PUBLIC INTEREST ADVOCACY CENTRE
LE CENTRE POUR LA DÉFENSE DE L'INTÉRÊT PUBLIC

285 McLeod Street, Suite 200, Ottawa, ON K2P 1A1

December 15, 2021

Dubi Kanengisser
c/o Toronto Police Services Board
40 College St.
Toronto, ON, M5G 2J3

Dear Mr. Kanengisser:

Re: Public Consultation – Use of New Artificial Intelligence Technologies Policy
Comments of the Public Interest Advocacy Centre (PIAC)

Please find attached the comments of the Public Interest Advocacy Centre on the above-noted consultation.

We trust the enclosed to be in order, but if you require anything further, please contact me at 613-562-4002 ext.122 or by email at: ysai@piac.ca

Yours truly,

Yuka Sai
Digitally signed by Yuka Sai
Date: 2021.12.15 16:33:15 -05'00'

Yuka Sai
Staff Lawyer

SUBMISSION OF THE PUBLIC INTEREST ADVOCACY CENTRE ("PIAC")

TO:

Public Consultation of the Toronto Police Services Board:

"Use of New Artificial Intelligence Technologies Policy"

15 December 2021

Contents

I. Introduction

1. The Public Interest Advocacy Centre (PIAC) is pleased to provide comments to the Toronto Police Services Board (the "Board") public consultation on the TPSB's Use of New Artificial Intelligence Technologies Policy (the "Policy") that governs the use of new AI technologies in a manner that is fair, equitable, and does not breach privacy rights. The TPSB is developing the draft Policy to help the Toronto Police Service (the "Service") minimize privacy risks and ensure transparency while allowing the Service to contribute to positive community safety outcomes through effective and efficient investigations using AI technologies. As PIAC has long advocated for stronger privacy protections to protect individual interests in the use of new technologies, we welcome the opportunity to comment in this consultation.

2. PIAC commends the Board for taking the initiative to not only develop its own AI policy, but also to seek input from stakeholders and the public. PIAC understands that the Board has already received and incorporated into the draft Policy the preliminary feedback from important stakeholders like the Ontario Human Rights Commission (OHRC) and the Law Commission of Ontario (LCO). However, PIAC believes that the draft Policy as published must be further revised and improved. In recommending specific revisions to the Policy, especially those applying to higher risk AI technologies, PIAC looked to the European Commission's 2021 proposal for a Regulation laying down rules on artificial intelligence.

3. PIAC's comments below detail various issues we have identified in the Policy pertaining to risk categorization, public accountability, transparency, independent review, monitoring procedures, and public engagement. However, ultimately an internal policy is only as strong as the underlying regulatory framework governing police use of AI technology, which in our view, is lacking. As privacy laws will soon change, the Board and the Service should proceed with caution until reforms are finalized.

II. Risk categorization and criteria must be more comprehensive

4. With regard to the proposed risk categories in the Policy, further revisions and additions are necessary. PIAC advises that some criteria under the High Risk category should be considered Extreme Risk, and that others under the Moderate and Low risk categories should be recategorized as High Risk. Furthermore, PIAC recommends that the High Risk criteria be more aligned with what the European Commission deems to be "high-risk AI systems" particularly in the context of law enforcement, in its recent 2021 proposal to introduce harmonized regulations for AI in the EU.[1]

5. The following criteria should be moved from the High Risk to the Extreme Risk category:

[1] Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS, https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1623335154975&uri=CELEX%3A52021PC0206

- **The criteria "where training or transactional data is known to be of poor quality, carry bias, or where the quality of such data is unknown"** should be placed in the Extreme Risk category and therefore AI technologies using such data should not be adopted. Data that are known to be biased or of poor/unknown quality poses too high a risk of poor and biased outcomes. It does not make sense that the Policy prohibits the adoption of a biased or flawed AI system known or is likely to cause harm or have impact on individual rights, while training or transactional data, which operationalizes the AI system, having the same bias or flaws does not merit the same prohibition. On the matter of bias, the OHRC in their preliminary comments recommended that "any AI system known to carry bias which can cause potential harm or have an impact on an individual's rights, despite the use of mitigation techniques" should not be used.[2] The LCO noted that issues arise where "a Canadian police service proposes using racialized, historic data to train a high-risk AI system," further noting that AI tools have been pulled back in the US due to concerns about racialized data.[3] The LCO seems to therefore recommend a moratorium on the use of AI tools that use biased data.

- **The criteria "where a system is not fully explainable in its behaviour"** should be placed in the Extreme Risk category. PIAC submits that this is a clear bright line for prohibiting the use of such AI technologies. If the Service cannot understand or explain how an AI system works, they cannot effectively assess impact, propose harm mitigation strategies, or consider possible unintended consequences. PIAC points out that the initial reporting stage under the Policy requires that the Chief of Police report to the Board about "how the AI technology operates;" if the Chief cannot explain how the AI system works, the Board cannot form any legitimate basis for approval. The OHRC also recommends that "any AI system that cannot be fully explainable in its behaviour is an *Extreme Risk Technology*, and therefore should not be used."[4] The European Commission's 2021 proposal is also clear that even a high-risk AI system must be usable and comprehensible in a way where instructions, capabilities, limitations, known or foreseeable circumstances are clear to users.[5]

6. The following criteria should be moved from the Moderate or Low Risk category to the High Risk category:

- **The criteria "where the 'human-in-the-loop' may have difficulty identifying bias or other decision failures of the AI"** should placed in the High Risk category. PIAC submits that where human oversight may have difficulty identifying bias, this merits careful impact analyses and extensive harm mitigation measures (i.e. training, contingency planning, etc) that are more commensurate with the High Risk characterization. If the AI system is so complex or unknown that potential harm cannot be effectively identified and addressed by human oversight, then the technology should not be used.

[2] OHRC letter, page 4.
[3] LCO letter, page 5.
[4] OHRC letter, page 4.
[5] *Supra* note 1, at Article 13.

- **The criteria "where the process involved suggests an allocation of resources"** should be expanded upon and placed in the High Risk category. The provision as written in the current draft Policy is too vague, as allocations of resources within the Service may affect what programs are funded or unfunded, which may cause potential harm to marginalized and vulnerable groups. PIAC therefore recommends that the provision be moved into the High Risk category and expanded to state: "where the process involved suggests an allocation of resources or change in operations in a way that could cause potential harm or have an impact on an individual's rights."

- **AI technology that "assists Members in identifying, categorizing, prioritizing or otherwise making decisions pertaining to members of the public"** should be moved from Low Risk to at least the Moderate Risk category. This criterion is far too broad. There are potentially many ways bias can seep into processes that identify, categorize, prioritize, or otherwise make decisions about individuals. Therefore, the risk categorization for these activities should at least be Moderate, with the additional requirement that the Chief of Police must explain to the Board and to the public how the process does not carry risk of potential harm or impact on individuals.

7. The following criteria should be added to the High Risk category. These provisions are adopted from Annex III, Section 6 of the European Commission's 2021 proposal, which specifically details high-risk AI processes that may be used by law enforcement:

- Where the process involved is intended for making individual risk assessments about natural persons in order to assess the risk of a natural person for offending or reoffending or the risk for potential victims of criminal offenses;

- Where the process involved is intended to be used as polygraphs and similar tools or to detect the emotional state of a natural person;

- Where the process involved is intended to be used to detect deep fakes;

- Where the process involved is intended for the evaluation of the reliability of evidence in the course of investigation or prosecution of criminal offences;

- Where the process involved is intended to be used to profile natural persons in order to predict the occurrence or reoccurrence of an actual or potential criminal offence, or in the course of detection, investigation or prosecution of criminal offences;

- Where the process is used for crime analytics regarding natural persons, allowing Members to search complex related and unrelated large data sets available in different data sources or in different data formats in order to identify unknown patterns or discover hidden relationship in the data.

8. With regard to the Low Risk and Minimal Risk categories, the distinction is unclear. PIAC submits that, for clarity, the categories should be limited to Extreme, High, Moderate, and Low Risk, with

the lowest end of the risk spectrum pertaining to administrative and operational tools that are unlikely to affect any decision-making processes about individuals. Therefore, the Low Risk Technologies category should encompass any AI technology that both 1) Does not fall under the categories of Extreme High Risk, High Risk, or Moderate Risk, and 2) Assists Members as administrative or operational tools that are unlikely to cause potential harm or have an impact on individual rights.

9. PIAC submits that more concrete, specific descriptions of the types of AI systems and technologies applicable to each level of risk are critical to protect privacy rights and individual freedoms. Less discretion on the part of the Service is preferable to potentially too much discretion in using AI technologies. This is not a new notion. In the context of automated facial recognition technology, the South Wales Police (SWP) in the United Kingdom ran into legal trouble in their use of a privacy-invasive mass surveillance technology while relying on laws and policies of general application.

10. In that case, *R (on the application of Bridges) v Chief Constable of South Wales Police* ("Bridges"), the European Court of Human Rights (ECHR) determined that the laws governing police use of automated facial recognition technology did not sufficiently circumscribe police discretion regarding where and on whom to use the technology. Though the case was specific to a particular type of facial recognition technology that was used for mass surveillance, the lesson is applicable to all privacy-invasive technologies used by police that potentially impact individual rights and freedoms.

11. As the Board itself has stated, "[n]o current legislation fully regulates the use of AI technologies, and the Province has not yet developed comprehensive guidelines for the use of such technologies in policing." PIAC is not opposed to law enforcement agencies developing their own internal policies on the use of AI technologies, but in the absence of AI policing regulations, internal polices should be as precise and comprehensive as possible about the limits of police discretion. The draft Policy, in PIAC's view, fails to do so.

12. Furthermore, PIAC notes that there is no explicit obligation in the Policy to verify, both prior to use and on an ongoing basis, that the training and transactional data used with an AI system is free of bias, regardless of whether the Service or the private sector partner supplies the data. This requirement can be written into the Policy under Section 5, and under the Continuous Review sections. In Bridges, the SWP failed to take reasonable steps to satisfy themselves, either directly or by way of independent verification, that the facial recognition software does not have an unacceptable level of bias.

III. Detail the harm mitigation measures and key performance indicators associated with each level of risk

13. The draft Policy states that the Chief of Police will propose "indicators" to be tracked until at least 12 months after full deployment to determine whether the AI tech is achieving the intended goal

and whether its deployment has had any unintended consequences. Though PIAC acknowledges that the exact indicators pertaining to an AI technology or system may vary on a case-by-case basis, the Policy is silent on any possible indicators, and more importantly, what indicators are commensurate to each level of risk.

14. The Policy, under section 1(e), also requires that the Chief establish, in consultation with experts and stakeholders, "the harm mitigation measures required for each level of risk (e.g., training, contingency planning)." However, the current draft Policy provides no elaboration on the appropriate harm mitigation measures attached to each level of risk, nor any commitment to enshrine these measures in the Policy. The LCO noted in their preliminary comments that "the draft policy reviewed by the LCO was silent on this issue," and further that "[a] fundamental component of any AI policy is to explicitly identify the harm mitigation measures attendant to each level of risk."[6] Generally, PIAC submits that high risk technologies require enhanced mitigation measures, scrutiny, and oversight, which merits at least an explicit description of the indicators and measures tailored to all high risk technologies.

15. The LCO cited the federal government's Directive on Automated Decision-making as a good example of a policy that incorporates specific harm mitigation strategies into the policy itself. For example, the Directive requires that an appropriate qualified expert be consulted to review the Automated Decision System, with more stringent peer review standards the greater the potential impact on individuals and communities.[7] For the highest risk level, the Directive also requires re-occurring training courses and a means to verify training completion.[8] The Board's draft Policy should similarly detail the minimum harm mitigation measures commensurate with the level of risk.

IV. Incorporate public accountability at all stages of development, deployment, and monitoring of AI systems

16. Under the draft Policy, reporting on the reviews and assessments of new AI technologies seems to flow primarily between the Chief of Police and the Board. Although the Guiding Principles of the Policy provide that "[t]o the greatest degree possible, the Board must conduct such reviews in public," the actual Policy explicitly provides for limited public transparency at only a few stages. Although the Policy provides that the general Section 1 procedures are made available to the public on the Service's website, there are no provisions that clearly indicates to what degree the public is informed about the details of the Chief's report to the Board under Section 5.

[6] LCO letter, page 5.
[7] Government of Canada's Directive on Automated Decision-Making, Appendix C, online: https://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=32592
[8] *Ibid*.

17. Only Section 8 rather nebulously suggests that the Chief "[w]ill develop and implement a public engagement strategy, commensurate with the risk level assigned to the new AI technology, to transparently inform the public of the use of the new AI technology that collects data about the members of the public or assists Service Members in identifying, categorizing or otherwise making decisions pertaining to members of the public, prior to its deployment." The provision does not state what the risk threshold is for public transparency, and the amount of information disclosed depending on the risk level. The Chief of Police therefore seems to have a great deal of discretion about the "public engagement strategy."

18. Section 15 of the Policy lists types of information to be posted on the Service's website regarding AI technologies deemed to be of High, Medium, or Low Risk. However, the required info for High and Medium Risk tech is extremely limited in comparison to what the Chief of Police must provide the Board in order to approve new AI technologies. For example, the required public-facing details do not include why the level of risk was ascribed, and the rationale for use despite the associated risks. The public should also be informed about the applicable legislative authority, how the tech operates, the results of Privacy Impact Assessments, possible consequences, risk mitigation plans etc. The ability of the public to hold the Service accountable is directly proportionate to the level of public transparency.

19. Preliminary comments from the LCO and OHRC also seems to place emphasis on public disclosure and transparency. The LCO, in their letter, highlighted 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."[9] While the LCO acknowledges there may be operational reasons to keep some aspects confidential, they still strongly recommend the disclosure of an AI Impact Assessment for higher risk policing AI systems. The OHRC recommends that all disclosures be provided in plain language, including any documents used in the risk-assessment process, and the risk assessment reports themselves. OHRC also supports additional disclosures like "when that information is disposed of and what actions are taken to minimize discriminatory effects and outcomes."[10]

20. Best practices for public transparency can also be drawn from the European Commission's 2021 proposal. For example, where appropriate, natural persons who directly interact with AI systems should be informed at point of contact that they are interacting with AI, unless specific exceptions apply relating to the prevention of crime.[11] Other transparency obligations to incorporate include informing individuals when AI systems are used to recognize emotions, conduct biometric categorisation, or to generate deep fakes.[12]

[9] LCO letter, page 3.
[10] OHRC letter, page 4.
[11] *Supra* note 1, at Article 52(1).
[12] *Supra* note 1, at Article 52(2)-(3).

V. Engage independent monitoring bodies to ensure objective oversight

21. PIAC submits that there is a marked lack of engagement in the Policy with any independent reviewing or monitoring bodies that can add an important dimension of accountability. For one thing, there is no provision to seek pre-approval or even to provide notice to the Information and Privacy Commissioner of Ontario (IPC), only that consultations may be conducted with the IPC, the Ministry of the Attorney General, and various other stakeholders prior to use "as appropriate in light of the potential risks posed by the contemplated technology." Again, the Chief of Police is given a great deal of discretion about the risk threshold for prior consultation. PIAC submits that there must be a strict requirement to consult and seek feedback from the IPC for High Risk technologies, and give notice to the IPC of all new AI technology to be used by the Service, including Low Risk applications.

22. PIAC is certainly not alone in recommending independent monitoring of the Service's AI technology use. The LCO stated that "[e]xperience 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 population."[13] Accordingly, the LCO suggests that "best practices include reviews, audits, and validation of higher-risk AI systems prior to their deployment and regularly during their operation by some by some form of independent monitor." These practices are largely absent from the draft Policy.

23. An independent monitoring body, free of any potential motivations to downplay or cover up bias and discrimination issues, is essential to keep the Service accountable, and appropriately limit the Board's discretion. PIAC therefore submits that the Policy should be revised to involve independent oversight throughout the lifespan of High Risk AI technologies. As an example, the OHRC proposed an oversight mechanism that engages an independent monitoring body to audit all AI systems for bias and discrimination every 12 months, report on and address systemic issues, hold public reviews of the AI Policy and key performance indicators every 12 months.[14]

24. Also absent from the Policy is a provision for reporting incidents/malfunctions to authorities, as in, the Policy contains no obligation or commitment to report any issues with the AI technology to the appropriate bodies, like the IPC, OHRC or the Ontario Civilian Police Commission, or even to the public. Also problematic is the absence of any commitments to publicly disclose the results of the 12-month deployment reports under Section 11, which describes critical info like what concerns have been raised by the public, how the Chief of Police has addressed those concerns, and whether the use of the tech will continue. PIAC submits that the public is entitled to ongoing disclosure, rather than be left to assume from silence that nothing of concern is occurring with the AI technologies the Service has deployed. There is also no requirement for the Board to notify

[13] LCO letter, page 6.
[14] OHRC letter, page 5.

the public when these reports reveal potential or legitimate issues, which further grants a large degree of discretion to the Board.

25. PIAC also notes that the Policy does not require the Board to inform the public of the termination of any AI tech in use prior to the Policy due to the tech being deemed to be "Extreme Risk." PIAC submits that the public be notified that an AI system was in use – likely without the public's knowledge – and then decommissioned due to factors serious enough to likely, or actually, impact individuals' rights and freedoms. Neither the Service nor the Board should be allowed to, in effect, quietly sweep past practices under the rug, thus depriving potentially affected individuals the opportunity to challenge decisions made about them. In the wake of the recent Clearview AI incident, PIAC would surmise that public trust in law enforcement is low. This public consultation and development of the Service's AI Policy is a valuable opportunity to engender public trust and accountability. Incorporating public transparency to the greatest possible degree into the new AI Policy is a foundational part of this aim.

VI. Establish stronger monitoring requirements

26. PIAC submits that the monitoring and Continuous Review provisions in the draft Policy require substantial clarification and improvements. Section 10 of the Policy provides that for the first year of full deployment of High or Medium Risk tech, the Chief of Police will monitor the performance indicators approved by the Board under section 5(n). However, section 12 provides that the Chief of Police will "continue to track the indicators approved by the Board under section 5(m)5(n) until it is determined by the Board that no additional monitoring is required." Clarity is needed regarding whether the performance indicators are monitored for just one year or until Board deems tracking no longer required.

27. Under Section 11, the Policy also requires the Chief of Police to provide the Board with a report within 15 months of full deployment describing the first year of deployment, compliance with applicable laws, performance according to Section 5(n), concerns from the public, results from post-deployment consultations, and whether the use of the AI technology will change in the future. This seems to be the only detailed reporting requirement in the Policy after deployment.

28. The only monitoring requirements after the initial report are provided for under Section 18 and 19. Section 18 requires that the Chief of Police review the continued use of High or Medium Risk technology every 5 years, based only on the quality of the technology, its outputs, and Key Performance Indicators, and the need for continued use. Section 19 provides that the Police Chief must also review every 5 years High, Medium, or Low risk technologies to determine whether they have been used for a novel purpose or in novel circumstances, in which case the Chief must provide a new report under Section 5. PIAC submits that once every 5 years is far too infrequent, and imposing monitoring responsibilities only on the Police Chief is dangerous. The Policy does not require these reviews be transparent to the public, and there is no involvement of

independent monitoring bodies in even these infrequent reviews. Therefore, PIAC recommends a continuous review requirement of every 12 months by both the Police Chief and an independent monitoring body, on all metrics measured under Section 11. Both the OHRC and the LCO recommend independent oversight, as mentioned above, and the OHRC also recommends that an independent oversight body assess and audit the Service's AI systems every 12 months, for bias and discrimination, and then report on and address systemic issues.

VII. Due process must be built into public feedback procedures

29. Under Section 13, the Policy requires that the Board provide a way for the public to submit concerns through the Board's website. The Executive Director then reports to the Board on the summary of concerns raised. Section 13(c) further instructs that the Executive Director, if the "communication from a member of the public amounts to a complaint, will advise the individual of their right to file a complaint with the Office of the Independent Police Review Director or successor role, or forward the communication to the Chief of Police, as appropriate, and inform the complainant of this action." It is unclear what the distinction is between a "concern" and "complaint." More importantly, the Executive Director is given full discretion to conduct traffic control on public feedback.

30. PIAC submits that under this Policy, the public is provided no guarantees as to whether or when their concerns and complaints will be addressed or even seen – we agree with the OHRC that there needs to be a clear, public process for addressing concerns from the public. The OHRC recommended that as a best practice, the complaints process "must adopt principles of due process and procedural fairness in its design, including an appeals process for complaints." At the very least, PIAC submits that all concerns must be seen by both the Board and the Chief of Police, and responded to by whichever of the two is most appropriate, within a set, reasonable amount of time. Whether responses are provided to individual complainants or more broadly on the Board's website may depend on the nature and frequency of issues raised.

VIII. Proceed with caution before privacy reform is finalized and in absence of laws governing AI policing

31. Notwithstanding PIAC's recommendations above, the Board should exercise caution before any provincial or federal privacy reform is finalized, particularly if reform affects how AI policing is regulated. The OHRC suggested that "any further actions or steps taken related to AI, especially in high-risk applications such as facial recognition or predictive policing, should be taken with extreme caution until such legislation and regulations are enacted." PIAC strongly agrees with this position. An internal policy or guidance is effective only to the extent that the underlying, enforceable legal framework is clear on the obligations of institutions, and provides for accountability, contractual transparency, and record-keeping requirements. As an example of

where the legal framework failed, PIAC points out that weaknesses in the *Privacy Act* are partly to blame for enabling the RCMP's improper use of Clearview AI.

32. Furthermore, changes in private sector privacy laws are coming soon. Some of the discourse surrounding this impending reform includes concerns that the change may not be for the better. A revived federal Bill C-11, and/or new provincial laws that are modeled upon it, may make it easier for the Service's private sector partners to collect and process more personal information under much broader consent exemptions, thus eroding the privacy interests of individuals. PIAC suggests that it may be wise to avoid fully operationalizing a new internal policy until reform is finalized, at which point consultations may continue on the Policy.

33. As a patchwork of general laws of application, Canada's current legal framework is ill-equipped to govern the use of AI by police. Furthermore, the conception of privacy under the Canadian privacy framework falls short of explicitly recognizing a constitutional and standalone right to privacy. In Bridges, the applicant brought a claim for judicial review on the basis that law enforcement's use of automated facial recognition violated his human right to respect for private life under the European Convention on Human Rights. There is no directly comparable right in Canada.

34. Though many privacy advocates continue to argue for policymakers to enshrine a right to privacy in legislation, the current framework is still not rooted in an explicit, broad acknowledgement of privacy as a fundamental human right that carries with it the weight of intrinsically linked rights and freedoms like freedom of thought and expression. Without the weight of these rights, the balancing exercise is such that the claimed benefits of using AI systems can more easily outweigh individual rights and freedoms in proportionality analyses, and the standards of harm mitigation are set lower. PIAC submits that, as long as this is the case, individuals and vulnerable groups are in a considerably weakened position in comparison to that of the private and public actors who seek to use personal information to fulfill their aims.

IX. Conclusion

35. PIAC submits that the draft Policy that the Board has published in this consultation is a good first step demonstrating initiative and good will. However, PIAC is largely aligned with LCO and OHRC's preliminary comments that there are further improvements to be made. Too much discretion in the hands of the Board and the Service poses significant risks to the rights and freedoms of individuals. A strong AI policy must place clear, appropriate limits on police discretion regarding public transparency, monitoring, harm mitigation, self-governance, and public feedback. A critical way to circumscribe discretion is to be as specific as possible in what types of AI systems fall under each risk category, and strictly forbid AI technologies that are found to carry bias in its operation and in its training data.

36. In the meantime, PIAC urges caution in proceeding with AI technologies under a final Policy. The federal and provincial legislative framework for privacy, implicating law enforcement's use of AI and engagement with private partners, will very soon enter a period of flux. Until broader reforms are finalized and enforceable laws are clarified, the Toronto Police must exercise restraint in using AI technologies, particularly higher-risk applications.

***** END OF DOCUMENT *****