

AMS HEALTHCARE CONFERENCE · MAY 7, 2024

A Discussion Paper:

The Scaling of Responsible AI in Healthcare



TABLE OF CONTENTS

02. Executive Summary

03. Introduction

04. Barriers to the Scaling of Responsible AI in Healthcare

1. Outdated Governance and Health Data Standards that Prevent Data Sharing
2. Underinvestment in Ontario's Healthcare Digital Infrastructure
3. A Lack of Skills in the Healthcare Workforce
4. Insufficient Collaboration Among Diverse Stakeholders
5. A Lack of Transparency, Trust, Patient Autonomy and Accountability
6. Social Concerns

10. Risks of Scaling Responsible AI in Healthcare

1. Data and Algorithmic Biases
2. Data Privacy and Security Breaches
3. Algorithmics Performance Degradation

12. Practical Actions for the Scaling of Responsible AI in Healthcare

1. Develop Executive and Clinical AI Innovation Leaders
2. Foster Collaboration Among Diverse Stakeholders
3. Invest in Digital Infrastructure and Workforce Capacity
4. Promote Health Equity Through Appropriate Data Collection, Analysis, and Use
5. Improve Access to Data While Ensuring Privacy and Security
6. Improve Transparency in Order to Build Trust and Establish Accountability
7. Communicate Success and "Create a Buzz" about AI in Healthcare

19. Conclusion

20. Appendices

1. Appendix A: Terminologies
2. Appendix B: Stakeholder Interviews - What We Heard
3. Appendix C: Biases Created Using AI in Healthcare
4. Appendix D: The Artificial Intelligence and Data Act (AIDA)

29. References

EXECUTIVE SUMMARY

Artificial Intelligence (AI) holds immense transformative potential for Ontario's healthcare system. It can significantly improve medical diagnoses and treatments, enhance patient adherence to treatment plans, and boost operational efficiencies. When harnessed effectively, AI has the potential to transform the healthcare system into one that is more efficient, effective, and patient-centric.

While AI's potential in healthcare is promising, its adoption and scalability face significant barriers and risks. The most significant barriers are outdated governance and health data standards, underinvestment in digital infrastructure, a lack of AI skills in the healthcare workforce, insufficient collaboration among diverse stakeholders, a lack of transparency, trust, and accountability, and additional social concerns. The risks of scaling AI include, but are not limited to, data and algorithmic biases, data privacy and security breaches, and algorithmic performance degradation.

Taking into account these barriers and risks, this discussion paper proposes Practical Actions for the Scaling of Responsible AI in Healthcare. They include Developing AI Innovation Leaders, Fostering Stakeholder Collaboration, Investing in Digital Infrastructure, Promoting Health Equity, Improving Data Access While Safeguarding Privacy and Security, Enhancing Transparency and Accountability, and Communicating AI Success, all of which foster a sense of confidence and reassurance about the potential of responsible AI in healthcare.

This discussion paper is a key input into the AMS Healthcare Conference 2024. The conference aims to both discuss the issues associated with, and also build momentum for, the widespread scaling of compassionate, ethically sound, and equitable AI in healthcare. The adoption and scaling of responsible AI in healthcare will pave the way for a more innovative and responsive healthcare system throughout Ontario.

INTRODUCTION

Artificial Intelligence (AI) is being used in healthcare to improve the diagnosis and treatment of diseases, increase adherence to treatment plans, create operational and economic efficiencies, promote healthier lifestyle routines, accelerate research and development, improve the targeting of public health interventions, and reduce health inequalities (Hee Lee & Yoon, 2021). It is an innovative technology that has the potential to transform healthcare systems including Ontario's. However, despite the promise of this groundbreaking technology to reshape Ontario's healthcare landscape, there are significant barriers and risks to its scalability (OECD, 2024).

Combining a comprehensive literature review with the feedback from key stakeholder interviews, this discussion paper identifies the six significant barriers and the three consequential risks associated with the adoption and scaling of AI in healthcare across Ontario. Based on these findings, the paper proposes seven practical actions to enable the widespread scaling of compassionate, ethically sound, and equity-enhancing AI. The paper includes questions for discussion about the policy responses needed to execute these practical actions. It is an input into the AMS Healthcare Conference 2024, which aims to build momentum for the successful scaling of responsible AI across Ontario's healthcare system.

BARRIERS TO THE SCALING OF RESPONSIBLE AI IN HEALTHCARE

Drawing on the feedback from key stakeholder interviews and combining this with a comprehensive literature review, it is proposed that the six most significant barriers to scaling AI in healthcare across Ontario are not due to a scarcity of AI innovations. Instead, the barriers stem from not having enough quality data, difficulties in data sharing, concerns about privacy and security, outdated governance that includes unclear data policy, legislation regulation and standards, insufficient data capacity and ineffective data literacy (Ontario Health Data Council Report, 2023). The following section identifies the barriers to scaling and examines what has been done in Ontario to overcome these barriers.

1. Outdated Governance and Health Data Standards that Prevent Data Sharing

Health data governance refers to the legislation policies, procedures, processes and people in place to manage and oversee the collection, sharing, access, and use of health data. It also ensures data quality, security, privacy and retention (Ontario Health Data Council Report, 2023). Currently, outdated health data governance limits the ability to access and share health data across institutions and jurisdictions (OECD, 2024). Knowledgeable stakeholders have put forward that there is an urgent need to update governance to enable timely access to quality data that can be shared across multiple institutions without compromising privacy or security. Without updated governance, AI will lack the data to develop, implement and scale AI in healthcare (Evans et al., 2020).

According to the Ontario Health Data Council Report (2023), Ontario is “data-rich but information poor due to large amounts of dormant data within data silos” (Ontario Health Data Council Report, 2023). The Council has recommended that Ontario “establish system level trustworthy governance and policy's for health data as a public good” that would enable access to data without compromising privacy or security. Box A summarizes the recommended implementation actions to establish up to date health data governance.

Recommended Implementation Actions to Establish System-Level Trustworthy Governance and Policy for Health Data as a Public Good.

- Develop shared, responsive governance for the creation, implementation, and sustainment of a trustworthy, integrated, and accountable health data ecosystem using an iterative approach that engages all health system stakeholders.
- Foster a transparent approach to the governance of health data, including increased public reporting and transparency about the collection, use, and disclosure of health data.
- Set standard, permanent, and meaningful advisory roles for the public in health data governance structures, including but not limited to community governance tables.
- Review policy opportunities to identify where barriers could be removed and new enablers installed, including removal of legislative or bureaucratic obstacles to health data access.
- Strengthen data and digital literacy and capacity across the health system and Ontario at large by empowering, engaging, and educating the public about health

Standardization of information storage and retrieval will be essential to securely transferring data across multiple platforms (He et al., 2019). An example of how Ontario is improving the ability to access and share health data across institutions and jurisdictions is the Digital Health Information Exchange (DHIEX) policy initiative, which came into effect on January 1st, 2021. This initiative has increased interoperability and improved data movement across multiple platforms under the regulation of the Personal Health Information Protection Act (PHIPA) (Ontario Health Data Council Report, 2023).

2. Underinvestment in Ontario’s Healthcare Digital Infrastructure

Capacity for mass data storage has undergone enormous expansion both locally and at remote sites, such as the “Cloud” and improvements in computing power are enabling more efficient processing of very large amounts of data. However, observers suggest that the availability and accessibility of comprehensive datasets, required for the implementation and scaling of AI, remain elusive because Canada has underinvested, relative to other nations, in its healthcare data systems (Evans et al., 2020).

According to the AI for Health (AI4H) Task Force Report: Building a Learning Health System for Canadians, “The time is now for Canada to make strategic investments in a national AI for health strategy to leverage all of our strengths, foster collaboration and coordination across sectors in jurisdictions and deliver better health for Canadians.” This investment needs to address the key domains listed in Box B (Evans et al., 2020).

Finally, some point to the siloed nature of Ontario’s healthcare system, highlighting that financial incentives are poorly aligned to support the adoption of innovative technologies. A more integrated system is able to reconcile investment in one segment with savings in another to determine whether or not the net outcome is positive (Evans et al., 2020).

Box B: AI4H Task Force Report: Building a Learning Health System for Canadians

Key Domains of Canada’s Health Data System that Require Investment

1. **Scope:** The range of data from diverse sources that can be brought together, including structured, unstructured and simulated data.
2. **Skews:** Problems arise when data that are not representative of relevant populations or problems are used to build/train algorithms.
3. **Standards:** Although there exist robust criteria for assessing data quality and reliability they are not widely recognized or used.
4. **Sharing and Ownership:** There are many different data sharing and ownership models ranging from patient-owned to public, private, not-for-profit data trusts, or federated learning models.
5. **Trust and Privacy:** Protecting the privacy of personal data requires continuous technical scrutiny as well as greater public engagement to build trust.

3. Inadequate AI Skills in the Healthcare Workforce

The scaling of AI in healthcare requires a workforce with strong literacy in data management, statistics, computer science, data privacy, biases, and ethics (Fisher and Rosella, 2022). Healthcare professionals, for the most part, lack the necessary expertise to develop, evaluate or adopt AI-based technologies. The competencies of healthcare workers and administrators need to go beyond those of traditional informatics and current digital health (EMR) training.

According to the AI for Health (AI4H) Task Force Report: Building a Learning Health System for Canadians (Evans et al., 2020), “ More attention needs to be directed to the educational training needs of users of AI, including health professionals, administrators and the public at large, and while the demand for training and retraining is surging, the growth supply of good quality training programs to meet this demand does not appear to be keeping pace.” Training the healthcare workforce on AI-based technologies is imperative for successfully scaling AI in healthcare (Fisher and Rosella, 2022).

4. Insufficient Collaboration Among Diverse Stakeholders

Collaboration among diverse stakeholders will be critical to successfully scaling AI in healthcare (Fisher and Rosella, 2022). For example, many of the key stakeholders interviewed pointed out that it will require a diverse group of stakeholders (policy experts, scientists, industry representatives, healthcare workers, and patients) working collaboratively to create technically sound, ethically responsible, and broadly accepted data standards within the healthcare system. Without common data standards, the information used to train the AI algorithms will be unreliable (e.g. using different terms for the same diagnosis or using the same term for different diagnoses), and the outcome will be unreliable AI solutions (Evans et al., 2020).

Establishing common legal standards, such as legislation that provides protected access to data, will also require a collaborative effort among various stakeholders. Without legal standards, AI solutions will have varying quality and unclear accountability (OECD, 2024).

Additionally, many have suggested that more collaboration between the public and private sectors would help Ontario’s efforts to scale AI in healthcare. The private sector brings expertise, funding, infrastructure, experience in regulatory processes, and new perspectives and insights into healthcare challenges and solutions (Fisher and Rosella, 2022).

5. A Lack of Transparency, Trust, Patient Autonomy and Accountability

AI-generated medical decisions are characterized by some as “black box” decisions because the rationale behind the decisions is opaque and inscrutable to both healthcare providers and patients (OECD, 2024). This difficulty in understanding the AI decision-making process can grow into a lack of trust in its solutions. Transparency in AI refers to increased explainability regarding how an AI model performs. However, disclosing additional data to gain this trust makes AI innovations more vulnerable to hacker attacks. Scaling AI will require establishing a balance between providing enough transparency to build trust without enabling bad actors to hack the system (Cheng et al., 2021).

The perceived lack of transparency in AI innovations also means that patients are less able to make informed decisions, reducing their autonomy (Quinn et al., 2021). Respecting patient autonomy is a fundamental ethical principle that forms the basis of patient-centred care. Patient autonomy refers to the right of patients to make informed decisions about their medical care, including the right to consent or to refuse treatment based on their values, preferences and understanding of their medical condition. Enhancing the transparency of AI innovations is critical to upholding the value of patient autonomy.

Establishing accountability for AI systems' mistakes is challenging because it can be shared among various stakeholders involved in the AI lifecycle, including developers, data scientists, healthcare organizations, and clinicians. However, without accountability, concerns about liability may prevent healthcare institutions and providers from adopting AI solutions. Clear guidelines must delineate accountability to ensure AI innovations are used responsibly and ethically (Khan et al., 2023). Without transparency and accountability, it will be challenging to establish trust with healthcare providers and patients (Lopez, 2018). Without trust, it will be difficult to scale AI in healthcare (WEF, 2024).

6. Social Concerns

Compassion and empathy are essential for high-quality health care. Research has demonstrated that pro-social caring behaviours improve health outcomes (Morrow et al., 2023). AI innovations lack compassion and empathy, and there is concern that scaling AI in healthcare will dehumanize healthcare and detract from the quality of patient care (Morrow et al., 2023). It has also been well-documented that many people feel threatened by artificial intelligence, believing it might eliminate their jobs, negatively affecting the wider acceptance of this technology (Davenport & Kalakota, 2019). Perception of bias, actual or perceived, may also be a barrier to AI's uptake. Finally, overestimating AI capabilities may create unrealistic expectations regarding its efficacy, creating disappointment and impeding its uptake (khan et al., 2023).

Overcoming these barriers will require a collaborative effort from policymakers, healthcare professionals, technology developers, patients, and the public to change institutional practices and policies, modernize data governance, and increase stakeholder engagement, collaboration, and trust in AI (OECD, 2024).

RISKS OF SCALING RESPONSIBLE AI IN HEALTHCARE

AI has the potential to revolutionize healthcare and improve health outcomes. Still, there are risks associated with AI that must be addressed to ensure its benefits are shared safely and equitably. Some of the risks to be managed and mitigated are:

1. Data and Algorithmic Biases

Some deficiencies AI is meant to overcome, such as unwarranted and inequitable variation in medical practice, permeate the data used to train the AI algorithms. For example, race-based adjustments to equations used to calculate estimated glomerular filtration rate (eGFR) levels for determining the severity of kidney disease resulted in the under-treatment of approximately 3.3 million Black Americans who, without these corrections, would be more likely to receive earlier treatment for a range of complicating conditions (London, 2022). AI systems trained on data that included these adjustments might continue underdiagnosing certain populations. AI algorithms must not have inherent biases from historical data or flawed clinical practices (Ueda et al., 2024). For more detailed information about the broad range of biases created when using AI in healthcare, refer to Appendix C.

2. Data Privacy and Security Breaches

AI works best when it has access to large datasets. However, patient records are confidential, and institutions are reluctant to exchange health data due to privacy and security concerns (khan et al., 2023). This is a justified concern because if bad actors or incompetent users get access to that health data, even if the individually identifiable information has been removed, it is possible that there may be enough clues for an individual to be identified. Training data sets may also be targeted for cyber-attacks, rendering AI solutions unusable or alternatively changing the training data sets to purposely build bias into the algorithms, creating inequitable variation in medical practice (OECD, 2024).

3. Algorithmics Performance Degradation

The performance of AI solutions can degrade precipitously after transitioning from training to implementation (Khan et al., 2023). AI solutions may initially perform well in a specific setting, but performance can vary significantly across settings and over time. This is concerning because the harm done by a faulty algorithm has the potential to be exponentially higher than that of a single doctor-patient interaction (Topol, 2019).

Managing and mitigating these risks will require significant alterations to current clinical data gathering and data management practices (London, 2022) and robust risk management programs that are based on cooperative and collective action (OECD, 2024).

PRACTICAL ACTIONS FOR THE SCALING OF RESPONSIBLE AI IN HEALTHCARE

The scaling of responsible AI in healthcare has emerged as a complex socio-technical undertaking that faces significant barriers and risks. Successfully scaling AI in healthcare will require social, ethical, and economic considerations (OECD, 2024). Taking into account these barriers and risks and the resulting social, ethical and economic considerations, the following practical actions for the scaling of responsible AI in healthcare are put forward for discussion:

01 Develop Executive and Clinical AI Innovation Leaders

Healthcare leaders who support innovation and collaboration are widely recognized as key factors in successfully scaling AI innovations (Petersson et al., 2022).

It has been well documented that many healthcare leaders are risk averse. They often resist novel technologies, such as AI, due to concerns over patient safety, data security, and the disruptive impact on workflow systems (Petersson et al., 2022). A cautious approach to new technologies, while understandable, contributes to slow and variable uptake (Keown et al., 2014; McKinsey, 2018).

The scaling of AI in healthcare will require informed, passionate, and motivated executive and clinical leadership willing to undertake the risks, disruptions, and costs associated with embracing this new technology. They must also communicate their strategy effectively and nurture a culture of continuous improvement, innovation, and adaptability within their workforce (Bajwa et al., 2021; Keown et al., 2014).

Policy Challenge: How do we prepare executive and clinical leaders to advocate for the scaling of AI in healthcare, ensuring they can effectively influence both internal organizational culture and the broader social and political landscapes?

02 Foster Collaboration Among Diverse Stakeholders

Strong and enduring collaborations among healthcare institutions, providers, professional associations, patients, academics, ethicists, regulators, and industry are essential to scaling responsible AI in healthcare (Topol, 2019).

Collaboration among diverse stakeholders in healthcare is critical to scaling responsible AI for several reasons. First, bringing together diverse expertise ensures that the scaled AI innovations will be more responsive to the specific health requirements of community health and align with society's values (Evans et al., 2020). New policies and standards must be established to encourage and enable more sharing of health data within often siloed healthcare systems. For example, this could include updated policies that allow for the collection and sharing of upstream (socio-economic) data, which is instrumental in assessing the social determinants of health and enhancing health outcomes. The development and implementation of these new standards and policies would require a collaborative effort from a broad range of healthcare stakeholders.

Collaboration is also essential to uncovering and mitigating biases within AI innovations. For example, it is helpful to have multiple perspectives to uncover varied sources of bias, including historical biases in patient care, biases in the data collection process and biases in the interpretation of outcomes (Ueda et al., 2024). An example would be that clinicians might be best at identifying where biases affect clinical outcomes, data scientists can point out and adjust for biases in the data and algorithmic functioning, bioethics experts can address moral implications, and patients can provide insight into how these technologies affect their care.

Finally, collaborative efforts would ensure more transparency and accountability (OECD, 2024). Improved transparency and accountability help build public trust in AI innovations, which is critical to successfully scaling any healthcare innovation (OECD, 2024; Hee Lee & Yoon, 2021).

Policy Challenge: How do we ensure collaboration among key stakeholders within Ontario's siloed healthcare system?

03 Invest in Digital Infrastructure and Workforce Capacity

Scaling AI will require new policies and funding models that direct investment to build the digital platforms and fund the workforce development required to scale AI in healthcare.

Building the digital infrastructure and workforce capacity required to scale AI in healthcare will be expensive. Health systems have limited resources to create and test new innovations and lack a structured process and funding model to scale innovation (Moroz et al., 2020). Scaling AI innovations may require more flexible funding models. Like many innovations, AI may increase costs in one part of the system and realize benefits in another (MacNeil et al., 2019; Keown et al., 2014).

In addition, introducing AI into clinical settings can be disruptive and unsettling to healthcare workers. AI, like other clinical innovations, will require the deployment of the full suite of clinical change management tools, including training programs, stakeholder engagement, workflow redesign, communication strategies, leadership support, performance monitoring and feedback mechanisms. This will require an investment of time, resources and funding (Keown et al., 2014).

Policy Challenge: What funding models will align the long-term benefits of AI innovation with significant short-term investment in digital infrastructure and workforce disruption? Is there an opportunity for public/private partnerships?

04

Promote Health Equity through Appropriate Data Collection, Analysis and Use

The data used to train AI algorithms must be diverse and representative of the population. AI algorithms must be continuously monitored for performance, identifying potential biases, and updating the algorithms (OECD, 2024).

Data used to train AI algorithms must reflect the target population's demographics, characteristics, healthcare needs, and potential disparities, incorporating data from various patient populations, age groups, disease stages, cultural and socioeconomic backgrounds, and healthcare settings (Ueda et al., 2024). Any biases in the data used to develop AI will skew AI results. Furthermore, there is no guarantee that an AI algorithm with high performance will maintain its high performance in the future (Ueda et al., 2024). Many have suggested that validation studies are essential for verifying the effectiveness of AI in different patient populations and conditions.

Regulations that include independent audits and AI validation may help identify and address potential biases and ensure that AI algorithms remain fair, accurate, and effective in diverse healthcare settings. Establishing accountability for algorithm quality control to continuously monitor AI performance, identify potential biases, and update algorithms could help mitigate this risk.

Raising awareness about the potential biases created using AI in healthcare (Appendix C); sharing best practices to address them, and encouraging open discussions on the implications of AI in healthcare decision-making enables healthcare professionals to evaluate AI recommendations critically, weigh potential risks and benefits, and consider alternative sources of information when making patient care decisions. Patients aware of potential biases can engage in more informed and meaningful conversations with their healthcare providers regarding treatment options and play a more active role in their care if they know AI's limitations.

Policy Challenge: What organization(s) should be responsible for ensuring that there are processes in place to monitor and mitigate AI biases? Should this be provincial or pan-Canadian?

05

Improve Access to Data while Ensuring Privacy and Security

Up-to-date governance, privacy, security, ethical, and legal considerations will be required to improve access to data while ensuring that data is secure and privacy is protected (Evans et al., 2020).

Protecting confidential health data is at the core of all health innovations and is foundational to the ethical implementation of AI in healthcare (OECD, 2024). A nimble and responsive governance and regulatory environment will be required to support access to and sharing of data for AI innovations because developing AI algorithms requires access to vast quantities of up-to-date, deidentified health data. Policies that enable the collection and sharing of upstream (socio-economic) data will also be helpful.

Robust security measures, such as encryption and anonymization techniques, will be required to improve access while protecting patient data from unauthorized access, data breaches, and other cybersecurity threats. Strict access controls and audit mechanisms must be implemented to monitor and track data use, ensure accountability, and prevent data misuse.

Canada's Artificial Intelligence Data Act (AIDA) is the federal government's first attempt to regulate AI. Other governments worldwide have regulated AI, including the European Union's 2021 EU AI and the United States' 2022 Act. See Appendix D for more information about these regulations.

Policy Challenge: How do we assure patients and the public that AI applications serve their needs without threatening their rights? How and who will ensure that provincial legislation adapts to enable the scaling of AI in healthcare? Should there be an AI Health Centre of Excellence that would oversee governance and continuous monitoring?

Improve Transparency in order to Build Trust and Establish Accountability

Transparency and explainability are essential elements of responsible AI. They enable healthcare professionals and patients to understand the basis of AI-generated predictions, foster trust, and establish accountability (World Economic Forum, 2024).

The black box concept in AI refers to algorithms where the decision-making algorithm is not transparent to the physician or patient. In healthcare, this concept is particularly concerning because decisions related to patient diagnosis, treatment and prognosis are made without a clear understanding of how AI has arrived at these recommendations. This creates issues related to patient trust and clinical accountability.

Improving transparency and explainability will lay the foundations for building trust between AI innovations and users, including healthcare providers, patients, and the public (Lopez, 2018). It does this by demystifying the technology for both healthcare providers and patients, helping to ensure that AI systems make decisions based on unbiased information, and enabling external bodies to audit and oversee the technologies, ensuring that they are used safely and appropriately. Improving transparency also allows for more patient autonomy.

It has been proposed that clear guidelines will be helpful in addressing responsibility and accountability for errors, harmful outcomes, and biases in AI-generated predictions (Lopez, 2018). There are conflicting opinions about who should be accountable for AI-related medical failures. The public believes physicians should be liable. However, physicians believe that it is the healthcare organizations that should be liable if their healthcare organization develops the AI system (Lopez, 2018). Currently, there is limited jurisprudence on accountability and liability for using AI in healthcare (Lopez, 2018).

Policy Challenge: How do we assure patients and the public that AI applications serve their needs without threatening their rights? How and who will ensure that provincial legislation adapts to enable the scaling of AI in healthcare? Should there be an AI in the Health Centre of Excellence that would oversee the governance and continuous monitoring of AI innovations in healthcare?

07

Communicate Success and “Create a Buzz” about AI in Healthcare

The future of AI in healthcare requires a buy-in from healthcare providers and the public. This will require improved awareness, transparency, and literacy regarding the benefits and risks associated with using AI in healthcare.

“Many references to AI in popular culture have presented AI as villainous and to be feared. Several attempts to leverage AI have had embarrassing results, such as IBM’s Watson being ineffective in managing cancer patients, skin lesion scans not being effective for people with darker skin, and recent “hallucinations” produced by generative AI systems.” OECD, 2024

We heard that communicating the successes of AI innovations in healthcare is critical for its successful scaling. Communication strategies will help users overcome perceived risks and build trust in AI innovations (Sebastian et al., 2023). Clear and effective communication should not only showcase the benefits, such as improved diagnostic accuracy, enhanced treatment personalization, and increased efficiency in healthcare delivery, but it should also help to demystify the technology, highlighting its value in real-world settings and addressing professional concerns, most notably privacy concerns (Sebastian et al., 2023). The narrative must be communicated strategically. For example, it must emphasize not only its effectiveness and safety but also its role in augmenting, not replacing, human expertise and improving patient outcomes.

Policy Challenge: How do we change the narrative about the purpose of AI to deliver the message that AI technologies complement and augment the expertise of providers and do not replace them?

CONCLUSION

The ambitious goal of scaling responsible AI throughout Ontario's healthcare system will be foundational to the transformation of the system. AI is already proving its value across sectors with significant risk factors like aeronautics, sectors requiring high security such as finance, and sectors with large-scale production like consumer goods (OECD, 2024). To parallel the advancements realized in these fields, leaders in healthcare must engage in deliberate and strategic actions, such as the seven actions presented in this paper. The time is now to implement policies that enable the scaling of responsible AI in healthcare and improve the quality and efficiency of our healthcare system,



“While it will take time, leadership, will, effort, and investment to achieve and sustain benefits from AI in health, **urgent action is necessary**. Time is running out for policymakers to stay ahead of the curve and take control of the evolution of AI in health systems before technology dictates its own future.”

Collective Action For Responsible AI In Health, OECD, 2024.

APPENDIX A: TERMINOLOGIES

Artificial intelligence (AI) are computer systems that learn, reason, adapt, and self-correct. AI is a group of technologies rather than a single technology.

Machine learning (ML) is a subset of AI techniques that enables computers to learn and improve their performance through experiences (data) rather than programming by a human. The more data you provide to the ML algorithm, the better it gets. In healthcare, the most common application of ML is predicting what treatment protocols are most likely to succeed on a patient based on various patient attributes and treatment context. This is referred to as precision medicine.

Deep Learning (DL) is a subset of ML that involves neural networks with multiple layers/variables. The common application of DL in healthcare is Medical Imaging Analysis. DL improves the accuracy and speed of diagnosis, supporting healthcare professionals in making more informed decisions.

Natural Language Processing (NLP) is a subset of AI that focuses on enabling computers to understand, interpret, and generate human language. The most common use of NLP in healthcare is the creation, understanding, and classification of clinical literature and published research. NLP systems can also analyze clinical notes, prepare reports, record patient interactions, and conduct AI dialogue.

Large Language Models (LLM) are designed to perform a wide range of NLP tasks. They learn intricate patterns and contextual nuances from diverse linguistic data during training. They play a pivotal role in advancing the capabilities of NLP systems, providing solutions for various language-related tasks. Their ability to capture complex language patterns and contextual relationships has contributed significantly to the progress of NLP.

Robots, machines programmed to perform specific tasks, can be equipped with AI to make autonomous decisions. AI can make robots more intelligent, adaptive, and capable of performing complex tasks. For example, it can improve a surgeon's ability to see, create precise and minimally invasive incisions and stitch wounds.

Robotic Process Automation (RPA), which does not involve robots but computer programs on servers, is used for repetitive tasks like updating patient records. It can be combined with other technologies, such as image recognition.

APPENDIX B: STAKEHOLDER INTERVIEWS - WHAT WE HEARD

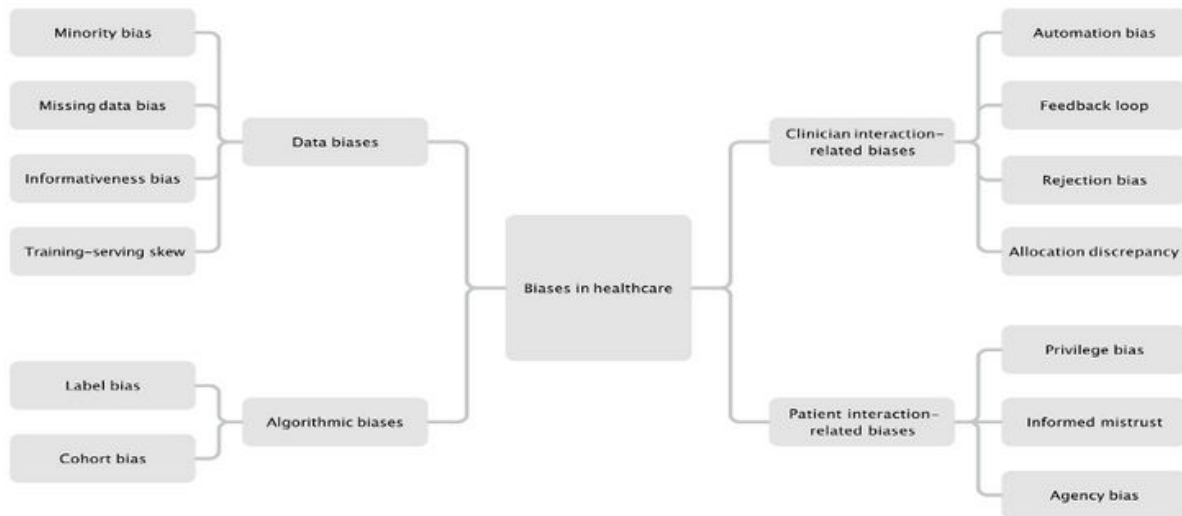
We interviewed key stakeholders to uncover their perspectives on the scaling of AI technologies in Ontario's healthcare system.

Stakeholder	Interview Comments: Scaling AI for Healthcare in Ontario
Research Expert	<ul style="list-style-type: none"> • The tech is there. Why don't we take a few easy cases, deploy them, and learn from them? • The government is not being innovative or brave enough to say, "Why don't we have a real-time data system that is managed by the province, not by individual hospitals?" because this is creating more silos. • Before regulating, we need to do a couple of large-scale pilots and learn about what you're going to regulate because if you regulate now, you're going in blind with only hypotheticals.
Research Institute	<ul style="list-style-type: none"> • Fostering public private partnerships is allowing some scalability to happen on a very small level. • We need to be more strategic about how we drive this forward, and this requires coordination amongst various partners. • Researched institutions create innovative technologies; the hospitals work to test and evaluate, get it up and running, and measure the benefit to their patients and their providers. Then the companies come in and play a part in the scalability. • We're putting safe things into the system, but what is more concerning is what happens once they are up and running. Do we have the investments we need in cyber security to make sure the systems are not vulnerable? We have created tools that are built on data and data is an international currency. Right now, more than money.
Bioethics Expert	<ul style="list-style-type: none"> • One of the challenges faced by larger-scale pilots is dealing with multiple research ethics approvals and the lack of institutional policies. • We need to get away from system-level competition and have more collaboration in AI innovation. • There is a need for transparency in order to build trust. Ironically, the public does not trust AI companies; however, they do trust Google and Facebook.

Stakeholder	Interview Comments: Scaling AI for Healthcare in Ontario
Hospital Executive	<ul style="list-style-type: none"> • We need strong, robust data platforms, frameworks, and people who understand data. • There needs to be centralized oversight because AI has its risks. • AI initiatives need to be around helping the population, not the provider. However, healthcare providers need to like it (AI innovations). AI innovation requires a huge investment, and the ability to roll it out depends on medical and clinical leadership support. At the end of the day, it's useless if they're not going to use it.
Government	<ul style="list-style-type: none"> • If AI is used in the back-office, it may not need regulation, but where it involves decisions about outcomes and allocations of resources, where it is more deterministic, it should be regulated. • The government is a natural body to do this, bringing the players together, putting the framework in place, and then setting expectations. • We are not moving fast enough to enable innovators to do what they want. They can develop the algorithm quickly but deploying it into the system and commercializing it is a challenge. • Governance and standards will do more to advance AI than a lot of other things because this frees people up from legal conversations that slow everything down.
Mental Health	<ul style="list-style-type: none"> • The data conversation is massive. There is a lot of data being replicated within silos. There are various players. The question is a very Canadian question is this a provincial or federal responsibility? • The mental health space is fraught with apps. It is unclear what level of evidence we are OK with allowing it to move along the spectrum of scale up? • There is a large cohort of executives in leadership roles that need some boot camps at a level that is appropriate for leadership executive. This will be very different to the training required for an entry to practice/ direct care person will need
Home Care	<ul style="list-style-type: none"> • Home care has no digital data. The systems are still manual. • Hospitals can show their wait time but home care cannot. • It is impossible, in real-time, to project service gaps; therefore, it is difficult to direct our very limited resources to the areas that are in the most need.

Stakeholder	Interview Comments: Scaling AI for Healthcare in Ontario
Patients	<ul style="list-style-type: none"> • I understand the tech. I think it's great, but I am not willing to give my data in return for an outcome. • First define the problem. It is not about researchers' ideas. It's about aligning with the values of health care in Ontario. • Transparency, education and engagement with patients and caregivers, removing the fear of the unknown is key to its success.
Healthcare Professional Association	<ul style="list-style-type: none"> • If a regulatory body asks where are the standards, I would ask the question where is the education for the clinicians that must meet these standards. • Professional associations approve academic programs and need to be an integral part of setting standards. • Health care providers can be on the lookout for biases and provide feedback

APPENDIX C: BIASES CREATED USING AI IN HEALTHCARE



Ueda et al., 2024

1. Data Biases

- **Minority Bias:** The number of minority group members in the dataset is insufficient for AI to learn accurate statistical patterns. For example, many cardiovascular risk prediction algorithms have a history of being trained primarily on male patient data. This has led to an inaccurate risk assessment in female patients with different symptoms and risk factors.
- **Missing Data Bias:** Data from groups are missing nonrandomly, making it difficult for AI to generate accurate predictions. For example, if patients in contact isolation have fewer vital sign records than other patients, the algorithm may struggle to identify clinical deterioration.
- **Informativeness Bias:** The features used for detection are not as apparent for certain groups, lowering their informativeness when predictions are made. For example, identifying melanoma from images of patients with dark skin is more challenging than those with light skin.
- **Training-Serving Skew:** There is a mismatch between the data used for AI training and those used during deployment. This can arise from non-representative training data due to selection bias or from the deployment of the model on patients with a population prevalence different from that of the training data. In a study training AI to diagnose pneumonia from chest X-rays, the performance of unseen data from the institution where the training data were collected was significantly higher than the performance of data collected from external hospitals.

2. Algorithmic Biases: Label and Cohort Biases

- **Label Bias:** This occurs when AI training uses inconsistent labels, which may be influenced by healthcare disparities rather than universally accepted truths, leading to biased decision-making based on inaccurate or inconsistent information in the AI algorithms. For example, significant racial bias has been observed in commercially available algorithms that predict patients' healthcare needs. The major contributing factor to this algorithm's bias was its design, which used cost as a proxy for healthcare needs, leading to an underestimation of the needs of Black patients compared with White patients with similar conditions.
- **Cohort Bias:** This occurs when AI is developed based on traditional or easily measurable groups without considering other potentially protected groups or varying levels of granularity. For example, mental health disorders have been underdiagnosed or misdiagnosed within lesbian, gay, bisexual, transgender, queer or questioning, intersex asexual, and other (LGBTQ +) populations. Algorithms often do not take the granularity of the LGBTQ + population into account and rely only on information about biological males and females. AI trained on such data may continue to overlook or misdiagnose mental health issues in these populations, potentially perpetuating existing disparities in mental healthcare.

3. Clinician Interaction-Related Biases: Automation Bias, Feedback Loop, Rejection Bias, Allocation Discrepancy

- **Automation Bias:** The tendency to overly rely on AI when tasks are transferred from healthcare professionals to AI programs. Overconfidence in algorithms can result in inappropriate actions based on inaccurate predictions.
- **Feedback Loop:** This occurs when clinicians accept AI recommendations even if they are incorrect, leading the algorithm to relearn and perpetuate the same mistakes.
- **Rejection Bias:** The conscious or unconscious desensitization to excessive alerts. Alert fatigue is a manifestation of this bias, as clinicians may ignore important alerts owing to an overwhelming number of false alarms.
- **Allocation Discrepancy:** When the positive predictive values for protected groups are disproportionately low, AI withholds necessary resources, such as clinical attention or social services. Such resource allocation discrepancies can exacerbate disparities in care and outcomes among the affected groups.

4. Patient Interaction-Related Biases: Privilege Bias, Informed Mistrust and Agency Bias

- **Privilege Bias:** When certain populations cannot access AI in care settings or when these algorithms require technology or sensors that are not available to all populations, leading to an unequal distribution of AI-driven healthcare benefits, potentially exacerbating existing healthcare disparities.
- **Informed Mistrust:** When some groups are skeptical about AI owing to historical exploitation and unethical practices in healthcare, leading these patients to avoid care or intentionally conceal information from clinicians or systems using AI.
- **Agency Bias:** When certain groups lack a voice in AI development, use, and evaluation. These groups may lack the access, resources, education, or political influence necessary to detect AI biases, voice concerns, and affect change. This lack of agency can result in AI being inadequate to consider the needs and perspectives of these groups, leading to biases and disparities in healthcare outcomes.

APPENDIX D: THE ARTIFICIAL INTELLIGENCE AND DATA ACT (AIDA)

The Artificial Intelligence and Data Act (AIDA): Regulating AI and Data in Canada

source: <https://www.fasken.com/en/knowledge/2023/03/a-roadmap-for-ai-regulation-in-canada-key-takeaways>

Bill C-27

This Bill contains three proposed Acts related to consumer privacy, data protection, and AI systems:

- *The Consumer Privacy Protection Act (CPPA)*
- *The Personal Information and Data Protection Tribunal Act (PIDPTA)*
- *The Artificial Intelligence and Data Act (AIDA).*

The Artificial Intelligence Data Act – AIDA

- *The Artificial Intelligence and Data Act (AIDA)* is the federal government's first attempt to regulate artificial intelligence comprehensively.
- AI regulation by other governments worldwide includes the European Union's 2021 EU AI Act and the United States' 2022 Algorithmic Accountability Act.

AIDA's Alignment with EU and OECD Standards

- Like the EU's AI Act, AIDA takes a risk-based approach to regulating AI. Canada categorizes AI based on whether it is "high impact," while the EU uses the language of "high-risk."
- AIDA is less prescriptive than the EU AI Act. The draft Act leaves room for provincial AI laws and further federal regulation.
- The government plans to make amendments to align with evolving international standards in the EU and the Organisation for Economic Co-operation and Development (OECD).
- Other planned amendments would impose more stringent requirements on high-risk systems at various stages of the system's lifecycle, including pre-commercialization, bringing AIDA closer to the EUAI Act.

Defining High-Impact Systems

- If passed, AIDA will regulate the design, development, and use of AI systems in the private sector, focusing on mitigating the risks associated with "high impact" AI systems.
- Much of the substantive content in AIDA is set to be established by future regulations, including the definition of a "high impact" AI system.
- The government's planned amendments to AIDA will set out an initial list of specific classes of high-impact AI systems*, which are those that are used for:
 - Employment-related determinations, such as hiring and remuneration.

- Determining whether to extend services to an individual, determining the costs and types of such services, and prioritizing the provision of such services.
- Processing biometric data for identification or determining an individual's behaviour or state of mind.
- Online content moderation on "online communications platforms", including search engines and social media, and the "prioritization of the presentation" (ie. recommendation of such content).
- Healthcare and emergency services.
- Decision-making by courts and administrative bodies.
- The exercise and performance of law enforcement power.

*The government can subsequently expand this list.

Obligations across the AI Value Chain

- In the AIDA Companion Document, the government's planned amendments would differentiate between obligations that apply to developers of machine-learning models that are intended for high-impact use, developers of high-impact systems, persons who make high-impact systems available for use, and persons who manage the operations of high-impact systems.
- The planned amendments would specifically regulate general-purpose AI systems, such as chatbots like ChatGPT, which may not be categorized as "high impact" but are widely used in various contexts.
- Developers of general-purpose AI systems would be required to perform specific risk assessments and mitigation testing during pre-market development. Once the system is on the market, a developer would need to make available a plain language description of the capabilities and limitations of the system and continuously monitor for harms and risks.
- Anyone who manages a general-purpose AI system must ensure individuals can identify AI-generated content.

REFERENCES

- Bajwa, J., Munir, U., Nori, A., & Williams, B. (2021). Artificial intelligence in healthcare: transforming the practice of medicine. *Future Healthcare Journal*, 8(2), e188–e194. <https://doi.org/10.7861/fhj.2021-0095>
- Cheng, L., Varshney, K. R., & Liu, H. (2021). Socially Responsible AI Algorithms: Issues, Purposes, and Challenges. <http://arxiv.org/abs/2101.02032>
- Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future Healthcare Journal*, 6(2), 94–98. <https://doi.org/10.7861/futurehosp.6-2-94>
- Evans, T., Naylor, D., Strome, E., Bernstein, A., & Goldenberg, A. (2020). Building a Learning Health System for Canadians Report of the Artificial Intelligence for Health (AI4H) Task Force. Retrieved from <https://cifar.ca/wp-content/uploads/2020/11/AI4Health-report-ENG-10-F.pdf>
- Fisher, S., & Rosella, L. C. (2022). Priorities for successful use of artificial intelligence by public health organizations: a literature review. In *BMC Public Health* (Vol. 22, Issue 1). BioMed Central Ltd. <https://doi.org/10.1186/s12889-022-14422-z>
- He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2019). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 25(1), 30–36. <https://doi.org/10.1038/s41591-018-0307-0>
- Health, V.-S. A. I., Drysdale, E., Dolatabadi, E., Chivers, C., Liu, V., Saria, S., Sendak, M., Wiens, J., Brudno, M., Hoyt, A., Mazwi, M., Mamdani, M., Allen, V., McGregor, C., Ross, H., Szeto, A., Anand, A., Wang, B., Alison Paprica, P., & Goldenberg, A. (2020). Implementing AI in Healthcare- Vector & Sick Kids. <https://vectorinstitute.ai/health-research/>
- Hee Lee, D., & Yoon, S. N. (2021). Application of artificial intelligence-based technologies in the healthcare industry: Opportunities and challenges. *International Journal of Environmental Research and Public Health*, 18(1), 1–18. <https://doi.org/10.3390/ijerph18010271>
- Keown, O. P., Parston, G., Patel, H., Rennie, F., Saoud, F., Al Kuwari, H., & Darzi, A. (2014). Lessons from eight countries on diffusing innovation in health care. *Health Affairs*, 33(9), 1516–1522. <https://doi.org/10.1377/hlthaff.2014.0382>
- London, A. J. (2022). Artificial intelligence in medicine: Overcoming or recapitulating structural challenges to improving patient care? In *Cell Reports Medicine* (Vol. 3, Issue 5). Cell Press. <https://doi.org/10.1016/j.xcrm.2022.100622>

- Lopez, N. (2018). Canadian Medical Association. Shaping the Future of Health and Medicine. Retrieved from <https://www.cma.ca/sites/default/files/pdf/Media-Releases/Shaping%20the%20Future%20of%20Health%20and%20Medicine.pdf>
- MacNeil, M., Koch, M., Kuspinar, A., Juzwishin, D., Lehoux, P., & Stolee, P. (2019). Enabling health technology innovation in Canada: Barriers and facilitators in policy and regulatory processes. In *Health Policy* (Vol. 123, Issue 2, pp. 203–214). Elsevier Ireland Ltd. <https://doi.org/10.1016/j.healthpol.2018.09.018>
- Moroz, I., Moroz, I., Archibald, D., Archibald, D., Breton, M., Cote-Boileau, E., Crowe, L., Crowe, L., Horsley, T., Hyseni, L., Johar, G., Keely, E., Keely, E., Keely, E., Burns, K. K., Kuziemy, C., Laplante, J., Mihan, A., Mihan, A., ... Liddy, C. (2020). Key factors for national spread and scale-up of an eConsult innovation. *Health Research Policy and Systems*, 18(1). <https://doi.org/10.1186/s12961-020-00574-0>
- OECD. (2024). Collective action for responsible AI in healthcare. Retrieved from <https://www.oecd.org/publications/collective-action-for-responsible-ai-in-health-f2050177-en.htm>.
- Ontario Health Data Council Report. A Vision for Ontario's Health Data Ecosystem (2023). Retrieved from <https://www.ontario.ca/page/ontario-health-data-council-report-vision-ontarios-health-data-ecosystem>
- Petersson, L., Larsson, I., Nygren, J. M., Nilsen, P., Neher, M., Reed, J. E., Tyskbo, D., & Svedberg, P. (2022). Challenges to implementing artificial intelligence in healthcare: a qualitative interview study with healthcare leaders in Sweden. *BMC Health Services Research*, 22(1). <https://doi.org/10.1186/s12913-022-08215-8>
- Sebastian, G., George, A., & Jackson Jr, G. (2023). Persuading Patients Using Rhetoric to Improve Artificial Intelligence Adoption: Experimental Study. *Journal of Medical Internet Research*, 25, e41430–e41430. <https://doi.org/10.2196/41430>
- Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. In *Nature Medicine* (Vol. 25, Issue 1, pp. 44–56). Nature Publishing Group. <https://doi.org/10.1038/s41591-018-0300-7>
- World Economic Forum (2023). Scaling smart solutions with AI in healthcare: Unlocking impact on high potential use cases. Retrieved from <https://www.weforum.org/publications/scaling-smart-solutions-with-ai-in-health-unlocking-impact-on-high-potential-use-cases/>
- Ueda, D., Kakinuma, T., Fujita, S., Kamagata, K., Fushimi, Y., Ito, R., Matsui, Y., Nozaki, T., Nakaura, T., Fujima, N., Tatsugami, F., Yanagawa, M., Hirata, K., Yamada, A., Tsuboyama, T., Kawamura, M., Fujioka, T., & Naganawa, S. (2024). Fairness of artificial intelligence in healthcare: review and recommendations. *Japanese Journal of Radiology*, 42(1), 3–15. <https://doi.org/10.1007/s11604-023-01474-3>