Human Rights AI Impact Assessment
Backgrounder

March 2025

LCO CDO
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 legislation and policies, and public engagement on important law reform issues. The LCO is independent of stakeholder interests and is committed to a public interest perspective for every project.

The LCO is located at Osgoode Hall Law School in Toronto.

More information about the LCO and this project is available at: https://www.lco-cdo.org .

LCO CDO
LAW COMMISSION OF ONTARIO
COMMISSION DU DROIT DE L'ONTARIO

Authors

Susie Lindsay, Counsel, LCO
Nye Thomas, Executive Director, LCO

Citation

Law Commission of Ontario, *Human Rights AI Impact Assessment Backgrounder*, (Toronto: March 2025).

Disclaimer

This Backgrounder accompanies the AI Human Rights Impact Assessment prepared by the LCO and Ontario Human Rights Commission (OHRC) with contributions from the Canadian Human Rights Commission (CHRC). The Backgrounder was written by the LCO and does not necessarily represent the views of the OHRC, CHRC, or the LCO's funders or supporters.

Contact

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

Tel: (416) 650-8406
E-mail: LawCommission@lco-cdo.org
Web: https://www.lco-cdo.org
LinkedIn: https://linkedin.com/company/lco-cdo
Bluesky: @lco-cdo.bsky.social
X: @LCO_CDO
YouTube: @lawcommissionofontario8724

Funders

Financial support for the Law Commission of Ontario is provided by the Law Foundation of Ontario, the Law Society of Ontario, and Osgoode Hall Law School.

1. Introduction

This Background Paper accompanies the recent release of the Law Commission of Ontario(LCO)/Ontario Human Rights Commission (OHRC) AI human rights impact assessment (HRIA).[1] The purpose of the HRIA is to help Canadian public and private organizations assess and mitigate the human rights impact of AI systems in a broad range of applications.

The context for the HRIA is the extraordinary growth in the use of AI by governments, public agencies, and the private sector across the world. The breadth and pace of AI systems reflects the perceived potential of AI to improve the accuracy, speed, and consistency of public and private decision-making. Notwithstanding this potential, use of AI is often problematic. There are many examples of private and public AI systems that have proven to be biased or discriminatory. The focus of this paper is on the LCO's work and role in the development of the HRIA.

Globally, many governments and organizations are developing or promoting "Trustworthy AI" frameworks to assure the public that AI systems are beneficial, lawful, and accountable. AI impact assessments have emerged as a leading strategy to promote "Trustworthy AI".

The LCO's 2022 *Accountable AI* report concluded that a made-in-Ontario AI human rights impact assessment would be an important step to promote "Trustworthy AI" in this province.[2] The LCO subsequently partnered with the OHRC to to create this tool, collaborating closely with the Canadian Human Rights Commission (CHRC).

In Canada, AI impact assessments are not yet specifically required by law, but legislation at both the provincial and federal level suggest that HRIA obligations could be introduced for many Canadian AI systems. For example, the proposed federal *Artificial Intelligence and Data Act* (*AIDA*) would require AI developers to identify, assess, and mitigate potential bias and harm in the design and development of "high impact" AI systems before they are deployed.[3] Similarly, Ontario's recently passed *Enhancing Digital Security and Trust Act, 2024* (*EDSTA*)[4] would require certain public sector entities to "develop and implement an accountability framework respecting their use of the artificial intelligence system."[5]

There are also many Canadian policy directives at the federal, provincial, and agency-level that require or suggest AI impact assessments, including the federal Automated Decision-making Directive (federal ADM Directive) and its companion, the Algorithmic Impact Assessment (federal AIA)[6], the Government of Ontario's "Responsible Use of Artificial Intelligence Directive" (the Ontario AI Directive)[7], and the Toronto Police Service Board's "Use of AI Technology Policy" (TPS AI Policy).[8]

Regardless of whether a human rights impact assessment is legally required or not, Canadian governments, agencies, and private sector organizations have a legal responsibility to comply with Canadian human rights law, including, where applicable, the *Canadian Charter of Rights and Freedoms*, the *Canadian Human Rights Act, the Ontario Human Rights Code*, and applicable human rights legislation in other jurisdictions. This responsibility applies to a wide range of AI systems that may be developed or deployed in Canada.

The LCO has learned that many Canadian organizations are developing (or have developed) dedicated AI governance models addressing a wide range of issues, including data security, privacy, and copyright. We have also learned that many of these initiatives do not address human rights or are based on international "ethical AI" norms or foreign laws rather than Canadian laws.[9]

The LCO/OHRC HRIA is the first AI human rights impact assessment based specifically on Canadian human rights law. It is a practical, step-by-step guide that public and private sector organizations can use to embed Canadian principles of "human rights by design" in Canadian AI systems.

The LCO emphasizes that the HRIA must be understood to be part of a larger strategy to promote human rights in AI systems. In the words of Data and Society and the European Center for Not-For-Profit Law, HRIAs are not a "silver bullet for addressing potential human rights impacts or for assessing the human rights, democracy, and rule of law impacts of AI systems."[10] HRIAs are one component of a sophisticated, multifaceted AI governance strategy that should include legislation, regulations, auditing requirements, disclosure requirements, oversight bodies, and remedies.[11]

2.
Purpose of this Paper

This Backgrounder accompanies the LCO/OHRC HRIA. It can be read as both an introduction to the HRIA and a summary analysis of AI impact assessments generally.

The LCO/OHRC HRIA is based on extensive research on AI regulation and AI impact assessment tools. The LCO consulted with a wide range of stakeholders and key informants, including government officials, lawyers, technologists, academics, NGOs, and representatives from both large and small private enterprises.

The Backgrounder summarizes the benefits, limitations, and strategic choices inherent in developing a HRIA. The LCO believes that policymakers and stakeholders will benefit from a comprehensive look at of the choices, opportunities, and challenges in this project.

3. Project Partners

The LCO/OHRC HRIA is the product of a unique collaboration between the Law Commission of Ontario (LCO) and the Ontario Human Rights Commission (OHRC), both of which are concerned about the proliferation and risks of AI systems in Ontario. The Canadian Human Rights Commission (CHRC) was a contributor to the HRIA as well. The Backgrounder is an LCO report.

3.1. Law Commission of Ontario

The LCO provides independent, balanced, and rigorous advice on complex and important legal policy issues. Through this work, the LCO promotes access to justice, evidence-based legislation and policies, and public engagement on important issues. The LCO is independent of stakeholder interests and is committed to a "public interest" perspective for every project.

The LCO has been working on AI law reform issues since 2019, focusing primarily on human rights, procedural fairness, and access to justice. Recent LCO reports or submissions include:

- Submission to Government of Ontario on Bill 194 (2024)
- *Accountable AI* (2022)
- *Regulating AI: Critical Issues and Choices* (2021)
- *Legal Issues and Government AI Development* (2021)
- The Rise and Fall of Algorithms in the American Justice System: Lessons for Canada (2020)

The LCO's AI and Human Rights Project Lead is Susie Lindsay. She can be contacted at slindsay@lco-cdo.org.

More information about the LCO's civil and administrative law AI projects is available here.

More information about the LCO's criminal AI projects is available here.

The LCO can be contacted at LawCommission@lco-cdo.org.

3.2. Ontario Human Rights Commission

The OHRC has a broad statutory mandate under the Ontario *Human Rights Code* to promote, protect and advance respect for human rights, and to identify and promote the elimination of discriminatory practices. The OHRC works in different ways to fulfill this mandate, including through education, policy development, public inquiries and litigation.

The OHRC has provided guidance to the Ontario government and other parties for a human rights-affirming approach to the use of AI, including:

- Submission to the Government of Ontario regarding Bill 194, Strengthening Cyber Security and Building Trust in the Public Sector Act, 2024
- Submission to the Standing Committee on Social Policy regarding Bill 149, Working for Workers Four Act, 2023
- Joint statement by the Information and Privacy Commissioner of Ontario and the Ontario Human Rights Commission on the use of AI technologies (2023)
- Submission on the Toronto Police Services Board's Use of Artificial Intelligence Technologies Policy (2021)
- Submission on Ontario's Trustworthy AI Framework (2021)

For information on the OHRC's actions on AI, visit ohrc.on.ca or write to info@ohrc.on.ca.

3.3. Canadian Human Rights Commission

The CHRC is Canada's national human rights institution, accredited "A-status" by the Global Alliance of National Human Rights Institutions. The CHRC was established by Parliament through the Canadian Human Rights Act (CHRA) in 1977. It has a broad mandate to promote and protect human rights.

The CHRC, pursuant to the CHRA, has jurisdiction to receive discrimination complaints against federal government departments and agencies, Crown corporations, First Nations governments, and federally regulated private sector organizations such as banks, airlines, and telecommunications companies. The CHRC also has a broad mandate to support research, raise public awareness, and issue reports on human rights issues.

The CHRC's work on AI and human rights includes:

- The CHRC's Annual Report to Parliament- Human rights and artificial intelligence (2023)
- The CHRC's written submission to the House of Commons Standing Committee on Access to Information, Privacy and Ethics on the use of facial recognition technology in policing (2022)

4. Artificial Intelligence Impact Assessments

An impact assessment is a process that involves documenting an undertaking, evaluating its impacts, and assigning responsibility for those impacts. There are many kinds of impact assessments, including environmental assessments, privacy assessments, and data security assessments, to name a few.

AI impact assessments are a distinct form of impact assessment. AI impact assessments have been defined as "emerging governance practices for delineating accountability, rendering visible the harms caused by algorithmic systems, and ensuring practical steps are taken to ameliorate those harms."[12]

Human rights AI impact assessments (HRIAs) are AI impact assessments focusing on human rights. HRIAs take many forms but are often used to identify bias in datasets; to assess the fairness or explainability of an AI system; or to identify broader human rights, equity, economic, or health impacts on users or communities potentially affected by an AI system.

AI impact assessments are an important AI governance strategy. AI impact assessments are increasingly being required by legislation or adopted by governments, public institutions, and private enterprises to demonstrate and implement their commitments to "trustworthy AI." Most notably, the European Union's *AI Act* requires a fundamental rights impact assessment be conducted for AI systems in the "high risk" category.[13]

The leading AI impact assessment tool in Canada is the federal government's Algorithmic Impact Assessment (federal AIA), a companion to the federal Automated Decision-making Directive (federal ADM Directive).[14] The federal AIA applies to many federal agencies who intend to introduce any automated decision-making tool that could impact legal rights, privileges or interests of clients, including individuals and businesses.[15] The federal AIA is one of the world's first AI impact assessments.[16] The LCO has written extensively about the federal ADM Directive and federal AIA in our reports.[17]

The Council of Europe adopted the first ever legally binding Framework Convention on AI and Human Rights, Democracy and the Rule of Law in 2024 to which Canada is a signatory. The Framework Convention presents principles for States to follow throughout the lifecycle of AI systems, and requires a risk and impact management framework be in place to identify,

assess, prevent and mitigate risks posed by AI systems. The proposed risk methodology is the Human Rights, Democracy and Rule of law Impact Assessment (HUDERIA) that offers processes to evaluate the socio-technical contexts of an AI system's development and use, and detailed insight into how to identify, assess and mitigate risks and potential impacts.

Sophisticated examples of AI impact assessments in other jurisdictions include:

- The Council of Europe Framework Convention on AI and Human Rights, Democracy and the Rule of Law's Human Rights, Democracy and the Rule of Law Impact Assessment (2024)[18]
- Australian Human Rights Commission's Human Rights Impact Assessment Tool: AI-informed Decision-making Systems in Banking (2023)[19]
- Microsoft's *Responsible AI Standard and Impact Assessment* (2022)[20]
- The University of Oxford capAI tool (2022)[21]
- The European Commission's *Assessment List for Trustworthy AI* (2020)[22]
- The Netherlands' Fundamental Rights and Algorithms Impact Assessment (FRAIA) (2021)[23]
- The Spanish Data Protection Agency's "Guide to Algorithmic Auditing" (2021)[24]
- Algorithm Watch's *Impact Assessment Tool for Public Authorities* (2021)[25]
- The ECP *AI Impact Assessment* (2018)[26]
- AI Now's *Algorithmic Impact Assessment* (2018)[27]

4.1. Structure of AI Impact Assessments

A review of AI impact assessments from other jurisdictions shows that there are many choices when designing an assessment tool.

AI impact assessments typically take the form of detailed, structured questionnaires addressing a long list of topics. More sophisticated tools will also identify risk categories and risk mitigation strategies. Some tools produce a compliance "score" or rating, while others simply identify areas for further development. The questions in an AI impact assessment can address every phase of an AI lifecycle or may be limited. Topics covered in an AI impact assessment often include:

- Purpose for the AI system
- Parties involved in designing the system
- Data source
- Validity and bias
- Impact on individuals or communities
- Harm mitigation strategies
- Explainability
- Human-in-the-loop
- Testing/reporting/ongoing evaluation

4.2. Benefits

The US research organization Data and Society and the European Center for Not-For-Profit Law state that AI impact assessments provide "an *ex ante* or *ex post* assessment of the potential or actual impacts of a technology, policy, or business practice."[28] Benefits include:

> *Creating a document of these impacts that can be shared with stakeholders, regulators, or members of the public.*
>
> *Providing a reflexive exercise for developers of a policy or technology to question what outcomes they hope to achieve, and what mitigative measures they may need to put in place to address potentially harmful outcomes.*
>
> *Providing a mechanism for developers and policymakers to engage with a range of stakeholders who may be affected by that policy or technology.*[29]

Many algorithmic justice advocates, scholars, technology companies and policymakers also believe AI impact assessments have the potential to address algorithmic harms far beyond narrowly constructed metrics and towards more substantive justice.[30]

Ex ante AI impact assessments are seen as particularly beneficial. This is because *ex ante* AI impact assessments have the potential to be effective at identifying and mitigating risks *before* widespread investment and deployment of an AI system. *Ex ante* AI impact assessments are also considered important tools to promote "trustworthy AI by design" and to embed human rights and procedural justice principles throughout the AI lifecycle.

Another important but perhaps underappreciated benefit of AI impact assessments is that they "offer a means to describe, measure, and assign responsibility for impacts without the need to encode explicit, scientific understandings in law."[31]

Finally, an impact assessment can be an important educational tool. The LCO's consultations revealed widespread inexperience with, and misunderstanding about, human rights laws in Ontario and Canada. As a result, these laws are often overlooked, ignored, or considered superficially. A HRIA can help people and organizations not trained in human rights law or policy consider human rights in the design of their AI systems.

4.3. Limitations and Challenges

Notwithstanding their benefits, many human rights organizations and others have warned that AI impact assessments and HRIAs are not a panacea. These organizations have noted that inadequate or superficial AI impact assessments may become "little more than a box-ticking exercise, a fake self-assessment, or a toothless list of recommendations"[32] and "smokescreens for human rights and other abuses."[33] Other limitations and challenges include:

- Neither AI impact assessments nor HRIAs create accountability by themselves. AI impact assessments and HRIAs produce a document which can enable a change in practice, but they do not create an inherent binding commitment for change to occur.
- There is no single AI impact assessment or HRIA that can adequately consider the human rights implications of every AI system. AI systems used for credit scoring, employment hiring, determining eligibility for government benefits, and policing can all potentially affect human rights, but the context and technologies are likely to be very different.
- AI impact assessments and HRIAs can be co-opted. Like any governance process, AI impact assessments and HRIAs run a risk of being conducted in bad faith to legitimize an existing practice or policy.

- A hallmark of AI systems, especially machine learning and complex neural networks, is their changeability. The unpredictability in AI systems can lead to unintended or unexpected outcomes. An impact assessment attempts to predict an AI systems impact at a specific point in time. As systems and training data evolve, however, a HRIA may quickly become out of date unless it is updated regularly.

4.4. Where Do AI Impact Assessments Fit?

The LCO agrees with Data and Society and the European Center for Not-For-Profit Law, that AI impact assessments and HRIAs are not a "silver bullet for addressing potential human rights impacts or for assessing the human rights, democracy, and rule of law impacts of AI systems."[34]

In our view, AI impact assessments and HRIAs are best understood as one component of a sophisticated, multifaceted AI governance strategy that should include legislation, regulation, policies, and standards.[35] The LCO believes legislation and regulations should provide an overarching legal framework while standards, policies, and impact assessments provide more specific guidance on how to operationalize these obligations.[36]

AI-specific legislation has been slow to develop in Canada.[37] Although the governments of Canada and Ontario have signaled an intention to introduce statutory *ex ante* AI accountability requirements, there is currently a wide gap in Canadian AI legislation and regulations, especially for human rights.[38] Until that gap is filled, the LCO/OHRC HRIA will help public and private sector organizations proactively address their existing human rights obligations.

Many private and public entities are creating their own internal AI governance frameworks to support the development of responsible AI. To assist with this, the Vector Institute released "Principles in Action: a Playbook for Responsible AI Product Development" as a guide to help organizations develop trustworthy and safe AI. This playbook encourages the use of mitigation strategies, including conducting meaningful impact assessments, to address potential risks and harms of AI. Another example is the IRCC's "Policy Playbook on Automated Support for Decision-making" which is intended, in part, to help the IRCC comply with the Federal Directive and AIA.

5.
Designing the LCO/OHRC HRIA

The format and approach of the LCO/OHRC HRIA was informed by our review of many AI impact assessments in the United States, Europe, and Canada.[39]

We also studied the growing international literature on AIAs.[40]

Most importantly, however, the LCO consulted with stakeholders and key informants across Ontario. Our consultees included data scientists and technologists, government representatives, public agency representatives, AI project managers and compliance officers, human rights experts, lawyers, academics, NGOs, and representatives from large and small private enterprises. The format, structure, and substance of the HRIA was strongly influenced by the thoughtful and candid feedback we received during these consultations.

This part of the Backgrounder summarizes the most significant issues and choices involved in developing the HRIA. The LCO believes that stakeholders will benefit from a comprehensive look at the HRIA to appreciate the choices, challenges, and issues the LCO and OHRC faced during this project. The LCO releases this backgrounder to help organizations use the HRIA effectively.

Issues considered in this section include:

- Should the HRIA be mandatory?
- Who are the assessors?
- What triggers the HRIA?
- Which AI systems are subject to an HRIA?
- Will the HRIA be disclosed?
- Is the HRIA too narrow?
- Is the HRIA too general?
- How detailed should the HRIA be?
- Does the HRIA give a score, passing grade, or licence?
- How does the HRIA align with current or pending legislation?

5.1. Should The HRIA Be Mandatory?

Where does or should the authority for the HRIA come from? Should it be required by statute or regulation, set out in some kind of directive, or a voluntarily assumed "ethical AI" guideline or best practice?

Many consultees suggested the LCO/OHRC HRIA had to be mandatory to be effective. Absent a mandatory legal obligation, these commentators believed the HRIA would not be implemented comprehensively and thoughtfully. These consultees also worried about "ethics washing", fearing that some organizations will use the HRIA to promote AI ethics without complying with human rights law in a meaningful or comprehensive way.

Other commentators disagreed, believing it was best to allow organizations (particularly private sector organizations) to adopt and adapt the HRIA voluntarily. These commentators were worried about prescriptive, detailed requirements that could deter innovation or were not applicable to specific contexts. These commentators also pointed out that guidelines, policies and standards tend to be nimble and adaptable to rapidly changing technologies. Finally, these commentators asked about how detailed, prescriptive requirements could be incorporated into existing internal compliance procedures or policies.

Many of the AI impact assessments the LCO researched are voluntary.[41] Our research revealed, however, that AI impact assessments are increasingly being required directly or indirectly by AI legislation or regulations. Most notably, the EU *AI Act* requires developers of "high-risk" AI systems to conduct a fundamental rights assessment before deploying the system to the public.[42] The EU *AI Act* does not include a specific impact assessment tool but provides a range of necessary procedural requirements.[43]

Another recent example is the new *Colorado Artificial Intelligence Act* (Colorado AI Act), which requires deployers of "high-risk systems" to conduct impact assessments annually and within 90 days of any substantial changes. The Colorado AI Act also requires developers to provide documentation required for the deployer to conduct an impact assessment.[44] The Colorado AI Act does not provide a specific impact assessment, but it does require, among other things, that the impact assessment include an analysis of risks of algorithmic discrimination and steps taken to mitigate the risks.[45]

In Canada, there are several examples where AI impact assessments appear to be mandatory legal or administrative requirements. For example, the federal AIA is both mandatory and explicitly articulated in a Treasury Board Directive.[46] Other examples of *ex ante* impact assessment requirements include the proposed federal *Artificial Intelligence and Data Act* (*AIDA*), which would require AI developers to identify, assess, and mitigate potential bias and harm in the design and development of "high impact" AI systems before they are deployed[47]; Ontario's recently passed *Enhancing Digital Security and Trust Act*, *2024 (EDSTA)*[48] and *Use of Artificial Intelligence Directive*[49] (Ontario AI Directive), and the Toronto Police Service Board's *Use of AI Technology Policy* (TPS AI Policy).[50] The Council of Europe's Framework Convention on AI, to which Canada is a recent signatory, requires parties adopt measures to ensure AI systems are consistent with obligations to protect human rights.

The LCO believes the mandatory vs. voluntary HRIA issue is not a binary choice. It is possible to both legislate mandatory human rights obligations and ensure sufficient flexibility to adapt to different sectors and evolving technology.

As a starting point, the LCO does not believe a completely voluntary HRIA is sufficient to promote and protect human rights in Canadian AI systems. As a result, the LCO believes some form of AI impact assessment or HRIA requirement should be adopted into federal and provincial legislation or regulation.[51] A statutory framework, including prohibitions, is necessary to establish the level of public and legal accountability commensurate with the issues and rights at stake.[52]

Specifying an *obligation* to complete an AI human rights impact assessment does not predetermine the *form* of that obligation. Indeed, the LCO would be wary of enshrining the details of the LCO/OHRC HRIA or any other AI impact assessment into legislation or regulations. In our view, legislation or regulations should provide an overarching legal obligation to conduct a HRIA, while standards, policies, directives, and "playbooks" can be used to operationalize the legal obligation. In this manner, a HRIA can be adapted and adopted by organizations to address their specific issues and context.

In Canada, this "law/standard" model has been largely adopted in the federal ADM Directive and federal AIA.[53] Elements of this model were also incorporated into *AIDA*, but that legislation is unlikely to be passed.[54] Ontario's *EDSTA* and AI Directive also offer some potential to establish this model.[55]

In our view, the best approach would be for both the federal and provincial government to enshrine an explicit legislative mandate for public and private sector organizations to conduct a human rights impact assessment. A less preferable but still welcome alternative would be to establish a human rights impact assessment requirement in regulations. The key is to establish a binding legal obligation to assess human rights compliance. Once established, this legal authority would promote the use of the LCO/OHRC HRIA (or an equivalent instrument adapted to specific contexts) and strongly mitigate the risk of "ethics washing."

In the meantime, the LCO/OHRC HRIA is not a mandatory or legal obligation. Nevertheless, under existing human rights legislation and the *Charter*, both the public and private sectors have an obligation to take reasonable measures to ensure that the services and products they provide do not discriminate.[56] The HRIA is an important tool to help developers and organizations comply with these obligations. Accordingly, the HRIA should be seen as an opportunity to identify human rights concerns in AI systems, promote good practices, and comply with existing human rights legislation in Ontario and Canada.

5.2. Who Are The Assessors?

Data and Society notes that "[e]very aspect of an impact assessment is deeply connected with who conducts the assessment."[57] This comment highlights the potential risk to the accuracy and credibility of a HRIA if the assessor is neither skilled nor independent.

Many reports suggest that assessors ought to include technical experts, social science experts, human rights experts, and community members, particularly community members from marginalized and vulnerable groups.[58]

Virtually all consultees, including private sector representatives, supported the need for multidisciplinary AI assessments. Many consultees commented, however, that it was both impractical and unnecessary to incorporate the full range of potential assessors into each assessment. Many commentators from large organizations also believed the best approach was to incorporate a HRIA into an organization's existing regulatory compliance framework and processes.

The LCO/OHRC HRIA does not dictate who should be involved in completing the HRIA in all circumstances. Rather, the HRIA incorporates a risk-management and best practices approach that recommends a process for how the tool should be completed and who should be involved. Our objective was to help organizations develop internal risk management and quality control measures consistent with human rights protection.

The LCO often heard that AI development and deployment is a "team sport." Successful implementation of AI systems will involve a diverse set of stakeholders, including developers, program managers, senior executives, compliance officers, and experts in privacy, data security, and human rights. We also heard that the AI development team should include persons with lived experience and community members who may be subject to the AI system. In our consultations, we heard of many examples of this approach, especially in large organizations in heavily regulated sectors.

Finally, we were often asked who had responsibility for reviewing or evaluating the HRIA. Do governments, agencies, or firms complete the HRIA without any independent or external review or oversight? If so, does the lack of external review raise the potential for "ethics washing"? Or is an external review requirement unnecessary given that organizations will always bear the legal responsibility for their AI systems once they are deployed?

In our view, external review, evaluation and oversight of the HRIA will significantly improve the effectiveness of the HRIA and the accountability of entities deploying it, especially for high-risk AI systems. Self-evaluation can be effective in low-risk contexts, but self-evaluation is not sufficient when human rights are at stake.[59] As a result, the LCO believes external evaluation of an AI system, or at least an external review of the HRIA, is necessary for high-risk AI systems.

This raises the question of how an external review should be conducted. The LCO notes the growing industry of established firms and tech start-ups positioning themselves as compliance or verification experts to review and evaluate AI systems for compliance issues, including human rights compliance. Although these firms and start-ups may provide a useful service, there is a concern that the privatization of regulation can create issues with independence, capture, and lax oversight.[60] The LCO believes that for issues as significant as human rights, some form of government oversight is necessary at the very least for high risk AI systems.

5.3. What Triggers The HRIA?

We were often asked about the "trigger" for the HRIA. At what point in the development or deployment of an AI system is there an obligation to undertake the HRIA? The trigger for a HRIA could arise in several circumstances. For example, the HRIA could be required or suggested at predetermined stages or fixed milestones of the AI development lifecycle. Alternately, the HRIA could be required if an AI system is to be deployed in potentially "high risk" applications, such as policing or the judicial system.

Many consultees noted that human rights issues can be quite different when an AI system is at the idea stage versus beta testing or deployment. Many developers (from both the public and private sectors) were concerned about regulatory or procedural requirements that applied too early in AI development, particularly when AI models were being "workshopped" or tested.

The LCO is mindful of the legitimate questions and concerns raised by stakeholders about the appropriate timing for the HRIA.

The LCO's key message is that human rights need to be considered throughout the AI development and deployment lifecycle. In the LCO's view, the best approach is to consider the human rights implications of an AI system at the following stages:

- When the idea for the AI system is explored and developed.
- Before the AI system is made available to external parties (e.g.: before a vendor makes a model or application available to purchasers, or service providers deploy an AI technology for customer service).
- Within ninety days of a material change in the system; and
- Yearly as part of regular maintenance and reviews.

The LCO wants to emphasize that our goal is to embed "human rights by design" principles throughout the lifecycle of an AI system. Accordingly, the LCO believes developers and organizations should be

mindful of potential human rights implications even at the earliest stages of AI development. By way of illustration, we often heard that privacy is "top of mind" during AI development well before any formal privacy impact assessment is undertaken. We are hopeful that the HRIA will embed human rights awareness in a similar manner.

5.4. What About Low Risk AI Systems? Exemptions?

Many consultees wanted to know whether an organization needed to complete the HRIA for every AI system the organization was developing. These questions tended to take two forms. First, we were asked if some AI systems did not need to be evaluated because they could not reasonably affect human rights or influence decisions affecting a person. Second, we were asked if there were some activities or sectors that were pre-emptively excluded from a HRIA.

Risk-Based HRIA Exceptions

The first question raises important operational questions for every organization developing AI systems. AI regulatory models typically triage the need for impact assessments through risk-based regulations or policies. In many legislative or regulatory models, the need for a HRIA is tied to the risk the AI system poses to individual rights. For example, the EU *AI Act* requires conformity or fundamental risk assessments to be conducted on AI systems that are deemed "high risk". The same obligation does not apply to AI systems deemed "minimal risk" or "limited risk".[61] Similarly, the Colorado *AI Act* attaches requirements only to "high risk" AI system. [62] By comparison, the federal ADM Directive requires that any government system, statistical model, or tool used to make an administrative decision must be subject to an algorithmic impact assessment.[63]

The LCO supports risk-based regulation as an effective way to mitigate the potential harms of AI. The LCO notes, however, some important cautions or qualifications to risk-based regulation in the context of human rights.

Human rights are immutable and, with rare exceptions, paramount. As a matter of law, human rights take priority over other interests, such as corporate efficiency. Further, there will invariably be AI systems that are not captured by categories described as "high risk".[64] *AIDA*, for example, would not apply to most government services and leaves out areas where discrimination is an ongoing issue, such as accommodations and education.[65] Similarly, the federal ADM Directive currently excludes biometric AI systems, behaviour assessment AI tools, and AI systems designed for fraud detection when these AI technologies are used outside the context of administrative decision-making. These AI systems have proven to be some of the most serious violators of human rights. As of February 2025, Treasury Board Secretariat (TBS) is undergoing the fourth review of the Federal Directive and Algorithmic Impact Assessment. As part of the review, TBS is considering bans on certain AI systems. Finally, the Canadian Human Rights Commission notes that

> *…some AI applications that are classified as low risk because of the seemingly benign nature would not be subject to proper oversight and safeguards. This also has the potential for unintended systemic discriminatory impacts from multiple, cumulative, or intersecting impacts of many different "low-risk" systems.*[66]

There are no circumstances where human rights law is not relevant or applicable. Nor should organizations assume their AI systems do not have human rights implications. As a result, the LCO believes that all AI systems should be subject to some form of human rights impact analysis. In no way does this mean that all AI systems must be treated the same. The LCO/OHRC HRIA, like most sophisticated AI impact assessments, identifies and sorts "low risk" AI systems through a step-by-step risk management framework that helps organizations determine where their system falls on a continuum of human rights risks. If the system is likely to be low risk, no further action is recommended.

HRIA Exemptions or Carve-outs

The question of exemptions or carve-outs for AI impact assessments is very controversial. Some AI legislative AI governance models specifically exclude certain AI systems or uses from their accountability requirements, including AI impact assessments. For example, Ontario's *EDSTA*, the provincial government's public sector AI legislation, includes exceptions for AI technologies used in the criminal justice system (including policing) and by courts and tribunals.[67]

The LCO opposes carve-outs or exemptions for organizations or sectors from conducting a HRIA. The *EDSTA*'s exception for AI systems deployed in the criminal justice system or court system is especially concerning. This carve-out means that some of the highest risk, most controversial AI systems (such as facial recognition surveillance systems, predictive policing AI systems) could be effectively unregulated. These are systems where there is the greatest need for a sophisticated HRIA.[68] The LCO's criminal AI project addresses these risks in more detail.[69]

5.5. Will the HRIA be Disclosed?

The LCO believes that requiring meaningful disclosure and transparency of key elements of a HRIA is a crucial step towards AI accountability and public trust. Transparency and disclosure are widely acknowledged features of "trustworthy AI" governance.[70] As a result, transparency and disclosure of AI impact assessments is a major focus in emerging AI regulations and assessment tools.[71]

Virtually all legislative AI governance models include some form of disclosure. For example, the federal ADM Directive requires that departments release the results of the federal AIA in an accessible format in both official languages on the Open Government Portal.[72] The EU *AI Act* requires "high risk AI systems" to be registered in an EU database prior to being placed in the market or into service.[73] Finally, Ontario's *Working for Workers Four Act* requires employers who use AI to screen, assess or select applicants for a publicly advertised position to disclose the use of AI in the job posting.[74]

The need for some form of disclosure was widely acknowledged in our consultations. Consultees also generally agreed that public sector AI systems raised different legal, policy, and operational disclosure issues than private sector systems. Not surprisingly, there was less agreement about level of detail or comprehensiveness of disclosure that should be required in public and private AI systems.

Many consultees tied the disclosure principle to enforcement and asked whether parties would disclose a HRIA if they were not required to. We were also asked about disclosure obligations in cases where the HRIA identifies a human rights concern: can human rights issues in AI systems be addressed if parties do not disclose the results of the HRIA? Many private sector commentators raised the opposite issue: will organizations complete the HRIA if they must disclose human rights risks, including risks that have been mitigated? In these circumstances, some worried that a disclosure obligation created a perverse incentive *not* to undertake a HRIA.

As a starting principle, the LCO agrees that HRIA disclosure requirements differ between public sector and private sector AI systems. Public sector actors are subject to the *Charter* and administrative law requirements.[75] Governments and public agencies also properly have higher expectations for public accountability and transparency. As a result, the LCO believes there should be a mandatory legal obligation to post at least a summary of the complete HRIA on a public AI registry. The LCO stresses that this obligation should be mandatory. Experience demonstrates that governments and public agencies may not disclose AI systems, even when they have voluntarily committed to do so.[76]

The private sector generally has fewer legal disclosure obligations, although market forces may encourage private entities to disclose the HRIA or other policies and procedures that protect privacy, the environment, etc. That said, problems can arise with private sector AI systems if there were no incentive to disclose some information about human rights issues and mitigation strategies. Ideally, private sector HRIA disclosure could perhaps be tied to the level of risk of the AI system, as is done under the EU *AI Act*.[77]

5.6. Is the HRIA Too Narrow?

The HRIA is a dedicated tool to assess bias and discrimination as defined in Canadian and Ontario human rights legislation.

During early consultations, human rights experts raised concerns that an AI human rights impact assessment that focused exclusively on bias and discrimination could mean other human rights issues would be undetected or unaddressed.

This is a legitimate concern, as AI systems have been found to pose risks to human rights above and beyond bias and discrimination. For example, many biometric surveillance systems used by police services have been strongly criticized for their risk to privacy, freedom of speech, and freedom of association.[78] Nevertheless, the LCO and OHRC decided to focus the HRIA on bias and discrimination issues for several related reasons.

First, many existing AI impact assessments already address data security, data quality, and privacy, but do not address bias and discrimination comprehensively.[79]

Second, focusing on bias and discrimination allowed us to address a particularly complex and evolving issue comprehensively, effectively, and efficiently. An AI impact assessment that tried to address too many issues could either be too general or too long to be effective.

Third, the HRIA is intended to apply to both public and private sector organizations. Several human rights, such as freedom of expression, religion or association apply specifically to government and not to private entities.

Finally, our focus on bias and discrimination kept our project within the mandate of the OHRC and human rights law in Ontario.

During our consultations, we learned of the wide and deep range of AI impact assessments and risk matrices used by governments and the private sector. In particular, the LCO was advised that many large organizations already have "responsible AI" frameworks that include privacy and data security assessments, among other risk assessment measures. Many consultees emphasized the need to ensure the LCO/OHRC HRIA could be adapted operationally into these existing compliance mechanisms. Based on our consultations, we are confident the LCO/OHRC HRIA can and will be integrated into other AI impact assessments over time.

5.7. Is the HRIA Too General?

An ongoing question in AI impact assessment policy development is whether it is better to develop a single "generalist" tool or to develop dedicated tools for specific sectors or activities.

This issue was raised frequently in our consultations, where it arose in two forms: 1) We were told that AI systems in different contexts raise different issues. For example, AI used in policing is distinct from AI used in banking, and 2) We were told the HRIA is likely to be too onerous for small organizations. As a result, we were told there is a risk the HRIA could create a competitive advantage for large organizations that have existing compliance departments.

The overarching theme of these questions is whether a series of specific, targeted HRIAs would be preferable to a single, generalist tool.

In the end, the LCO and OHRC determined that a single, generalist HRIA is the best option at this stage of AI governance in Ontario. We came to this conclusion for several reasons.

First, the LCO and OHRC encourage organizations and sectors to adapt the HRIA to their specific contexts and purposes. An excellent illustration of this approach is the HRIA developed by the Australian Human Rights Commission and National Bank of Australia.[80] The Australian HRIA includes both general

questions and specific questions on whether the AI system is making a decision related to specific banking activities (e.g.: creditworthiness of an individual, the pricing of banking products, automated customer advice etc.).[81] This precedent is a good example of how sectors or organizations can adopt and adapt a human rights analysis for their own targeted purposes. The LCO encourages Canadian organizations and sectors to consider how the LCO/OHRC HRIA can be adapted in a similar manner.

Second, there is a need for consistency in AI human rights assessment, irrespective of the specific context or scale of the AI system. A generalist tool is the foundation for this consistency.

Finally, the LCO/OHRC HRIA will help organizations comply with *existing* legal human rights obligations. The LCO/OHRC HRIA does not introduce or create new legal obligations. All organizations in all sectors are required to uphold human rights.

5.8. Does the HRIA Provide a Score, Passing Grade, or License?

Many consultees asked whether the HRIA would provide a score or passing grade upon completion.

Many private sector consultees suggested the HRIA would have a much greater impact if it led to some kind of certification or license that could be publicized or relied upon in the event of a human rights challenge to their AI system.

Not surprisingly, human rights experts and academics advised against providing a score or license. They pointed out that human rights law is too complex and fact-specific, and AI systems are too unpredictable, to certify or warrant the pre-approval of an AI system. They also worried that a certification model would potentially absolve parties of liability should problems occur down the road.

The LCO does not believe the HRIA can or should provide a passing grade or license. We understand how a HRIA certification or license would promote use of the tool. At the same time, we do not believe an *ex ante* impact assessment can or should absolve organizations of potential liability for human rights violations once the system is in operation.[82]

In the end, we settled on a reasonable, practical compromise: while the HRIA does not provide a score or passing grade, Part A of the HRIA provides six categories that an AI system can be slotted into depending on their potential human rights issues, and Part B provides recommendations on several widely acknowledged mitigation strategies in the event a human rights issue is discovered.

5.9. Is the HRIA Consistent with Existing or Proposed Legislation and Regulations?

Governance models to regulate AI are being developed, debated, and implemented around the world rapidly.[83] In Canada, there are a host of AI proclamations, policies, and frameworks at the federal, provincial, and municipal levels. Although there is more work to be done on AI specific legislation and regulations, both the provincial and federal governments have signaled a need for *ex ante* assessments of certain AI systems.

In our consultations, many private sector commentators told us that creating an additional AI impact assessment could create regulatory overburden, conflict with existing impact assessments, and frustrate innovation. Many commentators also asked whether a made-in-Ontario HRIA would overlap or be inconsistent with legislative requirements in other jurisdictions (i.e. the European Union and United States). Finally, we were also asked how the HRIA fits with privacy laws or existing, new, or proposed AI legislation at the federal and provincial levels.

At a general level, questions about how the HRIA fits with the EU *AI Act*, American AI requirements, *AIDA*, *EDSTA,* or privacy laws are uncomplicated. The HRIA is based on requirements under existing Canadian human rights law. As a matter of law, the *Charter*, *Canadian Human Rights Act* and *Ontario Human Rights Code* take precedence over other laws and regulations in the event of a conflict. Foreign laws or legal obligations do not supersede Canadian human rights law, and Canadian laws such as *AIDA* or *EDSTA* must conform to it. The need to comply with Canadian human rights law exists irrespective of whether there is an independent or specific legal obligation to undertake a HRIA.

The more specific answer to this question is that the LCO/OHRC HRIA was designed to be consistent with all current or proposed AI legislation in Canada and Ontario.

For example, if passed, *AIDA* would have introduced *ex ante* human rights obligations for AI developers, including obligations to identify and mitigate potential bias and harm in the design and development of AI systems.[84] *AIDA* would have also created documentation and transparency requirements. The LCO/OHRC HRIA could be used to fulfill these requirements. All or parts of the LCO/OHRC HRIA can also be incorporated into the federal ADM Directive and AIA.

The same is true for recent provincial legislation and Directives. The *EDSTA* does not include reference to human rights but indicates that certain risk management measures such as an "accountability framework", "disclosure" and "reporting" obligations and prohibitions will apply to AI systems used in some contexts. The LCO/OHRC HRIA could be used, in part, to fulfill these requirements. Similarly, the purpose of Ontario's AI Directive is to "set out the requirements for the transparent, responsible and accountable use of Artificial Intelligence".[85] Amongst other requirements, the Directive includes AI risk management requirements[86] and disclosure and transparency obligations,[87] both of which could be fulfilled in part by the HRIA.

6. Conclusion and More Information

The LCO/OHRC HRIA is the first AI human rights impact assessment based specifically on Canadian human rights law. It is designed to assist developers and administrators of AI systems identify, assess, minimize, or avoid discrimination and uphold human rights obligations throughout the lifecycle of an AI system.

Every private sector and public sector organization in Ontario and Canada must ensure they do not discriminate pursuant to human rights legislation. The HRIA is a practical, step-by-step guide that organizations can use to embed principles of "human rights by design" in Canadian AI systems.

The LCO wants to hear from a broad range of Ontarians about AI impact assessments, human rights, and AI governance generally. The LCO is committed to sharing ideas and building constructive dialogue on these issues across Ontario and Canada. As a result, we encourage any interested individual or organization to contact us. We are also available to discuss the HRIA and other AI and human rights issues with interested organizations or groups.

Project Lead and Contacts

The LCO's AI and Human Rights Project Lead is Susie Lindsay. She can be contacted at slindsay@lco-cdo.org.

More information about the LCO's civil and administrative law AI projects is available here. More information about the LCO's criminal AI projects is available here. The LCO can be contacted at LawCommission@lco-cdo.org or

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

LawCommission@lco-cdo.org

Endnotes

1 The LCO/OHRC Human Rights Impact Assessment was released in November 2024 and is available at https://www.lco-cdo.org/en/our-current-projects/ai-adm-and-the-justice-system/human-rights-ai-impact-assessment/.

2 Law Commission of Ontario, *Accountable AI, (2022) [Accountable AI], online:* https://www.lco-cdo.org/en/our-current-projects/ai-adm-and-the-justice-system/ai-and-adm-in-the-civil-administrative-justice-system/.

3 *Digital Charter Implementation Act*; *1st Sess. 44th Parliament, 2022, Part 3 "Artificial Intelligence and Data Act",* [*AIDA*], online: *https://www.parl.ca/legisinfo/en/bill/44-1/c-27*, s. 8. AIDA ceased to exist when Parliament was prorogued. We discuss AIDA in this paper as a sign of how the Federal government views potential AI legislation.

4 *Enhancing Digital Security and Trust Act*, S.O. 2024, c. 24, [*EDSTA*], online: https://www.ontario.ca/laws/statute/24e24.

5 *EDSTA*, section 5(3).

6 Canada, Directive on Automated Decision-Making (2019) [Canada ADM Directive], online: https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592; and Algorithmic Impact Assessment Tool [Canada AIA], online: https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html.

7 Government of Ontario, Ministry of Public and Business Service Delivery and Procurement, *Responsible Use of AI Directive*, December 1, 2024 [Ontario AI Directive]. The Directive does not explicitly require an impact assessment. However, it includes AI risk management requirements. For example, at section 6.1 parties must explain the use of the AI system, identify and assess risks of the AI system, plan and implement methods to control those risks, and report and monitor the efforts to assess and control the risks.

8 Toronto Police Services Board, "Use of Artificial Intelligence Technology" (February 28, 2022; updated January 11, 2024) [TPS AI Policy], online: https://tpsb.ca/policies-by-laws/board-policies/195-use-of-artificial-intelligence-technology.

9 Documents on file with the LCO.

10 Data and Society and the European Center for Not-For-Profit Law, R*ecommendations for Assessing AI Impacts to Human Rights, Democracy, and the Rule of Law* (2021) [Recommendations for Assessing AI Impacts], online: https://datasociety.net/announcements/2021/11/22/recommendations-for-incorporating-human-rights-into-ai-impact-assessments/?_rt=MnwxfGFzc2Vzc2luZyBhaSBpbXBhY3RzfDE3NDAwNjg0NjQ&_rt_nonce=3855b5a8a8 at 5.

11 For a good description of the components of a comprehensive "Trustworthy AI" legislative framework, see the following LCO submissions and reports: *Bill 194, Law Commission of Ontario Submission (2024) [LCO Bill 194 Submission], online:* https://www.lco-cdo.org/en/lco-releases-bill-194-submission/; *Accountable AI; and Regulating AI: Critical Issues and Choices, (2021) [Regulating AI], online:* https://www.lco-cdo.org/en/our-current-projects/ai-adm-and-the-justice-system/regulating-ai-critical-issues-and-choices/.

12 Jacob Metcalfe, quoted in accessnow, Human Rights Impact Assessments for AI: Analysis and Recommendations, [accessnow Human Rights Assessments], online at: https://www.accessnow.org/human-rights-impact-assessment-ai/ at 8.

13 *Artificial Intelligence Act*, Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 [EU *AI Act*], online: https://eur-lex.europa.eu/eli/reg/2024/1689/oj. Article 6 of the AI Act sets out classification rules for high-risk AI systems. Article 27 requires that deployers of AI system must complete a fundamental rights impact assessment (FRIA) before using the system.

14 Canada ADM Directive, s. 5.1.

15 Canada ADM Directive, Appendix A and "Guide on the Scope of the Directive on Automated Decision-Making", online: https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/guide-scope-directive-automated-decision-making.html.

16 The Canada ADM Directive took effect April 1, 2019 with compliance required by April 1, 2020. By comparison, the EU *AI Act* started consultations in 2018 and was passed in March 2024.

17 See LCO reports *Accountable AI* and *Regulating AI*.

18 Council of Europe, *Framework Convention on Artificial Intelligence* (2024), online: https://www.coe.int/en/web/artificial-intelligence and the Human Rights, Democracy and the Rule of Law Impact Assessment (December 2024) [HUDERIA], online: https://www.coe.int/en/web/portal/-/huderia-new-tool-to-assess-the-impact-of-ai-systems-on-human-rights.

19 Australian Human Rights Commission, *Human Rights Impact Assessment Tool: AI-informed Decision-making Systems in Banking* (2023) [AI-informed Decision-making Systems in Banking], online: https://humanrights.gov.au/our-work/technology-and-human-rights/publications/hria-tool-ai-banking#:~:text=The%20aim%20of%20the%20HRIA,for%20any%20human%20rights%20violations.

20 Microsoft, "Responsible AI Impact Assessment Template" (2022), online: https://blogs.microsoft.com/wp-content/uploads/prod/sites/5/2022/06/Microsoft-RAI-Impact-Assessment-Template.pdf; and "Microsoft Responsible AI Impact Assessment Guide", (June 2022), online: https://blogs.microsoft.com/wp-content/uploads/prod/sites/5/2022/06/Microsoft-RAI-Impact-Assessment-Guide.pdf.

21 Floridi, Luciano and Holweg, Matthias and Taddeo, Mariarosaria and Amaya, Javier and Mökander, Jakob and Wen, Yuni, *capAI - A Procedure for Conducting Conformity Assessment of AI Systems in Line with the EU Artificial Intelligence Act,* (March 23, 2022), online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4064091.

22 European Commission, "Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment, (July 17, 2020), online: https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment.

23 Government of the Netherlands, "Fundamental Rights and Algorithms Impact Assessment" (July 31, 2021), online: https://www.government.nl/documents/reports/2021/07/31/impact-assessment-fundamental-rights-and-algorithms.

24 Eticas Research and Consulting under supervision of the Spanish Data Protection Agency, "Guide to Algorithmic Auditing", (January 2021), online: https://eticas.ai/wp-content/uploads/2024/04/Guide-to-Algorithmic-Auditing-EN.pdf.

25 Algorithm Watch, "Automated Decision-Making Systems in the Public Sector – An Impact Assessment Tool for Public Authorities", (June 22, 2021), online: https://algorithmwatch.org/en/adms-impact-assessment-public-sector-algorithmwatch/.

26 ECP, "Artificial Intelligence Impact Assessment (English Version) (2018), online: https://ecp.nl/publicatie/artificial-intelligence-impact-assessment-english-version/.

27 Dillon Reisman, Jason Schultz, Kate Crawford and Meredith Wittaker, "Algorithmic Impact Assessments Report: a Practical Framework for Public Agency Accountability", (April 9, 2018), online: https://ainowinstitute.org/publication/algorithmic-impact-assessments-report-2.

28 Data and Society and European Center for Not-For-Profit Law, *Recommendations for Assessing AI Impacts to Human Rights, Democracy, and the Rule of Law* (2021) [*Recommendations for Assessing AI Impacts*], online: https://datasociety.net/announcements/2021/11/22/recommendations-for-incorporating-human-rights-into-ai-impact-assessments/ at 3.

29 *Recommendations for Assessing AI Impacts.*

30 Data and Society, *Assembling Accountability* (June 29, 2021) [*Assembling Accountability*], online: https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest/ at 4.

31 *Assembling Accountability* at 5.

32 *Recommendations for Assessing AI Impacts* at 5.

33 accessnow, H*uman Rights Impact Assessments for AI: Analysis and Recommendations* (2022) [accessnow AI Analysis] ([] in italics), online: https://www.accessnow.org/human-rights-impact-assessment-ai/ at 4.

34 accessnow, "Why we need human rights assessments for AI" (2023), online: https://www.accessnow.org/human-rights-impact-assessment-ai/.

35 See generally, LCO Bill 194 Submission, *Accountable AI*, and *Regulating AI* reports.

36 See generally, LCO Bill 194 Submission, *Accountable AI*, and *Regulating AI* reports.

37 As of February 2025, Canada still does not have national AI legislation. The Government of Ontario's *EDSTA* was only passed in late 2024.

38 See *Accountable AI* pages 40-55 where the LCO discusses gaps in human rights law and LCO *Bill 194 Submission* pages 4-7 where the LCO submits that Bill 194 should be amended to acknowledge the province's human rights obligations. The Federal Government's *AIDA* legislation would have addresses some AI human rights issues but is unlikely to pass.

39 See generally the list of impact assessments on page 8 of this Backgrounder.

40 There is a growing international literature addressing AI impact assessments. Notable examples including accessnow AI Analysis; Miranda Bogen, Centre for Democracy and Technology, Assessing AI, (January 20250, online: https://cdt.org/insights/assessing-ai-surveying-the-spectrum-of-approaches-to-understanding-and-auditing-ai-systems/;

Recommendations for Assessing AI Impacts; Assembling Accountability; Lara Grove, Ada Lovelace Institute, "Realizing the potential of algorithmic accountability mechanisms" (February 2023), online: https://www.adalovelaceinstitute.org/blog/algorithmic-accountability-mechanisms/; Ada Lovelace Institute, *Examining the Black Box*, (April 202), online: https://www.adalovelaceinstitute.org/report/examining-the-black-box-tools-for-assessing-algorithmic-systems/; and the Ada Lovelace Institute's *Algorithmic impact assessment: user guide* (February 2022), online: https://www.adalovelaceinstitute.org/resource/aia-user-guide/.

41 See generally the list of impact assessments on page 8 of this Backgrounder.

42 EU *AI Act,* Article 27.

43 EU *AI Act*, Article 27, s. 1(a)-(f), which states, for example, that assessments must consist of "categories of natural persons and groups likely to be affected" (c); "specific risks of harm likely have an impact on categories of natural persons or groups" and (d); description of "human oversight measures".

44 Colorado *Artificial Intelligence Act,* Senate Bill 24-205 (signed into law on May 17, 2024 and set to take effect February 1, 2026) 6-1-1702.2; 6-1-1703.3 [*Colorado AI Act*], online: https://leg.colorado.gov/bills/sb24-205.

45 *Colorado AI Act,* 6-1-1703 (b)(II).

46 Canada ADM Directive, s.6.1 and Canada AIA.

47 *AIDA*

48 *EDSTA*

49 Ontario AI Directive.

50 TPS AI Policy.

51 LCO Bill 194 Submission.

52 See generally the analysis in the LCO's Regulating AI report discussing the differences between "ethical AI" and "hard law" approaches to AI regulation, especially at pages 23-28.

53 The Canada ADM Directive mandates that all ADM systems within its scope conduct the AIA.

54 As drafted, *AIDA* required persons responsible for AI systems to assess whether their systems were high-impact and to "establish measures to identify, assess and mitigate the risks of harm or biased output" in accordance with *AIDA* regulations. *AIDA*, sections 7 and 8.

55 See generally, LCO Bill 194 Submission. *EDSTA* legislates the potential to create *ex ante* mitigation measures. However, what those measures are, what systems they apply to, and whether or when the regulations providing these details are created is unclear.

56 For more information about public and private sector legal obligations to comply with human rights law, including the *Charter*, see the LCO/OHRC HRIA and *Accountable AI* at 40-56.

57 *Recommendations for Assessing AI Impacts* at 13.

58 See generally, *Accountable AI* pages 39 and 50; *Legal Issues and Government AI Development: Workshop Report*, March 2021 5; *Regulating AI*.

59 Druthi Suresh, "The Case for Mandatory Regulation: Jurisprudence Showcases the Need to Move On from Self-Regulation in Tech", September 2024; online: https://www.business-humanrights.org/en/blog/the-case-for-mandatory-regulation-jurisprudence-showcases-the-need-to-move-on-from-self-regulation-in-tech/.

60 Gillian Hadfield and Jack Clark, "Regulatory Markets: The Future of AI Governance", (April 2023), online: https://ideas.repec.org/p/arx/papers/2304.04914.html. This paper provides a thoughtful overview of the benefits and concerns with private regulators. See pages 18-20 for a discussion of the concerns.

61 EU *AI Act*, Articles 27 and 47.

62 *Colorado AI Act*, 6-1-1703.3.

63 Canada ADM Directive, s. 5.1. Note that the Directive applies only to federal departments subject to Treasury Board Secretariat.

64 See, for example, Access Now, "The EU should regulate AI on the basis of rights, not risks.", (February 17, 2021), online: https://www.accessnow.org/eu-regulation-ai-risk-based-approach/.

65 Teresa Scassa, "AI, Human Rights and Canada's Proposed AI and Data Act" (March 2024), online: https://www.teresascassa.ca/index.php?option=com_k2&view=itemlist&task=tag&tag=AIDA&Itemid=80https://www.teresascassa.ca/index.php?option=com_k2&view=itemlist&task=tag&tag=AIDA&Itemid=80.

66 Canadian Human Rights Commission Submission to the UN OHCHR Application of digital technologies in the administration of justice (March 2024).

67 See, for example, *EDSTA* sections 1(1) and 5(1). The LCO's Bill 194 Submission discusses this issue in detail.

68 LCO Bill 194 Submission at 8-10.

69 See generally, the LCO's AI in Criminal Justice System project at https://www.lco-cdo.org/en/our-current-projects/crimai/.

70 *Accountable AI, Regulating AI*, and LCO Bill 194 Submission.

71 See generally the literature on AI impact assessments cited in footnote 40 above.

72 Canada ADM Directive, s.6.1.4.

73 EU *AI Act*, Article 49.

74 *Working for Workers Four Act*, 2024, S.O. 2024, c.3, online: https://www.ontario.ca/laws/statute/s24003. Note, there is significant criticism over the limited and vague amount of disclosure necessary. See Teresa Scassa "Ontario proposes to introduce provision regarding use of AI in hiring.", online: https://www.teresascassa.ca/index.php?option=com_k2&view=item&id=378:ontario-proposes-to-introduce-provision-regarding-use-of-ai-in-hiring&Itemid=80.

75 See generally *Accountable AI.*

76 See the LCO's discussion of the Government of Ontario's Data Catalogue in LCO Bill 194 Submission at 10-11

77 Under the EU *AI Act*, there are different levels of disclosure obligations depending on the level of risk. See generally, EU *AI Act*, Articles 49 and 50.

78 The benefits and risks of AI in the criminal justice system are discussed in the LCO's forthcoming AI in the Criminal Justice System Issue Papers. The second paper in this series, *Use of AI by Law Enforcement*, written by LCO Counsel Ryan Fritsch, discusses police AI systems and their impact on privacy, freedom of speech, and freedom of association in detail. The LCO will be releasing these papers in March 2025 and they will be available at https://www.lco-cdo.org/en/our-current-projects/crimai/.

79 For example, the Algorithm Watch Impact Assessment tool has "bias" under the umbrella of "justice and fairness" (p. 10) and includes two questions that directly address bias. The Oxford University capAI assessment addresses bias and discrimination tangentially through a few questions. The HRAIA is 99 pages and addresses inaccuracy, ineffectiveness and fundamental rights. These AI assessment tools tend to address a broad array of potential issues in AI systems. The Federal Directive AIA currently has bias issues under the section "data quality" and includes two questions. As of February 10, 2025, Treasury Board Secretariat is undergoing the fourth review of the Federal Directive and Algorithmic Impact Assessment. A key focus of the review is to address issues of bias and discrimination. The suggested changes address bias and discrimination in a more in-depth and meaningful way. See Overview 4th review of the Directive on Automated Decision-Making: https://wiki.gccollab.ca/File:Overview_4th_review_of_the_Directive_on_Automated_Decision-Making.pptx; and text changes to the Directive: https://wiki.gccollab.ca/File:Text_changes_to_the_Directive.pdf Text changes to the Directive.pdf- wiki; and text changes to the AIA tool: https://wiki.gccollab.ca/File:Text_changes_to_the_AIA_tool.pdf.

80 AI Informed Decision-making Systems in Banking.

81 AI Informed Decision-making Systems in Banking at 18, Question Q1.2.

82 Human rights enforcement in Ontario is an *ex post* assessment of what transpired and whether it was legal discrimination. The purpose of the HRIA tool is different. Rather than assessing human rights issues after AI systems have been implemented and deployed, the HRIA is intentionally designed to be an *ex ante* human rights assessment process to help developers and organizations assess and address AI human rights concerns during the AI development process. Accordingly, the HRIA tool should not, and cannot, be understood to protect or absolve developers and organizations of human rights liability in the future.

83 There are many resources to track the development of AI legislation, regulations, and policies. Two notable examples are the "AI Watch: Global Regulatory Tracker" developed by the global law firm White and Case, available at https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker#introduction and the IAPP "Global AI Governance Law and Policy: Jurisdiction Overviews" (July 2024), online: https://iapp.org/resources/article/global-ai-governance-

jurisdiction-overviews/. For an interesting summary of the breadth and volume of AI legislation in the United States, see National Conference of State Legislatures, "Artificial Intelligence 2024 Legislation" (Sept 2024), online: https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation and related materials.

84 *AIDA*, s.8 states "A person who is responsible for a high-impact system must, in accordance with the regulations, establish measures to identify, assess and mitigate the risks of harm or biased output that could result from the use of the system."

85 Ontario AI Directive, s. 2.

86 Ontario AI Directive, s. 6.1 states that parties must explain the use of the AI system, identify and assess risks of the AI system, plan and implement methods to control those risks, and report and monitor the efforts to assess and control the risks.

87 Ontario AI Directive, sections 6.2 and 6.3 state that if the public interacts directly with an AI system (such as a chat bot), or if an AI system is involved in decision-making directly affecting a member of the public, ministries/agencies must publicly disclose the AI use as part of the process, service or program. Ministries and agencies must also provide an accessible avenue for the public to seek information about the use of AI in a process, service or program.

LAW COMMISSION OF ONTARIO
COMMISSION DU DROIT DE L'ONTARIO

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