Global ATP Association of Test Publishers
The intelligent voice for testing
601 Pennsylvania Ave., N.W.
Suite 900
Washington D.C. 20004
+1.717.755.9747
www.testpublishers.org

BEFORE THE OFFICE OF THE PRIVACY COMMISSIONER OF CANADA

Re: Comments on Consultation on Proposals for the Proper Regulation of Artificial Intelligence

The Association of Test Publishers ("ATP") submits these comments expressing the views of the testing industry, especially members from Canada, in response to the request from the Office of the Privacy Commissioner of Canada ("OPC") for feedback on its recommendations for the regulation of artificial intelligence ("AI").[1] This submission is being made by the required date of March 13, 2020.

The ATP is the international trade association for the testing industry, including Canadian members, comprised of hundreds of publishers, sponsors,(i.e., owners of test content, such as certification bodies) and test delivery vendors of tests used in various settings (including healthcare, employment, education, e.g., academic admissions, clinical psychology, and certification/ licensure/ credentialing), as well as businesses that provide testing services or administer test programs ("Members"). Since its inception in 1987, the Association has advocated for the use of fair, reliable, and valid assessments, including ensuring the security of test content and test results. Our activities have included providing expertise to and lobbying the US Congress and state legislatures in the United States on proposals affecting the use of testing in employment, as well as representing the industry in regulatory matters and litigation surrounding the use of testing.

In the area of individual privacy protection, the ATP has provided specific education on the GDPR to its US and EU Members – we published a "Checklist for EU-US Privacy Shield Registration" (2016) and a "Compliance Guide for the EU General Data Protection Regulation" (2017). Thereafter, ATP submitted extensive comments to the European Data Protection Board on its proposed guidelines for the use of video surveillance (September 2019). Addressing various privacy regulations in the United States, the ATP twice submitted comments to the California Attorney General on its proposed regulations for implementing the California Consumer Privacy Act (December 2019 and February 2020), and comments to the Massachusetts Joint Committee on Consumer Protection on pending privacy legislation (January 2020).

[1] For ease of understanding, the ATP uses the term "AI" to include any form of artificial intelligence, automated decision-making, machine learning, profiling, and algorithmic software.

The ATP appreciates the opportunity provided by the OPC to give feedback on its proposals for the enhancement of the Personal Information Protection and Electronic Data Act ("PIPEDA"). We strongly believe that there are specific circumstances commonly found in the testing industry where the use of AI is both appropriate and necessary, where its use is justified when balanced against the rights of individual test takers, and where this technology should be allowed within the existing constraints of PIPEDA. Therefore, we request that the OPC carefully craft proposed language consistent with these explanations clarifying how those proposals should be written for incorporation into PIPEDA.

Many testing events occur in today's society, which greatly benefit society in general, along with test users and individual test takers. Canadian citizens are no exception to the vast – and growing – use of assessments by individuals to help themselves to advance personally or professionally. For that reason, it is vitally important that testing programs are able to ensure its tests are fair to all test takers – in so doing, testing organizations today use AI for the development of software for the delivery of assessments, as well as for the development of scoring rubrics. Testing organizations also rely on AI to develop items for use in assessments. Further, AI is currently playing a role in assisting testing organizations and end users of assessments, for a variety of purposes, including but not limited to: (1) assisting employers identifying candidates who meet their job-related needs; (2) providing doctors with data for diagnosing and treating physical diseases and mental disorders; (3) enabling certification bodies to ascertain if an individual has mastered specific competencies; and (4) performing test security analyses to detect cheating by test takers. Thus, it is clear that diverse uses of AI have already become an indispensible element of the assessment process.

The ATP would refer the OPC to the GDPR requirements around the use of AI, which offer useful guidance. Although we will address the GDPR language around profiling and AI in greater detail in our specific responses, we note that the right of individuals concerning the use of "automated decision-making" is not that wide-ranging (*see* EU Charter of Fundamental Rights, Articles 4[4] and 22). An individual's right is limited to preventing activities that are based *solely* on automated decision-making and that produce legal or similarly significant effects. Current guidance by Working Party 29 (now the EDPB) provides that "**meaningful** human intervention" takes an action outside of this right. Furthermore, use of pseudonymous (or anonymous) data may also remove the activity from application of the GDPR requirements. Consequently, use of automated decision-making and other functionality are not prohibited, but are tightly restricted.

The California Legislature is currently considering a bill (SB 1241), the Talent Equity for Competitive Hiring (TECH) Act. The bill establishes guidelines for employers to follow that allow them to modernize their recruiting processes using technological tools that reduce bias, leading to a more diverse workforce. The bill states that assessment technologies, including AI, "will be considered in compliance with anti-discrimination rules if: 1) they are pre-tested for bias before being deployed and found not likely to have an adverse impact on the basis of gender, race or ethnicity; 2) outcomes are reviewed annually and show no adverse impact or an improvement of hiring among underrepresented groups; and 3) their use is discontinued if a post-deployment review indicates adverse impact."[2]

[2] In September, 2019, the California Assembly passed a resolution (ACR 125), also known as the Fair Hiring Resolution, urging adoption of legislation that would tackle implicit racial and social biases in corporate hiring by creating clear rules of the road for how employers can use these smart technologies. Nevertheless, the resolution recognized that, "Innovative technologies for hiring and promotion, including artificial intelligence and algorithm-based technologies, have the potential to reduce bias and discrimination in hiring and promotion based on protected characteristics, such as socioeconomic status or status as a formerly incarcerated person….At the same time, these technologies can help employers reach larger and more diverse pools of qualified talent and better identify candidates with the right skills and abilities to succeed."

Because of the unique elements of testing programs, the ATP urges the OPC to take into account the special circumstances surrounding the use of AI in testing to accomplish legitimate goals for individuals as well as businesses. We encourage the OPC to recognize that this analysis establishes the balancing of the privacy rights of every individual with the rights of the test sponsor and the testing service organization that provides the testing services, as required by the PIPEDA.

In order to best inform the OPC and provide an introduction to the issues specific to testing, the ATP will address the proposals and questions set forth by the OPC in its request. These positions have been developed with input from our Canadian Members.

Proposals for Consideration

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

The ATP firmly agrees with the OPC's own acknowledgement that "PIPEDA is technologically neutral and is a law of general application." As such, we submit that it would be inappropriate to add definitions expressly relating to AI, automated decision-making, machine learning, automated processing, or profiling. However, as suggested elsewhere in the OPC proposals, there is a need for specific guidance to cover certain uses of AI, now and in the future, which would support clarification as to when and how such rules would apply.

Discussion questions:

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?

1. The ATP believes that under PIPEDA, an AI system should be held to the same standards as any other system for processing personal information or decision making about individuals. The laws, rights, and obligations related to the processing of personal information are not different because the technology is novel or new. Moreover, it is fundamental to our understanding that an AI system must be built, trained, used, and be maintained under human supervision (i.e., an organization or individual).
2. Accordingly, the ATP contends that no new or unique rules should apply only to AI, because the ability to define AI or to describe the future requirements surrounding a specific application of technology makes enforcement very difficult.

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

As cited by the OPC (*see* fn 11), the 2019 Resolution of Canada's Federal, Provincial and Territorial Information and Privacy Commissioners, has already determined that AI technologies must be "designed, developed and used in respect of fundamental human rights, by ensuring protection of privacy principles such as transparency, accountability, and fairness." The ATP has no disagreement with this statement.[3]

However, in order to ensure such protection, the ATP contends that PIPEDA should be read from a rights-based perspective that recognizes privacy in its proper breadth and scope, yet balances those rights with the rights of businesses under PIPEDA, and uses that balanced perspective to provide direction on how its provisions should be interpreted. Such an approach will clarify rights in PIPEDA and ensure that automated decision-making receives a proper focus.

Discussion question:

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?

1. As stated above (*see supra.* at page 3), we submit that an AI system should be held to the same standards as any other system that processes personal information in order to make decisions. The challenges would generally remain the same, but privacy requirements remain subject to PIPEDA. As such, notice of the use of AI and its purpose(s) must be given and an individual should have a reasonable opportunity to withdraw consent – which may result in the inability of a business to provide goods and services to that individual. That result occurs because the rights of the individual must be balanced against the rights of the business to protect its intellectual property. Finally, the legitimate interest of the business must be given equal consideration. Thus, PIPEDA should establish a framework that will be used to evaluate these factors in a neutral and transparent manner, judged in the same manner as any other processing system. Specific to the testing industry, the approach advocated by the California Legislature in SB 1421 (*see, supra.* at page 2), of providing a safe harbor for a testing organization that is able to document its evidence of fairness in the development of an AI system would provide an incentive for the testing industry to ensure that the use of technological solutions meet appropriate scientific research standards. Thus, in another testing example, an individual test taker generally will not be permitted to object to the use of AI if the person is trying to cheat on a test – the legitimate interests of all other test takers and the testing organization depend on such an interpretation.

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

[3] The ATP recognizes that the GDPR also incorporates a human rights-based approach to privacy within the EU's data protection legislation. In its various recitals, the GDPR makes repeated references to fundamental rights of individuals in relation to data processing.

As discussed above (*see supra.* at page 3), Article 22 of the GDPR grants individuals the right not to be subject to automated decision-making, including profiling, except when an automated decision is necessary for a contract; authorized by law; or explicit consent is obtained. However, Article 22 also provides an exception that where significant automated decisions are taken based on a legitimate basis for processing, or where the public interest (or official authority) exists, that can override the rights of the individual even though the individual still has the right to obtain human intervention, to contest the decision, and to express his or her point of view (*see* Article 21).

We believe that PIPEDA should be required to balance those same rights. Accordingly, the controller/business should be allowed to continue its use of AI by showing that there is a compelling reason that overrides the individual's right, including because the business can demonstrate that there is the establishment, exercise, or defense of legal claims.[4].

For these reasons, the ATP supports adoption of limited rights associated with restrictions on the use of AI through a right to object in PIPEDA, subject to a balancing of rights, parallel to those provided under the GDPR.[5] It seems that such a generic right and its balancing against the business's rights to protect its IRP could be added to Principle 4.3.8 of PIPEDA (i.e., an individual may withdraw consent at any time, subject to legal or contractual restrictions under the rights of the business and reasonable notice of those legitimate interests).

Discussion questions:

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?

1. The ATP contends that a generic right to object to any processing is all that is appropriate or required. To establish a right that is specific to an AI system would expand current rights under PIPEDA and convert it into a very different kind of law than it is today. Every individual has the right to decide if s/he wants to engage in an activity or purchase goods or services from any business – giving individuals notice that AI is being used, how it is used, and the purpose of such use, and the basics of how the decision was reached, is sufficient to enable every person to make that choice. There should be no alternative standard for AI.

   The ATP stresses that bias can exist in any decision-making system, even those that are 100% human controlled and contain no AI. Each country has laws and/or regulations governing those situations. Merely adding AI as an element to what is essentially a human decision-making process does not inherently make the process more suspect – or any more biased. Indeed, in most of these situations, a

[4] Under the GDPR, there are no exemptions or grounds to refuse an individual's objection to the use of automated decision-making (i.e., profiling) when it involves "direct marketing." Marketing uses are not involved in the purposes for which AI is used in testing as described above (*see supra.* at page 2).

[5] In this vein, the ATP notes that many of its Members conduct international testing operations. Thus, having consistent and harmonized laws, regulations, and interpretations would allow businesses to adopt a single set of privacy policies and procedures for use throughout the world.

human being remains the ultimate decision-maker.[6] The ATP does not endorse or support the concept that AI should be used exclusively for any decisions.

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

The ATP points out that even the GDPR recognizes that a business may have intellectual property rights ("IPR") that come into play when an individual seeks information about the collection and use of personal information. *See* Working Party 29 Guidelines on Article 15 (IRP and other intellectual property (e.g., trade secrets) that are central to the controller's business

model must be respected). This fact is equally relevant when it comes to AI technologies.[7] Thus, while an individual is entitled to have notice about the collection and use of their personal information, including the right to know if any automated decision-making is used and for what purpose it is used, that right does NOT give an individual the right to access the IPR of the business.

Regarding an individual's right to obtain information about the algorithmic logic used for AI, Working Party 29 has expressed its opinion that a controller only needs to provide "the rationale behind, or the criteria relied on" in reaching a decision without disclosing the entirety of the scientific basis, which is usually part of the business's IPR." Article 29 Working Party, 'Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679', WP251rev.01, adopted on October 3, 2017, as last revised and adopted on February 6, 2018. *See* hhttp://ec.europa.eu/newsroom/just/item-detail.cfm?item_id=50083. [8] This position is consistent with the ATP's view that disclosure of AI must not jeopardize any IPR.

We believe PIPEDA should include a right to explanation that would provide an individual interacting with an AI system the basic reasoning underlying any automated processing of their data, and the consequences of such reasoning for their rights and interests, subject to the controller's right not to be required to disclose its IPR. This would also help to satisfy PIPEDA's existing obligations of providing individuals with rights to access and correct their information.

---

[6] The ATP endorses the comments of Multi-Health Systems on this point, which note that, "A doctor who consults an AI system for diagnostic assistance remains the decision maker. A financial transaction denied for suspected fraud by an automated system, can be appealed by calling one's bank, yet the value of the many attempted fraudulent transactions that are caught by the bank's intelligent AI system exceeds the capacity of an organic process alone. We should not seek to enshrine in law or policy xenophobic practice that we otherwise would not permit if we substituted gender or race for machine or human."

[7] Not only does AI and related automated technology, including machine learning and the development of software algorithms, often involve patents or copyrights (or both), but the developer often has trade secrets associated with the technology. Although Canada has no trade secret law, the courts nevertheless recognize a business's right under common law to protect its valuable business information, including "new technology" by keeping it secret. *See* the Canadian Intellectual Property Office publication: https://www.ic.gc.ca/eic/site/cipointernet-internetopic.nsf/eng/wr03987.html (downloaded on March 8, 2020).

[8] We acknowledge that different interpretations exist with respect to whether the GDPR requires only an explanation of system functionality or extends to include the "rationale for the logic, significance and consequences of specific decisions." *See* Malgieri, G., "Automated decision-making in the EU Member States: The right to explanation and other 'suitable safeguards' in the national legislations." *Computer Law & Security Review* 35 (2019) (cited in the OPC Request, fn. 18). However, the ATP asserts that those interpretations completely ignore the role of the business's IPR, and therefore, must be discounted.

To assist in achieving this balanced outcome, the ATP would suggest that modifications to PIPEDA include the requirement for a Privacy Impact Assessment (PIA), including an assessment relating to the impacts of AI processing on privacy and human rights. The published content would be based on a minimum set of requirements that would be developed in consultation with the OPC. We view this result as similar in effect to the proposed California legislation (SB 1241) (*see supra.* at page 2).

Discussion questions:

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?

1. The right to explanation, balanced against the business's right to protect its IRP, should be based on the controller providing evidence about its AI, without any need to expose its IPR. PIPEDA should carefully establish a set of requirements by which a business can provide evidence that allows for evaluation, assessment, and an explanation of the basics for how a decision was arrived, but that stops short of any requirement to disclose IRP.
2. Even so, the ATP cautions that algorithmic transparency must not impede innovation and the recognition of the value of innovation and invention. We believe the OPC should strive to establish accountability and trust through regulations; but trust is more than trust in the way a business handles personal information – it is equally about the trust a person has in the underlying business. In testing, that means that test takers must also trust the program to deliver accurate scoring and ensure that its scores (or related outcomes, such as credentials) maintain the reputation and value for each person. So, an AI system used for detecting cheating is as important to the individual test takers as it is to the testing organization. Ironically, an AI system is simpler to assess for bias and determinant principles then an intelligent human. The ATP contends that transparency in the law/regulations will lead to greater trust that an AI system creates fair use of data and protection of privacy.

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

The GDPR, specifically Recital 78 and Article 25, requires a business to meet compliance standards in relation to its treatment of personal information by designing technology, services, and products to achieve maximum compliance and security (termed "privacy and data protection by design"), as well as requiring that the strictest privacy setting on products and services be set by default without any action by the consumer (termed privacy by default).

The privacy by design framework was first published in 2009 and then adopted by the International Assembly of Privacy Commissioners and Data Protection Authorities in 2010. Privacy by design is also encouraged indirectly by other global privacy laws. In general, privacy by design is a methodology, not an absolute requirement (e.g., under the GDPR, privacy by design is qualified by what available technology is considered "state of the art," by the cost of implementation, and by the nature, scope, context, and purposes of processing, as well as by the risks for the individuals whose personal information is being collected and processed.

With this caveat, then, the ATP believes that privacy by design principles should be incorporated into PIPEDA in a manner that is consistent with the GDPR.

Discussion questions:

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?

1. The ATP submits that Privacy by Design should be implemented by a requirement for a PIA, but such a requirement must not include specific prescriptive elements in the PIA. This would make the concept more of an enforceable standard, which would be less open to interpretation. An advantage of AI is that the system can be evaluated in a repeatable fashion to expose bias or decision criteria that run counter to expectation of a standard or regulatory framework.
2. Requiring a PIA as part of a "privacy by design" requirement would make it equally appropriate and feasible for manufacturers and any other businesses to test their systems and processes. This is a reasonable expectation.

Proposal 6: Make compliance with purpose specification and data minimization principles in the AI context both realistic and effective

It is true that data minimization is generally at odds with the underlying tenet behind AI – which is predicated on having a maximum amount of data available from which to train the AI engine and then utilize the power of an AI system to analyze data. Similarly, the notion of specification of purpose arguably creates a practical problem for an AI system where the purpose(s) of data usage may not be known (or appreciated) until after huge amounts of data are collected and analyzed. As noted in the Information Accountability Foundation paper cited by the OPC, "the insights data hold are not revealed until the data are analyzed, consent to processing cannot be obtained based on an accurately described purpose." (OPC Request, fn 27). In essence, the ATP concludes that restricting the use of AI in advance would require the business to know in advance what is actually going to be determined – which is impossible. Thus, the challenge would be to limit the very personal information that is appropriate and needed for AI purpose(s).

Discussion questions:

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?

1. An AI system requires training to produce its reasoning models. In most cases that training occurs through the evaluation of data to discover relationships. While it is possible to constrain an AI system to only focus on certain relationships, this would essentially create a reinforcement of the expected outcomes of the entity that created the restrictions. Similar, while research is working on methods to learn from smaller data sets, and concepts like transfer learning allow for training based on prior knowledge outside the present data set, the reality is that a full range of data is what is required for

development of an AI system. Accordingly, the ATP submits that a balance must be struck so that the legal principles are not used to overly-restrict the multi-faceted functioning of an AI system.

2. One-sided legal restrictions on the use of data, provided other core principles related to the use of personal information as discussed throughout these comments are respected, would reduce the likelihood that AI innovation would continue in Canada as compared with other countries.
3. Alternatives and safeguards have been outlined in other responses in this document.

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

Although the ATP concurs that affirmative consent is the primary basis for collection and use of personal information under the GDPR and PIPEDA, we hasten to point out that there is a growing sense that the consent model may not be viable in all situations. In the ATP's opinion, that concern now includes uses of AI. This result is due in part to the inability to obtain meaningful consent when businesses are unable adequately to inform individuals of the purposes for which their information is being collected, used, or disclosed in sufficient detail.

The OPC's Report on Consent (*see* fn. 34) acknowledges that alternate grounds to consent may be acceptable in certain circumstances, specifically when obtaining meaningful consent is not practicable and certain preconditions are met. The ATP is unsure at this stage of the discussion about AI whether meaningful consent should be required at all – the experience of testing organizations so far seems to indicate that in many instances consent is rendered of no value to the resolution of determining the balance of the privacy rights of the consumer with the legitimate interests of the testing organization.[9]

Discussion questions:

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?

[9] The ATP agrees that the use of non-identifiable or de-identified data, such as through the application of pseudonymization and anonymization methods, also may be a factor in determining whether certain other grounds for processing such as legitimate or public interest should be authorized under the Act.

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?

The testing industry has realized that legitimate interest often provides the most practical approach to the collection and use of personal information, because it establishes the proper framework from which to evaluate the balance of rights and interests that exists between test takers (i.e., consumers) and testing organizations. Based on this experience, the ATP submits that the OPC should develop its proposals based on the GDPR approach to legitimate interests. We will reserve further comments until specific language is provided.

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

De-identification is achieved through processes that remove information that can identify individuals from a data set so that the risks of re-identification and disclosure are reduced to low levels. The ATP submits, that re-identification of such information is only a concern if the ability to re-identify is actually possible by the controller/business; if such re-identification is purely theoretical (e.g., the database of tokenized personal information is owned and under the exclusive control of another entity, then the controller/business has no realistic opportunity to conduct the re-identification.

Discussion questions:

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?

1. The ATP contends that de-identification and differential privacy – as well as opaque processing/encrypted processing – are all concepts where AI can provide benefits, where AI systems can manage risk, and where future value can be obtained. The risk of re-identification should be managed as suggested above through a question of control of the database.
2. Application of de-identification should result in a determination that such data is entirely outside of the scope of PIPEDA, subject to its re-identification by the controller or where the controller regains access to the original personal information.
3. The ATP believes that re-identification of personal information can arise in different ways and contexts, and result in different impacts. Any AI system working with de-identified or obfuscated data must be subject to analysis to determine if any re-identified information becomes associated with the AI system. As long as the controller/collecting business has no access to the re-identified information, the de-identified information should continue to be considered non-personal. However, if the controller/business regains use of the original personal information, then it automatically becomes subject to PIPEDA.

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

The ATP contends that principles of accountability, accuracy, transparency, and data minimization (as well as access and correction), support some limited tracing of the source of AI system data.[10] This result is especially appropriate for AI data that is NOT collected directly from individuals, but is supplied from other sources and combined into the AI analysis. We agree with the OPC, citing the OECD *Principles on Artificial Intelligence* (OPC Request, fn. 40): "AI actors should ensure traceability, including in relation to datasets, processes and decisions made during the AI system lifecycle, to enable analysis of the AI system's outcomes and responses to inquiry, appropriate to the context and consistent with the state of art." And as referenced, the IEEE has stated that, "algorithmic traceability can provide insights on what computations led to questionable or dangerous behaviors." However, it is completely premature for the OPC to rely on proposed legislation in the US, the *Algorithmic Accountability Act* (AAA), which is not likely to be enacted, at least in its current form.

Discussion question:

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?

1. To the extent AI technology makes data tracing possible – and so long as the balance between individual privacy rights and intellectual property protection is considered – this proposal may have merit. Moreover, the ATP also submits it is critical that any requirement must strike a balance between administrative burden and risk to avoid stifling innovation in an undue fashion.

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

ATP may be able to support the OPC's recommendation that a more robust accountability principle related to AI be included in Principle 4.1 of PIPEDA, depending upon how this notion is implemented. A requirement that a business maintains a record of its internal evidence demonstrating AI accountability for the personal information under its control could be appropriate if exercised through an independent audit requirement or as part of an enforcement activity. But a business should only be required to conduct an audit of its AI system once after it is fully operational and following any major changes to the system. However, the ATP would NOT support such a requirement if it means that evidence would have to be provided on demand to every individual consumer– that sort of requirement is likely to encourage many thousands of requests that would be extremely burdensome on businesses – and would likely create the potential for individuals to seek access to the business's IPR (*see* discussion related to Proposal 4, pages 6-7).

[10] The concept of AI tracing has limited application. For example, the OPC focuses specifically on the use of AI in test scoring when it cites (OPC Request, fn. 40) a 2014 law article arguing that "aggrieved consumers [test takers] … could challenge mischaracterizations and erroneous inferences that led to their scores." The ATP urges the OPC to recognize that the GDPR does NOT support the notion that derived scores even qualify as personal information. Moreover, the GDPR's right to rectify errors in a person's test scores does not apply to test item answers that were answered incorrectly during the test or to non-completed items of timed tests, or even to omitted responses on some test items (e.g., non-cognitive items) – these test responses are NOT personal information. *See* Opinion of Advocate General Kokott in *Peter Novak v. Data Protection Commissioner*, Case C-434/16 (July 20, 2017). Thus, it is very misleading to discuss the need for AI tracing as it relates to test scores.

As for shifting liability, the ATP fails to understand what the OPC is recommending – all violations of PIPEDA, including any enhanced accountability requirement, are issued against the covered business, not to a machine. In the ATP's opinion, it would be totally inappropriate to fine individual employees of the business, when they were acting within the scope of their employment. Similarly, developing regulatory incentives would apply to the business (e.g., for adopting demonstrable accountability measures).

Discussion questions:

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?

1. ATP suggests that record-keeping that would enable the business to recreate a prediction or action should be part of any system or process, AI or otherwise. Such records are required for accountability and for legal defense. Nevertheless, any suggested requirement of third party audits, proactive inspections by the OPC or other means to force demonstration of accountability should only be allowed if the perceived value is balanced against costs to society as a whole and businesses that would have to comply. The ATP might be willing to support AI audits, which like audited corporate financial statements, may well have value despite potential risks of over-reliance on audits that could dissuade businesses from investing in innovation.
2. Traceability, process transparency, and auditability have the potential to create some level of accountability and the trust, but we question whether those results would occur at a level that justifies the effort. At no point would the ATP support the establishment of a privacy right covering engagement with an AI system – beyond what is expected from engagement with any system or process (*see* discussion of Proposal 1). Nor do we believe that the OPC should mandate the creation of an "AI Auditor" position.
3. The ATP position can be best explained by examining what responsibilities and accountabilities may have been lost where an individual or an organization utilizes an AI system in a process or service delivered to a third party. We submit that businesses deploy AI systems as they do organic systems and any other process or system and they are responsible for the decisions reached and the actions taken.

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

The ATP understands the concerns of the OPC with the privacy risks that could be posed by AI systems, particularly if enforcement for organizations found to be non-compliant is not meaningful. However, the ATP disagrees with the OPC that Canada has "fallen significantly behind" other jurisdictions in terms of enforcement, especially if that statement is intended to incorporate the need for "enforcement mechanism that ensure individuals have access to a quick and effective remedy for the protection of their rights…." We assume the OPC is referencing the use of private right of action provisions to replace or augment government enforcement efforts.

Indeed, the ATP does not support the OPC's recommendation that, with respect to AI, it should be given authority to make binding orders and impose consequential penalties for non-compliance with the law. Again, this proposal would treat AI very differently than any other method or process for the collection and use of personal information – that result is not warranted.

Discussion questions:

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

1. The ATP submits that all eligible businesses are already responsible for their actions – and actions of agents, processors, and service providers are equally covered – under PIPEDA. An AI system is not different from any product, service, or even some new manufacturing technique in terms of being covered by the law and regulations. In fact, an AI system is in many ways more transparent then historic decision-making because it is possible to explore the biases and deterministic results in ways that cannot be accomplished with an organic subject matter expert. Future innovation is unlikely to be limited to mathematical algorithms, but may rest on other scientific developments (e.g., in neuroscience, robotics, quantum computing). PIPEDA should not penalize AI by treating it differently than all other types of decision-making, current or future.
2. The ATP firmly believes that AI-specific regulation will only serve to blur the conversation of what is AI as compared with other methods – this problem will hamper innovation.

CONCLUSION

In summary, the ATP encourages the OPC to align its final proposals related to the use of AI with the GDPR. We believe it is critical that AI not be treated any differently than other decisions made without the benefit of those technologies. Otherwise, we are concerned that innovation will be stifled.

Thank you for your attention to the important issues about the Proposals for modifying PIPEDA raised by affected members of the testing industry located within and outside of Canada. The ATP would be pleased to answer any questions the Office of the Privacy Commissioner may have in response to these comments, including to do so in a face-to-face meeting. For any follow up, please contact our General Counsel at the number or email address shown below.

Sincerely,

ASSOCIATION OF TEST PUBLISHERS

William G. Harris, Ph.D.
CEO
601 Pennsylvania Ave., NW
South Bldg., Suite 900
Washington D.C. 20004

Hazel Wheldon (Member, ATP Board of Directors)
CEO
MHS, Inc.
3770 Victoria Park Avenue
Toronto, Ontario
Canada M2H 3M6

Alan J. Thiemann
General Counsel
Han Santos, PLLC
700 12th Street, NW, Suite 700
Washington, DC 2005
(202) 904-2467
alant@hansantos.com