DR. KWEKU OPOKU-AGYEMANG

WRITTEN RESPONSE: BUSINESS CATEGORY

OFFICE OF THE PRIVACY COMMISSIONER OF CANADA

NOTICE OF CONSULTATION AND CALL FOR COMMENTS:

PRIVACY GUIDELINES ON FACIAL RECOGNITION FOR POLICE AGENCIES

OCTOBER 15, 2021

Commissioner Therrien, Deputy Commissioner of Compliance Homan, and Members of the Legal Services Directorate,

Thank you for this opportunity for me to provide stakeholder feedback which I saw on your website. Facial recognition within the law enforcement domain may be the most polarizing application of artificial intelligence. Some believe it can enhance police investigations with better information and help our police agencies to be better at keeping Canadians safe. Others consider facial recognition algorithms to be a significant threat to citizen and community well-being due to algorithmic bias and other disadvantageous social impacts.

I am of the personal view that the privacy, ethical and social constraints of facial recognition information systems must be taken as seriously as possible for police agencies to meet or exceed community expectations and minimize bias. On careful study, I believe that the PIPEDA (Personal Information Protection and Electronic Documents Act) guidance from the Office of the Privacy Commissioner of Canada meets this challenge and will go a long way to help the municipal, regional, provincial and federal police agencies better serve and protect all Canadians.

Overall, I find that the guidance rigorously touches on the important points of Canadians' rights to privacy and the need to protect personal data, while providing a healthy sense of flexibility towards organizations that collect, use and disclose information for legitimate purposes.

My responses to the request for feedback proceed on page 3. I summarize my responses to the provided questions on the draft guidance. I focus mostly on the policy impact aspects to complement other submissions you may have received.

This document will be emailed to the Office of the Privacy Commissioner of Canada to hopefully assist with their information-gathering processes from stakeholders. Please be advised that this document will also be published as is online. This response is understood to be my personal opinion and represents no other person or organization. My contact information is at the end of the response as requested.

Thank you again for the opportunity to contribute to the discourse on the future of machine learning and privacy in Canada.

Best wishes,

Kweku Opoku-Agyemang, Ph.D.

Toronto, Canada

Contact: kweku@machinelearningxdoing.com

1. **Will this guidance have the intended effect of helping to ensure police agencies' use of FR is lawful and appropriately mitigates privacy risks? If you don't believe it will, why?**

Yes, but there is room for improvement. The guidance should have the intended effect of helping the police in terms of using FR lawfully and mitigating privacy risks. However, I anticipate significant heterogeneity or differences in terms of how well the guideline performs with respect to the lawful use of FR and any associated mitigation of privacy risks. For example, some police agencies are more diverse than others in terms of their demographics, social and economic characteristics. I expect this variation to be detected in the impact of the guidelines over time, but I expect these effects to be small. The more salient heterogeneity may have to do with the variation in how familiar police agencies are with the social impact of FR. Federal or municipal police agencies may be more aware than territorial agencies, on average. This perception is in line with the observation that "one size does not fit all": the guidelines themselves were clearly written with flexibility in mind.

2. **Can this guidance be practically implemented?**

Yes. The guidance can be practically implemented. The guidelines do seem rather clear and well-thought-out. Furthermore, police agencies already have certain guidelines they must follow and PIPEDA may be thought of in a similar vein.

In my answer to this question, I shall try to focus on how a hypothetical agency might consider implementing PIPEDA. Please note that I am not referring to any real-life agency, but how any organization in any sector might wish to implement any new guideline or policy to be practical.

One simple way to ensure that the guidelines are practically implemented would be to track agencies and officers within them over time to understand any issues that may arise over time. Since every agency is unique in its own way, a perhaps more sophisticated way to study implementation is to use internal pilot programs, where a sample of units or officers are exposed to the guidance then tracked over time in terms of their adherence and impact on lawful FR use and privacy risks. These would be compared to a group that is exposed at a different time (e.g. perhaps some agencies may use workshops to train officers on PIPEDA guidance). Whatever lessons emerge from the group exposed earlier would then be baked into the roll-out as the guidance is scaled throughout the agency. This is the core idea behind policy experiments (where a group is exposed to the guidelines and compared with a group that is exposed to the guidelines later). In many cases, however, police agencies are too resource-constrained to pursue a fully-fledged experiment. In such cases, some might consider pursuing "natural" experiments. Natural experiments are situations where the PIPEA policy can be evaluated within an agency even when it is not feasible to randomly assign access to the policy. This acclaimed statistical technique of natural experiments is the subject of the recently-announced 2021 Nobel Prize in economics (Royal Swedish Academy of Sciences, 2021), partially won by a Canadian.

These proposals only refer to changes in how the guidance is rolled out in a tailored manner to suit a unique police agency's context, but the response to this question does not require changes to the guidance itself.

3. **Are the recommendations in the "accuracy" section sufficient to help ensure police agencies meet their accuracy obligations in FR initiatives?**

The recommendations in the accuracy section seem well-thought out and with an eye towards the constraints. By "threshold", the section is referring to a user setting for FR. The acceptance or rejection of a match depends on the score of the match being higher or lower than a defined threshold, which may be adjusted to suit the localized context.

As far as I am aware, best practices are simply not yet available at the time of writing. There are trade-offs that must be acknowledged. If matching thresholds are set too high, even good matches are likely to be rejected. On the other hand, if the matching threshold is too low, police agencies become more likely to pick up face "matches" that are incorrect (that is, false positives). As inherently assumed in the accuracy section (Office of the Privacy Commissioner of Canada, 2020), the ability to meet accuracy obligations would ultimately depend on the algorithm deployed as well as the person that ultimately designs and implements it.

What seems a helpful complementary approach is to mandate a thorough human review of FR responses for agencies. Deferring to the algorithms is insufficient, simply because the goal should be to augment human work, not replace it.

4. **Can the recommendations in the guidance concerning the retention and disposal of personal information collected and used during a FR initiative be appropriately operationalized in a law enforcement context? If not, why?**

Yes. I believe they can, with the response from question 2. in mind.

5. **What measures or practices can police agencies implement to help ensure any third parties involved in FR initiatives operate with lawful authority?**

Police agencies can encourage third parties they engage in FR initiatives to append a section based on the PIPEDA guidelines to their collaboration agreement to maximize the likelihood of compliance. Vendors of FR software or stakeholders in control of faceprint databases accessed by police can be asked to provide a description of the circumstances surrounding the data collection process that are relevant.

6. **Do you foresee any negative consequences arising from the recommendations outlined in this guidance, and if so, what are they?**

Any policy, no matter how well-intentioned, may have some unintended consequences of a negative variety. The PIPEDA guidelines document is rather clear and thorough in its communication, which should minimize the incidence of these negative outcomes. A few possible candidates that come to mind are briefly described. For one, it is possible that some of the guidelines will inadvertently slightly reduce the efficiency of police agencies as they learn them for the very first time. This is not an unusual outcome in any policy realm and may

particularly occur in low-resource agencies in economically vulnerable areas (which may be more prone to criminal activity to begin with). However, I do not expect of the guidelines to have a significantly negative impact under any circumstances. As noted, any adverse effects are likely to be very short-lived as agencies become more comfortable with the guidelines over time.

7. **Is police use of FR appropriately regulated in Canada under existing law? If not, what are your concerns about the way police use of FR is currently regulated, and what changes should be made to the current legal framework?**

I believe that FR regulation in Canada would benefit from minor additions, but that these should build on existing privacy laws and not necessarily be a standalone regulatory framework.

One addition worth considering is that FR transcends facial recognition (that is, face identification and face verification). Some aspects that do not receive sufficient regulatory attention in my private view are photo clustering, race analysis, and real-time tracking to name a few. Ultimately, all forms of face recognition are potential threats to Canadian social ideals such as privacy.

The most commonly-deployed use of facial recognition is what is known as face matching or face identification. Here, the goal is to match two or more faceprints to investigate whether or not they represent the same human. In the law enforcement application, a faceprint from a security camera may be compared with a government database of ID photos. This appears to be the main use case covered by PIPEDA if I am not mistaken.

However, it is not the only technique, as the writers of the PIPEDA police are likely aware. Even without attaching faces to images, face matching can track a person in real-time, and this process (generally known as face tracking), may also be worthy of attention from the Office of the Privacy Commissioner in the future. I believe that these scenarios would benefit from PIPEDA guidance as well.

8. **What protections should be granted to individuals whose biometric information is included in a faceprint database?**

The protections granted to individuals whose biometric information is contained in a faceprint database should (1) be transparent to such individuals and (2) not be inordinately burdensome for police agencies that are using such information for purposes that are part of their roles and responsibilities. In my view, one of the best approaches worth considering in the medium-to-short term is via Institutional Review Board (IRB) frameworks that are commonly used in academic research institutions and computer science departments when research covers human subjects. Here, academics write out their hypotheses and data analysis and algorithmic agenda prior to collecting data.

There is no official registry for research ethics boards in Canada at the time of writing (which police agencies would be privy to), but those available to university researchers at our universities can serve as a model framework. Under certain circumstances, certain police agencies may create their own internal IRB offices that approve protocols concerning the use of

faceprint data on a case-by-case basis prior to any such predictive analytics being implemented. Such protocols can be renewed every period, assuming that investigators are following the PIEPA guidelines to a satisfactory degree. (This note is a mere summary response to the question and insufficient for policy in and of itself).

9. **Should police use of FR, including the collection of faceprints, be limited to a defined set of purposes (such as serious crimes or humanitarian reasons, e.g. missing persons)? Should they be able to use or retain faceprints beyond those of individuals who have been arrested or convicted?**

I believe that there must be constraints on the use and collection of faceprints such as the scraping of public images from the internet without user permission; not necessarily in terms of the severity of criminal allegations, but with respect to providing validation for standard methods of identifying perpetrators. That is, they should augment the professional work of police agencies and not attempt to replace it.

I feel I should explain this position a little further. Limiting the use of FR to a pre-defined set of purposes may introduce bias in the types of alleged criminal activity that receives algorithmic attention (for example, it is often the case that petty shoplifting tends to have different ethnic dimensions from white collar crime, and the latter is relatively understudied). Overall, I find it difficult to ascertain what police agency investigation scenarios may or may not require the use of FR, and would defer to police agency legal officers' responses to this particular question.

Although I am not a police officer, one must keep in mind that some criminals would be aware of such requirements if they were to be public information and at least some would adjust their behavior accordingly (for example, one might expect some to focus on high-reward crimes that are known to be exempt from FR). I also do not find it very pragmatic to only use or retain faceprints in cases where the alleged perpetrator has already been arrested or convicted. FR may be most useful to police agencies and the public as a whole when the alleged perpetrator is not yet in custody in many cases. However, this position would only be feasible when the technology matures further. Some of this would be more easily answered with access to databases from police agencies for an independent analysis of how effective they are. I do not have a full answer but hope the response was informative.

10. **Are there any other important policy issues that should be addressed in relation to police use of FR?**

I have no further policy issues to suggest. Thank you.

REFERENCES

Office of the Privacy Commissioner of Canada (2020). “PIPEDA Fair Information Principle 6—Accuracy.” https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/p_principle/principles/p_accuracy/

Royal Swedish Academy of Sciences (2021). “Natural Experiments Help Answer Important Questions.” The Prize in Economic Sciences 2021: Popular Science Background https://www.nobelprize.org/uploads/2021/10/popular-economicsciencesprize2021-3.pdf