



# Agentic AI is evolving, and so is primary care in Canada.

March 2026

# Primary care is stretched thin. AI can help.

Canada's primary care system is at an inflection point — not only in the pressures facing care teams, but in the technologies now available to support them. Advances in artificial intelligence (AI) are creating new opportunities to reduce administrative and cognitive burden, allowing clinicians and interdisciplinary teams to focus more fully on comprehensive, relationship-based care.

In particular, AI agents have the potential to support longitudinal clinical workflows, as does agentic AI, the next stage of their evolution, provided they are developed with clear boundaries and clinician oversight.

From a clinical perspective, this moment is shaped by the growing administrative and cognitive load placed on care teams, driven by rising patient volumes, increasing clinical complexity, and persistent resource constraints. Against that backdrop, **this paper explores how emerging AI capabilities are beginning to translate into practical, clinician-led solutions for primary care** — augmenting team-based workflows while preserving clinical judgment.

To start, it's essential to understand and contextualize the urgency for primary care providers in Canada. Clinicians want to ensure the best possible patient outcomes, which means providing personalized, continuous, preventive care. This work coincides with the realities of an aging population and a growing prevalence of complex, multimorbid patients.

Demand for primary care frequently exceeds available capacity<sup>1</sup>, contributing to widespread clinician burnout and placing additional strain on already limited team and system resources.

It's imperative to find solutions that can fulfil two critical functions: first, relieve the pressure on health systems and care providers; and second, empower care teams to do more with less. AI holds significant potential to address both these areas<sup>2</sup>, and adoption among physicians is rising.






Early adoption has been strongest in tools that reduce documentation burden. Already, two in five general practitioners and two in three specialists report using AI tools to help with tasks, according to an independent study conducted by MD Analytics. There's also an overwhelmingly optimistic perception and reception of AI tools, with 92 per cent of general practitioners and 89 per cent of specialists confident that they will have a net positive impact on their practices in the coming years.<sup>3</sup>

According to a study by Canada Health Infoway<sup>4</sup>, Ambient medical scribes already save time for 94 per cent of physicians surveyed. Not only do they report feeling more present during patient visits, but 89 per cent note a reduced administrative burden and 85 per cent report a reduced cognitive load.

Yet AI tools in their current state haven't solved the main frustrations facing care teams, with clinical documentation responsibilities being a leading cause of burnout for 90 per cent of physicians. While chatbots, AI-enabled research tools, and ambient scribes mean fewer keystrokes, they still operate within existing workflows or are bolted on. These tools must be initiated by physicians, and therefore don't support them with recall, follow-up, and the system-level cognitive burden that comes from trying to provide quality care for hundreds or thousands of patients.

**On average, Canadian physicians report that they're still spending approximately nine hours per week on admin and that 47 per cent of these tasks are unnecessary<sup>5</sup> — yet these tasks cumulatively consume some 19.8 million hours, almost the equivalent of 9,100 extra physicians working full time.<sup>6</sup>**


In fact, administration and paperwork remain some of the areas with the fewest use cases for current AI, according to MD Analytics. General practitioners and specialists agree that within the next five years<sup>3</sup>, they want to see AI services that help them write emails, letters, and patient handouts, as well as support with translation, during and after consultations.



More holistically, it's necessary that technology serves to connect the dots and close the gaps in primary care so clinicians and staff are no longer forced to act as the integration layer. While electronic medical record (EMR) systems have yielded significant benefits for clinics, they generally haven't been designed to meet the overwhelming demands facing physicians today. Care teams now have hundreds or thousands of patients and are increasingly expected to retain detailed longitudinal insights, track evolving patterns over time, and independently determine when follow-up or intervention is warranted.

Revolutionizing the EMR with AI is therefore not just a matter of efficiency or convenience. **It's imperative for enduring stability and sustainability for Canada's primary care providers**, as well as quality and continuity of care for patients.

With AI agents, EMR tools can move beyond document creation and information retrieval. It's time to embed tools that can work autonomously to detect trends, recognise and flag care gaps, help prepare before patient visits and follow up after, and coordinate teams. These are the types of functions and workflows that become unmanageable for even the best care teams when dealing with hundreds or thousands of patients. This is why they're excellent use cases for AI agents, as long as they're designed with integrity, deployed with intention, and always include a human in the loop to finalize decisions and authorize actions.



“

The ultimate goal is to reduce the heavy administrative and cognitive load on clinicians. There's so much overhead in navigating software, gathering information, and documenting everything. If we get this right, it's a real paradigm shift — the software works for you instead of you working through the software.

**Aaron Brady**

Director of Product Development  
TELUS Health

## AI Terms 101

**AI Agent:** A purpose-built workflow component that autonomously performs a single, well-defined task (e.g., booking agent, billing agent, patient data agent).

**Multi-Agent System:** Multiple agents working together, each completing discrete tasks within a larger workflow.

**Agentic AI:** A more autonomous, orchestrating layer that coordinates multiple AI agents to achieve an end-to-end clinical or operational outcome – with a clinician always in the loop.

**Generative AI (GenAI):** Tools that produce new content in the form of text, images, video, or other media, based on prompts from users. GenAI is often integrated into AI agents to allow them to interact with users through natural, intuitive language.

**Large language model (LLM):** The underlying algorithms that have ingested huge volumes of training data and been tuned to specific use cases, enabling AI tools to generate novel outputs.

## AI Agents and Agentic AI



AI agents are autonomous components that can reason, make decisions, and take actions to achieve a goal.

Agentic AI is the broader approach where software is designed to operate with this kind of goal-directed autonomy – often coordinating multiple agents and tools to complete complex workflows.

A simple travel analogy illustrates the difference:

**AI Agent:** Like a GPS, you enter a destination, and it calculates the best route adjusting for traffic or roadblocks, but always working toward the specific goal you set.

**Agentic AI:** Like a travel concierge, you share an outcome (“a relaxing beach weekend under \$2,000”), and it coordinates multiple steps: researching options, booking travel, managing changes to deliver a complete plan, while you retain final approval.


# AI agents as a force multiplier for health professionals.

Generative AI and ambient scribes arrived at the perfect moment – a moment when clinicians were overwhelmed and overburdened by the low-value tasks that demand constant clicking, typing, collating, and categorizing as they strived to collect and synthesize patient information.

When properly designed, AI agents represent a paradigmatic shift in clinical workflows. Trained and tuned by medical professionals to perform highly specific tasks with precision and accuracy, they can become trusted collaborators in care delivery. Rather than spending time navigating complex systems, care teams can rely on technologies that actively support and augment clinical work.

Like other AI capabilities, AI agents use statistical correlations between text and concepts to flexibly provide answers and perform actions. Unlike other tools, however, they don't have to be initiated by clinicians. Beyond documentation and transcribing use cases, **this allows AI agents to enhance overall efficiency by identifying patterns and trends** across workflows – helping care teams see the bigger picture rather than only focusing on individual details or data points.

Working continuously through a perceive-reason-act cycle, AI agents monitor for inputs, apply medical protocols and rules, and then autonomously execute their specific tasks. Consider the hypothetical case of a preventive care agent that not only flags overdue screenings and vaccinations across a patient panel, but generates reminders and queues them up for physician approval.



Used appropriately, this allows AI to function as a force multiplier for care teams. Information can be sorted and categorized automatically, trends and patterns can be surveilled, delegated tasks can be routed to the correct care team member, and clinicians can receive the appropriate preparation before appointments. AI agents can do all of this without siphoning the precious time and resources of primary care providers. Clinicians are free to bring their professional knowledge and judgment to bear as they make decisions, set priorities, and invest more of their energy and effort into relationship-based care.

Rather than replacing healthcare professionals, AI agents strengthen them: increasing throughput, reducing burnout, and enabling much more meaningful patient interactions. **Patients could benefit from better outcomes as well, since AI agents can detect gaps and health patterns, enrich longitudinal records, ensure follow-ups, and enhance continuity of care.** By enabling clinicians to do more with less, there's tremendous potential to ensure that more Canadians can access primary care as well.

## The entire health system also stands to gain from the responsible development of AI agents.

These systems can support conformance and alignment of federal, provincial, and territorial priorities. Also, as a source of aggregated, anonymized, and secure data, they could inform population-level health insights, planning, and reporting. They have the potential to enable early identification of care gaps and inefficiencies, scale care delivery without linear increases in the workforce, and uncover opportunities to optimize spending and saving.



# Understanding what AI can and can't be trusted with.

Of course, there are well-documented risks for AI systems that rely on LLMs, including biases and hallucinations.<sup>7</sup> This poses a challenge for implementing autonomous AI processes in any industry, but healthcare in particular demands high levels of accuracy and accountability. It's crucial for AI in this sector to be built with integrity and rigour so that their functions are transparent, defensible, and compliant with data privacy and protection requirements.

Clinicians and physicians know that accountability and explainability are crucial as they incorporate AI systems more frequently into their everyday practice.

Research by MD Analytics shows that 80 per cent of general practitioners and 89 per cent of specialists say they would need any AI to be transparent about how and why it made certain recommendations or claims.<sup>3</sup> Additionally, 82 per cent of general practitioners and 90 per cent of specialists say that in order to feel comfortable using AI, they would need more information about how its performance is monitored.<sup>3</sup>

The most compelling current **use cases for AI agents are narrow, well-defined tasks**, and the successful implementation of this technology will require a risk-stratified approach. There are areas that are low-stakes and low-risk, yet high-value; these include inbox sorting and labelling, visit and patient summaries, billing and coding support, and delegated task routing. Sick notes are another example; this is the type of task that drastically increases administrative burden, as each note takes approximately 10.4 minutes to complete, yet only 24 per cent of physicians are compensated for this extra labour despite writing an average 136 notes annually.<sup>6</sup>





These make compelling use cases for AI agents, as they can provide much-needed cognitive, administrative, and operational support. Prioritising them will help ensure early wins for the application and adoption of this technology in primary care settings.

At present, higher-risk domains include areas such as diagnostic reasoning, treatment selection, prognostication, and clinical decision support, where accuracy, transparency, and oversight are paramount. Advancing AI in these areas requires careful design, clear boundaries, and staged implementation, with a deliberate human-in-the-loop model rather than fully autonomous use.

When developed responsibly, AI can help clinicians see patterns more clearly, surface relevant information, and reduce cognitive burden, provided there is transparency about data sources, model behaviour, and limitations. In this context, clinical judgment and final decision-making must remain firmly in the hands of medical professionals.

## Use case

### Proactive management of chronic kidney disease (CKD) for patients living with Diabetes in primary care

This case study provided by Dr. Daniel Ngui, illustrates how agentic AI can be deployed safely in a high-stakes clinical domain by strictly constraining its role and preserving clinician authority at every decision point.

#### Dr. Daniel Ngui

Family Physician, Clinical Professor in the Department of Family Practice at the University of British Columbia in Vancouver, and Medical Director at Fraser Street Medical in Vancouver

[See Appendix for the full case study.](#)

## AI can transform the EMR into a clinical teammate.

As discussed, clinicians face an overwhelming volume of patient data, incoming documents, and administrative tasks. Navigating this information while delivering high-quality care contributes to cognitive overload and workflow inefficiencies. At TELUS Health, AI within the EMR is evolving from a documentation tool into an intelligent clinical workflow layer to help clinicians see what matters, track meaningful trends, and focus on patient care without taking over clinical decision-making.



AI agents embedded throughout the EMR support this shift. Family History and Social History Agents continuously analyze encounter notes alongside structured patient records, identifying gaps and suggesting updates to help maintain complete, accurate longitudinal documentation without manual input.

AI-powered Inbox Triage automatically matches faxes and incoming documents to new and existing patients in an average of 1.2 seconds, dramatically reducing time spent on manual routing and allowing clinic teams to focus on higher-value clinical and patient-facing work.

Referral workflows and follow-up processes are also being reimagined. When a referral is referenced in an encounter, the Referral Letter Agent drafts a suggested letter and title directly within the EMR for clinician review, edit, and approval, maintaining a human-in-the-loop approach.

Task management agents assign follow-up work to the appropriate team members, helping ensure operational processes are completed reliably while clinicians retain full oversight and control.



# The next-wave of agents has the potential to fundamentally change how clinicians practice.

Still in early development, the Patient Chart Query Agent unlocks past encounters and vast volumes of patient documentation, allowing providers to ask targeted questions of the chart — from medication histories and lab orders to longitudinal trends in PHQ-9 or GAD-7 scores and preventive care gaps — surfacing clinically relevant signals instantly and accurately.

Rather than requiring clinicians to navigate multiple sections of the chart, these agents synthesize longitudinal context in real time. This enables earlier identification of emerging risks, more precise monitoring of patient progress, and more confident clinical decision-making during the encounter. In this model, AI doesn't replace clinical judgment; it strengthens it by ensuring the right information is visible at the right moment.

Early experimentation with billing agents further demonstrates how AI can enhance clinical workflows without adding friction. These agents identify appropriate billing codes, calculate allowable units, and surface encounter optimization opportunities, helping ensure physicians are reimbursed for services provided while reducing administrative cognitive load.

Taken together, these capabilities represent a shift from the EMR as a passive system of record to an active clinical teammate. By continuously organizing information, surfacing meaningful signals, and reducing administrative and cognitive burden, embedded AI agents help clinicians stay focused on delivering high-quality care while maintaining full control and accountability.



Our vision for the EMR is to create a clinical teammate that leverages agentic AI, working quietly in the background to help providers and clinics with administrative work so they can stay on top of everything and spend more time with patients. Traditionally, EMRs were like filing cabinets. We want to flip that so the EMR helps with cognitive load, patient care, and the administrative work that keeps clinicians busy.

## **Victoria Phan**

Senior Product Manager  
Strategy & Innovation - Collaborative Health  
TELUS Health


## Advancing toward agentic AI requires trust and security.

AI can't be integrated into healthcare systems without careful oversight and due diligence. Personal health information is some of the most sensitive data an individual has, and AI technologies should only interact with it if they adhere to the highest standards. MD Analytics research found that 76 per cent of general practitioners and 86 per cent of specialists need more reassurance about it if they're going to use AI in their workflows.<sup>3</sup>

This is why **TELUS Health is proud to be at the forefront of building these solutions.** From a sovereign fibre route that keeps data within the country, to the TELUS Sovereign AI Factory (SAIF)<sup>8</sup> – Canada's fastest and most powerful supercomputer, fully Canadian-controlled and designed for highly regulated environments<sup>9</sup> – TELUS Health delivers the foundation for secure, compliant, high-performance AI in healthcare. Fully domestic development ensures compliance with PHIPA and PIPEDA while supporting ethical, responsible innovation aligned with Canadian societal needs.

TELUS is the first company in the world to be awarded the ISO 31700-1 Privacy by Design certification, granted to organizations that embed powerful, world-class privacy protections into their systems from the ground up. This marked a historic milestone in ethical innovation. It was also the first service provider in North America to join the NVIDIA Cloud Partner Network, reinforcing its leadership in AI infrastructure.<sup>10</sup>

There are also a range of national and global protocols that TELUS has committed to, including Canada's Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems, and the Hiroshima AI Process (HAIP) Reporting Framework.



At the core of it all are the principles and philosophy that drive AI design at TELUS. First, AI should always augment the intelligence and capabilities of users, not replace them. Second, ensure that trust is earned, with privacy and security embedded into every system and process: the inputs that AI is trained on are transparent, and so too are the outputs produced. And third, enable seamless onboarding to AI features for customers, since the entire purpose of AI is to simplify workflows and free up users' time and creativity.

With a longstanding commitment to robust oversight, structured governance, and privacy by design, TELUS is already **trusted by more than 44,000 clinicians to manage their patient healthcare records.**



“

Many public cloud services are subject to foreign jurisdiction, which can raise questions about data access even when data centres are on Canadian soil. With our sovereign AI factory in Canada, TELUS can run AI workloads on Canadian-controlled infrastructure, under Canadian law, giving governments and health systems stronger assurances around privacy and data sovereignty.

**Mohamed El-Demerdash**

President  
TELUS Health

“

Technology should adapt to the clinicians, not the other way around. That is why we embed AI directly into the workflow, making it simple to adopt and quick to deliver value, all while remaining trusted and secure. Our breadth of knowledge and experience across EMRs, pharmacy, and claims helps us to build AI that reduces administrative burden without adding complexity for care teams.

**Ratcho Batchvarov**

Vice-President  
Provider Solutions  
TELUS Health

# Unlocking a future of connected, AI-powered care in Canada.

Canadians deserve care that is timely, proactive, and centred on the whole person — not reactive or episodic. Achieving this requires a fully connected healthcare ecosystem, with secure, governed data flow across EMRs, pharmacy systems, virtual care platforms, insurers, and health authorities.

Such an ecosystem benefits care providers as well as patients. Enabling interoperability across data and records has been noted as a leading solution to administrative burden by 71 per cent of physicians.<sup>6</sup>


The enabling condition for this future is interoperability and trust. AI cannot deliver system-level impact if it is confined to isolated tools or weakened by governance gaps. At TELUS Health, the approach is to build and partner in ways that create an interoperable ecosystem, ensuring innovation reaches real clinical workflows quickly and meaningfully.

## When AI can add value

In the near term, this means focusing AI where it can deliver immediate value by reducing administrative burden, supporting patient navigation, and strengthening coordination within care teams. Over time, it means enabling secure, governed data flows that support seamless collaboration across settings and disciplines while meeting Canadians' expectations for privacy, safety, and accountability.

**True progress is simple. Technology should adapt to humans, not the other way around.**

AI succeeds when it's trusted, secure, and embedded in workflows; when it connects points of care to improve coordination and continuity; when it removes repetitive tasks so clinicians can focus on complex decision-making; and when it closes gaps so that no patient or process falls through the cracks.



Agentic AI is not about replacing clinicians or accelerating care at all costs. It is about making care safer, more reliable, and more humane by reducing unnecessary cognitive burden and enabling teams to function at the top of their scope.

### Dr. Daniel Ngui

Family Physician, Clinical Professor in the Department of Family Practice at the University of British Columbia in Vancouver, and Medical Director at Fraser Street Medical in Vancouver

In this future, primary care physicians, pharmacists, and the broader care team collaborate seamlessly across a shared, continuously updated view of the patient. The EMR becomes the keystone of this ecosystem; not simply storing information, but actively helping teams coordinate care, surface risks earlier, and support better decisions across the care journey.

The promise of agentic AI is not faster documentation or smarter automation. It's better care experiences for patients and more sustainable careers for clinicians. It's a future where providers spend less time navigating systems and more time helping people heal, manage chronic illness, and live healthier lives.

If Canada can build this future responsibly, with trust, safety, and patient-centred design at its core, agentic AI will not just change how healthcare is delivered. It will help redefine what patients and clinicians can expect from the healthcare system itself.

[Learn more about TELUS Health AI-powered features](#) →



The decisions we make now about how we design and embed agentic AI within our EMR ecosystem will define the next generation of care. Done thoughtfully, we can create a continuously learning network that connects clinicians, pharmacists, and patients through shared intelligence — improving outcomes while preserving the human judgment at the heart of medicine.

**Michael Cohen**

Vice-President of Products  
TELUS Health

# Appendix



A theoretical example of agentic AI in chronic kidney disease care for persons living with Diabetes: clinical principles for the safe use of agentic AI

**Use case:** Proactive management of chronic kidney disease (CKD) for patients living with Diabetes in primary care<sup>11</sup>

**Clinical context:** Team-based care guided by national recommendations (e.g., Diabetes Canada CKD and KDIGO Guidelines)

This case study provided by Dr. Daniel Ngui, illustrates how agentic AI can be deployed safely in a high-stakes clinical domain by strictly constraining its role and preserving clinician authority at every decision point

## Dr. Daniel Ngui

Family Physician, Clinical Professor in the Department of Family Practice at the University of British Columbia in Vancouver, and Medical Director at Fraser Street Medical in Vancouver

### 1. Clinical intent is set by humans before AI is introduced



Before any automation occurs, clinicians meet to establish shared clinical intent. This includes agreement on:

- The strength and limits of the medical evidence for the screening/intervention/care gap
- Which CKD populations are likely to benefit
- How co-morbidity, frailty, prognosis, and patient preference should modify action

Guidelines inform care, but they don't determine it. AI isn't permitted to infer intent. Clinicians are more likely to trust an AI agent if they understood how the Agent was trained, on what information and what were the parameters.

### 2. The workflow is explicitly designed by clinicians, not learned by an algorithm

Clinicians design the workflow, specifying:

- Which CKD stages qualify for proactive recall
- Laboratory thresholds that trigger review eg. Both ACR and eGFR
- Explicit exclusions (e.g., limited prognosis, prior intolerance, patient refusal)
- A defined team member to verify the recalls, who must review outputs before any patient contact

### 3. Agentic AI is assigned a narrow, targeted role with a human in the loop

Within these clinician-defined constraints, agentic AI is used only to:

- Scan the patient panel against predefined criteria
- Generate a draft recall list



- Assemble relevant structured data (ACR and eGFR trends, BP values, classes of medications or documentation of prior trials) and measures like A1c, LDL cholesterol and blood pressure
- Identify missing cardiorenal protective medication's
- Identify co-morbidities where missing cardiorenal protective medication's can change outcomes
- Route outputs to the appropriate clinician or care-team member for adjudication

**The AI Agent should be explicitly prohibited from deciding:**

- Whether treatment is indicated
- Whether competing priorities outweigh benefit
- Whether a patient should be contacted, or when

These decisions must remain with the clinician/teams as **mandatory clinician verification preserves accountability**

Every AI-generated output should be reviewed by a clinician, who:

- Confirms appropriateness
- Removes patients for whom recall would be burdensome or unsafe
- Accounts for prior discussions, contraindications, and preferences
- Documents rationale where relevant

This step ensures that automation never bypasses professional responsibility.

**4 . Team-based care recalls activate only after adjudication**

- Patient outreach could be initiated by anyone of the team, medical office assistants, nurses, or pharmacists
- A visit should be booked for more assessment, patient education and informed shared decision-making
- There should be outcome tracking and review for the team to enhance future campaigns

Here, AI functions as a force multiplier, supporting safe, guideline-based recalls that can be scaled and standardized for consistency

**Why this matters**

This real-world example demonstrates that agentic AI can safely support longitudinal, preventive care only when:

- Clinicians define intent and workflow based on the best medical evidence
- AI is constrained to specific, low-judgment tasks but searches the entire EMR to bring the right information to the care team
- Human verification is mandatory
- Accountability remains clear and human-owned

In this model, AI assists and clinicians decide. The result is safer automation, higher specificity, and care that remains relational, contextual, and both scientifically and ethically grounded.

## Reference list

1. Zafar, A. (2025). "More Canadians have a family doctor, but 'concerning gaps' found". CBC News. Retrieved from <https://www.cbc.ca/news/health/primary-care-ourcare-family-doctors-9.7007471>
2. Hannay, C. (2025). "Doctors look to AI to cut burdensome administrative work". The Globe and Mail. Retrieved from <https://www.theglobeandmail.com/business/economy/article-ai-artificial-intelligence-doctors-physicians-hospitals/>
3. MD Analytics (2025). "AI in the Doctor's Office". Independent Study. Retrieved from <https://drive.google.com/file/d/19hT1TkDQC4JF6LH72G-iFcredRa0MrRn/view>
4. Canada Health Infoway (2026). "AI scribe program: Streamlining documentation and reducing administrative burden for primary care clinicians". Retrieved from <https://www.infoway-inforoute.ca/en/featured-initiatives/ai-scribe-program>
5. Goodyear, S. (2026). "Canadian doctors say they're losing 20 million hours a year to unnecessary paperwork". CBC Radio. Retrieved from <https://www.cbc.ca/radio/asithappens/cma-doctor-paperwork-report-9.7063539>
6. Canadian Medical Association & Canadian Federation of Independent Business (2026). "Losing doctors to desk work: Canadian physicians lose 20 million hours each year to red tape". Ottawa: The Association. Retrieved from <https://digitallibrary.cma.ca/link/digitallibrary1478>
7. Ngui, D. & Boivin, M. (2025). "Harnessing Artificial Intelligence (AI) Tools in Primary Care: The Promise of Being Smarter, Safer, and More Present". Canadian Primary Care Today. Vol 3, Issue 3, Fall 2025. Retrieved from [https://canadianprimarycaretoday.com/article/view/3-3-Ngui\\_et\\_al/pdf\\_en](https://canadianprimarycaretoday.com/article/view/3-3-Ngui_et_al/pdf_en)
8. TELUS (2025). "TELUS Opens Canada's First fully Sovereign AI Factory". Retrieved from <https://www.telus.com/en/about/news-and-events/media-releases/telus-opens-canadas-first-fully-sovereign-ai-factory>
9. TELUS (2025). "TELUS Sovereign AI Factory named Canada's fastest and most powerful supercomputer". Retrieved from <https://www.telus.com/en/about/news-and-events/media-releases/TELUS-Sovereign-AI-Factory-named-Canadas-fastest-and-most-powerful-supercomputer>
10. TELUS (2025). "TELUS and League will partner to support Canadian healthcare innovation with sovereign AI". Retrieved from: <https://www.telus.com/en/about/news-and-events/media-releases/TELUS-and-League-will-partner-to-support-Canadian-healthcare-innovation-with-sovereign-AI>
11. Tobe, S. W., Bajaj, S. H., Tangri, N., Jain, R., Pham, T., Beaudin, V., & McFarlane, P. (2025). "Chronic Kidney Disease in Diabetes: A Clinical Practice Guideline". Canadian Journal of Diabetes. Retrieved from <https://guidelines.diabetes.ca/getmedia/a400d52c-e65b-4b59-9193-dac724679284/Ch29-Chronic-Kidney-Disease-in-Diabetes-2025-update.pdf>

