

11-28-2025 2:30 PM

Using Artificial Intelligence in Canadian Healthcare: Legal Reasoning and Governance Challenges in Privacy, Consent, and Transparency

Earlel Thiyagaratnam

Supervisor: Stark, Charles L., *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Master of Health Information Science degree in Health Information Science

© Earlel Thiyagaratnam 2025

Abstract

The growing use of artificial intelligence (AI) in Canadian healthcare raises significant governance challenges for legal systems that were largely developed before data-intensive and automated technologies became widespread. This thesis examines how Canadian law currently responds to issues created by health-related AI, including transparency, informed consent, data governance, and professional accountability.

Because there are very few Canadian court cases that deal directly with AI in clinical settings, this research analyzes court and tribunal decisions that address closely related issues involving digital technologies, complex information systems, and sensitive health data. These decisions are treated as practical sites where governance problems relevant to health AI already appear, even when artificial intelligence is not explicitly named.

Using a qualitative legal case study approach, the thesis analyzes a set of Canadian decisions to examine how courts reason about transparency, patient understanding, data flows, consent, and responsibility in technologically complex contexts. The analysis shows that Canadian courts frequently encounter difficulties when applying existing legal principles to systems that obscure how decisions are produced, rely on extensive data reuse, or embed expertise within technical infrastructures rather than individual professionals.

Across cases, courts struggle to balance demands for transparency with confidentiality obligations, to assess whether consent remains meaningful in complex information environments, and to allocate responsibility when automated or system-level tools influence decision-making. The thesis argues that these recurring legal tensions foreshadow deeper governance challenges as AI becomes more integrated into healthcare delivery. Rather than offering a statutory critique, the research highlights how judicial reasoning reveals structural pressures on concepts such as transparency, consent, and accountability in AI-adjacent contexts.

The thesis concludes by identifying areas where clearer governance approaches will be necessary, including improved consent practices, more robust expectations around transparency and explainability, clearer understandings of professional and institutional responsibility, and oversight mechanisms capable of adapting to rapid technological change. The findings demonstrate that while Canadian law provides important guiding principles, their application to health-related AI will require careful development to protect patient rights, maintain public trust, and support ethical healthcare innovation.

Keywords

Keywords: Artificial Intelligence (AI), Healthcare Law, Informed Consent, Privacy and Data Governance, Algorithmic Transparency, Professional Regulation, Accountability, Judicial Reasoning

Summary for Lay Audience

Artificial intelligence, often called AI, is increasingly used in Canadian healthcare to analyze medical information, support diagnoses, manage health data, and guide decision making. While these tools can be useful, they raise important questions about patient rights, privacy, consent, and responsibility when decisions are influenced by complex technologies.

This thesis examines how Canadian courts and tribunals engage with issues that are closely connected to the use of AI in healthcare. Because there are very few Canadian legal cases that directly involve AI in clinical care, the research analyzes related cases involving digital health tools, complex data systems, and the handling of sensitive health information. These cases show how legal challenges relevant to AI already arise in practice, even when artificial intelligence is not explicitly named.

The analysis focuses on judicial reasoning about transparency, meaning how clear and understandable systems are, informed consent, meaning whether individuals can meaningfully understand and agree to data use, and accountability, meaning how responsibility is assessed when decisions rely on technical systems. The findings show that courts often struggle to apply existing legal principles in situations where information flows are complex and expertise is embedded in technology rather than solely in human decision makers.

These judicial tensions highlight governance challenges that are likely to intensify as AI becomes more integrated into healthcare. Rather than offering a direct evaluation of specific statutes, the thesis uses these patterns in legal reasoning to anticipate areas where future governance approaches may require greater clarity, adaptability, and institutional support. Overall, the research demonstrates that Canadian law provides important starting principles, but that ongoing attention to transparency, consent, and responsibility will be essential as healthcare technologies continue to evolve.

Acknowledgments

Thank you everyone for your support, guidance, and mentorship throughout this research journey. I also want to thank you, readers, for reading my thesis.

Table of Contents

Abstract	ii
Keywords	iii
Summary for Lay Audience	iv
Acknowledgments	v
Table of Contents	vi
List of Tables	xi
List of Figures	xii
List of Appendices	xiii
Chapter 1	1
1 Introduction	1
1.1 Problem Context: The Emergence of Artificial Intelligence in Healthcare.....	1
1.2 Motivation for Examining Health and Artificial Intelligence Governance Through Legal Decisions.....	2
1.3 Research Contribution	3
1.4 Objectives	4
1.5 Research Questions.....	5
1.6 Thesis Organization	6
Chapter 2.....	8
2 Background	8
2.1 Artificial Intelligence in Healthcare: Key Concepts and Applications.....	8
2.2 Human-Artificial Intelligence in Clinical Contexts.....	8
2.2.1 The role of explainability.....	9
2.2.2 Transparency in clinical decision support.....	10
2.2.3 Human interpretation, trust, and uncertainty	11

2.3	Information Governance in Health AI	11
2.3.1	Privacy, consent, and data flows	12
2.3.2	Data quality and model behaviour	12
2.4	Existing Artificial Intelligence Governance Approaches	13
2.4.1	International and Canadian frameworks relevant to health contexts	14
2.4.2	The problem of uncertainty and rapid technological change	14
	Chapter 3	16
3	Literature Review: Foundations and Gaps in Health–AI Governance	16
3.1	Introduction	16
3.2	AI in Healthcare: Opportunities and Constraints	17
3.3	Transparency, Explainability, and Data Governance in Health AI	19
3.4	Synthesis: Gaps in Current Legal and Policy Knowledge	22
3.5	Conclusion	23
	Chapter 4	25
4	Methods	25
4.1	Rationale for Using Legal Decisions to Study Health–AI Governance	25
4.2	Research Questions	26
4.3	Methodology	26
4.3.1	Search strategy	28
4.3.2	Case selection process	30
4.4	Results	31
4.4.1	Case Selection Criteria	32
4.4.2	Final Case Set	37
4.4.3	Limitations of the Approach	38
	Chapter 5	39
5	C. J. v. M.W. (2020)	39

5.1 Background.....	39
5.2 Introduction.....	40
5.3 Case Background	41
5.4 Analysis of Legal Reasoning	41
5.4.1 Transparency and understanding digital assessments.....	41
5.4.2 Data governance and information sharing	42
5.4.3 Consent as a continuous process.....	43
5.5 Implications for Health AI Governance.....	43
Chapter 6.....	45
6 Doobay v. Cohen (2023).....	45
6.1 Background.....	45
6.2 Introduction.....	46
6.3 Case Background	47
6.4 Analysis of Legal Reasoning	48
6.4.1 Transparency in professional reasoning.....	48
6.4.2 Documentation and information sharing	48
6.4.3 Interpretation of complex information.....	49
6.5 Implications for Health AI Governance.....	49
Chapter 7.....	51
7 Lam v. Flo Health Inc (2024).....	51
7.1 Background.....	51
7.2 Introduction.....	52
7.3 Case Background	53
7.4 Analysis of Legal Reasoning	53
7.4.1 Transparency of data practices.....	53
7.4.2 Consent and user understanding	54

7.4.3	Data flows and third party involvement.....	54
7.5	Implications for Health AI Governance.....	55
Chapter 8	56
8	Cross-Case Thematic Analysis	56
8.1	Introduction.....	56
8.2	Theme 1: Transparency and Explainability	56
8.3	Theme 2: Data Governance, Privacy, and Consent	57
8.4	Theme 3: Consent as a Dynamic and Contextual Process	57
8.5	Theme 4: Human Interpretation and Reliance on Automated Systems	58
8.6	Theme 5: Judicial Reasoning and Anticipated Governance Gaps	58
Chapter 9	60
9	Discussion	60
9.1	Introduction.....	60
9.2	Implications for Health AI Governance in Canada.....	61
9.2.1	Transparency.....	62
9.2.2	Data flows	62
9.2.3	Consent and user understanding	63
9.2.4	Explainability and communication	63
9.3	Interdisciplinary Connections	64
9.4	Implications COVID-19 and the Increased Usage of AI.....	65
9.5	Strengths and Limitations	66
9.6	Further Research	68
9.7	Conclusion	70
References	73
Appendices	77
A	List of Abbreviations.....	77

B Chapter 4: Methodology Extended Information	78
Curriculum Vitae	87

List of Tables

Table 4.1: Case dates in years.....	35
Table 4.2: Legislation used in cases.	35

List of Figures

Figure 4.1: PRISMA flow diagram.....	33
Figure 4.2: Jurisdictions of cases.....	34
Figure 4.3: Place of case hearing.....	35
Figure 4.4: Resolution type of cases.....	36
Figure 4.5: Theme of main legal issues.....	37

List of Appendices

A	List of Abbreviations	77
B	Chapter 4: Methodology Extended Information	78

Chapter 1

1 Introduction

1.1 Problem Context: The Emergence of Artificial Intelligence in Healthcare

Artificial intelligence (AI) is increasingly used in healthcare to analyze large amounts of health information, such as medical records, diagnostic images, and genetic data. Many of these systems rely on machine learning (ML), which allows computers to identify patterns in data and make predictions based on those patterns. In healthcare, AI systems are used to support diagnosis, predict how diseases may progress, help plan treatments, and improve how healthcare systems operate overall (Johnson et al., 2021; Alowais et al., 2023). These technologies are often described as tools that can detect illness earlier, tailor care to individual patients, and make healthcare delivery more efficient.

ML plays a central role in many healthcare AI systems. By learning from past data, these systems can identify trends that may not be immediately visible to human clinicians, especially when the data are large or complex. In practice, ML is used to monitor chronic conditions, estimate future health risks, and support clinical decision-making by offering recommendations based on similar cases (Alowais et al., 2023). When used appropriately, these tools are intended to support clinical judgment by providing additional information, rather than replacing the role of healthcare professionals (Johnson et al., 2021).

At the same time, the growing use of AI is changing how health information is collected, processed, and used. Many AI systems require continuous access to large volumes of personal health data, including highly sensitive information such as reproductive and genetic data. These systems often rely on complex algorithms and data structures that are difficult to understand, even for experts. As a result, it can be unclear how decisions are produced, what data are being used, and how AI-generated outputs should be interpreted. These concerns are particularly serious in healthcare settings, where decisions may affect diagnosis, treatment options, or access to care (Gerke et al., 2020).

AI systems also affect relationships within healthcare institutions. As automated tools take on tasks that were traditionally performed by clinicians, decision-making processes change. Some forms of expertise become embedded in technology, while other forms remain with healthcare professionals or institutions. This shift can affect professional roles, reduce opportunities for patient involvement, and concentrate technical knowledge and control within technology companies rather than healthcare providers (Chui & Francisco, 2017). As Brynjolfsson and McAfee (2014) explain, automation does not simply replace human work but reorganizes labor, authority, and power. In healthcare, these changes are especially important because patients are often vulnerable and must rely on trust, professional judgment, and clear communication.

Therefore, the use of AI in healthcare raises governance questions, not just technical ones. Existing legal frameworks governing privacy, consent, and information use were largely developed before automated and data-intensive systems became widespread. As AI becomes more deeply integrated into healthcare delivery, it is unclear whether current laws can adequately address issues such as transparency, data governance, and meaningful human oversight of AI-supported decisions (Gerke et al., 2020).

1.2 Motivation for Examining Health and Artificial Intelligence Governance Through Legal Decisions

There are relatively few Canadian legal cases that deal directly with artificial intelligence in healthcare. This does not mean that governance problems related to health AI do not exist. Instead, it reflects the speed at which technology is developing compared to the slower pace at which legal disputes arise. Many of the key issues raised by health AI, including privacy, consent, data use, and transparency, already appear in Canadian legal decisions involving digital technologies and large-scale information systems.

Canadian courts and tribunals regularly consider whether individuals were properly informed about how their personal information was collected, used, or shared, whether organizations complied with privacy legislation, and whether institutional decision-making processes were sufficiently transparent. These legal questions closely resemble those raised by AI systems in healthcare, even when artificial intelligence is not mentioned explicitly. Judicial reasoning in

such cases therefore provides insight into how Canadian law applies governance principles in complex technological settings.

Recent cases involving digital health platforms illustrate this connection. In *Lam v. Flo Health Inc.* (2024), the court examined how a health application collected, used, and disclosed highly sensitive reproductive health data. Although the case did not involve clinical decision-making, it demonstrated how unclear or misleading data practices can undermine user understanding and trust in health-related technologies. The court's reasoning emphasized the importance of transparency, accuracy, and meaningful disclosure when sensitive health information is processed through complex technological systems.

More broadly, Canadian legal decisions addressing privacy, access to information, and technology-based decision-making show how courts deal with situations in which organizations control complex data systems and individuals have limited ability to understand them. These cases often address unequal access to information, the limits of meaningful consent, and the responsibilities of institutions that manage sensitive health data. Even when artificial intelligence is not explicitly discussed, the same governance concerns that arise in health AI contexts are often present.

This thesis uses Canadian legal decisions as a way to examine health AI governance. Court cases are treated as real-world settings where governance challenges become visible, even in the absence of AI-specific disputes or legislation. By analyzing how courts reason about data practices, transparency, and accountability, the thesis identifies gaps and tensions in existing legal frameworks and highlights how current law may struggle to respond as AI becomes more deeply embedded in healthcare.

1.3 Research Contribution

This thesis is part of an interdisciplinary degree program involving the Faculty of Health Sciences (FHS) and the Faculty of Information and Media Studies (FIMS).

In the field of **health sciences**, particularly clinical health, the contributions include:

1. Identifying legal and ethical challenges by mapping the main governance problems Canada faces when AI is used in healthcare. These include protecting patient privacy,

ensuring meaningful informed consent, reducing bias in data-driven systems, and clarifying responsibility when AI affects health decisions. This contribution is based on close analysis of seven Canadian court cases that raise issues similar to those found in health AI.

2. Connecting care principles to AI governance by showing how core healthcare values such as trust, transparency, accountability, and patient understanding influence whether AI use in healthcare is accepted and seen as legitimate. The analysis explains how Canadian courts deal with problems like lack of transparency and unequal access to information, and how existing laws support or weaken these care principles.

In the field of **information and media studies**, with a focus on information-seeking behaviors, the contributions are:

3. Creating transparency standards by developing practical ways to judge how clear and understandable AI systems used in healthcare are. Drawing on Canadian case law and governance research, these standards assess whether AI-related decisions and data practices can be accessed, explained, and challenged by patients, healthcare professionals, institutions, and regulators.
4. Highlighting information and power issues by showing how control over information in AI-assisted healthcare connects to broader problems related to data governance, changing work roles, and fairness. The cross-case analysis shows that when transparency is limited and technical expertise is concentrated within private organizations, it shapes who can access information and who has real influence over healthcare decisions.

1.4 Objectives

This thesis explores how Canadian law can help guide the responsible use of artificial intelligence in healthcare, even though there are very few Canadian court cases that deal directly with AI in medical settings. Because of this lack of direct case law, the thesis looks at Canadian court and tribunal decisions that address similar legal issues. These include transparency in decision making, how sensitive data is collected and used, how consent is obtained and respected, and how responsibility is assigned when decisions rely on complex

or automated systems. While these cases may not explicitly mention artificial intelligence, they raise legal and governance questions that closely mirror those created by AI in healthcare.

The purpose of the thesis is to understand how existing Canadian legal principles are likely to apply as AI becomes more common in healthcare. By examining how judges reason through related problems in other contexts, the thesis seeks to identify patterns, limits, and gaps in current governance frameworks. This approach allows the research to look ahead and assess whether existing laws and legal reasoning are well suited to address the challenges posed by health-related AI.

The thesis has three objectives:

Objective 1:

To examine how Canadian legal decisions address issues of transparency, data governance, and consent that are relevant to the use of artificial intelligence in healthcare. This includes analyzing how courts handle situations where decision processes are difficult to understand, where large amounts of personal or health data are involved, and where meaningful consent may be unclear or contested.

Objective 2:

To analyze judicial reasoning in health-adjacent cases in order to identify governance challenges linked to data-driven and automated systems. This objective focuses on how judges think about accountability, discretion, and oversight when decisions are influenced by technical systems rather than solely by human judgment.

Objective 3:

To synthesize insights across multiple legal cases to anticipate potential gaps in the governance of health-related artificial intelligence in Canada. By comparing patterns across decisions, the thesis highlights areas where existing laws may be silent, ambiguous, or strained when applied to AI-enabled healthcare.

1.5 Research Questions

This thesis addresses the following research questions:

1. How do Canadian legal decisions reveal challenges related to transparency and information governance that are relevant to artificial intelligence in healthcare?
2. How do existing legal approaches to privacy, data use, and consent speak to issues that also arise in health AI systems?
3. What patterns in judicial reasoning help anticipate gaps in the governance of health AI in Canada?

1.6 Thesis Organization

This thesis is organized in an integrated-article format.

Chapter 2 provides background on artificial intelligence in healthcare, focusing on human–AI interaction, transparency, data governance, and consent. It situates these concepts within existing legal and ethical discussions relevant to health-related data and automated systems.

Chapter 3 reviews the interdisciplinary literature on AI in healthcare, identifying key themes and gaps in Canadian-focused research. This chapter establishes the conceptual foundations for the thesis.

Chapter 4 outlines the research methodology. It explains the rationale for using Canadian legal decisions as qualitative case studies, describes the case selection process, and details the analytic approach used to identify governance-relevant themes.

Chapters 5, 6, and 7 each present an in-depth analysis of one Canadian legal decision. Each chapter is dedicated to a single case and examines how the court’s reasoning engages with issues relevant to health AI governance, including transparency, data governance, consent, and the interpretation of complex technological systems. While the cases do not always involve healthcare AI directly, they address information practices that closely parallel those used in health AI contexts.

Chapter 8 conducts a cross-case thematic analysis. Drawing on the findings from the three case chapters, this chapter identifies recurring patterns in judicial reasoning and examines how these patterns reveal strengths and gaps in current approaches to health AI governance in Canada.

Finally, **Chapter 9** concludes the thesis by summarizing the findings and discussing their implications for the future governance of artificial intelligence in Canadian healthcare. The conclusion reflects directly on the research objectives and questions and outlines directions for future research.

Chapter 2

2 Background

2.1 Artificial Intelligence in Healthcare: Key Concepts and Applications

AI refers to computer systems that are built to carry out tasks that people usually associate with human thinking. These tasks include recognizing patterns, making predictions, and supporting decisions. In healthcare, AI is most often used through ML methods. These methods examine very large and complex sets of data, such as medical images, electronic health records, genetic information, and data reported directly by patients. Many of these systems are now built into everyday clinical settings, where they help with diagnosis, treatment planning, risk assessment, and operational decisions. AI has shown strong performance in areas such as radiology, oncology, cardiology, and pathology. In these fields, clinicians routinely work with complex data where detecting patterns is essential (Esteva et al., 2019; Jiang et al., 2017). Predictive models are also used to estimate patient outcomes, guide how resources are distributed, and identify patients who may be at higher risk of complications or harm (Topol, 2019). In addition to direct clinical uses, AI tools are applied in drug discovery and health system management, where they can speed up research processes and improve efficiency (Holmes et al., 2015). Even with these developments, AI systems do not replace human judgment. Their results must be reviewed, interpreted, and applied by clinicians within existing organizational and professional structures. When AI is introduced into healthcare, it changes how medical decisions are created, explained, and defended. This shift raises basic questions about how AI systems work, how transparent and understandable they are, how data are used, and how human judgment fits into the process. These are not mainly technical problems. They are governance issues that arise at the intersection of healthcare, law, and information systems.

2.2 Human-Artificial Intelligence in Clinical Contexts

AI is changing healthcare by helping doctors make more accurate diagnoses, creating personalized treatment plans, and improving how hospitals and clinics operate. AI can look at large amounts of medical information, such as images, genetic data, and patient records, to

find patterns that humans might miss. This can lead to earlier and more accurate diagnoses, especially in areas like radiology, cancer care, and heart disease (Esteva et al., 2019; Jiang et al., 2017). AI can also predict patient outcomes, help hospitals use resources more efficiently, and support doctors in making decisions (Topol, 2019). It even speeds up the development of new drugs (Holmes et al., 2015). Even so, AI does not make decisions on its own. It is a tool that doctors use, and its usefulness depends on how well they understand and trust its results. Hospitals and clinics often face uncertainty and incomplete information, so doctors must use their judgment. AI adds another layer by turning patient information into probabilities or categories. How doctors and AI work together is key to whether AI helps or harms patient care. Studying how humans interact with AI is important for making sure it is used safely and effectively. As AI becomes more common in healthcare, it raises questions about how clear and understandable it is, and how much humans should rely on it. Transparency means explaining AI in ways that both patients and healthcare providers can understand, not just sharing technical details. This is important for building trust, using AI ethically, and making sure data privacy and human oversight are protected (Stanfill & Marc, 2019). These issues must be carefully managed to make sure AI is used responsibly and fairly in healthcare.

2.2.1 The role of explainability

Explainability is about how well people can understand how an AI system comes to its decisions. In healthcare, this is very important because medical decisions affect patients' lives and need to be explained to patients, doctors, and healthcare organizations. ML, a key part of AI, is changing patient care by analyzing large amounts of medical data to find patterns that help with diagnosis and treatment. ML algorithms look at patient histories, lab results, and medical images to predict how diseases might progress, spot health problems early, and suggest personalized treatments. For example, ML can predict flare-ups in chronic conditions so doctors can act early and improve outcomes (Holmes et al., 2015). ML can also help doctors choose the most effective treatments for each patient (Topol, 2019). ML can find complex patterns in data and make predictions that humans might miss (Esteva et al., 2019). But many high-performing ML systems are "black boxes," meaning it is hard to see how they turn input data into predictions. This makes it harder for doctors to interpret results, especially if the AI suggestion goes against their experience or usual medical practice. As ML systems learn from new data, their predictions get better over time (Esteva et al., 2019).

Explainability helps people use AI wisely instead of blindly trusting it. Doctors need to judge if a prediction makes sense for a patient, and patients need to understand how AI affects their care. Being able to see and check how ML makes decisions builds trust, keeps patients safe, and helps ensure fair treatment (Lipton, 2018). Explainability is especially important because patients and healthcare workers are usually not experts in AI. At the same time, ML raises concerns about data privacy, bias in algorithms, and keeping the human touch in care (Stanfill & Marc, 2019). These issues must be handled carefully so that ML is used safely, fairly, and responsibly in healthcare.

2.2.2 Transparency in clinical decision support

Transparency in AI means more than just being able to explain how it works. It includes being clear about how AI systems are designed, how they are used, and how information from them is shared in healthcare. Transparent AI systems show where the data comes from, what the limitations are, and how uncertain their recommendations might be, instead of acting like they are always correct. This is important because AI can influence how clinicians make decisions. If recommendations are given without enough context, doctors might rely too much on AI or follow it without question. On the other hand, if AI is not clear enough, people might ignore useful advice (Rajkumar et al., 2018). Using AI in healthcare also affects the work, expertise, and authority of healthcare professionals. AI can help doctors and nurses by analyzing complex data and spotting patterns that might be missed, which can lead to better decisions and outcomes. But there is a risk of relying too much on AI, which could make human expertise less valued and change the balance of power in healthcare (Davenport & Ronanki, 2018). Transparency helps prevent this by showing how decisions are made and making it clear that AI is a tool to support, not replace, human judgment. If AI is seen as replacing professionals, important human aspects of care could be lost, and existing inequalities could get worse. As AI takes over tasks like diagnosing conditions, human judgment and experience might be undervalued. Companies that control AI could gain too much influence over decisions and access to important information, which could make inequality worse (Chui & Francisco, 2017). To use AI fairly and responsibly, it is important to set rules that make AI a supportive tool, helping human decision-making while keeping accountability, transparency, and trust, especially for patients (Brynjolfsson & McAfee,

2014; Frey & Osborne, 2017). This way, AI and healthcare professionals can work together, with technology boosting human skills without replacing them.

2.2.3 Human interpretation, trust, and uncertainty

Trust in AI systems is not something fixed. It develops through use, depends on the situation, and is influenced by professional and institutional expectations. In healthcare, trust in AI should be balanced. Clinicians need enough trust to use AI tools, but not so much that they stop questioning or thinking critically about the system's recommendations. Human judgment is essential for dealing with uncertainty. AI systems produce predictions based on probabilities and data patterns, but medical decisions often require clear, yes-or-no choices. Clinicians interpret AI outputs using their own experience, knowledge of the patient, and the specific circumstances of care factors that AI systems cannot fully account for. This shows the limits of automation in healthcare. AI can support clinicians by identifying patterns and making predictions, but it cannot replace the ethical judgment, emotional understanding, and contextual awareness involved in patient care. For AI to be used effectively, healthcare systems must preserve the role of human decision-making and be open about uncertainty rather than hiding it.

2.3 Information Governance in Health AI

Information governance is about the rules and practices that guide how data are collected, used, shared, and understood. In healthcare AI, this is especially challenging because AI systems rely on constant streams of data and complex processes that are often hidden from users. Health data are very sensitive because they include personal details about people's bodies, behaviors, and identities. Using these data in AI raises ongoing concerns about privacy, consent, and responsible management, while also creating new issues around large-scale use, reusing data, and analyzing it for other purposes. Human-centered AI focuses on making sure AI systems support people and their values. Instead of only looking at technical performance, it considers how AI affects patients, clinicians, and others. Human-centered AI aims to design systems that are easy to understand, not "black boxes," so people can see how decisions are made. By aligning technology with human expertise, this approach builds trust, fairness, and accountability in healthcare. In this way, human-centered AI shows why clear

rules, ethical practices, and user-friendly design are essential in today's fast-changing healthcare environment.

2.3.1 Privacy, consent, and data flows

AI systems in healthcare rely on collecting large amounts of data from different sources over time. This makes it hard to get meaningful consent because people often cannot know in advance how their data will be used. This creates a gap between traditional consent approaches and the way data-driven AI works. Transparency helps address this problem. People need to understand not only what data are being collected, but also how it may be used, shared, or combined. Canadian privacy laws stress that consent must be meaningful and that people should have clear expectations about how their information is handled. Legal cases involving digital health apps show that data can move in ways users do not expect, especially when third-party software or analytics are involved, highlighting the challenges of managing complex data flows. Transparency is also a key part of making AI human-centered. In healthcare, patients and providers need to understand how AI systems make decisions, including the data used, the algorithms applied, and the reasoning behind predictions. Simple explanations, visual tools like decision trees or heat maps, and plain-language summaries can help. Transparency supports informed consent, helps ensure fairness, and lets users question decisions that affect their care (Lipton, 2018). Many AI systems, especially deep learning models, are considered "black boxes" because their decision processes are hard to interpret. This raises concerns about fairness, accountability, and bias (Sokol & Flach, 2020). Explainable AI (XAI) techniques aim to make these models easier to understand and reduce bias while promoting ethical use (Thunki et al., 2021). Despite progress, fully transparent AI remains difficult, and balancing clear explanations with model performance is still a major challenge (Iyer et al., 2018).

2.3.2 Data quality and model behaviour

AI models learn from the data they are trained on. In healthcare, if the data is incomplete, biased, or not representative, the model may perform unevenly across different patient groups, which can reinforce existing inequalities. Understanding how a model behaves means looking at both its technical design and the context in which it is used. Being clear about data sources, limitations, and uncertainty helps doctors and hospitals decide whether an

AI system is suitable for their patients. Without this clarity, it is hard to judge how reliable a model is or to spot systematic errors (Arrieta et al., 2020). Transparency is especially important in healthcare because AI models often act like “black boxes,” making decisions that are difficult to understand. If clinicians or patients cannot see how decisions are made, they may be hesitant to trust AI recommendations (Lipton, 2018; Rajkomar et al., 2018). Tools from XAI, such as SHapley Additive exPlanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), and Gradient-weighted Class Activation Mapping (Grad-CAM), can help, but it is still challenging to make these explanations clear and useful for everyday clinical practice (Arrieta et al., 2020). True transparency involves more than technical reports. It means clearly communicating the limits, biases, and uncertainties of AI systems so that clinicians and patients understand not just the “how” but also the “why” behind decisions (Esteva et al., 2019; Lipton, 2018). Good data governance supports transparency by focusing not only on privacy but also on data quality and interpretability. Simply having a technically accurate model is not enough; healthcare organizations need ways to understand, monitor, and oversee AI systems. Working together, AI developers, healthcare professionals, and legal experts, helps create standards for clear, accurate, and practical explanations. This approach makes AI more trustworthy, safer, and ethically responsible in healthcare settings (Arrieta et al., 2020).

2.4 Existing Artificial Intelligence Governance Approaches

AI governance in healthcare has developed through international guidance, professional standards, and local regulations. It focuses on creating rules and policies to make sure AI is used ethically, transparently, and in ways that protect people’s rights while supporting the public good (Binns, 2018; Brynjolfsson & McAfee, 2014; Williamson & Prybutok, 2024). Most approaches rely on broad principles rather than strict rules, highlighting transparency, human oversight, fairness, and flexibility. This is because AI technologies change quickly, and governance systems need to handle uncertainty about how AI will evolve and interact with laws and ethical standards (Taeiagh, 2021). Effective governance also considers social, ethical, and legal factors, using continuous monitoring, regular updates, scenario planning, and impact assessments to balance innovation with public safety (Binns, 2018; Siala & Wang, 2022).

2.4.1 International and Canadian frameworks relevant to health contexts

International organizations have been important in guiding how AI is used in healthcare. The World Health Organization (WHO) released guidance in 2021 on the ethics and governance of AI in health. This guidance, developed with experts in ethics, technology, law, and human rights, stresses that AI systems should be ethical, transparent, and accountable, while protecting human rights (World Health Organization, 2021). In 2023, WHO, together with the International Telecommunication Union (ITU) and the World Intellectual Property Organization (WIPO), started the Global Initiative on AI for Health (GI-AI4H). This initiative focuses on creating long-term support for AI in healthcare, especially in low- and middle-income countries, to ensure fair access to AI tools and health innovations (World Health Organization, 2023). In 2024, WHO issued further guidance on large multi-modal AI models in healthcare, giving over forty recommendations to make sure these complex systems are used safely and responsibly (World Health Organization, 2024). Overall, these international efforts show the need for clear rules and strong oversight as AI becomes more complex. In Canada, health-related AI governance relies on existing privacy laws, health regulations, and emerging AI policies. There is currently no single law specifically for AI in healthcare. However, rules about data protection, informed consent, and professional standards still apply and provide guidance. Canadian courts are increasingly interpreting these rules in cases involving digital health tools, showing how traditional laws are adapting to new technologies.

2.4.2 The problem of uncertainty and rapid technological change

Health AI governance faces a major challenge: uncertainty. AI technology develops very quickly, often faster than laws, regulations, and ethical guidelines can keep up. ML models are constantly improving, and new or unexpected uses of AI can disrupt healthcare, affecting diagnosis, treatment, and patient trust. Traditional regulatory systems, built for slower, more predictable developments, struggle to manage these fast changes (Binns, 2018; Taeihagh, 2021). To address this, governance needs to be flexible and adaptable. Policies should include regular reviews and updates, so rules can change as AI evolves. Tools like scenario planning, impact assessments, and iterative policy development help anticipate problems and guide decision-making without stopping innovation (Binns, 2018; Taeihagh, 2021). AI's impact goes beyond healthcare. It can affect privacy, jobs, and fairness in society, making

regulation even more complex (Siala & Wang, 2022). Policymakers should work together internationally to set standards that ensure AI develops in an ethical, transparent way that matches societal values (Floridi et al., 2020). Governance frameworks must balance flexibility with clear guidance, involve experts from many fields, and stay proactive to make sure AI benefits patients and the public while still allowing innovation (Brynjolfsson & McAfee, 2014).

Chapter 3

3 Literature Review: Foundations and Gaps in Health–AI Governance

3.1 Introduction

The existing literature on AI in healthcare covers a broad range of topics, including technical progress, ethical issues, and emerging legal challenges. However, much of the research simply summarizes key studies without critically analyzing their contributions or addressing the gaps that still need attention. This chapter aims to build on the existing literature by not only reviewing previous studies but also assessing their strengths and weaknesses. In doing so, it highlights where current legal and ethical frameworks are inadequate in keeping up with the rapid pace of AI development. One major area of focus is the inconsistent conceptualization of transparency, explainability, and human-AI interaction, with definitions varying widely across different studies. This lack of agreement makes it difficult to compare findings and develop clear policy recommendations. Additionally, transparency is often treated as a purely technical issue, but it is also a socio-legal challenge, particularly for non-expert stakeholders, such as patients and clinicians. These stakeholders need clear and understandable information about how AI systems make decisions. Unfortunately, transparency is expected as a given, without fully exploring its practical implications for building trust and accountability. By synthesizing these existing studies, chapter three highlights the gap between AI's potential in healthcare and its current application within outdated regulatory and ethical frameworks. This gap demonstrates the need for further research that connects technological innovation with effective regulatory oversight, a need that this thesis seeks to address.

To support this analysis, qualitative case study research is used to combine and analyze multiple individual cases to identify broader patterns, insights, and conclusions. Unlike traditional case studies, which typically focus on a single example, qualitative case study research brings together findings from different contexts to explore common themes, trends, or outcomes (Eisenhardt, 1989; Gerring, 2007). This method is particularly valuable in complex fields like healthcare, where practices and results can vary widely across different

settings. By synthesizing data from diverse environments, such as large hospitals and smaller rural clinics, researchers can gain a more comprehensive understanding of how AI technologies affect clinical workflows, diagnostics, patient monitoring, and treatment planning (Rajkomar et al., 2018; Esteva et al., 2019). This approach not only helps identify patterns related to AI's effectiveness and scalability but also sheds light on common barriers to its adoption, such as regulatory hurdles, data privacy concerns, and challenges in integrating AI systems into existing workflows (Obermeyer et al., 2019). Ultimately, a qualitative case study research approach provides valuable insights into best practices for AI implementation and highlights areas where further innovation and policy development are needed to ensure that AI can fully realize its potential to transform healthcare delivery (Topol, 2019).

3.2 AI in Healthcare: Opportunities and Constraints

This literature review provides an overview of the current state of research on the application of AI in healthcare. It examines key theories, technological advancements, and empirical studies, while identifying gaps and opportunities for further investigation in this rapidly evolving field. The literature reviewed here was selected because it covers both the technical potential and the legal-ethical challenges of AI in healthcare, while also revealing key gaps this research seeks to address. While there is a large amount of work on AI in healthcare, much of it either concentrates narrowly on technical performance or discusses ethics broadly without linking these ideas to legal and policy issues, particularly in Canada.

The chosen works provide in-depth analysis on explainability and trust in AI, the role of AI in urgent crises like COVID-19, and legal questions about liability and accountability. They also include thorough ethical analyses that highlight recurring concerns such as fairness, transparency, inclusivity, and human-centered design. Together, these studies offer a balanced view of the field, addressing human-AI interaction, clinical workflow integration, pandemic-driven innovation, legal liability, and ethical governance. Notably, they also reveal a lack of jurisdiction-specific research, especially in Canada, where the relationship between AI's technical capabilities and existing legal frameworks remains understudied. Their depth, diversity, and identification of research gaps make them crucial foundations for this thesis.

For instance, Mehrotra et al. (2024) examines how to foster appropriate trust in human-AI interaction by analyzing trends, challenges, and opportunities. The authors highlight that while trust in AI is crucial to prevent misuse (over-reliance) or disuse (under-reliance), there is no unified definition of "appropriate trust." They propose a framework called Belief, Intentions, and Actions (BIA) to categorize trust-related concepts: beliefs (human perceptions of AI trustworthiness), intentions (plans to interact with AI), and actions (actual reliance behaviors).

Loh et al. (2022) reviews the role of XAI in healthcare, with a focus on how enhancing the transparency of AI predictions can build trust and promote adoption in clinical settings (Dong et al., 2021). The review highlights critical areas, such as detecting abnormalities in one-dimensional biosignals and interpreting clinical notes, which require additional attention from the XAI research community to improve healthcare applications (Loh et al., 2022).

Similarly, Nazar et al. (2021) examines the emerging field of XAI and its impact on healthcare, emphasizing the challenges of integrating XAI into Human-Computer Interaction (HCI). Their research reviews key aspects of ML, identifies various XAI techniques, and highlights the specific challenges faced when applying XAI in healthcare contexts, particularly in enhancing user trust and improving system interactions.

The COVID-19 pandemic significantly accelerated the adoption of AI, as many industries, including healthcare, shifted to virtual environments for service delivery. For example, Sarker et al. (2021) explores the role of robotics and AI-based technologies in healthcare during the pandemic, focusing on their contributions to diagnosis, risk assessment, telehealthcare, disinfection, vaccine distribution, and mental health support. Through a systematic review of 147 studies, this research shows how these technologies have eased the burden on frontline healthcare workers and expedited vaccine development and distribution.

Cestonaro et al. (2023) examines the legal and ethical challenges of using AI in medical diagnostics, focusing on who should be held responsible if AI errors harm patients. AI improves diagnostic accuracy, reduces workloads, and enables personalized treatments, but it also raises concerns about biased algorithms, privacy issues, and the "black box" problem, where even developers cannot fully explain how AI reaches decisions. However, existing regulations are insufficient, and the authors stress the urgent need for updated legal

frameworks to address AI's unique risks, ensure transparency, and maintain trust in healthcare systems. They conclude that clear guidelines and specialized insurance for AI-related risks are essential as AI becomes more integrated into medicine.

The ethical issues surrounding the use of AI in healthcare are another important area of concern. Siala and Wang (2022) investigates these ethical challenges by reviewing 253 articles on AI ethics in healthcare. Their study aims to identify key themes for the responsible implementation of AI in healthcare settings. The authors propose a framework called SHIFT, which stands for Sustainability, Human-Centeredness, Inclusiveness, Fairness, and Transparency, and highlight the ethical challenges that need further exploration in future research on AI adoption in healthcare.

In addition, Tang et al. (2023) conducts a systematic review of 36 empirical studies on the ethics of medical AI, examining topics such as stakeholder attitudes, factors influencing acceptance, and the correction of biases in AI technologies. Their study points out the gap between high-level ethical guidelines and empirical research, urging for closer collaboration between ethicists, AI developers, clinicians, and other relevant stakeholders to address the ethical issues in medical AI more effectively.

Similarly, Karimian et al. (2022) performs a systematic scoping review of the ethical issues surrounding the use of AI in healthcare, focusing on key principles like autonomy, fairness, explainability, and privacy. Their review identifies gaps in the literature, such as limited exploration of ethical considerations during the design and deployment of AI, a lack of practical tools for ensuring ethical compliance, and the absence of diverse stakeholder perspectives in the conversation about AI ethics in healthcare.

3.3 Transparency, Explainability, and Data Governance in Health AI

Building the right level of trust between humans and AI is still a major challenge. The problem is made harder by unclear definitions, which often confuse terms like "appropriate trust," "calibrated trust," and "appropriate reliance." Current ways to measure trust are also inconsistent, using self-reports or behavior metrics without considering the full context (Mehrotra et al., 2024). Most research looks only at AI performance, such as accuracy and reliability, while ignoring other key factors like integrity, goodwill, and user-focused

intentions. Many studies also fail to test AI in realistic, high-risk situations or track how trust changes over time. Tools meant to improve transparency, like explanations or confidence scores, have mixed results and can sometimes lead to over-trust. To tackle these issues, Mehrotra et al. (2024) suggest the Belief, Intentions, Actions (BIA) framework. This model breaks trust into three parts: human perceptions, planned reliance, and real-world behavior. They also recommend shifting from just XAI to "contestable AI," where users can question AI decisions. Additionally, they call for clearer definitions and ethical guidelines, especially for critical fields like healthcare.

XAI is important for gaining the trust of clinicians and patients by making AI predictions easy to understand. Methods like SHAP are commonly used to analyze structured clinical data, such as finding key predictors using models like XGBoost. For medical images, techniques like GradCAM work well by creating heatmaps using CNNs (convolutional neural networks). Attention mechanisms are useful for text data, such as clinical notes (Loh et al., 2022; Nazar et al., 2021). Beyond clarity, XAI helps with checking model accuracy, adjusting settings, and comparing models trained on private datasets. However, there are still major challenges in using XAI for 1D biosignals (like ECG) and clinical text. Many current models do not meet user needs because their interfaces are not user-friendly. Other issues include ensuring accuracy for specific medical fields, avoiding misleading explanations, and solving broader problems like data security and compatibility. Developing strong, flexible XAI models that work across different fields remains a key challenge (Loh et al., 2022; Nazar et al., 2021).

The COVID-19 pandemic sped up the use of AI in many areas, especially in healthcare. AI helped quickly analyze medical images like Computed Tomography (CT) scans and X-rays using advanced tools like ResNet and DenseNet. It also supported contactless sample collection through robotics, improved testing schedules, expanded telehealth services (including mental health care), and helped with vaccine development and supply chain management (Sarker et al., 2021). These uses eased the workload for frontline workers and sped up pandemic responses. However, there were concerns. AI sometimes lacked the accuracy of traditional methods, faced resistance from healthcare workers and the public, and was expensive to implement. Production limits made it hard to use in low-resource areas. Many AI solutions were still in early testing, so large-scale use was limited. Additionally, the

constant changes in the COVID-19 pandemic made it difficult for AI models to stay effective, showing the need for systems that can adapt over time (Sarker et al., 2021).

Determining who is responsible for harm caused by AI involves complex legal issues that are not yet fully resolved. Many AI systems operate like "black boxes," making it hard to understand how they make decisions and difficult to prove what caused the harm.

Additionally, AI trained on biased data can worsen health inequalities (Cestonaro et al., 2023). Liability could fall on different parties, such as doctors (for supervising AI use), hospitals (for their employees' actions), developers (for faulty products), or even AI systems themselves if they are treated as legal entities. Some suggest shared responsibility among these groups, but there are no clear laws to support this. Key problems remain, including the lack of rules for transparency, safety standards, and reducing bias. Patients are often not properly informed about how AI is used or its risks. Insurance options for AI-related harm are also underdeveloped. Long-term effects, like doctors losing skills or patients receiving less empathy, are not well understood but are important for using AI responsibly (Cestonaro et al., 2023).

Siala and Wand (2022) outline key ethical principles for AI in healthcare. These include Sustainability (ensuring professionals receive AI ethics training), Human-Centeredness (using AI to support, not replace, practitioners while protecting patient-provider relationships), Inclusiveness (preventing corporate control of data and involving diverse stakeholders), Fairness (reducing bias through diverse datasets and local testing), and Transparency (with strong data governance and clear explanations). Patients and clinicians frequently raise concerns about AI's impact on human connection, autonomy, privacy, data security, and fairness. Clinicians, in particular, worry about losing professional independence (Tang et al., 2023; Karimian et al., 2022). Research shows that AI systems often perform worse for marginalized groups, including women, racial minorities, and publicly insured patients. This highlights the need for testing across diverse populations. Media coverage tends to focus on AI's benefits while overlooking ethical risks, and incomplete demographic data in training sets can reinforce bias (Tang et al., 2023). A major gap exists between ethical principles and real-world implementation. AI systems sometimes disregard individual preferences, and their lack of transparency makes it difficult for clinicians to explain decisions to patients. Privacy concerns remain unresolved, and the principle of "preventing

harm" is the least explored. Few tools exist to assess whether AI systems follow ethical guidelines or ensure fairness. Additionally, the perspectives of vulnerable groups and health IT leaders are often overlooked (Karimian et al., 2022; Tang et al., 2023). Proposed solutions include ethical frameworks based on virtue ethics and post-deployment reviews (Siala and Wand, 2022), "embedded ethics" (where ethicists, developers, clinicians, and patients collaborate throughout AI development), and practical evaluation methods aligned with guidelines like the European Union's (EU) Ethics Guidelines for Trustworthy AI (Karimian et al., 2022).

Cestonaro et al. (2023) found several problems in dealing with medical malpractice involving AI. One major issue is that current laws do not clearly state who is responsible when AI makes a mistake, making it hard to assign blame between developers, doctors, and hospitals (Cestonaro et al., 2023). Another problem is the "black box" nature of AI, which makes it difficult to understand how decisions are made and hard to prove what caused an error. This leaves both patients and healthcare providers without clear ways to seek justice. Additionally, AI algorithms can be biased because they are often trained on data that doesn't represent all groups fairly. This can worsen healthcare inequalities, but there are still few solutions to fix these biases (Cestonaro et al., 2023). The review also points out that regulations for AI transparency and safety are lacking, and current proposals don't fully handle the challenges posed by "black box" systems. There are also gaps in informed consent, as patients are usually not properly told about AI's role or risks in their care. Finally, there are no specialized insurance models or guidelines to handle liability issues, and little research exists on long-term effects, such as doctors losing skills or patient interactions becoming less empathetic (Cestonaro et al., 2023). These gaps show an urgent need for experts from different fields to work together to update policies and ensure trust in AI-driven healthcare.

3.4 Synthesis: Gaps in Current Legal and Policy Knowledge

A major gap in current research is the lack of jurisdiction-specific legal analysis, particularly in Canada, regarding how courts handle AI-related disputes involving liability, informed consent, privacy, safety, and fairness. Existing privacy laws and health regulations have not been thoroughly examined for their ability to govern AI, creating uncertainty. There is also little exploration of how courts balance innovation against patient rights and public safety

(Loh et al., 2022; Nazar et al., 2021; Siala & Wand, 2022). Another pressing issue is the disconnect between theory and practice in AI ethics and XAI). While ethical principles exist, practical tools to implement them are lacking. XAI solutions remain technically focused rather than user-centered. More robust methods are needed to ensure real-world fairness, privacy, and transparency. There are no comprehensive trust frameworks that combine technical capability with ethical integrity, and no legally viable accountability models for AI errors. The "black box" problem also undermines transparency and informed consent (Mehrotra et al., 2024; Cestonaro et al., 2023).

Although algorithmic bias risks are widely recognized, there are few scalable strategies to mitigate them. Marginalized groups are often excluded from AI governance discussions. Additionally, disciplinary barriers between technologists, clinicians, ethicists, legal experts, and patients prevent holistic solutions. Effective AI integration requires sustained, interdisciplinary collaboration (Tang et al., 2023; Karimian et al., 2022). Existing research on AI in healthcare highlights its transformative potential, particularly during the COVID-19 pandemic, in areas like diagnostics, telemedicine, and supply chain optimization (Sarkar et al., 2021). However, significant gaps remain, especially regarding Canada's legal and regulatory framework. Few studies examine how Canadian courts address AI-related legal disputes, leaving key questions unanswered. While ethical concerns such as privacy, safety, and fairness are well-documented, there is limited analysis of how they appear in Canadian case law. It is also unclear how courts weigh innovation against patient rights and public safety, which is critical for understanding judicial priorities. Although privacy laws and healthcare regulations are acknowledged as important for AI governance, the literature rarely examines gaps or ambiguities in Canada's legal system. This creates uncertainty about whether current laws can effectively regulate AI in healthcare.

3.5 Conclusion

AI's potential to revolutionize healthcare is undeniable, yet its realization is hampered by misaligned ethical, legal, and human-factors frameworks. Overcoming conceptual fragmentation in trust and transparency, resolving liability ambiguities, closing the ethics implementation gap, ensuring equitable outcomes, and fostering genuine collaboration are imperative. Addressing these gaps is essential to understand how Canadian law responds to

AI's ethical and regulatory challenges in healthcare. This requires both deeper analysis of existing legal frameworks and greater interdisciplinary collaboration to develop practical solutions that balance innovation with patient protection. Future research must prioritize adaptable governance models, clinically actionable XAI, practical ethics assessment tools, and strategies for building holistic trust. Only through addressing these gaps can AI's promise be fulfilled responsibly and equitably.

Chapter 4

4 Methods

This thesis uses a qualitative case study research approach to analyze Canadian legal cases involving the integration of AI technologies in healthcare. The main goal is to explore how AI is reshaping patient care, identify the legal challenges that arise, and evaluate how existing regulatory frameworks are addressing these changes. Key issues such as transparency, accountability, and the dynamics of human-AI interaction are central to this investigation. Here, "human-AI interaction" refers to the ongoing process where humans and AI systems exchange information and influence decision-making in both clinical and administrative contexts. This definition is based on a review of the literature that highlights the technical aspects as well as the relational elements of AI integration. By focusing on qualitative analysis, the approach provides an in-depth exploration of judicial reasoning and real-world implications, which are often overlooked in purely quantitative studies. It also highlights the complex relationship between human expertise and AI-driven decision-making, offering a comprehensive framework for understanding the benefits and challenges of AI in healthcare. These insights are crucial for shaping future policies that promote AI systems that are transparent, accountable, and effectively complement human decision-making (Eisenhardt, 1989; Gerring, 2007; Rajkomar et al., 2018; Esteva et al., 2019; Obermeyer et al., 2019; Topol, 2019).

4.1 Rationale for Using Legal Decisions to Study Health–AI Governance

AI is being used more frequently in Canadian healthcare, supporting clinicians in diagnosis, treatment planning, and administrative tasks. At the same time, AI introduces complex challenges related to transparency, explainability, privacy, consent, and the handling of patient information. While few Canadian court cases deal directly with AI in healthcare, decisions involving healthcare professionals, digital health tools, and privacy issues offer useful insights into how the law addresses similar governance challenges. Examining these cases allows this thesis to explore how courts reason about transparency, informed consent, and data governance in technologically complex contexts. The analysis identifies recurring

patterns, tensions, and potential gaps in governance, providing a foundation for understanding how AI can be used responsibly in healthcare while safeguarding patient trust and equitable access to services.

4.2 Research Questions

I address three research questions to identify Canadian cases that provide issues of lived experiences issue with AI in healthcare:

1. How do Canadian legal decisions reveal challenges related to transparency and information governance relevant to AI in healthcare?
2. How do existing laws concerning privacy, data use, and consent apply to issues raised by AI in health systems?
3. What patterns in judicial reasoning help anticipate governance gaps when AI is used in Canadian healthcare settings?

4.3 Methodology

Case studies are often criticized for not having the same level of rigor as other research methods, which typically follow more clearly defined guidelines for data collection and analysis. Critics argue that researchers must be transparent about their decision-making throughout the research process and clearly justify their choices (Meyer, 2001). In response to these concerns, several methodologies for conducting case study research have been proposed (Ebneyamini & Sadeghi Moghadam, 2018), each offering a different structure and level of flexibility. Yin and Dhameja (1997) structured approach to case study research outlines four main stages: designing the study, collecting data, analyzing the data, and drawing conclusions and recommendations. On the other hand, Stake and Walker (1996) advocates for a more flexible approach, introducing the concept of "progressive focusing," where the researcher refines their focus as new data is gathered and insights emerge. Tellis (1997) follows a similar four-stage process based on Yin's framework, as discussed by Yin and Dhameja (1997), but provides more specific detail at each stage, including designing a case study protocol and using tools such as interviews and questionnaires for data collection. Merriam (1998) offers a detailed, step-by-step guide that includes reviewing the literature,

developing a theoretical framework, refining research questions, and selecting a sample. Voss et al. (2002) highlight important steps such as defining the research framework, selecting cases, and conducting fieldwork. Finally, Dul and Hak (2008) propose a three-phase approach: preparation, research, and reporting. Each of these methodologies offers different levels of structure and flexibility, allowing researchers to choose the approach that best fits the needs of their study

This thesis uses Merriam's (1998) approach to case study research because it offers clear steps but also enough flexibility to work with complex topics like AI in healthcare. Merriam's method starts with looking at what has already been written on the topic to understand the background and identify what is missing. Next, a framework is created to guide the study and help make sense of the findings. The research questions are then adjusted as new ideas and patterns appear during the process. Cases are chosen on purpose, based on how much useful and detailed information they can provide. Information is gathered from different sources, such as documents so that the findings are based on rich, detailed evidence. The analysis looks for patterns that emerge from the data and also compares the cases to find similarities and differences. To make sure the results are trustworthy, Merriam recommends checking the information against more than one source (triangulation), asking other experts to review the work, and describing the cases in enough detail for readers to understand them fully. The findings are then explained in a way that helps readers gain a deeper understanding of the topic.

In this thesis, Merriam's method was adjusted to focus only on secondary sources, such as court decisions, laws, and legal commentary. The research began with a review of existing studies and the creation of a framework linking AI in healthcare with legal and ethical issues. The research questions were refined as themes emerged. Cases were chosen from the Canadian Legal Information Institute (CanLII) and other legal databases using clear criteria to make sure they were relevant and full of useful information. Data was collected through document analysis, and steps like triangulation, peer review, and careful documentation were used to keep the results accurate and reliable. This approach made it possible to explore in detail how Canadian courts are handling the legal challenges of AI in healthcare.

Following these guidelines, this thesis adheres to Merriam's (1998) systematic process for designing qualitative research, while also incorporating the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) Checklist (Appendix B.1) to support the case synthesis.

4.3.1 Search strategy

Sampling in qualitative research differs from statistical sampling in that it does not focus on selecting a random, representative sample. Instead, theoretical sampling prioritizes selecting cases that offer rich, informative data, which can contribute to or challenge developing theories (Eisenhardt, 1989; Gerring, 2007). In this study, the aim was to identify Canadian legal cases that shed light on the intersection of AI and healthcare, particularly those that address issues like transparency, accountability, and human-AI interaction. As Patton (1999) and Crabtree et al. (1995) point out, the quality and depth of the data are more important than the sheer number of cases. By purposefully selecting cases that provide detailed insights, this approach allows for a deeper understanding of the legal frameworks surrounding AI in healthcare.

To support this research, a comprehensive search strategy was developed to capture both Canadian legal cases and scholarly literature that discuss the integration of AI in healthcare. A key part of this strategy involved clearly defining important concepts such as “human-AI interaction,” “transparency,” and “explainability.” For example, human-AI interaction was defined as the mutual process by which humans and AI systems exchange information and influence decision-making in clinical settings. This understanding of transparency, where AI processes are communicated in clear, accessible language for non-experts like patients and clinicians, shaped the development of the search terms used in this study (Eisenhardt, 1989; Gerring, 2007).

The search was conducted using CanLII, a primary Canadian legal database that provides access to all publicly available legal cases in Canada. A combination of keywords and Boolean operators was used to generate a broad yet focused set of results. Search terms included phrases such as “artificial intelligence,” OR “machine learning,” AND “healthcare,” OR “health,” OR “primary care,” OR “tertiary care,” OR “health care.” An iterative approach was used, where the initial search results were reviewed carefully, and search terms were

refined to improve the accuracy and relevance of the results. This iterative process was essential for ensuring that the search strategy effectively identified cases and literature that directly addressed the regulatory challenges and ethical issues surrounding AI in healthcare.

In addition to the automated search process, manual screening played a crucial role in the strategy. The results were filtered based on specific inclusion criteria, such as relevance to AI applications in healthcare, significant legal challenges (e.g., privacy or accountability concerns), and a focus on Canadian jurisdiction. This careful screening process helped address potential limitations of relying solely on CanLII, such as the risk of missing cases from less prominent jurisdictions or newly emerging cases. Documenting each search iteration and providing clear reasons for excluding certain search terms further strengthened the transparency and replicability of the methodology.

The comprehensive search strategy ultimately provided a valuable dataset for qualitative case study research, while also revealing gaps in the existing literature, particularly concerning how legal frameworks are adapting to the fast-paced evolution of AI technologies. While the focus on Canadian legal cases has certain limitations, this approach establishes a solid foundation for future research, which could expand to include international databases and emerging digital sources (Eisenhardt, 1989; Gerring, 2007). To ensure the robustness and appropriateness of the search strategy and methods, a librarian from Western Health was consulted for guidance in planning, conducting, and reporting the research.

An overview of the search strategy for CanLII is provided in Appendix B.2, which includes a table of key concepts and an example of the search process. This documentation was used to pilot the search and refine the strategy, ensuring that all relevant articles were included. The retrieved cases were then reviewed and evaluated directly on CanLII.

4.3.1.1 Limitations of Canadian Legal Information Institute

CanLII is a crucial tool for accessing Canadian case law, statutes, and other legal materials. However, it has some important limitations that should be considered when conducting a case synthesis. While CanLII covers a wide range of Canadian case law, it may not include every case from all courts or jurisdictions. For example, decisions from certain lower courts or specialized courts might be missing. Additionally, some cases may be excluded due to

privacy concerns, legal restrictions, or incomplete submissions from courts. Unpublished decisions, including those not considered to have precedential value or those where publication was withheld for other reasons, are also not included in CanLII's database. Although CanLII is an invaluable resource for free access to Canadian case law, these limitations must be kept in mind when selecting cases for research. To fill these gaps and ensure a thorough review, CanLII's database was supplemented with more detailed, specialized, or up-to-date information from other paid legal databases or secondary resources as needed.

4.3.2 Case selection process

References were managed using Excel to streamline the process of removing duplicates, screening litigation histories, assessing relevant cases, and extracting necessary data. To visually track the screening process, the PRISMA flow diagram was used. The first step involved eliminating duplicate records from the results in Excel. Next, a pilot screening was conducted by two reviewers who worked together to ensure that the selected litigation histories aligned with the predefined inclusion and exclusion criteria. After reaching consensus, the pilot phase was completed. Following this, two reviewers independently screened the litigation histories of the remaining cases to identify those that met the inclusion criteria and exclude those that did not. Any disagreements between the two reviewers were resolved through discussion. Cases with insufficient information in the litigation history were moved to the full case screening stage for further evaluation.

Following Merriam's (1998) approach to in-depth case study research, full-text screening was carried out to ensure that only cases offering detailed, valuable insights into the integration of AI in healthcare were selected. Two reviewers independently examined the full texts of each case to determine if they met the inclusion criteria or should be excluded based on predefined exclusion criteria. This independent review process is crucial for ensuring internal validity through triangulation and peer review, as recommended by Merriam (1998). When disagreements occurred between the reviewers, they were resolved through discussion, and, if needed, a third reviewer was consulted to reach an agreement. Cases that seemed to lack sufficient information were not immediately excluded; instead, they were kept for further data extraction and analysis to ensure that no potentially valuable insights were missed.

Following the screening process, data extraction was carried out independently by two reviewers using an Excel spreadsheet specifically designed for this thesis. The spreadsheet was structured to capture key characteristics relevant to our research questions and objectives, ensuring that all extracted data aligned closely with the thesis focus on legal challenges, transparency, and human-AI interaction in healthcare. This method is in accordance with the principles of qualitative case study research as outlined by Merriam (1998), emphasizing the need for "thick descriptions", as described as detailed, context-rich explanations, and comprehensive, context-rich data. Any disagreements between reviewers during data extraction were resolved through discussion or by involving a third reviewer, further enhancing the reliability of the findings. Cases that did not contain the necessary information were excluded from further analysis to maintain the integrity of the dataset. The detailed process of data extraction and the resulting extraction table are documented in Appendix B.2, which serves as a transparent record of how the data was synthesized to support robust, theory-driven insights.

4.4 Results

The case selection process for this study was designed to ensure both thoroughness and transparency, in line with the qualitative case study methodology recommended by Merriam (1998). All identified cases were first screened for duplicates and then evaluated against predefined inclusion and exclusion criteria. The screening was carried out in two stages: litigation history screening and full-text screening by two independent reviewers. In the initial pilot phase, the reviewers worked together to confirm the clarity and appropriateness of the criteria. Once consensus was reached, the reviewers proceeded independently to review the remaining cases. Disagreements were resolved through discussion. This two-reviewer process was intended to strengthen internal validity by reducing bias and ensuring that decisions were not based on a single perspective. The same two-reviewer approach was applied to the data extraction stage, using a structured Excel spreadsheet to capture case characteristics directly relevant to the research questions. Any discrepancies in extracted data were again discussed. This careful, iterative process ensured that the final seven cases that met all eligibility criteria was both information-rich and methodologically robust, providing a strong foundation for the in-depth qualitative analysis that follows.

4.4.1 Case Selection Criteria

A total of 89 cases were initially captured through our comprehensive search strategy. Of these, 88 cases were unique and met the eligibility criteria for litigation history screening based on the predetermined inclusion and exclusion guidelines. Following a detailed evaluation process, guided by the principles of theoretical sampling and case study research as described by Merriam (1998), only 7 cases ultimately met the final inclusion criteria. This rigorous selection process ensured that the chosen cases were particularly rich in information and highly relevant to the investigation of AI integration in healthcare.

These 7 cases, which date back to 2020 at the earliest, have been systematically organized according to four key data extraction categories. These categories were developed to align with our methodological framework, ensuring that each case contributes to a comprehensive understanding of the legal challenges associated with AI in healthcare. The process reflects a deliberate and descriptive approach to case study research, emphasizing in-depth analysis and the use of "detailed, context-rich explanations to capture the complexities of the phenomenon (Merriam, 1998; Eisenhardt, 1989). A detailed list of these 7 references is provided in Appendix B.4, and the screening process is illustrated in Figure 4.1.

By selecting these cases through a careful, iterative evaluation, we have not only ensured the internal validity of our findings through techniques like triangulation and peer review but also addressed external validity by focusing on information-rich examples. This approach has allowed us to construct a robust dataset that supports a nuanced analysis of how AI is reshaping legal frameworks in healthcare, and it lays a solid foundation for the insights and recommendations that follow.

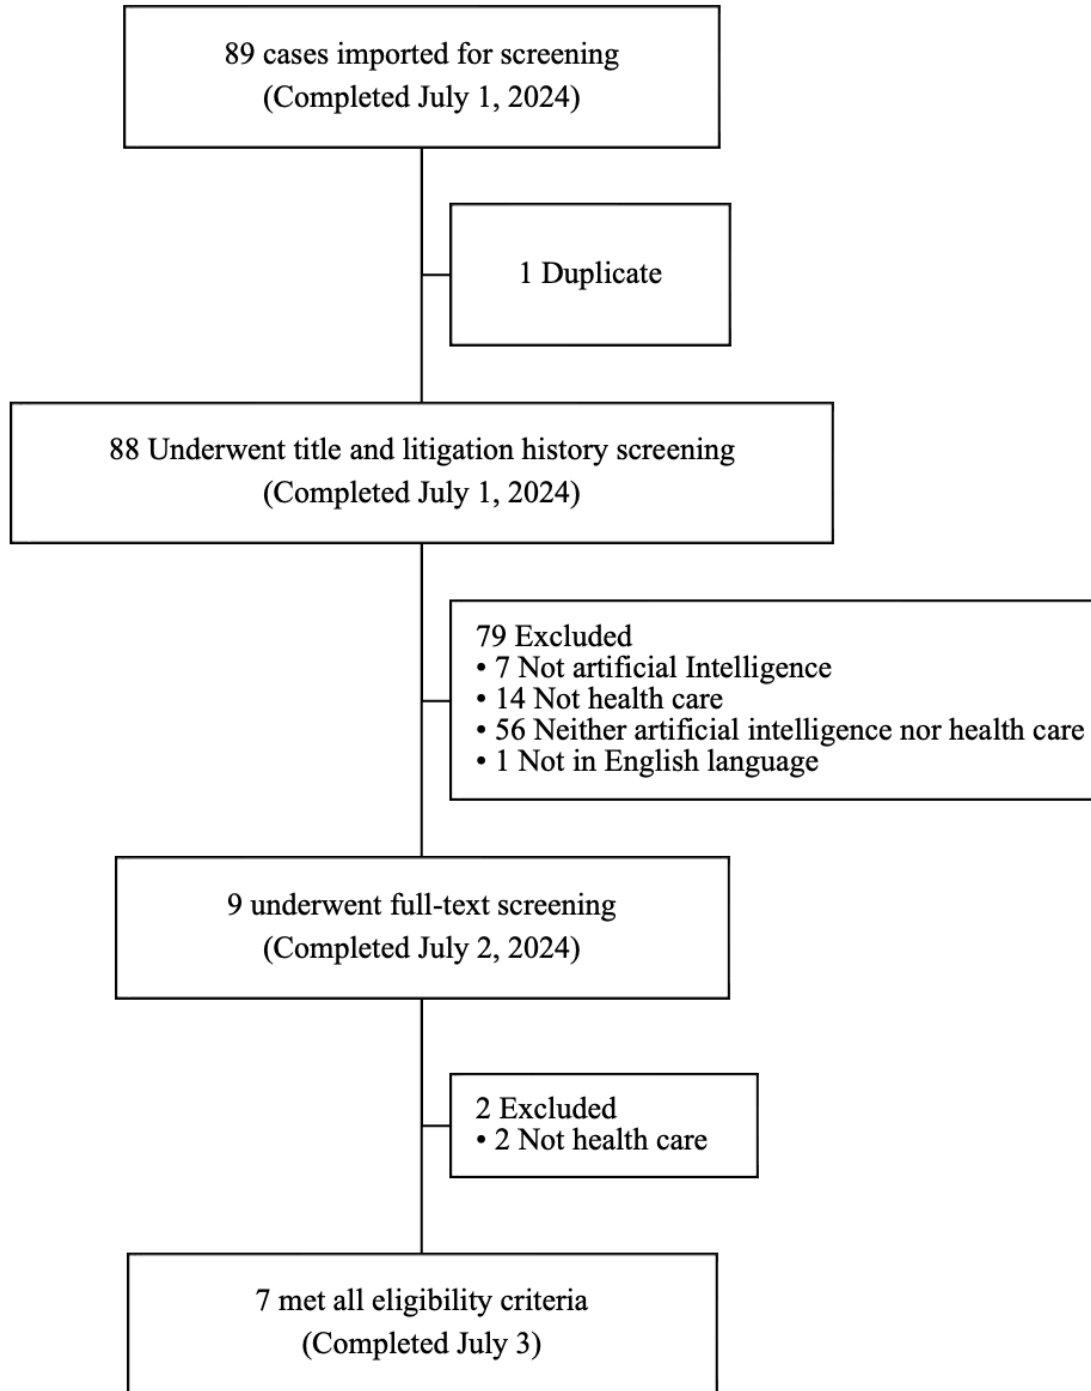


Figure 4.1: PRISMA flow diagram.

The majority of the cases (6 cases, or 85.7%) were resolved in Ontario, and therefore, were governed by the province's civil law. The remaining case (1 case, or 14.3%) was resolved in British Columbia, and was subject to the civil law of that province.




Figure 4.2: Jurisdictions of cases.

The court cases were sourced from the Supreme Court of British Columbia (1 case, or 14.3%) and the Superior Court of Justice in Ontario (2 cases, or 28.6%). Tribunal cases were drawn from Ontario-based tribunals, including the Information and Privacy Commissioner of Ontario (2 cases, or 28.6%) and the Health Professional Appeal and Review Board of Ontario (2 cases, or 28.6%).

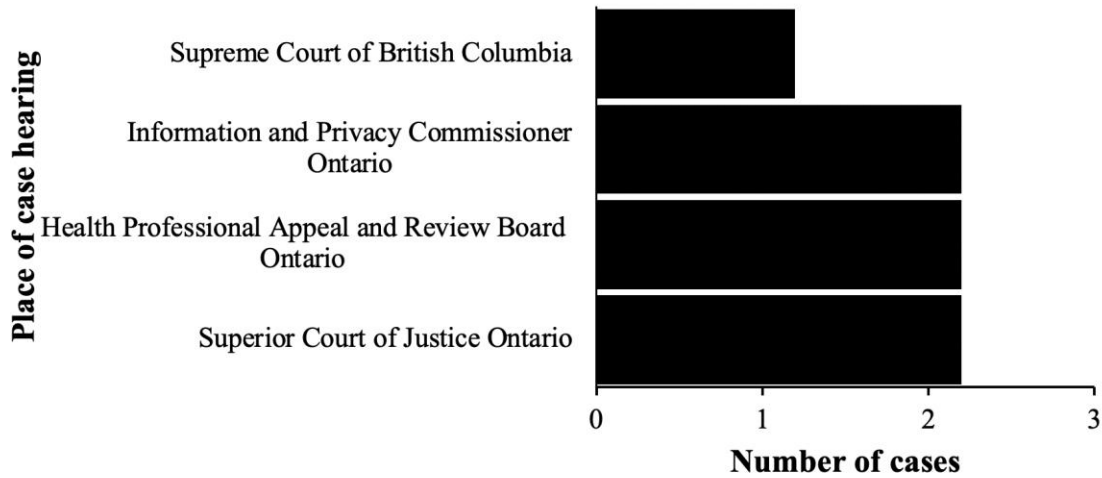


Figure 4.3: Place of case hearing.

The dates of the cases represent the years in which they occurred. The earliest cases involving AI and healthcare were from 2020 (2 cases, or 28.6%). Additional cases were identified in 2022 (1 case, or 14.3%), 2023 (2 cases, or 28.6%), and 2024 (2 cases, or 28.6%).

Table 4.1: Case dates in years.

Date	Number of cases (%)
2020	2 (28.6)
2022	1 (14.3)
2023	2 (28.6)
2024	2 (28.6)

Six key pieces of legislation were identified in the cases that met the inclusion criteria. These legislations, which were cited in the cases, include: the Freedom of Information and Protection of Privacy Act (FIPPA) (2 cases, or 28.6%), the Personal Information Protection and Electronic Documents Act (PIPEDA) (1 case, or 14.3%), the Personal Health Information Protection Act (PHIPA) (1 case, or 14.3%), the Competition Act (1 case, or 14.3%), Bill C-27 (1 case, or 14.3%), and the Regulated Health Professions Act (RHPA) (1 case, or 14.3%).

Table 4.2: Legislation used in cases.

Legislation	Number of cases
-------------	-----------------

FIPPA	2
PIPEDA	1
PHIPA	1
Competition Act	1
Bill C-27	1
RHPA	1

The majority of cases (4 cases, or 57.1%) were decided in favor of the plaintiff or applicant, meaning the committee decided to take further action based on the complaint. The remaining cases (3 cases, or 42.9%) were decided in favor of the defendant or respondent, meaning the committee reviewed the complaint but chose not to take any further action.




Figure 4.4: Resolution type of cases.

The most frequently occurring themes across the cases were privacy (2 cases, or 28.6%), discipline (2 cases, or 28.6%), and practice and procedure (2 cases, or 28.6%). Other notable themes included labor and employment (1 case, or 14.3%), contracts (1 case, or 14.3%), and torts (1 case, or 14.3%).

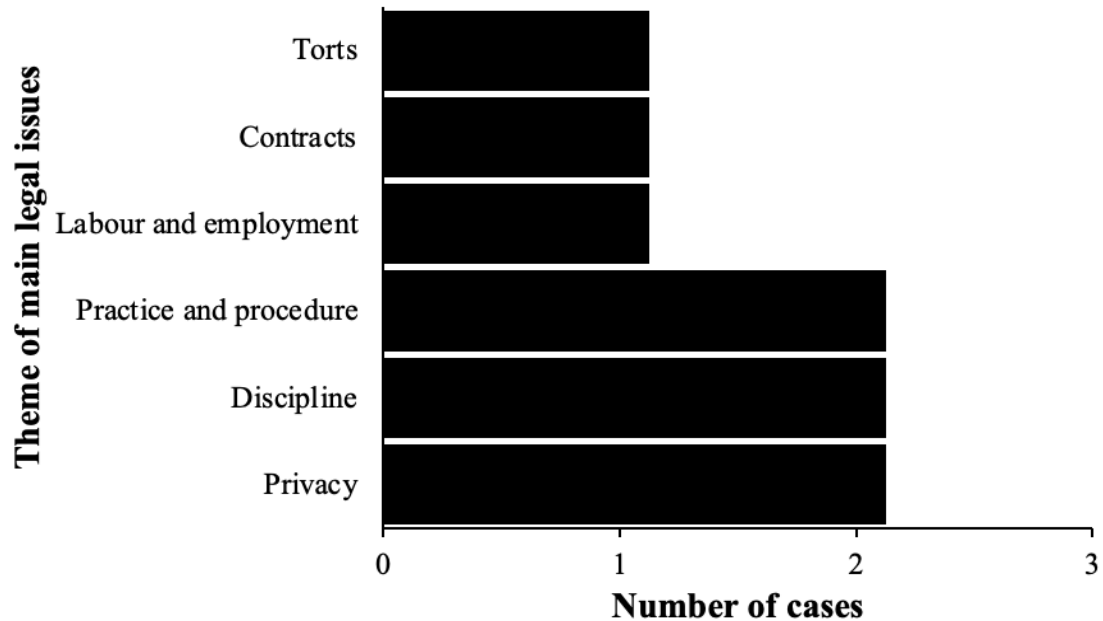


Figure 4.5: Theme of main legal issues.

4.4.2 Final Case Set

The final set of cases analyzed in this study was selected based on their direct relevance to AI or digital decision-support tools in health or health-adjacent contexts in Canada. The set includes two cases decided by the Health Professions Appeal and Review Board (HPARB) and one case decided by the Supreme Court of British Columbia (BCSC). The three cases, *C. J. v. M.W.* (2020 CanLII 57706 (ON HPARB)), *Doobay v. Cohen* (2023 ON HPARB 111813), and *Lam v. Flo Health Inc.* (2024 BCSC 391), included demonstrate intersections between AI, data governance, and professional practice, providing insights into transparency, consent, and the management of personal health information.

In *C. J. v. M.W.* (2020), an occupational therapist used BrainFx, a software tool with AI-driven analytics, to assess an employee's cognitive function. The case illustrates how Canadian law approaches consent, documentation, and disclosure when AI-assisted assessments are used, emphasizing the importance of ongoing consent and transparency in interpreting AI-generated results.

Doobay v. Cohen (2023) involved a professional conduct complaint that included concerns about the use of AI and data-driven decision-making in a mental health and social services

setting. This case demonstrates how regulatory bodies evaluate information handling, transparency, and fairness in contexts where AI tools are used, even if AI is not the main issue, highlighting the overlap between professional governance and AI-enabled practices.

Lam v. Flo Health Inc. (2024) concerned the collection, processing, and sharing of sensitive reproductive health data through an AI-driven mobile application. The class action focused on potential violations of privacy laws and contractual obligations related to personal health information. The case underscores the challenges of system-level transparency and consent in AI-mediated health applications, showing how courts assess compliance with privacy regulations, the clarity of consent processes, and the accountability of digital health platforms managing large-scale personal data.

4.4.3 Limitations of the Approach

The case study approach gives detailed, qualitative insights, but it has some important limitations. There are very few Canadian cases that directly involve AI in healthcare. This thesis therefore focuses on cases where AI or digital decision-support tools were either central or closely involved. Because of this, the findings may not apply broadly to all situations. Some cases, like *Doobay v. Cohen* (2023), only involve AI in a minor way. Lessons about health-AI governance are drawn from broader rules about professional regulation, privacy, and transparency, rather than from rulings that specifically deal with AI. AI tools and health applications are changing quickly. The cases examined here reflect the legal situation up to 2024, but this may shift as new technologies appear and courts address AI-related disputes more directly. This thesis does not cover liability, malpractice, or wider ethical debates beyond transparency and data governance. Limiting the focus this way helps narrow the analysis but means other legal or ethical issues are not addressed. Court decisions summarize reasoning but may not include full technical details about the AI systems. This limits the ability to evaluate technical accuracy or algorithmic fairness directly. Despite these limitations, the selected cases provide useful insight into how Canadian law interacts with AI tools in healthcare, particularly regarding transparency, data use, consent, and regulation.

Chapter 5

5 C. J. v. M.W. (2020)

Chapter 4 describes the methodology for the thesis. It explains the qualitative case study approach used to analyze Canadian legal cases involving AI in healthcare, why that method was chosen, and how it was applied. It covers the search strategy, the criteria for case selection, the data extraction process, and the measures taken to ensure rigor, such as triangulation, peer review, and detailed documentation. The chapter also outlines the results of the search process, including how many cases were initially identified, how many were excluded, and the final set of cases selected for in-depth analysis. Overall, Chapter 4 is about how the research was systematically designed and conducted to gather valid, detailed, and context-rich information about AI-related legal issues in Canadian healthcare.

The goal of Chapter 5 is to apply that methodology to a detailed case study analysis. It takes the cases selected in Chapter 4 and examines them in depth, focusing on how Canadian courts and tribunals have addressed legal questions arising from AI in healthcare. This includes analyzing judicial reasoning, the statutory frameworks applied, and how these cases compare to each other in their treatment of issues like privacy, informed consent, professional conduct, and accountability. The chapter's aim is to uncover recurring patterns, legal themes, and interpretive approaches, as well as to highlight where the law aligns or fails to align with the practical and ethical realities of AI in healthcare. By doing so, Chapter 5 builds a bridge between abstract legal principles and their real-world application in Canadian jurisprudence.

5.1 Background

As explained in Chapters 2 and 3, how people interact with AI systems and how those systems explain their outputs are major governance issues in healthcare. These concerns become especially important when decisions about individuals are based on complex digital tools that are difficult to understand. In such situations, people are still expected to understand the results, agree to their use, and make decisions based on them, even though the technology itself may be opaque or proprietary. This chapter uses *C. J. v. M.W. (2020)* as a case study to show how these governance issues appear in real legal practice. The decision

does not directly mention artificial intelligence. However, it involves a computerized assessment tool whose results influenced professional judgment and affected what information was shared. For this reason, the analysis looks closely at how the adjudicative body dealt with questions about whether the assessment was understandable, how data were managed, and whether consent was meaningful when a complex technological system was used. Rather than analyzing how the tool was technically designed, the chapter focuses on the court's reasoning. This approach shows that existing legal frameworks already address concerns that closely relate to AI governance in healthcare, even when AI is not explicitly named.

5.2 Introduction

This chapter looks at *C. J. v. M.W.* (2020), a decision of the HPARB. The case is used to show how Canadian decision makers deal with complex digital tools used in health-related professional practice. The case does not directly involve artificial intelligence. However, it raises many of the same governance problems that appear in healthcare AI. These problems include uncertainty about how digital tools work, how their results are interpreted, how consent is obtained and maintained, and how personal information is shared across different organizations. As explained in Chapter 2, AI systems in healthcare are usually designed to support professional judgment rather than replace it. These systems influence decisions by producing scores, categories, or assessments based on data. They often rely on proprietary technology, meaning their internal logic and datasets are not fully visible to patients or even to professionals (Shneiderman, 2022). Because of this, governance concerns go beyond whether the technology is technically accurate. They also include whether people can understand how the technology is being used, how information moves through institutions, and how personal health information is managed over time. *C. J. v. M.W.* (2020) provides a concrete example of how Canadian decision makers address these issues in a real professional setting.

This chapter focuses on three related governance themes discussed in Chapters 2 and 3. These are transparency in assessments that rely on technology, governance of data and information flows, and informed consent understood as an ongoing process rather than a single event. By examining the Board's reasoning in detail, the chapter shows that existing

legal and professional frameworks already address issues that are central to health AI governance, even when artificial intelligence is not explicitly mentioned or regulated.

5.3 Case Background

In *C. J. v. M.W.* (2020), the case began with a complaint by C.J. about the conduct of an occupational therapist, M.W. The therapist carried out a workplace accommodation assessment using BrainFx, a computerized cognitive assessment tool. BrainFx is a proprietary digital system that measures cognitive functioning by comparing an individual's performance to a reference dataset. The results of the assessment were later shared with a third party, specifically the complainant's employer, as part of a broader workplace accommodation process.

C.J. raised several concerns about how the assessment was handled. These concerns included claims that the assessment was not properly explained that consent was not meaningful or ongoing, that personal information was shared inappropriately, and that the therapist went beyond professional boundaries. The complaint was first reviewed by the Inquiries, Complaints and Reports Committee of the relevant regulatory college. That committee chose to issue guidance rather than impose discipline. C.J. appealed that decision to HPARB.

The Board was not asked to decide whether BrainFx was scientifically valid or clinically accurate. Instead, the dispute focused on how information produced by a complex digital tool was explained, interpreted, documented, and shared in professional practice. This focus reflects the framework discussed in Chapter 3. That framework emphasizes that many governance problems in health AI arise not from the technical design of algorithms, but from how information is managed around their use, including consent, interpretation, and institutional data flows (Gerke et al., 2020).

5.4 Analysis of Legal Reasoning

5.4.1 Transparency and understanding digital assessments

One of the main issues before the Board was whether the occupational therapist gave C.J. enough information to understand the BrainFx assessment. The Board looked at whether C.J. was told what the tool measured, why it was being used, how the results would be

interpreted, and how those results fit into the overall evaluation process (*C. J. v. M.W.*, 2020). The Board did not require the therapist to explain the technical details of how BrainFx worked or to reveal proprietary information about its algorithms or datasets. Instead, the key question was whether the assessment was understandable in practical terms. This approach is consistent with academic discussions about explainability in health AI. These discussions argue that explanations should be meaningful for the person receiving them and appropriate to the context, rather than technically detailed or exhaustive (Doshi-Velez & Kim, 2017; Lipton, 2018). As discussed in Chapter 2, many health AI systems operate as “black boxes” in practice. Their internal logic is often inaccessible, yet their outputs can strongly influence decisions. The Board’s reasoning suggests that Canadian governance expectations may focus on functional transparency. This means ensuring that people understand what a tool does and how its results affect decisions, even if they cannot examine how the tool works internally. This point is directly relevant to healthcare AI, where trust and accountability depend on clear explanations of how algorithmic outputs are used in professional judgment (Holzinger et al., 2017).

5.4.2 Data governance and information sharing

The Board also considered whether it was appropriate to share the BrainFx assessment results with a third party. In reviewing this issue, HPARB examined whether the disclosure was covered by the consent that had been given and whether it complied with professional and privacy obligations under the relevant regulatory frameworks (*C. J. v. M.W.*, 2020). Rather than treating disclosure as a narrow legal compliance issue, the Board took a broader view of data governance. It looked at how information was created, recorded, interpreted, and shared, as well as how the limits of the assessment were communicated. This reflects an understanding of data governance as covering the entire lifecycle of information, from collection to interpretation to disclosure. This approach aligns with the discussion in Chapter 2, which emphasizes that AI-generated information often moves across organizational boundaries and can influence decisions far beyond the original context in which data was collected. The case shows how governance problems can arise when outputs from digital tools are treated as authoritative without adequate explanation or context. For health AI, this highlights the importance of being able to track how data and algorithmic outputs move through systems and shape later decisions (Gerke et al., 2020; Murdoch et al., 2022).

5.4.3 Consent as a continuous process

Consent was another central issue in *C. J. v. M.W.* (2020). C.J. argued that consent for using the BrainFx assessment and for sharing its results was inadequate. The Board treated consent not as a single moment or form, but as an ongoing process. This process requires providing relevant information at appropriate times as circumstances change (*C. J. v. M.W.*, 2020). This view of consent is consistent with broader ethical and legal thinking in healthcare, which emphasizes that informed consent depends on ongoing communication rather than one-time disclosures (O'Neill, 2003). This point is especially important in the context of health AI, where systems may be introduced gradually and may produce new types of information over time. As discussed in Chapter 3, existing scholarship highlights how difficult it is to achieve meaningful consent for complex digital systems through standard consent forms or generic privacy notices. The Board's reasoning suggests that Canadian decision makers may expect consent practices to evolve as technologies are used and as their implications become clearer. This insight is directly relevant to AI governance in healthcare, where continued explanation and engagement may be necessary to maintain trust and legitimacy (Gerke et al., 2020; Murdoch et al., 2022).

5.5 Implications for Health AI Governance

The reasoning in *C. J. v. M.W.* (2020) provides several important lessons for health AI governance in Canada. First, transparency is understood as practical understanding rather than full technical disclosure. Health AI systems should be governed in ways that allow patients and users to understand how outputs are generated and how they affect decisions, even when the systems are complex or proprietary. Second, data governance is seen as covering the entire path of information, including interpretation and sharing. For health AI, this means governance frameworks must address not only privacy and security, but also how algorithmic outputs are explained and incorporated into professional judgment. Third, consent is treated as an ongoing responsibility. This is especially important for AI systems that produce continuous assessments, predictions, or recommendations. Consent practices must be flexible and responsive to changing information uses rather than relying on one-time approval. Therefore, these points support the argument made in Chapter 1 that legal decisions that do not explicitly focus on AI can still offer valuable insight into how Canadian law is

likely to address governance challenges at the intersection of healthcare and artificial intelligence.

Chapter 6

6 Doobay v. Cohen (2023)

Chapter 6 aims to build on the earlier case study by closely examining a second Canadian legal decision, using the same qualitative framework introduced in Chapter 4. The purpose is to apply the same method to a different legal and factual situation so that the cases can be meaningfully compared. Rather than restating the methodology, the chapter concentrates on how a different court or tribunal reasoned through issues that matter for health AI governance. These include how sensitive information is handled, how complex data practices are understood, and how human judgment operates when decisions are shaped by technological systems.

The chapter focuses in particular on how legal reasoning deals with fairness, transparency, and consent in settings where decisions depend on advanced information systems. Even though the case does not directly involve healthcare AI, the legal principles applied closely resemble those that arise in AI-supported healthcare settings. The analysis looks at which statutory frameworks are involved, how technical and evidentiary material is assessed, and how the decision-maker balances institutional procedures against individual rights.

By examining a second case in depth, Chapter 6 strengthens the analysis of how Canadian legal decision-makers respond to governance challenges posed by complex systems that are not fully transparent. The chapter supports the overall argument of the thesis by identifying both similarities and differences when compared with the case discussed in Chapter 5. This comparison helps establish broader patterns across cases and reinforces the conclusion that governance issues relevant to health AI are not unique to a single context, but appear repeatedly across different areas of law.

6.1 Background

As discussed in Chapters 2 and 3, problems about who controls data, how information is made visible, and whether people have given meaningful consent do not arise only when clinicians make decisions about patients. These same problems also appear in regulatory and institutional settings, where digital information systems shape how professionals are

monitored and held accountable. In these contexts, decisions are often made on the basis of complex records, documents, and data flows rather than face-to-face interactions.

This chapter uses *Doobay v. Cohen* (2023) as a case study to show how a decision-making body evaluates these kinds of information practices. The case illustrates how adjudicators reason through complicated informational arrangements when deciding whether a process was adequate and fair. Even though the case does not involve artificial intelligence directly, it raises governance issues that closely resemble those associated with AI-enabled systems. In particular, it deals with how digital information is handled, how it is interpreted, and how it is communicated to those affected by decisions.

The analysis concentrates on how the court examined the movement and presentation of information within the decision-making process. It looks at whether the information was clear, consistent, and understandable, and whether it supported fair reasoning. By focusing on this assessment, the chapter shows what the court implicitly expects in terms of transparency and procedural fairness in health-related governance, even when advanced technologies like AI are not formally part of the case.

6.2 Introduction

This chapter examines *Doobay v. Cohen* (2023), a decision of the HPARB, as a second case study on governance problems that arise when professional decisions depend on complex information systems. Like *C. J. v. M.W.* (2020) discussed in Chapter 5, this case does not involve artificial intelligence in a technical sense. There is no algorithm or automated system under review. However, the case raises closely related issues about how information is handled, recorded, interpreted, and explained. These are the same issues that commonly arise in healthcare settings where AI-supported tools are used. As explained in Chapter 2, health AI systems operate within broader information environments. These include electronic records, emails, institutional databases, and administrative processes. Governance problems in these settings often do not come from algorithms alone. They arise from how information is collected, interpreted, documented, and evaluated by human professionals. *Doobay v. Cohen* (2023) provides a useful example of how these information practices are assessed, especially in relation to professional discretion and the limits of regulatory oversight. This chapter focuses on three governance dimensions developed in Chapters 2 and 3. transparency

in professional reasoning, documentation and information management, and the interpretation of complex information. By examining the Board's reasoning, the chapter shows how Canadian legal institutions assess information practices that are increasingly common in digitally mediated healthcare, including settings where AI-supported systems are involved.

6.3 Case Background

Doobay v. Cohen (2023) arose from a complaint filed by Kathleen Doobay against Rodney Victor Cohen, a registered psychotherapist regulated by the College of Registered Psychotherapists and Registered Mental Health Therapists of Ontario. The complaint alleged several forms of misconduct, including failures to follow diversity, equity, and inclusion principles, coercive behavior, conflicts of interest, psychological harm, and misuse of confidential information.

The Inquiries, Complaints and Reports Committee reviewed the complaint, the practitioner's responses, and the applicable professional standards. The Committee concluded that there was not enough evidence to support findings of professional misconduct. It emphasized Cohen's experience working with diverse communities and found no reliable evidence of coercion, conflicts of interest, or misuse of confidential information. Doobay requested a review by HPARB, arguing that the investigation failed to properly address ethical concerns and problems related to how information was handled.

HPARB upheld the Committee's decision. It found that the investigation was thorough and that the conclusions were reasonable. At the same time, the case exposed broader governance challenges. This included difficulty assessing professional conduct outside formal client relationships, limited regulatory tools for addressing systemic or equity-related concerns, and weak mechanisms for handling privacy-related complaints. These issues are directly relevant to health AI governance, where decision-making often takes place in informal or data-mediated contexts that do not fit neatly within existing regulatory frameworks.

6.4 Analysis of Legal Reasoning

6.4.1 Transparency in professional reasoning

A key issue in the Board's analysis was whether the regulator's decision was sufficiently transparent. HPARB considered whether the Committee's reasons showed that it had genuinely engaged with the evidence and provided a clear explanation for its conclusions (*Doobay v. Cohen*, 2023). The Board explained that transparency in professional regulation does not mean responding in detail to every allegation or ethical concern raised by a complainant. Instead, transparency requires that the reasoning process can be understood and followed. A decision must show how the evidence was assessed and why the conclusions were reached. This approach matches broader discussions in health AI governance, where transparency is often understood as the ability to understand how a decision was made, rather than as full disclosure of technical systems or institutional processes (Lipton, 2018). As discussed in Chapter 5, Canadian adjudicative bodies generally take a practical approach to transparency. Applied to AI-enabled healthcare, this suggests that governance frameworks may focus on clearly explaining how AI outputs are used in professional decision-making, rather than requiring full technical explanations of how AI systems function internally. The Doobay decision reflects this approach by emphasizing the clarity of professional reasoning within the regulator's mandate, rather than engaging with broader ethical debates that fall outside its authority.

6.4.2 Documentation and information sharing

Documentation was central to HPARB's assessment of whether the investigation was adequate. The Board reviewed the records relied on by the Committee, including correspondence, clinical materials, and investigative documents, to determine whether they provided a sufficient basis for evaluating the practitioner's conduct (*Doobay v. Cohen*, 2023). This focus on documentation is highly relevant to health AI governance. As discussed in Chapter 2, AI-supported healthcare systems generate large volumes of data that must be integrated into clinical and regulatory records. Governance problems arise when it is unclear how data are interpreted or how they influence professional judgment. The Board's reasoning highlights documentation as a core governance mechanism. Proper records allow oversight bodies to reconstruct how decisions were made and to assess whether professional standards

were followed. In AI-enabled healthcare, this implies that governance should not focus only on data accuracy or system performance. It should also require clear documentation showing how AI-generated information is interpreted and applied by professionals (Gerke et al., 2020). Without this, meaningful oversight and public trust become difficult to sustain.

6.4.3 Interpretation of complex information

The case also shows how Canadian regulatory bodies handle interpretive judgment. The complainant disagreed with the practitioner's interpretation of information relating to clinical interactions and ethical duties. HPARB recognized that healthcare decisions often involve complex, incomplete, or ambiguous information, especially in situations that fall outside formal therapeutic relationships. Rather than replacing the practitioner's judgment with its own, the Board asked whether the practitioner's interpretations fell within a reasonable range, given professional standards and the available evidence. This reflects an acceptance that uncertainty is an unavoidable part of healthcare decision-making. As discussed in Chapter 3, AI systems do not remove interpretive uncertainty. Instead, they often add new layers of complexity through probabilities, risk scores, and predictive outputs. The reasoning in *Doobay v. Cohen* (2023) suggests that Canadian legal institutions may assess AI-assisted decisions by focusing on how professionals interpret and use complex information, rather than by examining the technology in isolation.

6.5 Implications for Health AI Governance

The decision in *Doobay v. Cohen* (2023) offers several lessons for health AI governance. First, transparency is understood as clarity in professional and institutional reasoning, not as technical disclosure. In AI-enabled healthcare, this suggests that governance frameworks should prioritize explaining how AI outputs influence professional decisions. Second, documentation is shown to be a critical governance tool. When AI systems affect healthcare decisions, records should clearly explain how data and system outputs are used in professional judgment. This is essential for accountability, regulatory oversight, and trust. Third, the Board's treatment of interpretive discretion confirms the continuing importance of human judgment in complex decision-making environments. Health AI systems should therefore be governed in ways that support professional interpretation and contextual understanding, rather than replacing or obscuring them.

Comparative scholarship reinforces these conclusions. Rueda et al. (2022) stress the importance of procedural fairness and explainability in high-stakes health decisions. Vogt (2025) raises similar concerns about bias and equity in algorithmic healthcare systems. Boudreau LeBlanc (2022) further observes that Canada's complaint-based regulatory model may struggle to address systemic or forward-looking governance challenges associated with digital health and AI. Taken together, these perspectives highlight the need for governance approaches that address both human judgment and data-driven systems within existing legal structures.

Chapter 7

7 Lam v. Flo Health Inc (2024)

The purpose of Chapter 7 is to complete the detailed case study analysis by applying the same qualitative framework to a third Canadian legal decision. Following the methodological approach set out in Chapter 4, the chapter examines how another court or tribunal dealt with issues that are directly relevant to the governance of artificial intelligence in healthcare. These issues include how decision-makers understand complex information systems, how data is managed and governed, and how institutional decision-making affects individual rights.

Chapter 7 looks closely at how legal reasoning operates when information practices are technically complex, difficult to fully explain, or shaped by automated or data-intensive processes. The analysis focuses on how the decision-maker evaluated whether affected individuals received adequate explanations, how sensitive or personal data was handled, and how human judgment was used to interpret outputs produced by complex systems. Although the case does not explicitly involve healthcare AI, the legal problems it addresses closely resemble those that arise when AI tools are used in healthcare environments.

By adding a third case, Chapter 7 strengthens the overall analysis by allowing comparisons across different legal settings. The chapter examines where judicial reasoning is consistent and where it differs, particularly in relation to transparency, consent, and data governance. This additional case deepens the empirical basis of the study and helps ensure that the findings are not based on a single example. As the final case study chapter, Chapter 7 prepares the ground for the cross-case thematic analysis in Chapter 8 by showing that the identified themes emerge across multiple independent decisions rather than from one isolated case.

7.1 Background

Building on the ideas set out in Chapters 2 and 3, this chapter uses *Lam v. Flo Health Inc.* (2024) as a case study to examine how digital health data is governed in consumer-facing technologies. Unlike the earlier cases, this decision focuses directly on how sensitive health

information is collected, used, and shared through a digital platform that depends on data-driven analysis. The case offers a clear example of how courts think about consent, transparency, and whether users can reasonably understand what happens to their health data when it moves through complex technological systems. By closely examining the court's reasoning, this chapter shows how existing legal rules are applied to everyday data practices that underpin AI-enabled health applications, even when the underlying algorithmic processes are not visible or easily understood by users.

7.2 Introduction

This chapter looks at *Lam v. Flo Health Inc.* (2024) decision of the Supreme Court of British Columbia, as an example of how courts deal with governance problems in data-heavy digital health platforms. Unlike the cases discussed in Chapters 5 and 6, this case does not involve hospitals, clinicians, or traditional healthcare settings. Instead, it focuses on a consumer mobile app that people use on their own, outside the clinical system. Even so, the app processes large amounts of highly sensitive health information. This makes the case especially useful for understanding governance issues where health data, automated processing, and user-facing technologies intersect.

As explained in Chapter 2, many modern health-related AI systems are built into digital platforms. These platforms depend on ongoing data collection, algorithmic analysis, and data-sharing arrangements that are often difficult for users to see or understand. Common concerns include whether data practices are transparent, whether consent is meaningful, and whether users can realistically understand how their information is used. *Lam v. Flo Health Inc.* (2024) addresses these concerns directly and shows how Canadian courts assess information practices that closely resemble those used in health AI systems.

This chapter examines how the Court approached questions of transparency, consent, and data movement within a digital health platform. It focuses on how these issues are handled using existing legal concepts and how the Court's reasoning can help guide future approaches to governing health AI in Canada.

7.3 Case Background

The case began as a proposed privacy class action against Flo Health Inc., the company behind the Flo Period Tracker app. The app allows users to record menstrual cycles and other reproductive health information and then provides personalized feedback based on the data users enter. The plaintiffs claimed that Flo Health collected sensitive personal and health data and shared that data with third-party analytics and marketing companies, even though the company told users their information would remain private.

The plaintiffs brought several legal claims, including claims under the British Columbia Privacy Act, breach of confidence, intrusion upon seclusion, and unjust enrichment. The Court decided that the Privacy Act and breach of confidence claims were legally viable and could go forward. It rejected the intrusion upon seclusion claim because that tort is not recognized in British Columbia. The unjust enrichment claim was also dismissed because the required legal elements were not met. At the certification stage, the Court allowed the case to proceed in a more limited form, concentrating on the main privacy and confidentiality issues (*Lam v. Flo Health Inc.*, 2024).

Although the Court did not directly evaluate artificial intelligence models, the Flo app uses automated data processing and predictive features to produce personalized outputs. As discussed in Chapter 3, systems like this can be understood as AI-enabled technologies because they convert personal health data into algorithm-driven insights. For this reason, the governance issues raised in *Lam* closely resemble those raised by health AI systems, especially when it comes to how data is governed and how well users understand what is happening to their information.

7.4 Analysis of Legal Reasoning

7.4.1 Transparency of data practices

A key issue for the Court was whether Flo Health clearly explained its data practices to users. The Court looked closely at the privacy policy, terms of use, and how information was presented within the app itself. The question was whether an ordinary user would reasonably understand how their health data was collected and shared (*Lam v. Flo Health Inc.*, 2024).

The Court made clear that transparency is not achieved simply by having a privacy policy. What matters is whether the information is meaningful in practice. Users must be able to understand what data is being shared and what the consequences of that sharing might be. This reasoning reflects the discussion in Chapter 2, which explains how long, technical, or legalistic disclosures can hide important information instead of clarifying it. From the perspective of health AI governance, the decision shows that transparency is about user understanding, not just formal disclosure. For AI-enabled health platforms, this means governance frameworks should look not only at whether information is disclosed, but also at whether it is presented in a way that helps users make sense of complex data practices.

7.4.2 Consent and user understanding

The Court also examined whether users gave meaningful consent to the collection and sharing of their health data. Consent was treated as an ongoing informational process rather than a simple checkbox or contractual formality. The Court asked whether the way information was structured and explained actually allowed users to understand how their data would be processed and shared with third parties (*Lam v. Flo Health Inc.*, 2024). This approach is consistent with the analysis in Chapters 5 and 6, where courts raised concerns about the limits of consent in complex technological systems. In *Lam*, vague language and broad assurances were seen as weakening consent, especially given the highly sensitive nature of reproductive health data. These issues reflect wider debates in the health AI literature about how difficult it is to obtain informed consent for systems that rely on constant data collection and automated analysis (Gerke et al., 2020). The Court's focus on user understanding suggests that Canadian courts may increasingly question whether consent mechanisms in AI-enabled health technologies truly allow people to understand how their data is used.

7.4.3 Data flows and third party involvement

The Court also considered how user data moved beyond the app itself. The plaintiffs claimed that Flo Health shared data with external analytics companies in ways that users would not reasonably expect. The Court examined whether information about third-party involvement was clear enough for users to understand where their data might go after it left the app. This part of the decision reflects governance challenges discussed in Chapter 2. Many health AI

systems operate within complex data networks that involve multiple organizations. From the user's perspective, there may be only one visible platform, even though their data travels through a larger ecosystem of service providers and technical partners. For health AI governance, the Court's reasoning highlights the need to make data flows visible and understandable. Governance approaches that focus only on the main platform or data controller risk ignoring the wider infrastructure that shapes how health data is processed and reused.

7.5 Implications for Health AI Governance

The reasoning in *Lam v. Flo Health Inc.* (2024) offers several lessons that are directly relevant to governing AI in healthcare. First, the Court treats transparency as something centered on users, not as a box-ticking exercise. Health AI platforms need to explain data practices in ways that people can realistically understand, especially when dealing with sensitive health information. Second, consent is linked to the quality and context of the information provided. For AI-enabled health systems, this means consent processes need to change as systems evolve, rather than relying on static and generalized disclosures. Third, the Court's focus on third-party data sharing shows the importance of governing entire data ecosystems, not just individual technologies. Health AI systems are part of distributed networks, which makes it essential to address how data is used downstream and shared with others. Taken together, these points support the argument made in Chapter 1 that effective health AI governance depends on close attention to information practices, system design, and institutional context, rather than focusing narrowly on how well algorithms perform.

Chapter 8

8 Cross-Case Thematic Analysis

Chapters 5 to 7 analyzed three Canadian legal decisions. Each case arose in a different setting and involved different technologies, but all of them dealt with information practices that closely resemble those used in health-related artificial intelligence systems. This chapter brings those case analyses together to identify shared patterns in judicial reasoning that matter for how AI in healthcare is governed.

This chapter does not restate the facts or outcomes of the individual cases. Instead, it examines how courts and tribunals reasoned about issues such as transparency, data governance, consent, and the interpretation of complex systems. This approach directly addresses the research questions set out in Chapter 1 and builds on the conceptual framework developed in Chapters 2 and 3.

By looking across the cases, this chapter shows that Canadian legal reasoning already addresses many of the problems that are likely to arise as AI systems become more deeply integrated into healthcare.

8.1 Introduction

8.2 Theme 1: Transparency and Explainability

Transparency appeared as a central concern in all three cases, but it was not treated as a demand for technical disclosure. Courts and tribunals did not ask for detailed explanations of how systems worked internally. Instead, they treated transparency as a question of whether the system was understandable in practice. In *C. J. v. M.W.* (2020), the HPARB focused on whether the person being assessed could understand why a computerized tool was used and what role it played in decision making. In *Doobay v. Cohen* (2023), the Board examined whether the professional's reasoning was explained clearly enough to allow others to review and understand it. In *Lam v. Flo Health Inc.* (2024), the Court considered whether users could reasonably understand data practices based on how information was presented to them.

Taken together, these cases show that Canadian legal institutions emphasize transparency that allows affected individuals to make sense of what is happening. They do not require disclosure of technical details such as algorithmic logic or system design. This approach is consistent with the discussion in Chapter 2, which distinguished between technical explainability and transparency that is centered on the needs of users. For health AI governance, this suggests that legal expectations will likely focus on whether AI-assisted decisions can be meaningfully explained to patients, healthcare professionals, and regulators, rather than on whether systems meet abstract or highly technical standards of transparency.

8.3 Theme 2: Data Governance, Privacy, and Consent

A second theme running across the cases is that data governance was treated as broader than data protection or confidentiality alone. Courts and tribunals looked closely at how information was created, interpreted, and shared within larger institutional systems. In *C. J. v. M.W.* (2020), governance concerns arose from how assessment results were disclosed to third parties and how those results were explained in context. In *Doobay v. Cohen* (2023), attention focused on record-keeping practices and whether documentation made it possible to reconstruct professional decision making. In *Lam v. Flo Health Inc.* (2024) the Court closely examined how sensitive health data moved through third-party analytics and related systems. These cases show that data governance is understood as a process that unfolds over time. Legal reasoning follows information as it moves between actors and institutions, rather than concentrating only on how data is first collected. This supports the argument made in Chapter 3 that many governance problems in health AI arise from data flows and secondary uses, not from data breaches or security failures. For AI in healthcare, this theme highlights the need to govern entire data environments, including what happens to information after it has been processed by algorithmic systems.

8.4 Theme 3: Consent as a Dynamic and Contextual Process

In all three cases, consent was treated as something that develops over time and depends on context, rather than as a one-time or purely formal requirement. In *C. J. v. M.W.* (2020), consent was assessed in relation to how assessment data was later used and disclosed. In *Doobay v. Cohen* (2023), the adequacy of consent was indirectly linked to how clearly information was communicated and documented. In *Lam v. Flo Health Inc.* (2024), the Court

directly questioned whether consent obtained through privacy policies and interface design could be considered meaningful, given the complexity of the underlying data practices. Across the cases, courts emphasized that consent depends on the quality, timing, and clarity of the information provided to individuals. As discussed in Chapter 2, this is especially important for health AI systems, which often involve ongoing data collection and evolving uses over time. The cases suggest that Canadian courts may be cautious about consent models that rely only on broad or generalized disclosures. This has clear implications for how AI-enabled health technologies should be designed and governed.

8.5 Theme 4: Human Interpretation and Reliance on Automated Systems

Another shared theme is the continued importance of human judgment and institutional structures in environments shaped by complex technologies. In *Doobay v. Cohen* (2023), the Board emphasized that professional judgment requires interpreting complex information within established standards of practice. In *C. J. v. M.W.* (2020), the outputs of the assessment tool informed professional reasoning but did not replace it. In *Lam v. Flo Health Inc.* (2024), user understanding and organizational design shaped how automated data processing affected individuals. Together, these cases reject the idea that technology operates on its own, separate from human or institutional influence. Instead, they show that governance concerns arise from how technological outputs are used within decision-making systems. For health AI governance, this reinforces the argument made in Chapter 1 that AI should be understood as embedded in social and institutional contexts. Legal reasoning therefore focuses on human interaction with technologically generated information, rather than treating technology as an independent actor.

8.6 Theme 5: Judicial Reasoning and Anticipated Governance Gaps

Across all three cases, courts and tribunals identified points where existing legal frameworks have difficulty accommodating increasingly complex information practices. These difficulties were not framed as legal failures. Instead, they appeared as areas where established concepts require careful interpretation when applied to modern digital systems. The decisions drew attention to recurring challenges, including determining what level of disclosure is sufficient in complex and opaque digital environments, understanding the limits

of consent in systems that depend on large volumes of continuously generated data, and tracing how information moves across multiple actors and institutional boundaries. As discussed in Chapter 3, these challenges are likely to intensify as AI systems become more widely embedded in healthcare. The cross-case analysis demonstrates that Canadian legal reasoning is already grappling with these issues, even in situations where artificial intelligence is not explicitly named or directly at issue.

Chapter 9

9 Discussion

Chapter 8 brought together the findings from the case studies in Chapters 5 to 7. It identified common themes in how Canadian courts and tribunals reason about governance when they deal with complex information systems. Across the cases, decision makers repeatedly focused on transparency, how data are collected and managed, whether consent is meaningful, and how much human judgment remains involved in interpreting system outputs. The chapter showed that existing legal oversight has clear strengths, particularly in insisting on accountability and procedural fairness. At the same time, it revealed gaps that could become more serious as AI is used in healthcare. These gaps include limited visibility into how decisions are made, constantly changing data practices, and ongoing difficulties in ensuring that consent is informed and genuinely voluntary.

Chapter 9 builds on this analysis by drawing overall conclusions and directly answering the research questions. It explains how Canadian legal reasoning can shape the governance of health AI, with particular attention to transparency, consent, and the role of institutions in mediating technological use. The chapter then considers what these findings mean for policy design, regulatory frameworks, and professional practice. It shows how principles already present in Canadian case law can help guide the responsible use of AI in healthcare settings. Finally, the chapter points to areas for future research, especially where law, health policy, and technology governance will need to change in response to AI-driven developments, so that innovation does not undermine patient rights, safety, or public trust.

9.1 Introduction

Chapter 9 concludes the thesis by drawing together the findings developed across the earlier chapters and returning directly to the thesis research questions. This chapter synthesizes how Canadian legal decisions illuminate governance challenges relevant to the use of artificial intelligence in healthcare, particularly in relation to transparency, data

governance, consent, and judicial reasoning in the face of complex technological systems. The analysis demonstrates that even where AI is not explicitly addressed, existing legal reasoning provides important insight into how Canadian governance structures are likely to respond to health AI. This conclusion situates the findings within a broader legal and policy context, showing how issues identified in Canadian case law parallel concerns raised in international discussions on AI in healthcare. While artificial intelligence offers significant potential to improve diagnostic accuracy, clinical efficiency, and personalized care, the thesis has shown that these benefits are accompanied by persistent governance challenges. These include difficulties ensuring meaningful transparency, managing health-related data across institutional boundaries, and maintaining informed consent in environments characterized by automation and technical complexity. Importantly, these challenges are not unique to AI but are intensified by its scale and opacity. Finally, the chapter outlines directions for future research, emphasizing the need for continued empirical and doctrinal study as health AI becomes more widespread. As AI-specific disputes emerge in Canadian courts, further analysis will be needed to assess how existing legal principles evolve in response. Ongoing interdisciplinary research will be essential to ensure that the integration of artificial intelligence into healthcare supports innovation while remaining aligned with ethical, legal, and social values central to the Canadian health system.

9.2 Implications for Health AI Governance in Canada

Looking at the cases *C. J. v. M.W.* (2020), *Doobay v. Cohen* (2023), and *Lam v. Flo Health Inc.* (2024) reveals patterns that matter for governing AI in healthcare. These patterns go beyond the specific details of each case. Together, the cases show that the main legal and governance challenges of AI in health do not come from whether the algorithms themselves are accurate or innovative. The risks arise from how information is created, recorded, shared, and explained to people affected by AI. In other words, governance should focus on making AI-driven information understandable, traceable, and contestable. Four key areas stand out which are transparency, data flows, consent and user understanding, and explainability and communication.

9.2.1 Transparency

All three cases show that transparency is a core principle of governance, though it looks different depending on the context. In *C. J. v. M.W.* (2020), the court emphasized that people who undergo assessments must understand how decisions about them are made. It was not about seeing the inner workings of the tool but about knowing what the results mean for them. In *Doobay v. Cohen* (2023), transparency focused on whether professional records and reasoning were clear and reviewable, showing that transparency relies on documentation and explanation rather than revealing technical methods. In *Lam v. Flo Health Inc.* (2024), transparency concerned how companies present information about data practices to users so it is understandable and not hidden in legalistic terms. These cases suggest that transparency should be functional rather than technical. The goal is not for everyone to know how an algorithm works internally, but for people, professionals, and regulators to understand how AI outputs are created, used, and relied upon. Governance frameworks should require clear documentation, understandable explanations, and accessible descriptions of system operation without forcing disclosure of proprietary algorithms. Transparency acts as an indirect way to ensure accountability by making information practices visible and assessable.

9.2.2 Data flows

Managing data across organizations and systems is another key concern. In *C. J. v. M.W.* (2020), the case highlighted how assessment information moves beyond where it was first collected, particularly when shared with others or used in later decisions. This shows that governance responsibility continues beyond data creation. In *Doobay v. Cohen* (2023), the focus was on internal data handling, where clear, complete, and continuous records were essential. Even though no AI was involved, it shows that poor data practices can create governance risks, risks that become bigger when AI combines data from multiple sources. *Lam v. Flo Health Inc.* (2024) highlighted how personal health data moves to third parties for analytics or advertising, illustrating the complexity of modern digital health systems. Therefore, these cases indicate that governance must cover the entire data lifecycle, from collection to downstream use. AI systems often rely on distributed data ecosystems rather than isolated databases, which can create gaps in

responsibility. Effective governance requires traceable and auditable practices, clear roles among data controllers, platform operators, and third-party processors, and alignment with legal, professional, and privacy standards.

9.2.3 Consent and user understanding

Consent is not just a one-time step but an ongoing, context-dependent process. In *C. J. v. M.W.* (2020), consent had to evolve as the use and sharing of assessment information changed. In *Doobay v. Cohen* (2023), consent was assessed indirectly through clear documentation and transparency, focusing on whether individuals could understand the professional process. *Lam v. Flo Health Inc.* (2024) emphasized whether users could realistically understand data collection and third-party sharing. This shows that consent in health AI cannot rely only on accepting terms once. As AI systems evolve, add new data sources, or create new uses, consent must be revisited and supported with understandable information. Governance should include ongoing disclosure, context-specific explanations, and chances for users to reconsider participation. Educating clinicians and patients about how AI produces outputs is essential for meaningful consent.

9.2.4 Explainability and communication

Explainability and communication connect transparency, data flows, and consent. In all three cases, courts looked at whether information produced or mediated by digital systems was communicated in ways that supported understanding and proper use. *C. J. v. M.W.* (2020) highlighted the need for clear feedback from assessment processes. *Doobay v. Cohen* (2023) showed that professional decisions must be clearly explained. *Lam v. Flo Health Inc.* (2024) extended this to consumer health technologies, where explaining data practices clearly is critical. Moreover, these cases show that explainability is not just technical. It also includes user-friendly interfaces, clear narratives, educational resources, and institutional support to help people make sense of AI outputs. Governance frameworks should ensure AI outputs are interpretable in practice, with explanations of context, limitations, and proper use. Good communication reduces misinterpretation, supports informed decisions, and builds trust in AI-enabled healthcare systems.

9.3 Interdisciplinary Connections

The findings of this research have the potential to significantly shape the adoption of AI in healthcare, influence healthcare policies, and build public trust both within Canada and internationally. By addressing critical challenges such as algorithmic bias, accountability, and data privacy, this study offers a roadmap for fostering trust in AI systems. For healthcare providers and patients, transparent consent models and strong data protection measures can reduce resistance to AI technologies. On a national level, the research encourages policymakers to establish clear standards that align AI development with ethical principles, helping to build a culture of trust and confidence. Globally, the findings contribute to the creation of best practices for the ethical integration of AI, enabling international stakeholders to harmonize governance frameworks and promote fair access to AI-driven healthcare innovations.

The research also highlights the need for updated healthcare policies that address the unique complexities of AI. Policies that require regular audits of AI algorithms, ensure robust data security, and support interoperability with existing healthcare systems are essential for maintaining public trust and improving healthcare delivery. By proposing adaptive regulatory frameworks, the study stresses the importance of policies that can evolve alongside advancements in AI. These recommendations also have relevance for international healthcare systems, offering a model for cross-border collaboration on AI governance, particularly in areas such as privacy, fairness, and accountability.

Public trust is fundamental to the successful integration of AI in healthcare. This research demonstrates that ethical governance, combined with transparency and equitable access, can foster trust among patients, healthcare providers, and the broader public.

Emphasizing informed consent and empowering patients to understand and engage in AI-driven decisions strengthens the relationship between technology and its users. By addressing concerns such as data misuse and accountability gaps, the research provides practical strategies for building trust at the national level, which can also be adapted to other countries and regions seeking to integrate AI into their healthcare systems responsibly.

This thesis makes valuable contributions to several disciplines, establishing critical connections between law, ethics, healthcare, and technology. In the legal field, it identifies key gaps in current frameworks, such as those related to informed consent and algorithmic accountability and offers insights for developing new legal principles to guide AI governance. From an ethical perspective, the research underscores the importance of embedding principles like fairness, transparency, and inclusivity into AI systems. By linking theoretical ethical debates with practical recommendations, it bridges the gap between abstract concepts and real-world applications.

From a healthcare standpoint, the findings highlight how AI can enhance diagnostic accuracy, improve resource allocation, and personalize treatments, while also warning against potential risks, such as reduced human oversight and the erosion of patient-provider trust. The research advocates for a collaborative approach to AI integration, ensuring that AI supports and enhances human decision-making rather than replacing it. For technologists and AI developers, the study provides practical guidance on creating systems that adhere to ethical and legal standards. Recommendations on transparency, data governance, and accountability offer a framework for developing AI technologies that prioritize user trust and societal benefits, ensuring that AI aligns with the broader goals of justice, equity, and public well-being. By demonstrating the relevance of its findings across these various disciplines, this research not only addresses the challenges posed by AI integration in healthcare but also offers actionable insights for advancing equitable, ethical, and effective healthcare systems at both national and global levels. It positions Canada as a potential leader in the responsible adoption of AI in healthcare, providing a model for other countries to follow.

9.4 Implications COVID-19 and the Increased Usage of AI

The cases discussed here, while not all directly related to COVID-19, reflect broader issues that became particularly prominent during the pandemic, especially the rapid deployment of AI technologies in healthcare and other sectors. In *C. J. v. M.W.* (2020), the use of BrainFx as a cognitive assessment tool highlights the growing trend of adopting AI systems to improve efficiency in healthcare. However, the applicant's concerns about the tool's accuracy and appropriateness underscore the risks associated

with deploying AI without sufficient validation or transparency. This mirrors broader ethical concerns raised during the pandemic, where the urgency of AI adoption sometimes conflicted with important issues like privacy, accuracy, and patient rights.

In *Doobay v Cohen* (2023), while not addressing AI applications related to the pandemic, touches on broader societal issues, such as income disparity and healthcare access, which were exacerbated by COVID-19. The case reflects ongoing ethical scrutiny of AI's role in healthcare, especially in terms of its potential to address or worsen inequities. The Board's dismissal of the complainant's concerns about systemic issues mirrors challenges faced during the pandemic in balancing the urgency of AI deployment with the goal of achieving equitable outcomes.

Lastly, in *Lam v. Flo Health Inc.* (2024), the case, though predating COVID-19, highlights risks associated with rapid deployment of digital health tools, a trend that was accelerated during the pandemic. The case focuses on privacy concerns arising from the misuse of data by Flo's app, which was used for remote health monitoring. This mirrors debates about contact-tracing apps, where urgency often conflicted with ethical data practices. The case demonstrates how poorly governed AI-driven health tools can erode public trust, a critical lesson for future responses to public health crises.

9.5 Strengths and Limitations

This thesis's main strength lies in its combination of Merriam's (1998) qualitative case study framework with established methods used in systematic reviews. Following Merriam's guidelines, the study focuses on a small number of carefully selected legal cases, provides detailed and contextual descriptions, and emphasizes practical relevance. This approach supports a clear, context-based understanding of how courts address legal issues related to AI in healthcare. The use of the PRISMA checklist further strengthened the study by improving transparency, reducing selection bias, and supporting peer evaluation through a systematic and replicable process for case identification, screening, and analysis, thereby enhancing internal validity.

The review used a comprehensive search strategy with no restrictions on publication dates and broad eligibility criteria. This allowed the inclusion of both early and more recent cases that reflect emerging legal issues at the intersection of technology and healthcare. This approach is consistent with Merriam's (1998) emphasis on producing detailed explanations that clarify complex phenomena, in this case the governance challenges associated with AI in healthcare.

Several techniques drawn from Merriam (1998) were used to support validity, although not all were applied in a traditional manner. Triangulation was achieved through the use of multiple data sources, including judicial decisions from CanLII, relevant academic literature, and iterative refinements of the search strategy. This reduced the risk that identified patterns were the result of a single source of data. While formal member checking through interviews or observations was not possible given the documentary nature of the study, informal discussions with legal experts during peer review helped assess the plausibility of interpretations of judicial reasoning. Although prolonged observation was not applicable, the study adopted a longitudinal perspective by examining cases decided from 2020 onward, allowing for analysis of changes in judicial approaches to AI-related issues in healthcare over time.

Peer review played an important role in maintaining research quality. The study's methodology, case selection, and data extraction processes were reviewed by academic peers and advisors. To reduce researcher bias, the study documented the selection process in detail, refined search terms through repeated testing, and applied clear inclusion and exclusion criteria. Although the analysis relied primarily on documentary sources rather than interviews or observations, this approach is consistent with Merriam's (1998) case study methodology when the focus is on legal analysis.

Despite these strengths, the study has several limitations. First, it included only English-language sources, which may have excluded relevant legal decisions and academic work published in other languages. This may limit perspectives from non-English-speaking jurisdictions. Second, the lack of shared definitions for key terms such as AI and ML made case identification more difficult and required repeated adjustments to search terms

and eligibility criteria. Future research could address this issue by involving interdisciplinary experts to help develop clearer and more consistent definitions.

Another limitation is the reliance on publicly available Canadian case law, which may not capture disputes resolved through private settlements or informal processes. Future studies could complement document analysis with interviews or surveys involving key stakeholders, such as healthcare providers, technology developers, policymakers, and patients, to gain broader perspectives. The rapid pace of change in AI technologies and related regulations also poses challenges, as legal analysis may become outdated. Longitudinal studies and comparative research involving other jurisdictions, such as the European Union or the United States, could help track these developments and identify broader governance trends.

The scope of this thesis is further limited by its focus on judicial reasoning rather than on a detailed analysis of legislative frameworks. As a result, the study does not assess the design or adequacy of specific statutes. Future research could build on the governance issues identified in this analysis by examining how existing and proposed legislation addresses concerns such as accountability, data governance, and bias. In addition, the focus on the Canadian healthcare context does not fully address socioeconomic factors, such as differences in access to digital technologies or levels of digital literacy, which may affect how AI is used in practice. These issues could be explored in future research through community-based and localized studies.

In conclusion, this research contributes to understanding how courts respond to legal and ethical issues raised by AI in healthcare. At the same time, its limitations point to the need for continued research that is longitudinal, comparative, and interdisciplinary in order to support the responsible and equitable use of AI in healthcare systems.

9.6 Further Research

Future research should adopt a comprehensive, interdisciplinary approach to examine how AI is reshaping healthcare, with a focus on legal accountability, ethical governance, and equitable implementation. A critical starting point is the analysis of legal cases

involving AI in healthcare, incorporating perspectives from plaintiffs, defendants, and policymakers. Studies should investigate how courts interpret liability, negligence, and professional accountability in AI-driven clinical decisions, particularly in cases where AI systems fail, exhibit bias, or contradict medical expertise. For example, in Canada, reviewing precedents under laws such as the RHPA and the PHIPA could clarify how existing frameworks address emerging AI challenges. Comparative research on international regulations including the EU's General Data Protection Regulation (GDPR) or the U.S. Health Insurance Portability and Accountability Act (HIPAA) could also identify best practices for balancing innovation with patient rights.

A deeper examination of AI's technical design and decision-making processes is equally vital. Research must move beyond algorithmic performance to assess real-world impacts, including how AI influences power dynamics between clinicians and patients, healthcare delivery, and systemic inequities. Transparency and interpretability of AI systems should be prioritized to address concerns about “black box” algorithmic processes, ensuring clinicians and patients understand AI-driven decisions. Qualitative studies capturing stakeholder experiences from healthcare providers and patients to developers can reveal whether AI enhances or disrupts care, as well as unintended consequences such as bias or disparities in access. For marginalized groups, particularly Indigenous communities or low-income populations, research must evaluate whether AI entrenches existing inequities due to unrepresentative data or uneven adoption.

Ethical and regulatory gaps demand urgent attention. Privacy law reforms are needed to address AI-specific challenges, including dynamic consent models for evolving algorithms and safeguards for secondary data use in AI training. Studies should explore how to adapt frameworks like PHIPA to govern biometric data from wearables and third-party vendor compliance. Concurrently, research must propose governance models that integrate technical, legal, and ethical safeguards, such as hybrid rule-based and ML systems that align predictive accuracy with regulatory clarity. Finally, socioeconomic research should assess how liability frameworks might allocate responsibility among clinicians, developers, and vendors, ensuring accountability without stifling innovation.

By combining case law analysis, stakeholder insights, and comparative policy review, future research can bridge gaps between AI's potential and its safe, equitable integration into healthcare. As Merriam (1998) emphasizes, heuristic, context-aware methodologies will be essential to develop frameworks that protect patient rights while fostering responsible AI advancement.

9.7 Conclusion

This review provides a broad and interdisciplinary examination of Canadian legal cases that shed light on governance challenges linked to the use of AI and AI-like technologies in healthcare. It uses a qualitative case study approach based on Merriam (1998), which emphasizes close, detailed analysis of real-world situations. This method allows the thesis to describe and examine legal and ethical issues that arise when healthcare increasingly relies on digital or technology-mediated tools. By analyzing court and tribunal decisions, the review identifies patterns in how judges and regulators reason about complex technologies and how established legal principles are applied in these contexts. A central focus is how issues of transparency, consent, and data governance are experienced by patients, healthcare professionals, and others affected by digital or AI-supported health systems. Unlike abstract or purely theoretical discussions of AI ethics, this analysis is grounded in Canadian case law and reflects Merriam's emphasis on explanations that are closely tied to real-world context.

The thesis examines key governance questions through three Canadian legal cases that illustrate challenges linked to AI-like technologies in healthcare: *C. J. v. M.W.* (2020); *Doobay v. Cohen* (2023); and *Lam v. Flo Health Inc.* (2024). Although *C. J. v. M.W.* (2020) does not directly involve artificial intelligence, it raises important concerns about digital assessment tools, particularly in relation to the interpretation of results, informed consent, and information handling. *Doobay v. Cohen* (2023) builds on these concerns by examining how complex professional records and documentation are governed, highlighting how human judgment operates alongside technology-mediated processes. *Lam v. Flo Health Inc.* (2024) shifts the focus to a consumer-facing digital health application and examines transparency, consent, and data sharing in contexts where health-related data are processed automatically and at scale. Taken together, these cases

show that Canadian courts continue to rely on established legal principles, such as privacy, consent, and professional standards, when addressing technology-related healthcare disputes, even as new forms of digital mediation become more prominent.

Several common legal and governance themes emerge across the cases. Transparency is understood not as technical disclosure of system design, but as whether information and decision-making processes are understandable to those affected. Data governance is treated broadly and includes how information is collected, interpreted, shared, and managed, including through third parties. Informed consent is framed as an ongoing and context-dependent process rather than a one-time event. Difficulties arise when explanations are vague or overly general and do not clearly convey how data will be used, as illustrated in *C. J. v. M.W.* (2020) and *Lam v. Flo Health Inc.* (2024). Across the cases, courts grapple with how existing legal concepts apply to emerging technological practices, particularly when algorithmic or automated outputs influence clinical or health-related decision-making.

Overall, the analysis shows that existing legal principles address many risks relevant to AI and digital health technologies, including concerns about privacy, fairness, transparency, and professional responsibility. At the same time, the cases reveal areas of uncertainty and strain in legal reasoning as courts confront increasingly complex and opaque technological systems. *C. J. v. M.W.* (2020) demonstrates challenges in explaining and obtaining meaningful consent for digital cognitive assessments. *Doobay v. Cohen* (2023) underscores the importance of careful interpretation and oversight of professional records created or managed through technology. *Lam v. Flo Health Inc.* (2024) highlights the significance of clarity and accountability in large-scale digital health data practices. Together, these patterns suggest that effective governance of health AI depends on continued dialogue across legal, clinical, and technical domains.

In conclusion, AI and digital health technologies hold significant potential to influence healthcare delivery and patient outcomes in Canada. Realizing these benefits will depend on governance approaches that emphasize transparency, meaningful consent, and responsible information practices, while remaining attentive to the ways technology

reshapes decision-making and professional roles. Courts, policymakers, healthcare professionals, and technologists all play a role in shaping how these tools are integrated into care. By grounding governance discussions in real legal reasoning and lived disputes, this thesis contributes to a clearer understanding of the challenges that lie ahead as AI becomes more deeply embedded in healthcare practice.

References

- Alowais, S. A., Alghamdi, S. S., Alsuhebany, N., Alqahtani, T., Alshaya, A. I., Almohareb, S. N., Aldairem, A., Alrashed, M., Bin Saleh, K., Badreldin, H. A., Al Yami, M. S., Al Harbi, S., & Albekairy, A. M. (2023). Revolutionizing healthcare: The role of artificial intelligence in clinical practice. *BMC Medical Education*, 23(1), 689. <https://doi.org/10.1186/s12909-023-04698-z>
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Binns, R. (2018). Algorithmic Accountability and Public Reason. *Philosophy & Technology*, 31(4), 543–556. <https://doi.org/10.1007/s13347-017-0263-5>
- Boudreau LeBlanc, A., Williams-Jones, B., & Aenishaenslin, C. (2022). Bio-Ethics and One Health: A Case Study Approach to Building Reflexive Governance. *Frontiers in Public Health*, 10, 1–16. <https://doi.org/10.3389/fpubh.2022.648593>
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.
- Cestonaro, C., Delicati, A., Marcante, B., Caenazzo, L., & Tozzo, P. (2023). Defining medical liability when artificial intelligence is applied on diagnostic algorithms: A systematic review. *Frontiers in Medicine*, 10, 1–12. <https://doi.org/10.3389/fmed.2023.1305756>
- Chui, M., & Francisco, S. (2017). Artificial intelligence the next digital frontier. *McKinsey and Company Global Institute*, 47(3.6), 6–8.
- Crabtree, B. F., Miller, W. L., & Swenson, M. M. (1995). Doing Qualitative Research. *Nursing Research (New York)*, 44(4), 254. <https://doi.org/10.1097/00006199-199507000-00011>
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- Dong, H., Suárez-Paniagua, V., Whiteley, W., & Wu, H. (2021). Explainable automated coding of clinical notes using hierarchical label-wise attention networks and label embedding initialisation. *Journal of Biomedical Informatics*, 116, 103728. <https://doi.org/10.1016/j.jbi.2021.103728>
- Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv*. <https://arxiv.org/abs/1702.08608>
- Dul, J., & Hak, T. (2008). *Case study methodology in business research*. Routledge.
- Ebneyamini, S., & Sadeghi Moghadam, M. R. (2018). Toward developing a framework for conducting case study research. *International Journal of Qualitative Methods*, 17(1), 1–11. <https://doi.org/10.1177/1609406918817954>
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532–550. <https://doi.org/10.5465/amr.1989.4308385>
- Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in

- healthcare. *Nature Medicine*, 25(1), 24–29. <https://doi.org/10.1038/s41591-018-0316-z>
- Floridi, L., Cowsls, J., King, T. C., & Taddeo, M. (2020). How to design AI for social good: Seven essential factors. *Science and Engineering Ethics*, 26(3), 1771–1796. <https://doi.org/10.1007/s11948-020-00213-5>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting & Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Gerke, S., Minssen, T., & Cohen, G. (2020). Ethical and legal challenges of artificial intelligence-driven healthcare. In *Artificial intelligence in healthcare* (pp. 295–336). Academic Press. <https://doi.org/10.1016/B978-0-12-818438-7.00012-5>
- Gerring, J. (2007). Is There a (Viable) Crucial-Case Method? *Comparative Political Studies*, 40(3), 231–253. <https://doi.org/10.1177/0010414006290784>
- Holmes, J. H., Peek, N., Bellazzi, R., & Sacchi, L. (2015). Artificial intelligence in medicine: 15th Conference on Artificial Intelligence in Medicine, AIME 2015, Pavia, Italy, June 17-20, 2015. Proceedings (Vol. 9105). Springer. <https://doi.org/10.1007/978-3-319-19551-3>
- Holzinger, A., Biemann, C., Pattichis, C. S., & Kell, D. B. (2017). What do we need to build explainable AI systems for the medical domain? *arXiv*. <https://arxiv.org/abs/1712.09923>
- Iyer, R., Li, Y., Li, H., Lewis, M., Sundar, R., & Sycara, K. (2018). Transparency and explanation in deep reinforcement learning neural networks. Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, 144–150. <https://doi.org/10.1145/3278721.3278776>
- Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., Wang, Y., Dong, Q., Shen, H., & Wang, Y. (2017). Artificial intelligence in healthcare: past, present and future. *Stroke and Vascular Neurology*, 2(4), 230–243. <https://doi.org/10.1136/svn-2017-000101>
- Johnson, K. B., Wei, W., Weeraratne, D., Frisse, M. E., Misulis, K., Rhee, K., Zhao, J., & Snowdon, J. L. (2021). Precision Medicine, AI, and the Future of Personalized Health Care. *Clinical and Translational Science*, 14(1), 86–93. <https://doi.org/10.1111/cts.12884>
- Karimian, G., Petelos, E., & Evers, S. M. A. A. (2022). The ethical issues of the application of artificial intelligence in healthcare: a systematic scoping review. *AI and Ethics (Online)*, 2(4), 539–551. <https://doi.org/10.1007/s43681-021-00131-7>
- Lipton, Z. C. (2018). The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *ACM Queue*, 16(3), 31–57. <https://doi.org/10.1145/3236386.3241340>
- Loh, H. W., Ooi, C. P., Seoni, S., Barua, P. D., Molinari, F., & Acharya, U. R. (2022). Application of explainable artificial intelligence for healthcare: A systematic review of the last decade (2011–2022). *Computer Methods and Programs in Biomedicine*, 226, 107121. <https://doi.org/10.1016/j.cmpb.2022.107121>
- Mehrotra, S., Degachi, C., Vereschak, O., Jonker, C. M., & Tielman, M. L. (2024). A Systematic Review on Fostering Appropriate Trust in Human-AI Interaction: Trends, Opportunities and Challenges. *ACM Journal on Responsible Computing*, 1(4), 1–45. <https://doi.org/10.1145/3696449>

- Merriam, S.B. (1998). *Qualitative Research and Case Study Applications in Education*. Jossey-Bass.
- Meyer, C. B. (2001). A Case in Case Study Methodology. *Field Methods*, 13(4), 329–352. <https://doi.org/10.1177/1525822X0101300402>
- Murdoch, B., Jandura, A., & Caulfield, T. (2022). Privacy considerations in the Canadian regulation of commercially-operated healthcare artificial intelligence. *Canadian Journal of Bioethics*, 5(4), 44–52. <https://doi.org/10.7202/1094696ar>
- Nazar, M., Alam, M. M., Yafi, E., & Su'ud, M. M. (2021). A Systematic Review of Human-Computer Interaction and Explainable Artificial Intelligence in Healthcare With Artificial Intelligence Techniques. *IEEE Access*, 9, 1–33. <https://doi.org/10.1109/ACCESS.2021.3127881>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
- O'Neill, O. (2003). Some limits of informed consent. *Journal of Medical Ethics*, 29(1), 4–7. <https://doi.org/10.1136/jme.29.1.4>
- Patton, M. Q. (1999). Enhancing the quality and credibility of qualitative analysis. *Health Services Research*, 34(5 Pt 2), 1189–1189.
- Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring Fairness in Machine Learning to Advance Health Equity. *Annals of Internal Medicine*, 169(12), 866–872. <https://doi.org/10.7326/M18-1990>
- Rueda, J. D., Chou, S., Sharp, R. R., & MacKay, C. R. (2022). Procedural fairness in high-stakes health decisions: Evidence and ethics. *Health Policy*, 126(6), 472–480. <https://doi.org/10.1016/j.healthpol.2022.03.004>
- Sarker, S., Jamal, L., Ahmed, S. F., & Irtisam, N. (2021). Robotics and artificial intelligence in healthcare during COVID-19 pandemic: A systematic review. *Robotics and Autonomous Systems*, 146, 1–18. <https://doi.org/10.1016/j.robot.2021.103902>
- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press. <https://doi.org/10.1093/oso/9780192845290.001.0001>
- Siala, H., & Wang, Y. (2022). SHIFTing artificial intelligence to be responsible in healthcare: A systematic review. *Social Science & Medicine (1982)*, 296, 1–15. <https://doi.org/10.1016/j.socscimed.2022.114782>
- Sokol, K., & Flach, P. (2020). One Explanation Does Not Fit All: The Promise of Interactive Explanations for Machine Learning Transparency. In *KI. Künstliche Intelligenz (Oldenbourg)* (Vol. 34, Number 2, pp. 235–250). Springer Berlin Heidelberg. <https://doi.org/10.1007/s13218-020-00637-y>
- Stake, R., & Walker, R. (1996). [Rev. of *The art of case study research*]. *Evaluation (London, England. 1995)*, 2(2), 231–235.
- Stanfill, M. H., & Marc, D. T. (2019). Health Information Management: Implications of Artificial Intelligence on Healthcare Data and Information Management. *Yearbook of Medical Informatics*, 28(1), 56–64. <https://doi.org/10.1055/s-0039-1677913>
- Taeihagh, A. (2021). Governance of artificial intelligence. *Policy & Society*, 40(2), 137–157. <https://doi.org/10.1080/14494035.2021.1928377>

- Tang, L., Li, J., & Fantus, S. (2023). Medical artificial intelligence ethics: A systematic review of empirical studies. *Digital Health*, 9, 1–22. <https://doi.org/10.1177/20552076231186064>
- Tellis, W. (1997). Application of a Case Study Methodology. *Qualitative Report*. <https://doi.org/10.46743/2160-3715/1997.2015>
- Thunki, P., Reddy, S. R. B., Raparathi, M., Maruthi, S., Dodda, S. B., & Ravichandran, P. (2021). Explainable AI in Data Science-Enhancing Model Interpretability and Transparency. *African Journal of Artificial Intelligence and Sustainable Development*, 1(1), 1–8.
- Topol, E. (2019). *Deep medicine: how artificial intelligence can make healthcare human again*. Basic Books.
- Vogt, Y. (2025). Disability and algorithmic fairness in healthcare: a narrative review. *Journal of Medical Artificial Intelligence*, 8, 56–56. <https://doi.org/10.21037/jmai-24-415>
- Voss, C., Tsikriktsis, N., & Frohlich, M. (2002). Case research in operations management. *International Journal of Operations & Production Management*, 22(2), 195–219. <https://doi.org/10.1108/01443570210414329>
- Williamson, S. M., & Prybutok, V. (2024). Balancing Privacy and Progress: A Review of Privacy Challenges, Systemic Oversight, and Patient Perceptions in AI-Driven Healthcare. *Applied Sciences*, 14(2), 1–47. <https://doi.org/10.3390/app14020675>.
- World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. World Health Organization. <https://www.who.int/publications/i/item/9789240029200>
- World Health Organization. (2023). *Global Initiative on AI for Health*. <https://www.who.int/initiatives/global-initiative-on-ai-for-health>
- World Health Organization. (2024, January 18). *WHO releases AI ethics and governance guidance for large multi-modal models* [Press release]. <https://www.who.int/news/item/18-01-2024-who-releases-ai-ethics-and-governance-guidance-for-large-multi-modal-models>
- Yin, R. K., & Dhameja, N. (1997). [Rev. of *Case study research: design and methods*]. *Indian Journal of Public Administration*, XLIII(1), 113–114.

Appendices

A List of Abbreviations

AI: Artificial Intelligence

BCSC: Supreme Court of British Columbia

CanLII: Canadian Legal Information Institute

CNN: Convolutional Neural Networks

CT: Computed Tomography

EU: European Union

FIPPA: Freedom of Information and Protection of Privacy Act

GDPR: General Data Protection Regulation

GradCAM: Gradient-weighted Class Activation Mapping

HCI: Human-Computer Interaction

HIPAA: Health Insurance Portability and Accountability Act

HPARB: Health Professions Appeal and Review Board

ICRC: Inquiries, Complaints, and Reports Committee

ITU: International Telecommunication Union

LIME: Local Interpretable Model-agnostic Explanations

ML: Machine Learning

PHIPA: Personal Health Information Protection Act

PICO: Population, Intervention, Comparison, and Outcome

PIPEDA: Personal Information Protection and Electronic Documents Act

PRISMA: Preferred Reporting Items for Systematic Reviews and Meta-Analyses

RHPA: Regulated Health Professions Act

SHAP: SHapley Additive exPlanations

WIPO: World Intellectual Property Organization

WHO: World Health Organization

XAI: Explainable Artificial Intelligence

B Chapter 4: Methodology Extended Information

B.1 PRISMA Checklist

Section and Topic	Item #	Checklist item	Reported (Yes/No)
TITLE			
Title	1	Identify the report as a systematic review.	Yes
BACKGROUND			
Objectives	2	Provide an explicit statement of the main objective(s) or question(s) the review addresses.	Yes
METHODS			
Eligibility criteria	3	Specify the inclusion and exclusion criteria for the review.	Yes
Information sources	4	Specify the information sources (e.g. databases, registers) used to identify studies and the date when each was last searched.	Yes
Risk of bias	5	Specify the methods used to assess risk of bias in the included studies.	No
Synthesis of results	6	Specify the methods used to present and synthesise results.	Yes
RESULTS			
Included studies	7	Give the total number of included studies and participants and summarise relevant characteristics of studies.	Yes
Synthesis of results	8	Present results for main outcomes, preferably indicating the number of included studies and participants for each. If meta-analysis was done, report the summary estimate and confidence/credible interval. If comparing groups, indicate the direction of the effect (i.e. which group is favoured).	Yes
DISCUSSION			
Limitations of evidence	9	Provide a brief summary of the limitations of the evidence included in the review (e.g. study risk of bias, inconsistency and imprecision).	Yes
Interpretation	10	Provide a general interpretation of the results and important implications.	Yes
OTHER			
Funding	11	Specify the primary source of funding for the review.	No
Registration	12	Provide the register name and registration number.	No

From: Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ* 2021;372:n71. doi: 10.1136/bmj.n71. This work is licensed under CC BY 4.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Figure B.1.1 PRISMA 2020 for Abstracts Checklist

Section and Topic	Item #	Checklist item	Location where item is reported
Title	1	Identify the report as a systematic review.	37
Abstract	2	Specify for each outcome the effect measure(s) (e.g. risk ratio, mean difference) used in the synthesis or presentation of results.	
Abstract	7a	See the PRISMA 2020 for Abstracts checklist.	169, 170
Abstract	7b	Describe the processes used to decide which studies were eligible for each synthesis (e.g. tabulating the study characteristics, comparing the characteristics of the studies to the eligibility criteria).	
INTRODUCTION			
Rationale	3	Describe the rationale for the review in the context of existing knowledge.	34, 35
Objectives	4	Provide a clear, explicit statement of the objective(s) or question(s) the review addresses.	35, 36
METHODS			
Eligibility criteria	5a	Specify the inclusion and exclusion criteria for the review and provide a rationale for the choice of them. If meta-analysis was performed, describe the model(s), method(s) to identify the presence and extent of statistical heterogeneity, and software package(s) used.	33, 42, 93, 169, 170
Information sources	6a	Specify all databases, registers, or organisations, reference lists and other sources searched or consulted to identify studies. Specify the date when each source was last searched or consulted.	39, 42
Search strategy	7a	Present the full search strategies for all databases as assessed and describe them, including any filters and limits used.	170, 171
Reporting bias assessment	14	Describe any methods used to assess risk of bias due to missing results in a synthesis (arising from reporting biases).	
Selection process	8	Specify the methods used to decide whether a study met the inclusion criteria of the review, including how many reviewers screened each record and each report retrieved, whether they worked independently, and if applicable, details of automation tools used in the process.	30-42, 169, 170
Certainty assessment	15	Describe any methods used to assess certainty (or confidence) in the body of evidence for an outcome.	
Results	9	Specify the methods used to collect data from reports, including how many reviewers collected data from each report and whether they worked independently; any processes for obtaining or confirming data from study investigators; and if applicable, details of automation tools used in the process.	40-42
Study selection	10a	Describe the results of the search and selection process, from the number of records identified in the search to the number of studies included in the review, ideally using a flow diagram.	41-50
Data items	10b	List and define all outcomes for which data were sought. Specify whether all results that were compatible with each outcome domain in each study were sought (e.g. for all measures, time points, analyses), and if not, the methods used to decide which results to collect.	43-50, 42
Study characteristics	10b	List and define all study and person variables that were sought (e.g. participant and intervention characteristics, funding sources). Describe any assumptions made about any missing or unclear information.	43-46, 171-181
Risk of bias in studies	11	Present the results of risk of bias for each included study, including details of the tool(s) used, how many reviewers assessed each study and whether they worked independently, and if applicable, details of automation tools used in the process.	30-42
Results of individual studies	19	For all outcomes, present, for each study: (a) summary statistics for each group (where appropriate) and (b) an effect estimate and its precision (e.g. confidence/credible interval), ideally using structured tables or plots.	
RESULTS			
Results of syntheses	20a	For each synthesis, briefly summarise the characteristics and risk of bias among contributing studies.	
	20b	Present results of all statistical syntheses conducted. If meta-analysis was done, present for each the summary estimate and its precision (e.g. confidence/credible interval) and measures of statistical heterogeneity. If comparing groups, describe the direction of the effect.	
	20c	Present results of all investigations of possible causes of heterogeneity among study results.	
	20d	Present results of all sensitivity analyses conducted to assess the robustness of the synthesized results.	
Reporting biases	21	Present assessments of risk of bias due to missing results (arising from reporting biases) for each synthesis assessed.	
Certainty of evidence	22	Present assessments of certainty (or confidence) in the body of evidence for each outcome assessed.	
DISCUSSION			
Discussion	23a	Provide a general interpretation of the results in the context of other evidence.	43-46, 171-181
	23b	Discuss any limitations of the evidence included in the review.	50
	23c	Discuss any limitations of the review processes used.	50
	23d	Discuss implications of the results for practice, policy, and future research.	51
OTHER INFORMATION			
Registration and protocol	24a	Provide registration information for the review, including register name and registration number, or state that the review was not registered.	
	24b	Indicate where the review protocol can be accessed, or state that a protocol was not prepared.	
	24c	Describe and explain any amendments to information provided at registration or in the protocol.	
Support	25	Describe sources of financial or non-financial support for the review, and the role of the funders or sponsors in the review.	
Competing interests	26	Declare any competing interests of review authors.	
Availability of	27	Report which of the following are publicly available and where they can be found: template data collection forms;	

Section and Topic	Item #	Checklist item	Location where item is reported
data, code and other materials		data extracted from included studies; data used for all analyses; analytic code; any other materials used in the review.	

From: Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ* 2021;372:n71. doi: 10.1136/bmj.n71. This work is licensed under CC BY 4.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Figure B.1.2: PRISMA 2020 Checklist

B.2 Search Strategies

Research Question: To what extent do Canadian legal frameworks, such as privacy laws and health regulations, effectively govern the use of AI technologies in healthcare, and where do gaps or uncertainties remain?

The Population, Intervention, Comparison, and Outcome (PICO) framework was used to guide the search strategy and concepts table.

Table B.2.1 Search Strategy and Eligibility Criteria

Criteria	Inclusion Criteria	Exclusion Criteria
Population	Healthcare professionals, patients, or organizations using AI technologies in healthcare in Canada	No relevance to health
Intervention	Judicial or tribunal decisions addressing digital health or AI-related information practices in Canada	No relevance to AI technologies
Comparison		
Outcome	Patterns of judicial	

reasoning regarding
 transparency, consent, data
 governance, and
 accountability;
 identification of governance
 gaps; implications for
 health AI oversight in
 Canada

Table B.2.2 Key Concepts Table

Component 1	Component 2
"Artificial intelligence"	"Health"
OR	OR
"Machine Learning"	"Healthcare"
OR	OR
AND	"Primary care"
"Machine Learning"	OR
"Machine Learning"	"Tertiary care"
"Machine Learning"	OR
"Machine Learning"	"Health care"

B.3 Case Summary

[C. J. v M. W.](#), 2020 CanLII 57706 (ON HPARB)

<https://www.canlii.org/en/on/onhparb/doc/2020/2020canlii57706/2020canlii57706.html?resultId=4e6fa2599fc04a75936a9984fa1837b4&searchId=2025-01-07T01:58:30:074/4f0381ca3808411597b3c6851c927f0f>

Who are the parties?

The applicant is C.J. The defendant is M.W., an occupational therapist.

What court or tribunal is hearing the case?

The case was heard by the Health Professions Appeal and Review Board (HPARB).

What is or are the main legal issues that the lawsuit is about?

The lawsuit involves complaints about the professional conduct of M.W., including alleged misuse of assessment tools, breaches of privacy, acting beyond her scope of practice, and failing to meet professional standards during a workplace assessment.

What is the plaintiff/applicant arguing?

C.J. argued that M.W. misused the BrainFx assessment tool; failed to disclose the potential impact of the assessment results on his employment; breached privacy laws (PHIPA) by sharing sensitive information; acted outside her professional scope in conducting neuropsychological assessments; and colluded with his employer, which led to his job termination.

What is the defendant/respondent arguing?

M.W. responded that she used the BrainFx tool appropriately and did not make diagnoses outside her scope of practice; the assessment results were shared in compliance with consent provided by C.J.; and she followed professional standards and acted in C.J.'s best interest to secure accommodations.

Who won the case?

The defendant, M.W., won the case, as the HPARB confirmed the Committee's decision to take no further disciplinary action against her beyond offering professional guidance.

How did the court/tribunal reach the decision? In other words, what is the reason the winning party was successful?

The Board found that the Committee conducted an adequate investigation and made a reasonable decision; M.W. appropriately obtained and documented consent for her actions; and no evidence supported C.J.'s claim that M.W. acted in bad faith or outside her professional scope.

[Doobay v Cohen](#), 2023 CanLII 111813 (ON HPARB)

<https://www.canlii.org/en/on/onhparb/doc/2023/2023canlii111813/2023canlii111813.htm?resultId=0141dab6cc524c7294c5eb2b003c42f9&searchId=2025-01-07T02:04:52:345/83a1c887f17e44f6bbf3618ae8f840fa>

Who are the parties?

The plaintiff is Kathleen Doobay. The defendant is Rodney Victor Cohen, Registered Psychotherapist.

What court or tribunal is hearing the case?

The Health Professions Appeal and Review Board (HPARB) in Ontario, Canada, is reviewing the case.

What is or are the main legal issues that the lawsuit is about?

The main legal issues include whether the Respondent's conduct lacked diversity, equity, and inclusion; whether the Respondent engaged in coercive or abusive practices; and whether the Respondent had conflicts of interest.

What is the plaintiff/applicant arguing?

The Applicant, Kathleen Doobay, argued that the Respondent acted unprofessionally by lacking inclusivity and fairness in his interactions; using coercive practices; having conflicts of interest related to political and business affiliations; broader concerns about systemic issues like racism, privacy violations, and misuse of artificial intelligence.

What is the defendant/respondent arguing?

The Respondent, Rodney Victor Cohen, denied all allegations, stating that the Applicant was not his psychotherapy client; his conduct was professional and inclusive; he did not engage in any conflicts of interest or coercive practices; and he provided evidence about his work with diverse populations and his professional ethics.

Who won the case?

The respondent, Rodney Victor Cohen, won the case.

How did the court/tribunal reach the decision? In other words, what is the reason the winning party was successful?

The tribunal found that the investigation by the regulatory body, Inquiries, Complaints, and Reports Committee (ICRC), was thorough and collected all relevant evidence; there was insufficient evidence to support the Applicant's claims; the Respondent's conduct did not meet the threshold for professional misconduct under the law; and the tribunal upheld the decision to take no further action, as the allegations were either unsupported by evidence or outside the jurisdiction of the regulatory framework.

[Lam v Flo Health Inc.](#), 2024 BCSC 391

<https://www.canlii.org/en/bc/bcsc/doc/2024/2024bcsc391/2024bcsc391.html?resultId=b3069b5cfca44ca98937a053a011c422&searchId=2025-01-07T02:06:18:027/90d8ef8e1e10414a88183a9c65337367>

Who are the parties?

The plaintiff is Jaime Cah Kate Lam, representing Canadian users of the Flo Period & Ovulation Tracker app. The defendant is Flo Health Inc., the company that created the app.

What court or tribunal is hearing the case?

The case is being heard by the Supreme Court of British Columbia.

What is or are the main legal issues that the lawsuit is about?

The lawsuit focuses on whether Flo Health improperly shared users' sensitive health data with third parties in breach of privacy laws, its own privacy policies, and user consent.

What is the plaintiff/applicant arguing?

The plaintiff argues that Flo Health violated users' privacy rights, breached contractual obligations in its privacy policies, and was negligent in safeguarding their sensitive health information.

What is the defendant/respondent arguing?

Flo Health denies the allegations, claiming it followed its privacy policies, and that the data sharing was either permitted or anonymized. It also argued that the case should not proceed due to clauses in its user agreement that limited lawsuits and required disputes to be resolved under California law.

Who won the case?

The plaintiff, Jaime Cah Kate Lam, won the certification phase, meaning the court approved the case to proceed as a class action.

How did the court/tribunal reach the decision? In other words, what is the reason the winning party was successful?

The court found that the plaintiff met the requirements for certifying a class action under the Class Proceedings Act. It rejected the defendant's arguments about jurisdiction and

clauses in the user agreement (like class action waivers), calling them unconscionable and against public policy. The court held that Flo Health's privacy policies and actions raised serious issues of potential breaches that should be resolved in a trial.

B.4 References

[C. J. v M. W.](#), 2020 CanLII 57706 (ON HPARB)

[Doobay v Cohen](#), 2023 CanLII 111813 (ON HPARB)

[Lam v Flo Health Inc.](#), 2024 BCSC 391

Curriculum Vitae

Name: Earlel Thiyagaratnam

Post-secondary Education and Degrees:

Brock University
St. Catharines, Ontario, Canada
BSc

York University
Toronto, Ontario, Canada
BHS

Western University
London, Ontario, Canada
MHIS

Western University
London, Ontario, Canada
MSc

Western University
London, Ontario, Canada
PhD HIS

Honours and Awards: Province of Ontario Graduate Scholarship
2024-2025

Related Work Experience Teaching Assistant
The University of Western Ontario
2022-2025

Publications:

Ajanaku, O. J., & Thiyagaratnam, E. (2025). Generative AI: A disruptive game changer in higher education. In T. R. Merlo (Ed.), *AI-driven revolution: Transforming the business landscape* (pp. 1–17). World Scientific.
https://doi.org/10.1142/9781800616578_0009