June 2019

Automation Not Domination: Legal and Regulatory Frameworks for AI

Adam Goldenberg and Michael Scherman, McCarthy Tétrault LLP

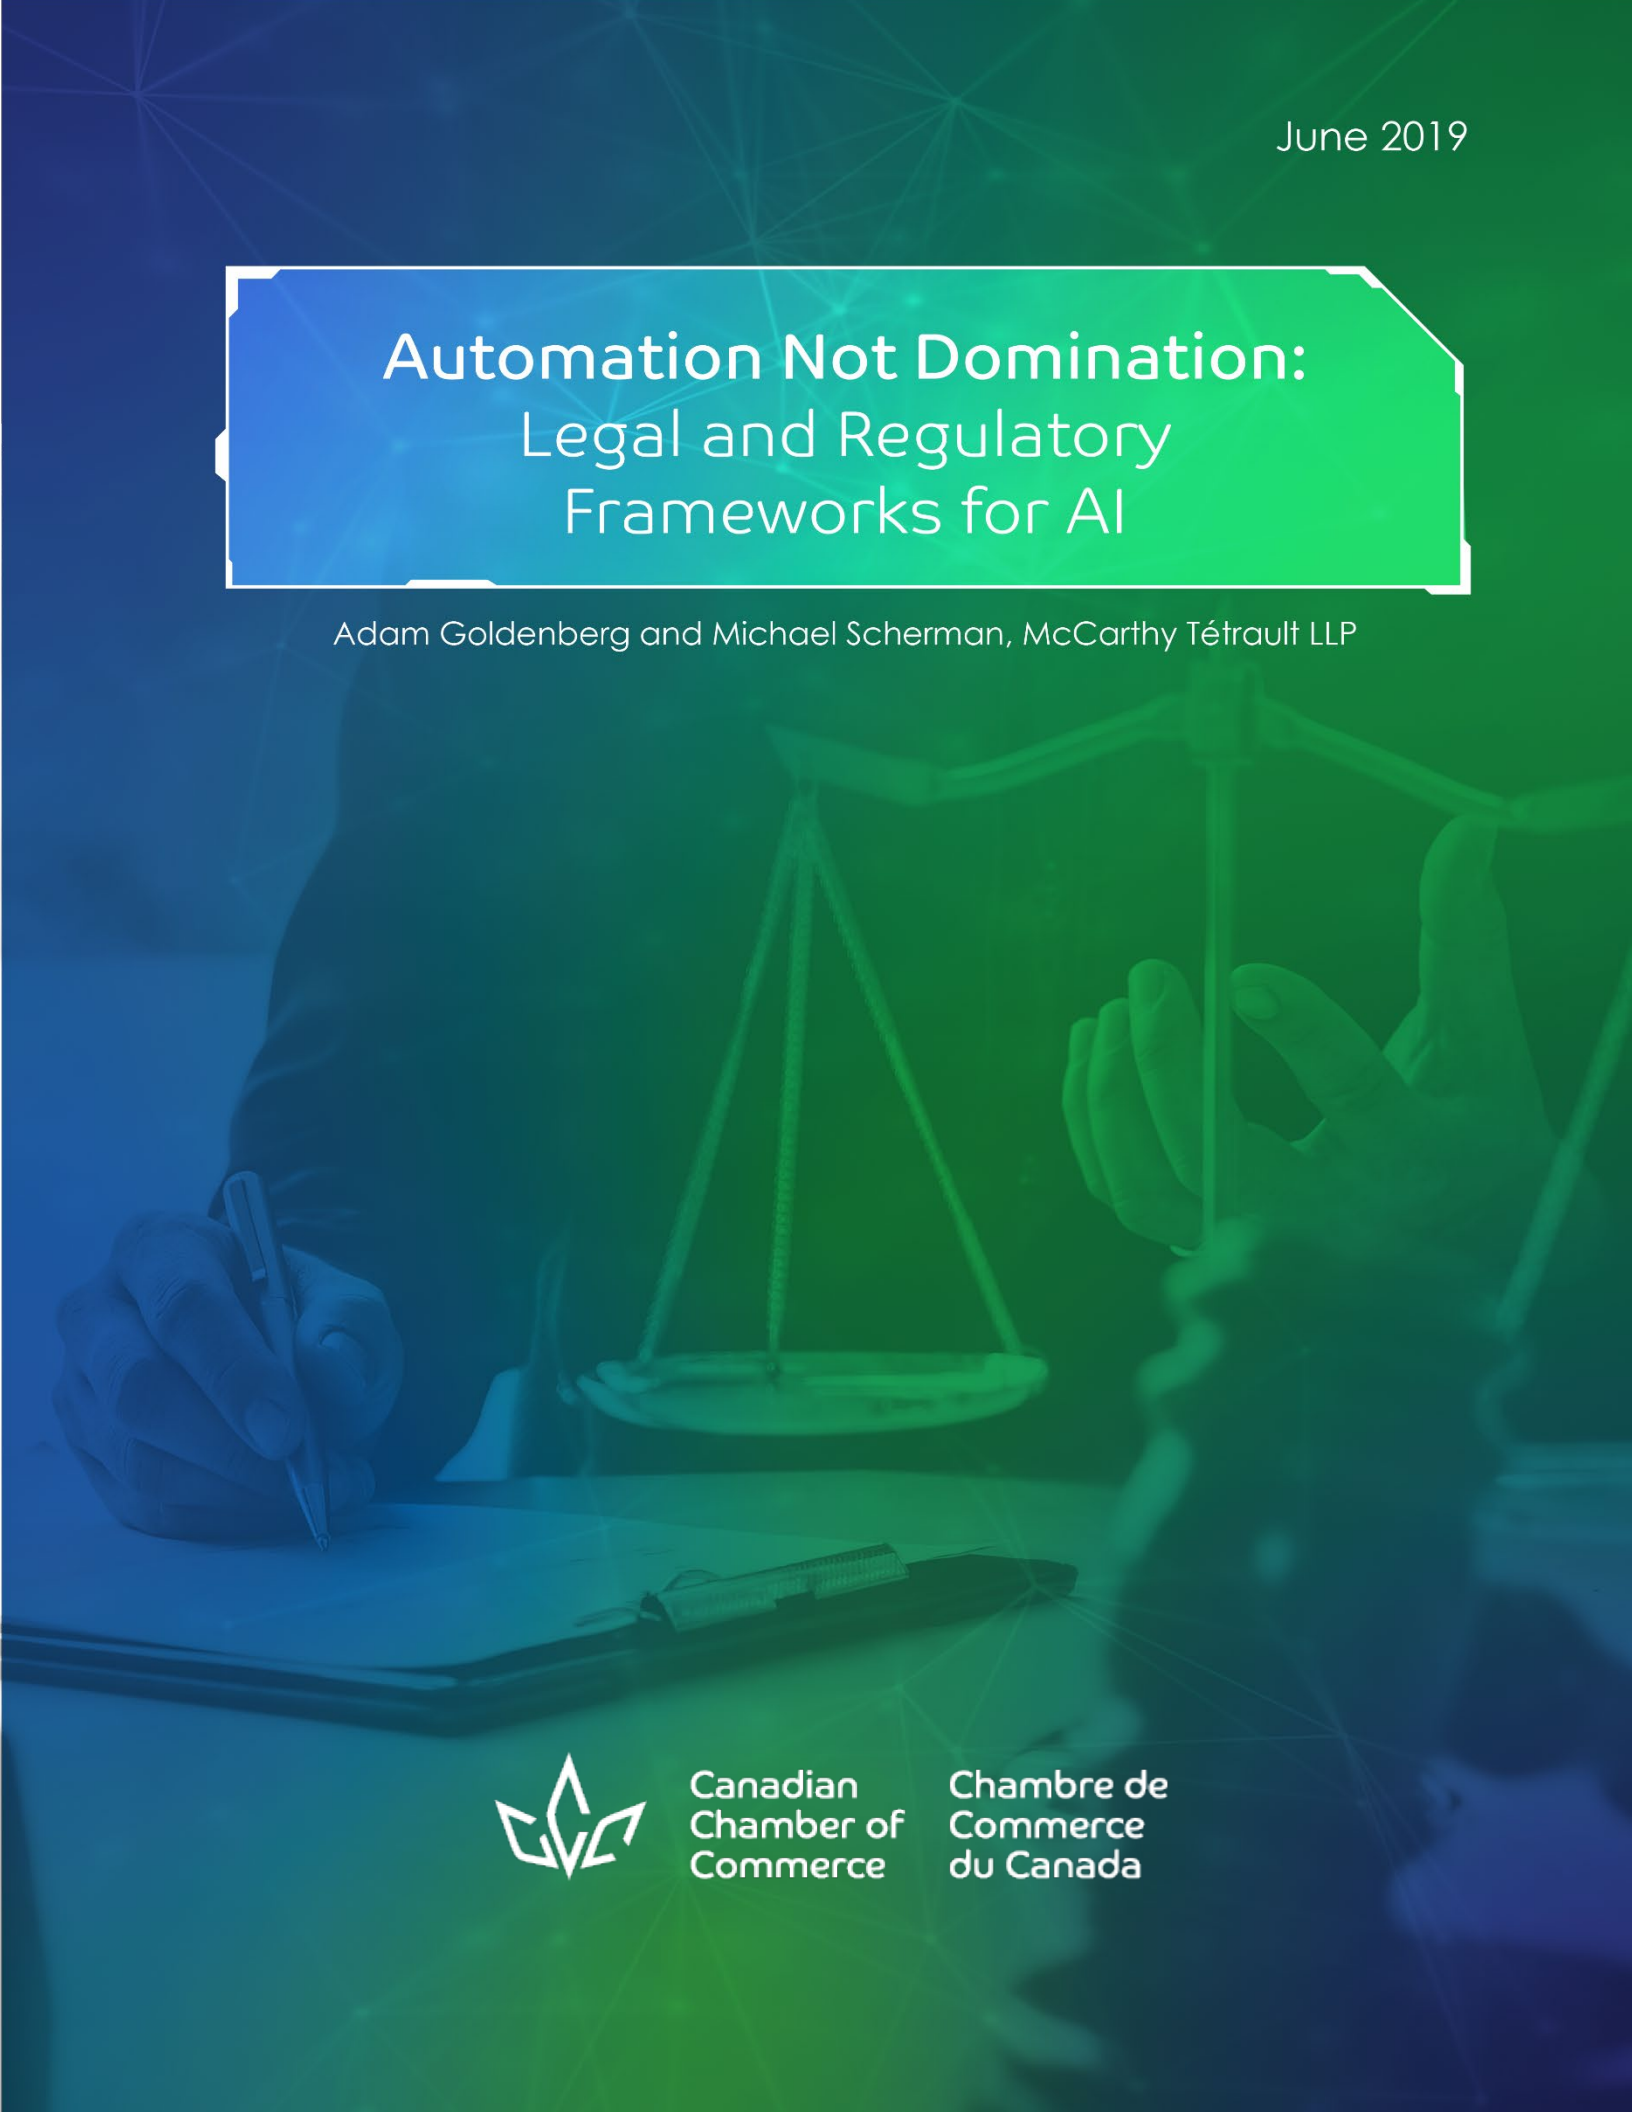

Canadian Chamber of Commerce
Chambre de Commerce du Canada

Background

AI-enabled systems are increasingly prevalent across all sectors. Their growing ubiquity implicates a range of legal and policy frameworks — including those that pertain to ethics and fairness, competition, privacy and intellectual property. These frameworks were not designed with AI in mind. They are, however, increasingly being adapted and applied to this new technology.

This process of adaptation and application makes AI a tantalizing target for policymakers. The state of public understanding (and misunderstanding) of AI's practical and social implications makes it tempting to think of its regulation as an unprecedented challenge, one that will require entirely novel approaches.

It is not and it will not. Rather, what lies ahead for Canada and other countries, as we harness AI's potential, is an ordinary process of considering these innovative technologies within established legal regimes and policy principles. There are, of course, significant differences between AI systems and other technologies. Some of these will require legislative and regulatory change. Yet, these changes can and should be evolutionary — even if the technologies to which they pertain are revolutionary. This report discusses how Canada's government can steward this evolution in a manner that protects Canadians' interests and ensures our shared prosperity in an AI-enabled world.

Ethical, explainable, accountable, safe and fair

Automated decision-making systems based on machine learning present a novel dilemma: the spectre of "algorithmic bias." Algorithmic bias ingrains human bias and discrimination in automated decision-making systems. It further develops as machine-learning models apply large data sets to generate probabilistic outcomes. Since these data sets are based on historical data and past results, a machine-learning system may perpetuate discrimination that previous outcomes reflect. Issues of accountability, transparency and non-discrimination quickly emerge.

The discussion around how to address these issues is well underway, both in Canada and abroad. In April 2019, the Government of Canada issued its *Directive on Automated Decision-Making* (the "**Directive**").[1] The Directive applies to all automated decision-making systems used by the federal government to provide services to external clients or to make administrative decisions and recommendations about external clients. It incorporates procedural fairness, due process and transparency into government AI-enabled decision making by way of impact assessments, public reporting and verifiability and auditing requirements, among other tools. As the use of AI proliferates across government, the impact of this Directive will increase. Notably, however, the Directive does not apply to automated decision-making systems in the private sector.

Other governments are also considering the issue, in relation to both the public and private sectors. The *Algorithmic Accountability Act* was recently introduced in both the U.S. Senate and the U.S. House of Representatives, and the European Commission's High-Level Expert Group on Artificial Intelligence has just published its *Ethics Guidelines for Trustworthy AI*.[2]

[1] Government of Canada, "Directive on Automated Decision-Making" (5 February 2019).
[2] European Commission, "Ethics guidelines for trustworthy AI" (8 April 2019).

AI systems can cause physical injuries, too. Autonomous vehicles can crash and harm passengers or pedestrians. AI-based tools can misdiagnose medical conditions and, thus, misinform treatment.[3] It is, of course, critical that organizations that use AI put in place strong risk management programs so that any such risks are well understood — and limited or eliminated to the extent possible.

Still, human activities are seldom risk-free. Individuals, communities and governments constantly assess and accept risks they consider reasonable in order to enjoy the benefits of technology. For example, highway speed limits are set at levels that reduce, but do not eliminate, the risk of physical injury. As Jonathan Sumption, formerly of the U.K. Supreme Court, recently put it:

Think about road accidents. They are, by far, the largest source of accidental, physical injury in this country. We could almost completely eliminate them by reviving the Locomotive Act of 1865, which limited the speed of motorised vehicles to four miles an hour in the country and two in towns. Today, we allow faster speeds than that, although we know for certain that it will mean many more people being killed or injured, and we do this because total safety would be too inconvenient. Difficult as it is to say so, hundreds of deaths on the roads and thousands of crippling injuries are thought to be a price worth paying for the ability to get around quicker and more comfortably. So, eliminating risk is not an absolute value, it's a question of degree.[4]

The potential safety risks of AI must be evaluated in the same light. The question should not be, "does this technology pose a risk?" but rather, "how much risk are we willing to accept in order to garner the advantages to individuals and our society that this technology offers?" The legislative and regulatory frameworks for AI must reflect the latter inquiry, not the former. We must decide, as a country, not *whether* we will accept the costs of AI but rather *which* costs are reasonable in view of the benefits we stand to gain, individually and collectively. The perfect should not become the enemy of the good.

Governments are rightly interested in ensuring ethical, explainable, accountable, safe and fair AI. Legislation and regulations must, nonetheless, be carefully considered so that they do not unnecessarily limit AI's potential to make our society safer, fairer, more liveable and more prosperous overall.

Competition

International competition/antitrust regimes aim to prevent businesses from distorting competition in the free market.[5] Competition laws generally focus on delivering outcomes that are beneficial to consumers and are sector neutral. There is, however, ongoing debate in the competition law community about whether this approach needs to be re-tooled to adapt to the new digital economy. This debate reflects broad recognition of the growing importance of the digital economy to consumers.

AI systems depend on massive amounts of data. The acute need for data means that proprietary datasets are an increasing source of value for participants in a wide range of industries. This demand is likely to increase as AI system quality is more closely tied to quantitative training.

[3] McKinsey Quarterly, "Confronting the risks of artificial intelligence" (April 2019).
[4] J. Sumption, "The Reith Lectures 2019: Law and the Decline of Politics — Lecture 1: Law's Expanding Empire" (May 21, 2019), *BBC Radio 4*.
[5] Competition Bureau Canada, "Intellectual Property Enforcement Guidelines" (13 March 2019).

The importance of data is only increasing as AI technologies move towards deep learning and neural networks — both of which require huge volumes of inputs.

A key question in competition law is the extent to which concentrating these datasets in the hands of a few companies could lessen, prevent or distort competition in a manner that would harm consumers. In particular, the potential harm associated with big data is that it becomes a pre-requisite to competing in a given market, preventing new entrants from gaining a foothold. In the context of AI, this would allow a limited number of companies to control AI development by virtue of collecting and monetizing essential raw data.

The Competition Bureau (the "**Bureau**") is looking at big data and associated issues closely. It released a consultation paper entitled *Big Data and Innovation: Implications for Competition Policy in Canada* in September 2017.[6] The report canvassed competition issues raised by the acquisition and use of big data. In February 2018, the Bureau released the results of its review based on the discussion paper *Big Data and Innovation: key themes for competition policy in Canada*.[7] In it, the Bureau takes the position that:

The emergence of firms that control and exploit data can raise new challenges for competition law enforcement but does not, in and of itself, necessitate an immediate cause for concern. There is little evidence that a new approach to competition policy is needed although big data may require the use of tools and methods that are somewhat specialized and, thus, may be less familiar to competition law enforcement. The fundamental aspects of the analytical framework (e.g., market definition, market power, competitive effects) should continue to guide enforcement.[8]

The Bureau's *2018-19 Annual Plan* identifies the digital economy as an enforcement priority. The Bureau has appointed a chief digital enforcement officer, whose role is to support the Bureau's enforcement activities in this area, including with respect to big data.[9]

More recently, in May 2019, the Minister of Innovation, Science and Economic Development ("**ISED**") released an open letter to the Commissioner of Competition that laid out various challenges and considerations related to competition law and policy in the data-driven economy.[10] In particular, the Minister asked the Bureau to consider critical competition law issues presented by the digital and data-driven economy, including "the impact of digital transformation on competition; the emerging issues for competition in data accumulation, transparency, and control; the effectiveness of current competition policy tools and marketplace frameworks; and the effectiveness of current investigative and judicial processes."[11]

[6] Competition Bureau Canada, "Big data and Innovation: Implications for Competition Policy in Canada" (September 18, 2017).
[7] Competition Bureau Canada, "Big data and innovation: key themes for competition policy in Canada" (February 19, 2018).
[8] *Ibid* at page 4.
[9] Competition Bureau Canada, "Building Trust to Advance Competition in the Marketplace" (30 May 2018).
[10] Minister of Innovation, Science and Economic Development Canada, "Letter from Minister of Innovation, Science and Economic Development to the Commissioner of Competition" (May 2019).
[11] *Ibid* at para. 11.

Privacy

Many AI systems rely on large datasets. Where this is so, system development will often involve the use of personal information, or de-identified data, that was derived from personal information. Access to large, useful datasets is crucial to the continued development of AI systems. Yet, such access will come with significant implications for privacy.

In May 2019, the Government of Canada announced a new *Digital Charter*, along with a number of proposals from ISED to modernize the *Personal Information Protection and Electronic Documents Act* ("***PIPEDA***").[12] Among other proposals, the *Digital Charter* contemplates alternatives to consent for the use of personal information in "standard business activities." It also proposes a risk-based approach to regulating the de-identification of personal information.

It is often challenging to obtain the necessary consents and to de-identify personal information in a way that moves the data beyond the scope of Canada's current privacy laws. These challenges can cripple business innovation and the development and deployment of AI. Well-structured alternatives to consent, together with a balanced and clear regime for the de-identification of personal information, would help organizations to acquire and use the datasets that are required to develop and operate AI-based technologies, while simultaneously protecting the privacy interests of individuals. The *Digital Charter* and the accompanying ISED proposals appear to recognize the need to strike this important balance.

Privacy considerations also arise when AI is used to process personal information and to make important decisions about individuals. In February 2018, the Standing Committee on Access to Information, Privacy and Ethics of the House of Commons released a report on its review of *PIPEDA*. The report included a recommendation that "the Government of Canada consider implementing measures to improve algorithmic transparency."[13] This recommendation was made on the basis that, where personal information is processed using complex algorithms — such as in an AI system — informed consent from an individual requires the individual to know how the algorithm will use the data and to what end. In its response to the Committee's report, the government agreed with the call for algorithmic transparency as a means of protecting privacy and enhancing innovation. It committed to holding consultations on the potential benefits and impacts of incorporating algorithmic transparency into Canada's privacy framework.[14]

ISED reiterated these transparency concerns in its May 2019 proposals to modernize *PIPEDA*. As ISED put it, "a lack of transparency around automated decision-making processes and the resulting decisions increases individuals' concerns related to bias and potential discrimination."[15] In its current form, however, *PIPEDA* (and privacy laws in Canada generally) apply to big data analytics and AI mainly through the lens of individual consent and data security requirements.

---

[12] Innovation, Science and Economic Development Canada, "Strengthening Privacy for the Digital Age: Proposals to modernize the Personal Information Protection and Electronic Documents Act" (May 2019).
[13] Standing Committee on Access to Information, Privacy and Ethics, "Towards Privacy By Design: Review Of The Personal Information Protection And Electronic Documents Act" (February 2018) at p. 25.
[14] Government of Canada, "Government Response To The Twelfth Report Of The Standing Committee On Access To Information, Privacy And Ethics" at p. 7.
[15] *Supra* note 12.

Intellectual Property

Intellectual property ("**IP**") law strives to protect innovators without stifling future innovation. Unique challenges arise in dealing with AI systems in the IP context. These include the issues of how to allocate rights in works created by AI systems, and what level of protection to afford to the datasets used by AI systems.

As AI proliferates, it is becoming increasingly common for AI-based technologies to themselves generate IP. This raises a number of novel issues that existing IP regimes do not expressly address. For example, when AI technologies are involved, identifying the inventor of a patentable invention or the author of a work under copyright law, in not always straightforward. These questions are even more complicated where there are numerous parties involved; *e.g.*, the developer of the AI algorithm, the company that made the input data available for training the AI system and the user of the AI system at the time the IP is developed. These remain open questions, both in Canada and elsewhere.

The House of Commons Standing Committee on Industry, Science and Technology has recently recommended "[t]hat the Government of Canada consider amending the *Copyright Act* or introducing other legislation to provide clarity around the ownership of a computer-generated work."[16] Governments, here and elsewhere, must carefully consider any such changes to existing regimes; they will have material implications on the development and use of AI.

The amount of protection that IP law affords to datasets will have a similar effect on AI development and use. Large datasets are frequently necessary, and companies often make substantial investments to assemble valuable datasets. Some jurisdictions, such as the EU, have regimes that expressly protect collections of information. Canada has not enacted such a regime, and this type of protection does not currently exist here.

To benefit from some level of protection under the current *Copyright Act*, datasets may be analogized to "compilations" if they meet an originality threshold. Cases have been brought before the courts to define whether a given collection of data meets the requirements to benefit from protections against infringement of copyright. In considering changes to the level of protection afforded to datasets, governments must bear in mind that any protection afforded to datasets cuts both ways; AI developers may seek to safeguard their own datasets, but their (or their AI system's) use of data may also infringe on others'.

[16] House of Commons, *Statutory Review of the Copyright Act: Report of the Standing Committee on Industry, Science and Technology* (June 2019), at p. 51.

What We Heard…

On March 19, 2019, at a roundtable discussion hosted by the University of Waterloo, the Canadian Chamber of Commerce and McCarthy Tétrault LLP brought together experts and business leaders from across multiple industries and academia to discuss public policy questions in connection with AI from a legal and regulatory perspective. This roundtable was part of a series to gather insight from experts and business leaders across industries and academia on the future of AI in Canada. The discussion spanned a broad range of challenges that will bear on Canada's ability to realize AI's full potential across our entire economy and geography.

The following is an overview of the discussion as we heard it. These findings were subsequently discussed and further developed at a roundtable discussion hosted by McCarthy Tétrault LLP in Vancouver on May 22, 2019.

I. Embrace some regulatory "lag"

We Asked: **Is regulatory "lag" good for AI innovation and uptake?**
We Heard: **Yes, but not at the expense of regulatory certainty.**

Experience shows that legislation introduced with the right intentions can have unintended consequences for technology uptake and, thus, for economic growth. In British Columbia, for example, the data residency provisions of provincial privacy legislation[17] act as a major barrier to the public sector's adopting cloud computing services. If the federal government moves too quickly to regulate too much, it risks undermining Canada's burgeoning AI ecosystem.

Society has seen economic revolutions before. AI could spur such a revolution. However, it will not do so all at once. Any attempts at "forward thinking" or "real time" regulation could quickly be overtaken by advances in AI technology. Canada's priority, at this juncture, should be to encourage those advances not seek to limit them prospectively through regulation.

Canada should, instead, tell the world that its jurisdiction will be a friendly one for AI innovation. We should rely on existing legislation and regulation, create an accessible inventory of those frameworks, provide guidance on how they apply to AI technology — perhaps through a "single window" approach — and seek to limit redundancies and to close gaps. But we should give businesses based in Canada a wide berth to develop AI technologies and create the jobs and foster the economic growth that will come with it.

There are circumstances in which regulatory "lag" is undesirable, however. It could, for example, discourage AI companies from operating or making investments here if our current regulations (or lack thereof) introduce uncertainty or do not allow for the development or use of a technology. An example would be if existing motor vehicle laws did not allow for the AI-enabled automation of certain driving decisions. In such a case, assuming the technology's introduction would benefit society overall, the government would be well advised to appropriately amend legislation as early as possible.

[17] *Freedom of Information and Protection of Privacy Act*, R.S.B.C. 1996, c. 165, s. 30.1.

Whether to regulate, and to what degree, should turn on the objective of promoting innovation and encouraging investment. Governments that attempt to anticipate the trajectory of AI technologies and that base regulation on anticipation risk undermining the public interest in the name of protecting it. Regulatory certainty is necessary, but rushed regulation could be fatal. Canada's governments must tread carefully; it is no overstatement to say that our economic future depends on it.

A number of specific potential regulatory approaches were discussed during our roundtable at the University of Waterloo:

Roundtable participants were generally not supportive of a "sandbox" model. Their concern was that the bureaucratic processes associated with operating such a "safe" environment could be sufficiently cumbersome to push AI innovation to other jurisdictions.

It was suggested that algorithms could be vetted for "safety" by a regulator, subject to confidentiality protections and similarly to pharmaceuticals. Some roundtable participants saw this approach as bound to be limited by the government's technical capacity and also as an unnecessary step that could serve as a brake on economic growth.

It was also suggested that the government support a standards-based approach, including binding professional standards for technologists similar to those imposed on members of other regulated professions. This proposal was generally well received, though implementing it (in the absence of effective self-regulation) would likely require provincial, rather than federal, action.

II. Respect consumer attitudes

We Asked: **What do Canadians want?**
We Heard: **The public expects principled guidance to prevent harm to individuals and an ambitious agenda for global leadership.**

"AI" is not one thing. Government has a role in educating Canadians about what different AI technologies are, what they can do and how they will affect our society. A workable definition of AI is a necessary starting point for official guidance on how different AI systems may be developed and deployed.

Canadians do not expect a complex regulatory framework. We do, however, wish to advance our collective desire to encourage the responsible and ethical development of AI in Canada. The federal government can do so by developing clear guidance to businesses on how their existing legal obligations apply to AI innovation.

Guidance should be oriented around the principle of avoiding harm. Like "AI," the term "harm" will require precise definition. That definition may largely be derived from existing legislative and regulatory frameworks, such as those that guard against discrimination. In this way, the guidance need not be legally binding in itself in order to build Canadians' trust that AI innovation will not harm them. Avoiding harm does not, however, mean regulating AI to a standard of perfection. Software will crash. Systems will fail. Embracing AI's potential requires accepting reasonable missteps.

The purpose of guidance must be to ensure AI developers and users understand their legal obligations and to orient their efforts towards the objective of avoiding legally cognizable harms. Eliminating such harms, however, will not be possible. We should not pretend otherwise. The price of preventing harm completely is not one that Canada (or any country) can afford to pay.

We want Canada to be a global leader in AI. We want companies based in Canada to be at the forefront of innovation in this space. But we also want the application of AI in business processes not to undermine our legal rights and interests. Effective, principled guidance can provide a bridge between these two desires, to ensure they do not operate in conflict with one another.

III. Encourage competitiveness and growth

We Asked: **What do Canadian businesses need?**
We Heard: **Canadian businesses need to be able to compete globally.**

One roundtable participant highlighted how China is using its citizens' heath information to create a massive data set to train algorithms. This illustrates the global competitiveness of AI innovation. Businesses based in Canada have to compete in this environment. Governments must bear this in mind as they approach AI regulation.

Canada needs the ability safely to generate and maintain data sets that will allow AI algorithms to learn. Our legal and regulatory frameworks must allow these data sets to be generated and accessed. Existing and proposed limitations on AI development and deployment should be assessed against this policy objective.

We also need governments to work in partnership with the private sector and with one another. Private sector organizations should be allowed space to develop self-regulatory approaches within the ambit of principled guidance from government. Those self-regulatory approaches should inform future government action. And, because AI regulation crosses jurisdictional boundaries between orders of government, the federal government's collaboration must extend to the provinces.
Working together, Canada's governments and businesses are well positioned to decide, in partnership, how existing legal and regulatory frameworks apply to AI innovation. The goal should be to provide effective guidance that bridges the divide between AI development and AI deployment — and to create an environment that encourages economic growth.

Recommendations

The Government of Canada should:

- **Implement public education initiatives.** The public needs to understand what AI is, and what it can (and cannot) do. Informed citizens can advocate for themselves and their communities. The circle of stakeholders involved in shaping the legislative and regulatory frameworks that govern AI must be as broad as possible to ensure sustainable "buy in." The government has a crucial role in developing, funding and promoting the public education initiatives that will make broad, democratic engagement possible.

- **Offer clear guidance for businesses.** Developers and users require clear guidance about their legal obligations related to AI. As legal and regulatory systems adapt, the government should ensure businesses understand their evolving obligations. The government can do so by, among other things, establishing a "single window" approach for businesses that seek to develop and deploy AI in Canada. Crucially, the government can encourage businesses to use the single window by proposing legislation that would shield business's interactions with officials from future disclosure or use in enforcement activities.

- **Monitor algorithmic "safety."** Accountability and transparency will be key to avoiding unreasonable harm in AI development and deployment. Principled guidance from government should inform thoughtful self-regulation by industry — and both should strive to achieve accountability and transparency. This will ensure Canada remains a global leader in AI, and that Canadians benefit from the advancements and improvements that AI will offer, all at a cost that the public understands and accepts.

- **Promote professionalism.** The government should promote the adoption of industry or professional standards aimed to ensure professionalism among the designers and implementers of AI systems.

- **Consider responsive regulation.** The government should focus on addressing existing problems and creating regulatory certainty, not on prospective regulation. In doing so, the government should avoid fencing off unexpected areas of innovation. It should also avoid creating new regulations, which quickly become outdated as technology evolves. The government should, however, act quickly where existing regulation is preventing development or use of a new technology, or where a regulatory gap is causing uncertainty that discourages innovation or investment in Canada.

- **Tailor sector-specific solutions and provide guidance on compliance, in competition law.** The government should preserve the Competition Bureau's principles-based approach to competition law and continue to take a case-by-case approach when dealing with business sectors. The government should also provide guidance on competition law compliance to companies implementing AI systems. As the use of AI within each industry grows, it will benefit from systems that were designed with competition law in mind.

- **Enable innovation and growth while protecting privacy**. The government should revise privacy laws in a way that enables innovation and the use of datasets while protecting privacy rights. A risk-based and balanced regime for the de-identification of personal information, and subsequent use thereof by technology companies, is one example.

- **Support Canadian innovators.** The government should support Canadian AI developers in order to foster innovation and Canada's profile on the world stage. It should streamline and clarify the IP regimes that relate to AI in a way that balances the need for businesses to protect their investment in AI technologies with the need to foster future innovation.

- **Consider alternative models.** The government should facilitate industry consultations and working groups that can consider alternative data access models. Such models may include data trusts and data-sharing agreements. The government should recognize, however, that such models could have adverse consequences, including by undermining the willingness of large technology companies to invest in Canada and by imposing a disproportionate burden on SMEs. Government-supported, industry-led consultations should consider both the benefits and the potential adverse consequences of alternative approaches to governance.

Thank You to our Partner

mccarthy tetrault

Premier Sponsors

aws

Google

Grant Thornton
An instinct for growth™

mccarthy tetrault

Microsoft

Supporting Sponsors

Desjardins

IBM

UNIVERSITY OF WATERLOO