ITI Promoting Innovation Worldwide

28th November 2023

ITI Response to Bill C-27 – Digital Charter Implementation Act

The Information Technology Industry Council (ITI) is the premier voice, advocate, and thought leader for the global information and communication technology industry. Founded in 1916, ITI is an international trade association that promotes public policies and industry standards that advance competition and innovation worldwide. Our members include the world's leading innovation companies, with headquarters worldwide and value chains distributed around the globe. These companies are leading Internet services and e-commerce companies, wireless and fixed network equipment manufacturers and suppliers, computer hardware and software companies, and consumer technology and electronics companies. Artificial Intelligence (AI) is a priority technology area for many of our members, who develop and use AI systems to improve technology, facilitate business, and solve problems big and small.

ITI supports Canada's overall goal of building an ecosystem of trust in AI through forward-thinking approaches to governance, recognizing the need to ensure that AI systems are developed and deployed in a fair, transparent, accountable, and privacy-respecting manner. Finding the right balance in a legal framework is especially challenging when considering a novel technology that is constantly evolving.

ITI and its members share the firm belief that building trust in the era of digital transformation is essential and agree that there are important questions that need to be addressed regarding the responsible development and use of AI technology. As this technology evolves, we take seriously our responsibility as enablers of a world with AI, including seeking solutions to address potential negative externalities and helping to train the workforce of the future. Our members are aware of and are taking steps to understand, identify and treat the potential for negative outcomes while leveraging benefits that may be associated with the use of AI systems.

ITI is actively engaged in AI policy around the world. In 2021, we issued a set of *Global AI Policy Recommendations*, aimed at helping governments facilitate an environment that supports AI while simultaneously recognizing that there are challenges that need to be addressed as the uptake of AI grows around the world.[1] In 2022, we also published our *Global Policy Principles for Enabling Transparency of AI Systems*[2] where we underscore that transparency is a critical part of developing accountable and trustworthy AI systems while avoiding unintended outcomes or other harmful impacts. We have also actively worked to inform the efforts of the National

[1] Our complete *Global AI Policy Recommendations* are available here: https://www.itic.org/documents/artificial-intelligence/ITI_GlobalAIPrinciples_032321_v3.pdf

[2] Our complete *Policy Principles for Enabling AI Transparency of AI Systems* are available here: https://www.itic.org/documents/artificialintelligence/ITIsPolicyPrinciplesforEnablingTransparencyofAISystems2022.pdf

Global Headquarters
700 K Street NW, Suite 600
Washington, D.C. 20001, USA
+1 202-737-8888

Europe Office
Rue de la Loi 227
Brussels - 1040, Belgium
+32 (0)2-321-10-90

info@itic.org
www.itic.org
@iti_techtweets

Institute of Standards and Technology (NIST) to foster trust in AI technology in the U.S. [3] ITI has also been engaged in AI regulatory efforts in other regions of the world. In the EU, we have been engaged in the development of the EU AI Act since 2020 and continue to weigh-in as trilogue negotiations progress. In the UK, we have participated in consultations on the UK's approach to AI and interacted with key legislators. In Latin America, we have actively engaged in different AI-related consultations and public hearings held by the Brazilian Congress and will continue to contribute to Brazil's legislative efforts as Congress seeks to pass legislation.

In 2023, ITI launched the AI Futures Initiative, a task force of AI policy and technical experts from ITI's member companies, which is focused on producing actionable policy recommendations that support the responsible use and deployment of AI.

We are grateful for the opportunity to provide feedback to the Standing Committee on Industry and Technology, House of Commons (Committee) on the Digital Charter Implementation Act. Below we provide comments and feedback to the Consumer Privacy Protection Act and the Artificial Intelligence and Data Act (AIDA).

A. Consumer Privacy Protection Act (CPPA)

ITI acknowledges the comprehensive approach of the proposed Consumer Privacy Protection Act (CPPA) in its focus on protecting the privacy rights of individuals while balancing the operational needs of organizations. Notably, the CPPA recognizes the necessity of providing flexibility to allow businesses to process data for legitimate purposes. This ensures that organizations can continue to innovate and operate efficiently, aligning the protection of privacy with the dynamism in the realm of data and technology. We are also encouraged by the introduction of a safe harbor program under the code of practice, an initiative that underscores the Act's adaptability and attentiveness to varied organizational needs.

However, we remain concerned about several provisions which we have included below:

1. We recommend revisiting the private right of action provision to ensure that it does not inadvertently encourage litigation and hinder innovation. A balanced approach, promoting compliance and education, should be the cornerstone of the enforcement mechanism.

2. We also recommend reconsidering linking the assessment of penalties to global revenue calculations. Civil penalty frameworks should not be broadly tied to global revenue but should instead be more closely proportionate to harms to affected individuals. In jurisdictions that have adopted a global revenue-based approach, it has operated as an excessive deterrent that ultimately has chilled innovation and related consumer benefits. With this in mind, the proposed approach to penalty assessment should be reconsidered.

[3] See ITI response to RFI on AI RMF Concept Paper here: ITI Comments on AI RMF Concept Paper FINAL.pdf

3. We recommend slight revision to the definition of "anonymization" to ensure technical feasibility. We are concerned that the definition creates an impractical standard rather than a risk-based definition that seeks to reduce the risk of identifiability to a negligible level, taking into account generally accepted best practices.

4. We also recommend amending the data re-identification provisions so that organizations can use re-identified information where they can rely on consent or an exception to consent. Data de-identification and re-identification is a common practice to safeguard and minimize data and could not be used in these ways if it is an offense to re-identify the de-identified information.

5. Additionally, given the substantive changes and enhancements introduced by the CPPA, we recommend allowing organizations ample time to transition to the new privacy regime. A phased implementation will facilitate a more seamless adoption, ensuring that businesses can effectively integrate and uphold the enhanced privacy standards.

B. Artificial Intelligence and Data Act (AIDA)

General Feedback

The benefits of AI are vast and multifaceted. AI driven medical diagnostics can alert doctors to early warning signs to help them treat patients. AI systems are capable of monitoring large volumes of financial transactions to identify fraud. Small and medium-sized enterprises (SMEs) can gather new insights and improve their businesses by using AI and data analytics. AI-powered cybersecurity solutions help organizations actively defend, automatically respond to, and recover from unmanaged Information Technology (IT) infrastructure risks, without manual work. Recommender systems driven by AI provide content to consumers that is useful and valuable to them. Such technological advances can bring innumerable benefits to Canada. At the same time, promoting the responsible growth of AI is key while guarding against potential negative impacts. The AI ecosystem is global, and the technology is not developed in regional siloes. Therefore, Canada has an opportunity to take an international leadership role in promoting advances in AI.

Below, we offer some general thoughts to Part 3: Artificial Intelligence and Data Act (AIDA) of C-27.

We urge the Committee to consult with all stakeholders before updating the AIDA draft – including the amendments proposed by the Minister of Science, Innovation, and Industry ('Minister') on October 3, 2023. We understand that the Minister has proposed significant amendments to the AIDA section of C-27 in light of the rapidly changing AI environment. While the proposed amendments are still being considered by the Committee, we have provided our feedback and addressed portions of the amendments in the appropriate section. In order for

stakeholders to provide more fulsome feedback on proposed amendments, we request the Committee consult all stakeholders before updating the draft.

We urge the Committee to more clearly articulate a risk-based, pro-innovation approach in AIDA. We encourage governments around the world to take a risk-based approach with respect to AI regulation, as not all applications or uses of AI will pose the same level of risk and require the same level of regulatory oversight. An approach that also prioritizes innovation will allow AI to address the most pressing societal challenges in areas such as healthcare, public security, cybersecurity, economic growth, climate change and disaster management. Promoting these advances is no less important than managing the challenges. We agree with Canada's assertion that regulators must focus on high-risk situations that result in serious harm to individuals or their interests associated with the development, deployment and use of AI. However, in order to provide greater clarity, we request the Committee modify Section 4 (Purposes) to clearly state that it intends to take a risk-based approach, where obligations are focused on high-risk AI systems. While we appreciate the Minister's efforts to define what constitutes a high-impact system, we recommend that this list be filtered further to more narrowly focus on uses within the outlined classes that could have a significant negative impact on people, especially related to health, safety, freedom, discrimination, or human rights. We also encourage the Committee to consider how it might work with international counterparts pursuing similar efforts to ensure that approaches are as interoperable as possible.

The Committee should allow for flexibility in its approach. An approach that prioritizes a flexible, risk-based application of obligations will enable innovation in Canada and allow for adjustments as the technology evolves. A regulatory approach that focuses on outcomes, as opposed to one that mandates prescriptive requirements, will enable flexibility, allowing organizations to approach obligations in the way that is most appropriate for them. AI regulation should support and build on ongoing efforts to establish best practices in the field of responsible AI development and deployment, rather than risk undermining these efforts by prematurely codifying inflexible mandates in a field that is rapidly evolving.

We encourage the Committee to continue stakeholder engagement as it advances its approach to AI – particularly with developers, designers, deployers, and others as well as those impacted by these technologies. The Committee should maintain ongoing and consistent conversation with those stakeholders who are developing, designing, deploying and operating, as well as those using AI technologies as it seeks to further hone its approach. Such conversations will provide a robust, well-rounded understanding of the landscape and allow for exchange amongst all relevant stakeholders. This consultation is a helpful first step in facilitating that exchange.

The Committee should reference global, voluntary, industry-led technical standards in AIDA. AI technical standards are essential to increase interoperability, harmonization, and trust in AI systems. They can inform AI regulation, practical implementation, governance and technical requirements. Governments should work to support global, voluntary, industry-led standards,

and safeguard the work and processes of international standards development bodies. Broad contributions to and adoption of international standards reduces market access barriers. Standards work for the net benefit of the international community and should be developed and applied without prejudice to cultural norms and without imposing the culture of any one nation. Standards work should also be technology neutral (avoiding preferential treatment for any specific technical approach). We therefore request the Committee include text in AIDA that recognizes the important role of voluntary, industry-led technical standards.

Specific Feedback

1. **Section 2 (Definitions)**

a. Artificial Intelligence System

The Committee should avoid creating a broad definition of AI, which could encompass nearly all modern software systems, including automation. However, just as a definition should not be too broad, it should also not be too narrowly focused on a detailed and prescriptive description of the underlying technical elements of AI and machine learning. These are dynamic and continuously evolving fields, and any attempt to encapsulate their technical details will inevitably and rapidly become outdated.

An Artificial Intelligence system is a broad framework that encompasses an AI model along with the necessary infrastructure and components to develop, implement and utilize the model effectively.[4] For example, generative AI applications such as ChatGPT, Google Bard and Anthropic's Claude are AI systems consisting of an underlying AI model, as well as a user interface, an API and various other features such as plugins that interact with the model.

One definition of AI worth considering is the updated November 2023 definition of AI system offered by the Committee on Digital Economy Policy at the Organization for Economic Cooperation and Development (OECD):

"An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.".[5]

This definition is notable in that it addresses the need to focus on learning systems, while also taking into account the AI development lifecycle. We appreciate that the proposed amendment from the Minister aims to align the definition of AI with the EU AI Act and the OECD. However, it also seems to define an AI model in the context of a "technological" system. To reconcile this

[4] https://www.itic.org/documents/artificial-intelligence/ITI_AIPolicyPrinciples_080323.pdf
[5] https://oecd.ai/en/ai-principles

difference, we strongly urge the Committee to consider adhering more closely to the above OECD definition of an AI system.

2. Section 4: Purpose of the Act

ITI supports the overall goal of building a thoughtful, proportionate, and risk-based approach to AI governance. However, it is important to note that AI – like many other technologies – is constantly evolving. A regulatory structure for AI must be able to dynamically recalibrate so as not to stifle beneficial innovations.

A clearly defined risk-based approach is required to avoid a "one size fits all" solution that may have unintended consequences on a diverse range of AI use cases. This approach has also largely been accepted in the EU in its upcoming AI Act and the UK's White Paper on AI. We therefore request the Committee reconsider the purpose of the Act, clearly including language that reflects its intent to take a risk-based approach.

Additionally, the purpose of 'regulating interprovincial trade and commerce in AI' seems somewhat adjacent to what we understand to be the overall purpose of AIDA – preventing harms that may come from the application of AI systems in Canada.

3. Section 5 (Definitions)

a. Biased output

While the global conversation around bias has been largely focused on the negative impacts of bias, as has been noted in NIST's Special Publication 1270, some types of bias "are purposeful and beneficial."[6] For example, machine learning systems that underpin AI applications can in certain instances model implicit biases to improve online shopping experiences or recommend preferred content. However, we agree that biased outputs can lead to harms (discussed in the next section) by potentially negatively impacting an individual's health, safety, freedom, or human rights. To ensure there is clarity around the definition of biased output, the Committee should align the meaning of biased output with existing Canadian Law – including the Canadian Human Rights Act by replacing the word "adversely" with "materially and unlawfully."

b. Harm

Policymakers, in close consultation with industry and other stakeholders, should consider how to characterize "high-risk" or "high impact" applications of AI, including by identifying the appropriate roles for AI developers, deployers, and users in making risk determinations. In general, we prefer the use of the term "high-risk" as this aligns with the way in which other regulatory and policy approaches are being construed, and as stated elsewhere, we think using

[6] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf

the same terms across jurisdictions will help to solidify a common lexicon and foster interoperable approaches. We reiterate our view of a "high-risk" system: an AI system is high-risk when a negative outcome could have a legal or similarly significant impact on people—especially as it pertains to health, safety, freedom, discrimination, or human rights. In thinking about high-risk applications, focusing on "sectors" may lead to overly broad categorizations – it is important to use a sufficiently targeted and well-outlined classification to ensure this criterion does not become irrelevant.

The current draft includes "psychological" harm which may have an impact on an individual's well-being through the use of AI. However, additional clarity is needed with regard to what legal standard or process would be used to demonstrate "psychological harm."

The Committee should therefore revise its definition of harm to include situations when there is a negative outcome which could have a significant impact on an individual's health, safety, freedom, or human rights. Harm should mean an unintended effect from the use of a high-risk AI system which poses a significant impact on an individual's health, safety, freedom or human rights.

c. High impact system

As highlighted earlier, we encourage the Committee to take a risk-based approach with respect to AI regulation, as not all applications or uses of AI will pose the same level of risk and require the same level of regulatory oversight. We agree with the assertion that regulators should focus on systems that pose a high-risk concern rather than hypothetical or low risks associated with AI. However, in order to provide greater clarity, we request the Committee provide a definition of high impact systems within AIDA. In particular, the Committee should establish a clear and precise definition of "high impact system" that is closely tailored to risk, clarifies the harms to individuals that policymakers are intending to mitigate, and provides legal certainty to all parties in the AI value chain regarding the development, deployment and use of AI systems instead of leaving it to future regulation.

We appreciate that the proposed amendments by the Minister attempt to provide a list of classes that would be classified as "high impact" due to their effect on health and safety of users. In addition to describing certain classes as high impact, it is essential that the Committee consider the context in which systems are deployed before classifying whole activities as high impact. For example, the proposed class to provide 'services' to individuals will need to be appropriately scoped as the term is extremely broad.

We are particularly concerned with the inclusion of broadly defined systems that "prioritize" content as high impact. By classifying this activity as "high impact," there may be unintended consequences on innovating to address critical use cases, including the use of automated systems for safety and integrity purposes. In addition, the inclusion of those systems among the list of

high impact AI systems does not appear to be in alignment with the other AI systems included in the list of high-impact systems. For example, the high-impact list includes AI systems in areas such as employment, identification of an individual, health care services, administrative decisions about an individual and assistance of peace officers, in which a negative outcome could have a significant legal or similar impact on people's human rights, safety, or access to services. In our view, automated systems used for "content moderation" or "prioritization of content" are not comparable to those other high-impact automated decisions. Additionally, the Committee should consider that organic content prioritization is already regulated by general data protection laws.

Further, with respect to the use of AI for content moderation, we recommend that measures that target at fostering transparency around the content moderation processes be consistent with the nature of the service, consider the context and risk profiles of services, provide for fair and non-discriminatory application, and strike the right balance with freedom of expression, rights of third parties, or potential impacts on other players for a healthy online ecosystem.

Finally, as referenced earlier, we strongly urge the Committee to further filter the overarching list to focus on uses within the outlined classes that **could have a significant negative impact on people, especially related to health, safety, freedom, discrimination, or human rights**, as there may be benign use cases even within the outlined classes. This will allow impacted stakeholders to assess whether their system might fall into the "high impact" category using clear pre-defined criteria.

d. *Person responsible*

In the current draft, the definition for "person responsible" is quite broad, and it is not clear how the requirements would apply to a person that is designing, developing, or deploying an AI system, as opposed to one that is "managing" it. For example, it is unclear who the responsible party would be in an instance where an organization "develops" a system and then sells it to someone who "manages" it: would both parties be responsible for meeting the requirements under the legislation, or would only one party be responsible? If so, which party?

Stakeholders throughout the value chain play a role in the development and deployment of AI in a responsible manner. Responsibilities should be allocated among actors based on their role and function in the AI value chain and it would be beneficial for the Committee to more clearly delineate which obligations fall on the developer and which obligations fall on the deployer. Access and use of AI systems can also be supported by open innovation or open-source models, but this requires clear and reasonable rules regarding developer and deployer obligations.

To help the Committee refine its definition of 'Person Responsible', we encourage lawmakers to reference ISO/IEC 5339 IEC 5339 Information technology—Artificial intelligence— Guidelines for AI applications, an international standard which is currently under development, and which provides guidelines for identifying the context, opportunities, and processes for developing and

applying AI applications. The guidelines provide a macro-level view of the AI application context, the stakeholders and their roles, relationship to the lifecycle of the system, and common AI application characteristics, properties, and considerations.

4. **Section 7 (Assessment – High Impact System)**

The current draft requires that the 'person responsible for an AI system' conduct an assessment as to whether it is a high impact AI system. Much of this assessment will depend on the criteria that the bill uses to determine what constitutes a 'high-impact system,' which is presently left up to future regulation. As we mention above, for legal certainty, it is critical for the Committee to include criteria within AIDA that specifically delineates what constitutes a high-impact AI system.

5. **Section 11 (Publication of description – making system available for use & managing operation of the system)**

The current draft requires the 'person who makes available for use a high-impact AI system' to provide a specific set of information about the high-impact AI system. While transparency is an important principle in fostering trust in the use of AI systems, it is not an end in itself — it is a means by which to enable accountability, empower users, and build trust. In designing transparency requirements, it is important that the Committee consider its objectives, and what the best way to achieve said objectives are. Importantly, in considering any future transparency requirements, the Committee should keep in mind the need to protect sensitive IP, trade secrets and ensure that any publicly available information does not reveal information or data that could undermine cybersecurity.

6. **Section 8 and 9 (Measures related to risks and monitoring of mitigation measures)**

The current draft requires the 'person who is responsible for a high-impact system' to manage risk and monitor mitigation measures in accordance with regulations that may be established under separate rulemaking. There are multiple stakeholders in the AI value chain who play a role in the development and deployment of AI in a responsible manner. Responsibilities should be allocated among actors based on their role and function in the AI value chain.

We appreciate how that the Minister's proposed amendments seek to provide clarity around the allocation of responsibility across the AI value chain by replacing Sections 8 and 9 with new text that would delineate which responsibilities 'persons responsible' should undertake and by distributing those responsibilities out to different actors in the AI value chain.

In the section titled "Aligning AIDA with the EU AI Act and the OECD by making targeted amendments to key definitions," the Minister suggests adding sections that would ensure that persons developing a machine learning model intended for high impact use take appropriate measures before it goes on the market (either by itself or as a part of a high-impact system).

These responsibilities can be classified as responsibilities of a "developer." This section also lays out responsibilities for the "deployer", which includes "persons placing on the market…would be responsible for ensuring that necessary measures with regard to the development were taken prior to the system entering the market. We recommend the Committee replace the term "development" with "deployment" so as to clearly reflect the different roles and responsibilities of those in the AI value chain, including developers and deployers. We note that the amendments also differentiate between those "developing" high-impact systems, those "making available high-impact systems," and those "managing the operations of high-impact systems." In some cases, the developer may also be the person that is making available and managing the operations of the AI system, a deployer may be both the person making available and managing the operations of the AI system, or there may be distinct actors undertaking each individual activity. We encourage Canada to further define these terms and also note that in certain instances these parties may be the same.

In particular, we further suggest the Committee define terms such as "deployer" and "developer" so that roles and responsibilities are clearly demarcated. We recommend the following definitions:

A developer (sometimes used interchangeably with producer) is the entity that is producing the AI model or system. The developer of an AI model or system is in control of certain information and decisions, e.g., how the model's training data is selected and used, what kind of testing and validation is performed on the model, etc. Accordingly, developers are best positioned to manage model-level risks and understand the capabilities and limitations of a particular model. In many instances, an AI model can be built into other products that are then deployed by a different entity.

A deployer (sometimes also called a provider) is the entity that is deciding the means by and purpose for which the AI system or model is ultimately being used. Deployers often have a direct relationship with the consumer. While developers are best positioned to assess, to the best of their ability, and document the capabilities and limitations of a model or system, deployers, when equipped with necessary information from developers, are best positioned to document and assess risks associated with a specific use case.*

Furthermore, in seeking to understand and further delineate roles and responsibilities in the context of future legislative or regulatory approaches, we encourage the Committee to reference ISO/IEC 5339 Information technology—Artificial intelligence— Guidelines for AI applications, an international standard which is currently under development, and which provides guidelines for identifying the context, opportunities, and processes for developing and applying AI applications. The guidelines provide a macro-level view of the AI application context, the stakeholders and their roles, relationship to the lifecycle of the system, and common AI application characteristics, properties, and considerations.

7. Section 29 and 30 (Administrative Monetary Penalties and Contraventions)

We recognize that policymakers are concerned about liability for AI systems. At the same time, it is a complex topic area that is not always effectively addressed via product liability law given the unique characteristics of AI systems. Indeed, software in general, and AI in particular, relies upon complex supply chains that include multiple actors throughout its lifecycle, which include developers, deployers, and potentially others (i.e., distributor, producer, professional or private user). Some of these actors may not be aware of the existence/role of other actors or may be unaware of the ways in which another actor might be using their products or services. As such, we recommend the Committee revise Section 29 and Section 30 to provide a safe harbor for organizations that have conducted due diligence and to ensure penalties are proportionate to the contraventions.

In keeping with the concept of proportionality, the Committee should review the amount of the proposed sanction, which is in our view excessive and could serve to jeopardize other policy goals, such as promoting innovation. In an area that is still evolving in many aspects and where the world is still developing guidelines and frameworks, Canada's approach could create a disproportionate sanction, endangering the development of technologies locally. The sanction amount, coupled with the lack of clarity within the definitions and the lack of clarity related to the distribution of responsibilities in the value could very significantly harm innovation.

Further, criminalizing artificial intelligence, particularly when a robust set of consensus-based industry standards have not yet emerged, will undermine the development, commercialization and adoption of AI technology in Canada. We recommend amending criminal liability provisions so that they apply only in the case of intentional use of an AI system to cause serious physical harm or serious fraud.

8. Sections 13, 14 and 18 (Ministerial Orders and Publication)

The current draft provides powers to the Minister to require covered entities to provide information, particularly for those systems whose uses are classified as high impact. We request the Committee to include language that will ensure that there is no mandatory disclosure of source code, intellectual property or any information that could jeopardize cybersecurity.

9. Proposed Amendment (Distinct obligations for general purpose AI systems)

We encourage the Canadian government to review ITI's recently published *Guide on the AI Value Chain & Foundation Models*,[7] which delves more deeply into definitions and offers key policy considerations associated with this AI technology. In creating additional obligations for "general purpose AI systems," we encourage policymakers to consider what **unique risks** they are trying to address in doing so, particularly because not every application of a general-purpose

[7]https://www.itic.org/documents/artificial-intelligence/ITI_AIPolicyPrinciples_080323.pdf

AI system is going to be high-risk. The Committee should additionally recognize that risk management is a shared responsibility between multiple stakeholders in the value chain. **Information-sharing and transparency in the value chain is critical.** We agree that developers of foundation models can take certain measures prior to release to increase safety and security, increase transparency about the model's capabilities, and contribute to reducing risks from their downstream application to the extent possible. However, it is crucial to recognize that effective risk-management is ultimately context specific and, in many cases, depends on the circumstances of the deployment or use of the AI system. Because foundation models can be used for myriads of downstream tasks, foreseeing potential risks arising from every use and determining whether such uses constitute a high-risk application will not be feasible or possible. In keeping with a risk-based approach, foundation models and/or general-purpose AI systems should not be treated as de facto high-risk AI applications.

Given that the landscape around controls and guardrails for foundation models is evolving quickly and standards are being developed, an overly prescriptive approach to risk management at the development stage for foundation models would undermine ongoing research and risk stifling innovation. In seeking to levy obligations on foundation model developers, the Canadian government should seek to support overarching outcomes that developers of foundation models should meet in order to increase the overall safety and effectiveness of their models. Outcomes could include ensuring that the data used to train the models is relevant and high quality to the extent possible, the assessment and management of known risks, including related to bias, providing transparency about data and model design, and establishing appropriate technical documentation for downstream providers without compromising sensitive IP.

Ultimately, any obligations placed on the foundation model developer should be calibrated to the level of risk it poses, be practicable, and only extend to what a foundation model developer can reasonably address during design and development, rather than all potential risks that downstream application may present.

***

ITI Member Companies