LAW COMMISSION OF ONTARIO
COMMISSION DU DROIT DE L'ONTARIO

Dubi Kanengisser, PhD
Senior Advisor, Strategic Analysis and Governance
Toronto Police Services Board
40 College Street
Toronto, ON M5G 2J3
dubi.kanengisser@tpsb.ca

September 22, 2021

Dear Dr. Kanengisser,

The Law Commission of Ontario (LCO) welcomes the opportunity to provide comments on an early draft of the Toronto Police Service Board's *Use of Artificial Intelligence Technologies Policy* ("the draft AI policy"). The LCO commends the Toronto Police Service (TPS) and Toronto Police Services Board (TPSB) for taking proactive steps to create a policy to govern police use of AI systems.

About the LCO

By way of background, the LCO is Ontario's leading law reform agency. The LCO provides independent, balanced and authoritative advice on complex and important legal policy issues. Through this work, the LCO promotes access to justice, evidence-based law reform and public debate.

The LCO has unparalleled experience analyzing AI, regulation and the public sector. Recent LCO reports addressing these issues include:

- *Regulating AI: Critical Issues and Choices* (April 2021)
- *Legal Issues and Government AI Development* (March 2021)
- *The Rise and Fall of Algorithms in the American Justice System: Lessons for Canada* (October 2020)

Many of the suggestions in this letter are drawn from these reports. This work is part of the LCO's ongoing AI, Automated Decision-Making and the Justice System project.

Introduction

The use of AI, algorithms and automated-decision making are expanding in police services across the world. This expansion raises new and crucial questions about equality, bias, access to justice, due process, and fundamental rights. AI and algorithms offer many benefits, including the potential to provide consistent, "evidence-based" and efficient predictions. However, the introduction of AI systems into police service also raises the potential for police to unintentionally cause harm to vulnerable, racialized and marginalized populations. Experience demonstrates the risk of adopting unproven and under-evaluated technologies too quickly to address long-standing, complex and structural problems.

The LCO and many other organizations have identified significant criticisms of AI and algorithmic tools. Moreover, AI or algorithmic tools used in policing have been the subject of extensive review and commentary, particularly in the United States, where the use of these tools is widespread. Police use of biometric tools (including facial recognition) and predictive policing systems have been criticized for being racist, opaque, and illegal.[1] In Canada, the most significant analysis of AI and algorithmic policing technology to date is *To Surveil and Predict*, a 2020 report written by The Citizen Lab and the University of Toronto's International Human Rights Program.[2]

Comments on the Draft AI Policy

The LCO offers some preliminary comments and observations about the draft policy based on our experience with public sector AI systems and regulation. These comments are based on what we understand is an early draft of the policy. The LCO has not made specific comments on the language or the draft policy for two reasons: First, we understand an updated text will be circulated publicly shortly. Second, as will be discussed below, we believe many important provisions of the policy should be discussed and debated publicly with a wide range of stakeholders.

At the outset, the LCO compliments the TPSB and TPS for addressing AI regulation and governance issues proactively and publicly. This is an important step that addresses some of the concerns that have arisen in other jurisdictions when police services have deployed this technology without public knowledge or discussion.[3] The TPSB and TPS have given themselves a great opportunity to avoid many of earlier mistakes in this field.

The LCO also wants to acknowledge the challenge and complexity of AI regulation, particularly in policing contexts. For example, AI systems are evolving rapidly, and applications can range from low-risk administrative tools to high-risk investigative systems. AI regulation is also difficult because resources may be limited. And perhaps most importantly, AI *policing* regulation is a particularly complex area that

---

[1] For example, see the work of Professor Andrew Ferguson including *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement* and *Policing Predictive Policing* and related analysis by the NYU Policing Project, among others.

[2] Kate Robertson, Cynthia Khoo, and Yolanda Song, To Surveil and Predict: A Human Rights Analysis of Algorithmic Policing in Canada (2020), Citizen Lab and International Human Rights Program, University of Toronto.

[3] The LCO's *The Rise and Fall of Algorithms in American Criminal Justice* report discusses the introduction of algorithmic risk assessments in American bail proceedings as a cautionary example. The LCO can provide many additional examples.

necessarily combines AI and algorithmic technical issues, policing operational objectives, legal rights, and the need for public accountability.

Fortunately, much has been learned about how to design, develop, implement, and evaluate public sector and policing AI and algorithmic systems. Many organizations, technologists and academics have developed best practices and/or legal regimes to better govern the use of these tools. The LCO and other organizations discuss many of these practices and recommendations in the above referenced reports.

Before beginning, the LCO wants to highlight a key finding of our AI-related work: There is a strong consensus amongst policymakers, operational staff, academics, technologists, legal professionals, civil society organizations, and community representatives that AI regulation must be a deliberative and multidisciplinary. As a result, the LCO's first recommendation is that the TPSB and TPS commit to establishing an appropriate process and timeframe for this work. There are many examples of AI or algorithmic technology being introduced too quickly, or without appropriate governance measures.

In the LCO's view, the objective of the TPSB AI policy should be to establish a transparent and comprehensive framework governing the consideration, development, deployment and evaluation of AI and algorithmic tools by the Toronto Police Service.

What follows below are the LCO's comments on the issues the TPS AI policy should address. As noted above, at this point the LCO will largely defer on the specific language or text that should be included in the policy.

In our view, the TPS AI policy should include or address the following elements and issues:

- **Objectives, Purpose and Benefit of AI Tools.** A clear first principle of any AI governance policy, be it in policing or any other area of public administration, is the need for the agency or organization sponsoring or considering an AI tool to clearly articulate the objective, purpose and claimed benefits of that tool. What legislative, policy or operational objectives will the tool help fulfill? Why is the tool needed? How will the tool improve on current practices?

- **Public Engagement** – Public participation and engagement is crucial for the development of trustworthy and reliable AI. Public engagement is also crucial to ensure human rights and due process are protected. As a result, public engagement should be enshrined in the policy as a key principle. Public participation should occur during the design, deployment, and evaluation of AI systems. In the policing context, experts in policing, human rights, privacy, criminal law, information technology, data science and community representatives should be involved. Releasing a draft AI policy for public comment is an important first step.

- **Disclosure/Transparency.** Current leading practices in AI regulation (the Ontario's Trustworthy AI Framework "no AI in Secret" principle, the federal government's Directive on Automated Decision-making, the proposed European Commission AI rules, and the LCO's *Regulating AI* report) emphasize the need for comprehensive public disclosure/transparency of AI systems to promote public understanding and accountability of AI systems, particularly for high-risk systems/activities.

The draft AI policy reviewed by the LCO includes commitments to public disclosure/ transparency. These commitments are important, but the final policy should be more explicit and detailed to better ensure "Trustworthy AI." In this regard, it is important to note that disclosure/transparency has two dimensions: 1) *How* to disclose the existence of an AI system, and 2) *What* to disclose about a system.

Regarding the first question, the LCO strongly recommends the final TPS AI policy include a commitment to develop a Toronto Police Service "AI Register." The objective, content and range of AI Registers is discussed extensively in the LCO's *Regulating AI* report.[4]

The second question is more complicated, especially in policing. As a starting point, the LCO strongly recommends the disclosure of an AI Impact Assessment for higher-risk policing AI systems. (AI Impact Assessments are discussed below.) The LCO believes there are compelling reasons to require very broad disclosure of many policing AI systems. For example, the history of racialized data in criminal justice confirms the need for extensive disclosure of the data used to train policing AI systems. At the same time, there may be legitimate operational reasons to keep aspects of a policing AI system confidential.

Experience in other jurisdictions proves that policing AI disclosure issues are both controversial and consequential. As a result, the decision about what kinds of information to disclose (or to *not* disclose) must be deliberate, public and multidisciplinary.

- **Risk Categories.** The draft AI policy reviewed by the LCO includes four risk categories, ranging from "extreme risk" to "low risk." Creating a sliding risk scale is a best practice in AI regulation, but the rules and details distinguishing between risks are important and often controversial.

  To promote public accountability and "Trustworthy AI", the draft TPS AI policy must explicitly address the following questions:

  - What risk categories does the policy include?
  - What criteria are used to assess the risk level of AI systems or applications?
  - Who decides which risk category is appropriate for specific AI systems or applications?
  - What are the regulatory consequences of each risk level?

  The accountability, credibility, and effectiveness of the TPS policy will depend on these questions being addressed comprehensively and transparently. Risk identification has proven to be one of the most controversial aspects of AI regulation. In this respect, the draft policy reviewed by the LCO takes important steps.

  Risk categorization is a complex topic, involving technical, legal, policy and political considerations. Care must be taken to ensure broad participation in these discussions. A simple example demonstrates this point: The draft policy prohibits AI systems that result in "mass surveillance." This is an appropriate prohibition, but what is its scope? Which technologies

[4] *Regulating AI* at 29-33.

would be included? Are there exceptions to this rule?[5] The credibility and effectiveness of the TPS policy will depend on having a thoughtful public discussion about "red line" technologies that have such pervasive and systemic risks that they should be *a priori* prohibited.

- **Impact Assessments.** From an operational perspective, the TPS and TPSB will be challenged to organize the risk analysis of so many disparate kinds of AI systems. This challenge is heightened by the high level of public scrutiny and accountability demanded of police services. Fortunately, there are tools available to address this task. AI Impact Assessments are essentially a form of evaluation framework designed to comprehensively identify and assess the *objectives, details, benefits, risks, and mitigation strategies* of an AI system. Impact Assessments take many forms and have varying levels of detail. The LCO discusses Impact Assessments in detail in our *Regulating AI* report.[6] Based on that analysis, the LCO strongly recommends the TPSB AI policy include a commitment to develop a robust Impact Assessment for policing AI systems. The LCO further recommends that TPS Impact Assessments be mandatory, detailed, and transparent.

- **Harm Mitigation.** A fundamental component of any AI policy is to explicitly identify the harm mitigation measures attendant to each level of risk. The draft policy reviewed by the LCO was silent on this issue. The federal government's *Directive on Automated Decision-making* is a good example of a policy that incorporates harm mitigation strategies into the policy itself. Importantly, mitigation strategies should acknowledge the range of AI systems and be tailored to risk: High risk/impact systems should require much greater scrutiny and oversight than systems with minimal or no risk/impact.

  Like risk categorization, AI harm mitigation raises complex technical, legal, policy and political considerations. As a result, these issues should be discussed in a deliberate, public and multidisciplinary process.

- **Data Bias, Accuracy, Reliability and Validity.** Data bias, accuracy, reliability, and validity are significant and ongoing concerns with all AI systems, especially in policing and justice applications. The LCO report, *The Rise and Fall of Algorithms in American Criminal Justice: Lessons for Canada*, discusses data discrimination issues extensively.[7] In the United States, many algorithmic and AI tools used in criminal justice (notably predictive policing tools and algorithmic risk assessments used in bail and sentencing) have been pulled back due to unresolved concerns about racialized policing data. These issues will no doubt arise in this country when a Canadian police service proposes using racialized, historic data to train a high-risk AI system.

---

[5] For example, Article 5 of the European Commission's recently proposed AI rules would prohibit "real-time" remote biometric identification systems in publicly accessible spaces for the purpose of law enforcement subject to several exceptions, including

- Where strictly necessary to search for a missing child,
- To prevent a specific and imminent terrorist threat, or
- To detect, locate, identify or prosecute a perpetrator or suspect of a serious criminal offence.

See generally, European Commission, Proposal for A Regulation Of The European Parliament And Of The Council Laying Down Harmonised Rules On Artificial Intelligence, (2021).,

[6] *Regulating AI* at 33-36.

[7] *The Rise and Fall of Algorithms in American Criminal Justice* at 20-25, 26-27 and 32-33.

The TPS and TPSB should publicly identify and discuss if, or how, the TPS intends to use (and mitigate, if possible) historic and biased datasets to train TPS AI systems. The TPSB and TPS should also explicitly commit to adopt best practices in data collection, disclosure, monitoring and evaluation.

- **Legal Compliance, Human Rights, Due Process and Evidence.** The TPSB's AI policy should explicitly state that the TPS's use of AI systems will comply with the *Charter,* Ontario Human Rights Code, and privacy legislation. These are important public commitments in addition to being legal requirements.

  The TPSB's AI policy should also acknowledge the potential use of some policing AI systems in criminal investigations and prosecutions. In the event AI-generated evidence is sought to be tendered at trial, the use of that evidence will have to meet criminal law due process and evidential requirements. The extent of these obligations in individual cases may be difficult to discern. In other circumstances, policing AI tools may generate information that is not used in individual investigations or prosecutions.

  Determining the extent of due process and evidential requirements raises complex legal, operational, and technical issues that require considerable analysis. As a result, the LCO recommends the TPS AI policy acknowledge the potential use of AI tools in criminal prosecutions and the need to work with stakeholders to address these questions appropriately.

- **Independent Monitoring and Oversight.** Experience suggests that self-governance is not sufficient oversight for an AI system that affects individual rights or has the potential to cause harm to vulnerable populations. As a result, best practices in AI regulation include reviews, audits, and validation of higher-risk AI systems prior to their deployment *and* regularly during their operation by some form of independent monitor. Monitoring and oversight are important steps to promote trustworthiness and system effectiveness.[8]

The LCO appreciates the opportunity to participate in this important consultation. The LCO is committed to working with the TPSB, TPS and others to develop a thoughtful AI TPS policy.

Sincerely,

Nye Thomas
Executive Director
Law Commission of Ontario

Cc LCO Board
Susie Lindsay, LCO Policy Counsel, Lead Civil AI Project
Ryan Fritsch, LCO Policy Counsel, Lead Criminal AI Project

[8] See generally *Regulating AI* at 41 and 45-47 and *The Rise and Fall of Algorithms in American Criminal Justice* at 32.