



Treasury Board of Canada
Secrétariat

Secrétariat du Conseil du Trésor
du Canada

Canada

4th review of the Directive on Automated Decision-Making

Overview and proposed modifications

Fall 2024

Purpose

- Provide an overview of the 4th review of the Directive on Automated Decision-Making
- Seek feedback on policy recommendations and proposed modifications
- Discuss proposed updates

Background

- Providing better programs and services for Canadians sometimes involves the use of artificial intelligence (AI) to support sound decision making
- The Government of Canada (GC) is committed to ensuring that the government's use of AI is governed with clear values, ethics and laws, and in accordance with human rights
- Treasury Board of Canada Secretariat (TBS) is responsible for providing government-wide direction on information and data governance, information technology, security, privacy and access to information.
- The Directive on Automated Decision-Making (directive) sets rules for how federal departments can use automated systems (including AI) to make or support decisions that impact the legal rights, privileges or interests of clients. For example:
 - determining eligibility for permits and benefits
 - assessing eligibility for entry to Canada
 - deciding to hire an individual to work in the public service
 - granting market authorization for a pharmaceutical product
 - launching an investigation into an individual's conduct

Directive on Automated Decision-Making

The directive seeks to ensure transparency, accountability and procedural fairness in the use of automated decision systems in the federal government.

It requires departments to:

- assess the impacts of automated decision systems
- be transparent
- ensure quality
- provide recourse on decisions
- report publicly on system effectiveness and efficiency

The directive came into effect in April 2019 and applies to systems developed or procured after April 2020.

Overview of the Directive on Automated Decision-Making

Directive on Automated Decision-Making Requirements

Algorithmic Impact Assessment

Understand

- AIA before production
- AIA when scope or functionality changes
- Publication of AIA results

Transparency

Communicate

- Notice before decisions
- Explanation after decisions
- Access to components
- Release of source code
- Documentation of decisions
- Publication of results in meeting program objectives

Quality assurance

Prevent

- Testing and monitoring of outcomes
- Data quality
- Data governance
- Peer review
- GBA Plus
- Employee training
- Continuity management
- Security
- Consultation with legal services
- Human intervention

Recourse

Correct

- Recourse options to challenge decisions

Algorithmic impact assessment (AIA)

- Mandatory risk assessment tool
- Determines the impact level of an automated decision-system
- 51 risk and 34 mitigation questions
- Scores based on many factors including systems design, algorithm, decision type, impact and data
- Developed based on best practices in consultation with both internal and external stakeholders
- Developed in the open, and available to the public for sharing and re-use under an open license

The screenshot shows the 'Algorithmic Impact Assessment' tool interface. At the top, it features the Government of Canada logo and the text 'Government of Canada' and 'Gouvernement du Canada'. Below this is a dark header with the title 'Algorithmic Impact Assessment'. A breadcrumb trail shows 'Home > Open Government'. The main heading is 'Algorithmic Impact Assessment'. A light blue informational box states: 'Information in the AIA is only stored locally on your computer, and the Government of Canada does not have access to the information you place into the tool. If you wish to keep your work, please save the data locally for future use by using the 'Save' button. You can reload a previously saved AIA form using the 'Upload JSON File' button.' Below this are three buttons: 'Save' (green), 'Upload JSON File' (grey), and 'Start Again' (grey). A section titled 'Navigate to a Specific Page (Out of 13)' contains a dropdown menu set to 'Section 8: Impact Assessment' and a progress bar showing 'Page 8 of 13'. The 'Impact Assessment' section contains three questions with radio button options:

- Will the system only be used to assist a decision-maker?
 Yes
 No
- Will the system be replacing a decision that would otherwise be made by a human?
 Yes
 No
- Will the system be replacing human decisions that require judgement or discretion?
 Yes
 No

The final question is: 'Is the system used by a different part of the organization than the ones who developed it?' with an Yes option.

Reviewing the Directive

The directive must be reviewed every 2 years to keep the instrument relevant and reflect the evolving technology and regulatory landscape.

1 st review (2020-21)	2 nd review (2021-22)	3 rd review (2022-23)	4 th review (2024-25)
<ul style="list-style-type: none"> ✓ Strengthen transparency and quality assurance ✓ Update references to policy instruments ✓ Clarify requirements 	<ul style="list-style-type: none"> ✓ Author guidelines supporting interpretation of requirements 	<ul style="list-style-type: none"> ✓ Expand scope ✓ Strengthen transparency and quality assurance ✓ Enable inclusive approaches ✓ Improve coherence with other policies ✓ Assess reasons for automation ✓ Assess impacts on persons with disabilities ✓ Clarify requirements 	<p>Underway summer 2024 to 2025</p> <ul style="list-style-type: none"> <input type="checkbox"/> Support effective implementation <input type="checkbox"/> Strengthen client protections <input type="checkbox"/> Enhance assessment of impacts

4th review intended outcomes

Implementing the proposed updates would:

- Improve protections to clients and federal institutions
- Strengthen compliance with the directive
- Reinforce commitments to transparency and accountability
- Clarify and improve understanding of requirements and AIA questions to align with intent and interpretation in practice
- Reduce redundancies and misalignment across the TBS policy suite

Overview of key themes and issues

Three themes with 7 topics have been identified to address in the 4th review. Each theme has a goal and recommendations to achieve it.

Support effective implementation	Strengthen client protections	Enhance assessment of impacts
<ul style="list-style-type: none"> • Monitor policy implementation • Reduce the number of organizations excluded from directive • Adopt internationally recognized definition of AI 	<ul style="list-style-type: none"> • Clarify obligations and enhance impact assessment of human rights • Strengthen protections and assessment of impacts for persons with disabilities • Identify banned uses 	<ul style="list-style-type: none"> • Clarify and enhance the AIA

In addition to the changes across 3 themes, additional changes are proposed to improve clarity, reduce redundancies and align with other policy instruments. All proposed changes to the Directive and AIA are available on our [GCwiki page](#)

Support effective implementation

	Monitoring policy implementation	Excluded organizations	Definition of AI
Goal	Increase and verify departmental compliance with the directive to support improved outcomes for clients, federal institutions and Canadian society.	Increase the number of organizations that are subject to the directive to expand protections and reduce risks to clients, federal institutions, and Canadian society.	Align the GC definition of AI with a more recent and internationally-recognized definition, increasing compatibility and facilitating understanding.
Recommendations	<ul style="list-style-type: none"> • Add a new reporting requirement for departments to submit a report signed by the responsible assistant deputy minister (ADM) to TBS confirming compliance with the Directive. • Add clarity to the role of the CIO of Canada related to compliance monitoring. • Add a responsibility for TBS to publish an annual summary of the compliance reports on the Open Government Portal. • Add a requirement for the responsible ADM to approve the completed AIA prior to its publication. 	<ul style="list-style-type: none"> • Remove the Excluded organizations subsection 9.1.1, such that the directive would apply to Agents of Parliament 	<ul style="list-style-type: none"> • Remove the definition of AI from Appendix A of the directive • Modify the definition of AI in Appendix A of the Policy on Service and Digital to align with the OECD definition

Strengthen client protections

	Human rights	Persons with disabilities	Bans
Goal	Clarify obligations and enhance impact assessment of human rights	Strengthen protections and assessment of impacts for persons with disabilities, raising awareness of the impacts of automated decision systems	Identify explicit limits or define circumstances in which automated systems pose an unacceptable risk
Recommendations	<ul style="list-style-type: none"> • Modify the testing and monitoring requirements in the directive to more clearly reference human rights • Add targeted questions to the AIA to strengthen the consideration of impacts to a broader range of people 	<ul style="list-style-type: none"> • Add a requirement to document system failures and take corrective actions • Add new AIA questions to capture compliance with existing accessibility standards and broaden the consideration of impacts 	<ul style="list-style-type: none"> • Add a requirement in the Policy on Service and Digital to provide parameters for use • Add a requirement and an appendix to the Directive on Service and Digital and prepare supporting guidance to state uses of AI considered unacceptable by the GC

6. Examples of unacceptable AI uses

Unacceptable use

Example of unacceptable use case

To manipulate or deceive in a way that alters behaviour resulting in harm or impacts to individual autonomy and fundamental freedoms

- Using deepfake technology that can spread misinformation, manipulate public opinion, and lead to increased polarization.

To score or classify people in a way that leads to unjustified censorship or surveillance or that impacts freedom of expression, privacy, and autonomy

- Collecting social behaviour and inferred, predicted or observed personal characteristics over time to create a score that impacts individuals' and groups' ability to function in society.

Biometric categorization to infer personal information about individuals such as their race, political affiliation, religion, sexual orientation, gender identify and expression, and disability status

- Using physical characteristics such as hair and eye colour to infer a person's ethnicity.

Untargeted facial recognition scraping with the intent to create or expand databases

- Taking publicly available images from the internet for identity verification

Real-time and remote biometric identification systems used in public spaces, with limited exceptions when the risk of harm is outweighed by the benefit

- Using real-time facial recognition to identify individuals that are subject to a regulatory fine.

Emotion recognition

- Relying on AI to infer emotions with the goal of determining risk, capabilities or skills or to make a decision without human oversight.

Determining the risk of a person or group committing an offence based solely on AI profiling or AI assessment of personality traits

- Relying on the use of predictive policing techniques to forecast criminal activity.

Enhance assessment of impacts

	AIA modifications
Goal	Increase clarity and thoroughness of the AIA tool
Recommendations	<ul style="list-style-type: none">• Add questions in AIA sections where gaps exist or that support other areas of the 4th review• Modify and add questions to respond to feedback and clarify intent• Editorial changes

All changes to the AIA can be found [here](#)

Tell us what you think

- We are seeking your input on the recommendations and directive edits as part of the 4th review of the directive.
- Please respond to the survey where you will be asked to:
 - confirm if the recommendations help to achieve the topic goals
 - share specific edits to the updated text of the directive and AIA
 - identify any concerns or gaps
 - provide input on the approach to the "bans" topic
- We recommend that you refer to the following documents as you complete the survey:
 - Overview of the 4th review of the Directive on Automated Decision-Making
 - Text changes to the Directive
 - Text changes to the AIA tool
- The survey will be open from November 19, 2024 to January 8, 2025

Next steps

- Review the 4th review materials and respond to the survey by January 8

Timeline

Fall 2024	Winter 2025	Spring 2025	Ongoing
Consultation			
Analysis & Incorporate Feedback			
	Governance & Approvals		
		Publish	
			Support Departments

Questions for discussion

- Do the proposed updates appropriately address the goals for the 7 identified topics? Why or why not?
- Are there any proposed changes to the Directive or AIA that are concerning?
- Is there a theme or topic that is missing?

Questions?

Please reach out to the TBS
Responsible Data and AI team

ai-ia@tbs-sct.gc.ca

