Check against delivery

Keynote by Patricia Kosseim, Information and Privacy Commissioner of Ontario
Ontario Bar Association Privacy Law Summit
October 1, 2024

Balancing AI innovation with privacy in today's digital world

Introduction

- Good morning, everyone.

- Thank you to the Ontario Bar Association for the kind invitation for me to be here today, along with Assistant Commissioner Warren Mar.

- I am also joined by my colleagues Daniel and Nanditha at the IPC exhibit booth, whom I encourage you to visit if you have any questions about our office or would like to get access to information and resources from the IPC.

- I want to extend my warmest congratulations to Molly Reynolds, this year's recipient of the Karen Spector award. It's a very well-deserved honor for all her impactful work in the field of privacy law, her mentorship among the bar, and her tireless efforts to promote diversity and inclusivity by welcoming and integrating members with true lived experiences.

- The potential and promise of artificial intelligence (AI) have captivated the imagination of media, policymakers, and regulators around the world — not to mention individual Ontarians in everyday life.

- Today, I'd like to focus my remarks on how we can foster the responsible adoption of AI in the public sector while ensuring privacy, accountability, transparency, and ethical responsibility.

- I'd also like to share some of the work we've been doing at the IPC to advance our vision of a modern and effective regulator with real world impact, particularly in AI.

AI in Ontario's public sector

- More and more, we are seeing the positive impacts of AI innovation in the public service.

- It's clear that AI offers opportunities for more efficient, effective, and responsive government, helping to improve services for residents in countless ways.

- We're seeing more and more intelligent chatbots being developed and used to respond to public inquiries on a timelier basis, capable of translating multiple languages in real time.

- Sensors and video are being combined with weather patterns and/or usage data to predict citizens' needs and deploy municipal services on a more efficient basis (whether road repair, salting, snow plowing or garbage removal).

- We're seeing tools like Chat GPT being used to synthesize and analyze large volumes of laws and regulations to accelerate finding preliminary answers to questions or sifting through reams of comments or feedback received through public consultations.

- The province of Ontario recently announced a pilot project involving more than 150 primary care physicians using AI scribes. These automated scribes summarize encounters with patients and include these as electronic medical notes to reduce the doctor's paperwork burden, improve the quality of their interactions with patients, and help them get to their next patient faster.

- AI tools are now regularly used by public institutions 24/7 to detect anomalous behavioral patterns or other suspicious activity that may be consistent with potential cybersecurity threats. These are then triaged for further human analysis, helping avert crippling impacts on essential services and critical infrastructure.

Risks of AI

- But as with any powerful technology, AI comes with its share of risks too.

- And these risks, without proper oversight and regulation, can undermine the very benefits we hope to achieve through AI.

- Privacy risks can include:

    - the possibility that the data used to train AI systems has not been collected with lawful authority;

    - that such data is flawed or biased in the first place, increasing the risks of erroneous inferences being made about someone's employability, insurability, credit worthiness, access to health, education or housing, innocence or even guilt; such consequential decisions have the potential of upending peoples' lives and exacerbating the adverse impacts experienced by vulnerable and historically disadvantaged groups;

    - the ability of AI tools to influence and even nudge people's behaviours in certain ways — to buy things, to say, believe and even, do things — undermining our sense of autonomy, creativity and human intuition;

    - the grave dangers of generative AI being used by malevolent individuals or organizations to spread misinformation and create deep fakes for exploitative means;

    - highly sophisticated phishing attacks being brought to a whole new level, thanks to the much more targeted, personalized messages made possible through AI.

- And this is to say nothing of the disruptive impacts AI may have on the labour and employment market, intellectual property rights, the environment, and its truly catastrophic potential to be used in biological arms and autonomous weapons.

- Considering the present and future uses of AI in the public sector, and extrapolating from these types of risks and beyond, it emphasizes a crucial point — without proper oversight, the promise of AI can quickly turn into peril.

IPC Advocacy for Responsible AI

- This leads us to an important question: how can we more effectively foster innovation while also ensuring that privacy rights are protected?

- Legislators worldwide have been modernizing laws, and in some cases, adopting new ones, to set out guardrails within which organizations must operate when developing or deploying AI.

- The primary aim, of course, is to protect individuals from risk of harm, but also to encourage social acceptance of these new technologies and enhance public trust in the organizations that use them.

- We also know that regulatory certainty can help improve innovation by providing frameworks in which organizations, governments, and businesses can operate with clarity and confidence.

- My office has been urging Ontario to be proactive in developing and/or future-proofing provincial laws and policies to meet these needs.

- Last year, my office issued a joint statement with the Ontario Human Rights Commission, calling on the provincial government to develop and implement effective guardrails for the use of AI technology in the public sector, addressing safety, privacy, accountability, transparency, and human rights.

- We said these rules are necessary for Ontario to fully derive the benefits of AI technologies in a manner that is ethically responsible,

sustainable, and supported by public trust.

- Recognizing the need for a harmonized approach across the country, we joined our provincial, federal, and territorial counterparts in collectively publishing *Principles for Responsible, Trustworthy, and Privacy-Protective Generative AI Technologies*.

- These FPT principles are intended to help organizations build privacy protection into the design of generative AI tools and throughout their development, provision, adoption, and downstream use.

- The principles are devised to mitigate privacy risks and promote the safe creation of AI technologies.

- Particular consideration is given to protecting vulnerable and historically marginalized groups, and ensuring ultra transparency with individuals to inform them when they are interacting with AI, or when content they are seeing has been created by a generative AI tool.

- At an international level, our office joined data protection and privacy authorities from around the world at the 45th Global Privacy Assembly to raise awareness of the need for core data protection and privacy principles to govern the development, operation, and deployment of existing and emergent AI systems.

- The IPC co-sponsored a resolution on the use of AI in the workplace, as well as a resolution on generative AI technologies.

- Both resolutions were adopted unanimously by data protection authorities worldwide and are worthy of attention.

- All of these and other advocacy efforts culminated in a strong recommendation in my annual report urging the Ontario government to adopt clear and effective guardrails around the use of AI.

Bill 194, the *Strengthening Cyber Security and Building Trust in the Public Sector Act*

- In May 2024, the Ontario government tabled Bill 194, the *Strengthening Cyber Security and Building Trust in the Public Sector Act*.

- Schedule 1 of the bill seeks to regulate the use of AI by public sector entities, among other activities.

- It proposes to set out, *through regulation*, requirements with respect to transparency, accountability, risk management, technical standards and oversight, as well as certain prohibited uses.

- While this represents an important first step, my office filed a submission with the Legislative Assembly outlining a number of recommendations on how the bill could be improved. Our submission is available on our website.

- In it, we advocate for clear statutory guardrails around the use of AI technologies, rather than leaving such fundamental matters to regulation alone.

- We recommend that the development and deployment of AI must be:

    - **Valid and reliable**: These are fundamental features of trustworthy AI tools. If an AI technology does not exhibit valid and reliable outputs for the purpose it was designed, used, or implemented, its use would not be responsible.

    - **Safe:** AI technologies must be designed to support life, physical or mental health, economic security, and the environment, with robust human monitoring and cybersecurity measures in place to ensure safety.

    - **Privacy protective:** AI technologies should be developed and adopted using a privacy by design approach. Institutions should

take measures to build in privacy and security protections, while also supporting the right to access information.

- **Transparent:** Transparency involves adopting policies and practices that make the operation of an AI tool visible and understandable. Institutions should prioritize traceability and explainability.

- **Accountable:** Institutions should establish a clear internal governance structure for the development, deployment, use, repurposing, and decommissioning of AI technologies. There should be human review of any decisions having potential consequences for individuals and they should be subject to independent oversight.

- **Human rights affirming:** AI tools should respect and affirm human rights for individuals and communities. They should aim to address and redress historical discrimination and bias, and to promote fairness and equity.

- We also recommended that certain prohibited practices, or no-go zones, be codified in the law.

- This is especially important when there is a clear threshold of risk or level of harm beyond which we all universally agree that we do not want to venture as a society.

- Another key recommendation is for a risk-based regulatory approach.

- As part of this approach, higher requirements and stronger oversight and enforcement measures would be imposed commensurate with higher levels of risk or potential harm to foster public trust in government's use of AI. There should be public consultation and engagement with impacted groups to ensure AI serves and benefits *all* Ontarians.

- These recommendations are generally in line with several other statutes worldwide, for example, the EU *AI Act*, Colorado's *Consumer Protections for Artificial Intelligence*, and *Canada's Artificial Intelligence and Data Act* in Bill C-27. None of these would apply to Ontario's public sector, hence the critical need for Bill 194.

- That said, however, in today's fast-paced digital world, we need to ensure a harmonized approach, with other national and international regulatory regimes to avoid a regulatory patchwork across regions, countries, and even provinces.

- A lot of folks today, businesses and governments alike, may be hesitant about integrating or using AI tools, in this landscape of regulatory uncertainty.

- Public sector institutions may also not understand how to engage with private industry to build these tools into their programs and service delivery.

- These concerns can cause reticence to innovate and perpetual stagnation, leaving Ontario and Ontarians behind.

- A strong regulatory framework can lead to greater clarity about the responsible use of AI and help boost multi-sectoral innovation and build public trust.

AI applications before our tribunal

- In addition to our policy advocacy work, we've also begun addressing specific AI applications in our privacy investigations.

- Last March, my office investigated the use of AI-enabled proctoring software at McMaster University.

- We analyzed the university's compliance with existing law, and recommended stronger measures to protect students' personal

information and ensure an approach that balances academic integrity and student privacy rights.

- We also made additional recommendations to address the broader privacy and ethical risks of the university's use of AI, which the university has been working to address.

- You can find a copy of our investigation report on our website too.

Public education efforts

- We've also stepped up our public education efforts around AI, to help inform Ontarians about the risks and benefits of AI, and broaden the conversation around the dinner table, so to speak.

- For example, last January for Data Privacy Day, we organized a panel of experts from academia, research, business, government, and civil society, to discuss artificial intelligence in the public sector, which you can still watch on the IPC's YouTube channel.

- We have also dedicated several Info Matters podcast episodes to the topic of AI — in the health and law enforcement sectors, which I invite you to listen to. They're available on our website, or wherever you listen to your podcasts.

- There are also a number of blogs I have written that take a deeper dive into some AI and privacy related issues, including one called *Privacy and Humanity on the Brink* — which foreshadows what I will come back to in the conclusion of my talk. These blogs are available on our website.

- A few other things of note…

Third party contracting guidance

- More and more, public sector institutions rely on third party vendors to process data on their behalf. They are increasingly turning to vendors for AI solutions too.

- Our office recently released guidance on third party contracting practices that addresses the privacy and access concerns associated with outsourcing.

- The guidance is intended to help organizations exercise due diligence when contracting out data processing, including use of AI, rather than leave things up to chance or simply relying on the third party's verbal assurances.

- Institutions must be able to demonstrate the measures they have taken to ensure privacy and access issues are addressed throughout the procurement process from planning, tendering, vendor selection, negotiation, and agreement management.

- As we often say, institutions can outsource data processing functions, but they cannot outsource accountability.

Tribunal modernization

- Also, you should know that effective September 9th, the IPC revised its code of procedure for handling appeals under FIPPA and MFIPPA.

- Our code had not been updated in any significant way since 1994, so it was time!

- We also updated related practice directions and policies that are all up on our website at ipc.on.ca.

- As a modern and effective regulator, the IPC is committed to providing Ontarians with fair and just consideration of appeals, while

being transparent about our appeal procedures, improving their timeliness, and making the most efficient use of public resources.

- Among other revisions to the code, there are now new disclosure requirements for parties using AI tools when preparing submissions to the IPC, such as:
    - the fact that AI was used;
    - the type of AI used; and
    - how AI was used.

- In addition, parties using AI tools when making representations to our office must review the accuracy and content of legal references or analyses contained in their representations that are created or generated by AI.

- They must also certify in writing to the IPC that they have completed such review.

Different plausible futures of AI

- So, where to from here?

- A lot has been written about the future of AI and its implications, and reading all these different perspectives can be quite dizzying to follow.

- Everything from the glass half full, to the glass half empty, to the glass completely shattered.

- Among the optimistic viewpoints, Bill Gates, co-founder of Microsoft has written that “Generative AI has the potential to change the world in ways that we can’t even imagine. It has the power to create new ideas, products, and services that will make our lives easier, more productive and more creative. It also has the potential to solve some of the world’s biggest problems, such as climate change, poverty and disease”.

- Yet he, among other top AI scientists and experts, including two of the three godfathers of AI, Geoffrey Hinton and Yoshua Bengio, signed on to the following joint statement: “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”

- Warning of an existential catastrophe, they call for a global moratorium on large scale AI training runs by AI labs “until scientists consider it safe to proceed” and they find a way of aligning between AI goals and human values. They call for urgent international coordination among government policymakers around the world to prohibit the creation of artificial general intelligence that could risk human extinction.

- Indeed, many scientists believe that the extinction of humans by superhuman computers is a plausible scenario, right up there with global warming and nuclear weapons.

- Historian and philosopher, Yuval Noah Harari, is among those most vocal sounding the alarm bell, saying that “potentially, we are talking about the end of human history — the end of the period dominated by human beings.”

- Similarly, Stephen Hawking once predicted that “success in creating AI would be the biggest event in human history. Unfortunately, it could also be the last.”

- Cynics like Andrew Ng of Stanford University, and former chief scientist at Chinese internet giant Baidu, have called this level of alarm so remote, “it’s like worrying about overpopulation on Mars.”

- While others criticize these end of the world scare tactics as a way of deflecting regulators’ attention away from actual AI risks here and now; by coalescing around such statements and calling for moratoria, they create a form of regulatory capture that will do nothing to redistribute the concentration of power and wealth from monolithic giants.

- Regardless of which view we hold today; it is likely to change over time as we learn more about these issues.

- As Eliezer Yudkowsky, AI researcher and founder of the Machine Intelligence Research Institute wisely said, “by far, the greatest danger of artificial intelligence is that people conclude too early that they understand it.”

- I think it’s clear we need to think carefully and simultaneously about both the short and long-term impacts of AI, to understand as best we can, its potential implications — for better and for worse, both real and existential — so we can proactively shape the destiny we want to see not only for ourselves, but for future generations.

- As the lasting words of Abraham Lincoln remind us, we “cannot escape the responsibility of tomorrow by evading it today.”

- Thank you.