

The AIDA companion document: A roadmap for Canada's artificial intelligence law

March 30, 2023

Introduction

The [Artificial Intelligence Data Act](#) (the AIDA), which is being introduced by the Canadian government alongside the [Consumer Privacy Protection Act](#) (CPPA) as part of Bill C-27, the Digital Charter Implementation Act, 2022, is Canada's first attempt at regulating artificial intelligence (AI). On March 14, 2023, Innovation, Science, and Economic Development Canada (ISED) published the [AIDA companion document](#) (the companion document), which is particularly helpful since the AIDA draft currently being considered by the government leaves many core concepts to be defined through regulation (e.g., what may be considered as a high-impact system and triggers the implementation of risk mitigation measures).

In this article, we summarize the companion document and analyze how it supplements our understanding of the AIDA.

The AIDA: Purpose, scope and applicability

The AIDA's two stated proposed purposes are:

1. To regulate international and interprovincial trade and commerce in AI systems by establishing common requirements, applicable across Canada, for the design, development and use of AI systems.
2. To prohibit conduct that may result in serious harm to individuals or their interests.

These purposes are quite broad on their own, and even more considering the AIDA's broad definition of "AI systems". Specifically, it defines an AI system as "a technological system that, autonomously or partly autonomously, processes data related to human activities through the use of a genetic algorithm, a neural network, machine learning or another technique in order to generate content or make decisions, recommendations or predictions" (s. 2).

Despite this broad definition, the companion document makes important clarifications regarding the limits of AIDA’s scope. First, as discussed further below, the companion document clarifies that the AIDA would only apply to fully-fledged AI systems and not models or other open-source tools that could be used to develop an AI system.

Further, the companion document clarifies that the AIDA is intended to work with - not supplant - existing legislation that would apply to AI systems. For example, the [Canada Consumer Product Safety Act](#) can be a source of exposure to manufacturers of consumer products incorporating AI technologies, as it is meant to address and prevent dangers to human health or safety that are posed by consumer products in Canada. Health Canada has also [offered guidance](#) affecting digital health products and medical devices under the rubric of Software as a Medical Device, an area in which AI technology will have significant impact.

It can be expected that the general approach to regulating AI with the AIDA will eventually carve out specific areas of application, perhaps in the automotive and medical device spaces. The companion document also clarifies that the AIDA’s proactive, risk-based approach to AI regulation is intended to be compatible and align with foreign legislation, like the EU’s proposed [AI Act](#). It was informed by and reflects AI industry norms, such as the guidelines published by the US National Institute of Standards and Technology (NIST) [Risk Management Framework](#) (RMF).

High-impact systems, potential harms and biased output

Despite applying to virtually all advanced systems, the AIDA’s more stringent requirements target “high-impact systems”. Anyone responsible for an AI system - namely a person who designs, develops, or makes available for use the AI system or manages its operation - must assess whether it is a high-impact system in accordance with regulations that are yet to be drafted. The companion document notes that ISED considers the following to be the key factors that persons will use in making this determination:

- Evidence of risks of harm to health and safety, or a risk of adverse impact on human rights, based on both the intended purpose and potential unintended consequences.
- The severity of potential harms.
- The scale of use.
- The nature of harms or adverse impacts that have already taken place.
- The extent to which for practical or legal reasons it is not reasonably possible to opt-out from that system.
- Imbalances of economic or social circumstances, or age of impacted persons.
- The degree to which the risks are adequately regulated under another law.

ISED also notes that it is particularly interested in the following systems due to their potential impacts:

- **Screening systems impacting access to services or employment** : ISED notes that these systems are intended to make decisions, recommendations, or predictions for purposes relating to access to services, such as credit or employment. As such, they can potentially produce discriminatory outcomes and economic harm. Class actions have already been filed in the US against AI-

powered recruitment systems, and New York City has proposed [legislation to regulate automated employment decision tools](#).

- **Biometric systems used for identification and inference:** these are systems that use biometric data to make predictions about people, such as identifying a person remotely or making predictions about the characteristics, psychology, or behaviours of individuals. ISED notes that they have the potential of having significant impacts on mental health and autonomy and appears to be making [reference to organizations like Clearview AI Inc., which Canada’s four privacy commissioners investigated](#) in 2021.
- **Systems that can influence human behaviour at scale:** ISED notes that applications such as AI powered online content recommendation systems have been shown to have the ability to influence human behavior, expression and emotion on a large scale, and their potential impacts include harm to psychological and physical health. This category may be a reference to content moderation systems used by popular social media platforms.
- **Systems critical to health and safety :** this category refers to AI applications integrated in health and safety functions, such as automated driving systems and systems making triage decisions in the health sector, which could cause direct physical harm.

The factors that ISED lists as part of the assessment of AI systems show that it is balancing its regulatory obligations and the harms that could stem from large or risky AI systems with its desire to foster innovation and economic development. That said, the examples of systems that ISED is interested in regulating demonstrates that it considers certain systems to be inherently riskier than others, regardless of their scale. Further, it reflects concerns that have been addressed in other legislation recently enacted, like [Québec’s Bill 64, An Act to modernize legislative provisions as regards to the protection of personal information](#), which regulates certain [biometric systems](#) and automated decision-making tools outlined in the **Automated Decision Making** section below.

While ISED is broadly concerned about both individual and collective harms that could be caused by high-impact AI systems, it is particularly concerned about systems whose outputs are biased. Moreover, while it recognizes that certain adverse differentiations are unavoidable, it also recognizes that it must guard against systems that use proxies for prohibited grounds to differentiate people or amplify existing underlying correlations.

Regulated activities

The scope of application of the AIDA’s regulatory requirements is generally determined by the definition of the term “regulated activities”. Notably, this definition determines the scope of the anonymization requirement further discussed in the **Anonymization** section below. The definition is quite broad as it encompasses the entire lifecycle of an AI system, starting from the processing or making available for use of any data relating to human activities for the purpose of designing, developing, or using an AI system, to designing, developing, or making available for use an AI system or managing its operations (s. 5(1)).

The companion document provides examples for each category of regulated activity and recognizes that depending on the specific context, multiple businesses could be involved in carrying out the same and/or different regulated activities for a single AI system, ultimately demonstrating the need to dissect and differentiate the measures that

each of the involved businesses must take to comply with the AIDA's regulatory requirements. Consult the **Assessment, Mitigation, and Monitoring Obligations** section below for more information on how the obligations of each person involved across a high-impact system's lifecycle will be adjusted to the regulated activities they perform.

It's important to note that the activities listed in the AIDA are classified as "regulated activities" only when conducted within the context of international or interprovincial trade and commerce, meaning that all activities involving the development or use of an AI system performed outside of this context are excluded. This is particularly relevant for researchers who may not be subject to compliance obligations under the AIDA when conducting research or developing methodologies, as confirmed by ISED. However, since AI research and industry are closely linked, researchers should keep in mind that they could still be subject to the legislation when they release the outcomes of their work in a commercial context.

Open source software: the companion document provides important specifications for the open source publication of AI systems, a popular and valued practice amongst the AI community. The companion document notably highlights the fact that publishing an AI system open source (possible under various licenses) does not necessarily mean that the person behind the publication must automatically comply with the obligations that would apply to a person making the AI system available for use. In addressing this issue, ISED makes a distinction between persons who publish models or other tools, for example, as open source software, and those who publish a fully-functioning high-impact AI system. The former persons publish their models for others to use as a basis for developing their own AI systems based on their data and objectives, while the latter will have to comply with the obligations that apply to persons making the AI system available for use. With these comments, ISED seems to be focused on addressing the concerns of the AI research community.

However, this distinction can be indeterminate, and further clarification may be necessary to determine what constitutes a fully functioning high-impact AI system. For instance, a system like Stability's Stable Diffusion appears to be a "fully-functioning high-impact AI system" that can be used on its own, but because it is published under an [Open Rail license](#), it also allows users to [fine-tune the system](#) based on their own data and objectives. Hence, determining whether such a system amounts to making an AI system available for use might require a case-by-case evaluation of the system's downstream task.

Person responsible

The AIDA includes a broad definition of who may qualify as a person responsible for ensuring the compliance of an AI system by including in its definition persons involved in **the system's design, development, deployment for use, and management of its operations** (s. 5(2)). Note that this definition only applies to persons who carry out these roles in the course of international or interprovincial trade and commerce, which may again offer a caveat to those who work in research.

The companion document indirectly clarifies the application of the concept of "person responsible" by segmenting the responsibilities of such a person according to the regulated activity in which it is involved. See the **Assessment, mitigation, and monitoring obligations** section for more information on this approach.

Regulatory requirements

The AIDA includes a set of regulatory requirements that each differ in their application according to the role of the person concerned and which regulated activity they perform. This section of the article provides an overview of these requirements and ends with a table summarizing the clarifications provided in the companion document.

Anonymization

Under the AIDA, any person who conducts a regulated activity and processes or provides access to anonymized data during the course of that activity is required to establish measures in accordance with the regulations regarding the process of anonymizing data and the management and use of anonymized data and keep records of such measures (ss. 6 and 10). This obligation applies generally and is not limited to high-impact systems. Although the companion document does not explicitly address the anonymization requirement under the AIDA, it is evidently linked to the guiding principle of Accountability, which, as described by ISED, promotes the proactive document of policies, processes, and measures implemented.

Assessment, mitigation, and monitoring obligations

Assessment of any AI system: under the AIDA, a person who is responsible for an AI system must evaluate whether it is considered high-impact and keep records describing the reasons supporting their assessment (ss. 7 and 10). This means that everyone **within the definition of “person responsible” for an AI system is required to conduct an independent assessment of whether the AI system is a high-impact one.** In a scenario where multiple persons fall under this definition, this may include both the developer who made the AI system available for use as a service, as well as the person using the service and managing the operation of that same AI system. The specific timing of the assessment and the consequences of divergent conclusions remain unclear. However, specific guidelines for conducting this assessment are to be included in future regulations. In the interim, those responsible for an AI system may wish to conduct a preliminary assessment of whether it falls under one of the high-impact areas or involves any of the high-impact considerations identified by ISED, as discussed in the **High-Impact Systems** section.

Risk mitigation and monitoring of high-impact systems: the AIDA provides a set of requirements that aim to limit the risk that the use of a high-system result in harm or biased output (ss. 8, 9, and 10). The companion document suggests that to comply with these requirements, persons will be expected to institute appropriate accountability mechanisms, such as internal regulatory compliance governance processes and policies, appropriate to the risks associated with the regulated activities.

Moreover, the companion document provides examples of the mitigation measures that **should be implemented at each stage of an AI system’s lifecycle and concurrently illustrates how the AIDA’s risk mitigation requirement will apply to persons differently according to the regulated activity they are performing.** For example, ISED provides a set of mitigations measures for persons that develop a high-impact system that heavily reflect the companion document’s “Validity & Robustness” principle. Specifically,

persons developing high-impact systems will be expected to document the datasets used, build in mechanisms for human oversight and monitoring, document appropriate uses and limitations, and perform evaluations and validation. Consult the **table** at the end of this section for a summary of how the AIDA’s risk mitigation requirement will apply to different persons.

Transparency

The AIDA imposes transparency requirements regarding the intended use and actual use of high-impact systems. The companion document provides further guidance on what is expected in terms of transparency, namely providing the public with appropriate information about how high-impact AI systems are being used, in a way that allows the public to understand the capabilities, limitations and potential impacts of the systems. The “Transparency” principle ties in with the “Human Oversight” and “Monitoring” principle introduced in the companion document, which asserts that meaningful oversight necessitates interpretability that is appropriate to the given context.

Disclosure of intended use : persons who make available for use a high-impact system are required to comply with transparency obligations that center around the intended use of the system (s. 11(1)). This ensures that their transparency responsibilities are confined to the purposes for which the system was intended to be used and not, for example, for any unauthorized use that may occur online or otherwise.

Disclosure of actual use : a similar transparency obligation is applicable to persons responsible for operating a high-impact system. However, their requirement differs slightly, reflecting the distinct position of these persons across the lifecycle of an AI system (s. 11(2)). Accordingly, while most of the disclosure requirements might seem similar, persons involved in managing the operation of a high-impact system must provide their own description of the actual content generated or outcomes produced, just as they must establish their own set of mitigation measures to comply with the risk management obligation outlined in section 8 of the AIDA.

Automated decision making : it will be interesting to see if the government follows-up with regulation on how these transparency and accountability requirements will interact with related requirements in privacy laws. Indeed, privacy laws are beginning to address these concerns by requiring organizations to provide individuals with information about how automated decisions are made and what personal information was used in these decisions.

For example, starting on September 22, 2023, Québec's recently amended [Act respecting the protection of personal information in the private sector](#) (the Private Sector Act) will require organizations to inform individuals about the personal information used to render a decision based exclusively on automated processing of such information, as well as the reasons, principal factors, and parameters that led to the decision (s. 12.1 Private Sector Act). Similarly, the CPPA requires organizations to make available in their privacy policy a general account of their use of automated decision systems and to provide an explanation of the prediction, recommendation, or decision made by the system and that could have a significant impact on individuals, upon request by the affected individual, which is outlined in subsections 62(2)(c) and 63(3) of the CPPA. Hence, the disclosure requirements in the AIDA may complement privacy laws by providing individuals with a broader range of information about how the AI system used

to make an automated decision is intended to be - or is used - and what efforts are made to mitigate associated risks of biased output and harm.

Compliance and AMPs

Under the AIDA, the minister will have the power to designate an Artificial Intelligence and Data Commissioner to assist with overseeing compliance. The Commissioner will study the systemic effects of AI systems to inform administrative and policy decisions, as well as assist businesses with voluntary compliance.

Once the AIDA comes into force, the minister will also have the power to order the production of records to demonstrate compliance and to order independent audits, which ISED notes will be performed by qualified independent auditors. In situations where there is a risk of imminent harm, the minister may also order the cessation of use of a system, or publicly disclose information regarding contraventions of the AIDA or for the purpose of preventing harm. The minister may seek an order from the federal court to enforce its orders.

In addition to creating regulatory offenses and crimes relating to AI systems, AIDA also gives the minister the authority to issue administrative monetary penalties (AMPs) for any violations of the AIDA. Although the AMP scheme will be set out in regulations which have yet to be published, the companion document provides insight into how the scheme will operate. Notably, it notes that the AMP scheme will not likely come into force at the same time as the rest of the AIDA. Rather, ISED indicates that it will create **this AMP scheme once the “ecosystem and regulatory framework have sufficiently matured”**.

ISED notes that AMPs will be intended to serve as a flexible enforcement tool designed to encourage compliance. Beyond noting that the scheme will take into account the size of firms and whether other efforts to encourage compliance have failed, the companion document does not provide specific criteria that the minister will use to compute the quantum or appropriateness of an AMP. Finally, ISED appears to recognize that it will need to review the effectiveness and appropriateness of its enforcement actions. It plans to recruit external experts to analyze its administration and enforcement of the AIDA and will appoint a committee to advise the minister.

Takeaways and what to expect

The Canadian government expects that if Bill C-27 is adopted in the following months, the AIDA will not come into force before 2025. This two-year gap will allow for the development of regulations before AIDA comes into effect. While further changes should be expected, the AIDA companion document provides insight into the types of high-impact systems targeted by the current government.

The Canadian government is clearly prioritizing the development of agile legislation and utilizing regulation to achieve this objective to ensure that the legislation stays current and effective in a rapidly changing environment. Along with promoting flexible legislation through regulations, the government has assigned the minister of ISED, along with the AI and Data commissioner, to oversee the alignment of policy and enforcement as technology advances.

Following the adoption of Bill C-27, the Canadian government intends to initiate a consultation on various subjects, which are relatively extensive, to determine how to implement and prepare draft regulations. The consultation will cover the following topics:

1. Identifying the kinds of systems that should be classified as high impact.
2. Determining the types of standards and certifications required to ensure that AI systems meet the expectations of Canadians.
3. Establishing priorities for the development and enforcement of regulations, including the creation of an AMP framework.
4. Specifying the responsibilities of the AI and Data Commissioner.
5. Forming an advisory committee.

The regulatory landscape for AI is moving very quickly around the world and persons designing, developing or making AI systems available for use in Canada should expect to implement monitoring and mitigations practices, and should start documenting their risk evaluation process as part of their preparation to implement an AI risk management program. This will help ensure that they comply with any future regulations and are able to demonstrate accountability for their use of AI.

By

[Marc Vani](#), [George R. Wray](#)

Expertise

[Cybersecurity](#), [Privacy & Data Protection](#), [Artificial Intelligence \(AI\)](#)

BLG | Canada's Law Firm

As the largest, truly full-service Canadian law firm, Borden Ladner Gervais LLP (BLG) delivers practical legal advice for domestic and international clients across more practices and industries than any Canadian firm. With over 725 lawyers, intellectual property agents and other professionals, BLG serves the legal needs of businesses and institutions across Canada and beyond – from M&A and capital markets, to disputes, financing, and trademark & patent registration.

blg.com

BLG Offices

Calgary

Centennial Place, East Tower
520 3rd Avenue S.W.
Calgary, AB, Canada
T2P 0R3

T 403.232.9500
F 403.266.1395

Ottawa

World Exchange Plaza
100 Queen Street
Ottawa, ON, Canada
K1P 1J9

T 613.237.5160
F 613.230.8842

Vancouver

1200 Waterfront Centre
200 Burrard Street
Vancouver, BC, Canada
V7X 1T2

T 604.687.5744
F 604.687.1415

Montréal

1000 De La Gauchetière Street West
Suite 900
Montréal, QC, Canada
H3B 5H4

T 514.954.2555
F 514.879.9015

Toronto

Bay Adelaide Centre, East Tower
22 Adelaide Street West
Toronto, ON, Canada
M5H 4E3

T 416.367.6000
F 416.367.6749

The information contained herein is of a general nature and is not intended to constitute legal advice, a complete statement of the law, or an opinion on any subject. No one should act upon it or refrain from acting without a thorough examination of the law after the facts of a specific situation are considered. You are urged to consult your legal adviser in cases of specific questions or concerns. BLG does not warrant or guarantee the accuracy, currency or completeness of this publication. No part of this publication may be reproduced without prior written permission of Borden Ladner Gervais LLP. If this publication was sent to you by BLG and you do not wish to receive further publications from BLG, you may ask to remove your contact information from our mailing lists by emailing unsubscribe@blg.com or manage your subscription preferences at blg.com/MyPreferences. If you feel you have received this message in error please contact communications@blg.com. BLG's privacy policy for publications may be found at blg.com/en/privacy.

© 2024 Borden Ladner Gervais LLP. Borden Ladner Gervais LLP is an Ontario Limited Liability Partnership.