OSLER

Finding the balance: how Canada should approach AI regulation

MARCH 13, 2026 8 MIN READ

Related Expertise

- Artificial Intelligence
- Privacy and Data Management
- Technology

Authors: Michael Fekete, Simon Hodgett, Sam Ip, Jordan Geist

Key Takeaways

- AI is the fastest-adopted technology in human history, and Canada is at an inflection point on how to regulate it.
- Currently, Canada lacks a comprehensive AI regulatory framework, relying instead on sector-specific laws and general regulations.
- The upcoming national AI strategy presents an opportunity for Canada to establish a measured, harm-based regulatory framework that addresses emerging risks without deterring AI investment and innovation.

As artificial intelligence continues to transform industries and reshape daily life, Canada's approach to AI governance has reached an inflection point. The federal government has recently undertaken public consultations and a "30-Day AI Sprint" to inform the development of a renewed national AI strategy.

The public consultations, led by Innovation, Science and Economic Development Canada (ISED), sought input from industry, academia and civil society on the future of Canada's AI strategy and regulatory environment. These consultations are unfolding alongside broader global debates about AI safety and governance, including the release earlier this year of the *International AI Safety Report 2026*, led by Canadian AI pioneer Yoshua Bengio and amid heightened public attention to AI safety and digital platform responsibility in Canada.

Taken together, these developments have renewed debate about how Canada should regulate AI while maintaining its position as a global leader in AI research and innovation.

The current state of AI regulation in Canada

Canada does not currently have a comprehensive AI law. The *Artificial Intelligence and Data Act* was introduced as part of Bill C-27 in 2022, but it ultimately died when Parliament dissolved in 2025. As a result, Canada lacks the kind of comprehensive, economy-wide AI regulatory framework that some other jurisdictions have pursued.

In practice, however, AI technologies are already regulated in Canada through a combination of sector-specific legislation and laws of general application.

Quebec's Law 25, for example, imposes requirements on organizations where decisions are made based exclusively by automated processing of personal information, including obligations related to transparency and the right to have decisions reviewed by a human. Other sector specific frameworks govern particular AI-enabled technologies, the federal *Motor Vehicle Safety Act* and provincial highway traffic acts provide frameworks for regulating autonomous vehicles. Likewise, the *Ontario's Working for Workers Act*, 2023 mandates transparency when employers use AI to screen, assess, or select job applicants — provisions which came into force on January 1 of this year.

Beyond these targeted measures, Canada possesses a body of laws of general application that capture many potential AI-related harms. The *Personal Information Protection and Electronic Documents Act* governs the collection, use, and disclosure of personal information in the course of commercial activities, and its requirements apply regardless of whether data processing is performed by humans or AI systems. Federal and provincial human rights legislation prohibits discrimination on protected grounds, including when AI systems are involved in decision-making that affects individuals. Similarly, labour and employment laws continue to govern workplace relationships and employer obligations, even as AI transforms how work is performed and how employment decisions are made.

Taken together, this patchwork of legislation means that Canada already regulates many aspects of AI deployment, albeit indirectly.

Lessons from international approaches to AI regulation

Regulating artificial intelligence is exceptionally complex. Policymakers must balance competing considerations including harm prevention, broader social impacts, economic growth, and geopolitical positioning.

Three key factors contribute to this complexity. First, AI is a general-purpose technology, akin to electricity, the steam engine, and the Internet. Historically, societies have never attempted to regulate a general-purpose technology at the macro level in the manner now being contemplated for AI. Second, generative AI is the fastest-adopted technology in human history, with tools like ChatGPT achieving mainstream adoption in under three years. That pace poses creating significant challenges for regulatory frameworks that typically evolve over much longer time horizons. Third, there is a legitimate policy interest in encouraging AI development and adoption to capture the economic and social benefits the technology can deliver.

Against this backdrop, jurisdictions around the world have begun experimenting with various approaches to AI regulation and governance.

The European Union's *Artificial Intelligence Act* (EU AI Act) represents one of the most comprehensive approaches to AI regulation currently in force. The Act introduces a risk-based framework that calibrates regulatory requirements based on the potential for harm. In practice, however, the EU AI Act is both highly prescriptive and operationally onerous. Compliance costs are substantial, and the penalties for non-compliance are extraordinarily severe, creating significant financial exposure for businesses operating in the AI space.

The United States, by comparison, has pursued a more decentralized and agency-driven approach to AI governance, emphasizing AI leadership and innovation. Rather than introducing a comprehensive federal AI framework, U.S. policy relies on existing laws, federal executive actions, including Executive Order 14365, and sector-specific oversight by federal agencies. Individual states, meanwhile, have begun introducing targeted legislation addressing specific AI risks, including algorithmic discrimination, AI transparency and the

deployment of high-risk AI systems.

These differing approaches highlight the core policy tensions facing governments. Prescriptive regulatory frameworks provide clarity but risk imposing significant compliance burdens and slowing innovation, while decentralized approaches may support technological development but can create regulatory fragmentation and uncertainty.

A framework for balanced, harm-based regulation

As Canada charts its own approach to AI governance, policymakers must strike a careful balance between addressing emerging risks and keeping the country an attractive environment for AI research, investment and commercialization, a stated priority of the current government.

Feedback from the federal AI consultations highlights this tension and the need to balance innovation with responsible governance. Many participants emphasized the importance of transparent, risk-based regulatory frameworks that promote public trust in AI systems. Submissions also underscored the importance of regulatory clarity and predictable standards, particularly as Canadian organizations seek to scale AI deployment. At the same time, some stakeholders cautioned against regulatory approaches that could impose disproportionate burdens on AI developers and deployers, particularly given the global competition for AI investment and talent.

These considerations underscore the importance of calibrating Canada's regulatory approach carefully. An overly cumbersome framework risks inadvertently discouraging investment in and adoption of AI technologies, hampering needed productivity gains. This concern is amplified by the competitive landscape, particularly the comparatively hands-off approach to AI regulation currently prevailing in the United States. A framework that is overly burdensome could drive AI development and deployment to jurisdictions with lighter regulatory requirements, potentially undermining Canada's ability to compete in the AI global economy.

In light of these considerations, Canada should adopt a measured, harm-based regulatory approach built around four core principles.

First, the starting point is the enforcement of existing laws. Canada already possesses a substantial body of legislation covering consumer protection, privacy, human rights, product liability, and professional standards. Many potential harms arising from AI applications are already prohibited or regulated under these existing frameworks. Before creating new AI-specific rules, policymakers should ensure that existing laws are being effectively applied and enforced in the AI context.

Second, where genuine gaps exist in the current legal framework, targeted measures may be appropriate. However, policymakers must be careful not to conflate distinct issues or create overbroad responses to specific concerns. For example, online safety and AI regulation, while related, address different categories of risk and should not be treated as interchangeable. Recent discussions following the tragic incident in Tumbler Ridge, British Columbia illustrate this challenge. Public debate has included questions about whether AI developers or platforms should identify and report potentially harmful user behaviour. But as some commentators have noted, such proposals raise broader questions about intermediary responsibility, privacy and platform governance generally rather than AI regulation per se. Broad reporting obligations on AI providers could encourage expanded surveillance of user interactions and undermine privacy protections. These discussions illustrate the importance of ensuring regulatory responses remain tailored to the specific risks at issue and account for

unintended consequences.

Third, Canada should leverage its existing sectoral regulators where available rather than creating entirely new oversight mechanisms. Agencies such as Health Canada, the Office of the Superintendent of Financial Institutions, and Transport Canada possess deep domain expertise and established relationships with regulated entities. These regulators should have the latitude to apply generally applicable rules to AI applications within their sectors or to develop new, sector-specific requirements where warranted. This approach allows regulation to reflect the distinct risk profiles and operational realities of different industries.

Finally, Canada should consider enacting a harm-based backstop law that provides a flexible framework for responding to unforeseen harms or newly apparent risks. Rather than prescribing rules for every possible AI application, such legislation would establish principles and mechanisms for rapid regulatory response when significant harms emerge. The framework could also identify certain categories of AI use that are prohibited, while enabling additional restrictions, safeguards or compliance obligations to be introduced through regulation as new risks become apparent as the technology evolves. This approach would allow Canada to maintain a lighter regulatory touch during periods of stability while retaining the capacity to act decisively when circumstances require.

A look ahead

With the anticipated release of Canada's renewed national AI strategy expected in the near term, policymakers have an opportunity to clarify the country's approach to AI governance.

A balanced framework that prioritizes enforcement of existing laws, targeted interventions where genuine regulatory gaps exist, and sector-specific oversight may allow Canada to address emerging risks while continuing to foster a climate of innovation and investment in AI.

Striking this balance right will be critical for Canada to maintain its reputation as a global leader in artificial intelligence while ensuring that the deployment of these technologies remains aligned with public expectations around safety, fairness and accountability.