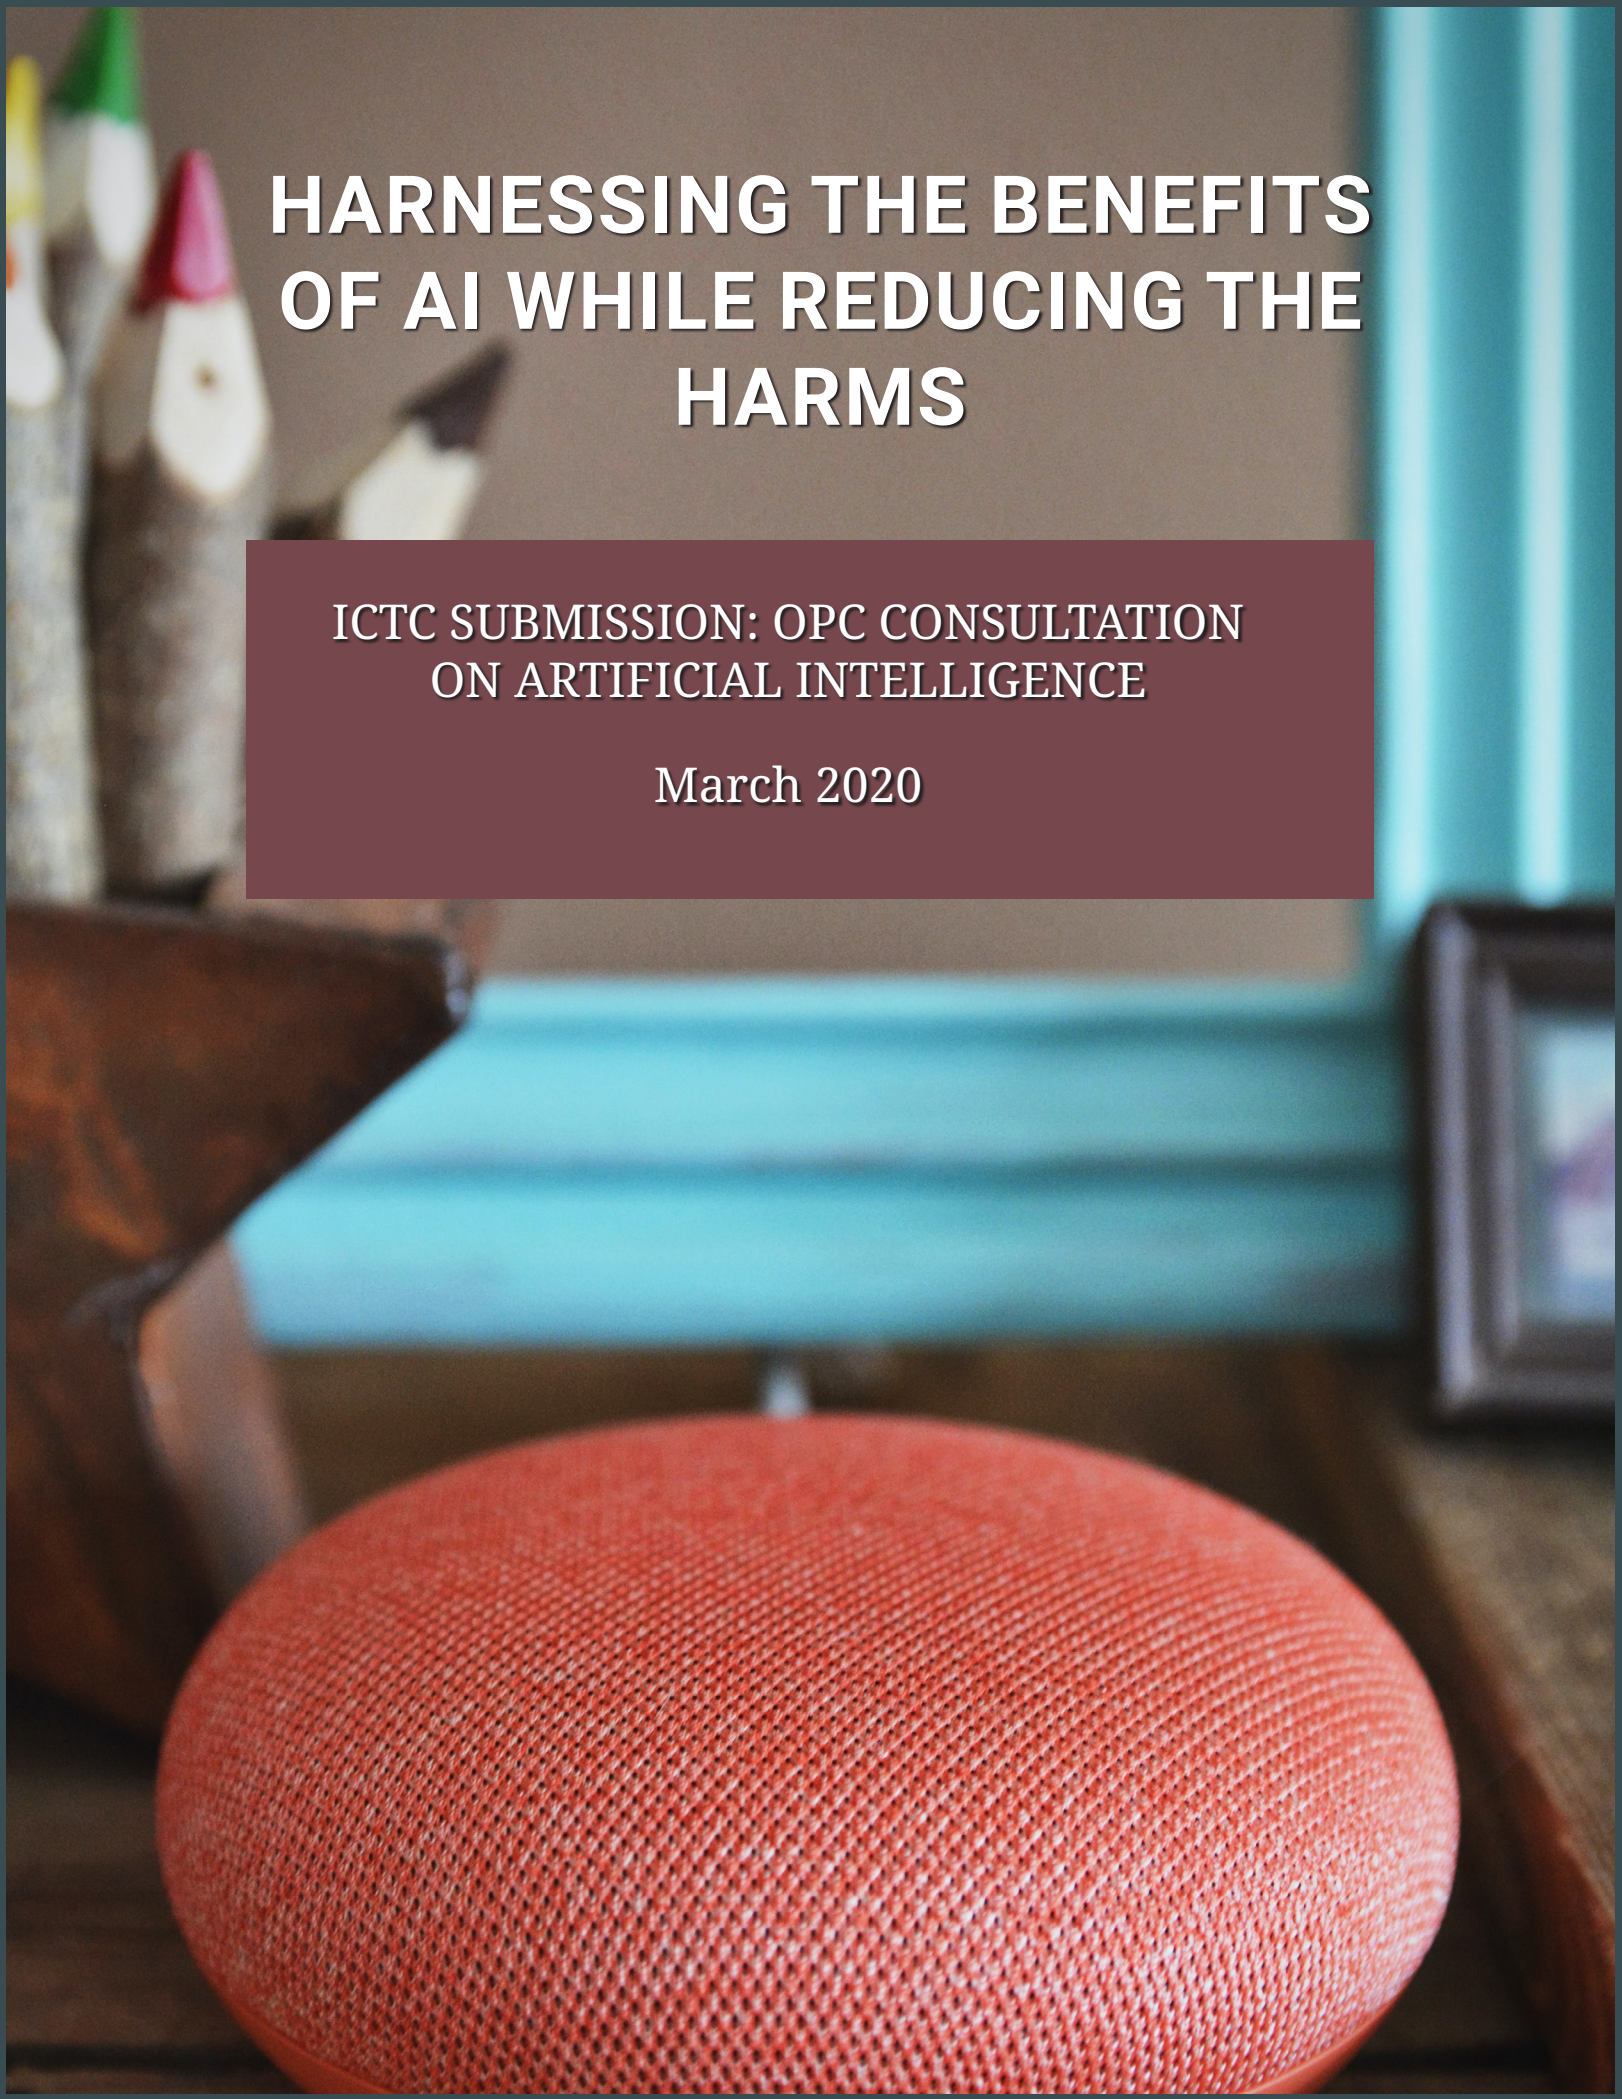

HARNESSING THE BENEFITS OF AI WHILE REDUCING THE HARMS

ICTC SUBMISSION: OPC CONSULTATION ON ARTIFICIAL INTELLIGENCE

March 2020

Preface

As a not-for-profit, national center of expertise, ICTC strengthens Canada's digital advantage in a global economy. Through trusted research, practical policy advice, and creative capacity-building programs, ICTC fosters globally competitive Canadian industries enabled by innovative and diverse digital talent. In partnership with a vast network of industry leaders, academic partners, and policy makers from across Canada, ICTC has empowered a robust and inclusive digital economy for over 25 years.

To cite this paper:

Harnessing the Benefits of AI While Reducing the Harms. Information and Communications Technology Council (March 2020). Ottawa, Canada.

Researched and written by Rob Davidson (Manager, Data Analysis and Research), Kiera Schuller (Research and Policy Analyst), and Mairead Matthews (Research and Policy Analyst).

Information and Communications Technology Council
116 Lisgar Street, Suite 300, Ottawa, ON K2P 0C2
(tel) 613-237-8551 (fax) 613-230-3490

Conseil des technologies de l'information et des communications
116, rue Lisgar, pièce 300, Ottawa, ON K2P 0C2
(tél) 613-237-8551 (téléc) 613-230-349

TABLE OF CONTENTS

INTRODUCTION

The Office of the Privacy Commissioner (OPC) is correct that PIPEDA falls short in its application to artificial intelligence (AI).[1] Commercial organizations throughout Canada's vertical industries are introducing AI to replace and/or supplement human decision-making and analysis. At the same time, AI requires vast amounts of personal information to perform well and return promising results. For these reasons, AI is profoundly impacting the way we use personal information, both in terms of our policies and practices, and the types of activities we use personal information for.

Nonetheless it is important to stay prudent in our approach to regulating AI. Overregulation would have serious ramifications for innovation in Canada, limiting the potential benefits AI has to offer, and hampering current efforts to establish Canada as an international leader in AI. Likewise, an inadequate regulatory response would leave individuals without the explicit tools and levers needed to protect themselves and their personal information in the context of AI.

At the very least, ICTC proposes that we must clearly establish the following rights and obligations:

1) A requirement for proactive and responsible disclosure around the use of automated and semi-automated decision-making systems, so that individuals may be aware of and understand the implications associated with the intended use of their data.

2) The right to be informed when subject to automated and semi-automated decision-making.

3) The right to access commercial organizations' policies and practices regarding the use of personal information in automated and semi-automated decision-making.

4) The right to request and access a privacy impact assessment, and a parallel requirement for commercial organizations to conduct privacy impact assessments for certain kinds of automated and semi-automated decision-making systems.

---

1 "***Artificial Intelligence (AI):*** *a multi-disciplinary subject, involving methodologies and techniques from various fundamental disciplines such as mathematics, engineering, natural science, computer science, and linguistics, to name a few. Over the last few decades, AI has evolved into a number of technological areas such as planning, natural language processing, speech processing, machine learning, vision recognition, neural networks and robotics, among others.*"

McLaughlin, Ryan; Quan, Trevor. On the Edge of Tomorrow – Canada's AI Augmented Workforce. Information and Communications Technology Council (December 2019). Ottawa, Canada. https://www.ictc-ctic.ca/wp-content/uploads/2020/02/canadas-ai-workforce-FINAL-ENG-2.24.20.pdf

5) The right to access specific information about automated and semi-automated decision-making systems, such as: the degree of human involvement in decision-making, the degree of decision traceability, and key characteristics of the training data, including potential biases.

6) The right to have personal information forgotten–also known as the right to erasure. This is particularly important given that AI may collect and use inaccurate data, or even create data about individuals based on erroneous or biased algorithms.

As PIPEDA may not be the right venue to conduct all of this work, we must continue to explore other methods to ensure respect for the rule of law, human rights, diversity, and democratic value in the context of AI in Canada.

PROPOSAL 01

Proposal: 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, or should certain rules be limited to AI due to its specific risks to privacy and, consequently, other human rights?

Short form: There should be no mention of specific technologies such as AI in PIPEDA, even if those technologies pose specific risks to privacy or other human rights. Legislative reform should instead seek to maintain a technology-neutral approach by regulating specific activities linked to AI that feature novel forms of data processing, such as autonomous or semi-autonomous decision making.

Discussion:

PIPEDA is technology-neutral

One of PIPEDA's primary strengths is that it is a technology-neutral law. PIPEDA does not regulate specific kinds of technology, but instead specific activities– namely, the collection, use, and disclosure of personal information. This technology-neutral approach has enabled PIPEDA to stay relevant and largely effective throughout numerous revolutionary changes in tech.

Explicitly defining AI in PIPEDA would render Canada's federal privacy legislation partial to technology. The potential harmful outcomes of this decision are well-known:

1. New innovations in technology may render the explicit definition in PIPEDA outdated, impractical, or ineffective.
2. The explicit definition may be too vague, creating confusion for businesses, increasing the regulatory burden associated with other types of data use, and restricting innovation.
3. The explicit definition may be too specific, rendering the regulation less effective and creating loopholes in the law.

A Challenging Definition: AI is a broad and controversial term

Regulating activity rather than AI, specifically, is important because there remains considerable disagreement in the AI community as to what AI actually is. Some individuals subscribe to a very broad conception of AI that includes everything from robotic process automation to computer vision to deep learning algorithms. Others have a much narrower approach, viewing only machine learning or deep learning algorithms as AI. Considering the level of disagreement present among technical AI experts, it may be extremely difficult for regulators to establish a legal definition for AI that is both functional and accurate.

Building upon this, it is also important to note that not all AI has the same type and/or levels of impact. Depending on the type of AI and context of its use, there are diverse range of potential outcomes and impacts. This is true even for those types of AI that use personal information. Consider, for example, the following diverse applications of AI involving personal information:

1. AI that uses client information to assist financial advisors in giving financial advice.
2. AI that uses customer information to facilitate video recommendations on a streaming platform.
3. AI that uses personal information to assist judges in making court-related decisions.
4. AI that uses customer information to suggest additional products at checkout on an e-commerce platform.
5. AI that uses biometric information to facilitate facial recognition and identify high-ranking or important guests to hotel employees.

For each of the above applications, there are varying levels of risk and sensitivity, both with respect to the kinds of personal information involved and the potential impacts on individuals. Arguably, AI that recommends videos to viewers or products to customers, should not be subject to the same level of regulation as AI used to inform court decisions or facilitate facial recognition.

For these reasons, it may be more appropriate to leave direct reference to AI out of federal privacy legislation, and to instead regulate certain activities linked to AI that are particularly sensitive, impactful, or high-risk.

Examples of a technology-neutral approach

The approach of regulating activity rather than specific technology can be found in both the Government of Canada's Directive on Automated Decision-Making (GOC Directive) and the European Union's General Data Protection Regulations (GDPR).

First, in the GOC Directive, AI is mentioned only in the preamble of the text to explain the Directive's overall purpose (see below).

> *"The Government of Canada is increasingly looking to utilize artificial intelligence to make, or assist in making, administrative decisions to improve service delivery. The Government is committed to doing so in a manner that is compatible with core administrative law principles such as transparency, accountability, legality, and procedural fairness. Understanding that this technology is changing rapidly, this Directive will continue to evolve to ensure that it remains relevant."*[2]

From that point on, the GOC Directive refers only to "Automated Decision Systems" in place of AI. Automated Decisions Systems are defined as:

> *"Any technology that either assists or replaces the judgement of human decision-makers. These systems draw from fields like statistics, linguistics, and computer science, and use techniques such as rules-based systems, regression, predictive analysis, machine learning, deep learning, and neural nets."*[3]

Second, the GDPR also makes no explicit mention of AI, referring in text only to Automated Decision-Making and Profiling. Automated Decision-Making is defined as:

> *"A decision which is made following the processing of personal data that has been conducted solely by automatic means, where no humans are involved in the decision-making process,"*[4]

while Automated profiling is defined as:

> *"any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements."* [5]

2 Directive on Automated Decision-Making. TBS, 2019. https://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=32592#appA
3 Directive on Automated Decision-Making. TBS, 2019. https://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=32592#appA
4 General Data Protection Regulation. Eur-Lex, 2016. https://eur-lex.europa.eu/eli/reg/2016/679/oj
5 General Data Protection Regulation. Eur-Lex, 2016. https://eur-lex.europa.eu/eli/reg/2016/679/oj

Despite minor differences in wording, all three of the above definitions focus not on AI as a technology, but on the activity/process of automated or semi-automated decision-making, in addition to the potential impacts on the individuals involved. The GOC Directive governs the *use* of Automated Decision Systems *in* decision-making, whether entirely automated or human-assisted, while the GDPR governs the *act* of automated decision-making or profiling itself.

According to S.3 of PIPEDA, its purpose is to govern the collection, use, and disclosure of personal information by commercial organizations.[6] Defined as two possible types of data use, automated and semi-automated decision-making may thus find a functional and appropriate regulatory home in PIPEDA.

2. If the latter, how should we define AI?

As discussed, there should be no mention of specific technologies such as AI in PIPEDA. Legislative reform should instead seek to maintain a technology-neutral approach by regulating specific activities linked to AI that feature novel forms of data processing, such as autonomous or semi-autonomous decision making.

6 Government of Canada, 2020 https://laws-lois.justice.gc.ca/ENG/ACTS/P-8.6/page-1.html#h-416931

PROPOSAL 04

Proposal: 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?

Discussion:

The Right to be Informed and an Obligation to Inform

First, this response assumes pre-existing support for an individual's right to be informed when their personal information is used to facilitate either (1) automated or semi-automated decision-making or (2) the training of an algorithm. Without such information, individuals will be hindered from knowing when (or how) to exercise the right to seek an explanation, how to look out for human rights infringements, or how to seek to correct information.

In addition to granting individuals a right to be informed, commercial organizations should be *obligated* to inform individuals when their personal information will be subject to automated processing. Similarly, all practices and policies regarding the use of personal information in automated and semi-automated decisions and to train algorithms should abide by a principle of openness and be fully disclosed.

The Right to an Explanation

Individuals should be granted the right to an explanation when subject to automated or semi-automated decision-making, particularly where the potential outcomes of the decision may be impactful, or where highly sensitive personal information is used. The right to an explanation could depend on factors including:

1. The level of sensitivity of the personal information being used
2. The type of decision being made; or
3. The severity of potential outcomes for the individual, should a decision be made.

Considering the right to an explanation

When it comes to implementing the right to an explanation, regulators should strive to establish clear, practical requirements for commercial organizations. Established requirements should not be vague and confusing, nor impossible to fulfill.

"Explainable AI" is a controversial and often misleading topic; despite many attempts to define "explainable AI," the AI community still lacks a common and agreed upon definition. Critically, not all types of AI lend themselves kindly to clear explanation. Furthermore, not all commercial organizations will be able to explain exactly why, how, and according to what variables their AI models make decisions. At the very least, however, commercial organizations should be required to explain key attributes of their decision-making systems. These include:

1. The degree of human involvement in decision-making. Was the decision semi- or fully-automated? If there was a human involved, what was their role, and how did they rely on the decision-making system to come to their final decision? What other factors did they consider?
2. The degree of decision traceability. How transparent is the model? Can it be explained? Are you able to trace exactly why, how, and according to what variables it makes decisions?

Key characteristics of the training data, including potential biases. How diverse was the data used to train the AI model? Are there any potential biases resulting from the training data? What are those biases, and how might they affect the current decision?

Traceability and Audits:

At a minimum, elements such as data training, algorithms, resulting decisions, and their relationships must be logged and persisted so that robust and effective audits can be performed. Without linking decisions and algorithms to the data used to train models, it is near impossible to provide any transparency and explanation.

Commercial organizations should also be required to explain how their automated-decision systems work, or in other words, the variables and weights involved, and the underlying logic or rational.

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

Existing traditional transparency measures such as audits and regular enforcement actions by regulators may be applicable to the AI context. However, they may not provide the sufficient clarity, emphasis or focus on unique specificities of AI that would be required in order to enable the substantive manifestation of these rights. New transparency measures that more narrowly or robustly address specificities and concerns related to AI may provide value by improving privacy protections and providing for more explicit accountability or provide value through pure signalling. However, additional transparency measures will also produce regulatory burdens for commercial organizations and governments alike, thus all new measures should be assessed on their potential benefits as well as potential resulting regulatory burdens.

Two possible measures are discussed below, to illustrate the potential benefits as well as drawbacks of specific additional measures.

A) Privacy impact assessments

Privacy impact assessments build upon existing privacy protections by putting the onus to assess privacy impacts on commercial organizations, rather than individuals. Under the current system, this responsibility falls on individuals, as they are the party responsible for reading privacy policies and providing their consent. This has proven overburdensome and impractical in today's complex privacy landscape, where individuals are likely to have tens, if not hundreds, of privacy policies to read each day. As individuals become less able and less likely to read privacy policies before providing consent, their consent is in turn rendered increasingly less meaningful.

As a matter of best practice, privacy impact assessments should be clear, easy to understand, easy to access, and publicly accessible. This could be achieved through universal standards governing the layout and content of privacy impact assessments, similar to way nutrition labels are governed today. Similarly, it could be beneficial to establish clear, universal rules regarding their accessibility and dissemination.

To reduce the overall compliance burden for commercial organizations, and regulatory burden for government, privacy impact assessments could be required only in some instances of data use, perhaps dependent on the sensitivity of personal information being used, type of decision being made, or severity of potential outcomes for the individual.

B) Public Filings

Mandating public filings for AI, in contrast, would likely result in overregulation and bureaucratic burden for both organizations and governments. Moreover, public filings would be ineffective in the context of continuous-learning or adaptive algorithms, where the algorithm's functionality, output, and associated risks actually change with each new use.[7] The desired results could be better achieved through the previously discussed right to be informed, obligation to inform, and obligation to conduct a privacy impact assessment.

7 See FDA discussion paper on the proposed regulatory framework for modifications to AI/ML as a medical device. FDA, 2019. https://www.fda.gov/media/122535/download

PROPOSAL 05

Proposal: 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?**

Discussion:

In this rapidly evolving industrial, technological and political landscape, governments - including the Canadian government – are seeking to identify methods by which to adequately regulate artificial intelligence (AI), data practices and related activities. Within this, global privacy discussions are at the forefront. In Canada, it has been widely recognized that existing privacy legislation is inadequate to provide the Office of the Privacy Commissioner of Canada (OPC) with the tools to either meaningfully enforce existing privacy laws in this new context, or to fully address the new, emerging issues arising in this context, particularly around AI.

As such, PIPEDA requires an upgrade with a clear focus on regulating and addressing AI, from each the angles of privacy, human rights and democracy. This response will focus on "Privacy by Design," which can be considered one way of enhancing this capacity. However, there are several elements to consider when evaluating whether Privacy by Design would potentially be a useful legal requirement under PIPEDA.

Central Aim:

The central aim of mandating "Privacy by Design" or by Default is to give data subjects more agency, decision-making power and control "over the collection and use of their personal data."[8] By mandating privacy-protecting considerations as a default part of early design stages of product development or strategy, this approach seeks to give individuals stronger control from the beginning over the types, amount and use of their gathered data. There is already an explicit legal obligation for privacy under the GDPR in Europe, which underscores the potential significance of privacy as a core legal and human rights-focused concept "in Europe and, by extension, Canada."[9] Mandating Privacy by Design (by Default)

8 Information and Privacy Commissioner of Ontario, Privacy by Design, The 7 Foundational Principles, January 2011.
9 Krebs, 2018. https://www.mondaq.com/canada/Privacy/753378/Implementing-Privacy-By-Design

is also, functionally, a mandate of Human Rights by design, given that privacy is increasingly recognized as a fundamental human right in several jurisdictions.

The Value of Design-Thinking:

Privacy, while long a historical social norm, has undergone transformations in its legal and social standing in recent years. On the one hand, there has been an increasing focus on "design-thinking," a conceptual framework and worldview that focuses on proactive design to pre-emptively overcome problems and constraints through holistic, interdisciplinary, integrative, and innovative design. Privacy can be approached through a design-thinking perspective, which seeks to incorporate privacy "into networked data systems and technologies by default," and render privacy "integral to organizational priorities, project objectives, design processes, and planning operations."

What is Privacy By Design?

"Privacy by design" embodies functional data protection through organizational and technological design. It's core notion is that data protection and privacy can be best achieved when it is integrated as a central principle during the early design and creation/production of a technology.[4] It requires designing any system, product or process "in a manner that protects the privacy rights of individuals, rather than considering the associated privacy implications of a system or process only after deployment."[10] It seeks to "protect personal information by implementing measures proactively and preventively." [11]

Developed in Canada in the 1990s by Ann Cavoukian (then Information and Privacy Commissioner of Ontario), Privacy by Design has seven foundational principles, which the Committee recommends PIPEDA includes where possible.[12]

7 Principles of Privacy by Design:

1. "Proactive not Reactive; Preventative not Remedial: *The goal of privacy by design is to take preventative action by implementing measures to reduce the risk of privacy infractions*."[13]

The approach is characterized by proactive rather than reactive measures, which seek to anticipate and prevent privacy problems before they happen. Rather than waiting for privacy risks to materialize or offering remedies for resolving privacy

10 Krebs, 2018. https://www.mondaq.com/canada/Privacy/753378/Implementing-Privacy-By-Design
11 House of Commons, 2018. https://www.ourcommons.ca/Content/Committee/421/ETHI/Reports/RP9690701/ethirp12/ethirp12-e.pdf
12 Ibid.
13 Information and Privacy Commissioner of Ontario, Privacy by Design, The 7 Foundational Principles, January 2011.

infractions after they have happened, this approach seeks to prevent them from occurring; for example, preventing data breaches from happening in the first place. As such, it comes "before-the-fact," rather than "after-the-fact." It can be applied to "information technologies, organizational practices, physical design, or networked information ecosystems," and "begins with an explicit recognition of the value and benefits of proactively adopting strong privacy practices, early and consistently." This principle entails:

- A clear commitment to establish, uphold and enforce high standards of privacy – often of higher standard than those established by global laws and regulation,
- A commitment to privacy that is "demonstrably shared throughout by user communities and stakeholders,"
- Defined methods to recognize and anticipate poor privacy designs, practices and outcomes, and to correct any negative impacts before they occur through "proactive, systematic, and innovative efforts."[14]

2. Privacy as the Default Setting: "*The default setting for all products and services should be to protect personal information so that an individual's privacy is automatically protected without any action being required by the individual*."[15]

The approach aims to provide the most privacy by "ensuring that personal data are *automatically* protected in any business practice or IT-related setting. No action should be required on the part of the individual to protect their privacy – it is built into the system, by default." This principle could entail elements such as:

- *"Purpose Specification – the purposes for which personal information is collected, used, retained and disclosed shall be communicated to the individual (data subject) at or before the time the information is collected. Specified purposes should be clear, limited and relevant to the circumstances.*
- *Collection Limitation – the collection of personal information must be fair, lawful and limited to that which is necessary for the specified purposes.*
- *Data Minimization – the collection of personally identifiable information should be kept to a strict minimum" and "wherever possible, identifiability, observability, and linkability of personal information should be minimized.*
- *Use, Retention, and Disclosure Limitation – the use, retention, and disclosure of personal information [should] be limited to the relevant purposes identified to the individual, for which he or she has consented, except where otherwise required by law. Personal information shall be retained only as long as necessary to fulfill the stated purposes, and then securely destroyed.*

14 Banks, 2018.https://iapp.org/news/a/legislating-privacy-by-design-in-canada/
15 Information and Privacy Commissioner of Ontario, Privacy by Design, The 7 Foundational Principles, January 2011.

- *Default: Where the need or use of personal information is not clear, there shall be a presumption of privacy and the precautionary principle shall apply: the default settings shall be the most privacy protective."*[16]

3. Privacy Embedded into Design: "*The protection of personal information should be an integral part of information systems and business practices; it should not be an add-on."*[17]

The approach must be incorporated into the design and architecture of IT systems and business practices, with systematic, principled approaches that "use accepted standards and frameworks" and "which are amenable to external reviews and audits. Fair information practices should be applied with equal rigour, at every step in the design and operation. Wherever possible, detailed privacy impact and risk assessments should be carried out and published, clearly documenting the privacy risks and all measures taken to mitigate those risks, including consideration of alternatives and the selection of metrics. The privacy impacts of the resulting technology, operation or information architecture, and their uses, should be demonstrably minimized, and not easily degraded through use, misconfiguration or error."[18]

4. Full Functionality – "Positive-Sum, not Zero-Sum: *Privacy by design should be considered a benefit; there should be no trade-offs with other features to achieve this goal."*[19]

Positive-Sum Privacy by Design aims to accommodate all relevant, legitimate interests and objectives in a positive-sum "win-win" manner. Privacy is often presented in a zero-sum manner, in competition other legitimate interests, objectives, and capabilities, Privacy by Design takes an alternative approach, recognizing legitimate diverse non-privacy objectives and addressing them in a holistic manner. It refuses inaccurate and fake dichotomies such as privacy versus security; instead, it emphasizes that it is possible and desirable to achieve both. To this end, all considered "interests and objectives should be clearly documented," standard "metrics agreed upon and applied, and trade-offs rejected as often being unnecessary, in favour of finding a solution that enables multi-functionality."[20]

16 Banks, 2018. https://iapp.org/news/a/legislating-privacy-by-design-in-canada/
17 Information and Privacy Commissioner of Ontario, Privacy by Design, The 7 Foundational Principles, January 2011.
18 Banks, 2018. https://iapp.org/news/a/legislating-privacy-by-design-in-canada/
19 Information and Privacy Commissioner of Ontario, Privacy by Design, The 7 Foundational Principles, January 2011.
20 Banks, 2018. https://iapp.org/news/a/legislating-privacy-by-design-in-canada/

5. End-to-End Security – "Full Lifecycle Protection: *The protection of personal information must extend throughout the system's entire life cycle."*[21]

Full lifecycle protection ensures security throughout the entire duration of the lifecycle of data collected, including collection, storage, and destruction in a timely manner. To ensure end-to-end, cradle to grave, secure lifecycle management of information and data, there cannot be gaps in protection or accountability. Security is particularly important, as without robust security, there can be no privacy. Organizations must comply with relevant standards established by "recognized standard-development bodies," and assure "the confidentiality, integrity and availability of personal data throughout its lifecycle including, inter alia, methods of secure destruction, appropriate encryption, and strong access control and logging methods."

6. Visibility and Transparency – "Keep it Open: *Transparency is important to ensure that systems and practices are truly able to protect user privacy; independent verification must always be possible*".[22]

The approach aims to assure all stakeholders that the business practice or technology involved is "operating according to the stated promises and objectives." All decisions and operations should remain visible and transparent, to all users and providers alike, and should also be subject to trusted independent verification. This aligns with general fair information practices, but for the purpose of *robust auditing, particular emphasis could be placed on:*

- *"Accountability – The collection of personal information entails a duty of care for its protection. Responsibility for all privacy-related policies and procedures shall be documented and communicated as appropriate, and assigned to a specified individual. When transferring personal information to third parties, equivalent privacy protection through contractual or other means shall be secured.*
- *Openness – Openness and transparency are key to accountability. Information about the policies and practices relating to the management of personal information shall be made readily available to individuals.*
- *Compliance – Complaint and redress mechanisms should be established, and information communicated about them to individuals, including how to access the next level of appeal. Necessary steps to monitor, evaluate, and verify compliance with privacy policies and procedures should be taken."*[23]

21 Information and Privacy Commissioner of Ontario, Privacy by Design, The 7 Foundational Principles, January 2011.
22 Information and Privacy Commissioner of Ontario, Privacy by Design, The 7 Foundational Principles, January 2011.
23 Banks, 2018. https://iapp.org/news/a/legislating-privacy-by-design-in-canada/

7. Respect for User Privacy – "Keep it User-Centric: *Above all, privacy by design entails putting individuals' interests first.204 In the EU, the principles of data protection by design have been written into Article 25 of the GDPR.*[24]"[25]

Privacy by Design requires that respect for the interests of the individual is a top priority. This is achieved through the provision of measures such as strong privacy defaults, proper notice, and supporting diverse user-friendly options. It is recognized that empowering and enabling data subjects to have agency and an active role in the managing their own data is one of the most effective checks against the misuse of personal data and abuse of privacy. Respect for User Privacy involves the following elements:

- "*Consent – The individual's free and specific consent is required for the collection, use or disclosure of personal information, except where otherwise permitted by law. The greater the sensitivity of the data, the clearer and more specific the quality of the consent required. Consent may be withdrawn at a later date.*
- *Accuracy – personal information shall be as accurate, complete, and up-to-date as is necessary to fulfill the specified purposes.*
- *Access – Individuals shall be provided access to their personal information and informed of its uses and disclosures. Individuals shall be able to challenge the accuracy and completeness of the information and have it amended as appropriate.*
- *Compliance – Organizations must establish complaint and redress mechanisms, and communicate information about them to the public, including how to access the next level of appeal. Respect for User Privacy goes beyond these FIPs, and extends to the need for human-machine interfaces to be human-centered, user-centric and user-friendly so that informed privacy decisions may be reliably exercised. Similarly, business operations and physical architectures should also demonstrate the same degree of consideration for the individual, who should feature prominently at the centre of operations involving collections of personal data.*"[26]

<u>Accountability: the key to embedding privacy in the design process</u>
Key to ensuring that organizations and data controllers substantively incorporate

24 General Data Protection Regulation, Reg (EU) 2016/679, article 25; see also paragraph 78 of the preamble
25 House of Commons, 2018. https://www.ourcommons.ca/Content/Committee/421/ETHI/Reports/RP9690701/ethirp12/ethirp12-e.pdf See more at: "Privacy by Design - The 7 Foundational Principle."s https://iapp.org/resources/article/privacy-by-design-the-7-foundational-principles/ "This document, authored by former Information and Privacy Commissioner of Ontario Ann Cavoukian, provides readers with additional information, clarification and guidance on applying the seven foundational principles of privacy by design. This guidance is intended to serve as a reference framework and may be used for developing more detailed criteria for application and audit/verification purposes."
26 Banks, 2018. https://iapp.org/news/a/legislating-privacy-by-design-in-canada/

privacy considerations early in their design processes is accountability. Europe's GDPR, as an example, not only ensures organizations are responsible for complying with privacy principles, but also that they must be able to *demonstrate* their compliance. Because of this, Deloitte has posited, organizations under this framework will benefit most by having an active transparent, accountable privacy strategy.[27] This includes making choices early in development about how privacy is handled and executed within the legal boundaries, including conducting an initial "Privacy Impact Assessment" (PIA)[28] which identifies privacy risks in a new design, and provide rationale for an approach to privacy protection. A legal principle of Privacy by Design would push organizations to consider ethical aspects of the data collection – how much data is necessary and why, why types/limits of data and why, and how to ensure data transparency. Finally, it would force organizations to consider the issues of data security, quality, protection and retirement. For example, adequate security measures in protect data, how to ensure the quality of data, and what happens to the data once "the product or service retires."[29]

Mutual Benefits: Advantages for businesses and data controllers

Privacy by Design not only provides a way to ensure privacy protection, but also offers advantages for organizations, businesses and data controllers. Some mutual advantages, for example, include enhancing business efficiency and promoting trust among data subjects. Under the existing Directive in Canada, all data controllers are already required "to implement appropriate technical and organisational measures to protect data against unlawful processing"; however, currently, this can be done at the very end of product's development. In contrast, if privacy considerations are required at the earliest stage of development, as a key ingredient of a product or service rather than an element considered at the end, this would be much more efficient for businesses to comply rather than seeking to apply privacy considerations after a design has been fully developed.[30] By "considering the legitimate needs, purpose/use and collection of personal data prior" to starting, businesses reduce the chance that at a later stage, they discover that "embedding the required level of privacy protection within the technology is "challenging, expensive of even impossible."[31] Identifying early what data is needed, how to legitimately collect it, and how to give data subjects the choice of how their data is used also enables better *transparency,* which is crucial for developing trust to gather data in the first place, thereby making the process of technological development more efficient and effective.

27 Danon, 2020. https://www2.deloitte.com/ch/en/pages/risk/articles/gdpr-privacy-by-design-and-by-default.htm

28 Ibid.

29 Ibid.

30 Ibid.

31 Ibid.

A Model: Article 25 of the GDPR

Article 25 of the GDPR, meaningfully titled "Data Protection by Design and by Default," provides a useful framework that can be applied to AI systems in Canada. Article 25 holds that organizations must "consider privacy at the initial design stages and throughout the complete development process of new products, processes or services that involve processing personal data."[32] Article 25 of the GDPR engages with numerous elements of an obligation of "Data protection by design," such as implementing "appropriate technical and organizational measures designed to implement the data protection principles" as well as to "safeguard individual rights and freedoms." Article 25 also notes that "an approved certification mechanism" may be used to demonstrate compliance." Article 25 of the GDPR specifically notes three parts: first, it takes into account "the state of the art, the cost of implementation and the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for rights and freedoms of natural persons posed by the processing," the obligation holder must "both at the time of the determination of the means for processing and at the time of the processing itself, implement appropriate technical and organisational measures […] which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects."[33] Examples of such measures include, for instance pseudonymization. Second, the obligation holder must "implement appropriate technical and organisational measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are processed." This "applies to the amount of personal data collected, the extent of their processing, the period of their storage and their accessibility."[34] Moreover, "such measures shall ensure that by default personal data are not made accessible without the individual's intervention to an indefinite number of natural persons."[35] Finally, "an approved certification mechanism pursuant to Article 42 may be used as an element to demonstrate compliance with the requirements."[36]

As the European Data Protection Supervisor, Giovanni Buttarelli, stated: "privacy by design and by default are no longer mere recommendations, but are now legal and clear obligations for all data controllers. Beyond obligations, there is now a system to ensure that "designers, producers and developers comply with obligations and design systems to be "less invasive and more user-friendly."[37] Article 25 is, however,

32 Ibid.
33 Intersoft Consulting, 2020. https://gdpr-info.eu/issues/privacy-by-design/
34 Ibid.
35 Ibid.
36 Ibid.
37 House of Commons, 2018 https://www.ourcommons.ca/Content/Committee/421/ETHI/Reports/RP9690701/ethirp12/ethirp12-e.pdf. See: ETHI, Evidence, 1st Session, 42nd Parliament, 13 June 2017, 1240 (Giovanni Buttarelli, Supervisor, European Data Control Supervisor).

qualified by a "risk-based approach and reasonableness standard," meaning that "the more sensitive the information or the higher the risk to rights of individuals, the greater the obligation on the data controller to take measures to protect that data and to show this was considered and effected at the time of design."[38] It is also limited in application only to data *controllers* who determine the means and processing of personally identifiable information (and, to some degree, processors), not to the manufacturers of the technology. Yet, in practice, this "places an indirect or commercial obligation on manufacturers of technology, including Canadian organizations who supply technology to others subject to the GDPR," as controllers will be more inclined to choose suppliers that will ensure they comply with the law.[39]

Article 25 of the GDPR has already had an influence on Canada. The recent Report of the Standing Committee on Access to Information, Privacy and Ethics in Canada recommended that Privacy by Design be made an "explicit part of Canadian privacy law" and a "central principle" of PIPEDA. The Committee determined that the "integrated, proactive approach" of Privacy by Design was "an effective way to protect the privacy and reputation of Canadians," ensuring that privacy considerations are taken into account at all stages of development, including the design, marketing and retirement of a product."[40]

Furthermore, "the practical consequences for data controllers and manufacturers of technology" will quickly evolve, as the enforcement of Privacy by Design as "a legal obligation in Europe" will compel "Canadian companies with operations or customers in Europe" to be aware of European "legal obligations" as well as "related requirements from a commercial and reputational perspective."[41] It is highly possible that Privacy by Design will also play a role in determining whether Canada's privacy regime is "considered 'adequate' under European data transfer rules."

Challenges and Limitations

While Privacy by Design is considered by many to be a means of meaningfully preserving privacy rights, the approach also has important limitations. Foremost, it is unlikely to be sufficient on its own. First, a critical stance could argue that this is simply a weaker, less-enforceable version of Proposal 2, as a rights-based approach but with less legal "teeth." Indeed, despite the legal component, the obligation of Privacy by Design ultimately lies in complying with a standard rather than a law. As such, there is potential for it to be masked by jargon and used merely for signalling and "ethic-washing" activities, which give the impression of active

38 Intersoft Consulting, 2020. https://gdpr-info.eu/issues/privacy-by-design/
39 Krebs, 2018. https://www.mondaq.com/canada/Privacy/753378/Implementing-Privacy-By-Design
40 House of Commons, 2018. https://www.ourcommons.ca/Content/Committee/421/ETHI/Reports/RP9690701/ethirp12/ethirp12-e.pdf
41 Krebs, 2018. https://www.mondaq.com/canada/Privacy/753378/Implementing-Privacy-By-Design

privacy protection where there is in fact a lack. Such activities may, in fact, further protect actors from future liability, rather than exposing them to it, for even if data controllers can claim – or even provide evidence – that time and effort was spent considering and integrating privacy protection in their product or service, that does not ensure the product or service does so effectively.[42]

Considering the above, the core challenge lies in being able to test or evaluate the efforts or process by which organizations implement Privacy by Design. Can there be a way to evaluate and assess the adequacy of such efforts, particularly taking into account the potential inability for design to anticipate all problems related to privacy? Additionally, is there potential that creating such standards or tests may create bureaucracy that prevents smaller companies from being able to participate, despite seeking the same levels of Privacy by Design as larger companies? An additional danger would be producing impact assessments or adequacy tests that become a mechanism by which large or resourceful actors can take to their advantage, by producing evidence of benign or positive intentions or considerations in their designs, which may or may not reflect reality, where small companies can not.[43]

Other considerations would include: Would an approved certification mechanism be possible to demonstrate compliance with requirements?

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?**

It would be challenging 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 Canadian market. Such an obligation would likely produce additional bureaucratic burden on both government and business, and could have potentially negative impacts on Canada's markets. While not unfeasible, critical elements to consider include: the impacts of the obligation on manufacturers seeking to access the Canadian market, the bureaucratic burden on both business and governments, and the impacts on the market, supply and demand, and prices in Canada.

<u>Factors to Consider</u>

Among others, several key factors are:

- What tests/procedures would be accepted as standard, by what accreditation process?

42 Hirsh, 2020. https://metaviews.substack.com/p/call-to-action-regulation-of-ai
43 Hirsh, 2020. https://metaviews.substack.com/p/call-to-action-regulation-of-ai

- Costs of establishing and enforcing the standard
- Impact on foreign manufacturers seeking to access Canada's market: effect of reducing interest in Canadian market and adding burden to business:
    - Ability for manufacturers to access and conduct tests (financial limitations, adequate resources)
    - Differential impacts on small vs. large companies
- Impacts on Canadian market prices and access to innovation
- Fairness between foreign and domestic manufacturers

<u>One Model: EU 'Product Requirements' and Safety Standards for Market Entry</u>

Arguably, the EU 'Product Requirements' standards, which sets standards for access to the European market, provide a comparable example. The EU has established rigorous "product safety standards, health and environmental standards, sector-specific standards, standards in international trade."[44] These include based EU product standards, chemical safety standards, medicinal product safety standards, food safety standards, and basic safety requirements for goods in the EU market. To take the example of Food, the EU provides a 'General Food Law' which contains principles, requirements, procedures for good law, as well as labelling and packaging rules and fitness checks, to ensure the quality of food within the EU. For imported and exported goods, the EU sets technical requirements for goods entering the EU, including health, safety and environmental requirements which, for example, tackle plant, animal and public health threats. Customs implement and enforce a wide range of legislations which, for example, protect the environment and public health, and check products that enter or leave the EU, controlling to "prevent unsafe or non-compliant products from entering the EU market."[45] Moreover, specific requirements apply for particular products, such as "medicines for human use and organic products."[46] These have not significantly harmed the European markets, and created negative incentives.

<u>Text vs. Practice:</u>

It would be possible to achieve the desired outcomes, in practice, through less direct methods. For example, Article 25 of the GDPR is distinctly limited in application only to data controllers, who determine the means and processing of personally identifiable information (and, to some degree, processors); it is not designed to apply to the manufacturers of the technology. Yet, in practice, Article 25 does arguably place an indirect obligation on technology manufacturers to assess privacy and human rights impacts of their products, for data controllers will be more inclined to choose suppliers that will enable them to fully comply with the law.[47] This applies to "Canadian organizations who supply technology to others

44 EU, 2020. https://ec.europa.eu/info/business-economy-euro/product-safety-and-requirements/eu-product-requirements_en
45 Ibid.
46 Ibid.
47 Krebs, 2018. https://www.mondaq.com/canada/Privacy/753378/Implementing-Privacy-By-Design

subject to the GDPR."[48] This could occur in reverse, with foreign and international manufacturers seeking to supply technology and products to Canada. As such, legislating such a requirement as suggested per above may be unnecessary to achieve the desired outcome in practice.

Conclusion: Considering the above, this proposal would not be recommended.

48 Krebs, 2018. https://www.mondaq.com/canada/Privacy/753378/Implementing-Privacy-By-Design

PROPOSAL 06

Proposal: Make compliance with purpose specification and data minimization principles in the AI context both realistic and effective

1. **Can the legal principles of purpose specification and data minimization work in an AI context and be designed for at the outset?**

The legal principles of purpose specification and data minimization are in direct conflict with the underlying goals and ideology behind big data and AI. To be functional in the context of AI, data minimization and purpose specification would have to exist as some sort of default standard, allowing for exceptions based on either consent or alternative grounds for processing beyond consent. For example, an individual could consent to:

A. Their personal information being stored and used for other purposes similar to that of the original use – e.g. a similar rationale for using the data, similar end goal, or similar context.

B. Their personal information being used for other specific purposes – e.g. a specific kind of health research, or environmental research.

C. Their personal information being used for other purposes, so long as a specific set of pre-requisites are fulfilled – e.g. their personal data being turned into synthetic data with a low risk of re-identification.

Alternative grounds for processing beyond consent are discussed in the following response.

2. **If yes, would doing so limit potential societal benefits to be gained from using AI?**

Unconstrained requirements for purpose specification and data minimization in the context of AI would likely limit the potential societal benefits to be gained from using AI. With the appropriate default standards and exceptions however, there may be a way to satisfy both objectives.

PROPOSAL 07

Proposal: 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?**

Obtaining meaningful consent must be the default before "consentless" alternatives are permitted. Such a law will need to articulate the regulatory requirements–for example, public filing, individual notice, or justification–organizations will have to meet before processing data outside of meaningful consent. Cost constraints cannot be an adequate justification for bypassing consent, organizations should be required to prove the collective benefit to society and the individual outweigh the need for meaningful consent.

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?**

Consent/no consent is a false dichotomy, there are other meaningful consent models that can be explored beyond direct individual consent. In personal finance, people can employ accountants to represent them for tax and finance preparations. For consent, there is an opportunity to create a personal data protection agent role, that can work on behalf of people. To ensure satisfactory protection of personal data, these agents should be:

- Industry certified, similar to Chartered Professional Accountants;
- Required to stay up to date with all data privacy and consent developments;
- Able to provide individuals with sound advice regarding consent and meaningful consent;
- Able to act as a proxy for individuals to grant and revoke consent.

3. **Requiring consent implies organizations can define purposes for which they intend to use data with sufficient precision for the consent to be meaningful. Are the various purposes inherent in AI processing sufficiently knowable so that they can be clearly explained to an individual at the time**

of collection in order for meaningful consent to be obtained?

If the purposes of the AI processing cannot be explained with sufficient precision, then the request to do so by the organization should be declined by regulators.

4. **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?**

As discussed, meaningful consent must be the default for all situations. As such, PIPEDA reform must focus on promoting and fostering increased personal privacy, and not increased organizational entitlement to use AI.

5. **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)?**

The first step should be to assess whether protected personal information is necessary for the AI purposes. Synthetic data and differential privacy processes should be mandatory consideration before exploring the need for processing protected personal information. Then and only then are grounds like GDPR's legitimate interests considered.

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

Incentives would be an interesting policy tool to encourage organizations to obtain meaningful consent and use synthetic data and differential privacy properly.

PROPOSAL 08

Proposal: 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 and other privacy techniques like synthetic data need to be looked at from the perspective of risk management. These techniques are valid, but they are simply elements of larger risk management plans, and not solutions unto themselves. Risk management plans include detailed scenarios that identify the likelihood of reidentification for individuals and groups of individuals, in addition to the probable level of harm if data is reidentified. Importantly, risk management plans need to be transparent and accessible so that individuals can properly assess the risk of reidentification and meaningfully object to processing by the underlying AI.

2. **Which PIPEDA principles would be subject to exceptions or relaxation?**

Any exceptions or relaxation of the rules must be dependent on the likelihood of re-identification and the significance of potential harm if reidentification occurs. For exceptions or relaxation to occur, both the likelihood of re-identification and the potential harm of re-identification should be low. Under a low likelihood and harm scenario, reduced or eliminated fines for re-identification breaches could be considered, with response focused on mitigation and correction rather than financial penalties.

3. **What could be enhanced measures under a reformed Act to prevent re-identification?**

Similar to the EU model, a reformed PIPEDA ACT should incorporate the need for a Data Protection Officer. This officer should be certified in the proper use and monitoring of emerging industry-standard de-identification techniques, such as synthetic data and differential privacy.

PROPOSAL 09

Proposal: 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?**

We should, at a minimum, require AI decisions to be traceable to specific instances of algorithmic models and training data. This information should be kept and stored on an immutable database and should be detailed enough to enable past decisions to be recreated as part of the audit or review process.[49]

49 Rob Davidson: https://www.linkedin.com/pulse/framework-components-ai-governance-rob-davidson-/

PROPOSAL 10

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?**

If robustly designed and subject to constant review, the proposed enhanced accountability measures would provide a good foundation for enforcement. As mentioned in response to Proposal 8, requiring commercial organizations to appoint a Data Protection Officer (or similar) would provide a locus for accountability.

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

See Proposal 9, response #1.

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

As discussed, requiring commercial organizations to appoint a Data Protection Officer (or similar) would provide a locus for accountability. In addition to this, however, it would align regulatory and legal approaches to AI by ensuring humans remain accountable for AI decisions.

PROPOSAL 11

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, however, the enforceable penalties must be significant enough to act as a deterrent against non-compliance with the law.

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

No. The current rush to implement and deploy AI is driven by profit and funding objectives. Without meaningful and enforceable penalties, some organizations may willfully ignore compliance with the law, managing symbolic sanctions as a cost of doing business.