Check against delivery

Keynote by Patricia Kosseim, Information and Privacy Commissioner of Ontario
Public Sector Innovation Show
September 10, 2024

INTEGRATING AI FOR ENHANCED CITIZEN EXPERIENCE: BALANCING PRIVACY AND INNOVATION IN PUBLIC SERVICES

Introduction

- Good afternoon, everyone. It's a real pleasure and honour to be here today.

- I am very excited to be engaging with Ontario's public sector about how we can responsibly adopt AI in the public sector to accelerate innovation while also ensuring accountability, transparency, privacy and security.

Innovation and AI

- Innovation and modernization offer opportunities for more efficient, effective and responsive government.

- Many new technologies have the exciting potential to help institutions improve their service delivery to the public.

- AI technologies, in particular, are rapidly transforming our world faster than we know it.

- Ontario's Secretary of the Cabinet, Michelle DiEmanuele, in her annual report to the premier, calls on public servants to embrace flexibility, continuous improvement and sustainability by adopting new technologies, including AI.

- She also reminds public servants of their commitment to serve the people of Ontario with integrity, adding that as Ontario's public sector seizes the opportunities of AI, those efforts should be informed by good governance and protection of the public interest.

- I tend to agree with the latter point.

- Our enthusiasm to adopt AI must be matched by an unwavering commitment to accountability, transparency, privacy, and security, all in the public interest.

Benefits and Risks of AI

- AI technologies are now mainstream, due to the widespread popularity of apps we use every single day, including search engines, facial recognition tagging, real-time translation and geolocation services — even the smart devices in our homes, like our fridges, are fitted up with AI.

- And now, large AI-powered language tools like ChatGPT, have become a household name.

- Apple is now also working on its own competitive AI product, Apple GPT.

- Driven by the promise of improved efficiency, policy makers are increasingly seeing the benefits of AI for delivering services to the public.

Benefits of AI

- For example, earlier this year, the Ontario government requested OntarioMD to conduct an evaluation pilot of AI scribe technology that uses AI to automatically transcribe patient visits with their family physicians — generating medical notes and summaries, and automating follow up reports, such as referrals to specialists.

- The hope is for AI scribes to reduce the significant time that family physicians in Ontario spend on documentation, allowing them to focus more attention on their patients, improving both the quantity and quality of their services.

- A relatively new chatbot or AI assistant called GovAI is also making inroads into Canada's public sector, with nine provinces and 89 municipalities already using it across the country.

- Public institutions can tailor this tool so that administrators can ensure it is working in support of their missions and mandates.

- For instance, OPS's internal AI chatbot, called EVA (Enterprise Virtual Assistant) was developed to answer over 3,000 IT and HR related questions, involving almost 100,000 engagements in 2023-24, most of which took less than a minute each.

- External facing chatbots, like the Driver and Vehicle contact chatbot and the Interactive Voice Response system for social assistance clients are fielding tens, even hundreds of thousands of calls and chats, significantly reducing call and wait times for citizens.

- Other projects are being explored or piloted to accelerate the procurement process, reduce red tape and synthesize public consultation input.

- The Town of Innisfil, Ontario, uses an AI-powered chatbot to assist residents with inquiries and make municipal services more accessible, such as providing multilingual support or assisting individuals with disabilities. Wait times are also reduced, improving overall customer experience.

- AI tools are now table stakes for public institutions mounting a first level of automated cybersecurity defence against malicious attacks, running 24-7. By detecting anomalous behavioral patterns or other suspicious activity that may be consistent with potential threat actors, and triaging them for further human analysis, these AI tools can help avert crippling impacts on essential services and critical infrastructure.

Risks of AI

- While AI holds great potential to improve government services, we must also address its risks and implications, especially given its reliance on huge volumes of personal information.

- It's well documented that AI can carry significant risks of discrimination based on biased datasets on which algorithms are trained.

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

- In a recent study, researchers asked ChatGPT to explain how it ranked resumes. In one case, it claimed that a resume referencing an autism leadership award demonstrated less of a leadership role compared to others, implying that people with autism don't make as good leaders.

- Another study found that AI systems used by banks to assess creditworthiness were biased by using proxies, such as postal codes and other sociodemographic data, leading to lower credit scores for those from marginalized communities.

- Recently, the police department in Frederick, Colorado, claimed it was the first law enforcement agency in the world to use a chatbot that summarizes police officers' interactions with individuals from the audio on their body-worn cameras.

- However, given the evidentiary weight of police reports in legal proceedings and the importance of police accountability in our society, civil rights experts have warned that chatbots are known to make mistakes, such as confusing jokes or mistaking words. This can result in wrongful arrests, reinforcing bias, or even covering up potential abuse.

- Legal experts have suggested that any court proceeding based on information from such chatbots should include information about how the models were trained, what information was provided, what information was excluded, and how the models were tested.

Need for Guardrails and Legislation

- As the public sector in Ontario moves towards adopting AI technology, we must establish legal and ethical guardrails around its development and deployment.

- We are seeing many global efforts to establish proper regulatory frameworks around AI safety and security in both public and private sectors.

- In Europe, the EU *AI Act*, which came into effect on August 1, of this year, establishes legal obligations for providers and users depending on level of risk. It also prohibits certain AI practices altogether, such as behavioural manipulation, the scraping of facial images from the internet, the use of social scoring, and the biometric categorization of individuals or groups.

- In 2023, the U.S. White House established a Blueprint for an AI Bill of Rights, calling for safe and effective systems, protection against algorithmic discrimination, data privacy, notice and explanation, and the right to opt out in favor of human alternatives and timely reconsideration of automated decisions.

- Several states have followed suit.

- In California, the *AI Accountability Act*, prohibits the state from contracting AI services unless the vendor meets certain standards, and it requires notifying the public when they are interacting with AI.

- Colorado's *AI Act* 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. The bill passed second reading in the House of Commons and is still under review by the Standing Committee on Industry and Technology.

- Here at home in Ontario, the government has tabled Bill 194, which seeks to regulate the use of AI by public sector organizations. The bill provides for regulation making authority in respect of transparency, accountability, risk management, technical standards, and oversight, as well as certain prohibited uses.

- While this is a promising first step, my office has submitted recommendations to the legislative assembly on how the bill could be improved.

- Among other things, we believe the law should enshrine clear statutory guardrails around the use of AI technologies and not leave such fundamental matters to regulation.

- The Ontario Human Rights Commission, the Law Commission of Ontario, and academic experts have all made similar recommendations.

- In our submission, we went on to recommend a set of principles that should guide public sector organizations as they develop or deploy AI systems.

    - Before AI technologies are adopted, they should have to meet independent testing standards to ensure they are ***valid and reliable*** and work as intended.

    - AI systems should be ***safe*** and designed for the physical and mental well being of people, our economic security, and the good of our environment, and they should continue to be monitored throughout their lifespan.

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

    - ***Transparent*** policies and practices should inform people when they are interacting with AI and when decisions have been made about them using AI.

- An ***accountable*** governance structure is necessary so that individuals can challenge the accuracy of decisions made about them and seek recourse. Public sector organizations should also be subject to review by an independent oversight body with authority to enforce these principles.

- Most importantly, 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.

- Another important issue I want to address here is the role of humans, or human-in-the-loop oversight.

- AI systems should be designed so that those using these systems can understand and explain the data, criteria, and reasoning they exercise in producing an output.

- Of course, many of these systems make mistakes or are not yet wholly transparent.

- Whenever an output could have a significant effect on Ontarians' lives, it's essential that organizations carefully determine how transparent and knowable these AI-generated outputs truly are and take a risk-based approach in assessing whether an AI system should even be adopted.

- And when a system is being adopted, organizations must ensure that they are providing adequate time, resources, information, and capacity needed for well-trained humans to effectively review automated outputs, predictions or decisions.

AI in Ontario and IPC's Ongoing Involvement in AI

- 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.

- That being said, at the IPC, we're not just waiting to see the outcome of Bill 194.

- We've been actively working and advocating for the adoption of guardrails around the responsible use of AI for a couple of years now.

- 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, addressing safety, privacy, accountability, transparency, and human rights.

- My office also took up the cause on a national level by joining our federal, provincial, and territorial counterparts to release *Principles for Responsible, Trustworthy, and Privacy-Protective Generative AI Technologies*.

- Then we went international by co-sponsoring two 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*.

- My office has also been active in getting Ontarians from all walks of life involved 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 also 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.

- We have addressed AI issues in some of 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.

- Most recently, my office issued a revised code of procedure for processing FOI appeals under FIPPA and MFIPPA that came into effect yesterday.

- 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 appeal procedures, improving their timeliness, and making 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.

- Also, 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 and certify in writing to the IPC that they have completed such review.

Third Party Outsourcing Guidance

- Before I conclude, I'd like to tell you about another recent guidance document we issued, called *Privacy and Access in Public Sector Contracting with Third Party Service Providers*.

- This is a timely publication given that public institutions often engage third-party vendors to help them deliver services to the public and these arrangements often involve extensive processing personal information, increasingly through the use of AI tools that are difficult to see and understand.

- No matter the arrangement with a third party vendor, public institutions must ensure full compliance with Ontario's access and privacy laws.

- As we like to say at the IPC, "You can outsource services, but you can't outsource accountability."

- This guidance provides practical advice to identify access and privacy considerations when contracting with third-party service providers.

- It includes best practices and recommendations to support proper due diligence and accountability throughout the procurement process — from planning, to tendering, to vendor selection, to contracting, right up to and including, close-out of the agreement.

- I also encourage you to visit the IPC booth here today and see all these and other great resources we have to offer.

Conclusion

- To conclude, I'd like to point out that my office is not just here to call out non-compliance with Ontario's access and privacy laws — we are also a resource.

- Our mandate includes offering comment on the access and privacy implications of proposed government programs — we can help

mitigate potential risks before they become real-world problems.

- If your institution is exploring the novel use of AI technologies, I encourage you to reach out to my office.

- You can learn more about our policy consultation process on our website under Guidance for Organizations.

- While the temptation to rush AI adoption will be strong and at times overwhelming, I encourage you to take the time needed to ask the tough questions from the start.

- Do the upfront work — privacy impact assessments and/ or algorithmic impact assessments — to assess the potential legal, ethical and social impacts of an AI tool, and amend your plan as needed to mitigate the risks involved.

- I promise you it will be time well spent. Otherwise, as J.R.R. Tolkien once said, “shortcuts make for long delays.”

- Conversely, many institutions will hesitate to innovate because they are risk averse, they’re scared of making mistakes, and don’t know what the rules are — a phenomenon that’s called reticence risk.

- As Brene Brown has succinctly put it, “So many leaders fail to realize that without vulnerability there is no creativity or innovation … there is nothing more uncertain than the creative process, and there is absolutely no innovation without failure. Period.”

- There will be failure, but it’s important to fail fast, fail small and fail forward. Which is why that up front planning work is absolutely critical as is an agile, continuous learning approach.

- In my blog, *Privacy and Humanity on the Brink*, I took a more existential look at AI. I wrote that life as we know it will never be the same.

- Through the rapid adoption of AI and other new technologies, we are creating a legacy we have yet to fully understand.

- One that will challenge our right to privacy like never before and even our fundamental sense of human agency for ourselves and for generations to come.

- But with change comes opportunity.

- As public servants committed to serving Ontarians with integrity, you still have the time and the ability to shape the future, ensuring that AI technologies are safe, transparent, accountable, and ethically responsible.

- As former U.S. Secretary of State, Madeleine Albright, once said, “What people have the capacity to choose, they have the ability to change.”

- So, let’s choose responsible AI — and let’s be thoughtful and deliberate about how we need to change our laws and policy frameworks to ensure we have clearly articulated guardrails that will help preserve public trust in the exciting benefits that innovation and technology have to offer.

- Thank you.