

Regulating AI: Issues and Choices

Nye Thomas
Executive Director
Law Commission of Ontario

March 26, 2024

AI systems offer significant potential benefits to governments, the private sector, and the public. Many believe that these tools can “crack the code of mass adjudication”, improve the accuracy and consistency of decision-making, improve public and private services, and reduce backlogs. At the same time, public and private sector use of AI is controversial. There are many examples of AI systems that have proven to be biased, illegal, secretive, or ineffective.

Many governments – including the Governments of Canada and Ontario – have adopted or are developing “Trustworthy AI” frameworks to assure the public that AI development and use will be transparent, legal, and beneficial.

Achieving “Trustworthy AI” depends on a complex series of policy, legal and operational questions that go far beyond public statements of principle. Addressed thoughtfully, the answers to these questions will help governments, the private sector, and other institutions agencies maximize AI’s benefits and minimize its harms.

AI regulation in Canada and elsewhere is advancing quickly, but inconsistently. Every day, it seems like someone has introduced a new statute, policy, standard, or rule of practice to govern a widening array of AI systems.

Not surprisingly, AI regulation raises many of the same questions and issues that arise in many other areas of regulation. As a result, it can be helpful to analyze AI governance debates in terms and language that are familiar to legal policymakers, judges, and lawyers.

In that spirit, I offer a summary of five important issues that Law Commission of Ontario (LCO) believes will define the shape and effectiveness of AI regulation in Canada in the coming months and years.

“Ethical AI” vs. “Hard” Law

As a generalization, many private sector proponents and AI developers believe ethical AI/best practice guidelines are needed to ensure rigid government regulations do not stifle innovation

or economic development. Ethical AI guidelines are also often seen as better suited to rapidly changing technology.

Ethical AI approaches are frequently criticized. Critics state that ethical guidelines are insufficient to mitigate the harms caused by AI due to their lack of specificity and reliance on voluntary compliance. More pointedly, many critics believe that ethical AI guidelines – without more – are a form of “ethics washing.”

Notwithstanding these criticisms, ethical AI policies and guidelines are ubiquitous. Examples include Ontario’s Beta Principles on Ethical AI, the Biden administration’s *AI Bill of Rights*, the U.S. National Institute of Standards “Risk Management Framework” and the Organization for Economic Cooperation Development’s (OECD) “Principles on Artificial Intelligence.” Many corporations have also developed (or are developing) their own “in-house” ethical AI frameworks.

National, Local, or Targeted Legislation

In many jurisdictions, the AI governance debate has evolved from *whether* to regulate AI to more specific conversations about *how* to regulate AI.

Examples of national legislation include the European Union’s recently passed *Artificial Intelligence Act*, Canada’s proposed federal *Artificial Intelligence and Data Act (AIDA)*, and the draft *Algorithmic Accountability Act of 2022* in the United States.

The EU’s *Artificial Intelligence Act*, like the General Data Protection Regulation before it, is setting the global standard for national AI regulation. The Act is a comprehensive 100+ page statute that applies to anyone who develops or deploys any type of AI, including public and private sector AI systems and both large and small enterprises. In contrast, *AIDA* is both less comprehensive and narrower in scope. The roughly ten-page Bill is “shell” legislation that leaves most details to be determined through regulations. The UK’s national AI framework is different again. There is no national AI legislation in the UK. Rather, the UK government’s approach set out in a recent white paper, *AI regulation: a pro-innovation approach*, lists five principles to guide existing regulators to govern the responsible use and development of AI within their own mandate.

In addition to national legislative frameworks, many jurisdictions have enacted targeted legislation or policies to govern AI in specific locations or contexts. For example, the Government of Canada’s *Directive on Automated Decision Making* is a government directive that applies to federal AI systems to ensure procedural fairness in federal algorithmic government decision-making. Yet another approach is emerging in the United States, where there are many examples of state or local legislation targeted to specific AI applications, such as New York City’s legislation governing employment AI systems and the more than 20 U.S. jurisdictions that have banned or restricted the use of police facial recognition systems.

The choice of national, local, or targeted AI legislation obviously has important consequences for AI governance. On the one hand, national or “framework” legislation is potentially comprehensive and inclusive of all AI systems in a jurisdiction. The challenge, however, is that

generic legislation does not respond to specific AI issues or risks that arise in specific contexts, such as policing or government decision-making.

These questions are particularly important in Canada. *AIDA* is an important advance, but it is not comprehensive federal legislation. Most notably, *AIDA* does not apply to most public sector AI systems. More importantly, but less obviously, *AIDA* must be supplemented by dedicated provincial legislation. Absent dedicated provincial, there is a possibility that some of the most consequential potential uses of AI systems used by provinces, municipalities, police services, child welfare agencies and other important public institutions will be under- or unregulated.

Regulating by Risk

Risk-based AI governance models assume that different AI systems will have different level of risk and that regulatory obligations should be tailored proportionately. The risks addressed in these models typically include privacy risk, data security risk, bias risk, and, increasingly, other forms of human rights risk.

Risk-based AI regulation is explicit in the EU *AI Act*, *AIDA*, Canada's Federal Directive, the U.S. *AI Bill of Rights*, and most other AI governance regimes.

Not all risk-based models are the same: *AIDA* includes two levels of risk: high and low. By comparison, the EU has three levels of risk: unacceptable risk (i.e. prohibited systems), high risk and low risk.

Risk-based models embed important assumptions and choices that may ultimately determine the success or failure of AI regulation, including the criteria used to distinguish risk levels, who assesses risk, and how is risk mitigated.

Risk-based models also often include outright or limited prohibitions on some forms of AI, including facial recognition and biometric identification systems, predictive policing systems, and social network behavioural analysis systems.

Many AI governance regimes also explicitly exempt whole areas of AI use, such as AI systems used in national security applications, from governance requirements. The breadth or scope of regulatory exceptions obviously has important AI governance implications.

Disclosure and Transparency

It is widely acknowledged that some form of disclosure and transparency should be a feature of AI governance models. That said, jurisdictions and models have considerable differences between them regarding the level of disclosure required, when disclosure is required, and who disclosure is made to.

Important Questions About AI Governance

Purpose, Definition, and Scope

- Will AI regulation promote innovation, rights protection. or both?
- How should AI be defined?
- What institutions or activities will be regulated?
- Exemptions?

Legislation and Ethical AI Risk-Based Regulation

- What form(s) will regulation take?
- Is there a statutory framework?
- Commitment to comprehensive regulation?
- Are there regulatory “gaps”?
- Need for dedicated rules, regulations, or practices to govern specific contexts or applications?

Risk-Based Regulation

- Will regulation be risk-based?
- What are the risk categories and who created them?
- Who assess risk of an AI system?
- Who reviews the risk assessment?
- Are AI impact assessments mandatory?
- How is risk mitigated?
- Ongoing risk assessment over lifespan of AI system?
- Are independent audits and evaluations required?

Disclosure and Transparency

- Is there a commitment to comprehensive accountability and transparency?
- Mandatory AI registers?
- What will be disclosed, to who, and when?

Bias, Privacy, and Fairness

- How do regulations address bias and discrimination?
- How do regulations ensure procedural fairness and privacy?
- How is compliance with human rights code, privacy legislation, administrative law and the *Charter* ensured?

Oversight and Remedies

- Are there dedicated remedies for rights violations?
- Are there dedicated procedures to ensure access to justice?
- Is there independent oversight of an AI system?

Enforcement and Remedies

Similarly, there is a wide divergence in whether, or how, enforcement and remedies are addressed in AI governance frameworks. For example, *AIDA* include provisions creating an AI and Data Commissioner who is largely responsible for enforcing the Act. In contrast to some American AI statutes, *AIDA* does not create a dedicated right of action or procedures for individuals who may have been harmed by an *AIDA*-regulated AI system.

Conclusion

As noted above, AI regulation in Canada and elsewhere is advancing quickly but inconsistently. Despite many shared assumptions about “Trustworthy AI” principles and practices, there is a need for more legislative coordination between initiatives at the federal, provincial, and local levels. There are also many unanswered questions that will ultimately determine the success or failure of AI governance.

This inconsistency is understandable given that we are still in the early days of AI regulation. Nevertheless, the LCO and many other organizations have identified the key elements of thoughtful, comprehensive AI regulation, including:

- Baseline requirements for all government and private sector systems, irrespective of risk.
- Risk-based requirements that could include:
 - Strong protections for AI transparency, including disclosure of the existence of an AI system and a broad range of data, tools, and processes used by the system.
 - Auditing and evaluation requirements.
 - AI impact assessments
 - Explainability requirements.
- Mandatory “AI Registers.
- Explicit requirement that AI systems will comply with human rights legislation, privacy legislation, and the *Charter* (for government systems).
- Data standards.
- Access to meaningful procedural protections and remedies.
- Creation of an independent AI oversight office.

The LCO believes the best approach to AI regulation is to adopt a mixture of “hard” and “soft” law instruments, tailoring each to their appropriate purpose and context. This view is sometimes called a “smart mix” or “mixed model” of AI regulation.

In the LCO’s view, federal and provincial framework legislation is clearly needed to provide the foundational governance framework for AI systems in Canada. A legislative framework would provide consistent direction and accountability requirements to the actors, governments, agencies, and private enterprises within its scope. Finally, legislation would establish a level of public and legal accountability commensurate with the issues and rights at stake.

That said, the LCO does not discount the use or importance of ethical guidelines, directives, “playbooks” or best practices. Indeed, the LCO believes these instruments have significant potential to supplement or expand upon mandatory legal obligations and requirements.