April 2026

Developing an Artificial Intelligence Framework for the Health and Wellbeing of Older Canadians

NATIONAL INSTITUTE ON AGEING

Table of Contents

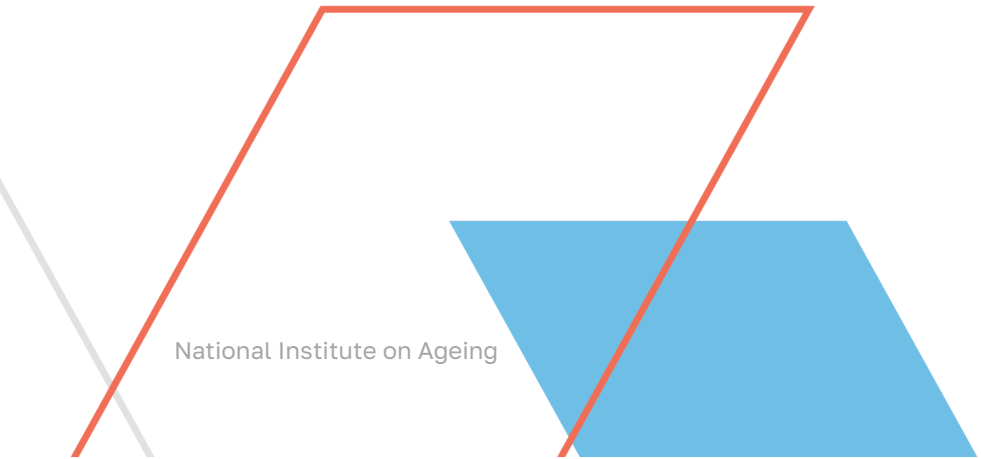

Introduction

The widespread and growing use of Artificial Intelligence (AI) is prompting governments to consider how AI-enabled technologies and applications should be regulated. AI-enabled technologies can potentially deliver meaningful benefits, but they also pose real risks that are sometimes difficult to detect.[1]

AI-enabled technologies are already being widely adopted in Canada in ways that disproportionately affect older adults. For example, hospitals are using AI-enabled technologies to support clinical decision-making by allowing physicians and other health care providers to input patient information to help make diagnoses;[2] AI-enabled sensors are being deployed to automatically detect falls, wandering or changes in the daily activities of older adults living in long-term care homes;[3] and in-home AI-enabled technologies are being developed to monitor older adults in their living spaces and provide early detection of emerging health concerns.[4]

These uses can directly support older adults by monitoring their health conditions and providing early warning when additional care may be needed. AI can also help health care providers interpret complex data (e.g., diagnostic medical imaging and lab results), generate clinical notes and support diagnosis and treatment decisions. This can in turn provide health care providers with more time for direct patient and client care.[5]

At the same time, AI-enabled technologies can hallucinate to provide incorrect, misleading or fabricated responses or incorporate and amplify systemic biases and inequities. AI used in health care trained on historically biased data may also entrench discriminatory treatment on grounds such as race, sex and sexual orientation.[7] This is not unique to the health care context. AI used in employment faces similar challenges in embedding ageist and discriminatory assumptions that can lead to barriers to employment and wage gaps.

AI-enabled technologies are increasingly being used to make significant decisions affecting Canadians, including in health care and other essential services, prompting an urgent need to develop a regulatory AI framework that is ethical, safe and accountable.[8] Where the consequences of AI-enabled technologies remain uncertain, safeguards should be implemented at an early stage to prevent harm before it becomes embedded, particularly for older Canadians and other groups that may already face discrimination.

This brief explores emerging best practices globally and puts forth considerations for regulating AI in Canada to ensure that AI benefits all Canadians, while minimizing potential harms for older adults in particular.

Recent AI Legislative Measures

In 2022, Parliament introduced the *Artificial Intelligence and Data Act* (AIDA) as part of Bill C-27. The AIDA was intended to support AI innovation while promoting its responsible development and use. Despite the AIDA's intentions, the legislative process drew significant criticism for being exclusionary and insufficiently democratic. Critics also pointed to substantive weaknesses in the bill including limited enforcement mechanisms, exemptions for public-sector uses of AI and the deferral of key definitions and oversight details to future regulations. Bill C-27 ultimately died on the order paper when Canada's parliament was prorogued in early 2025.

Despite the AIDA's shortcomings, a robust regime for AI regulation is vital. In response to the growing role of AI and its societal impact, Canada appointed its first Minister of Artificial Intelligence & Digital Innovation in 2025 to better address the issue of AI regulation. The government has indicated its intentions to take a renewed approach to regulating AI, potentially one that prioritizes legislative amendments across multiple laws over a single European-style AI Act.

As of 2026, the EU AI Act is considered to have the most robust and comprehensive regulatory framework for AI in the world. The Act, which became law in mid-2024, classifies and regulates AI-enabled technologies based on the risk they pose, imposes stringent obligations on high-risk applications, including those in health care, and prohibits the development and use of harmful AI-enabled technologies. The EU AI Act mandates transparency, human oversight and post-market surveillance, positioning the EU at the forefront of global AI governance.

The urgency to establish a principled and robust AI regulatory regime that supports older adults in Canada arises from similar concerns to those affecting populations who are vulnerable to societal harms more broadly. Although AI-enabled technologies hold considerable potential to enhance the lives of older adults, the potential vulnerabilities of this group can leave them open to abuse from the deployment and use of inadequately regulated AI-enabled technologies. This underscores the need for proactive governance to mitigate risks before they become entrenched, rather than waiting for unequivocal evidence of harm.

Potential AI Harms for Older Adults

If left unchecked, AI-enabled technologies may continue to perpetuate societal biases, undermining the quality of health care and social supports as well as the overall wellbeing of older adults. Evidence shows that older adults are routinely under-represented in the datasets used to build AI technologies; AI models trained on these datasets poorly reflect older populations in their analysis and outputs. A review of 92 facial analysis datasets found only five that explicitly included people over 65 and just one with any participants over 85, a gap that predicts poorer performance for older users.[9] The World Health Organization (WHO) also notes that datasets used to train AI models used in health care often exclude older adults, despite likely being the largest group that uses health care services.[10]

AI-enabled technologies can reinforce ageist stereotypes by portraying older adults as monolithic, dependent and resistant to technology. Research has found that some AI-enabled technologies favour the perspectives of certain demographic groups and produce outputs that diverge from the views of older adults, even when age is expressly identified in the prompt. These findings suggest that age-related bias can emerge in AI-enabled technologies that were not designed with inclusivity as a consideration.[11] This risk arises within a broader social context where ageism remains entrenched. The NIA's 2025 Ageing in Canada Survey found that 70% of older adults reported experiencing everyday ageism, while only 13% said older adults are generally respected and valued in Canadian society and 47% said older adults are sometimes overlooked or dismissed. In that context, AI-enabled technologies risk not only reflecting existing ageist attitudes but also reproducing and exacerbating them.[12]

Outside the health care sphere, AI technologies utilized for employment have been demonstrated to produce ageist exclusion from work and income opportunities. For example, in 2023, the U.S. Equal Employment Opportunity Commission filed a suit against iTutorGroup, which had programmed their application software to automatically reject female applicants aged 55 or older and male applicants aged 60 or older. iTutorGroup rejected more than 200 qualified applicants because of their age. The decree settling the suit required $365,000 to be distributed amongst applicants who were automatically rejected due to their age.[13]

These examples of ageism manifest both in the development and use of AI-enabled technologies through interconnected pathways that researchers have identified as allocation and representation harm. Allocation harm occurs when AI-enabled technologies extend or withhold opportunities, resources, or information for certain subgroups, such as older adults due to, for example, AI training data derived from younger populations leading to incorrect medical diagnoses for older aults, or when

software algorithms systematically exclude individuals over 40 from employment opportunities through age-discriminatory targeting. Continued utilization of unrepresentative data or data skewed by ageist stereotypes, prejudice or discrimination can lead to limitations in the design of tools and systems that exacerbate existing, or create new, barriers for older adults.

Representation harm occurs when AI-enabled technologies reinforce problematic stereotypes portraying older adults as homogenous, dependent and technologically incompetent. The prevalence of these views can lead older adults to internalize harmful stereotypes, which may limit their adoption and use of technologies that could improve their health and wellbeing.[14]

To ensure that AI benefits all Canadians, including older adults, regulations must identify and eliminate biases and prejudices from the design, development, use and output of these technologies. Failure to act soon risks entrenching inequity and bias in critical public systems as unregulated AI-enabled technologies rapidly advance and become incorporated in these systems. By entrenching protections, accountability and stakeholder participation, Canada can build a regime that protects rights, supports innovation and ensures that AI serves and promotes the health, dignity and wellbeing of all Canadians.[15]

Recommendations for an AI Regulatory Framework that Enables the Health and Wellbeing of Older Canadians

In developing an AI regulatory framework that protects and promotes the health and wellbeing of older Canadians, Parliament can take its cue from both international and domestic experiences. The Organisation for Economic Co-operation and Development's (OECD) recommendations on AI call on states to adopt national policies and regulatory frameworks for the responsible stewardship of "trustworthy" AI by promoting and implementing the principles of sustainable development, human-rights, transparency, explainability and accountability.[16]

The United States currently has no overarching federal AI statute. AI governance relies on sector specific regulation and voluntary guidance such as the National Institute of Standards and Technology AI Risk Management Framework,[17] along with executive orders.[18] In contrast, China's approach to AI regulation involves

application-specific regulations on recommendation algorithms, deep-synthesis and generative-AI services made available in China.[19]

A starting point for ensuring "trustworthy" AI-enabled technologies across Canada is to align any future legislative and policy frameworks with widely accepted principles such as the OECD's requirement for AI-enabled technologies to respect AI Principles. These principles include promoting shared wellbeing and the protection of human rights.

Adoption of the OECD AI Principles on their own, however, is not enough. Equally important are the mechanisms used to implement those policies. In Canada, the Treasury Board's Directive on Automated Decision-Making already requires safeguards for AI-enabled technologies used in the federal public service, including the completion of an Algorithmic Impact Assessment.[20] Similar enforceable requirements should be extended and adapted to non-federal and private sector organizations to require meaningful assessments of potentially risky AI-enabled technologies during their development and use. This is particularly important for AI-enabled technologies used in the provision of health care, long-term care, social services and financial services, contexts where older Canadians may be especially affected.

To foster public trust in AI-enabled technologies, oversight, enforcement, risk mitigation and bias prevention should be treated as foundational requirements for any AI regulatory framework. The development of these elements should be informed by the evaluation and incorporation of best practices employed by jurisdictions, with robust frameworks to:

1. Implement clear legal standards for AI-based decisions that provide protection from systemic bias and discrimination.

2. Create enforcement mechanisms administered by an arm's length body with powers to audit and penalize non-compliance.

3. Create transparency in the use of AI by establishing a Canadian "right to know" around how AI-based decisions are made.

The following recommendations for an AI regulatory framework address the use of AI-enabled technologies, with a particular emphasis on supporting older adults. As noted above, any regime should evaluate and incorporate best practices and, at a minimum, include the elements described below.

A Centralized Federal AI Oversight Body with Specialized Divisions

To regulate AI-enabled technologies, particularly technologies used in the delivery of care and support, Canada should establish a dedicated federal AI oversight body with specialized divisions, including a care-specific division, that works in parallel with Health Canada, provincial/territorial regulators and existing privacy and human frameworks. A parallel oversight body is justified because AI-enabled technologies increasingly function across jurisdictions, evolve over time and create risks that are not managed consistently through Canada's current patchwork of rules.[21] The purpose of a parallel oversight body is to ensure consistent regulation across Canadian jurisdictions. The core functions of the body would be to:

1. **Maintain sector-specific expertise in care-related AI-enabled technologies**

The body should include divisions with the technical capacity to evaluate AI-enabled technologies over time, especially for high-risk uses affecting older adults, so it retains the capability to detect capacity-based bias and discrimination.

2. **Take a multi-stakeholder approach to governance**

The oversight body should consider formal input from developers, industry, academia, civil society and affected communities such as older adults, caregivers and care providers. Participatory governance also reflects international best practice for "trustworthy AI."[22]

3. **Coordinate with provinces while setting national standards**

The oversight body should coordinate with provincial/territorial governments in a manner analogous to the European AI Office and national authorities' shared enforcement model to provide coherent and consistent implementation, supervision and enforcement across jurisdictions.[23]

4. **Establish clear federal standards for promoting safety, data validation and protection**

The oversight body should publish enforceable standards for AI used in public settings based on the risk posed. This should include minimum validation requirements, bias identification and mitigation, risk management, transparency, accountability and incident reporting.

An AI oversight body should not replace Health Canada's mandate over regulated medical devices. Instead, it should extend regulatory coverage to AI-enabled technologies that fall outside of current regulated medical-device classifications but still influence the delivery of care. Health Canada's recently published guidance on

machine learning medical devices demonstrates the level of specificity required for pre-market review in clinical settings.[24] An oversight body would apply a comparable degree of rigor to AI-enabled technologies that operate with limited or no scrutiny.

A Risk Classification System for AI Development and Use

A tiered risk classification system analogous to the one utilized by the EU's AI Act,[25] which establishes a detailed risk-based regulatory model, should guide the regulation of AI-enabled technologies in Canada. This approach ensures that potential harms are identified early and that regulatory requirements scale appropriately with the level of risk posed by an AI-enabled technology. An AI risk classification system should include the following tiers:

Low-Risk Applications

Low-risk systems include tools such as voice-to-text transcription, appointment schedulers and basic health-information chatbots that do not provide clinical advice. The regulatory approach to low-risk applications should include:

1. Registration with a regulator (e.g. Health Canada) and notification of implementation.
2. Compliance with general safety standards.
3. Monitoring of complaints and tracking incidents.
4. Enforcement measures including corrective actions and administrative penalties.

Moderate-Risk Applications

Moderate-risk applications include technologies like clinical decision-making support tools, diagnostic image analysis and drug interaction checks. The regulatory approach to moderate-risk applications should include:

1. Approval processes involving clinical studies, literature review, expert opinion and comparison to approved systems.
2. Assessment of viability through mandatory third-party certification, approval and audit.
3. Requirement for gradual deployment, competency assessment, ongoing education and mandatory human supervision.

4. Enforcement measures including conditional approvals, financial penalties, compliance orders and application of clinical practice standards and penalties.

High-Risk Applications

High-risk applications include agentic AI-enabled technologies capable of autonomous diagnosis, treatment planning algorithms and life-critical monitoring systems. For AI-enabled technologies used for critical, high-risk purposes, advance safeguards must be established, with ongoing monitoring to identify emerging risks before they can escalate. The regulatory approach to high-risk applications includes those listed for moderate-risk applications, but additionally includes:

1. A validation and approval process involving rigorous clinical trials, comprehensive assessment and quality management compliance (e.g., ISO standards).
2. Post-market surveillance and continuous oversight analogous to that of a health care provider.
3. Pre-market technical documentation requirements and transparency obligations.
4. Enhanced enforcement measures, including corporate and executive accountability and liability, increased monitoring and surveillance, mandatory technical modifications and cross-border enforcement mechanisms.

Prohibited Applications

AI-enabled technologies that pose an unacceptable risk of harm must be prohibited. Examples include:

- Harmful subliminal manipulation.
- Exploitation of personal vulnerabilities based on age, disability or condition.
- Social scoring of patients for health care access.
- Biometric categorization used for discriminatory purposes.

These applications should be subject to statutory bans on development, deployment and use, with enforcement mechanisms that include criminal penalties, fines, imprisonment and import restrictions.

Algorithmic Bias Prevention Mechanisms

AI applications used in the provision of care and support may exhibit multiple forms of bias that would systematically disadvantage certain groups. Preventing such bias requires anticipating risks during system design, so that unfair outcomes are avoided before they become widespread. The mechanisms proposed below draw on international guidance from the WHO policy briefs "Ageism in Artificial Intelligence for Health,"[26] which calls for age-inclusive data collection and age-diverse data science teams, and "Ageism in artificial intelligence for health," which identifies inclusiveness and equity in design and data collection, transparency, and accountability as core principles for AI in health.[27] They are also consistent with the EU AI Act's risk-based approach.[28] A bias prevention strategy should include the following oversight mechanisms:

1. **Inclusive data**: Ensure training datasets represent all populations, including older adults.
2. **Algorithmic bias testing**: Evaluate system performance across demographic groups to identify patterns of disadvantage, including age-related and intersectional biases (e.g., age, gender, ethnicity).
3. **Fairness in AI design and metrics**: Require techniques that explicitly account for fairness during AI system design and training. Implement standardized evaluations to ensure equitable performance across populations.
4. **Diverse development teams**: Require AI system development teams to be trained in development tailored to diverse capacity needs.
5. **AI transparency**: Require understandable justifications for AI decisions and predictions.

Design and Validation Standards for AI-enabled Technologies Based on Capacity

The idea of "capacity" is not fixed or zero-sum. Many older adults experience varying degrees of limitations as they get older, and the same individual may interact with technology differently.[29] This means AI-enabled technologies used in care and support settings should be designed and validated based on accessibility and interactivity, with ongoing monitoring after deployment. AI regulation should require:

1. **Accessible design standards that appropriately consider vision, hearing and cognitive impairments.**

Design standards should conform with recognized digital accessibility requirements for user interfaces, such as WCAG 2.2 principles,[30] and hardware used in care environments, such as CAN/ASC–EN 301 549.[31]

2. **Alternative interaction methods including simplified interfaces.**

AI-enabled technologies should be required to provide multiple means of interaction so that diminished vision, hearing, dexterity, literacy or cognition do not become barriers to access.

3. **Performance testing and validation requirements for a diverse range of potential users.**

Pre-deployment validation of an AI-enabled technology should require the meaningful input of individuals with diverse capacities, particularly those with sensory or cognitive impairments. Validation steps should also reflect the intended environment in which an AI-enabled technology will be used.

Privacy and Informed Consent

AI regulation should coordinate with, but remain distinct from, privacy law due to fundamental differences in scope, objectives and enforcement mechanisms. While privacy laws aim to protect rights to personal data, regulation of AI aims to address technical and systemic risks even when personal data is not involved, such as algorithmic bias, fairness, transparency and safety. There is, however, some overlap between privacy concerns and AI regulation, since personal data may be utilized by AI-enabled technologies. In these situations, AI regulation should align with privacy laws to allow privacy concerns to be addressed without creating confusion. To ensure AI regulation remains legally distinct but is interoperable with privacy laws, AI-enabled technologies should be subject to:

1. **Enhanced consent models**: Ensuring clear, informed consent for AI-specific uses.
2. **Data minimization**: Using only necessary data for specified purposes.
3. **Purpose limitation**: Placing restrictions on any secondary use of personal data.
4. **User control**: Ensuring rights to access, correct and delete personal information.
5. **Substitute decision-maker and care provider access**: Including carefully governed mechanisms that allow designated substitute decision-makers and,

where appropriate, care providers access to relevant AI information, supported by protocols that uphold the older adult's autonomy, confirm legal authority and prevent coercion, undue influence or misuse.

6. **International data flows**: Ensuring protections for international data transfers and harmonization of standards.

7. **Review**: Requiring regular review and updating of consent frameworks as AI capabilities evolve.

Transparency and Accountability Requirements

There are significant challenges with AI-enabled technologies that may not be explainable, i.e., are unable to provide a transparent, understandable account of how the AI-enabled technology derived its output. This can create a lack of trust in AI responses and may undermine the ability to provide informed consent. Furthermore, a lack of transparency also creates challenges in establishing causation when identifying algorithmic bias or discrimination in outputs and training data. To address these issues, AI-enabled technologies should be subject to the following requirements: [32]

1. **Public transparency reporting:** Disclosure of AI training data and prediction methodologies. Where disclosure is limited by privacy or confidentiality, regulation should still allow for regulators to securely access the data needed for independent assessment of AI-enabled technologies.

2. **Disaggregated performance reporting:** Reporting of performance and error rates across groups with differing capacities. This matters because older adults can be under-represented in datasets and validation, which increases the risk of harmful bias.

3. **Performance metrics**: Public reporting on safety and effectiveness to ensure that AI-enabled technologies work as designed and to assess whether they have any detrimental impact on individual patients, clients, users or groups.

4. **Complaint mechanisms**: Simple, accessible means for older adults, caregivers and care providers to report AI-related concerns.

5. **Public education**: Education and training programs specifically tailored to increasing AI literacy among older adults, caregivers and care providers. Topics should include how to refuse AI use where appropriate and how to challenge AI outputs or decisions.

Public Participation and Stakeholder Engagement

Stakeholder and public participation during the development, approval and use of AI-enabled technologies is foundational in addressing concerns with both democratic legitimacy and practical implementation. Meaningful public engagement ensures that certain core societal principles and expectations are deployed towards creating trust and acceptance of opaque systems. Broad stakeholder engagement provides critical input to ensure oversight remains up-to-date and creates opportunities for early warning systems that help identify risks before they become systemic problems. An effective public and stakeholder engagement mechanism should involve:

- Meaningful participation in regulatory development, along with mandatory public consultation for high-impact AI systems.
- Multidisciplinary collaboration with affected communities. Engagement should involve older adult and patient/client/caregiver advocacy groups, care providers, technology developers, policymakers, academic researchers and ethics bodies.

Conclusion

Taken together, these recommendations outline a regulatory framework built on international best practices that is principled, pragmatic and responsive. An AI regulatory framework should balance innovation with accountability by establishing clear standards, ensuring meaningful stakeholder participation, and embedding safeguards against bias and discrimination at every stage of AI development and deployment.

By ensuring older adults, caregivers, care providers and health care professionals are actively engaged alongside AI developers, policymakers and academics, Canada can build a regulatory framework that is resilient and adaptive to emerging risks. With rigorous oversight, Canada can set a standard for responsible AI governance that supports innovation, strengthens public trust and protects vulnerable populations, while safeguarding social equity, fundamental rights and the overall health and wellbeing of all Canadians.

References

1. World Health Organization, "Ethics and Governance of Artificial Intelligence for Health: WHO Guidance" (2021), at online: https://iris.who.int/server/api/core/bitstreams/e9e62c65-6045-481e-bd04-20e206bc5039/content.
2. College of Physicians and Surgeons of Ontario, "Using artificial intelligence in clinical practice" (2023), online: https://www.cpso.on.ca/physicians/policies-guidance/advice-to-the-profession/using-artificial-intelligence-in-clinical-practice.
3. CBC News, *AI monitoring system could improve safety in long-term care homes, Waterloo researchers say*. (03 April 2023), online: https://www.cbc.ca/news/canada/kitchener-waterloo/ai-monitoring-system-long-term-care-homes-university-waterloo-1.6800965.
4. University of Waterloo, "New in-home AI tool monitors health of elderly residents" (2023), online: https://uwaterloo.ca/news/media/new-home-ai-tool-monitors-health-elderly-residents; CHAH Technology, "CHAH AI Care: Supporting Aging in Place with Dignity and Safety" (2025), online (pdf): https://chah.ai/wp-content/uploads/2025/11/CHAH-One-Pager.pdf; CHAH Technology, "Comprehensive Healthcare at Home: Redefining the Canadian Healthcare System" (2025), online (pdf): CHAH [https://chah.ai/wp-content/uploads/2025/11/CHAH-Whitepaper.pdf.
5. World Health Organization, "Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models" (2024) at 3-4, online: https://iris.who.int/server/api/core/bitstreams/f780d926-4ae3-42ce-a6d6-e898a5562621/content.
6. *Ibid* at x; Yoshua Bengio et al, "International AI Safety Report 2026" (February 2026) at 72, online (pdf): https://internationalaisafetyreport.org/sites/default/files/2026-02/international-ai-safety-report-2026_1.pdf.
7. Bradley Henderson, Colleen M Flood & Teresa Scassa, "Artificial Intelligence in Canadian Healthcare: Will the Law Protect Us from Algorithmic Bias Resulting in Discrimination?" (2022) 19:2 *Canadian Journal of Law and Technology* at 475-476.
8. WHO, *supra* note 1 at v.
9. Park, J. S., Bernstein, M. S., Brewer, R. N., Kamar, E., & Morris, M. R. (2021). Understanding the representation and representativeness of age in AI data sets. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)* (pp. 834–842), pp. 834–835, 838. https://doi.org/10.1145/3461702.3462590.
10. World Health Organization, "Ageism in artificial intelligence for health: WHO policy brief" (2022) at 6, online: https://iris.who.int/server/api/core/bitstreams/8ce470c0-44f2-4350-a5bc-2606f4f59b86/content.
11. Liu, S., Maturi, T., Yi, B., Shen, S., & Mihalcea, R. (2024). The generation gap: Exploring age bias in the value systems of large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing* (pp. 19617-19634). https://doi.org/10.48550/arXiv.2404.08760.
12. Tazeen Qureshi, Gabrielle Gallant, Talia Bronstein, Mark Hazelden, Alyssa Brierley & Samir Sinha, "Perspectives on Growing Older in Canada: The 2025 NIA Ageing in Canada Survey" (2026) National Institute on Ageing, Toronto Metropolitan University at 7, 10.
13. U.S. Equal Employment Opportunity Commission, "iTutorGroup to pay $365,000 to settle EEOC discriminatory hiring suit" (2023), online: https://www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit.
14. Charlene H Chu, Rune Nyrup, Kathleen Leslie, Jiamin Shi, Andria Bianchi, Alexandra Lyn, Molly McNicholl, Shehroz Khan, Samira Rahimi, Amanda Grenier, "Digital ageism: Challenges and opportunities in artificial intelligence for older adults" (20 January 2022) 62:7 *The Gerontologist* at 948, online: https://doi.org/10.1093/geront/gnab167.
15. *Ibid*.
16. OECD, "Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449), Section 1: Background Information" (2024) OECD Legal Instruments, online: https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449.
17. National Institute of Standards and Technology, "AI Risk Management Framework (AI RMF 1.0)" (2023), online: https://www.nist.gov/itl/ai-risk-management-framework.
18. The White House, "Executive Order 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, 88 Fed Reg 75191" (1 November 2023), online:

https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.
19. White & Case LLP., “AI Watch: Global regulatory tracker – China" (2024), online: https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-china.
20. Treasury Board of Canada Secretariat, “Directive on automated decision-making" (2023) Government of Canada, online: https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592.
21. Abdullah H Ishaque, Abdi Aidid & Gelareh Zadeh, "Lessons from the Failure of Canada’s Artificial Intelligence and Data Act" (2025) 2 *NEJM AI* at 5.
22. OECD, *supra*, note 16.
23. European Commission, "Governance and enforcement of the AI Act" (14 November 2025), online: https://digital-strategy.ec.europa.eu/en/policies/ai-act-governance-and-enforcement.
24. Health Canada, “Pre-market guidance for machine learning-enabled medical devices”, Cat No H164-354/2025E-PDF (Ottawa: Health Canada, 15 February 2025), online (pdf): https://www.canada.ca/content/dam/hc-sc/documents/services/drugs-health-products/medical-devices/application-information/guidance-documents/pre-market-guidance-machine-learning-enabled-medical-devices/pre-market-guidance-machine-learning-enabled-medical-devices.pdf.
25. The EU framework organizes AI systems into minimal, limited, high, and unacceptable risk categories, with escalating obligations tied to the potential harm they pose.
26. WHO, *supra* note 1.
27. WHO, *supra* note 10.
28. European Parliament & Council of the European Union, *Regulation (EU) 2024/1689 of 13 June 2024 on Artificial Intelligence (Artificial Intelligence Act)*, [2024] OJ L 168/1, arts 8–17. Online: https://eur-lex.europa.eu/eli/reg/2024/1689.
29. World Health Organization, “World report on ageing and health” (2015) at 65, online (pdf): https://iris.who.int/server/api/core/bitstreams/5174879e-b0dc-43fc-b3a8-b1db31c51d4c/content.
30. World Wide Web Consortium (W3C), “Web Content Accessibility Guidelines (WCAG) 2.2" (12 December 2024), online: https://www.w3.org/TR/WCAG22/.
31. Accessibility Standards Canada, “CAN/ASC–EN 301 549:2024 Accessibility requirements for ICT products and services" (May 2025), online: https://accessible.canada.ca/creating-accessibility-standards/canasc-en-301-5492024-accessibility-requirements-ict-products-and-services.
32. WHO, *supra* note 1 at 26-27,47-48.