

Check against delivery

Keynote by Patricia Kosseim, Information and Privacy Commissioner of Ontario
Association of Municipalities of Ontario (AMO)
AMO Annual Conference
August 20, 2024

Thinking Outside the Box: Artificial Intelligence in the Municipal Sector

Introduction

- Good afternoon, everyone. Thank you to AMO for the opportunity to be here today.
- The municipal sector represents such an important community for our office.
- The Information and Privacy Commissioner oversees and enforces Ontario's access and privacy laws. But it's not all about telling institutions what they must or must not do.
- We're also deeply committed to collaborating with regulated entities so we can work together to explore opportunities and overcome challenges that lie ahead.
- Today's panel will dive into the critical topic of AI adoption by municipalities.
- One of the strategic priorities of my office is *Privacy and Transparency in a Modern Government*.
- Our goal is to advance Ontarians' privacy and access rights by working with public institutions to develop bedrock principles and governance frameworks for the responsible use of digital technologies, including AI.
- As Sir Mark Walport, the former Chief Scientific Adviser to the UK government said, "Public trust is a vital condition for artificial intelligence to be used productively."
- I couldn't agree more.

- For municipalities to successfully improve the delivery of programs and services through AI, maintaining the public's trust will be paramount.
- And to maintain public trust, we have to address not only the benefits but also the risks of AI.

Benefits of AI

- Being an optimist, let me start with some examples of the concrete benefits.
- The [City of Vaughan](#) saved \$400,000 by using AI combined with weather forecasts and road conditions to optimize its road salting operations by identifying where it needed it the most, rather than just salting the entire city.
- The [City of Toronto](#) is piloting a project to fight traffic congestion using AI with cameras and sensors to adjust traffic lights and optimize traffic flows.
- The [Township of South Stormont](#) aims to improve its finances and help the environment by using AI and road network information to generate optimal routes for its garbage trucks.
- In the [City of Calgary](#), AI is being explored to prioritize road maintenance work using photo imagery from traffic cameras and the like, to detect cracks in the pavement and potholes, providing more responsive and efficient road repair services.
- Meanwhile, the [City of Edmonton](#), uses AI to identify and track the movement of wildlife in the city, to learn more about where humans and wildlife are likely to interact, and help reduce negative impacts.
- And a little further afield, the [City of Baltimore](#), has developed a next-generation 911 system, using AI to translate foreign languages in real

time and respond back and forth with callers in their native tongue through synthesized voice.

Risks of AI

- While AI holds a lot of potential to improve life in our towns and cities, we must also be mindful of the risks to our privacy and human rights, particularly when AI relies on immense volumes of personal information.
- AI can [can replicate and amplify real-world bias](#) and discrimination based on historical datasets that algorithms are trained on.
- This can lead to individuals from vulnerable and marginalized communities being unfairly treated or negatively targeted by flawed AI applications.
- For example, an algorithm used in [US hospitals](#) to predict who was more likely to require extensive medical care turned out to be heavily skewed in favor of white patients over black patients, based on historical spending figures that reflected who tended to have greater access, rather than need.
- You'll recall the algorithm used by [Amazon](#) to accelerate its recruitment processes was found to be inherently biased against women candidates, because of the 10 years worth of male dominated CVs it was trained on.
- In a more recent [study](#), researchers asked ChatGPT to explain how it was ranking resumes. In one case, it claimed that a resume referencing an [autism](#) leadership award demonstrated less of a leadership role, implying that people with autism don't make as good leaders relative to others.
- In another recent study, [researchers found](#) that depending on the dialect used for input into their AI model, the results could lead to prejudicial assumptions about people's character, employability, and criminal tendencies.

- Another risk municipalities must contend with is the deliberate use of AI by others for phishing and malware attacks.
- For example, cybersecurity researchers have demonstrated that it is possible for AI-generated malware (called [BlackMamba](#)) to mutate every time it runs, slipping through predictive cybersecurity software.
- AI has also [changed the nature of phishing attacks](#). Typical phishing attacks rely on mass emails sent indiscriminately, with the hope that a few people will click on a link or download an infected file.
- With the help of AI, criminals can now craft highly personalized and convincing phishing messages by analyzing social media profiles and online information.
- Bad actors can take [social engineering attacks](#) to new heights, drafting sophisticated emails in any language and make it look like a real human wrote them; they're now also capable of synthetically mimicking the voice of CEOs, tricking staff into sending them financial information.

Need for Guardrails / Legislation

- Whether AI is being adopted for good or for ill, we must establish legal and ethical guardrails around its development and deployment.
- We are seeing legislative proposals and frameworks emerge worldwide to address AI safety and security.
- In Europe, the EU *AI Act*, which came into effect on August 1, establishes legal obligations for providers and users depending on level of risk. It also prohibits certain AI practices altogether such as behavioural manipulation, indiscriminate scraping of facial images from the internet, the use of social scoring, and biometric categorization of individuals or groups.
- Colorado's *AI Act*, the first comprehensive AI law in the U.S., defines high-risk AI with a specific focus on bias and discrimination. Under

the law, developers must exercise reasonable care to protect against algorithmic discrimination.

- In Canada, the *Artificial Intelligence and Data Act*, part of Bill C-27, mandates measures to identify and mitigate risks of harm and monitor compliance. This federal bill passed second reading and is currently at the committee review stage.
- Closer to home, the Ontario government tabled Bill 194 that seeks to regulate the use of AI by public sector entities, including municipalities. 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 a very important first step, my office filed a submission with the Legislative Assembly, recommending how the bill could be improved.
- Among other things, we recommend that the law enshrine clear statutory guardrails around the use of AI technologies, and not leave such fundamental matters to regulation. For example, we recommend that the development and deployment of AI must be valid and reliable; safe; privacy protective; transparent; accountable; and human rights affirming.
- We call for a more transparent and inclusive public participatory process for developing AI governance frameworks. We recommend that certain prohibited practices, or no-go zones, be clearly laid out in the law. And we recommend there be a system of independent oversight to ensure accountability and help garner public trust in government's use of AI to serve and benefit *all* Ontarians.
- We are not alone. The Ontario Human Rights Commission, the Law Commission of Ontario and academic experts have made similar recommendations.

- I look forward to participating in an active public debate on these and other important matters related to Bill 194 when the Legislature resumes sitting in the fall.

AI in Ontario and the IPC's Ongoing Involvement in AI

- That said, we are not just waiting around to see what will happen with Bill 194.
- My office has been actively working for the past couple of years, advocating for the adoption of guardrails around the responsible use of AI.
- Last year, the IPC issued a joint statement with the Ontario Human Rights Commission urging the provincial government to develop and implement effective guardrails for the use of AI technology in the public sector.
- We took up the cause nationally, when we joined with our federal, provincial, and territorial counterparts to release *Principles for Responsible, Trustworthy, and Privacy-Protective Generative AI Technologies*.
- And then took it up a notch further, by co-sponsoring two international resolutions at the 45th Global Privacy Assembly that were unanimously adopted by data protection authorities around the world. One on *Generative Artificial Intelligence Systems* and the other on *Artificial Intelligence and Employment*.
- We've been actively trying to engage Ontarians in the conversation as well.
- For those of you who may have missed it, our Privacy Day event in January focused on AI in the public sector, featuring fascinating insights and different perspectives from an expert panel. You can still watch it on our YouTube channel for some great takeaways.

- We've dedicated several episodes of our Info Matters podcast to privacy and security issues arising from AI, including in the law enforcement and healthcare sectors. I invite you to have a listen.
- And we've started addressing AI issues arising in the context of some of our privacy investigations as well.
- In March, my office investigated the use of AI-enabled proctoring software at McMaster University. We recommended stronger measures to protect students' personal information and ensure an approach that balances academic integrity and student privacy rights. We also went on to make additional recommendations to address the broader privacy and ethical risks associated with the university's use of AI.
- And more recently, the IPC revised its code of procedure for processing appeals under FIPPA and MFIPPA.
- This is the first major overhaul of our code of procedure since its adoption over thirty years ago.
- 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 procedures, improving timeliness of the appeals process, and making most efficient use of public resources.
- As part of our public consultation, we received feedback from many interested parties, including municipalities, who provided us with valuable perspectives and insights, and I want to sincerely thank you for your input into this process.
- Revisions to the code include 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.
- Also, parties using AI tools when making representations to our office must review the accuracy and content of legal references or analysis contained in their representations that are created or generated by AI and certify in writing to the IPC that they have completed that review.
- We have made several other significant changes to our processes, and I encourage you all to review the new code of procedure that comes into effect on **September 9, 2024**.

Conclusion

- Two summers ago now, I wrote a blog entitled, Privacy and Humanity on the Brink, where I spoke of certain existential risks that remind us of our fragility as human beings. I believe AI is one of those.
- Without guardrails, AI can freely cross the boundary between predicting human behaviour and nudging our behaviour in ways that jeopardize our fundamental sense of human agency;
- It can wreak havoc on our individual capacity to decide for ourselves which educational opportunities to provide our children, what jobs or careers we pursue, what purchasing choices we make, and what political parties we vote for.
- AI can also nudge policy makers on how to allocate scarce resources, how to administer our criminal justice system and what investments to make now that will definitively shape our future.
- These decisions strike at the core of a fair, just and democratic society, not to mention the use of AI in warfare that goes to our global security and our very survival as a human species.
- For many of us, long-term planning three to five years out can be tough enough, let alone thinking through the broader societal implications that our actions today will have on our children's future.

- As Andrew Coyne said yesterday, we're generally not good at that. Which is why we need to learn from the Haudenosaunee people who teach us to think about more than current preoccupations and transpose us beyond the here and now. They remind us of our transient presence on earth and our responsibility to ensure a sustainable future by reflecting on how our decisions today will impact seven generations ahead.
- AI gives us tremendous hope and opportunity for an exciting and innovative future. And it is ours for the making.
- As Abraham Lincoln once said, ""The best way to predict the future is to create it." To which I would add one word... "responsibly".
- Thank you for your attention, and we'll now move into our panel session and hear from Shannon and Mike, before opening up the discussion to a Q and A.