Information and Privacy
Commissioner of Ontario
Commissaire à l'information et à la
protection de la vie privée de l'Ontario

Check against delivery

Keynote by Patricia Kosseim, Information and Privacy Commissioner of Ontario
Conference Board of Canada Council for Chief Privacy Officers
May 27, 2025

Privacy and the AI Regulatory Landscape

Opening

- Good morning, everyone.
- I'd like to thank the Conference Board of Canada for hosting this CPO council roundtable and inviting me to speak here today about privacy and the AI regulatory landscape in Ontario and beyond.
- Many conversations about artificial intelligence start with a reminder that these technologies are not new.
- We know that as far back as 1950, a pioneering mathematician and computer scientist by the name of Alan Turing began asking the bold question: "Can machines think?"
- Only a few years after that, Turing's seminal paper on Computing Machinery and Intelligence ushered in a whole new field of research, and the term artificial intelligence was born.
- While the origins of AI have been around for decades, I think it's fair to say that AI has only recently begun to enter the consciousness of everyday folks.
- Unbeknownst to many of us, we have been benefiting from AI applications for years, operating behind the scenes each time we asked Google to translate, complete or autocorrect something, or whenever we use Google maps to find the quickest directions to somewhere, accounting for real-time traffic and road conditions.
- But now, every Google search we make gives us an explicitly entitled "AI generated overview," instantaneously, right then and there before our very eyes.
- With the appropriate caveat that "AI responses may contain mistakes."

2 Bloor Street East
Suite 1400
Toronto, Ontario
Canada M4W 1A8

2, rue Bloor Est
Bureau 1400
Toronto (Ontario)
Canada M4W 1A8

Tel/Tél : (416) 326-3333
1 (800) 387-0073
TTY/ATS : (416) 325-7539
Web : www.ipc.on.ca

- In addition to Google Gemini, anyone can now run queries through several other large language models available at the ready, like Open AI's Chat GPT, Meta's Llama, or Microsoft's Copilot, to name a few.

- Times certainly have changed — allowing each one of us now to see and use the benefits of AI technologies every day, and better understand how they can transform our lives, not just in words, but in actions as well.

Introduction: Ontario budget 2025

- Last week, Ontario released its 2025 budget, called a Plan to Protect Ontario.

- In it, the government explicitly commits to maintaining Ontario's place as a global leader in AI development.

- Ontario wants to build further AI growth and innovation by leveraging its nearly 400 firms, over a thousand recent AI master's degree graduates, and more than $1.5 billion in AI-related venture capital investments.

- First, the government has committed to developing additional AI capacity through Ontario's post-secondary institutions and attracting top AI talent in the province by investing in world leading businesses using AI to create jobs, drive innovation, and accelerate solutions across Ontario.

- Second, the government has committed to fostering a strong AI ecosystem by investing in critical infrastructure and R&D, to support companies in developing, adopting, and bringing new AI applications to market.

- And third, through the *Strengthening Cyber Security and Building Trust in the Public Sector Act*, the government has laid the foundation for the safe and responsible use of AI by adopting a framework that will govern the use of AI technologies by Ontario's public and broader public service, committing to safeguarding public interest against new emerging risks, including discrimination, surveillance, and threats to personal privacy.

- The government itself, it says, is already using several AI applications to increase efficiency and improve services to Ontarians, including using AI to:

    - analyze data
    - monitor air and water quality to predict risks
    - support predictive soil mapping, and
    - develop chatbots, virtual assistants and AI scribes, and
    - accelerate red tape reduction.

- For our province to really become a global leader in AI, we need to start by understanding where Ontario currently stacks up compared to the rest of the world.

AI technologies in the public sector – international stats and trends

- A 2024 report by the Organization for Economic Co-operation and Development, (OECD) entitled Governing with Artificial Intelligence: Are Governments Ready? surveyed how countries around the world are using AI in the public sector to increase productivity, enhance responsiveness of public services, and strengthen the accountability of governments.

- This OECD report focuses on the role of governments not only as funders, enablers and regulators of AI, but also as users of these technologies. As users, governments assume direct responsibility for proper governance to AI and should mitigate any risks that are linked with the technologies.

- According to the OECD report, "The public sector has a special responsibility to deploy AI in a way that minimizes harm and prioritizes the well-being of individuals and communities, especially when deploying AI in sensitive policy domains such as law enforcement, immigration control, welfare benefits, and fraud prevention."

- I couldn't agree more with the special nature of this responsibility.

- The OECD report cites several examples of AI being deployed by public sector institutions around the world aimed at improving efficiency, responsiveness, and accountability.

- As examples:

    - The municipality of Nijmegen, Netherlands is using AI to monitor traffic and economic activity throughout the city to design more effective policies for ensuring road safety and providing entrepreneurial support.

    - In Queensland, Australia, the government is using AI enabled satellite imagery to map and classify land use features to improve their ability to detect threats to biosecurity and natural disaster events.

    - The Labour and Welfare Administration in Norway uses a conversational AI chatbot called Frida to help citizens access social benefits 24-7 and resolve up to 80 per cent of enquiries successfully without the need for civil servants to intervene.

    - In Japan, they're developing AI systems to detect errors when individuals register a bank account with the government to receive benefits.

    - Transport Canada uses a risk-assessment algorithm to identify potentially high-risk cargo for further screening and assessment before being loaded onto aircraft.

- And there are now countless other examples.

Risks of AI

- While AI has potential to improve government services, there are also serious risks, especially given its reliance on vast volumes of personal information.

- We've known about algorithmic bias for years, for example, documented cases where algorithms trained on biased data return flawed results in making bail assessments, allocating health resources, assessing creditworthiness or screening job applicants for positions.

- This can lead to individuals from vulnerable and marginalized communities being unfairly treated or negatively targeted.

- Numerous studies over the last several years, including research from Stanford University, have found that large language models tend to perpetuate racial biases in their responses, reinforcing stereotypes and producing systematically different outputs based on names or dialects detected in users' input cues.

- A new study from the Melbourne Law School warns against AI hiring systems that could be entrenching discrimination against women, people with disabilities, and culturally diverse applicants.

- In France, human rights organizations filed a legal challenge against the government's use of AI algorithms to detect welfare fraud, based on allegations that the algorithms used were disproportionately targeting single mothers and disabled individuals.

- It's clear from these examples, and many others, that for public sector use of AI systems to succeed, there needs to be a clear and robust AI policy framework.

AI Policy Framework in Ontario

- In Ontario, the government initiated the development of an AI governance framework when it passed Bill 194, the *Strengthening Cyber Security and Building Trust in the Public Sector Act* last November.

- The first schedule of the legislation, the *Enhancing Digital Security and Trust Act*, or EDSTA, empowers the government to regulate public sector entities in respect of:

    1. mandatory cyber security programs,
    2. responsible use of AI technologies, and;
    3. use of digital technologies affecting children and youth under eighteen.

- Regarding AI more specifically, the new law paves the way for regulations to address transparency, accountability, risk management, technical standards, and oversight.

- It also sets the foundation for prohibiting certain uses of AI technologies, or no-go zones.

- Although this is a highly laudable, relevant and timely initiative, our primary concern is that right now, EDSTA is no more than an empty shell.

- Key protections, directives, and standards have been left to future, as yet, undefined, regulations.

- For the time being, until regulations are drafted, it's unclear how the law will protect the personal information of Ontarians and guard against potential misuse of AI.

- My office provided the government with many recommendations on Bill 194, several of them focused exclusively on the AI-related provisions.

- Our submission is available on our website if you'd like to read it in full.

- Some of our key AI-related recommendations were that the statutory framework should be risk based, codify fundamental ethical principles, establish clear no-go zones, and provide for independent oversight.

- Unfortunately, none of our recommendations were adopted into the language in EDSTA prior to its adoption.

- However, there are signs that the government of Ontario is listening.

- The government's own Responsible Use of Artificial Intelligence Directive, which was published in early 2025 to guide the Government of Ontario's use of AI technologies, includes a set of principles that borrow heavily from the IPC's recommendations and mirror similar principles set out in the Canadian federal government's AI strategy for the public service.

- While a provincial public sector directive is a good first step, my office will continue to advocate strongly for the rights of Ontarians by working constructively with the government on the development of binding and enforceable regulations on provincial public sector entities and beyond, to include municipal institutions and the broader public sector as well.

Ethical principles

- In particular, we will continue to urge government to codify fundamental ethical principles guiding the development and deployment of AI systems by all public sector entities in Ontario.

- First, AI technologies should meet independent testing standards to ensure they are valid and reliable. Any tested technologies should demonstrably work as intended in the environments in which they will be used.

- Second, AI systems should be configured to be safe and support human life, physical and mental health, economic security, and the environment.

- They should be monitored and evaluated throughout their lifespan — from conception to decommissioning — to confirm they continue to support these objectives.

- Third, AI technologies should be privacy protective and developed using a privacy by design approach that anticipates and mitigates privacy risks to individuals and groups.

- This includes having clear lawful authority to collect, process, retain, and use personal data, including training data.

- This also means building in measures to ensure the accuracy of AI outputs, to protect the security of personal information, and where appropriate, to give individuals an opportunity to opt-out of automated decision-making.

- Fourth, transparency is paramount.

- Public sector entities should adopt policies and practices that make visible, explainable, and understandable how their AI systems work, and individuals should be informed when they are interacting with AI and when decisions have been made about them using AI.

- Fifth, public sector entities should be held accountable for their use of AI systems.

- They should assign clearly defined roles and responsibilities, including having an effective human-in-the-loop to allow for real-time intervention and validation of the outputs as needed.

- They should conduct algorithmic impact assessments, as well as PIAs, to identify and mitigate broader risks.

- They should document their design and application choices, as well as any consequential decisions made about individuals using AI outputs.

- Individuals must be able to challenge the accuracy of those decisions and seek recourse when they believe they have been negatively impacted by them.

- And public sector entities should be subject to review by an independent oversight body with authority to enforce these principles and require the organization to undertake remedial or corrective actions.

- And finally, AI technologies should be human rights affirming by being fair and equitable for all individuals and communities.

- This is especially important when considering historical discrimination, and bias against marginalized communities.

Consistency among principles

- These principles are aligned with the fundamental AI principles set out in a joint statement we released with the Ontario Human Rights Commission in May 2023.

- They are consistent with the Principles for Responsible, Trustworthy, and Privacy-Protective Generative AI Technologies my office issued together with Canada's federal, provincial and territorial privacy regulators in December 2023 to advocate for clear and binding guardrails around the use of AI technologies.

- And they are harmonized internationally, with two resolutions that our office co-sponsored, and were universally adopted at the Global Privacy Assembly also in 2023.

- These include a resolution on Generative Artificial Intelligence Systems, which emphasizes that AI systems must be based on the principles of data protection, privacy, human control, transparency and democratic values.

- And a resolution on Artificial Intelligence and Employment, which highlights important considerations for the use of AI technologies in the workplace.

- What's even more interesting, is that these principles align with principles that have been developed by other influential bodies beyond the international advocacy work of data protection authorities.

- This includes the OECD AI principles, which were the first intergovernmental standard on AI developed in 2019 (and updated since). These principles promote innovative, trustworthy AI that respects human rights and democratic values.

- The General Conference of UNESCO set out similar values and principles in its Recommendation on the Ethics of Artificial Intelligence in 2021.

- Building on these are the G7's Hiroshima Process International Guiding Principles for Organizations Developing Advanced AI System developed at the Hiroshima summit of 2023.

- Then, of course, there are the Ethics Guidelines for Trustworthy AI that were formally adopted in 2024 as part of the *EU AI Act*.

- All of these are consistent in supporting safe, secure, and trustworthy AI worldwide and helping nations and organizations seize the benefits of AI, while also addressing the risks and mitigating potential harms of these technologies through robust and responsible governance.

- In fact, until recently, I would have said that there was a very clear picture of universally harmonized AI principles emerging.

Recent wedge in global consensus

- However, more recent events have begun to drive a wedge between this growing body of consensus.

- Particularly since the new administration in the U.S., there has been some retrenchment in the global commitment to regulate responsible AI.

- We are seeing a shift from a safety-first approach towards AI technologies to one which is more focused on markets and rapid innovation.

- Specifically, I'd say up until 2024, there was a considerable mounting pressure globally to ensure the safety of AI first and foremost through a harmonized approach to trustworthy and responsible AI.

- Come 2025, this approach appears to be reversing with increasing pressure to eliminate barriers and put innovation at the forefront.

- Within days of taking office in January 2025, the current U.S. President reversed former President Biden's signed Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, replacing it instead with his own Executive Order on Removing Barriers to American Leadership in Artificial Intelligence. Key oversight policies have been eliminated, emphasizing an innovation-driven, industry-led, form of self governance instead.

- The third AI Safety Summit held in Paris in February 2025, tellingly renamed AI "Action" Summit, saw a marked departure from the momentum that had been building in previous summits.

- The resulting Statement on Inclusive and Sustainable Artificial Intelligence for People and the Planet, endorsed by over 60 countries, including Canada, was not signed by the U.S. or the U.K.

- In California, usually one of the most progressive states when it comes to privacy and AI regulation, recent bills aimed at supporting a safety-first AI approach have been vetoed.

- Even in Europe, the *EU AI Act* is facing pushback from industry groups that are expressing concerns over economic competitiveness.

- This includes efforts to dilute the Code of Practice on General Purpose AI, which would have the effect of weakening safeguards around transparency, accountability, and risk mitigation in the development of AI systems.

- More broadly, all of these international shifts in how to regulate AI technologies are linked with heightened geopolitical and competitive geopolitical tensions.

So where does this leave Canada and Ontario?

- So where does this leave Canada and Ontario?

- Prime Minister Mark Carney has appointed former broadcaster Evan Solomon, recently elected MP in Toronto Centre, as Minister of Artificial Intelligence — a newly created position.

- The federal mandate letter speaks of combining massive infrastructure build "with the transformative nature of AI to create opportunities for millions of Canadians to find new rewarding careers."

- The mandate letter goes on to say that: "Government itself must become much more productive by deploying AI at scale, by focusing on results over spending, and by using scarce tax dollars to catalyse multiples of private investment."

- Although this is an exciting new direction, we have yet to see what path this new role will take.

- Meanwhile, the prospects for Bill C-27, the *Artificial Intelligence and Data Act*, remains unclear.

- Given the uncertain stance of the federal government on AI governance, and the increasingly discordant discussion on AI governance globally, Ontario has a real opportunity to lead and set a model for an effective, innovative, and safe AI regulatory framework.

- The government has already established the skeleton of such a framework in EDSTA, and it's now time to put flesh on the bones.

- The coming months will be critical in terms of timing, if Ontario wants to seize the opportunity to become a responsible leader in AI.

- The overcorrection we are seeing towards deregulation south of the border misses the point entirely.

- Clear and robust regulatory policy that people can trust is an enabler of AI, not a detractor to innovation.

- According to a 2024 annual survey from Edelman Trust Barometer, only 31 per cent of Canadians trust these systems — 19 points below the global average.

- Canadians are concerned AI will take their jobs, mishandle their personal data and reinforce unfair biases in hiring and policing.

- They are also concerned about the role of AI technologies in spreading misinformation and undermining privacy.

- To foster greater trust in artificial intelligence, we need a robust regulatory framework.

- We need a policy environment that is supportive of the technology, yet safe.

- We need a principles-based and proportionate approach to risk, that provides the legal certainty and predictability public institutions need to innovate with confidence while maintaining the trust of the people they serve.

- Because ultimately, for public institutions to be successful in their adoption of AI technologies, the public must be able to trust that they are being used responsibly.

- Today's rapidly evolving digital landscape requires governments and regulators to be proactive, nimble, and solution oriented, if we are to realize the tremendous hope and opportunity of AI for an exciting and innovative future.

- And it presents a real opportunity for Ontario to step into the spotlight, seize the moment and become a world leader on the responsible and trustworthy adoption of AI systems, supporting a digital future where innovation thrives and transparency, privacy, and human rights are respected.

- AI has dominated news headlines, and the story is still unfolding. Let's make sure the story is a good one.

Looking ahead

- Looking ahead, I am excited to be embarking on my second mandate at the IPC.

- One of our priorities going forward, as I mentioned, will be working constructively with the government on regulations for the *Enhancing Digital Security and Trust Act*.

- We are particularly interested in opportunities to develop a regulatory approach for AI that is proportionate to the level of risk involved, principles-based, with clear no-go zones and a robust and independent system of oversight.

- We'll also be working to develop guidance that will help operationalize Schedule 2 of Bill 194, which brings into force mandatory privacy impact assessments and mandatory breach reporting for provincial institutions as of July 1, 2025.

- Before leaving you, I'd also like to tell you about a few things my office is currently working on that may be of particular interest to you.

- Within the next couple of weeks, my office will be releasing its 2024 annual report that highlights some of the major milestones of the past year with a particular focus on results achieved.

- This summer, we'll be releasing new guidance on the use of AI scribes in health care, designed to automate the administrative tasks of documenting patient consultations.

- Our new guidance includes privacy-related considerations that health information custodians should consider prior to procurement, implementation, and use of these AI scribes.

- And you can be sure that this guidance will be entirely consistent with the ethical AI principles I mentioned earlier.

- We’ll also be releasing updated de-identification guidelines this summer, as well.

- These are an update of our globally recognized de-identification guidelines that were first released in 2016, which we believe will be very helpful to all organizations.

- Key updates include new fact sheets, checklists, case studies, documentation requirements, and the integration of other privacy enhancing technologies such as synthetic data generation.

- Over the next few months, we’ll be consulting broadly on our strategic plan for 2025-2030. Whether that means a wholesale reorientation of our current plan or some refinements along our journey of becoming a modern and effective regulator with real world impact, we want to hear from you.

- We need to remain relevant, taking into consideration emerging risks and realities, and build a plan that’s future proof.

- Because one thing is clear. My office's commitment to Ontarians remains steadfast, and building on our past work, we will continue to do our utmost to protect the access and privacy rights of Ontarians, both now and in the future.

- Thank you.