Centre for Information Policy Leadership
HUNTON ANDREWS KURTH

12 March 2020

Comments by the Centre for Information Policy Leadership on the the Office of the Privacy Commissioner of Canada's Consultation on Proposals for Ensuring Appropriate Regulation of Artificial Intelligence

On January 28, 2020, the Office of the Privacy Commissioner of Canada (OPC) issued eleven proposals regarding the regulation of artificial intelligence (AI) and invited the public to comment by March 13, 2020.[1] The Centre for Information Policy Leadership (CIPL)[2] welcomes the opportunity to respond and submit the comments below as input for the final recommendations. CIPL agrees with the importance of the issues raised and hopes that our responses below are helpful in thinking through how to harness the immense benefits of AI without posing unnecessary risks to individuals.

CIPL is conducting extensive research on the interplay between AI and data protection through its project on "Delivering Sustainable AI Accountability in Practice."[3] This project aims to provide a detailed understanding of the opportunities presented by AI, its challenges to data protection laws, and practical ways to address these issues through best practices and organizational accountability. CIPL published its first report "Artificial Intelligence and Data Protection in Tension" in October 2018,[4] and its second report on "Hard Issues and Practical Solutions" in February 2020.[5] These reports are referenced in this submission to supplement CIPL's comments below.

Overall Comments and Summary of CIPL Key Recommendations

CIPL welcomes this Consultation and shares many of the concerns identified by the OPC. CIPL agrees that AI has immense potential for social and economic benefits. We are already seeing these benefits and capabilities of AI across a wide range of public and private sector stakeholders.[6] CIPL also shares the

[1] Consultation on the OPC's Proposals for Ensuring Appropriate Regulation of Artificial Intelligence: Seeking Views on the OPC's Recommendations to Government/Parliament, 28 January 2020, available at https://priv.gc.ca/en/about-the-opc/what-we-do/consultations/consultation-ai/pos_ai_202001/.

[2] CIPL is a global data privacy and cybersecurity think tank in the law firm of Hunton Andrews Kurth LLP and is financially supported by the law firm and 90 member companies that are leaders in key sectors of the global economy. CIPL's mission is to engage in thought leadership and develop best practices that ensure both effective privacy protections and the responsible use of personal information in the modern information age. CIPL's work facilitates constructive engagement between business leaders, privacy and security professionals, regulators and policymakers around the world. For more information, please see CIPL's website at http://www.informationpolicycentre.com/. Nothing in this submission should be construed as representing the views of any individual CIPL member company or of the law firm of Hunton Andrews Kurth.

[3] See CIPL Project on Artificial Intelligence and Data Protection: Delivering Sustainable AI Accountability in Practice, https://www.informationpolicycentre.com/ai-project.html.

[4] See CIPL First Report on "Delivering Sustainable AI Accountability in Practice: Artificial Intelligence and Data Protection in Tension", 10 October 2018, available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_ai_first_report_-_artificial_intelligence_and_data_protection_in_te....pdf.

[5] See CIPL Second Report on "Delivering Sustainable AI Accountability in Practice: Hard Issues and Practical Solutions," February 2020, available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_second_report_-_artificial_intelligence_and_data_protection_-_hard_issues_and_practical_solutions__27_february_2020_.pdf.

[6] CIPL's First AI Report detailed many of the public and private uses of AI across a wide range of sectors—health and medicine, transportation, financial services, marketing, agriculture, education and training, cybersecurity,

belief that promoting the responsible development and deployment of these systems is essential to allowing AI to achieve its full potential, and like the OPC, we are mindful of the privacy, data protection, and human rights risks that can be posed by poorly or irresponsibly implemented AI or other emerging technologies.

CIPL also agrees with the OPC's assessment that AI challenges traditional principles of data protection and cornerstones of privacy, such as purpose limitation, data minimization, transparency, fairness, and automated decision-making.[7] Those challenges are not insurmountable, however, and there is often sufficient scope in current data protection measures to overcome these challenges, although doing so requires creativity, flexibility, agility, cooperation, and continued vigilance from both organizations and regulators. Rather than craft AI-specific laws or regulations, CIPL believes that regulators and organizations should focus instead on applying existing accountability tools to AI applications.

Comments

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

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

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?*

CIPL encourages OPC to maintain its principle of technological neutrality in favor of regulating based on the impact of technology uses, as this will allow responsible organizations to experiment, learn, and grow as they develop best practices for implementing innovative uses of data rather than focus on whether or not that use of data falls within a specific technological category.

CIPL believes that preserving the principle of technological neutrality is important for ensuring the holistic protection of privacy and other human rights. Many of the challenges identified throughout the OPC Consultation predate AI and are posed by technologies other than AI. For example, collection limitation, purpose specification, data minimization, and transparency/explainability have been the subject of concerns with big data, and these will likely continue to be concerns with future technologies. The challenges and risks are bigger and broader than AI, and CIPL encourages regulators and organizations to craft solutions and regulations that reflect this. By choosing to remain technologically neutral in its recommendations, the OPC can further ensure that the solutions developed will not be narrowly confined to AI but rather can be applied more broadly.

---

public authorities/services and data protection. Many of these remarkable uses and benefits of AI are already being realized, although it can be expected that future advancements will continue to push the threshold of what is currently possible.

[7] CIPL's First AI Report explains these challenges at their surface, while its Second Report dives deeper into the specific issues of fairness, transparency, purpose specification and use limitation, and data minimization.

AI-specific rules, regardless of the definition selected, will inevitably result in debates about whether something qualifies as AI or not. As the OPC noted in Proposal 1, governments and regional organizations do not present consistent definitions of AI, and centering a regulation around one definition has an overwhelming potential to be too vague or too narrow. Organizations should not be able to escape their obligations simply by falling outside the definition of AI. Any technology can be used in a way that threatens privacy and human rights.

Furthermore, AI-specific legal structures or regulations could potentially deny society the benefits of properly implemented AI without addressing the underlying problem, which is often the impact of the decision rather than the technology itself. For example, the discomfort with automated decision-making is likely not the AI technology, but rather the fact that a machine is making a significant decision that could result in negative or legally significant impacts on an individual. In this case, the problem to be addressed is not the technology, but rather the role or absence of humans in decision-making. The type of technology is irrelevant; the impact of the decision made by that technology is the source of discomfort or distrust.

While technology-neutral regulation is favorable for managing the impacts of all emerging technologies on individual privacy and human rights, CIPL agrees with the OPC's insight that current rules governing processing may need to be enhanced to achieve these goals. Technology-neutral solutions may allow for innovative ways to help achieve these goals; technology applications may even help. For example, AI and all other innovative and powerful data processing technologies can enable new tools for responsible data governance. These new tools may in fact be necessary to assist in governing new technologies and the impacts of innovative data uses.

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

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

CIPL shares the OPC's commitment to ensuring that new technologies are deployed in ways that respect privacy, data protection, and other human rights. Rather than focusing on a strictly rights-based approach, however, CIPL encourages the consideration of a risk-based approach, which would focus attention on uses of data that pose the greatest risks for individuals and for society. This will allow the flexibility to consider privacy and data protection rights within a broader scope of rights and interests.

CIPL supports strengthening privacy rights in PIPEDA but would encourage further consideration of a risk-based approach to implement data protection principles rather than a strictly rights-based approach. This risk-based approach has been suggested or endorsed by other jurisdictions, including in the GDPR,[8] the

[8] This is apparent in the DPIA requirement in Article 35, among others. For an overview of risk-based provisions of the GDPR, see Gabriel Maldoff, "The Risk-Based Approach in the GDPR: Interpretation and Implications," International Association of Privacy Professionals, available at https://iapp.org/media/pdf/resource_center/GDPR_Study_Maldoff.pdf.

Singapore Model AI Governance Framework,[9] and most recently, the US Office of Management and Budget's (OMB) Guidance for Regulation of AI Applications.[10] There are subtle differences between the approaches, but a risk-based approach allows the flexibility to balance privacy and data protection against other human rights, such as those respecting life and health.

Emphasizing an analysis of impacts and risks to individuals does not diminish the obligation to comply fully with data protection law and does not diminish individual rights, but rather focuses protection of those rights in situations where the risk of harm is greatest. This focus can help determine how to allocate scarce resources by organizations and regulators; help assure data uses that pose greater risks receive greater considerations; help justify more burdensome mitigation processes when warranted by the level of potential harm; and help determine which precautionary or remedial measures an organization should implement to protect against risks to individuals. As explained by the Platform for the Information Society, "[t]he nature of the AI application and the context in which it is used, define to a great extent which tradeoffs must be made in a specific case...AI applications in the medical sector will partly lead to different questions and areas of concern than AI applications in logistics."[11]

Considering the potential impact and risk of harms for proposed processing on individuals allows organizations to efficiently uphold privacy rights without creating overly burdensome regulations for less risky applications of technology. For example, uses of AI that pose little risk of harm to individuals, either because the decision being made is trivial or because the likelihood of a harmful outcome is remote, may understandably warrant less scrutiny. In practice, this could mean that using AI for song or restaurant recommendations warrant less scrutiny than using AI for health diagnoses or transportation decisions.

III. 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?*

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

CIPL believes that the keys to upholding privacy and data protection in the context of automated decision-making are the same in other forms of processing: giving individuals information about the data used, how decisions are generally made, how to correct any inaccurate or false information, and how to seek redress in the case of erroneous or inappropriate decisions.

[9] Singapore Personal Data Protection Commission, "A Proposed Model Artificial Intelligence Governance Framework," (January 2019), available at https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Resource-for-Organisation/AI/AProposed-Model-AI-Governance-Framework-January-2019.pdf, at page 6-7.

[10] "[A] risk-based approach should be used to determine which risks are acceptable and which risks present the possibility of unacceptable harm, or harm that has expected costs greater than expected benefits. Agencies should be transparent about their evaluations of risk and re-evaluate their assumptions and conclusions at appropriate intervals so as to foster accountability." Russel Vought, "Draft Memorandum for the Heads of Executive Departments and Agencies: Guidance for the Regulation of Artificial Intelligence Applications," US Office of Management and Budget (7 January 2019), available at https://www.whitehouse.gov/wp-content/uploads/2020/01/Draft-OMB-Memo-on-Regulation-ofAI-1-7-19.pdf.

[11] "Artificial Intelligence Impact Assessment," Platform for the Information Society (2018), available at https://ecp.nl/wp-content/uploads/2019/01/Artificial-Intelligence-Impact-Assessment-English.pdf, at page 21.

CIPL cautions against a broad right to object to decisions based solely on automated processing, as this could potentially limit the benefits of AI for less critical or risky applications. As examined above, not all applications pose the same level of risk to individuals. CIPL believes that deploying a risk-based approach can help determine the parameters and conditions for when a right to object is appropriate. In some circumstances, a right to object may be sensible and productive to upholding the goals of data protection. In situations with a lower risk of impact, however, a right to object may be unnecessary or unproductive, and it may be preferable to instead provide a right to human oversight, or subsequent visible avenues of redress.

While the proposal and discussion questions did not explicitly mention a right to redress, it is likely that redress will play an essential role in the effective governance of AI. Redress allows individuals to contest and change an outcome they believe is inaccurate, unfair, or otherwise inappropriate. Even with the proper controls and constraints on algorithms, and even with allowing a right to object, it is unlikely that we will achieve the full potential of AI while also preventing all bad outcomes or even all harms. Rather than viewing the potential risk as a reason for shying away from these new technologies, we should instead strive to ensure that, particularly in the context of automated decision-making with a legal or similarly significant impact, individuals have an effective and efficient avenue for contesting outcomes and appealing decisions. Doing so will help protect not only data protection, but also other aspects of human dignity.

IV. 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?*

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

Transparency standards in the law should be generally applicable to all data processing, focusing on the delivery of understandable, actionable and relevant information to individuals. In particular, CIPL does not recommend algorithmic transparency if that term is understood to refer to disclosing the algorithms, as algorithms are proprietary, most individuals would not be able to understand complex algorithms, and in many settings it is important to protect against algorithms being manipulated or "gamed" inappropriately. However, individuals should have access to other information, such as the types of data that go into AI and automated decision-making models, how to correct false or outdated information, and how to remedy erroneous decisions.

CIPL does not recommend having a different standard for automated decision-making and processing than the standards in place for other forms of processing. While the standards and tools for transparency may need to evolve to reflect the new capabilities of AI and emerging technologies, these tools should focus on the overall goals of transparency. These goals are (1) to inform individuals about how their data is used to make decisions, (2) hold organizations accountable for their policies and procedures concerning AI, (3) help detect and correct bias, and (4) generally foster trust in the use and proliferation of AI and new technologies.

Transparency of AI applications has been a particularly difficult challenge, and it is often unclear what exactly transparency means in this context. But the same is often true of human decision-making. Humans are often unable to consistently and rationally explain their preferences for one option over another. Considering approaches to transparency in an offline world can be illustrative of what level and type of transparency to strive for when building regulations around AI and new technologies.

The complexity and changing nature of AI algorithms further complicates transparency expectations. One of AI's strengths is spotting complex patterns and finding inferences that had previously been missed, but this complexity is inherently difficult to explain in terms that are easily understood by humans. Advances in research have led to tools that can help developers better understand how their AI models work, but this requires investing the time and energy to interrogate models, which may not always be feasible (particularly for less risky applications). Furthermore, AI systems may be updated and retrained using additional inputs, so decisions may not be easily repeatable. Because of this complex and dynamic nature, CIPL believes that providing information about the algorithm does not serve the goals of transparency. Not only is disclosing the algorithmic code to individuals and regulators unlikely to be particularly useful for providing clarity about the decision, but algorithmic transparency could have the potentially harmful effects of disclosing trade secrets or helping individuals game the system.

The Consultation cites the Council of Europe's guideline that encourages mandatory disclosure when an individual is interacting with an AI application as well as an explanation of why AI is being deployed and what is expected from its use.[12] This might be helpful in some cases, but it will often be meaningless, overly burdensome, or otherwise ineffective in building trust. As explained in the comments for Proposal 1, it is ultimately not the AI technology that matters, but rather the fact that a nonhuman decision is having consequences on an individual in a way that he or she might not expect.

Though CIPL cautions against regulatory requirements for algorithmic transparency or mandatory disclosure of the use of AI, CIPL recognizes and agrees with the OPC's assessment that transparency requirements may need to be revamped or reconsidered in the context of AI in order to be meaningful. The following considerations may be helpful in outlining what could constitute meaningful transparency and openness:

Audience: Transparency may look different depending on the audience it is geared toward—the individual or category of individuals, the regulator, a business partner, or even for purposes of internal transparency to an oversight board or senior leaders. All of these different audiences imply different types and requirements of transparency that should be fulfilled appropriately. A regulator may need to know more details about an AI use-case in the context of an investigation or audit—the model, the data sets, inputs and outputs, measures to ensure fairness and absence of bias, etc. For individuals, this type of information may be too much and "missing the forest for the trees." Equally, an organization developing AI technology to be used by another organization may be unable to provide transparency to data subjects directly, but it may need to provide additional transparency about the technical measures to ensure a properly working model, bias avoidance, accuracy, documentation regarding tradeoffs, etc. Therefore, it may be hard to be

[12] See Consultative Committee of the Convention for the Protection of Individuals with Regard to Automatic Processing of Personal Data, Council of Europe, Guidelines on Artificial Intelligence and Data Protection, 25 January 2019, available at https://rm.coe.int/guidelines-on-artificial-intelligence-and-data-protection/168091f9d8, Section II(11).

categorical about the precise list of elements of transparency, as it very much depends on who the audience is and the specific purpose of transparency in a given context.

Context-Based Transparency: The level and method of transparency should ultimately be tied to the context and the purpose of AI applications. As highlighted in the UK ICO's Project Expl*AI*n, a survey of citizen juries empirically demonstrated that individuals facing AI healthcare scenarios cared more about accuracy than transparency, while transparency expectations were heightened for the use of AI in job recruitment and criminal justice scenarios.[13] This suggests that transparency, and the tools used to achieve it, may differ based on what the AI application is used for, what the consequences are, and what rights individuals have going forward.

To illustrate these different considerations for transparency, consider the use of facial recognition technologies by airlines to check boarding passes or by customs officials to allow individuals into a country. The decision made by the AI in these cases is very significant, but transparency regarding the fact that AI is being used or about the code itself is unlikely to be of concern to the impacted individual. Instead, the concern is with how to contest or change the decision, so facilitating the goals of transparency will require a greater emphasis on speedy and effective avenues of redress. Developing efficient and visible avenues for such review—whether before or after a decision—will be an important part of transparency in AI contexts.

V. 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?*

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?*

Privacy by Design and Human Rights by Design encourage organizations to be thoughtful and responsible about the way data is being processed. If these terms are understood to require organizations to develop processes that promote thoughtful innovation throughout the product or application lifecycle, CIPL supports these as legal requirements. We do, however, suggest keeping these in line with general principles of accountability rather than rigid processes, as this will allow organizations to find innovative ways to foster and implement responsible AI.

The OPC's proposal astutely recognizes the importance of monitoring applications throughout their product lifecycle. CIPL fully agrees that it is important to encourage accountable and thoughtful innovation throughout all stages of design, development, and deployment of AI. To the extent that privacy or human rights by design are defined as mechanisms to encourage organizations to pause and consider the impacts of their innovation, CIPL supports such a requirement. Depending on how this legal requirement is framed will determine how desirable or effective it will be, and there may be other tools that can help encourage this level of thoughtful consideration.

[13] "Project Expl*AI*n: Interim Report," UK Information Commissioner's Office (3 June 2019), available at https://ico.org.uk/media/2615039/project-explain-20190603.pdf, at page 15.

CIPL would like to highlight one important consideration of requiring privacy or human rights by design. It is necessary to consider who such a requirement would apply to in practice. Organizations that sell off-the-shelf applications of AI may not know how the technology is being deployed, so a requirement that manufacturers conduct privacy or human rights by design assessments likely offers little protection or is otherwise unhelpful to prevent violations. For this reason, it may be more prudent to focus on the use and potential impacts of a given application and requiring organizations to implement procedures to ensure accountability.

The goal is not to determine whether one particular application of AI is in compliance with privacy and human rights at one moment in time, but rather to know that all applications are being examined and monitored on an ongoing basis with the overarching objective of continuous improvement and risk mitigation. Privacy by design is one procedural tool to accomplish this objective, but there are a variety of other tools available as well, such as data protection impact assessments (DPIAs), AI-specific DPIAs, or data review boards (DRBs).

AI DPIAs: Many organizations today use DPIAs to comply with data protection and to demonstrate their compliance. Some have decided to use DPIAs in an even broader context than that required by law, partially to foster privacy by design and risk mitigation and partially to establish a common lexicon and methodology for assessing data uses across departments and geographies. These assessments may have additional value in the context of AI, and some organizations are developing AI-specific DPIAs, either as a supplement to the assessments required by law, or as an entirely separate assessment.

DRBs: Data review boards are another potential tool for organizations to structure how they conduct the balancing of interests between the impact of data uses and new AI applications. This emerging tool requires organizations to consider the impact of data uses and foster responsible decision-making. "The goal of a DRB is to facilitate better decision-making and responsible innovation, improve organizational accountability and create trust. DRBs will help organizations consider novel data uses in the context of the law, as well as organizational and societal values."[14]

Given the wide variety of tools available to foster accountability and responsible decision-making,[15] CIPL would encourage a requirement for organizations to design and implement processes that uphold privacy and human rights. However, we support a strong focus on a broader principle of accountability, as this will allow and even encourage organizations to develop innovative processes to uphold societal values and human rights.

[14] Rachel Dockery, Fred Cate, & Stanley Crosley, "Why Data Review Boards Are a Promising Tool for Improving Institutional Decision-Making," IAPP (28 February 2020), available at https://iapp.org/news/a/why-data-review-boards-are-a-promising-tool-for-improving-institutional-decision-making/#.

[15] CIPL's Second AI Report has a table in Appendix B that provides 67 possible tools and processes that organizations are implementing to foster the responsible and accountable deployment of AI. CIPL Second Report on "Delivering Sustainable AI Accountability in Practice: Hard Issues and Practical Solutions," February 2020, available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_second_report_-_artificial_intelligence_and_data_protection_-_hard_issues_and_practical_solutions__27_february_2020_.pdf, at page 34-35.

VI. 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?*

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

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

CIPL supports the OPC's efforts to modernize the principles of purpose specification and data minimization, as this will help society harness the economic and social benefits of AI. Taking a risk-based approach to these principles and considering the context in which data is collected and processed will help achieve the goals of data protection without compromising the benefits of AI.

Purpose specification and data minimization are two traditional data protection principles that sometimes conflict with the capabilities of emerging technologies. Similar to developing and implementing new tools for fostering meaningful transparency, doing the same for purpose specification and data minimization will first require understanding the goals of these principles and then finding the appropriate tools to achieve those goals.

The spirit of purpose specification requires that notice be precise, as "use for AI" alone would be neither specific nor precise enough to provide meaningful information to the individual. Instead of allowing purposes to become so broad as to be meaningless, data protection authorities have interpreted purposes narrowly, which highlights the need to provide flexibility for allowing further processing. The GDPR, like the OECD Privacy Guidelines, explicitly permits further processing for new, "not incompatible" purposes.[16] Further processing based on "compatibility" should be allowed for future uses that are consistent with, can co-exist with, and do not undermine or negate the original purpose. These uses must be backed by strong accountability-based safeguards, including benefit and risk assessments, to ensure that new uses do not expose the individual to unwarranted increased risks or adverse impacts.

Similarly, while the intention and goals of the data minimization principle are still possible in our technological landscape, achieving these goals will require more creative solutions and flexible interpretations. For AI, particularly at the development and training stages, what is necessary is a considerable amount of data, and having too little data can hinder the development of an algorithm. For instance, the collection and retention of significant amounts of data, including sensitive data, may be necessary to mitigate the risks and ensure fairness in certain AI applications. This is a contextual tradeoff which organizations will need to assess carefully in order to strike an appropriate balance between competing requirements. It may be necessary to collect and retain information about race and gender to balance an employment screening tool that is hiring only white male candidates due to inherent bias in

[16] "The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose." OECD Revised Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (2013), available at http://oecd.org/sti/ieconomy/oecd_privacy_framework.pdf.

the original training data set. This seems counterintuitive to the traditional understanding of data minimization, but in reality, having more data —in some cases— is necessary to reduce risk.

Some helpful considerations for purpose specification and data minimization in the context of AI and other emerging technologies might include:

Differentiating between data used for training AI versus deploying AI: The concept of a training phase is novel to AI, and data is often needed in greater amounts during the training phase than during deployment. In the training phase, where no individual decision-making occurs, the risk of harm to individuals by repurposing their data is lessened or eliminated entirely. As such, further processing in this phase should be deemed compatible with the original purpose. Additionally, by limiting data use in the deployment phase but providing more flexibility for data use in the training phase, organizations are managing the potential harm to individuals and thus upholding the original intention of the data minimization principle. While other accountability tools will be necessary to govern the training phase, distinguishing between the training and deployment phase for purposes of purpose specification and data minimization could help balance innovation while fostering better data protection for individuals.

Context-based data minimization: Organizations should proactively articulate and document the need to collect and process data (whether it is old data or data not on its face strictly necessary to the purpose of the processing), as well as what is expected to be learned or accomplished by processing the data. This would be especially helpful for the training phase, although it could be useful for both training and deployment. Determining what is adequate, relevant, and necessary will be dependent on the context, but this proactive and continuous assessment will serve to demonstrate that the data to be collected is relevant and not excessive in relation to the purpose for processing.

Deploying tools to minimize risks to individuals: Technological tools to help with data minimization are still in an early stage of development and are often expensive for smaller organizations to deploy, but their continued exploration should be encouraged. For example, in some cases, federated learning could enable AI algorithms to learn without data ever leaving a device and without the need to centralize large amounts of data in a single virtual location. Organizations may also consider the possibility of anonymizing or pseudonymizing data sets, although this may pose challenges of its own. At the same time, while further research and development efforts are needed to ensure proper de-identification, a flexible interpretation of notions of anonymous or pseudonymous data would go a long way to enable use of data for training of AI and to reduce the compliance risks for organizations.

Benefits of a risk-based approach: Lastly, the risk-based approach supported in Proposal 3 can be helpful in the purpose specification and data minimization contexts. The level of continued notice and the requirements necessary for further processing old data may be understood as a function of the risk of harm posed by that processing. "Data used in one context for one purpose or subject to one set of protections may be both beneficial and desirable, where the same data used in a different context or for another purpose or without appropriate protections may be both dangerous and undesirable."[17] Therefore, purpose specification and data minimization may be more effective if these principles rely less

[17] Fred H. Cate and Rachel D. Dockery, "Artificial Intelligence and Data Protection: Observations on a Growing Conflict," Seoul National University Journal of Law & Economic Regulation, Vol. 11. No. 2 (2018), at page 123.

on evaluating the appropriateness of an intended use by the original terms and instead focus on the risk and impact of the new use.

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

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

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?*

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?*

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?*

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)?*

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

CIPL supports and commends efforts to create innovative approaches to finding grounds for processing when traditional avenues of consent are not feasible. While notice and consent are one way to approach privacy protection, they are not the only way to empower individuals nor in many settings are they the best way to protect individuals. By creating alternative grounds for processing, the OPC will help to balance the need for privacy protection with the vast benefits—both social and economic—promised by AI and other emerging technologies.

CIPL encourages the OPC to further de-emphasize consent, as this has the potential to unreasonably burden individuals, is increasingly ineffective at protecting privacy and other rights, and can undermine legitimate, necessary, or beneficial processing activities. The three suggestions mentioned above (considering training data separately from deploying data, allowing organizations to demonstrate data minimization by proactive articulation of the goals of processing, and employing technological tools to enhance privacy) are all methods to decrease the need to return to individuals to get consent.

CIPL has previously encouraged the adoption of a legitimate interest exception under PIPEDA, noting that this exception "takes on a particularly important function in the fast moving, rapidly developing and

changing digital economy" because "it is capable of legitimizing any processing operations (including those that might be as of yet unknown or unimagined and thus not susceptible to specially-designed exceptions to consent) in which the legitimate interests of the business or a third party are not outweighed by the rights and freedoms of an individual, as determined by a risk/benefit assessment."[18]

CIPL believes that allowing for these alternative grounds for processing will ultimately help society reap the benefits of new technologies without unnecessarily burdening individuals or organizations, although we also recognize an increasing need for organizational accountability and data stewardship as we shift from the individual control model to allow other grounds for processing.

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

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

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

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

CIPL recommends creating a broad exception for de-identified information from all relevant statutory requirements (such as consent, data minimization, etc.). De-identification can facilitate responsible use of personal information to help train and deploy new and beneficial technologies while also upholding individual privacy. Therefore, incentives for de-identification may be helpful for facilitating the development of innovative technologies, such as allowing de-identified data to be used for internal research or for AI training without having to set pre-defined retention periods or the data being in scope for the exercise of individual rights such as access, correction and deletion.

The OPC Consultation rightfully points to the increasing ability to re-identify previously anonymized and de-identified information through sophisticated techniques and asks what protections are available to address this problem. CIPL believes that in light of the fact that complete and permanent anonymization or de-identification is increasingly difficult, technical anonymization techniques must in some contexts be complemented by enforceable administrative, technical, physical and legal safeguards that prohibit attempted re-identification of personal information except for certain permissible purposes.

One useful standard was articulated by the US Federal Trade Commission in 2012: "Personal information should be subject to fewer privacy protections or legal requirements if (1) the data is not reasonably identifiable; (2) the company publicly commits not to re-identify it; and (3) the company requires downstream users of the data to keep it in de-identified form." [19] This standard could be translated for

[18] CIPL Comments on Innovation, Science and Economic Development Canada's Proposals to Modernize the Personal Information Protection and Electronic Documents Act (27 September 2019), available at: https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipls_comments_on_iseds_proposals_to_modernize_the_personal_information_protection_and_electronic_documents_act.pdf.

[19] U.S. Federal Trade Commission report, "Protecting Consumer Privacy in an Era of Rapid Change,

the Canadian legal framework to mean that anonymization or de-identification requires reasonable technical anonymization or de-identification in light of the purpose for which the information is being used, coupled with appropriate contractual and legal safeguards that ensure an enforceable obligation not to re-identify the information.

Lastly, it is important to note that—in some specific and narrowly defined cases—re-identification is legitimate and must be protected by appropriate exceptions. For example, if security research aims to test security measures and techniques, re-identification of data that has been de-identified should not be subject to penalties. Those carrying out such genuine testing could be obliged to inform the company first before going public with their findings. This would mitigate the risk of people making public disclosures that could negatively impact individuals and claiming a defense of security testing.

IX. 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?*

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

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

CIPL fully supports demonstrable accountability as a governance model enabling trust in AI development and use. Accountability should be a cornerstone of modern data protection, as it allows the flexibility for innovation without compromising individual privacy or placing unnecessary burdens on individuals.

The rapid and widespread development of new technologies—including AI— has created a renewed need for greater organizational accountability and data stewardship. [20] Developing an impact- and process-

Recommendation for Business and Policymakers," March, 2012 at 22, available at https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protectingconsumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf.

[20] For a full discussion of organizational accountability in data protection, see CIPL white papers on "The Case for Accountability: How it Enables Effective Data Protection and Trust in the Digital Society", 23 July 2018, available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_accountability_paper_1_-_the_case_for_accountability_-_how_it_enables_effective_data_protection_and_trust_in_the_digital_society.pdf; "Incentivising Accountability: How Data Protection Authorities and Law Makers Can Encourage Accountability", 23 July 2018, available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_accountability_paper_2_-_incentivising_accountability_-_how_data_protection_authorities_and_law_makers_can_encourage_accountability.pdf; and CIPL Accountability Q&A, 3 July 2019, available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_accountability_q_a__3_july_2019_.pdf.

oriented approach to data protection will necessarily require that organizations become better data stewards. This will include the need for organizational risk management, improved processes, and better transparency. CIPL believes that continued attention and vigilance can be achieved through the development of collaborative standards, the sharing of best practices, investing in awareness raising, education, research, and training and through establishing demonstrable governance processes for all relevant actors.

An enhanced focus on data stewardship and organizational accountability is especially necessary in the context of AI. This is because of the challenges in providing individuals with meaningful disclosures about AI tools and algorithms that are difficult even for experts to understand. While a stewardship focus does not eliminate the need for disclosure and transparency, it recognizes that organizations have an obligation to make more thoughtful decisions, and to assume greater responsibility for the consequences of the products, services and technologies that they are developing, in situations where individuals are less able to make informed decisions of their own.

The CIPL Accountability Wheel has been used to promote organizational accountability in the context of building, implementing and demonstrating comprehensive privacy programs. This framework can be a useful tool for helping organizations develop, deploy, and organize robust and comprehensive data protection measures in the AI context and also to demonstrate accountability. The Accountability Wheel provides a uniform architecture with seven elements for organizations to build and demonstrate accountability: Leadership and Oversight; Risk Assessment; Policies and Procedures; Transparency; Training and Awareness; Monitoring and Verification; and Response and Enforcement. Organizational efforts to promote trustworthiness around AI can map to this wheel to ensure a holistic approach, as each element provides important protections for individuals.[21]

[21] See Appendix B of CIPL's Second AI Report, which lists 67 possible tools and processes that organizations are implementing to foster the responsible and accountable deployment of AI. CIPL Second Report on "Delivering Sustainable AI Accountability in Practice: Hard Issues and Practical Solutions," February 2020, available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_second_report_-_artificial_intelligence_and_data_protection_-_hard_issues_and_practical_solutions__27_february_2020_.pdf, at page 34-35.

Conclusion

CIPL is grateful for the opportunity to comment on the Office of the Privacy Commissioner of Canada's "Consultation on the OPC's Proposals for ensuring appropriate regulation of artificial intelligence." We look forward to further opportunities for dialogue on AI or other privacy and data protection matters.

If you would like to discuss any of the comments in this paper or require additional information, please contact Bojana Bellamy, bbellamy@huntonAK.com; Markus Heyder, mheyder@huntonAK.com; Nathalie Laneret, nlaneret@huntonAK.com; Sam Grogan, sgrogan@huntonAK.com; Matthew Starr, mstarr@huntonAK.com or Giovanna Carloni, gcarloni@huntonAK.com.