ITI Promoting Innovation Worldwide

March 12, 2020

ITI Comments to the Canadian Office of the Privacy Commissioner's Consultation on Privacy and Artificial Intelligence

ITI appreciates the opportunity to comment on the Canadian Office of the Privacy Commissioner's (OPC) Consultation on Privacy and Artificial Intelligence.

ITI represents the world's leading information and communications technology (ICT) companies. We promote innovation worldwide, serving as the ICT industry's premier advocate and thought leader in the United States and around the globe. ITI's membership comprises top innovation companies from all corners of the technology sector, including hardware, software, digital services, semiconductor, network equipment, and other internet and technology-enabled companies that rely on ICT to evolve their businesses.

Artificial Intelligence (AI) is a priority technology area for many of our members, who have developed and are using AI systems to improve technology, facilitate business, and solve problems big and small. ITI and its member companies believe that effective government approaches to AI clear barriers to innovation, provide predictable and sustainable environments for business, protect public and individual safety, and build public trust in the technology. We support the responsible development and deployment of AI in line with the preceding objectives. We also support the OPC's goals of ensuring that Canada's Personal Information Protection and Electronic Documents Act (PIPEDA) meets the needs of Canadian citizens, innovators, and policymakers as AI advances.

ITI and our member companies have long lauded PIPEDA for its flexible and forward-looking approach to data privacy, particularly because PIPEDA is both protective of privacy and supportive of economic growth and responsible innovation. The law can be applied to countless new and unforeseen commercial activities in a rapidly changing data environment, without the need for frequent legislative amendments. PIPEDA's effectiveness stems from the fact that it is principles-based, technology- and business-neutral, balanced, promotes organizational accountability and transparency, and focuses on meaningful control for individuals.

Given the transverse nature of AI, ITI believes that it is critically important that society, governments, and the technology sector work together to address the most complex issues it presents. We urge OPC to consider whether specific new regulations are necessary in order to adapt to a world with AI, or if OPC can achieve its policy objectives by adjusting PIPEDA to appropriately safeguard Canadian citizens, businesses, and innovators, without creating measures specific to AI.

As a general matter, we believe that changes in technology do not necessarily require changes to regulation or law, especially where existing regulations or other requirements may already adequately meet desired policy objectives. Similarly, given the rapidly evolving nature of AI technology and wide array of AI applications, regulation is difficult and prone to becoming obsolete quickly. As each AI application is different, there is no one-size-fits-all solution.

ITI also recommends that OPC's efforts be evaluated against a goal of ensuring that PIPEDA is interoperable with privacy frameworks deployed in other leading jurisdictions, including the

European Union. Interoperability does not require that the same exact obligations be adopted everywhere (for example, by replicating certain requirements), but rather by developing a regulatory framework compatible with global standards.

In analyzing the discussion proposals offered by OPC, ITI and its member companies offer the following recommendations and best practices:

Develop a Global Definition for AI

Regarding Proposal 1, it is important that OPC considers that currently there is no global consensus around a definition of AI in different areas such as business, policy, research and technology. AI systems are comprised of more than algorithms, and not all algorithms implicate AI; therefore, an essential factor in developing any AI strategy is to properly define AI, including the component parts of AI systems beyond algorithms, as well as to define related key terms such as machine learning. Since there is no single agreed-upon definition of AI, it will be important for policymakers to provide greater clarity if they plan to seek specific rules for AI functions, and to ensure that continued innovation and development in AI and other emerging technologies is not inadvertently stifled through well-intentioned, but premature, laws. Countries around the world are only just beginning to evaluate the use of AI in the context of personal information processing, and there is no consensus yet on if and how AI should be treated differently from other forms of technology. We accordingly recommend that the OPC should continue to monitor and participate in international work in this area, and as it reviews and fine-tunes approaches to personal information protection in this digital age, OPC should continue to consult with industry, who is best placed to advise on how regulation will affect the use of technology. To the extent that regulation or changes to PIPEDA are required, either should remain technology-neutral, so as not to inadvertently leave out one or more key components of the increasingly critical and dynamic technology landscape.

Focus on Activity rather than Function

Similarly, we encourage OPC to consider how it approaches certain activities, regardless of their basis in technology. For example, rules governing automated decision making should apply to all companies interacting with data via automated means, whether or not specific AI systems are employed. By focusing on outcomes, rather than motivations, Canada can ensure that PIPEDA continues to withstand the test of time and maintains sufficient flexibility to nurture innovation and emerging technologies without needing to be revisited each time there is a new technology.

Take a Sector- or Application-specific Approach

Both Proposals 2 and 6 should be considered in the context of a sector- or application-specific approach, rather than a horizontal one-size-fits-all approach across the economy. This type of approach allows for a more careful calibration of policy or regulatory measures to identified risks. In our view, this is a more agile approach - evaluating sector-specific legislation to identify what legislative gaps exist and the extent and manner in which any such gaps should be filled before assessing the need for upgrading the regulatory framework to enable AI to fulfil its potential.

Consider Context in Applying Consent as well as Alternative Legal Bases

We agree that explicit consent is often an appropriate basis for processing data, but as other global privacy regimes such as the GDPR acknowledge, obtaining consent should not be the sole basis or the presumed basis for processing data. There are many legal grounds for processing other than consent, including fulfillment of a contract, compliance with a legal obligation, consideration of individual or public health and safety, support for network and information security, conducting research and measurement, and prevention of fraud. Those types of legitimate or public uses of personal data should be considered as alternative bases for processing dependent upon the data processing context.

Determining when consent is required should consider the context of the interaction between the customer and business. Context is a foundational consideration for identifying appropriate policies, especially as related to Proposal 3 on automated decision-making and Proposal 7 on legal grounds for processing. Processing of personal information should be permitted where reasonable, relative to the context concerning the collection and processing of data and in accordance with key criteria, such as 1) the extent, frequency, nature, and history of interactions between individuals and covered entities providing data technologies/services; 2) expectations of reasonable individuals about how a covered entity collects data; 3) the extent to which the collection or processing of data is necessary for the performance of the product or service requested; or 4) the benefits of the collection and processing of data. For example, a legal basis other than consent must be established for processing personal data in AI models to the extent strictly necessary and proportionate for the purposes of ensuring network and information security. This is a good example of an instance where consent provides an inadequate legal ground for processing, both because there is a clear legitimate interest or public interest, and consent would not be meaningful or possible to obtain.

As with other privacy-related frameworks, including legal bases and exemptions to collect or process data for security and public safety purposes are important. The cybersecurity industry today is integrating AI/ML technologies into products and services to defend against evolving threats. Indeed, as the cost of computing continues to decline, digital adversaries are able to successfully conduct increasingly automated attacks at minimal cost. In response, network defenders are turning to automated, AI/ML-based defenses – bringing software to a software fight – to effectively defend against these attacks. Hindering the ability of cybersecurity providers to leverage AI technologies weakens their ability to defend organizations globally – including Canadian government agencies, customers, partner entities, and supply chains – against cyber adversaries.

Embrace Pseudonymization, De-Identification, and Other Privacy Enhancing Techniques
Proposal 7 discusses consent and options for encouraging meaningful consent for data processing. Pseudonymization and de-identification are critical tools for allowing the marketplace to work by enabling data processing without infringing on personal privacy.

Also, a technical area of research and development that offers promise for addressing the potential trade-off between data access/sharing for AI purposes and privacy is Privacy Preserving Machine Learning (PPML), and more broadly, Privacy Enhancing Techniques (PETs). PPML refers to a set of techniques that use cryptographic, mathematical and statistical tools to help perform AI tasks while limiting the access and sharing of personal information.

The ability to analyse and measure data – for example, to determine online ad delivery success and market concentration (in the latter instance, by making sure antitrust regulators have sufficient data to assess marketplace effects) – is also critical to modern day commerce and market regulation. Pseudonymization allows for measurement to occur in a way that benefits consumers, businesses, and regulators, and it should be preserved as an important component of consent and data processing.

Promote Reasonable Accountability Requirements

Our industry is committed to partnering with relevant stakeholders to develop a reasonable accountability framework. As leaders in AI technology, ITI members recognize the important role they play in making sure technology is built and applied for the benefit of everyone. Approaches towards regulating AI and other technology must be context- and risk-specific and should take into account that not all applications require an all-encompassing fundamental rights-based approach. Some basic AI uses have little or no impact on individuals' rights, such as in the context of industrial automation and analytics to streamline automobile manufacturing or to improve baggage handling and tracking at busy European airports. Many other uses – e.g. in medicine, financial services or transport – are often subject to significant sectoral regulation. While it is important to assess whether applicable sectoral rules are exhaustive, it should also be recognized that they cover many of the most common concerns such that further regulation is unnecessary.

Carefully Approach Explainability and Transparency

When dealing with the issue of explainability in Proposal 4, OPC should seek to ensure that they are not creating a system that establishes an environment where statistical outliers are viewed as a flaw in an overall AI system. Oftentimes, if an outlier is indeed an outlier, then an AI model will learn and dismiss it in later iterations. In many cases, finding a satisfactory explanation as to why an individual result was produced is not technically possible or relevant if the AI system dismisses the findings on its own. Guidelines should capture a statistically meaningful number of results to ensure uncertain results are actual concerns and not just isolated anomalies. Ongoing standardization work developed by international standardization organizations (ISO/IEC JTC1/SC42) appears promising in terms of defining parameters for explanation and will have increasing relevance in addressing explicability measures.

There are certain situations in which users and consumers can benefit from an AI system's explanation. Most of these scenarios involve systems that are either deployed in a commercial context or are making decisions that have a materially significant impact on a user's well-being, such as an individual's life or liberty. As such, AI systems deployed on the market that make determinations or recommendations with those types of potentially significant implications for individuals should be able to explain and contextualize their conclusions. Such situations are often highly context-dependent, however, and any proposed rules seeking to address the topic of explainability need to take account of the particular risks presented by an AI system's decision-making capabilities in the context of the specific AI application and its potential impact on an individual.

Similarly, the requirement in Proposal 4 (as well as 5 and 9) to disclose algorithms is extremely concerning. Algorithms are intellectual and proprietary property and closely guarded as trade secrets, and there are many legitimate reasons that companies may not want to reveal their AI algorithms. For example, in the security industry, if an adversary learns the details of an AI-based network defense capability, they will be able to defeat the security tool. If the models that security companies use are made public, digital adversaries could craft malware specifically to defeat it. In the financial services industry, as another example, unauthorised knowledge of trading system algorithms could lead to adverse and illicit gaming of the system.

Use Risk-Based Approaches to Privacy by Design and Bias

Proposals 5 and 9 deal with privacy by design and bias. It is imperative that policies in these areas continue to build on PIPEDA's existing risk-based approaches so they can be tailored to individual companies' business models, including how their technologies and services interact with individuals, without being overly prescriptive. In considering bias, it is important to understand how different companies deal with bias and take steps to prevent it. For example, machine bias may be introduced at various stages due to characteristics inhering in how the AI system is programmed to make decisions at the architecture, design and training stages. Training sessions, in particular, are potentially subject to the inherent biases of the humans who may not have been properly trained to mitigate those biases. Biases can also exist in the data sets used to train AI algorithms themselves. While bias cannot be removed completely due to the unavoidable human inputs to AI, it is important to regularly assess AI systems for bias and disclose it once identified. Once bias in AI systems can be better documented then policymakers can assess what level and what types of bias are acceptable. ITI encourages policymakers to consider that AI/ML-based applications, with proper supervision over input and outputs, have the capability to mitigate pervasive human bias in decision making, allocation of resources and benefits and adjudication. Responsibly used, AI has the potential to mitigate against disparate treatment and impact derived from human-based systems.

Thank you for your consideration of these comments, and we look forward to continuing to serve as a resource on these important policy matters now and in the future.

Sincerely,

Ashley E. Friedman
Senior Director, Policy
Information Technology Industry Council