Submission to Consultation on the Office of the Privacy Commissioner of Canada's Proposals for ensuring appropriate regulation of artificial intelligence

---

Prepared by Joyce Yan for the Office of the Privacy Commissioner of Canada

February 20, 2020

BC Freedom of Information and Privacy Association
103 – 1093 West Broadway
Vancouver, British Columbia, V5N 1E2
Phone: 604-739-9788 | Fax: 604-739-9148

FIPA would like to acknowledge the Law Foundation of British Columbia. Their ongoing support of our work in the areas of law reform, research and education makes submissions like this possible.

Introduction

FIPA is a non-partisan, non-profit society that was established in 1991 to promote and defend freedom of information and privacy rights in Canada. Our goal is to empower citizens by increasing their access to information and their control over their own personal information. We serve a wide variety of individuals and organizations through programs of public education, public assistance, research, and law reform.

We thank the Office of the Privacy Commissioner of Canada for this opportunity to discuss issues relating to appropriate regulation of artificial intelligence (AI).

While this submission focuses on informing the OPC's proposals for ensuring appropriate regulation of AI, it also contains discussions that could be applied to regulators more generally, as well as to industry and government.

Summary

The context in which we understand privacy is shifting in the current landscape of big data analytics, machine learning, and artificial intelligence. What was generally considered a broad topic with varying normative understandings is now at the forefront of debates and policy work as the varied stakeholders attempt to narrow its scope in the Canadian context. Simultaneously, we are seeing an increasing and more varied quantity of data being given to and collected by both private and public bodies in the name of technological innovation – most of which are fueled by data, thereby making privacy concerns more acute.

Resulting from this dichotomous relationship is a narrative where privacy is pitted against innovation; where privacy protection is seen as a check on innovation. This is a dangerous discursive framing of the issue because it implies that unless individuals want to remove themselves for this new form of a social contract, they must give up control over their personal information. AI represents a paradigm shift in technology: rather than an incremental expansion of existing methods and practices, we are seeing a revolution of Big Data, which has already had widespread and – in many cases – deleterious implication for privacy rights. The intersection of machine learning and Big Data has the potential to fundamentally alter what it means to have a 'reasonable expectation of privacy." The implications of AI transcend any meaningful distinctions between public and private sector (including privacy laws). Crucially, history has shown us that when new privacy-impacting technologies are adopted ahead of corresponding changes in law and regulations governing privacy, it can be difficult to 'roll back' established practice and undo erosions to privacy rights. Technology moves quickly from being novel and extraordinary to routine and normalized. This calls for a precautionary approach – one that involves restrictions on the adoption and use of privacy-impacting AI until such time that robust and updated legal and regulatory frameworks are in place.

We urge the policy makers to keep PIPEDA (and other legislation) technology-neutral. Privacy remains a broad concept that is not limited to technological contexts and that can also transcend any particular technology. Rather than writing a new legislation targeted specifically at AI, we

maintain that existing privacy laws need to be strengthened significantly to expand their scope to be able to govern AI within the existing legal framework.

We are pleased to see that the Office of the Privacy Commissioner (OPC) is moving towards adopting a human-rights based approach to privacy rights. This framing of privacy no longer sees the interplay between privacy and innovation as zero-sum; rather, it emphasizes the foundational role of trust [through privacy] to support the digital economy. Building on this notion, we emphasize the importance of keeping meaningful consent and transparency central to data privacy and AI governance.

Lastly, we fully support the numerous calls for expanding the Information and Privacy Commissioners' powers to support a robust enforcement regime where they can utilize order-making powers and administrative monetary penalties to enforce compliance. We have been calling for these changes for over a decade and believe that given the move towards more automation and increasing sophistication of data processing methods, this is a requirement *no*w more than ever.

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

We need robust and clear privacy laws to underpin AI and data sharing frameworks to ensure that they are consistent with fundamental values of transparency and ethics. As it stands now, PIPEDA and the current privacy discourse pits innovation against privacy protection where the latter is seen as a check on the former. We agree with the Privacy Commissioner of Canada, who argues that rather than needing to balance a data-driven economy, strong privacy protection is necessary to build "the trust we need to allow the digital economy to flourish, and the social license the government will need from Canadians to innovate with their personal data."[1]

In short, AI should be governed by the same rules as other forms of data processing, meaning that there needs to be a clear legislated definition of "artificial intelligence" and "privacy" incorporated into the federal privacy legislations. That being said, there are principles that need to be enhanced in the context of AI - the sheer volume of data required by AI, the potential for exacerbated social prejudices, and the inherent opacity of the processing are some of the many reasons AI requires additional considerations. While these principles will help render AI governable, they need to be equally applied to all data across all contexts.

Explainability: A distinctive aspect of AI and big data analytics is the way in which it processes massive amounts of data. It involves feeding massive quantities of data through non-linear neural networks that classify the data based on the outputs from each successive layer, this results in a 'black box' effect thereby rendering the ability to explain reasons for decisions nearly impossible.[2] Algorithmic accountability, as defined by the ICO as the ability "to check that the algorithms used and developed by machine learning systems are actually doing what we think they're doing and aren't producing discriminatory, erroneous or unjustified results," is a challenge due to the opaque and autonomous nature inherent to algorithms. This is particularly concerning given the increased potential for discriminatory decisions. AI governance needs to find ways to build discrimination detection into their machine learning systems and support them with robust privacy legislation to prevent such decisions.

Furthermore, understanding explanations behind AI decisions requires the ability to explain and translate complex decision-making rationale into an accessible and clear form or language geared

[1] "Submission: National Digital and Data Consultations," Office of the Privacy Commissioner of Canada submission to Innovation, Science, and Economic Development Canada, last modified December 5, 2018, https://www.priv.gc.ca/en/opc-actions-and-decisions/submissions-to-consultations/sub_ised_181123/

[2] "Big data, artificial intelligence, machine learning and data protection," ICO, last modified April 7, 2017, https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf

towards a lay audience. The ability to create a human-understandable narrative is imperative to AI governance.

Data minimization: A problematic aspect of AI is its need for massive quantities of data and its tendency to collect as much of it as possible. This creates tension with privacy laws, which seeks to limit the collection of personal information to that which is needed for the purposes identified by the organization. A complicating and conflicting factor is that this principle of limitation may introduce bias into algorithms: by limiting certain datasets, the algorithms are only being trained on subsets of the datasets. This emphasizes the importance of building discrimination detection into systems and with strengthened privacy legislation in place.

While we aren't in a position to provide guidance on the actual design and implementation of such systems, we maintain and strongly support limiting the collection and processing of personal information, as well as limiting the length of time data is kept. Here, we emphasize the need for clear restrictions on retention and use of personal data to avoid open-ended retention policies. Organizations need to be able to articulate at the outset whether the data they collect is necessary for the purpose of processing, which needs to include expectations of what is to be learned or achieved by processing that data.[3]

We have provided comment in the past to the OPC on the topic of the right to be forgotten where we cautioned policy and law makers in assigning responsibility for it. We highlighted a number of concerns:[4]

> For one thing, a system that requires content be taken down or obscured from certain major websites, search engines, or social media platforms while ignoring others will exacerbate existing differences between those who know where to look for certain information and those who do not.[5] And as the Electronic Frontier Foundation's Danny O'Brien warns, "Popular search engines will list the best of everyone, and be compelled to disappear other facts. Meanwhile, a new market is created for mining and organizing accurate public data out of the reach of … authorities."
>
> Another consideration is that allowing different websites to set up different processes for requesting or disputing a takedown may create confusion for users, and deter less technologically-literate people from using the system at all.
>
> As well, it is important to insure against erroneous or malicious requests to delete online content. While data shows that 95% of Google privacy requests are from citizens, rather

[3] Ibid., 40.
[4] BC Freedom of Information and Privacy Association, "Submission to Consultation on Online Reputation," April 28, 2016, https://fipa.bc.ca/wordpress/wp-content/uploads/2016/04/BCFIPA_OnlineReputation2016.pdf
[5] Eduardo Bertoni, "The Right to Be Forgotten: An Insult to Latin American History," *The Huffington Post,* September 2014. http://www.huffingtonpost.com/eduardo-bertoni/the-right-to-be-forgotten_b_5870664.html?utm_hp_ref=tw.

than criminals, politicians, or public figures,[6] it is vital that any request be considered carefully to avoid the deletion of information that is in the public interest.

Finally, we wish to highlight that jurisdictional issues may come into play. A Canadian body may not have the power to enforce its rules—such as in the Globe24h case, in which a Romanian website was ordered to take down information, but refused and faced no consequences[7]—or where a Canadian may want to hide information from parties living in another country. Technological solutions such as geoblocking may be available in the former case, but the latter case would require global cooperation. And, as Radsch remarks in her chapter, *Laws, Norms and Block Bots: A Multifaceted Approach to Combatting Online Abuse*, "in many parts of the world, including countries that do already have special mechanisms to address online abuse, law enforcement agencies are not equipped to deal with these complaints, and can even perpetuate the harm by requiring that offending content be further circulated." [8]

Meaningful consent: Consent must only be considered valid if it is truly "meaningful" – if it is reasonable to expect that individuals to whom a business' activities are directed would understand the nature, purpose and consequences of the collection, use or disclosure to which they are consenting. It must be given freely, informed and ongoing, and individuals must have both the right and the means to withdraw consent at any time. Furthermore, we believe that consent can only be legitimate if it is not coerced. In practice, this requires meaningful alternatives available to individuals who do not choose to consent to privacy-impacting practices.

Purpose limitation: The ICO sets out a two-part test before an organization can repurpose data: first, the purpose for which the data is collected must be specified and lawful; second, if the data is further processed for any other purpose, it must not be incompatible with the original purpose.[9] Exceptions to this exist in limited circumstances such as receiving the consent of the data subject, or by the authority of law.

We support the GDPR's position: in assessing compatibility, it is necessary to take account of any link between the original and the new processing, the reasonable expectations of the data subjects, the nature of the data, the consequences of the further processing, and the existence of safeguards.[10]

---

[6] Sylvia Tippmann and Julia Powles, "*Google accidentally reveals data on 'right to be forgotten' requests," The Guardian*, July 2015. http://www.theguardian.com/technology/2015/jul/14/google-accidentally-reveals-right-to-be-forgotten-requests.

[7] "*PIPEDA Report of Findings #2015-002*," Office of the Privacy Commissioner of Canada, last date modified March 23, 2018, https://www.priv.gc.ca/cf-dc/2015/2015_002_0605_e.asp

[8] "New challenges to freedom of expression: Countering online abuse of female journalists," Office of the Representative on Freedom of the Media Organization for Security and Co-operation in Europe, 2016, https://www.osce.org/fom/220411?download=true

[9] ICO, *Big Data*, 38.

[10] EU *General Data Protection Regulation (GDPR):* Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L119/1.

A key factor that underpins any decision is the principle of fairness.

Fairness: Article 5(1)(a) of the GDPR states that personal data must be "processed fairly, lawfully and in an transparent manner in relation to the data subject," which seeks to establish whether or not the processing of data is fair – fair in the sense that the processing must be inherently transparent, considerations of the effects of the processing on individuals, and their expectations as to how their data will be used.

In other words, data must be processed with respect for the data subject's interests, and data controllers must take measures to prevent discriminatory impacts on the individual. Companies and organizations looking to implement AI systems will need to learn how to mitigate biases within algorithmic models.

2. *If certain rules should apply to AI only, how should AI be defined in the law to help clarify the application of such rules?*

We maintain that, like the GDPR, PIPEDA should remain technology-neutral, while addressing the large-scale automated processing of data, which includes the aforementioned principles. By remaining technology-neutral, it ensures that the legislation will remain relevant and pertinent as technological advancements continue to develop. We argue that if the privacy legislation is amended with automated processing in mind, it will undoubtedly ensure that other forms of data processing will also be covered.

Proposal 2: Adopt a rights-based approach in the law, whereby data protection principles are implemented as a means to protect the 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?*

BC FIPA recognizes that a right-based approach to data protection and privacy is two-fold: the first as set out in article 12 of the *Universal Declaration on Human Rights* (UNDHR), which provides that:

> 12. No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation. Everyone has the right to the protection of the law against such interference or attacks.[11]

[11] "Universal Declaration of Human Rights." United Nations. United Nations. Accessed March 13, 2020. https://www.un.org/en/universal-declaration-human-rights/.

Similar provisions are found in the *International Covenant on Civil and Political Rights* (ICCPR),[12] and in the 2013 UN General Assembly's resolution on the *Right to Privacy in the Digital Age*.[13]

Secondly, we recognize that the right to privacy is also an enabling right where its protection is essential in allowing other rights and freedoms to be exercised, such as the freedom from discrimination and inequality. A human-rights based approach to privacy addresses the fundamental right to privacy, as well as the interrelationship between privacy and the ability of individuals to exercise their other rights and freedoms with autonomy and dignity.

Currently PIPEDA does not recognize privacy as a *human right*; rather, it addresses a right of privacy but in the context of balancing the right against the needs of organizations to collect use and disclose personal information.[14] This framing of the right to privacy as a data protection right is overly narrow as it effectively ties privacy rights exclusively to data protection, which is highly concerning.

This is evident in the framing of "personal information" by both the GDPR and PIPEDA. PIPEDA applies only to 'personal information,' which is defined as information about an identifiable individual.[15] This, then, precludes deidentified data; however, in the age of AI and machine learning, the processing of aggregated deidentified data may re-identify individuals. This begs the question of whether deidentified data can remain truly deidentified in an era of mass data collection and growing algorithmic sophistication.

On the other hand, the GDPR defines personal data as:

> …any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.[16]

The latter approach recognizes that data protection is an important aspect of privacy protection, but a human-rights based approach to privacy extends beyond data protection. The GDPR acknowledges, for example, the importance of "respect for private and family life, home and communications, the protection of personal data, freedom of thought, conscience and religion,

---

[12] United Nations. "International Covenant on Civil and Political Rights." *Treaty Series*, vol. 999, p. 171 and vol. 1057, p. 407. Accessed March 13, 2020. https://treaties.un.org/doc/Treaties/1976/03/19760323%2006-17%20AM/Ch_IV_04.pdf

[13] General Assembly Resolution, *Right to Privacy in the Digital Age*, A/RES/68/167 (18 December 2013), http://undocs.org/A/RES/68/167; Supplemental resolutions: General Assembly Resolution, Right to Privacy in the Digital Age, A/RES/69/166 (18 December 2014); General Assembly Resolution, Right to Privacy in the Digital Age, A/RES/71/199 (19 December 2016).

[14] *Personal Information Protection and Electronic Documents Act (PIPEDA), Statutes of Canada* 2000, c.5. s.3.

[15] PIPEDA, s. 2(1)

[16] GDPR, s. 4(1).

freedom of expression and information, freedom to conduct a business, the right to an effective remedy and to a fair trial, and cultural, religious and linguistic diversity."[17]

Practically speaking, the challenges that we foresee are grounded in the transactions costs that will be imposed onto companies and organizations working with AI. To become compliant with a human-rights based approach to privacy, there will be costs associated with compliance and risk assessment tools. With that being said, we emphasize the need to steer clear of the 'privacy as a check on innovation' discourse, and encourage organizations to be privacy-protective from the onset of working with AI. We have heard from organizations that the cost of becoming privacy compliant is a barrier to innovation, and greater privacy protection offered through 'privacy as a human right' will exacerbate those costs. Our response to that is that due consideration (beyond that of 'traditional' data processing) needs to be given when working with AI given its propensity to collect and process immense amounts of data, and more so for its potential to produce discriminatory, erroneous or unjustified results.

We also acknowledge that arguments regarding the innovation and business costs of other fundamental human rights are not taken as seriously or given as much attention in comparison. This is a peculiar feature of the privacy field, and it reflects the fact that developments in technology and policy that impact privacy are, in the contemporary context, often driven by the expansion of what Shoshana Zuboff rightly describes as "surveillance capitalism."[18] The front-end monetary costs associated with practices that uphold human rights are considerably lower than remedial costs associated with breaches, judgements, and suits in response to violations.

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?*

We agree that PIPEDA needs to include a right to object to automated decision making.

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

The parameters set out in Articles 21 and 22 of the GDPR provide a sufficient framing of such a right.

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?*

The right to an explanation fundamentally addresses issues of transparency and meaningful consent. Data subjects need to be informed if they interact with an AI application, and have a right to obtain information on the reasoning underlying any decision or processing operations applied to them. The problem with AI is its inherent opaqueness – the 'black box' analogy

[17] GDPR, *supra* note 25, Recital 4.

[18] Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: Public Affairs, 2019)

emphasizes the lack of transparency in decision-making processes, which has spurred calls for algorithmic transparency.

- Data subjects need to be notified when a decision about them is made based on automated processing.
- Individuals need to be given enough information to be able to understand what they are agreeing to and how their information will be used.
- Data controllers need to be able to explain and translate complex decision-making rationale into an accessible and clear form or language geared towards a lay audience. Communication to individuals must, therefore, be understandable, meaningful and actionable.

We argue that the right to an explanation needs to embrace the principle of plain-language drafting. Again, this touches upon the issue of meaningful consent where consent is only truly 'meaningful' if individuals can fully comprehend what is being asked of them, and how they can maintain control their personal information.

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

Kaminski (2019) argues that the centrality of transparency has gotten lost in the 'right to explanation' debate.[19] There have been calls for transparency in algorithmic decision-making in the form of notice towards individuals or providing options for third-party oversights; some calls have been ambitiously broad, while others respond by outlining the harms that could arise due to this level of transparency.[20] Regardless, transparency in some form or another has its place in algorithmic accountability governance.

When discussing transparency in this sense, perhaps it is useful to look at Frank Pasquale's "qualified transparency": a system of targeted revelations of different degrees of depth and scope aimed at different recipients.[21] Transparency, in this case, is not limited to public revelations, but "includes putting in place internal company oversight, oversight by regulators, oversight by third parties, and communications to affected individuals."[22] The level of access (revelations) would vary depending on the body or organization seeking that information.

In order to enhance transparency measures, our regulators (provincial and federal) need to have an expanded and robust enforcement regime where they can levy order-making powers and administrative monetary penalties.

---

[19] Margot E. Kaminski, "The Right to Explanation, Explained," *Berkeley Technology Law Journal 34, no. 1 (2019):* 209, https://scholar.law.colorado.edu/articles/1227/

[20] Ibid.

[21] Frank Pasquale, *The Black Box Society: The Secret Algorithms that Control Money and Information* (Cambridge, MA: Harvard University Press, 2016), 142.

[22] Kaminski, *The Right to Explanation,* p.210

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?*

We support that the 10 Privacy Principles of PIPEDA, also known as the 10 Fair Information Principles, originally found in the national standard CSA Model Code for the Protection of Personal Information be maintained as:

1. Accountability
2. Identifying Purposes
3. Consent
4. Limiting Collection
5. Limiting Use, Disclosure, and Retention
6. Accuracy
7. Safeguards
8. Openness
9. Individual Access
10. Challenging Compliance

In addition to maintaining these principles, the following should made be a legal requirement under PIPEDA:

- Proactive and preventative (as opposed to reactive and remedial);
- Privacy as the default;
- Privacy embedded into design;
- View privacy as positive-sum to avoid false dichotomies (i.e. privacy vs. innovation);
- End-to-end security (protecting data throughout the entire lifecycle);
- Visibility and transparency; and
- Respect for user privacy.

We also note that we share the GDPR's definition of 'Privacy by Design,' which is essentially "data protection through technology design."[23] In addition, we want to emphasize that privacy by design is not a panacea, and that restrictions on collection at the outset are often more privacy protective.

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. An obligation for manufactures to test AI prior to market access is comparable to the stringent processes that food and drugs are subject to under the Food and Drug Administration in the US and Canada's Food Inspection Agency. The similar notion is seen in the academic research context where research proposals undergo strict review between the researcher and

[23] "Privacy by Design," General Data Protection Regulation (GDPR), Accessed March 13, 2020. https://gdpr-info.eu/issues/privacy-by-design/

Research Ethics Boards. This obligation to test AI products prior to market access, however, cannot exist on its own: robust privacy protections and stronger regulatory powers must also be implemented.

Proposal 6: 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 fundamental nature of AI means that it requires massive quantities of data, its output is often unpredictable, and how it processes that data may not be transparent. This in and of itself puts AI in tension with privacy laws including PIPEDA. PIPEDA limits the collection of personal information to that which is necessary for the purposes identified by the organization; personal information may be retained only so long as necessary to fulfill the reasonable stated purpose for which it was initially collected.

The limitation principles found in privacy legislation are inherently incompatible with AI simply because AI requires all available information to learn and process. Limiting these datasets actually introduces risk; by providing only a subset of data for AI to process and learn from, it inevitably creates biases.

This then begs the question of how AI is being defined. If the discussion is solely focused on algorithms, we could argue that the principles of purpose specification and data minimization could work in the AI [algorithmic] context. If we, however, look at the entirety of big data analytics, which is the intersection of big data, AI, and machine learning, the picture becomes a lot less clear.

The issue with purpose specification ties back in with the complexities of meaningful consent and explainability in an AI context, where organizations need to be able to explain how data is being used in an accessible manner (explainability), and if that data is to be used for a different purpose beyond that of which has been consented to, new consent must be obtained.

We maintain that, by default, AI systems need to be subject to the legal principles of purpose specification and data minimization at the outset *prior* to market access.

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

NA – see below.

3. *If no, what are the alternatives or safeguards to consider?*

It is perhaps safe to assume that data scientists who design and build AI systems will not necessarily take any data minimization constraints into consideration. The onus then falls on organizations and regulators to ensure compliance with data collection laws. Organizations need to have a risk management practice in place to ensure that data minimization requirements and purpose limitation are fully considered from the design phase or as part of the procurement process due diligence.

The ICO has put forth a two-part test regarding purpose limitation: first, the purpose for which the data is collected must be specified and lawful; second, if the data is further processed for any other purpose, it must not be incompatible with the original purpose.[24] We would add that, in addition to the two-part test, there needs to be a strong meaningful consent model where individuals are given enough information about their data to make informed decisions.

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 and obtain consent in the first place, including through innovative models, before turning to other grounds?*

We believe that meaningful consent is free and voluntary, informed, specific, and ongoing. Individuals must have both the right and the means to withdraw consent at any time. If processing goes beyond the initial consent, new consent must be obtained.

We do agree that the current binary model of consent is inadequate. In order for informed consent to be legitimate, there needs to be meaningful alternatives available to individuals who do not choose to consent to privacy-impacting practices.

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?*

Any system that collects, uses and discloses personal information needs to obtain explicit consent from the data subjects from whom data is collected. In the context of AI, explainability is a central pillar where accessible and understandable information needs to be communicated to data subjects so that they can make decisions with autonomy and dignity. In any instance where an individual is incapable of providing their explicit consent, other forms of protection must be available, but the bottom line is that individuals must be given the chance to provide consent at the outset of any data collection.

Consent is fundamentally important (as we have noted throughout this submission), but it is also vital that it not be treated as a work-around or panacea that authorizes what would otherwise be regarded as illegitimate forms of collection, use, and storage.

3. *Requiring consent implies organizations are able to 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?*

[24] ICO, *Big Data,* 38.

While the specific purposes inherent in AI processing is beyond the scope of our work, we maintain that they need to be rendered explainable to the degree that any layperson would have a clear understanding of how their data will be collected, used, and disclosed prior to collection.

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?*

Consent should always be the first line of defense, especially where purposes are clear and directly relevant to a service. In instances where those purposes are less clear or indirectly relevant, this could be where provisions found in the GDPR and the OPC's Report on Consent come into play.

5. *How should any new grounds for processing in PIPEDA be framed: as socially beneficial purposes or more broadly, such as the GDPR's legitimate interests?*

A broader approach, such as the GDPR's legitimate interests, would be more widely applicable, which avoids precluding any other perspectives that fall outside the scope of "socially beneficial." Because the GDPR's legitimate interests is a wider net, there are a wide range of interests that fall within it – including the legitimate interests of the public.

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

As seen in other jurisdictions (i.e. the UK, the US) with contrasts between corporate governance by rule or principle, there is ample evidence that civil monetary penalties are necessary to compel behavioural changes from businesses regarding privacy practices. The incentive here, in this case, is to avoid monetary fines by using meaningful consent models. We believe that we should not be rewarding behaviors that should already be standardized; rather, there needs to be stricter mechanisms to penalize and hold organizations accountable for not having these in place sooner.

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 in achieving both legitimate commercial interests and protection of privacy?*

In the age of AI, machine learning and big data analytics, the sheer amount of personal data and variety that is collected is unprecedented. Where once de-identification was a legitimate strategy to avoid privacy implication, there is doubt whether anonymized personal information can remain truly anonymous in a time of vast data stores and an ever-growing sophistication of analytical tools.

Because of this, we argue that privacy legislation in Canada needs to expand the definition of 'personal information' to include de-identified data (as seen in the GDPR). For more, see our response under Proposal #2.

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

We emphasize the need to expand the definition of 'personal information' under PIPEDA to include de-identified data, but recognize that there are circumstances where a degree of flexibility can be awarded. We caution against relaxing consent requirements, and do not believe that de-identified data should be treated differently from personal information from a consent perspective. Because there is always a risk for re-identification, individuals need to be informed. De-identification could, however, be a factor in deciding whether alternative grounds for processing (i.e. legitimate interests) should be authorized.

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

First and foremost, create a robust enforcement regime where provincial and federal regulators can levy order-making powers and administrative monetary penalties to those found to be non-compliant.

Secondly, shifting the way privacy is framed as zero-sum to a human-rights based approach will empower data subjects to regain control over their personal information.

Third, require organizations and companies to adopt a privacy by design and by default for any systems that collect personal information prior to market access.

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, data traceability is necessary to ensure compliance with principles of data accuracy, transparency, access and correction, and accountability. According to Sanjay Srivastava (2018), traceability addresses several challenges in AI implementation including making answers more understandable by humans and helping to drive compliance in regulated industries by allowing companies to better understand the reasoning process.[25] On the consumer side, traceability will allow data subjects to get into AI decision loops and have the ability to stop or control its tasks when the need arises.

Going back to the fundamentals, in order to achieve meaningful compliance with these principles is to adopt a human-rights based approach to data privacy.

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

1. *Would enhanced measures such as those as we propose be effective means to ensure demonstrable accountability on the part of organizations?*

[25] Sanjay Srivastava, "The Path to Explainable AI," CIO, May 21, 2018, https://www.cio.com/article/3274566/the-path-to-explainable-ai.html

We support the requirement of organizations to be able to provide evidence of adherence demonstrating accountability on request. Inherent to this discussion are the issues explainability, traceability, and human and privacy rights impact assessments.

In addition, we have been long-standing champions of a legislated duty to document of deliberations, actions and decisions in a manner that is accessible and thorough. Previously, these calls were primarily directed at public bodies, but there is no reason that this principle should not be applied in this context to both private and public bodies.

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

The most important consideration is stronger and more comprehensive oversight by the regulators. The federal Privacy Commissioner needs to be given order-making powers, as well as the ability to impose administrative monetary penalties. Once these are implemented, independent third-party auditing throughout the lifecycle of the AI system would be a useful tool at the Commissioner's disposal.

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

Stringent policies and regulations that are enforceable through the Commissioner's office are mandatory. AI systems are not solely limited to the technologies or source codes that make up the system; the policies and procedures that govern it from its inception to design to execution are an integral component of the AI system as well. We cannot emphasize enough the importance of expanding an enforcement regime where the privacy commissioners can utilize enforcement mechanisms to the full extent that the law permits.

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, BC FIPA fully supports expanding commissioners' and regulators' powers so that they can have a robust enforcement regime where they can levy order-making powers and administrative monetary penalties.

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

No. At this time, it would be prudent to focus attention on expanding the powers of the regulators.

Bibliography

BC Freedom of Information and Privacy Association. "Submission to Consultation on Online Reputation." April 28, 2016. https://fipa.bc.ca/wordpress/wp-content/uploads/2016/04/BCFIPA_OnlineReputation2016.pdf

Bertoni, Eduardo. "The Right to Be Forgotten: An Insult to Latin American History." *Huffington Post*, September 2014. http://www.huffingtonpost.com/eduardo-bertoni/the-right-to-be-forgotten_b_5870664.html?utm_hp_ref=tw

EU *General Data Protection Regulation (GDPR):* Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L119/1.

General Assembly Resolution. *Right to Privacy in the Digital Age*. A/RES/68/167 (18 December 2013). http://undocs.org/en/A/RES/68/167

General Assembly Resolution. *Right to Privacy in the Digital Age*. A/RES/69/166 (18 December 2014). https://undocs.org/en/A/RES/69/166

General Assembly Resolution. *Right to Privacy in the Digital Age*. A/RES/71/199 (19 December 2016). https://undocs.org/en/A/RES/71/1699

Information Commissioner's Office. "Big Data, Artificial Intelligence, Machine Learning and Data Protection," last modified April 7, 2017, https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf

Kaminski, Margot E. "The Right to Explanation, Explained." *Berkeley Technology Law Journal* 34, no. 1 (2019): 189-218. https://scholar.law.colorado.edu/articles/1227/

Office of the Privacy Commissioner of Canada. "PIPEDA Report of Findings #2015-002." Last modified: March 23, 2018

Office of the Privacy Commissioner of Canada. "Submission: National Digital and Data Consultations - November 23, 2018." Last modified: December 5, 2018. https://www.priv.gc.ca/en/opc-actions-and-decisions/submissions-to-consultations/sub_ised_181123/.

Office of the Representative on Freedom of the Media Organization for Security and Co-operation in Europe. "New Challenges to Freedom of Expression: Countering Online Abuse of Female Journalists." 2016. https://www.osce.org/fom/220411?download=true

Pasquale, Frank. *Black Box Society: The Secret Algorithms that Control Money and Information.* Cambridge, MA: Harvard University Press, 2016.

Personal Information Protection and Electronic Documents Act. Statutes of Canada 2000, c.5. https://laws-lois.justice.gc.ca/ENG/ACTS/P-8.6/index.html

"Privacy by Design." General Data Protection Regulation (GDPR). Accessed March 13, 2020. https://gdpr-info.eu/issues/privacy-by-design/

Srivastava, Sanjay. "The Path to Explainable AI." CIO, May 21, 2018. https://www.cio.com/article/3274566/the-path-to-explainable-ai.html

Tippmann, Sylvia, and Julia Powles. "Google Accidentally Reveals Data and 'Right to be Forgotten' Requests." *The Guardian*, July 2015. http://www.theguardian.com/technology/2015/jul/14/google-accidentally-reveals-right-to-be-forgotten-requests.

United Nations. "International Covenant on Civil and Political Rights." *Treaty Series*, vol. 999, p. 171 and vol. 1057, p. 407. Accessed March 13, 2020. https://treaties.un.org/doc/Treaties/1976/03/19760323%2006-17%20AM/Ch_IV_04.pdf

United Nations. "Universal Declaration of Human Rights." United Nations. Accessed March 13, 2020. https://www.un.org/en/universal-declaration-human-rights/.

Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs, 2019.