

Check against delivery

Speech by Patricia Kosseim, Information and Privacy Commissioner of Ontario
University of Toronto Lecture Series
March 10, 2023

The future of data privacy: A world beyond imagination

1. Introduction

Good afternoon, thank you for having me here today.

It's a real pleasure to speak to such a diverse group of people in the continual pursuit of knowledge at this grand University of Toronto. As Henry Ford once said, "Anyone who stops learning is old, whether at twenty or eighty. Anyone who keeps learning stays young."

The life-long quest for learning and knowledge should serve as an inspiration for all of us here today whether we are fresh off the vine, or of a more mature vintage.

My talk today is about the *Future of data privacy: A world beyond imagination*. But before we get to that rather existential theme, let me begin with a few basics about the Office of the Information and Privacy Commissioner of Ontario (IPC) and what we do.

2. Who is the IPC and what do we do?

The IPC oversees the administration and enforcement of provincial access and privacy laws that provide the public with a right of access to government-held information and access to their own personal information. These laws also establish rules for the protection of personal information held by public institutions, health care providers, and child and family service providers.

Our Tribunal and Dispute Resolution Services Department hears appeals of institutions' refusal to provide requesters with access to information they are seeking. We work hard to resolve these appeals through early resolution and mediation in most cases, and the small percentage of cases that don't resolve informally move on to a more formal adjudication process where we issue decisions about the matter on appeal, and can order public institutions to disclose the requested information in appropriate cases. Our

Tribunal may also investigate certain privacy complaints brought by individuals and privacy breaches when they occur.

As with most regulators with limited resources, we are under a perennial challenge to respond to the growing demands of Ontarians seeking to exercise their information rights and the rising number of privacy breaches. We have recently undertaken serious efforts to modernize our tribunal and streamline our processes so that we can become more efficient in responding to Ontarians' appeals and complaints, and still reserve a portion of our resources to proactively advance other vital parts of our mandate, that include:

- reviewing privacy policies and information management practices of public institutions
- providing comment on proposed government legislation and programs
- engaging in research in the areas of access and privacy, and
- educating the public, media and others about Ontario's access and privacy laws and current issues affecting access and privacy rights.

When I began my five-year term as commissioner in 2020, one of the first things I did was undertake a broad consultation process to determine how best to focus our efforts given such a broad mandate. With the input of many interested parties and the public, we identified key access and privacy areas that would best serve the interests of Ontarians and where the IPC could make a positive and significant impact.

As a result, we adopted four [strategic priorities](#) that would guide the work of our office over the next five years.

These are:

- Privacy and Transparency in a Modern Government
- Children and Youth in a Digital World
- Next-Generation Law Enforcement
- Trust in Digital Health

Each of these priorities address key challenges facing the privacy and access rights of Ontarians in an increasingly data-driven world, where

organizations are rapidly accelerating their use of artificial intelligence technologies.

3. Artificial intelligence

Artificial intelligence of course, is all the rave today, bringing with it exciting new promise to solve many of the world's biggest problems across all sectors of society. The trend towards automated decision-making is exploding as we speak to: accelerate the delivery of government services, solve major public health issues, reconfigure our cities, improve public safety, respond to global emergencies, enhance commercial innovation and bolster our economy. AI has the potential to rock our world as we know it.

But with these game-changing advancements in technology, come important and serious implications that we need to think about very seriously and carefully.

Despite the many promises of AI, there are also many security dangers that lurk in the background. Malicious actors have already worked out how to attribute synthetically-generated voices to people unbeknownst to them. Synthetic voices are created by uploading a few minutes of someone's voice and then, through machine learning, automatically generating that same-sounding voice to "speak" some more using any other text of their choice. You could just imagine how this technology can be used to ruin people's reputations or spread falsehoods about them.

Fully synthetic digital avatars, which look and speak like humans, are already being used to spread propaganda. They are easy and very cheap to create. As this technology comes into greater use, we may see online forums and other platforms seeded with artificially created members who seem well-minded, talking about innocent hobbies like knitting, or cycling, or gardening, only to lure the trust of human members into their online community. Once real people are ensnared into the group, these avatars can surreptitiously collect sensitive personal information about them, creating a whole new breed of cybersecurity risks.

Even well-meaning applications of artificial intelligence can inadvertently cause harm. AI raises considerable risks of discriminatory decisions or suggestions being made about people as a result of biases inherent in the

datasets that algorithms are trained on. We've seen multiple examples of persons, particularly from marginalized groups, who have been unfairly treated or targeted as a result of erroneous conclusions or inferences that are perpetuated through flawed AI applications.

[Documented examples](#) include an algorithm used in U.S. hospitals to predict who are more likely to require extensive medical care, but turned out to be heavily skewed in favor of white patients over black patients. Another algorithm used in the U.S. court systems to predict the likelihood of recidivism, was found to return twice as many false positives among black offenders compared with white offenders. A third notable example came to light when Amazon's algorithm used to accelerate their recruitment process was found to be inherently biased against women based on the gender disparities of the ten years' worth of resumes that were used to train the AI system.

[Other examples](#) of AI gone awry include instances where European names are favored over African-American names, or where facial recognition systems fail to recognize darker skin tones.

These are examples that already exist today. Can we even begin to *imagine* the possible future implications of this technology? The space between our current reality and "science-in-waiting" — as some would put it — is closing in more rapidly than we think.

4. Privacy and humanity on the brink

In one of my blogs, [Privacy and humanity on the brink](#), I wrote about the existential threat that artificial intelligence poses to our species as human beings with free will and cognitive ability to make our own decisions.

Artificial intelligence is already crossing the boundaries between *identifying* existing patterns of human behaviour, to *predicting* future human behaviour with near-perfect accuracy, to *nudging* human behaviour in ways that jeopardize our very capacity to decide what is best and how to act for ourselves.

At what point will algorithmic predictions become a self-fulfilling prophecy? Are predictions about who we are or what we'll do actually nudging us to become or to do those very things? Are predictions about intellectual

aptitude influencing the educational path we choose for our children? Are AI-enabled credit-risk assessments foreclosing the pursuit of our economic opportunities? Are algorithms being used to accelerate recruitment processes actually steering the direction of our careers away from our real dreams and aspirations? Are dating apps unwittingly persuading us whom to marry? Are social media platforms unduly influencing who we vote for? Are algorithms being trained on some deeply-engrained parts of our brains to predict a predisposition toward certain behaviour actually making us act out in that very same way?

In her recent book, *Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology*, [Nita Farahany](#), an expert on the ethics of neuroscience and professor at Duke Law School, asks us to imagine a world where the average person can peer into their own mind to eliminate painful memories or cure addictions. Conversely, we are also asked to imagine a world where a person's brain can be interrogated by others to learn their political beliefs, or where their thoughts can be used as evidence of a crime.

For those of you who think we are still a long way from these possibilities, consider that in China there are some construction, manufacturing, and transportation companies that require workers to wear helmets fitted with brain sensors. These sensors use AI technology to monitor their level of fatigue or stress.

While this technology may currently be used to make certain workplaces safer, imagine a future where your government or your employer can monitor your emotions, predict your thoughts and gauge your reaction when the name of the CEO of the company you work for is mentioned, or when you see an image of the ruling political leader.

There is no doubt that neuroscience and neurotechnology can be of immense benefit to humanity, but without safeguards, they can also seriously threaten our right to privacy.

5. The right to privacy

By allowing us to create boundaries, privacy shields us from unjustified intrusions into our lives, giving us the space to grow as individuals and explore our thoughts and ideas without fear of judgement or discrimination.

Understood this way, the right to privacy has long been conceived as the gateway to other fundamental human rights, including the right to autonomy, the right to liberty and security of the person, the right to vote, the right to be free from unreasonable search and seizure, and the freedom of association.

And then there's freedom of expression. As renowned whistleblower Edward Snowden, once said, "Privacy is the fountainhead of all other rights. Freedom of speech doesn't have a lot of meaning if you can't have a quiet space to decide what it is that you actually want to say."

As we enter the world of AI, and neurotechnology, do we need to expand the contours of our fundamental right to privacy even further, to serve as a bulwark for our freedom of thought too? To protect us against the detection and manipulation of our thoughts and preserve our right to self-determination? Some have even touted stepping up the notion of "cognitive liberty" — the idea that individuals should have the right to refuse access to their brain or have it altered in some way.

Through the rapid adoption of information technologies, combined with biotechnologies, we have created a legacy we have yet to fully understand. One that will challenge our right to privacy like never before, and in some ways, our right to be human.

Like climate change, these are not distant issues that can eventually be addressed when we get around to it, or worse, relegated to the next generation to worry about. We need to recognize their immediacy now and be actively working towards solutions that preserve the future of our privacy.

The Haudenosaunee's Seventh Generation Principle teaches us to think much longer-term beyond the here and now. Simply put, the decisions we make today must take into consideration those who are not yet born, but who will inherit the world we create.

6. Strategic foresight

One way we can go about anticipating the long-term impacts of emerging technologies, is through the discipline of "strategic foresight." Strategic foresight is a structured and systematic way of gathering information about

possible future operating environments, and projecting forward through a process of informed imagination to make smarter decisions and be better prepared.

It's about imagining several plausible futures by identifying trends, risks, and emerging issues. Through this deliberative process we can gain useful insights and better orient ourselves towards the most desirable scenario through strategic planning and policymaking.

I host a podcast series called [Info Matters](#) that focuses on current privacy and access to information issues. In [season 2, episode 9](#), I spoke with our assistant commissioner, Eric Ward, who has a background in this area through his prior work with Policy Horizons Canada, the federal government's centre of expertise for strategic foresight. Our conversation explored how strategic foresight can be used to anticipate and address emerging data issues in Ontario.

One specific area we focused on was the future of law enforcement. While the movie *Minority Report*, came out more than twenty years ago, its concept of "Precrime" — where the occurrence of a crime can be anticipated before it happens — is particularly relevant given the new and powerful technologies now available to law enforcement. In fact, many sci-fi writers and creators are insightful futurists from whom we can learn a lot.

With the advent of investigative genetic genealogy, facial recognition and other emerging technologies being increasingly adopted by police services, how might we anticipate and develop solutions to the never before seen access and privacy challenges they pose?

Essentially, strategic foresight begins by building a truly collaborative and inclusive network, with law enforcement organizations and government policy-makers, as well as those who are most affected by community-police interactions. This includes Indigenous communities, racialized Ontarians, gender diverse people, mental health advocates and communities marginalized by poverty who interact in different and specific ways with law enforcement.

Strategic foresight is at its best when it really captures and engages as many different views as possible on an issue of shared importance. With appropriate expertise, the IPC plans to use this methodology to advance

knowledge and thought leadership on issues of [Next-Generation Law Enforcement](#).

7. Legal and ethical frameworks – A call to action

As we move towards a future enabled and possibly defined by AI, we need to take action now to preserve our ability to decide our own fate. We must put in place the legal and ethical frameworks that will establish the guardrails needed to protect our gateway right to privacy and the other related human rights and fundamental values we cherish and hold dear as a society.

In Canada, the proposed [Artificial Intelligence and Data Act](#) is part of a larger suite of data protection reforms contained in Bill C-27, that would regulate certain activities related to artificial intelligence systems. Far be it from me to comment on a bill outside my jurisdiction, but I will say that Ontario has the opportunity to develop its own novel approach for governing AI and related technologies that would not be covered by the federal bill.

As a hub of AI innovation, Ontario is in a fortunate position to take the lead in developing a legal and policy framework for the use of AI. At the very least, the government of Ontario must take steps to responsibly govern its own use of artificial intelligence to enhance delivery of government services and programs with clear and transparent guardrails that Ontarians can support as being socially and ethically acceptable. A few of these guardrails could include:

- defining harms more broadly than physical, psychological, property or economic harms to an individual, to also include group harms resulting from AI systems
- taking a broader and more integrated human rights approach that concerns itself with privacy rights *and* the right to be free from discrimination
- putting in place the appropriate oversight and governance mechanisms, including whistle-blowing protections, to ensure

transparency and public accountability of both developers and users of AI systems

- ensuring that less complex and lower-risk AI processes are adopted and evaluated, first, before government agencies deploy higher-risk AI systems which may have more significant or longer-lasting detrimental impacts on Ontarians
- requiring algorithmic impact assessments integrated with [privacy impact assessments](#) to determine what data will be used, how it will be processed, how Ontarians' privacy will be protected, and the potential impact of the decisions an algorithm makes
- providing Ontarians with the right to challenge automated decisions and outcomes affecting them
- soliciting meaningful input from all affected parties to formulate a thoughtfully-articulated, principled framework that balances fundamental ethical values of autonomy, dignity and integrity of persons or groups, with broader societal interests and public good considerations.

Following its public consultations on a [Trustworthy AI Framework](#) in 2021 and publication of what it heard from interested parties, the Ontario government is expected to come back with an updated draft AI Framework. Our office is anxiously awaiting next steps. We certainly intend to comment on the draft framework, and I urge all interested parties to do the same.

While there is a certain merit in starting with a policy framework that is a more flexible, nimble and agile form of regulation to govern this fast-changing new frontier, there are certain fundamental protections that nonetheless must be enshrined in law to serve as a basic and non-negotiable safety net for Ontarians' rights.

This means basic legal obligations for public institutions to be held accountable for the personal information they collect, use and disclose, including requirements to conduct privacy impact assessments, to put in place privacy management frameworks and to report data breaches when they happen and take responsibility for mitigating and correcting them,

This also includes a proper and explicit regime for Ontarians to be able to bring privacy complaints and have them investigated by an independent oversight body like my office.

And it includes extending statutory privacy protections to cover employees in Ontario who currently are not protected by any privacy law, and bringing not for profit organizations and political parties within the ambit of Ontario's privacy laws as well.

It is of vital importance that we act on this now. The technologies that we seek to govern are continually evolving. Without solid ethical and legal guardrails in place, they could very well one day become so amorphous and pervasive as to become completely uncontrollable.

8. Conclusion

So, let me bookend this talk with another rejuvenating quote that captures the sentiment of lifelong learning, this one by Mahatma Gandhi who once said, "Live as though you were to die today. Learn as though you were to live forever."

Indeed, taking responsibility for our actions today requires that we recognize the fleeting time we have on this earth and make the most of it, while also perpetually extending our quest for understanding the impact of our actions on the "forever" we leave behind for future generations.

Thank you.