Reprinted from *Canadian Healthcare Technology,* April 2026 issue

CANADIAN Healthcare Technology

www.canhealth.com

Conducting an artificial intelligence Privacy Impact Assessment (PIA)

BY PATRICK LO AND SHIRLEY FENTON

Artificial intelligence (AI) is moving quickly from pilot projects into day-to-day healthcare operations, supporting clinical documentation, diagnostics, scheduling, population health analytics, and patient communications.

While these tools promise efficiency and improved care, they also amplify privacy risk. In Canada, a Privacy Impact Assessment (PIA) remains one of the most effective mechanisms for ensuring that AI adoption respects patient privacy, complies with Canadian federal and provincial laws, and maintains public trust.

What is a Privacy Impact Assessment (PIA)? A Privacy Impact Assessment is a structured, documented process designed to identify, assess, and mitigate privacy risks associated with the collection, use, disclosure, and retention of personal information or personal health information (PHI).

In healthcare, a PIA typically records the purpose and legal authority for collecting PHI, outlines detailed data flows across systems and vendors, and highlights any identified privacy risks. Additionally, it describes the safeguards and mitigation measures implemented to address these risks, as well as any residual risks that have been accepted by accountable leadership.

PIAs are not theoretical exercises. They are practical risk management tools used to assess whether an organization exercised due diligence when introducing new systems or technologies.

Why PIAs matter in healthcare: Healthcare data is one of the most sensitive types of personal information. Improper use or exposure of this data can lead to stigma, discrimination, and a loss of trust, ultimately causing real harm to patients. The risks associated with this data are amplified by AI systems that operate at a large scale, rely on complex data processing, and can evolve over time.

For hospitals and health authorities, a PIA demonstrates that innovation has been balanced with legal compliance, ethical obligations, and patient expectations. In practice, PIAs help executives, project sponsors, privacy offices, IT, and clinical leaders ask the right questions early, before a tool is embedded in clinical workflows.

Shirley Fenton

Patrick Lo

Is an "AI PIA" different from a traditional PIA? Fundamentally, an AI PIA adheres to the same principles as any standard PIA: it involves data mapping, analysis of legal authority, identification of risks, and planning for mitigation. What differs is the extent and focus of the risk perspective.

The AI PIA risk analysis should consider the following use cases, addressing issues commonly associated with AI initiatives:

- Training and secondary use of data: Was PHI utilized for training or fine-tuning the model? If so, what authority was involved, what was the sample size, what was the population of the patient data (diversity, conditions, institutions) and what measures were in place to ensure safety?
- Inference and re-identification risk: AI-generated outputs can often expose sensitive characteristics or lead to conclusions about patients that extend beyond the data originally provided.
- Transparency and explainability: It is essential for both clinicians and patients to have a clear understanding of the capabilities and limitations of AI technologies. They should be aware of what the AI can accomplish, what it cannot, and how to properly interpret the outputs provided by these systems.
- Bias and equity: It is essential to evaluate and track performance disparities among different populations.
- Governance Risk: Weak governance can lead to several significant challenges, including regulatory violations, ethical shortcomings, unanticipated behaviors in models, and a decline in organizational trust. Addressing these issues is crucial for maintaining integrity and accountability within an organization.
- Ongoing change: Changes in model updates, prompt modifications, and vendor releases can significantly impact privacy risks over time. It's important to stay informed about these developments, as

they may alter how data is handled and protected. A continuous cycle of monitoring, detection, analysis, and remediation is needed to avoid data drift.

• Privacy Risk: Evaluates how the AI system collects, uses, stores, shares, and protects personal information/personal health information.

• Security Risk: Security risk looks at the AI system's protection from cyberattacks, manipulation, hacking, and unauthorized access.

Procuring AI vs. building AI – key PIA differences: The decision to procure an AI solution from an external provider or to develop one in-house has implications for the focus of the PIA. However, regardless of the approach taken, the requirement to complete a PIA remains unchanged.

• Procuring AI (vendor solutions, SaaS, embedded EMR tools): PIAs for procured AI emphasize vendor due diligence and contractual controls. Key considerations include data residency, cross-border transfers, subcontractors, restrictions on secondary use or model training, breach notification timelines, audit rights, and how updates are managed. Guidance from bodies such as the Information and Privacy Commissioner of Ontario and the Information Privacy Commissioner of Alberta is particularly relevant for hospitals navigating vendor-led AI deployments.

• Building AI (in-house or custom-trained models): When an organization is involved in building an AI solution, it assumes greater accountability throughout the process. The AI governance and documentation become critical. The Privacy Impact Assessment (PIA) must extend into the system development lifecycle, addressing several critical areas. These include data collection and quality control measures to ensure the integrity of the information used. Additionally, the PIA should cover model testing, validation, and documentation to establish the reliability of the AI systems. Ongoing monitoring for bias and potential misuse is essential to maintain ethical standards and effectiveness. Finally, effective change management and version control processes need to be implemented to adapt to new challenges and ensure the AI system remains relevant and compliant over time.

Canada's legislative and guidance landscape: There's currently no Federal AI-specific legislation in Canada. However, across Canada and at the international level, PIAs are embedded in privacy frameworks, with increasing attention to AI. The following are some of the PIA and AI guidance that will assist in the AI PIA process:

• Federal: Under the Privacy Act and Treasury Board policy instruments, federal institutions must complete PIAs. The federal Directive on Automated Decision-Making and Algorithmic Impact Assessment sets additional expectations for automated systems. The Office of the Privacy Commissioner of Canada has also issued guidance on responsible, privacy-protective AI. In addition, the Pan-Canadian AI for Health (AI4H) Guiding Principles outline person-centric, equitable, safe, and transparent adoption of AI in Canada's health systems.

A PIA remains one of the most effective mechanisms for ensuring that AI adoption respects patient privacy,

• Ontario: Hospitals operating as public institutions must comply with FIPPA, while also meeting PHIPA obligations for PHI. Recent amendments have made PIAs explicitly mandatory before collecting personal information. The IPC has published AI-specific guidance for healthcare contexts.

• Alberta: The Office of the Information and Privacy Commissioner of Alberta provides PIA guidance and has released AI-scribe-specific materials for custodians under the Health Information Act.

• Québec: Law 25 introduced a formal privacy impact assessment regime, with guidance from the Commission d'accès à l'information.

• Professional colleges: Clinical regulators such as the College of Physicians and Surgeons of Ontario and the College of Physicians & Surgeons of Alberta have issued guidance on the responsible use of AI, reinforcing professional accountability alongside privacy compliance.

• International: The OECD AI Principles provide a global foundation focused on human rights, fairness, transparency, and societal benefit. These principles align closely with the Pan-Canadian AI for Health (AI4H) guiding principles. The EU AI Act establishes a comprehensive set of requirements to effectively mitigate the risks associated with AI.

An AI PIA is not a one-time checkbox; it is a living governance artifact. Conducted early and maintained over time, it enables innovation while protecting patients, clinicians, and organizations alike.

If an organization cannot clearly explain how AI data flows work, what authority supports them, and how harms are mitigated, it is not yet ready to deploy. Privacy Impact Assessments are not merely a compliance checkbox; they represent a commitment to safeguarding the rights and dignity of individuals in a digital age.

By integrating privacy into every facet of data practices, organizations contribute to a world where data-driven innovation coexists harmoniously with privacy protection, enriching the digital experience for all.

Patrick Lo is chief executive officer of Privacy Horizons. Shirley Fenton is president of the National Institutes of Health Informatics and cofounder of Waterloo MedTech.