's a noticeable drop-off compared with the visa photos.

I think this is suggestive of the fact that if companies and agencies want to use these technologies and justify them on the grounds that, “Look, we tested them in the lab and they had a very high accuracy rate in the lab,” there has to be an attempt to qualify the difference between the settings in which these technologies are tested and the settings in which they are deployed. I think the committee is aware of this, but the index suggests that it is a pressing concern.

I will also add that we cite some research that was published a couple of years ago—and that I think has appeared in the committee as well—in the form of a 2018 paper by Timnit Gebru et al., entitled “Gender Shades”, which looks at the fact that a lot of facial analysis benchmarks tend to be overwhelmingly composed of light-skinned individuals, leading to subsequent bias, and that existing algorithmic systems tend to disproportionately misclassify darker-skinned females as the most misclassified group.

We allude to this, and I think there is a general sense in the research community that there should have been more work being done in this space, but I would be unable to comment as to the exact amount of investment that is being put into this particular field at the moment.

Mr. Parm Bains: In November 2020, HAI published a report, "Evaluating Facial Recognition Technology". I think you talked a bit about this in terms of the clarity of the images.

One of the concerns raised is that "FRT vendors may train their images with well-lit, clear images and with proper software usage from machine learning professionals", but when deployed by law enforcement, FRTs rely on images produced by body cameras and other sources in "suboptimal" conditions.

Is this a problem that can be corrected? With respect to body cameras and the law enforcement technology that they're using, how can that be improved?

Mr. Nestor Maslej: It might be outside of the scope of my area of expertise to identify ways in which they could be improved.

I would, however, say that in the paper you're citing, the issue they talk about there is "domain shift", which is the fact that very often the settings in which some of these algorithms are tested are radically different from the settings in which they are deployed.

At the minimum, there ought to be some kind of clarity and honesty in terms of the agencies that use these tools—whether it's companies or agencies—about the extent to which a difference exists between testing conditions and conditions in which these tools are actually deployed. I think it would be problematic if there were such a big discrepancy, that these tools were tested in one setting but then deployed in completely different settings. If there isn't a clear sense and a clear understanding as to whether this difference exists, then these technologies perhaps have a great likelihood of being misused and serving more nefarious purposes.

● (1655)

Mr. Parm Bains: That takes me to my next question, which is about human error. It's also a concern with FRT. The same report indicates that while Amazon Rekognition "recommends a 99% confidence threshold on identity matching for use in law enforcement", one of the sheriff's offices interviewed for the report stated, "We do not set nor do we utilize a confidence threshold."

Do you think any use of FRT requires a trained professional who understands the technology's structure and design?

Mr. Nestor Maslej: Again, I didn't contribute to that report directly, so I wouldn't be best suited to answer the question, but I think perhaps the issue that the report is getting at there is one of institutional shift.

You might potentially have these technologies used by different individuals in different parameters. Certainly, being trained in the usage of these systems can be important, but I think there's also a recognition here that unless there is some kind of set regulatory standard about what is an acceptable benchmark or an acceptable framework, you might have different jurisdictions using these technologies in different ways. On the question of what this acceptable benchmark is, again, that is outside my area of expertise. I will leave it to you policy-makers to crack that one.

I think the point being made is that if a threshold doesn't exist, it's a lot likelier that individual agencies will make these assertions themselves. There are reasons certain agencies might favour lower or higher thresholds, and this can lead to potential misuse with some of the technology.

The Chair: Thank you, Mr. Maslej.

[*Translation*]

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

Mr. René Villemure (Trois-Rivières, BQ): Thank you, Mr. Chair.

I thank the witnesses for joining us.

Ms. Polsky, if we define understanding as the ability to grasp everything that is at stake, I conclude from your statement that there is a great deal of misunderstanding among people, governments and users—in short, everyone who is involved to a greater or lesser degree in facial recognition.

I have three questions for you, Ms. Polsky.

In reality, the consent we give by clicking on "I accept" is not a choice. We have no choice but to consent. Is that right?

[*English*]

Ms. Sharon Polsky: That is, as I said in my remarks, the consent fantasy. Mr. Zuckerberg himself said to the U.S. Congress that even he doesn't read these things. The last time I counted—yes, I did count—the Google privacy policy, it was 38 pages long. Nobody is going to read it. As a result, they're clicking on...what? They don't know.

The problem or the catch there is that at least under Canadian legislation, a company or an organization is supposed to collect personal information only after they have informed consent. When even Mr. Zuckerberg acknowledges that nobody reads these things, they are collecting personal information without informed consent, contrary to the provisions of PIPEDA and, I think, all of the other privacy laws across Canada.

[*Translation*]

Mr. René Villemure: I think that those who wrote the policies have not read them.

In Europe, users can continue without clicking on "accept". Do you think that should also be implemented in Canada?

[*English*]

Ms. Sharon Polsky: If you're talking about getting rid of the cookie consents, that has become a farce, quite bluntly. You see them on so many websites. There are some websites where you can go in to adjust your cookie settings, but then you can't get past that. You must accept all cookie settings, which is contrary to the GDPR.

To say get rid of consent policies, well, the way to do that is to put the onus not on you and me and individuals to read through these library books, but on the organizations. Require them by law to stop collecting and distributing our information.

[*Translation*]

Mr. René Villemure: Do you trust the industry to self regulate?

[*English*]

Ms. Sharon Polsky: We've already seen that happening in the United States, where the big technology companies have literally written the legislation that is being passed in several states. They call that privacy law. It's not. It doesn't protect individuals. It doesn't give them any greater right to protection or privacy. That's why in my remarks I said to craft these laws without the direct or indirect involvement of industry.

● (1700)

[*Translation*]

Mr. René Villemure: You currently don't know anyone in the industry or any business that, in an effort to self regulate, would expand best practices, as we would like to see it done?

[*English*]

Ms. Sharon Polsky: There were many.... Not only am I the president of the Privacy and Access Council of Canada, but for roughly 30 years I've been doing privacy impact assessments and privacy consulting privately. Through that, I have been invited inside everything, from governments and public bodies to Fortune 100s.

I understand how they operate, the technologies they build and what they deploy. There are some that sincerely believe they're doing it right, but remember, without education, they're misled. They're maybe assessing their own understanding a little favourably.

You get into situations where it's a Canadian company and they store the information in Canada, but they use third parties in the United States to provide essential services for that app or that service. Companies don't recognize that as a problem. They don't notify users. They don't have any concern. They're ignorant, totally unaware that there is a privacy implication.

[*Translation*]

Mr. René Villemure: To your knowledge, is the RCMP using facial recognition? If so, does it understand what it entails?

[*English*]

Ms. Sharon Polsky: I have had the opportunity to speak with a couple of very senior members of the RCMP, and they had, I think, a solid understanding. They are genuinely concerned. Their hands, I might say, are tied sometimes. Sometimes the technology is bought or trialled by somebody, and nobody else knows it.

That goes on in private sector organizations also. Instead of going through an approval process, somebody goes and buys something and plugs it in. If you don't know it's there, you can't monitor it and you can't sanction it.

Whether the RCMP is actually using facial recognition, I don't know for sure, but I wouldn't doubt it.

[*Translation*]

Mr. René Villemure: Can you say a few more words about the required education? What kind of content is needed to raise awareness among people who are involved in that education?

[*English*]

Ms. Sharon Polsky: There was a professor at McGill a few years ago, with whom I was discussing developing some education to be rolled out to schools. There are media organizations and various privacy commissioners across the country, that have developed little courses, little programs. They're available; they're not mandatory, though.

I recognize that there's a problem, because it's federal, but education is provincial. However, to have, first of all—

The Chair: I'm very sorry, but we're getting to be quite a bit over time. We're quite pressed today in our schedule.

I am going to have to switch and give the floor to Mr. Green, for up to six minutes.

Mr. Matthew Green (Hamilton Centre, NDP): Thank you very much, Mr. Chair. My questions will be for Dr. Maslej.

The rapid expansion and development of AI technology comes with significant risks. Chapter 3 of your AI Index report outlines some of the harms of AI technologies, including “commercial facial recognition systems that discriminate based on race, résumé screening systems that discriminate on gender, and AI-powered clinical health tools that are biased along socio-economic and racial lines. These models have been found to reflect and amplify human social biases, discriminate based on protected attributes and generate false information about the world.”

Could you please elaborate on some of these harms and risks posed by the use of AI technologies, particularly at a time when investment in and development of technologies are so rapidly accelerating?

Mr. Nestor Maslej: Certainly. That's an excellent question. It speaks to the fact that although regulation of facial recognition technologies certainly matters, there are also other AI use cases that might merit regulatory attention.

One of the takeaways from the data we have on facial recognition is that facial recognition is in the middle of the pack in terms of total private investment. It's more invested than things like drones or legal tech, but it's behind NLP in medical and health care, which suggests, as well, that there are other AI use cases that might merit further attention.

As the data from the McKinsey survey suggests, facial recognition is not as embedded in certain business technologies and processes as are other AI systems. Again, this doesn't necessarily mean that we shouldn't care about facial recognition regulation; it's an issue of utmost importance. It's just that we should also be cognizant of the other problems AI might pose, especially at a time when AI is becoming increasingly ubiquitous.

You alluded to a couple of the different examples that we cited in the report. I'll speak to a couple of different ones.

We talked about the fact that there are résumé-screening systems that have been shown to discriminate based on gender. We cite evidence from a newspaper article that a couple of years ago, Amazon developed a machine-learning résumé-screening system that was shown to systematically downgrade the applications of women.

Again, it would be ideal for a lot of these companies, especially the very big ones, if someone gave them 100 résumés and they could just give them to a machine that says automatically that these are the best three candidates and just hire these three people.

The reason Amazon trained a biased system was that the system was ultimately trained on data from résumés that were submitted to Amazon previously. Overwhelmingly, the résumés that tended to be submitted to Amazon were submitted by men. This reflects the fact that the tech industry is mostly dominated by men. Given that men were mostly traditionally hired, the AI system learned to penalize the term “women”. That meant that if you included a resume that, for example, said, “I was captain of the women's swim team,” the algorithm saw that historically very few women have been hired at Amazon, this person has “women” in their résumé, so let's downgrade this résumé. Amazon claimed that this tool was not actually ever deployed to make hiring decisions, but the point stands that this bias remains.

We also talk about bias in multimodal linguistic models. We talk, as well, about bias in medical image segmentation. I could go on at length about that, but I'll perhaps give you the opportunity to pose additional questions.

● (1705)

Mr. Matthew Green: I do have additional questions on the types of regulatory frameworks that you think are needed. Very important to this committee is going to be our recommendation.

Doctor, if you could, talk about what frameworks are needed to protect Canadians when it comes to the use of these AI technologies, including the ones you just listed.

Mr. Nestor Maslej: I'll say a couple of things.

First, I'm not technically a doctor. Although I very much appreciate the title, I feel that I would be remiss not to correct that.

Second, it is perhaps outside my area of expertise to offer recommendations for the committee. I understand that they are very valuable and very essential, but I feel that I can comment most on the data and what impact—

Mr. Matthew Green: We can accept that.

Since we're on familiar terms, Nestor, if I could, given your subject matter expertise—and I would suggest, given your credentials, you have subject matter expertise—could you state for the record whether or not you support a moratorium on the use of facial recognition technologies and other forms of AI by law enforcement, until government is adequately able to catch up to its impacts?

Mr. Nestor Maslej: Again, I think answering that question is outside of my scope of expertise. I would defer more to the other witness on the panel, who I think has a bit more experience in this domain and can comment a bit more authoritatively.

Mr. Matthew Green: In your view, then, could you perhaps comment on best practices that Canada should learn from, in comparison to other AI legislation and jurisdictions?

Mr. Nestor Maslej: Perhaps not necessarily as a best practice but a point of reality, one of the big takeaways from this “AI Index” report is that AI is becoming increasingly ubiquitous in all of our lives.

Ten years ago, there were a lot of AI problems that were very difficult to solve. This meant that AI was something that was just being researched, whereas, if you move forward 10 years, AI is now one of those things that are coming out of the lab and moving into the real world. A lot of companies are very excited about using AI technologies, and you're going to start seeing them used more and more. Investment in AI is going through the roof, and the number of AI patents is going through the roof.

Very often, I would say, a lot of companies are quite keen to use AI before perhaps coming to terms with some of the negative ways in which it can be deployed. As a regulator, very often it might be worth asking, when should we care about this? When is the time to regulate? I would say that—

The Chair: I have to cut you off. I'm very sorry to do so, but we're getting further behind.

I'm going to have to cut the times for the subsequent rounds, and I still think we might end up having to squeeze a little past 5:30 to get in a few minutes of committee business.

We're going to go with four minutes each for Mr. Bezan and Mr. Fergus, two each for Mr. Villemure and Mr. Green, and then four each for Mr. Kurek and Ms. Hepfner.

Go ahead, Mr. Bezan, for four minutes.

● (1710)

Mr. James Bezan (Selkirk—Interlake—Eastman, CPC): Thank you, Mr. Chair, and I want to thank our witnesses for appearing today.

Mr. Maslej, you were reeling off quite a bit of data and hard percentages after looking at it, yet, in your "AI Index" report of 2021, you said there wasn't enough data out there. Have you collected enough data to help us as regulators develop the legislative framework to control artificial intelligence or to provide the right policy framework in which to move ahead on things like facial recognition?

Mr. Nestor Maslej: I would say yes.

AI is obviously something that changes day by day. I mean, 2022 has been a tremendous year of AI progress; it seems like every week there's a new model that's breaking ground. I don't think we're ever going to get to a point where we'll have data to sufficiently know the answer to every question, but we're getting more data, and an absence of absolute data does not mean that we shouldn't take action.

We know, for instance, as I stated earlier, that a lot of these facial recognition systems perform a lot worse on these kinds of wild photos, photos where individuals are not looking into the camera straight, or where lighting is not super good, and that might have important implications for how these technologies are regulated.

We're still far away from getting to a point where we're going to have data to answer every single question, but we are getting more data, and I think the data we have at the moment is sufficient to take action on certain different issues.

Mr. James Bezan: Through the committee here, we've heard quite a bit about the shortfalls in how the data has been accumulated and how the technology has been adapted, but with bias and prejudice. Do we feel we are in a position—in your case, coming from Stanford University—where things are more balanced on that side of the equation or...? I'll ask this in my final minutes here of Ms. Polsky as well: What are the chances for abuse, the false positives and, ultimately, those who want to definitely use this to further human rights abuses?

Mr. Nestor Maslej: I can perhaps go first and then I'll defer to Ms. Polsky.

I would say again that there are definitely questions that remain unanswered, but we do have a lot of data that says things that are difficult to dispute. As mentioned, the paper that I cited earlier shows that facial recognition systems can be biased, and I think that's a generally well-accepted fact. That can be something that a committee of regulators could act on, but I'll defer to Ms. Polsky for an additional answer.

Ms. Sharon Polsky: Thank you.

I'm not an academic. I leave that to you, sir, but I go back to news reports out of the Welsh police, where the senior-ranking officer said that facial recognition—and I'm paraphrasing—came up with something like 92% false positives, and he said that was okay, because no technology is perfect. That was in 2017. In 2020 or 2021, the chief of police of Chicago, I believe it was, said that facial recognition was something like 95% erroneous.

That can have profound implications on people's lives, because once you are identified as a person of interest, you're in the system, and then any time a cop looks at you, they run your name and you're already there. There's already a presumption that they should look a little more closely at you, because the facial recognition got it wrong.

The Chair: Thank you, Ms. Polsky.

With that, we'll go to Mr. Fergus for up to four minutes.

Hon. Greg Fergus (Hull—Aylmer, Lib.): Thank you, Mr. Chair, and I'd like to thank the witnesses for being here today.

Through you, Mr. Chair, I would have asked this question of both our witnesses, just following up on what Mr. Green asked, but given that our other witness didn't want to pronounce on this issue, I will ask this question of Ms. Polsky.

Ms. Polsky, in answering a question from my colleague, you indicated the problem of false positives and the extremely high percentage of false positives, on the order of 19 times out of 20. Given that you have pointed out those numbers, and given that a number of our witnesses before this committee have pointed out that we should place a moratorium on the use of facial recognition technology by the public and perhaps even by the private sector until a framework for this technology is put in place, do you feel that there should be a moratorium? Would you agree with those witnesses that there should be a moratorium on the use of FRT in public spaces as well as private spaces?

● (1715)

Ms. Sharon Polsky: The short answer is yes, considering that several years ago, when I did some research, Toronto already had 15,000 CCTV cameras in public use. That doesn't include what's in stores, cars, cellphones and all the rest of it. Calgary replaced its lamp standards with a new type of light, but the lamp standards themselves, 80,000 of them, are capable of having microphones and high-resolution cameras, watching and listening to everything and everybody.

All too often, we have public bodies not doing the facial recognition or any of these AI-embedded technologies themselves but engaging private sector organizations to do it and getting around the accountability. I would say there needs to be a moratorium on public and private usage of it.

Hon. Greg Fergus: Thank you very much, Ms. Polsky.

Mr. Maslej, going back to chapter 3 of your report.... I had an opportunity to read your report. I also had an opportunity to speak to you about the report in advance of this meeting. I'm wondering if you could talk about how you feel or what your report says about what types of measures need to be adopted by the community to eliminate the bias that you'll find in the algorithms so that we'd be able to promote better fairness and the ability of FRT to accurately reduce bias against females or against people of colour as much as possible. What's the report on progress on that? What have you seen over the last couple of years?

Mr. Nestor Maslej: I will say that our report doesn't make any concrete recommendations for steps that should be taken. It's more trying to take stock of where the AI landscape is. I will make a couple of points, though.

First, I think the report would clearly imply that there should be a greater consciousness that AI tools are going to become increasingly ubiquitous and that a lot of these tools are flawed. They're not perfect. Sometimes people are going to use these tools without being aware of what their flaws might be. Perhaps we should be asking ourselves how they might be flawed a lot sooner, before we actually use them.

On the second point, I will say that chapter 5 looks at legislator—

Hon. Greg Fergus: I'm sorry, Mr. Maslej. Let me interrupt you there. I have very little time—

The Chair: You have none.

Hon. Greg Fergus: Oh, drat. Well, this will be very quick.

You pointed out that there's a reason to give some greater consideration to the use of this technology, but wouldn't that lead you or the report to come to the conclusion that there should be a moratorium until greater certainty or greater accuracy can be brought to bear for the use of this technology? Can you answer yes or no, if possible?

The Chair: It will have to be yes or no. We're out of time.

Mr. Nestor Maslej: Again, I will politely decline to answer that question. I don't think the report—

The Chair: Okay. Thank you.

With that, we'll move to Monsieur Villemure for two minutes.

[*Translation*]

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

Ms. Polsky, we are often told that facial recognition data must be used to create a feeling of safety. However, it seems to me that mass surveillance as you describe it and as we understand it is likely to create a feeling of unsafety rather than safety. What do you think about that?

● (1720)

[*English*]

Ms. Sharon Polsky: I have to agree with you. Keep in mind that the people who are telling us that there is great demand for these new technologies are the vendors. They're the ones who will profit from it.

It's as simple as that.

[*Translation*]

Mr. René Villemure: Thank you very much.

Mr. Maslej, on page 62 of the Artificial Intelligence Index Report 2022, you talk about algorithm error rates. How can algorithms be refined by accumulating large amounts of data without it becoming surveillance?

[*English*]

Mr. Nestor Maslej: That is one of the challenges in this kind of endeavour. I would say, broadly speaking, that it's a matter of asking the question of how we're going about collecting that data. In the absence of a regulatory framework, it is easy for different companies to operate in different kinds of capacities. If rules are more clearly ironed out and identified, it is easier to have players operating on the same field.

It is a challenge. Data is essential in the operation of these systems, but just because data is essential—I would say this as an individual, not representing my institution—that does not imply that we shouldn't have any kind of regulations or—

[*Translation*]

Mr. René Villemure: I have to interrupt you, as I have only a few seconds left.

Thank you.

Ms. Polsky, could you tell us in writing what elements you find worthwhile in the European legislation on facial recognition and data protection?

[*English*]

Ms. Sharon Polsky: I'm sorry. I missed the beginning. If I would....?

[*Translation*]

Mr. René Villemure: Could you give us something in writing on worthwhile policies in the European Union's General Data Protection Regulation, GDPR? We could learn from it.

[*English*]

Ms. Sharon Polsky: I would be pleased to.

[*Translation*]

Mr. René Villemure: Thank you very much.

[*English*]

The Chair: Go ahead, Mr. Green.

Mr. Matthew Green: Thank you very much, Mr. Chair.

For my own edification, this is a question I'm always wondering about when we have expert witnesses.

Mr. Maslej, can you state whether Stanford and the human-centred artificial intelligence centre are funded by any AI companies? Are there any potential conflicts to put on the record?

Mr. Nestor Maslej: I'm not familiar with the financial situation that surrounds the institute, but I also want to clarify that I'm speaking here as an individual and, mostly, to present the data that the index has on facial recognition. My views are not those of Stanford.

Mr. Matthew Green: That's good. Okay.

The AI Index report indicates that language models are now more capable than ever before, but also more biased. Can you elaborate on that statement?

Why are these models becoming more biased as they become more advanced, and what risk does that pose?

Mr. Nestor Maslej: Part of the reason they're becoming more biased is that, typically, these models are being fed increasingly large numbers of data. For certain models, it is advantageous to have a lot of data. The more data you give the model, the more likely it is to get some data that is not ideal.

We saw this in the report with this model clip, which is a multimodal linguistic model. This model was asked to assign the probability of an American astronaut, Eileen Collins, being.... The model was asked, "What is this image of?" The model assigned a higher probability that this photograph was of a smiling housewife in an orange jumpsuit with the American flag than that it was of an astronaut with the American flag.

That's not our finding. That's a finding from a paper of Birhane et al., 2021. It's illustrative of the fact that when you give these data a lot of models, which might be required for higher performance, they might catch some conspiratorial and biased data. If we're not filtering that data proactively, it could be very likely that these models behave in toxic and problematic ways.

Mr. Matthew Green: Can these biases be mitigated? If so, how?

Mr. Nestor Maslej: One of the things I allude to—and we talk about this in the report—is the issue of filtration, whereby you can mandate that companies filter the kind of data they use to train their systems. It's been reported in different papers that there can be a filtration tax. That is, if you filter data before you apply it to a model, the model might not perform as optimally as if you gave it unfiltered data, because the more data a model has, typically, the better it can perform on certain tasks.

Filtration can be an avenue to do that, but it also might present some trade-offs for businesses.

The Chair: We're over again. I'm sorry.

We're going to end up going a little over time here. I promised our last two members four minutes each, so I'll stick with that and go to Mr. Kurek for four minutes.

Mr. Damien Kurek (Battle River—Crowfoot, CPC): Thank you very much, Mr. Chair, and thank you to our witnesses for their expertise and for coming before us today.

Ms. Polsky, with regard to recommendation 19 in your report, if a foreign national is negatively impacted by an in-Canada FRT project, what would you suggest is the action, corrective or otherwise, that should be taken?

● (1725)

Ms. Sharon Polsky: Forgive me, because I don't have the report on screen, but if it affects a foreign national in the same way, I'd say that immigrants to Canada, who are not yet landed immigrants or citizens, have constitutional rights and charter protection, and perhaps that needs to be afforded to them as well. That is something that needs further exploration for sure.

Mr. Damien Kurek: Thank you for that.

I want to talk about the second recommendation. I appreciate that it's very helpful, just as a note for witnesses, for recommendations to come forward, because that certainly helps committees in the structure of the reports we are able to put together.

On the second recommendation, when it comes to the broad, society-wide database, how could such a database cause harm and possibly result in abuse of those who may be found in that database? The inverse of that is, for those who are not in that database, how could they possibly be harmed as well?

Ms. Sharon Polsky: A simple example of something that's going on in Canada, unregulated once again.... We surveyed the residential tenancy boards across the country. It's an app that if you want to rent an apartment or a house, it's a tenancy app. You can't put pen to paper and fill out an application form. You must use this app. It also creates, if you will, a blacklist of tenants, because the landlords can put in any information or comment, like, "She was late on her rent by two days," or, "She has a kid that's loud." They can put in whatever they want, and other potential landlords can look at this and say, "I'm not renting to this person."

What happens to the person who wants to rent a home and doesn't know this comment exists? They have no recourse. Do they become homeless, as a result? In the United States, apparently, this is going on, and some of these also require that the applicants submit their facial biometrics and other biometrics, including very personal information that wouldn't be allowed to be requested elsewhere. These have profound implications.

Mr. Damien Kurek: It's interesting, because recently I was watching a television show, and that was one of the case studies with profound possible impacts.

I'd like to move from the tenancy example, if I could. There's been a collection of data at Toronto Pearson, the international airport that I'm sure all of us on this call have been through countless times. It includes, I can only imagine, an unbelievable amount of data about Canadians, individuals visiting our country, and everyone in between. I'm curious if there are any further thoughts, recommendations, or concerns that you would highlight, and how an organization like an airport, or another entity, with law enforcement—

The Chair: You've used up all of your time. We don't have any time for an answer. I'll allow Ms. Polsky maybe five seconds, if she has a very succinct point to make in response to that lengthy question.

Ms. Sharon Polsky: Between Beyond the Border, Preclearance, the Customs Act, and all the legislation that has an impact on this, we have to look at that not in isolation but as a whole for the entire system.

The Chair: Thank you.

Ms. Hepfner has the final four minutes.

Ms. Lisa Hepfner (Hamilton Mountain, Lib.): Thank you, Chair, and thank you to the witnesses for your testimony today.

Mr. Maslej, you talked about this in your statement, but I found it really interesting that in the research from the 2022 AI Index report, facial recognition technology in 2017 had a 50% error rate, and by 2021 there were no platforms with an error rate greater than 3%.

Can you reflect on that? Why have we seen such technological advancement? How has this happened? What are the ramifications and implications for the future based on that?

● (1730)

Mr. Nestor Maslej: Yes, that's a great question.

I would say that the reason you have seen this technological advancement is that AI systems are getting better across the board. There have been a lot of developments in the architecture that powers these systems. These systems now run on better hardware, which means they're able to operate at much faster speeds. There are a lot of academic and business reasons that these systems are operating better.

In terms of consequences, again, all of this portends and points to the fact that AI is going to become a part of our lives whether we like it or not and, as I have said to you today, there is a lot that AI systems can do, but there's also a lot they can do that we didn't expect them to do or didn't perhaps want them to do. Rather than just welcoming these systems into our lives with open arms, it is important to ask what kinds of effects they might ultimately have.

Again, if we live in a world where, for instance, FRTs are now having incredibly high success rates, it might be a lot easier for companies to justify that they ought to be used, but again, that doesn't necessarily imply that we shouldn't think critically about how they ought to perhaps be regulated or managed by government officials.

Ms. Lisa Hepfner: Thank you very much. That is a very helpful answer.

Ms. Polsky, you and I have spoken already, and you've spoken a bit about this today, but you make a really good case for the fact that we don't have enough education around facial recognition technology. I'm wondering what your suggestions are around education. How do we improve this? What do we need to do specifically around education?

Ms. Sharon Polsky: I think it will be important to mandate that for the Privacy Commissioner—and this goes for the individual provinces as well that have substantially similar legislation—all legislation require that, first, the privacy commissioners' offices be fully funded and that they have a separate fund, also fully funded, for education. They don't all have an education mandate. They do some work to educate, but there needs to be a much more formalized program, because that will translate into people being more aware of the legislation and their rights and responsibilities.

They can build that into the technology, and then, once they have the technology, one thing that could be done is to test it in a sandbox, a neutral sandbox run by the Privacy Commissioner, as an opportunity not only for the commissioners and civil society groups to examine it, but for the corporations to allow it to be examined in a trusted neutral setting that does not violate their copyright or their intellectual property. That way, it gets tested and approved before it's allowed for sale in Canada.

Also, fund the education through the Privacy Commissioner's office, and again, without the influence of industry, please.

Ms. Lisa Hepfner: Thank you very much.

I have only a few seconds left, so I'll just thank you all for your time today. It was very helpful and very interesting.

The Chair: Thank you.

With that, I thank our witnesses, although I must now ask them to disconnect as quickly as possible. We will be disconnecting this Zoom and beginning an in camera Zoom for, hopefully, just a couple of quick minutes of committee housekeeping business.

With that, the meeting is suspended pending the in camera Zoom call.

[*Proceedings continue in camera*]