

Legal Issues and Government AI Development

Workshop Report



LAW COMMISSION OF ONTARIO
COMMISSION DU DROIT DE L'ONTARIO

About the Law Commission of Ontario

The Law Commission of Ontario (LCO) is Ontario's leading law reform agency.

The LCO provides independent, balanced, and authoritative advice on complex and important legal policy issues. Through this work, the LCO promotes access to justice, evidence-based law reform and public debate.

The LCO evaluates laws impartially, transparently and broadly. The LCO's analysis is informed by legal analysis; multi-disciplinary research; contemporary social, demographic and economic conditions; and the impact of technology.

The LCO is located at Osgoode Hall Law School, York University, Toronto.

More information about the LCO is available at www.lco-cdo.org.

Authors

Susie Lindsay, Counsel, Law Commission of Ontario

Jesse Beatson, Consultant, Law Commission of Ontario

Nye Thomas, Executive Director, Law Commission of Ontario

Disclaimer

The opinions or points of view expressed in the LCO's research, findings and recommendations do not necessarily represent the views of the Ontario Digital Service, the Government of Ontario, funders (Law Foundation of Ontario, Osgoode Hall Law School, Law Society of Ontario) or supporters (Law Deans of Ontario, York University).

Citation

Law Commission of Ontario, *Legal Issues and Government AI Development: Workshop Report*, (Toronto: March 2021)

Contact

Law Commission of Ontario
2032 Ignat Kaneff Building
Osgoode Hall Law School, York University
4700 Keele Street
Toronto, Ontario, Canada
M3J 1P3

Email: LawCommission@lco-cdo.org
Web: www.lco-cdo.org
Twitter: @LCO_CDO
Tel: (416) 650-8406
Toll-free: 1 (866) 950-8406

Funders

The Law Commission of Ontario is funded by the **Law Foundation of Ontario**, the **Law Society of Ontario**, **Osgoode Hall Law School**, and **York University**. The LCO also receives significant in-kind support from Osgoode Hall Law School and York University.



Legal Issues and Government AI Development: Workshop Report



INTRODUCTION

Government interest in artificial intelligence (AI) and automated decision-making (ADM) systems is growing rapidly. This is because AI and ADM have tremendous potential to transform government decision-making and public services. According to a recent study (Engstrom et al, Appendix A) of American administrative agencies,

Rapid developments in AI have the potential to reduce the cost of core governance functions, improve the quality of decisions, and unleash the power of administrative data, thereby making government performance more efficient and effective.

This same study notes that in the US, AI and ADM tools are currently being used to assist government operations “across the full range of governance tasks”, including:

- Enforcing regulatory mandates centred on market efficiency, workplace safety, health care, and environmental protection;
- Adjudicating government benefits, from disability benefits to intellectual property rights;
- Monitoring and analyzing risks to public health and safety;
- Extracting useable information from the government’s massive data streams, from consumer complaints to weather patterns; and
- Communicating with the public about its rights and obligations as welfare beneficiaries, taxpayers, asylum seekers, and business owners.

However, along with great potential is great concern. Many government AI and ADM systems are criticized due to serious concerns about racial bias, lack of transparency and legal accountability, data issues and inadequate public engagement.

In many instances, AI and/or ADM systems are being used to either make or assist government decisions that impact individual rights, interests and obligations. For example, AI and ADM systems are currently being used by many governments to determine government benefits, immigration eligibility, assess the risk of child abuse, identify individuals likely to be criminals and prioritize access to housing, to name a few examples. As a result, justice system stakeholders in Canada and elsewhere are analyzing the legal principles, statutory rules and constitutional requirements that arise when government decision-making is assisted by machines. An important part of this analysis is to consider how “traditional” legal rules and statutes apply to these systems, and whether to create new laws or regulations.

The need to address these issues prompted the Law Commission of Ontario (LCO) and the Ontario Digital Service (ODS) to bring together Government of Ontario lawyers, policymakers, operational managers and technology experts with Law Commission counsel and advisors for an informal and collaborative discussion about artificial intelligence, automated decision-making and the law.

THE EVENT ORGANIZERS

The Law Commission of Ontario (LCO)

The Law Commission of Ontario (LCO) is Ontario's leading law reform agency. The LCO provides independent, balanced, and authoritative advice on complex and important legal policy issues. Through this work, the LCO promotes access to justice, evidence-based legislation and legal policies, and public debate on important and topical law reform issues. The LCO evaluates laws impartially, transparently and broadly. The LCO's analysis is informed by legal analysis; multi-disciplinary research; contemporary social, demographic and economic conditions; and the impact of rapidly changing technology.

LCO reports are an authoritative, long-term resource for policymakers, stakeholders, academics, and the general public. LCO reports have led to legislative amendments, policy changes and are frequently cited in judicial decisions, academic articles, government reports and media stories. The LCO's most recent report, *The Rise and Fall of Algorithms in American Criminal Justice: Lessons for Canada*, was released in October 2020. More information about the LCO is available [here](#).

The Ontario Digital Service (ODS)

The Ontario Digital Service (ODS) is a transformational organization within Ontario's provincial government. The ODS works to enable the adoption of digital channels and platforms for government. Values of human-centred design and open data are key pieces to how the ODS helps public servants across the province embrace new technologies and ways of working.

Within the ODS, the Data Stewardship Team focuses on responsible data sharing and use through our world class Open Data program which includes Ontario's Data Catalogue. This work is evolving to encompass the government's use of algorithms and other data-driven technology like artificial intelligence (AI). Surfacing early AI activities across the government is helping to set coordinated expectations around 'good' and 'responsible' use of AI by government.

The Data Stewardship Team was excited to co-lead these exploratory workshops with the LCO, looking into the use of AI by government to create advice and guidance collaboratively that is practical as well as aspirational. More information about the ODS' Data Stewardship work is available [here](#).

HOW DID WE ORGANIZE THE WORKSHOPS?

The LCO and ODS organized four workshops over several weeks in late 2020. The workshops were designed to be participatory. Each attendee came prepared to think and share ideas and opinions.

The workshops were grounded in a hypothetical case study about the use of AI and automated decision-making to administer social benefits and detect fraud. Each week we discussed a new stage of development of the hypothetical system and the legal and policy issues that could arise. We began each session with a brief presentation and then explored the various issues and worked through them as a group. Materials specific to each session were distributed a week prior.

The workshop considered four stages of development of an AI system: development/design, implementation/operation, oversight/evaluation and legal challenges. Each session considered four key legal/policy issues: disclosure, bias, due process and public participation.

The workshops included the following presentations:

DEVELOPMENT OF AI OR ADM SYSTEM

- **Matt Ross**, Data Scientist, Manager AI, City of London, Ontario
- **Benoit Deshaies**, Director of Data and Artificial Intelligence, Treasury Board of Canada
- **Amy Bihari**, Lead Open Data and Emerging Technologies, Treasury Board Secretariat, Ontario Digital Services

BIAS AND EXPLAINABILITY

- **Insiya Essajee**, Counsel, Ontario Human Rights Commission
- **Rim Khazall**, AI Analyst, Treasury Board Secretariat

DUE PROCESS AND PUBLIC PARTICIPATION

- **Teresa Scassa**, Professor of Law, University of Ottawa
- **Ryan Fritsch**, Counsel, Law Commission of Ontario

LITIGATION AND LEGAL RISKS

- **Susie Lindsay**, Counsel, Law Commission of Canada
- **Michelle Mann**, General Counsel, Justice Canada

The forum was conducted under the Chatham House Rule.

THEMES AND INSIGHTS


The workshops generated many ideas, discussions and questions. For clarity, throughout this report “AI system” is used as a catch-all term to include data, models, processes and the actual IT system.

Notwithstanding their different experiences and skills, it quickly became apparent that every participant shared two fundamental assumptions:

- While AI systems have great potential to improve public services, they also raise new, complex and multi-disciplinary issues.
- In order to be trustworthy, government AI systems must be legal and effective, minimize bias and maximize public participation.

With these assumptions in mind, the workshops generated eight major themes and insights:

1. Clear Legal Rules	Everyone involved in the design, development, deployment and oversight of AI systems will benefit from clearer legal rules and guidance.
2. Context-Sensitive Rules, Guidelines or Regulations	Some legal obligations apply consistently to all AI systems, but others must be sensitive to context.
3. Public Participation	Public participation and consultation is crucial to AI design and deployment.
4. Broad Expertise Throughout AI Life Cycle	A broad range of expertise needs to be involved at every stage of AI design, development, deployment and oversight.
5. Bias	There must be a commitment to evaluating for bias and disparate impact on vulnerable communities.
6. Culture Change	Development of AI systems requires a more iterative approach than traditional IT tools.
7. Disclosure, Transparency and Accountability	AI disclosure, transparency and accountability are important public values and legal obligations.
8. Risk Management	Steps can be taken to reduce risks, including litigation risk.



1. Clear Legal Rules

Everyone involved in the design, development, deployment and oversight of AI systems will benefit from clearer legal rules and guidance.

AI and ADM systems are currently being used or contemplated by many public institutions, but the legal rules governing these systems are unclear.

Workshop discussions revealed a great need and interest in establishing clearer legal rules and expectations for everyone involved in AI and ADM design, development, deployment and oversight.

Every participant wanted to ensure government AI systems complied with the *Charter of Human Rights and Freedoms*, the *Ontario Human Rights Code*, privacy legislation, procurement rules and any other applicable statutory, administrative law or internal government requirements. Not surprisingly, in many cases developers, policy makers, operational managers, government lawyers and frontline staff find it difficult to know when and how these rules apply to AI systems.

A lack of clear guidance about AI legal requirements can have serious consequences. Programmers and developers, for example, need clear guidance to help them design AI systems. Similarly, a lack of clear guidance may mean operational managers or policy staff are unable to clearly confirm and document that AI systems comply with applicable legal rules.

Notably, along with clear rules, there needs to be capacity to implement the rules. Hence, capacity-building is an important corollary to rule-setting.

2. Context-Sensitive Rules, Guidelines or Regulations

Some legal obligations apply consistently to all AI systems, but others must be sensitive to context.

Building on the first theme, participants discussed what rules, guidelines or regulations addressing provincial government use of AI systems might look like. Some of the issues discussed included:

- How to strike a balance between encouraging innovation and protecting rights;
- How to understand and address bias and discrimination issues in government AI systems;

- Government procurement; and,
- How to ensure compliance with procedural fairness and due process obligations.

Participants discussed the Government of Canada's Directive on Automated Decision-Making ("the federal Directive"). The federal Directive is one of the first examples of a dedicated regulation or directive specifically addressing government use of automated decision-making systems.

The federal Directive is modelled on administrative law principles and, subject to some exceptions, requires all federal automated decision-making systems to undergo a mandatory Algorithmic Impact Assessment. Significantly, the Directive establishes baseline requirements that apply to all ADM systems, regardless of their impact level. The Directive then establishes four levels of impact, judged by the impact on individuals or communities. The four levels range from Level I (an automated decision that "will likely have little to no impact" and that is "reversible and brief") to Level IV (an automated decision that "will likely have very high impacts" and is "irreversible and perpetual"). The Directive then establishes requirements for each impact level, including greater or lesser levels of notice, peer review, training and human intervention.

The workshop discussed the variable impacts of AI systems at length. The workshop learned that administrative law generally requires more extensive transparency and disclosure requirements when decisions have a strong impact on individuals. By the same token, legal requirements can be reduced where decisions do not impact such rights or are more administrative in nature. The goal is to tailor legal requirements to the potential risks or impacts on an individual. This is best achieved with a sliding scale of regulation, rules or guidelines relative to the potential impacts and harms of a given AI system. In light of these obligations, it was agreed that it is important to proactively design, implement and maintain AI systems that facilitate and promote procedural rights. It is also necessary to address administrative law and procedural fairness requirements for each application, as the legal requirements will vary depending on context.

The federal Directive is a useful starting point. However, participants agreed there is much work to be done to design systems of regulations, rules or directives that can further policy objectives and bring clarity and organization to the application of legal principles. For example, principles of procedural fairness need to go beyond the decision-making context. There are many ways in which AI systems can nudge people, influence behaviour, but fall short of "decisions". Further, other legal and policy obligations beyond procedural fairness need to be addressed.

3. Public Engagement

Public engagement is crucial to AI design and deployment.

Participants learned that a key to successful design and deployment of government AI systems is public input and consultation. This is especially true of systems that can impact rights and entitlements. Consultations should start early in the development phase.

Public engagement is crucial to building and maintaining effective, trustworthy and rights-respecting public services. Public engagement fosters "fairness by design", helps to ensure AI systems achieve their policy objectives, and helps avoid unforeseen obstacles and pitfalls. Public engagement is also an important part of the development of rules and regulations that govern AI systems.

The nature and degree of public engagement can be a sliding scale: systems with higher impact and risks should have more public engagement. Questions to be asked include: Who are the stakeholders? Who is potentially affected by the system? What are the benefits of the proposed system? What are the risks?

For public consultation to be meaningful, participants must be supported, given the resources and time to provide meaningful feedback, and their opinions must be given due consideration. Successful public consultation also requires a commitment to public education.

4. Broad Expertise Throughout AI Life Cycle

A broad range of expertise needs to be involved at every stage of AI design, development, deployment and oversight.

Most AI systems are designed and built by various technology and data practitioners. Technologists, however, should be teamed with policy experts, lawyers, operational managers and frontline staff – the people who will oversee and implement the system – through the entire life cycle of an AI project.

Early advice from lawyers and policy staff can help address challenging compliance issues such as due process, privacy and human rights. Lawyers can also help prepare impact assessments, explainability requirements, and inform recourse options. Policy staff, legal staff, operational staff and community representatives can also help to assess the validity, quality, completeness and accuracy of the data used to train and support an AI system.

Building an AI system to inform government decisions affecting rights or entitlements is more complicated than many technology projects. The breadth and depth of issues to be considered can be overwhelming. Creating a team with broad expertise will improve decision-making, be more efficient, and improve public services.

Broad expertise should remain part of the AI system through the entire life cycle, including development and implementation. For example, AI systems need to be constantly evaluated for effectiveness and bias. Systems can also “drift” over time. To address these issues, the same team approach and expertise is necessary post-implementation to consider relevant technological, legal and policy issues.

5. Culture Change

Development of AI systems requires a more iterative approach than traditional IT tools.

Using AI to successfully improve government policy and services is often more than simply adopting a new technological tool. AI is a combination of technology and policy. AI can introduce many new opportunities, issues and risks. Additionally, existing government development processes are outdated, rigid, and follow the waterfall method of completing each stage before advancing to the next. As a result, there is a need for a cultural shift in IT design.

Government policy or program development timelines can be rushed for many reasons. Rapid government AI development can be risky, however, especially for high-impact, complex systems. In

many jurisdictions, rapid government deployment of AI systems has led to significant public controversies, legal challenges, poor public services and in some cases direct harm to individuals. As a result, workshop participants discussed several “cultural” issues regarding AI development and implementation. Some of the recurring themes from these discussions include:

- AI design and implementation should be iterative. Accordingly, it is important to build feedback loops and continual consultation and monitoring into the AI system’s life cycle. In some cases, it may be challenging to ensure these loops, but this is a necessary part of quality assurance and ongoing effectiveness.
- It is important to establish reasonable expectations about what AI technology can achieve and consider whether AI is the right solution. There is a risk that AI or algorithmic systems are seen as a panacea for complex government issues. As a result, project scope and definition are keys to successful AI outcomes and effective public services. Finally, we need a governance framework that ensures user focus and user involvement at all stages.
- Data and data literacy are key to successful AI development and deployment. Data quality, reliability and validity are central concerns for any AI system. It is also important to educate all stakeholders on the benefits, limits and qualifications of data-driven predictive technologies, such as AI.
- Development of AI systems must be through a comprehensive, not purely technological, lens. All voices need to be included and awareness and literacy must be expanded for both technological and policy issues.

6. Bias

There must be a commitment to evaluating for bias and disparate impact on vulnerable communities

Many AI systems are criticized for having biases, hidden within their code or data sets, that can lead to discriminatory results. Most commonly, these systems are criticized as being racist or sexist, though a system can potentially be biased in any number of different ways. The issue of bias is particularly complex for individuals who are members of multiple, overlapping categories. Experience in other jurisdictions demonstrates that systemic biases can be unintentionally embedded in a system and have a disproportionate impact on marginalized and vulnerable communities.

Workshop discussions revealed that technologists and policy developers can have different perspectives on bias. However, all participants acknowledged the critical importance of taking steps to provide bias-free public services – not least of which is to uphold the *Charter of Rights and Freedoms* and the *Ontario Human Rights Code* – both of which require government activity to be free of discrimination.

In evaluating a system for bias, participants discussed the need for a multi-disciplinary and intersectional lens. Participants also discussed how evaluating and correcting systemic bias can be very challenging, particularly in the absence of clear legal rules and expectations. Developers noted that bias often has a much different meaning in IT and systems development than in human rights law. They also noted there may not be reliable or uniform metrics or methods for assessing bias within a system.

Some of the questions discussed include:

- How is bias measured?
- What is the standard for evaluating bias in an AI system?
- Is an AI system successful if it is *less* biased than human decision-making?
- Isn't there also potential for AI to reduce or expose bias in existing services, programs and/or datasets?

Notwithstanding these questions, all participants agreed that there is a need to evaluate bias in government AI systems. It was acknowledged, however, that not all systems will raise the same issues or have the same impact on Ontarians. As a result, risk mitigation strategies are likely to depend to a significant degree on context.

7. Disclosure, Transparency and Accountability

AI system disclosure, transparency and accountability are important public values and legal obligations.

Participants agreed that the principles of transparency and accountability need to be embedded in government AI systems from the outset. Transparency in government AI systems supports AI innovation, fosters “trustworthy AI”, improves government services and decision-making, is consistent with administrative law, and reduces the risk of discrimination and/or legal challenges. As a result, public institutions need to be wary of using AI systems that are an impenetrable “black box”.

Participants learned that, in most cases, the subject of an administrative, regulatory or legal decision is entitled to know how the decision-maker reached their conclusion. If the decision-maker is a machine, as opposed to a public servant, the system needs to be constructed so that it can answer for its determinations. The reasoning process must be “explainable”. This is not always an easy task. Standards need to be developed to provide clarity to developers, frontline staff and operational managers.

For transparency to be meaningful, a government AI system must be understandable to multiple audiences. This includes frontline staff, clients, policy and operational managers, politicians, lawyers, data scientists, and the general public. Robust internal transparency leads to good public-facing transparency and vice versa.

Participants were also mindful of transparency issues that arise when AI is procured from an external vendor. In these circumstances, there may be fewer opportunities for government officials to influence the design of the system, understand how it works, or ensure appropriate public safeguards. A further issue is that a private company may assert proprietary claims and resist the disclosure that would be expected of “public” systems.

In these circumstances, it may be difficult to test a system adequately or ensure meaningful disclosure, unless appropriate protections are secured at the contracting stage. It was noted that external procurement is unlikely to absolve the government of legal responsibilities or accountability.

Participants agreed that transparency is important in principle, but that the timing and extent of disclosure is a complicated issue. For example, what are the rules for protecting personal privacy and

complying with privacy legislation? Some participants also asked about the necessity or propriety of disclosing source code, due to concerns about security and “gaming the system.”

Participants agreed that these are valid concerns. It was noted, however, that many of these concerns are not new and that existing government privacy or procurement rules may provide guidance. Participants also agreed that the federal Directive is a useful starting point for ideas on how government departments can provide transparency in AI systems. Finally, there was a consensus that disclosure, transparency and accountability are important public values and legal requirements. As a result, participants agreed it is important that government officials consider disclosure and accountability issues at every stage of an AI life cycle.

8. Risk Management

Steps can be taken to reduce litigation risk.


Participants learned that some AI applications will be scrutinized heavily. Litigation is a real possibility, especially in high profile and high impact applications. This has occurred repeatedly in other jurisdictions.

Participants heard that thoughtful AI design will go a long way to reducing risk of costly, complex, high-profile and lengthy litigation.

If litigation is initiated, defendants (i.e., the government) need to be prepared to produce relevant information related to the AI system. As a result, it is a best practice to build systems in such a way that this information is being recorded, retained and is readily accessible. The question of what counts as relevant and proportionate disclosure is a question that courts will determine.

For an individual claimant, litigating a decision rendered by an AI system will likely be extremely costly and complex. AI systems also bring risks of test case litigation. What starts out as an individual claim could easily become a systemic challenge (a “test case”). As is often the case in “test case” litigation, other stakeholders may get involved as intervenors.

Workshop participants generally agreed that “regulation by litigation” is not satisfactory for individual plaintiffs or governments. This conclusion reinforces the desirability of clear legal guidance, interdisciplinary development and implementation teams, and transparency.



NEXT STEPS, MORE INFORMATION AND COMMENTS

The LCO encourages public institutions to engage in similar exercises when considering or developing AI systems. The LCO can provide supporting resources and advice to any organization considering such an exercise.

Appendix A includes a list of reading materials and resources on government use of AI and automated decision-making systems and general background information.

The LCO believes that successful law reform depends on broad and accessible consultations with individuals, communities and organizations across Ontario. For more information or questions about the LCO, this project or related issues, go to the LCO's webpage at www.lco-cdo.org or contact the LCO at lawcommission@lco-cdo.org.

Law Commission of Ontario
Osgoode Hall Law School, York University
2032 Ignat Kaneff Building
4700 Keele Street
Toronto, Ontario, Canada
M3J 1P3

Email: LawCommission@lco-cdo.org
Web: www.lco-cdo.org
Twitter: [@LCO_CDO](https://twitter.com/LCO_CDO)
Tel: (416) 650-8406
Toll-free: 1 (866) 950-8406

SELECT ADDITIONAL MATERIALS AND RECOMMENDED READINGS

LCO Reports

- *The Rise and Fall of Algorithms in American Criminal Justice: Lessons for Canada*, (October 2020).
- *Regulating AI: Critical Issues and Choices*, (Forthcoming, Spring 2021).
- *AI, ADM and Government Decision-Making*, (Forthcoming, Spring 2021).
- *Probabilistic Genotyping DNA Tools in Canadian Criminal Courts*, (Forthcoming, Spring 2021).
- *AI, Automated Decision-Making: Impact on Access to Justice and Legal Aid*, (2019).
- *AI for Lawyers: A Primer on Artificial Intelligence in Ontario's Justice System* with Element AI and Osgoode Hall Law School, (2019).
- *Roundtable on Digital Rights and Digital Society* with the Mozilla Foundation, (2018).

Government of Canada Directive on Automated Decision-Making

- Government of Canada, *Directive on Automated Decision-Making*, (2019).
- Government of Canada, *Algorithmic Impact Assessment*, (2019).

Additional Reports and Papers

- Teresa Scassa, *Administrative Law and the Governance of Automated Decision-Making: A Critical Look at Canada's Directive on Automated Decision-Making*, (2021).
- David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey and Mariano-Florentino Cuéllar, *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies*, (2020).
- Australia Human Rights Commission, *Human Rights and Technology Discussion Paper*, (2019).
- New Zealand Law Foundation, *Government Use of Artificial Intelligence in New Zealand*, (2019).
- New York City Automated Decisions Systems Task Force, *Automated Decision Systems Task Force Report*, (2019).
- Rashida Richardson, ed., *Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force*, (2019).
- AI Now Institute, *Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability*, (2018).
- Ryan Calo and Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, (2020).
- The Royal Society, *Explainable AI: The Basics*, (2019).
- Ellen P. Goodman, *The Challenge of Equitable Algorithmic Change*, (2019).