Artificial Intelligence

DRAFT

Purpose

This practice advice provides guidance to registered nurses (RNs) and nurse practitioners (NPs), herein referred to as **REGISTRANTS**,[1] on the safe, ethical and professional use of **ARTIFICIAL INTELLIGENCE (AI)**. The College of Registered Nurses of Alberta (CRNA) AI practice advice document includes a framework that supports informed, evidence-based adoption of AI that aligns with regulatory standards, professional judgement and patient-centered care. AI should enhance—not replace—the human connection at the heart of registrant practice. The purpose of this document is to define key terms and concepts related to AI, outline its associated benefits and risks and introduce eleven guiding principles within a framework. This framework provides practical guidance to support registrants in using AI confidently, safely and ethically in their practice.

"AI SHOULD ENHANCE—NOT REPLACE—THE HUMAN CONNECTION"

Introduction

The CRNA recognizes that AI in health care is rapidly evolving. AI technologies are used in many areas, including but not limited to medical devices, diagnostic algorithms, workflow software, automated documentation and decision-support systems that assist both patients and providers. The integration of AI into registrant practice is expected to accelerate, continuing to transform various aspects of care.

The principles that make the framework remind us that the human-to-human connection remains at the heart of patient care. Meaningful relationships between patients and providers are the foundation of trust, compassion and healing. Every health-care journey is built on connection, creating moments of understanding that bring value to both the patient and the registrant. As AI continues to evolve in what can be considered a second Renaissance, its role in health care should be to strengthen—not replace—this human connection. AI can support health care by taking on repetitive tasks or providing helpful

---

1 Words and phrases displayed in BOLD CAPITALS upon first mention are defined in the Glossary.

insights, allowing providers to spend more time focused on what matters most: the person in front of them.

What is AI?

AI is a branch of computer science focused on developing systems that can reason, learn and perform tasks traditionally requiring human intelligence. AI systems can also analyze and interpret data at a scale and speed that exceeds human capacity.[1] One of AI's strengths is its ability to carry out tasks that resemble human functions, learn from experience and adjust to new information and environments. The field of AI is broad, containing numerous applications that are used in different disciplines.

AI has been used in health care for many years, but the more recent innovations—combined with access to large volumes of data—have enabled the development of more advanced algorithms, driving the rise of **DEEP LEARNING (DL)**.[2] DL, a subset of **MACHINE LEARNING (ML)**, uses multi-layered artificial neural networks to process and analyze complex relationships within large datasets. The depth of these layers enables the system to extract increasingly high-level patterns, whether trained through supervised or unsupervised learning. In health care, DL has become the most widely used and successful uses of AI, driving many of the recent advances in medical applications.

One area of DL is **GENERATIVE AI (GenAI)**—a type of technology capable of creating new content, such as text, images, data, video and audio, in response to user prompts. Through the neural network architectures, AI systems are capable of utilizing DL models to learn patterns from data and generate new content. GenAI can analyze vast amounts of data and provide user-friendly responses through **NATURAL LANGUAGE PROCESSING (NLP)**. These kinds of AI systems have become prominent due to their ability to analyze and utilize data and give responses that are easy to read and access. Still, the reliability of GenAI is dependent on how users prompt (or ask) for a response. Therefore, high-quality prompting with AI systems is essential in order to mitigate an impartial or inaccurate output.

Examples of GenAI include

- ChatGPT, which generates human-like text responses;
- DALL·E, which creates AI-generated images;
- Jukebox, which produces music.

AGENTIC AI, also built on top of DL models, refers to AI systems that utilize autonomous agents—meaning they can make decisions, set goals and take actions in pursuit of those goals, often across multiple steps and with minimal human intervention. These systems are being introduced into the Canadian health-care space by assuming tasks that can be automated but also require a degree of data assimilation and decision-making. Some aspects of human resources, scheduling, patient navigation or even IT operations can be carried out by agentic AI.

While AI continues to evolve in health care, several areas of application are emerging as key focal points and warrant attention.[3]

1. Documentation: AI-powered documentation systems, or AI-scribes, may assist registrants through NLP. NLP and voice recognition allow the AI program to understand, interpret, and generate human language. Therefore, transcription of spoken clinical notes can be obtained contemporaneously and summarized for admission into a patient's medical record.
2. Remote Monitoring: AI technologies may remotely collect, analyze, and interpret patient data through different types of devices, such as wearables or smartphones. This data is shared by patients to be analyzed by the health-care provider in a secure location to determine health outcomes and potential interventions.
3. Personalized Care: AI can support personalized care by providing treatment strategies, which can be used to enhance health outcomes and reduce the likelihood of complications—thereby potentially reducing costs through more efficient use of resources.
4. Disease Detection and Diagnosis: AI may be used to support enhanced abnormality or disease detection in medical imaging, or it may assist in compiling the data from imaging, clinical examination, family history, environmental factors and other considerations for early disease detection.
5. Clinical Training and Education: AI tools can accelerate and enhance clinical education by summarizing current evidence for physicians, medical students and patients—offering accessible background information and the latest insights on interventions.

   In addition to these technologies in health care, other advances include (but are not limited to)

   - assistance with scheduling;
   - error reduction;
   - prediction of patient outcomes, discharge forecasts and repatriation logistics;
   - triage and differential diagnosis assistance, lab or diagnostic test recommendations;
   - genomics and personalized medicine.

AI Associated Risks

While AI holds considerable potential to support various aspects of health care, it is also associated with substantial risk. This section outlines these risks and responsibilities, providing context for the decision-making support found in the CRNA AI Framework.

Bias

AI models are trained on large datasets that include information obtained in books, academic writing, news, public or internal sources and the internet. When a user prompts a response from AI, the response is given based on how the AI was trained using large databases. It will not only analyze what it believes to be the meaning behind the input it

receives (the question asked) but also predict the most relevant words and response based on the patterns it learned. However, the quality and fairness of these responses depend on the data the model was trained on. If the training data contains biases—whether related to gender, ethnicity, culture or otherwise—the AI's outputs may reflect and even amplify those biases. This highlights the critical importance of human oversight: health-care providers must consistently assess AI-generated responses for bias, fairness and alignment with evidence-based care. Users should actively screen for potential bias and inappropriate outputs before applying them in clinical settings.

Reliability

Within GenAI are **LARGE LANGUAGE MODELS (LLMs)**, like ChatGPT or Grok, which are designed to generate natural, context-specific and human-like text responses.[4] LLMs are user-friendly, able to perform real-time knowledge retrieval and as such are particularly useful for processing and generating written information, making them accessible tools for health-care providers and patients seeking information.

Despite their capabilities, non-clinically oriented LLMs carry significant risks in health-care settings. LLMs are prone to producing **AI HALLUCINATION(S)**—incorrect or fabricated content presented confidently as fact. AI hallucination is a phenomenon where an LLM perceives data that is not there and seeks to fill the information gap by giving information or data that is nonexistent in a way that makes it seem believable, true and accurate.[5] Hallucinations may occur during text retrieval or speech-to-text conversion (such as transcriptions for a patient chart), and the responsibility of the user is to verify the accuracy of the AI output. The LLM may even generate false or misleading information, producing outputs that do not accurately reflect reality and/or the patient's true clinical findings.[14] The prevalence of AI error and hallucination is not known, and they may often be difficult to detect.

Additionally, any LLM may overgeneralize health conditions, inadvertently assume one diagnostic criterion relates to a specific health condition or hallucinate nonexistent data and provide inaccurate diagnoses, or similar. These types of issues may cause patients to believe they have diagnostic criteria, or they may inadvertently influence health care decision-making and clinical expertise causing adverse treatment outcomes and even patient harm.[14] As such, AI outputs must align with evidence-based care if they are used to support clinical decision-making.[6]

Registrants should seek to use clinically oriented LLMs for more accurate information in both research and clinical training and stay educated about the reliability of AI tools to support patients who may encounter these issues. AI-generated outputs must be checked for accuracy and validated as soon as possible. If an AI scribe is used for documentation, the registrant must validate the AI-output as soon as possible, seeking to remain contemporaneous in charting. When a registrant is using an AI tool in a workplace and does not have the ability to verify the reliability of that tool or output it provides, they should seek employer guarantees for reliability and trustworthiness.

Data, Privacy and Security

AI is emerging as a tool with the potential to support clinical decision-making by analyzing large volumes of data to detect patterns and trends. By processing patient information alongside broader clinical data, AI systems may assist in identifying treatment options tailored to individual needs, suggest possible diagnoses earlier in a patient's care journey and help anticipate disease progression or complications. These possibilities reflect how AI could complement a health-care provider's clinical reasoning by offering data-driven insights that may not be easily recognized through traditional means. However, this potential is still developing, and AI may still produce hallucinations or generate misleading information, seeing correlations that are not present.

Furthermore, when using AI to assist with automation (agentic AI), charting, disease progression or other possible tools, a patient's age, symptoms and other related factors may be enough information to pinpoint who the patient is, even if a name and personal health number is not used. At this point in AI development, not all AI tools are able to complement the health-care provider because the use of data-driven insights may end up in a breach of health information protection law. Even with good intentions, if a registrant misplaces a patient's protected health information into a public platform, they have risked exposing sensitive patient data. Use of a public AI tool could result in unintentional privacy breaches, which would lead to professional misconduct, or even an unlawful act and violation of the CRNA privacy standards, which are reportable offences.

AI support for clinical documentation requires the AI tool to be compliant with the health information regulation. AI chatbots, such as ChatGPT and similar public-facing models, are not designed to meet the strict privacy, security and regulatory requirements of health-care environments. These models are typically trained on large, publicly available datasets and operate on cloud-based platforms that are not configured to comply with health-care privacy regulation, such as the *Health Information Act* (HIA) in Alberta. Because they are hosted on public servers, they do not have the necessary data encryption, access controls, audit trails, privacy agreements with health-care providers or secure storage protocols required for safeguarding personal health information. The patient information could be stored or accessed outside the registrant's control, putting patient privacy at risk. As a result, using AI tools with patient-identifying information creates significant privacy risks and violates health-care privacy laws.

The Right AI for the Right Task

AI tools are built for different purposes, and using the wrong tool—or the right tool the wrong way—can create risks for patient care, safety and privacy. Remaining human-centric in the use of AI tools means asking if the right tool is being used for the right purpose, in the right setting and is reliable to produce the right output. For example, public LLMs pose a substantial risk to data privacy and security and must never be used for documentation assistance. If a registrant were to use a public GenAI system, such as ChatGPT or Gemini, to help with administrative tasks such as summarizing notes or drafting messages and utilized any of the patients' protected health information, this data would be exposed.

Consider other uses of AI, such as a tool for research. When seeking to fully research a subject, a basic GPT is not appropriate, nor is the AI response given in a Google search, though these can be an initial step. Many AI platforms are usable for deep research (such as ChatGPT-4.5 or Grok's DeepSearch), which will take longer to provide a response, but the results will assist in knowledge collection, especially when prompted to show the thought process and references for critical appraisal.

Responsibility, Accountability and Oversight

There are no AI laws in operation in Canada. However, in 2024, an Air Canada chatbot produced inaccurate information regarding fares to a customer. The Air Canada AI chatbot incident gained global attention as a cautionary example of what can happen when AI-generated content is not properly overseen. In this case, an AI chatbot provided a customer with incorrect information regarding airline policies. When the error was brought to light, Air Canada initially argued it was not responsible for the chatbot's mistake. However, a court ruled that the airline was accountable for the information provided by its AI system, reinforcing a critical principle: organizations are responsible for the outputs of the AI tools they deploy, regardless of whether those outputs were generated automatically or by a human. This incident raised global awareness about the risks of relying on AI and highlighted the need for clear human accountability, ethical oversight and governance in the use of AI systems across all industries, including health care.

While AI technologies have the potential to generate helpful and reliable information, they can also produce content that is partially or entirely inaccurate. For this reason, it is essential that registrants apply professional judgment and always verify the accuracy of AI-generated information before using it in clinical care. Keeping the use of AI human-centric means that the responsibility for decisions affecting patient care remains with the registrant, who is accountable for ensuring safe and evidence-based practice. Additionally, any business or organization using AI chatbots is legally responsible for the information these tools provide, reinforcing the need for careful oversight and ethical use of AI in health care.

Prompting

AI output reliability and trustability is directly correlated to the data it accesses and the prompts the user provides to arrive at the most appropriate output. Prompting AI is an art that should include all of the following: Who (assign the AI an identity or role), Why (provide the context of the query and the end goal of this query), What (give AI specifics as to what the output should sound like), How (tell AI to explain how it arrived at the output and provide resources).

Example: if a registrant is seeking to provide an AI-derived handout on managing post-operative pain at home, the prompt would look something like this:

"You are an RN working in a post-operative unit and you are discharging a healthy patient who will need to manage post-operative pain. The patient has never had surgery before and may not be used to managing pain but as the RN you are giving them something helpful to

read when they are home. Provide three paragraphs that summarize evidence-based, free, and well-known tips for managing pain at home. Keep the summary at a grade 6 reading level with a friendly and educational tone. Do not speak about medication management but use other tools such as ice and rest. Include a section that helps the patient know they can reach out to their health-care provider for help. Show your thought process and provide citations that specify where you obtained the information provided."

The final prompt, "How," empowers the AI user to check the thought process and citations given for accuracy, reliability and trustworthiness. Always critically appraise the sources provided and scrutinize the output for bias, stereotype, reliability, veracity and use critical thinking to edit and adapt the output.

Other prompting techniques include giving AI context in learning and examples of the kinds of information one seeks. For example, if a registrant is seeking to understand aspects of medication interactions, give the AI examples of credible, evidence-based literature already available. These sorts of prompts can help the AI provide better responses.

Transparency and Consent

A 2024 scoping review examined 37 articles exploring how patients view the integration of AI in health care.[7] The review found that many patients are cautious about the use of AI in health care, expressing concerns about informed consent, regulatory oversight and the overall trustworthiness of these technologies.[3] Patients also wanted to understand what steps they could take if an AI tool made an error in their care. In addition, they were interested in knowing which AI tools would be used, how those decisions would be made, and how their personal health information would be protected.[3]

Transparency and consent mean that a patient is informed of what AI tools will be used, for what purpose, how AI will use patient information, how this will benefit their health-care journey and of their right to opt out at any point during the care interaction. For example, if an automatic scribe is being used and the patient has consented at the beginning of the interaction, they may decide to refuse this automation during a clinical exam at the point where they feel the need to disclose sensitive information. At this point, another method of documentation will be used and the patient's rights honoured.

Clinic Owners and AI Adoption

Before introducing any AI technology into a clinic, it is important to understand how the tool is classified, what clinical evidence supports its use and what limitations it may have. Having this knowledge helps ensure AI is used responsibly, effectively and in line with both regulatory and professional standards.

For those seeking to adopt AI software, it is important to understand that not all AI systems available globally are designed to meet Canadian health-care requirements.[3, 9, 10] Some technologies may not comply with the safety, privacy or quality standards set by Canadian

regulations. In Canada, Health Canada is the federal authority responsible for regulating medical devices, and software as a medical device, including those that use AI. Health Canada released the draft guidance: "*Pre-market Guidance for Machine Learning-enabled Medical Devices (MLMD)*," and has had the "*Guidance Document: Software as Medical Device (SaMD: Definition and Classification*," which helps manufacturers understand how to meet the regulatory requirements when applying for approval of Class I, III and IV MLMD, and integrative SaMD.[8, 9, 10, 13] Only AI systems that align with Health Canada's guidance should be used in clinical practice. In addition, AI adoption must comply with the *Personal Information Protection Act* and the HIA.

In the age of cybersecurity, AI may heighten existing threats due to its ability to process large datasets and continuously adapt and evolve. Custodians and trustees of personal health information, such as clinic owners, must conduct a privacy impact assessment before implementing new AI technologies and ensure they are compliant and sufficiently secure.[16] In addition, clinic owners are responsible to provide reasonable safeguards, policies, procedures and training for staff.

Finally, clinic owners should ensure their information-management practices respect Indigenous data sovereignty.[14] The Canadian Institute for Health Information has released guidance in this area.[15]

The CRNA AI Framework

The differing types and functions of AI offer numerous enhancements to improving patient care. Whether the support of clinical decision-making, transcription services or robotic intelligence, AI is transforming areas of health care and presenting opportunities and important challenges. As registrants navigate the 'AI renaissance', what practice looks like now may change substantively within a few short years. While seeking to foster safe and ethical adoption, the CRNA also seeks to draw thoughtful attention to the safe and responsible use of AI using eleven principles. The eleven principles summarize and outline registrant responsibility with AI use, adoption and implementation.

Principles

1. Human-Centric

- At the heart of the CRNA AI practice advice document and framework is the commitment to place people first: AI must augment human capabilities, enriching the therapeutic relationships that build trust and facilitate healing in health care.
- AI must never replace human judgment, empathy or oversight, ensuring that decisions remain grounded in professional expertise and respect for individual dignity.
- Registrants should only use AI when this is in the best interest of the patient.

2. Effective and Appropriate

- AI tools should be selected for effectiveness and remain evidence-based to support desirable and meaningful patient or system outcomes.
- Outcomes should be routinely monitored and assessed to ensure AI tools are delivering reliable results and supporting intended goals.
- AI tools are not all designed to perform similar tasks. Registrants should be aware of the tool's intended purpose and its associated outcome, and assess whether it is the right tool, for the right setting, for the right care and the right purpose.

3. Ethical and Legal Use

- Use of AI must align with the code of ethics and standards of practice.
- Registrants should clearly communicate how they are using AI in clinical decision-making and support and be prepared to explain to patients how the tool(s) function, including their limitations, and respect the patient's right to decline or opt out of AI-supported care where appropriate.
- Registrants using AI for transcription services should have a data retention strategy. Once an AI transcription is validated for use in the patient chart, the original data source should be deleted.

4. Competence and Education

- Registrants should acquire AI knowledge and nomenclature. Be aware of the benefits, risks and limitations.
- Registrants at minimum should be competent and proficient in AI tools that have immediate and direct impact to their day-to-day activities. This is the ability to bring a reasonable degree of skill and knowledge to any task related to practice.
- If the registrant is a clinic owner, ensure adequate training is in place for employees.

5. Fair and Equitable

- Registrants should seek to identify and rectify AI biases or stereotypes that contribute to inequities.
- Registrants should ensure that the use of AI tools respect the dignity, values and preferences of patients, communities and providers.
- Wherever possible, registrants should help inform diverse representation in AI testing and validation.
- Wherever possible, registrants should support AI advances in health care for all population groups, ensuring AI deployment does not widen existing disparities.

6. Accurate and Accountable

- The registrant holds primary responsibility for decisions related to patient care.
- Registrants must review, critically appraise and validate any AI outputs for veracity and against bias.
- Registrants should seek to use clinically oriented AI tools for more accurate information in both clinical practice and research or training settings.
- In a situation where a registrant is not able to personally verify AI-outputs seek guarantees from the employer to ascertain reliability.

7. Patient Safety, Privacy, Confidentiality and Consent

- Registrants should always prioritize patient safety, including minimizing risk introduced by AI, such as AI hallucination or error.
- Registrants must never place patient identifying data or information into a public-facing AI tool. Patient data must be used, stored, accessed and transmitted through AI tools that comply with applicable privacy and security laws and the CRNA standards of practice.
- Registrants must obtain informed consent before each use of an AI tool, clearly explaining the purpose, potential benefits, limitations and risks involved. Signage throughout common areas should not be considered informed consent in most cases. Verbal consent is preferred, and patients should understand their autonomy to opt out at any point, even after consent has been given.
- Clinic owners who become custodians and trustees of patient health information are required to
    - implement reasonable safeguards, policies, procedures and training for staff (clinic owners are encouraged to review vendor contractual agreements and protections with legal counsel, such as the Canadian Nurses Protective Society, or otherwise),
    - ensure adoption and implementation comply with CRNA privacy requirements and standards of practice,
    - ensure implemented AI tools are reasonably protected from cybersecurity threats,
    - know how adopted AI systems or tools are classified and meet Health Canada requirements for use and procurement, and
    - know what clinical evidence supports the system or tool and what limitations it may have.

8. Oversight and Collaboration

- Registrants should use AI judiciously. Always combine AI outputs with professional judgment and critical thinking.
- If the registrant is the most responsible person to oversee the use of AI, such as an independent clinic owner who has adopted AI, they should ensure there is an established AI oversight and monitoring program of all AI tools, including those used by employees.
- AI should support, not replace, human judgment, critical thinking or clinical skills.
- Registrants are encouraged to maintain entry to practice skills, such as dosage calculation to confirm reliability of AI outputs or looking for drug-drug interactions, not becoming over-reliant on AI tools.
- Registrants are encouraged to participate in the development of AI tools and systems. While this does not mean they need to write code, their involvement in co-development, refinement, validation and implementation of AI technologies in health care is valuable—especially during this period of rapid AI advancement.[2]

9. Data Governance and Quality

- Registrants are encouraged to advocate for and participate in regular audits of data inputs and sources used in AI systems, ensuring they are accurate, ethically sourced

and aligned with professional standards and code of ethics. Remaining alert to indications of concern (e.g., incomplete information, unaddressed reports of malfunction or security) is part of an interdisciplinary process.

- Registrants are encouraged to collaborate with health authorities, IT specialists and interdisciplinary teams to establish and maintain robust data governance practices, including clear stewardship roles for ongoing data quality management.
- Clinic owners should consider Indigenous data sovereignty.

10. Transparency and Explainability

- Registrants should only utilize AI tools with built-in explainability features to enable AI output review.
- Registrants must be able to disclose use of AI-assisted decision-making.

11. The Iterative Process Surrounding the Framework

- The final piece of the framework is the iterative process: assess, monitor and evaluate. Registrants should consistently engage in the assessment, monitoring and evaluation of AI tools in health care, being proactive to voice concerns, engage in policy development and AI development and adoption.

The CRNA AI Framework

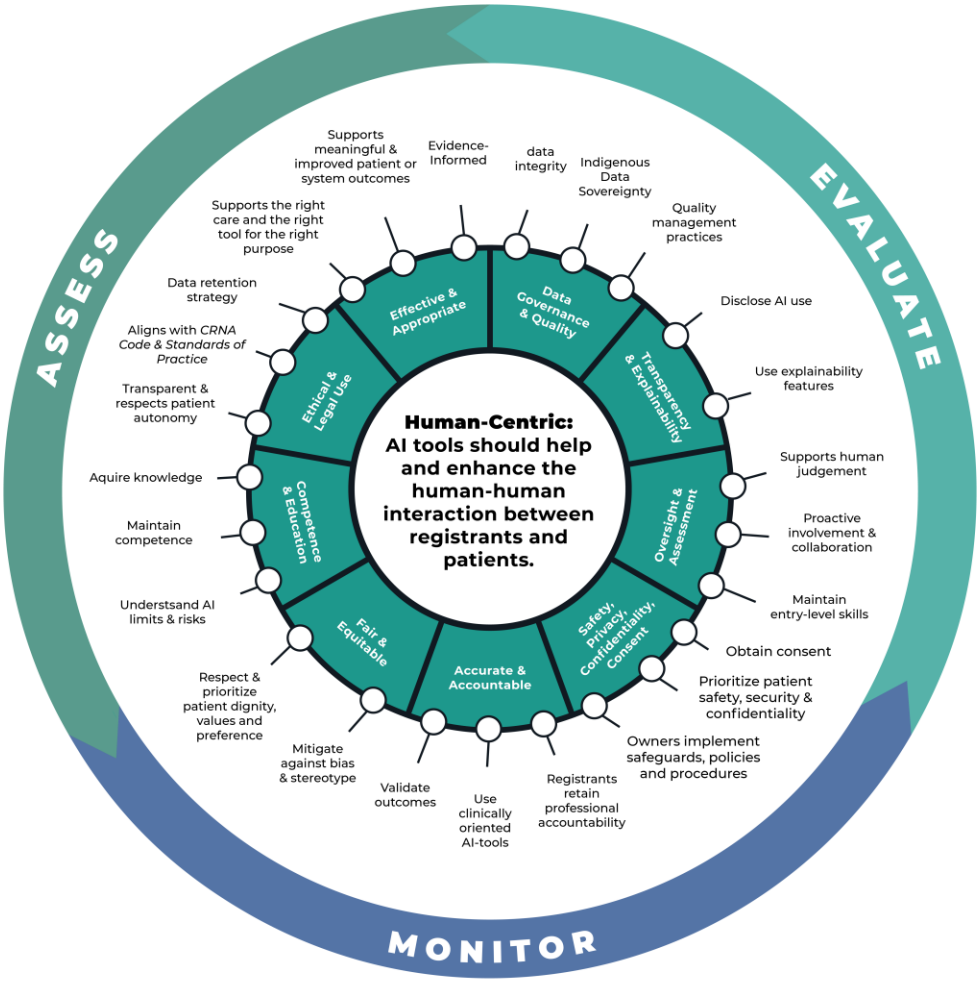

The CRNA AI Framework

The framework outlines eleven principles with human-centric as the focal point. Each principle includes actionable statements promoting safe, ethical, and evidence-informed AI use. The framework is embedded within an iterative process—assess, monitor, and evaluate—which ensures AI tools are responsibly implemented, safely used, and aligned with professional standards and patient needs. Registrants are encouraged to assess AI technologies before use, monitor their performance and impact, and regularly evaluate outcomes. This cycle promotes accountability, protects patient care, and supports the profession's human-centric values.

*Sources that have informed the development of this Framework and the infographic: Government of Canada. (2025)[10]; Health Quality Alberta. (2025)[11]; NextGen Invent Corporation. (2025)[12].

College of Registered Nurses of Alberta

Glossary

AGENTIC AI – refers to AI systems that can act as autonomous agents.

ARTIFICIAL INTELLIGENCE (AI) – A branch of computer science focused on developing systems that can reason, learn, and perform tasks traditionally requiring human intelligence.

AI HALLUCINATION – incorrect or fabricated content presented confidently as fact.

DEEP LEARNING (DL) – a subset of machine learning that uses multi-layered artificial neural networks to process and analyze complex relationships within large datasets.

GENERATIVE AI (GENAI) – a subset of AI that produces audio, visual, code or written text through pattern recognition in large databases and is usually based on deep learning methods.

LARGE LANGUAGE MODEL (LLM) – an advanced computational model of AI trained on vast amounts of text data. It is designed to understand, generate, respond to human language and perform tasks assigned to it. Based on deep learning, a LLM can recognize patterns in language, predict word sequences and perform tasks such as translation, summarization and question answering. LLMs are the foundation for many natural language tools, such as ChatGPT, Google's BERT and others.[3]

MACHINE LEARNING (ML) – a subset of AI that focuses on developing algorithms and systems capable of learning patterns from data and improving their performance over time without being explicitly programmed for each scenario. [3]

NATURAL LANGUAGE PROCESSING (NLP) – a branch of computer science and AI that enables computers to understand, interpret, generate, respond and communicate in human language.

REGISTRANT(S) – Includes registered nurses (RNs), graduate nurses (GNs), certified graduate nurses (CGNs), nurse practitioners (NPs), graduate nurse practitioners (GNPs), and RN or NP courtesy registrants on the CRNA registry.

References

1. Google. "What is Artificial Intelligence (AI)? Retrieved June 30th, 2025. https://cloud.google.com/learn/what-is-artificial-intelligence.

2. Royal College of Physicians and Surgeons of Canada. (2020). *Task Force Report on Artificial Intelligence and Emerging Digital Technologies*. file:///Users/g/Downloads/2020-task-force-report-on-ai-and-emerging-digital-technologies-e.pdf.

3. Canada's Drug Agency, March 2025. Horizon Scan: 2025 Watch List: Artificial Intelligence in Health Care. *Canadian Journal of Health Technologies*. https://www.cda-amc.ca/sites/default/files/Tech%20Trends/2025/ER0015%3D2025_Watch_List.pdf.

4. Jeff Mason, et al., *2022 Health Technology Trends to Watch: Top 10 List*, CADTH, March 2022, https://canjhealthtechnol.ca/index.php/cjht/article/view/er0012/er0012.

5. IBM, 2023. *What are AI Hallucinations*. https://www.ibm.com/think/topics/ai-hallucinations.

6. Topaz, Maxim, et al., 2025. *The Overlooked Dark Side of Generative AI in Nursing: An International Think Tank's Perspective*. Journal of Nursing Scholarship.

7. Moy S, Irannejad M, Manning SJ, et al. *Patient Perspectives on the Use of Artificial Intelligence in Health Care: A Scoping Review*. *J Patient Cen Res Rev*. 2024; 11(1);51-52. doi: 10.17294/2330-0698.2029. PMID: 38596349; PMCID: PMC11000703.

8. Health Canada. 2025. *Pre-market Guidance for Machine Learning-enabled Medical Devices (MLMD)*. https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/application-information/guidance-documents/pre-market-guidance-machine-learning-enabled-medical-devices.html#a1.

9. College of Physicians and Surgeons of British Columbia. (2024). *Interim Guidance: Ethical Principles for Artificial Intelligence in Medicine*. https://www.cpsbc.ca/files/pdf/IG-Artificial-Intelligence-in-Medicine.pdf.

10. Government of Canada. (2025). *Guide on the Use of Generative Artificial Intelligence*. https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/guide-use-generative-ai.html.

11. Health Quality Alberta. (2025). *The Alberta Quality Dimensions for Health*. https://hqa.ca/.

12. NextGen Invent Corporation. (2025). *Decision Making Framework for AI Based Product Businesses*. https://nextgeninvent.com/.

13. Health Canada. 2019. *Guidance Document: Software as a Medical Device (SaMD): Definition and Classification*. https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/application-information/guidance-documents/software-medical-device-guidance-document.html.

14. College of Physicians and Surgeons of Alberta. (2023). *Artificial Intelligence in Generated Patient Record Consent*. Advice to the Profession. https://cpsa.ca/wp-content/uploads/2023/08/AP_Artificial-Intelligence.pdf.

15. Canadian Institute for Health Information. (2020). *A Path Forward: Toward Respectful Governance of First Nations, Inuit and Métis Data Housed at CIHI*. https://www.cihi.ca/sites/default/files/document/path-toward-respectful-governance-fnim-2020-report-en.pdf.

16. Alberta Health, Government of Alberta. (2023). *Completing a Privacy Impact Assessment*. https://open.alberta.ca/dataset/5af38e0a-0aa1-4ac9-8b7a-900d3550d88c/resource/675df883-fefa-4b93-9f2f-bf56f49ded53/download/hlth-completing-privacy-impact-assessment-annotated-template-version-1-1-2023.pdf.