Comments on AI consultation paper

13 March 2020

Submitted to:

Office of the Privacy Commissioner of Canada
30, Victoria Street
Gatineau, Quebec
K1A 1H3

To Whom It May Concern:

The world is watching Canada. Federal research development and industrial strategy have positioned technologies commonly referred to as artificial intelligence (AI) as a centrepiece of Canada's digital economy. In the wake of these initiatives, calls for Canada to establish a strong rights-based regulatory framework have resulted in proposals such as the Digital Charter. Even though AI governance in Canada has lagged behind investment, the Federal government continues to be perceived as a world leader in the responsible development of AI. Consultations like this one set national and international precedents and cannot be taken lightly.

Privacy is only part of the solution to the challenge of responsible development of AI. Justice, equality, human dignity, and democratic norms all must be taken into consideration to achieve a stable and responsible context for the development of AI.

We seek to offer recommendations that lie within the scope of the OPC's current and future mandate. Our statement has been collectively drafted and reflects the shared views of an interdisciplinary group of scholars in Montreal. Together, we combine expertise in a variety of fields and subjects with a shared interest in the study of science and technology as well as insights gained from studying or working in AI development.

Before we respond to the questions, we wish to highlight the larger issues that must be part of the government's overall strategy. AI governance in Canada should:

1. Interpret the Canadian Charter of Rights and Freedoms as it applies to the development of artificial intelligence

Nationally, most frameworks for the responsible development of AI have been non-binding ethical declarations. While we welcome industrial efforts to develop shared norms, the idea that AI development exists in a regulatory void is a fiction. Instead, there needs to be immediate clarification of the charter as it relates to AI. Privacy rights are part of this interpretation.

2. Clarify the status of intellectual property with regard to algorithms trained using public knowledge, culture, and history as well as individual expertise

At a time when public websites like Reddit[1] and Wikipedia[2] or the everyday habits of people on the street are sources of AI training data, what is the relationship between these trained algorithms and their public sources? Uncertainty about the status of training data may facilitate the enclosure of common goods as proprietary black-boxed algorithms. "Intelligence" might be a proxy for common sense or public knowledge. The integration of workplace surveillance and AI may result in employees training their automated replacements, a process that converts individual expertise into corporate assets.[3] These translations cannot go unquestioned, and shared common goods and cultures should not be readily commodified by AI. Possible solutions such as data trusts should be pursued.[4]

Furthermore, the terms and conditions of use that social media users agree to (or have agreed to) may require renewed regulatory scrutiny. While platforms have historically claimed a legal right of ownership over user data, the development of new and complex uses for that data problematizes the assumption that users have already given informed consent. Users of a photo-sharing service, for example, could not have foreseen their photos being used to train a facial-recognition system years after they uploaded them (and prior to facial recognition becoming a viable technology).[5]

[1] "R/NoSleep, One of the Largest Subreddits On Reddit, Goes Dark In IP-Theft Protest," Slashdot, February 24, 2020, https://yro.slashdot.org/story/20/02/24/2245247/rnosleep-one-of-the-largest-subreddits-on-reddit-goes-dark-in-ip-theft-protest.

[2] Bernard Marr, "The Amazing Ways How Wikipedia Uses Artificial Intelligence," *Forbes*, August 17, 2018, https://www.forbes.com/sites/bernardmarr/2018/08/17/the-amazing-ways-how-wikipedia-uses-artificial-intelligence/.

[3] Solon Barocas, "Machine Learning Is a Co-Opting Machine." *Public Books* (blog), June 18, 2019. https://www.publicbooks.org/machine-learning-is-a-co-opting-machine/.

[4] "Report—International Meeting on Diversity of Content in the Digital Age," Canadian Heritage, May 29, 2019, https://www.canada.ca/en/canadian-heritage/services/diversity-content-digital-age/international-engagement-strategy/report.html.

[5] Shannon Liao, "IBM Didn't Inform People When It Used Their Flickr Photos for Facial Recognition Training," The Verge, March 12, 2019, https://www.theverge.com/2019/3/12/18262646/ibm-didnt-inform-people-when-it-used-their-flickr-photos-for-facial-recognition-training.

3. *Develop better Aboriginal-Crown relations in regard to Indigenous data sovereignty and epistemologies*

As settler allies, we recognize and value the existing knowledge of and emerging research on Indigenous approaches to data and AI.[6] In Canada, First Nations representatives have been leading efforts to assert data sovereignty over the past twenty years.[7] For instance, the First Nations Information Governance Centre (FNIGC), a nonprofit organization operating with a special mandate from the Assembly of First Nations' Chiefs in Assembly, has been working on questions around data sovereignty since the mid-1990s, and the First Nations principles of OCAP® (Ownership, Control, Access, and Possession) were developed in 1998.[8]

On the global level, various advocacy groups have put forward agendas and principles for Indigenous self-determination over data and data analysis (see e.g., the principles for Indigenous data governance by the Global Indigenous Data Alliance).

These approaches need to be understood, respected, and incorporated into any program for the responsible development of AI—any such framework should commit to actively improving Aboriginal-Crown relations.

4. *Establish better systems of public consultation and expression of the common interest*

Three years ago, Dr. Michael Geist warned that: "the volume of consultations runs the real risk of turning into 'consultation theatre,' where the government or agencies seek out public participation not as a mechanism to generate ideas or gauge public opinion, but rather as a validation exercise at best or as theatre with no intent to act on submissions at worst."[9] His warning seems even more apt today with regard to this primarily written consultation occurring more or less concurrently with major reforms in broadcasting, telecommunications, data, privacy, and copyright. The current model of consultation continues to be, as Geist notes, "simply unsustainable for all but the most deep-pocketed organizations, effectively excluding many public interest voices who lack the

[6] For more, see: John Taylor and Tahu Kukutai, eds., *Indigenous Data Sovereignty: Toward an Agenda* (Acton: Australian National University Press, 2016); Stephanie Carroll Rainie et al., "Data as a Strategic Resource: Self-Determination, Governance, and the Data Challenge for Indigenous Nations in the United States," *International Indigenous Policy Journal* 8, no. 2 (2017), https://doi.org/10.18584/iipj.2017.8.2.1.

[7] The First Nations Information Governance Centre, "First Nations Data Sovereignty in Canada," *Statistical Journal of the IAOS* 35, no. 1 (2019): 47–69, https://doi.org/10.3233/SJI-180478.

[8] For more details, see the First Nations principles of OCAP®.

[9] Michael Geist, "Too Much of a Good Thing: What Lies behind Canada's Emerging Consultation Crisis," October 25, 2017. http://www.michaelgeist.ca/2017/10/too-much-of-a-good-thing-what-lies-behind-canadas-emerging-consultation-crisis/.

resources to invest a growing proportion of their time on government consultations or hearings." Given that Geist made these comments three years ago, what solutions has the OPC put in place? We are concerned that the announcement and short consultation window for this framework and future policy may exclude the public from AI governance.

5. Monitor concentration, power imbalances, and the overall state of the AI industry

Responsible AI may be developed in irresponsible conditions, where, as in Germany, the lack of consumer choice means that consent to participate cannot be meaningfully given.[10] Likewise, market concentration, consumer lock-in,[11] and research capacity determine the limits of possibility for the responsible development of AI.

6. Recognize that artificial intelligence requires situated, public knowledge

There is a lack of public literacy about AI, leading to utopian and dystopian thinking about its future without understanding the contexts, perspectives, and applications of AI. M. C. Elish and danah boyd warn that without better public knowledge about AI in regulatory settings, the industry will be the only source of information about its capabilities and limitations:

> Through the manufacturing of hype and promise, the business community has helped produce a rhetoric around these technologies that extends far past the current methodological capabilities. . . . The ability to manufacture legitimacy has far-reaching implications. Not only does it trigger innovation and bolster economies, but it also provides cover for nascent technologies to potentially create fundamentally unsound truth claims about the world.[12]

Public understanding of AI is currently strongly influenced by actors whose primary goal is to sell the public, or at least investors, on the glories of AI.

The semantics of intelligence, agency, and accountability demand critical investigation. Even the term "artificial intelligence" suggests an equivalence between humans and machines that likely will confuse sound policy and mire policy recommendations in anthropocentric thinking. As Jason Edward Lewis, Noelani Arista, Archer Pechawis, and Suzanne Kite write, "a problematic aspect of the current AI debate is the assumption that AIs would be homogeneous when in fact every AI would be profoundly different, from a military AI designed to operate autonomous killing

[10] Emily Dreyfuss, "German Regulators Just Outlawed Facebook's Whole Ad Business." *Wired*, February 7, 2019, https://www.wired.com/story/germany-facebook-antitrust-ruling/.

[11] Matthew Scott Hindman, *The Internet Trap: How the Digital Economy Builds Monopolies and Undermines Democracy* (Princeton, NJ: Princeton University Press, 2018).

[12] M. C. Elish and danah boyd, "Situating Methods in the Magic of Big Data and AI," *Communication Monographs* 85, no. 1 (2018): 57–80, https://doi.org/10.1080/03637751.2017.1375130.

machines to an AI built to oversee the United States' electrical grid."[13] Indeed, there likely will be no one universal definition of AI and, instead, this consultation should better define the concerns that necessitate the generic term and how to understand its meaning in context.

With these broader concerns in mind, we now provide answers to the OPC's questions as provided in the Consultation. At a high level, we wish to stress that:

1. Privacy is one possible data ontology.

Data protection, assessment, and privacy require the production of a normative framework for the way data comes to be in the world. We consider matters of sustainability, justice, dignity, and diversity to be part of responsible data collection and use; therefore, they are a fundamental part of the responsible development of AI. Methods such as Privacy by Design, Human Rights Assessment, Algorithmic Impact Assessments, and data supply chains all need to be framed not merely as accountable, but as policy instruments to actualize better data for use by AI systems. Commons models, proactive disclosure, and access to information all must be considered as part of the systems that produce data.

2. Regulation through black-boxed algorithms poses real and fundamental challenges to accountable governance.

The turn toward automated decision-making increases complexity and creates new barriers to transparency that impede public accountability. Trade secrets, the inscrutability of algorithms under the proposed USMCA, as well as the inexplicability or unpredictability of AI systems may never be reconciled with democratic norms where laws and policies are publicly understood.

3. The development of individual privacy rights does not absolve the Privacy Commissioner from proactive enforcement and continuing investigation of the collective and societal impacts of AI.

The latter recommendations for expanded powers of the OPC are welcome remedies to today's status quo where obtuse individual privacy controls impede systemic remedies and improve public welfare.[14]

4. The OPC should be able to prohibit certain applications of AI based on the sources of training data or the assumptions in their inferences.

[13] Lewis et al.

[14] Woodrow Hartzog and Neil Richards, "It's Time to Try Something Different on Internet Privacy," *Washington Post*, sec. Opinion, December 20, 2018. https://www.washingtonpost.com/opinions/its-time-to-try-something-different-on-internet-privacy/2018/12/20/bc1d71c0-0315-11e9-9122-82e98f91ee6f_story.html.

Not all applications of AI have the same risks; nor do all AI systems make the same knowledge claims. For example, we are sceptical of technologies that claim to automatically detect emotion, mood, deception, criminality, or other behavioural qualities as well as race, gender, or sexuality from images of faces. There is ample evidence that the science behind these systems is disputed. While scientific debate is important, high-risk applications of disputed technologies, such as in the criminal, medical, or immigration systems, mandate greater scrutiny and should be temporarily prohibited if necessary.

5. Algorithmic Assessments need to be developed as robust, interdisciplinary research tools integrated into design, deployment, and ongoing monitoring.

The OPC needs to support the development of better assessment tools to improve the responsible development of AI. These tools need to be grounded in shared understandings, making them more akin to research methodologies than most of the current proposals that risk being determined by the person who completes the test. Experiments, research creation, ethnography, and new digital methods all need to be integrated into the interdisciplinary teams assessing AI.

6. Responsible decision-making should not simply be a matter of greater rationalization.

Responsibility requires systems to acknowledge uncertainty, fallibility, and error. These qualities can be attributed to humans and machines. The goal should not be to create a perfectly rational decision-making process, but to establish guidelines for deliberation in the face of uncertainty, including reviews and appeals as well as redundancy to ensure shared responsibility between AIs and humans-in-the-loop.

Much of the promise of automated decision-making (ADM) presumes an internal consistency between decisions, higher-level laws, and principles. This attitude clashes with the longstanding tradition of legal realism, i.e., the notion that the law should be grounded in real-world experience and evidence. ADM threatens to strip the law of *judgement*, or the capacity to make considered decisions and come to sensible conclusions. This element must be at the forefront of discussions around AI policy, including the ability to judge, individually and as a society, when not to use AI technologies.

These general comments reflect the shared opinions of our group. We look forward to supporting the OPC as it moves forward in establishing sound governance for the responsible development of AI.

Signed by (in alphabetical order),

Ana Brandusescu
Professor of Practice, Centre for Interdisciplinary Research on Montreal
McGill University

Janna Frenzel
PhD Candidate in Communication Studies
Concordia University

Nick Gertler
BA Student in Communication Studies
Concordia University

Robert Hunt
PhD Candidate in Communication
Concordia University

Fenwick McKelvey
Associate Professor, Communication Studies
Concordia University

Nicole Rigillo
Berggruen Fellow

Bart Simon
Associate Professor of Sociology, Milieux Institute for Arts, Culture and Technology
Concordia University

Luke Stark
Postdoctoral Researcher
Microsoft Research

Ceyda Yolgormez
PhD Candidate in Social and Cultural Analysis, Concordia University

Proposal 1: Incorporate a definition of AI within the law that would serve to clarify which legal rules would apply only to it, while other rules would apply to all processing, including AI

1. Should AI be governed by the same rules as other forms of processing, potentially enhanced as recommended in this paper (which means there would be no need for a definition and the principles of technological neutrality would be preserved) or should certain rules be limited to AI due to its specific risks to privacy and, consequently, to other human rights?

The term artificial intelligence invites ambiguity precisely because it has been the site of differing research agendas and projects for over fifty years. What advantages does the term offer us today?

Law should remain technologically neutral while recognizing the particular class of technologies with the emergent properties that we have come to refer to as "artificial intelligence." Artificial intelligence then refers to a general concern about unpredictable, complex technologies.

Decrease predictability

Artificial intelligence might in this case be a shorthand for systems that exhibit emergent behaviour and function unpredictably. The capacity for agency is often assigned to technologies that behave unpredictably or fail. In that sense, the fallibility of these technologies should be considered in the definition itself. The general inclination is toward thinking that when a job is outsourced to a machine, the machine will perform better than its human counterpart, resulting in fewer errors and problems. While this might be true for certain types of machinery, the history of AI shows us that these particular systems exhibit their agency via their unpredictability. The flexibility and uncertainty inherent to the functioning of AI systems introduce an irreducible element to the process of decision making/action. This is why edge cases should be treated as central concerns while thinking about policies around AI.

Increase complexity

AI might also require a refined definition of the complexity of technology before the law. Paying attention only to the algorithm might ignore material infrastructures (such as telecommunications infrastructure, data centres, and electrical grids), data collection practices (especially OSINT and repurposed data), and administrative procedures (such as augmenting or automating human decision-making).

Proposal 2: Adopt a rights-based approach in the law, whereby data protection principles are implemented as a means to protect a broader right to privacy—recognized as a fundamental human right and as foundational to the exercise of other human rights

1. What challenges, if any, would be created for organizations if the law were amended to more clearly require that any development of AI systems must first be checked against privacy, human rights and the basic tenets of constitutional democracy?

Data protection is fundamental to the responsible development of AI. Data protection needs to ensure that the possibilities of new statistical reasoning developed through machine-learning at least recognize constitutionally granted rights in Canada as well as acknowledge the social impacts of the design and development of new technology.

AI development is driven by poorly regulated and highly concentrated data collection that relies on largely US-based contract law to justify the global repurposing of communication and other data into training data sets. Indeed, "big data" may be more of an issue than AI. As AI Now cofounder and codirector Meredith Whittaker explained in her testimony before the United States House of Representatives Committee on Science, Space, and Technology:

> AI is not a new set of technologies, and many of the core techniques that power AI systems, including neural nets, have been around for decades. The biggest changes recently have not been wholly new AI techniques (although we have seen improvements and innovations). What has changed drastically is the availability of massive amounts of data and vast computational resources. It's this that is behind the AI boom we see today. These are assets that only a handful of major tech companies have, and very few others do. This is one of the reasons why the US government contracts with companies like Amazon, instead of building its own infrastructure and AI. Without the legal protections afforded the private sector around privacy, and existing market reach and infrastructural economies of scale, it's virtually impossible to obtain the resources needed to create AI from scratch.[15]

Many existing organizations have been built around the use of data and its analysis by artificially intelligent systems in ways that violate the right to privacy. Entire systems have developed around the routine violation of privacy, and users have become accustomed to the tradeoff of data for free or discounted services. The future of these kinds of systems would be largely dependent on how privacy and its violation is defined.

[15] https://ainowinstitute.org/062619-whittaker-house-testimony.pdf.

Data, however, should not be seen in the abstract. There is a difference between analyzing the logs collected by intelligent infrastructures versus records of human conversations. Data protection should focus on protecting, but not limited to, :

- data from common sense or the public domain;
- personal information, especially sensitive health information, even in the aggregate;
- shared cultural and historical records;
- Indigenous knowledge and specific National approaches to data sovereignty; and,
- data with significant public interest and utility.

Data protection needs to establish clear norms in the responsible collection and use of data, balancing both the needs of individual privacy as well as protecting against the commercialization of knowledge with significant public benefit.

Proposal 3: Create a right in the law to object to automated decision-making and not to be subject to decisions based solely on automated processing, subject to certain exceptions

1. Should PIPEDA include a right to object as framed in this proposal?

Yes, but the OPC should not understand a right to object as analogous to a right to withhold consent. This scheme is equivalent to the much-maligned "opt-in" vs. "opt-out" debates around cable packages in the 1990s—Canadians should not have to constantly opt out of this collection after the fact, particularly given the financial precariousness of start-up technology firms and the inherently fungible and easily transmissible nature of digital data. The OPC should include both a right to object and a right to proactively withhold consent.

2. If so, what should be the relevant parameters and conditions for its application?

The right to withhold consent should consider whether a decision is wholly or partly automated. We should distinguish between a fully automated decision, decisions made solely by humans, and decisions that are "augmented" by the inputs of artificially intelligent agents. Each of these raises unique challenges for bias, discrimination, and the extent to which human decisions can be "nudged" in particular directions by system designers. For example, the design of explainability interfaces and the kinds of information they present to users can shape user trust in AI systems as well as user perceptions of their capabilities.[16] The interface between humans and intelligent machines is not neutral—a notion that remains relatively new among the scientists and developers

[16] Carrie J. Cai, Jonas Jongejan, and Jess Holbrook, "The Effects of Example-based Explanations in a Machine Learning Interface," *Proceedings of the 24th International Conference on Intelligent User Interfaces* (2019): 258–62.

charged with building explainable AI.[17] Before concluding that the right to withhold consent is not necessary for augmented decisions, appropriate audits and safeguards need to be put in place.

Withholding consent should be the baseline, but being able to opt out is not as crucial as ensuring that the design and deployment of technologies respects human rights and dignity. Automated decision-making intrinsically concerns fairness, and without proper oversight it likely will exacerbate long-standing and systemic inequities in the application of the law. It is essential to ensure that automation is not designed to target impoverished or other at-risk populations. For example, *ProPublica*[18] recently reported that the USA's Internal Revenue Service tends to target low-value tax fraud because more complex cases cannot be detected with automated tools. Government needs to ensure that automated systems are not means of more intensely policing the poor and perpetuating systemic bias.[19]

The right to object should *not* be reserved only for Canadian citizens; it should be extended to anyone affected by an automated decision made in Canada. Without this clarification, the right to object to automated decision-making in Canada may be denied to those without Canadian citizenship (for, e.g., IRCC's use of AI and ADM in the processing of tourist visa applications of foreign nationals[20]).

Finally, we should not permit the use of "black box" models for high-stakes decisions. AI scientists like Cynthia Rudin have called for the limitation of high-stakes automated decisions to AI models that are *human interpretable*, or easily graspable without an additional machine-learning explanation method. This is both because explainability methods can themselves be erroneous,[21] and also because human interpretable explanations are inherently easier to understand and challenge.[22] Fully automated decision-making should **not** be used for high-stakes decisions. The grid used by the Treasury Board Secretariat to identify low, medium, and high impact AI use cases

[17] Tim Miller, Piers Howe, and Liz Sonenberg, "Explainable AI: Beware of Inmates Running the Asylum Or: How I Learnt to Stop Worrying and Love the Social and Behavioural Sciences," *arXiv* 1712.00547 (2017).

[18] Paul Kiel and Jesse Eisinger, "Who's More Likely to Be Audited: A Person Making $20,000—or $400,000?", *ProPublica,* December 12, 2018, https://www.propublica.org/article/earned-income-tax-credit-irs-audit-working-poor.

[19] Virginia Eubanks, *Automating Inequality* (London: St. Martin's Press, 2018).

[20] Petra Molnar and Lex Gill, "Bots at the Gate: A Human Rights Analysis of Automated Decision-Making in Canada's Immigration and Refugee System" (Toronto: International Human Rights Program and the Citizen Lab, University of Toronto, 2016).

[21] Julius Adebayo et al., "Sanity Checks for Saliency Maps," *Advances in Neural Information Processing Systems* (2018): 9505–15.

[22] Cynthia Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," *Nature Machine Intelligence* 1, no. 5 (2019): 206–15.

requires more elaboration and should include specific examples (e.g. what counts as a "low impact" decision in practice?) as well as a clearer set of criteria for determining exceptions.

Proposal 4: Provide individuals with a right to explanation and increased transparency when they interact with, or are subject to, automated processing

1. What should the right to an explanation entail?

A right to explanation is a necessary but insufficient element in an AI regulatory regime. Many of the suggested transparency mechanisms, while salutary, are technical and aimed at expert actors. For this reason, transparency will never be sufficient for ordinary individuals dealing with injustices propagated by AI systems. The Council of Europe's parameters are a good start, but rights to explanation are meaningless without the ability to be preemptively omitted from these systems of analysis or the wholesale banning of certain technologies in particular domains. The right to an explanation for an automated decision is largely a concern when adopting the use of black-box models for decision-making—and we advise against this. Avoiding the use of black box models for high-stakes decisions would obviate the need for explainability methods, which continue to be fallible, as mentioned above.

2. Would enhanced transparency measures significantly improve privacy protection, or would more traditional measures suffice, such as audits and other enforcement actions of regulators?

We recommend the OPC fund research into alternative forms of explanation. There are scant real-world examples of explanations being provided to end users affected by an automated decision made by a black box model. In a recent study of machine-learning models with explainability interfaces, explanations were largely provided to data scientists or area experts rather than end users.[23] In the relatively few cases where explanations were offered to end users, human experts served as mediators between the model and the user. However, in Canada, legal systems and notions of liability centered around humans as decision-makers will complicate the use of (and processes for objecting to) these explanations in practice. OPC should define protocols for audits of machine-learning models and forbid the use of black-box models for high stakes decisions in the interest of the public good.

Given existing barriers to the provision of faithful, trustworthy explanations, government should invest in interdisciplinary team research on explanations (including social scientists). Until

[23] Umang Bhatt et al., "Explainable Machine Learning in Deployment," *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (2020): 648–57.

explanations can be shown to be reproducibly faithful and trustworthy, black-box models should be avoided for high-stakes decisions altogether (e.g., extending credit, immigration, welfare, etc).

Proposal 5: Require the application of Privacy by Design and Human Rights by Design in all phases of processing, including data collection

1. Should Privacy by Design be a legal requirement under PIPEDA?

Yes, though subject to further consultation on implementation and scope.

2. Would it be feasible or desirable to create an obligation for manufacturers to test AI products and procedures for privacy and human rights impacts as a precondition of access to the market?

Yes, it is desirable to create an obligation for manufacturers to test AI products and procedures for privacy and human rights impacts as a precondition of access to the market. Feasibility might require limiting application to certain domains (e.g., health, financial decisions, etc.).

Proposal 7: Include in the law alternative grounds for processing and solutions to protect privacy when obtaining meaningful consent is not practicable

1. If a new law were to add grounds for processing beyond consent, with privacy protective conditions, should it require organizations to seek to obtain consent in the first place, including through innovative models, before turning to other grounds?

Yes.

2. Is it fair to consumers to create a system where, through the consent model, they would share the burden of authorizing AI versus one where the law would accept that consent is often not practical and other forms of protection must be found?

No.

3. Should consent be reserved for situations where purposes are clear and directly relevant to a service, leaving certain situations to be governed by other grounds? In your view, what are the situations that should be governed by other grounds?

4. How should any new grounds for processing in PIPEDA be framed: as socially beneficial purposes (where the public interest clearly outweighs privacy incursions) or more broadly, such as the GDPR's legitimate interests (which includes legitimate commercial interests)?

5. What are your views on adopting incentives that would encourage meaningful consent models for use of personal information for business innovation?

These concerns might largely be outside of the scope of the OPC because they refer to cultural and societal resources rather than personal information. The OPC might not be able to effectively rule on exemptions for artificial intelligence without factoring in the common resources required for the development of artificial intelligence that are difficult to capture within a system of individual rights.

Proposal 8: Establish rules that allow for flexibility in using information that has been rendered non-identifiable, while ensuring there are enhanced measures to protect against re-identification

1. What could be the role of de-identification or other comparable state of the art techniques (synthetic data, differential privacy, etc.) in achieving both legitimate commercial interests and protection of privacy?

De-identification has been proven to be ineffective at protecting privacy, particularly if multiple data sets are combined. De-identified data should be subject to the same privacy constraints as personally identifiable data. Synthetic data and K-anonymity are better mechanisms to protect privacy.[24]

Proposal 9: Require organizations to ensure data and algorithmic traceability, including in relation to datasets, processes and decisions made during the AI system lifecycle

1. Is data traceability necessary, in an AI context, to ensure compliance with principles of data accuracy, transparency, access and correction and accountability, or are there other effective ways to achieve meaningful compliance with these principles?

Yes. Standards for data logging should be developed that ensure inputs of data are identifiable and auditable. However, traceability is insufficient as a safeguard to long-term improvement of business practices. Companies should, as per the European Union's draft assessment list, be required to implement a series of checklists at every stage of the AI design appointment lifecycle. These checklists should be accompanied by other guidance to ensure they are being implemented

[24] Joseph Jermoe, "De-Identification Should Be Relevant to a Privacy Law, But Not an Automatic Get-Out-of-Jail-Free Card," *Center for Democracy and Technology* (blog), April 1, 2019. https://cdt.org/insights/de-identification-should-be-relevant-to-a-privacy-law-but-not-an-automatic-get-out-of-jail-free-card/.

appropriately, and are themselves not a panacea. However, traceability on its own will not solve the problems of unfairness in AI systems: many things can be visible without being immediately repaired.

Proposal 10: Mandate demonstrable accountability for the development and implementation of AI processing

1. Would enhanced measures such as those as we propose (record-keeping, third party audits, proactive inspections by the OPC) be effective means to ensure demonstrable accountability on the part of organizations?

Yes.

2. What are the implementation considerations for the various measures identified?

As with any enforcement of regulations, it will be important for companies to understand the proposed accountability mechanisms as tough but fair. These mechanisms should be drafted with an understanding of business practices, but should not let those practices dictate the values supported by enforcement. In other words, experts in the design and appointment process should work with the OPC to identify which claims by organizations are true tactical barriers to accountability and which are spurious. For example, the Partnership on AI's ABOUT ML initiative[25] aims to develop a standardized process/checklist for documentation throughout the machine learning system lifecycle (build, deploy, monitor).

3. What additional measures should be put in place to ensure that humans remain accountable for AI decisions?

One possible measure for small companies would be the creation of an OPC task force able to assist organizations in implementing appropriate procedures to ensure compliance with the new regulations. These regulations will be relatively easy for large companies to implement but will be more of a strain on smaller organizations. In order to make these regulations enforceable and effective, the OPC must find a mechanism to both ensure technical accountability and foster a sense of capacity and trust in smaller companies. In addition, the governance and legal structure of lobbying, especially corporate lobbying, needs to be revisited. In addition to the rationale of being too transparent, there is an existing and ongoing power imbalance where (tech) companies influence policymakers and policymaking for private interest and gain.

[25] Partnership on AI. ABOUT ML: https://www.partnershiponai.org/about-ml-get-involved/

Beyond regulatory measures, the OPC could support corporate transparency efforts, especially in legislation, such as creating a public beneficial ownership registry.[26] But more fundamentally, the availability of open data on companies in Canada needs improvement. In its efforts to zoom in on AI tech, OPC can strengthen ties with Innovation, Science and Economic Development Canada, who are leading the transparency and accountability work in the public sector on individuals (and companies) funding and managing this work.

See, for example, the implementation of public procurement reform policies and standards for transparency and accountability in AI companies that the government contracts and/or subsidizes like open contracting[27] and the Open Contracting Data Standard[28] that the federal government has adopted.[29]

Proposal 11: Empower the OPC to issue binding orders and financial penalties to organizations for non-compliance with the law

1. Do you agree that in order for AI to be implemented in respect of privacy and human rights, organizations need to be subject to enforceable penalties for non-compliance with the law?

Yes, this should apply to companies operating in Canada as well as Canadian AI companies operating internationally, including companies that have incorporated local offices in other jurisdictions. Penalties should extend beyond financial penalties.

2. Are there additional or alternative measures that could achieve the same objectives?

Yes. However, if the OPC will expand its powers, the Office should be required to have increased public oversight. For example, options may include establishing an academic advisory panel and a public interest advisory panel similar to recommendations in the Broadcasting and Telecommunications Legislative Review (BTLR) panel's report recommendations for the Canadian Radio-television and Telecommunications Commission (CRTC). The BTLR report proposes to change the name of the CRTC to the Canadian Communications Commission and to democratize appointment of commissioners, and it commits to establish "a transparent process for funding

[26] Strengthening Corporate Beneficial Ownership Transparency in Canada: https://www.ic.gc.ca/eic/site/142.nsf/eng/00001.html.

[27] Open Contracting Partnership. What Is Open Contracting? https://www.open-contracting.org/what-is-open-contracting/

[28] Open Contracting Partnership. Open Contracting Data Standard. https://standard.open-contracting.org/latest/en/

[29] Canada's Plans to the Open Government Partnership. End-of-Term Self-Assessment Report on Action Plan on Open Government 2014-2016. Action Plan Commitment 8: Open Contracting https://open.canada.ca/en/commitment/sar/2014-2016/action-plan-commitment-8-open-contracting

public interest participation" and a "Public Interest Committee funded by the CRTC."[30] These are all institutional reforms that initiate a longer journey toward better public understanding of media systems. That this might precede regulation of these systems may be a problem when the BTLR's recommendations are interpreted.

[30] Canadian Radio-television and Telecommunications Commission. CRTC written public submission to the Legislative Review Panel
https://crtc.gc.ca/eng/publications/reports/rp190110.htm