



# Guidelines for the Use of Artificial Intelligence in Canadian Courts

First Edition, September 2024

---

Prepared by Martin Felsky, Ph.D., J.D., and Professor Karen Eltis for the Canadian Judicial Council

## TABLE OF CONTENTS

<b>TABLE OF CONTENTS</b> .....	<b>2</b>
<b>PREAMBLE: NO DELEGATION OF JUDICIAL DECISION-MAKING</b> .....	<b>3</b>
<b>INTRODUCTION</b> .....	<b>4</b>
1. AWARENESS AND EDUCATION .....	5
2. CAUTION.....	5
3. LEADERSHIP .....	5
<b>GUIDELINES</b> .....	<b>6</b>
GUIDELINE 1: PROTECT JUDICIAL INDEPENDENCE.....	6
GUIDELINE 2: USE AI CONSISTENTLY WITH CORE VALUES AND ETHICAL RULES .....	7
GUIDELINE 3: HAVE REGARD TO THE LEGAL ASPECTS OF AI USE .....	7
GUIDELINE 4: AI TOOLS MUST BE SUBJECT TO STRINGENT INFORMATION SECURITY STANDARDS (AND OUTPUT SAFEGUARDS).....	8
GUIDELINE 5: ANY AI TOOL USED IN COURT APPLICATIONS MUST BE ABLE to PROVIDE UNDERSTANDABLE EXPLANATIONS FOR THEIR DECISION-MAKING OUTPUT.....	8
GUIDELINE 6: COURTS MUST REGULARLY TRACK THE IMPACT OF AI DEPLOYMENTS .....	9
GUIDELINE 7: DEVELOP A PROGRAM OF EDUCATION AND PROVIDE USER SUPPORT.....	9

## PREAMBLE: NO DELEGATION OF JUDICIAL DECISION-MAKING

*“It is fundamental to the independence, impartiality and integrity of the judiciary for a judge to exercise the powers of office without undue or unauthorized reliance upon non-judges.”<sup>1</sup>*

Judges hold exclusive responsibility for their judicial decisions. It must be unequivocally understood that no judge is permitted to delegate decision-making authority, whether to a law clerk, administrative assistant, or computer program, regardless of their capabilities.

Nonetheless, judges are encouraged to leverage available support systems to assist in their judicial responsibilities, including, for example:

- Consulting with a law clerk on legal queries.
- Requesting an administrative assistant to proofread and format draft decisions.
- Utilizing grammar and spell-check features in word processors.
- Employing speech recognition tools for dictation.

Such activities should not be misconstrued as judicial decision-making.

In recent years, the landscape of desktop and mobile technology has evolved significantly with the advent of artificial intelligence (AI).

AI encompasses computer systems that simulate human intelligence and decision-making processes. A specific subset of AI, known as generative AI, can autonomously create original content, such as text, images, video, and music, based on the data it has been trained on.

One type of generative AI is based on large language models (LLMs), which are designed to process, understand, and generate human speech and written language. They are powerful tools for researching, organizing, editing, translating and otherwise improving written work. Popular word processing and email applications from Google and Microsoft now include generative AI out of the box. Generative AI is in fact becoming a competitive feature of many platforms, systems and apps.

Prominent publicly accessible LLM engines in use today include ChatGPT (OpenAI), Claude (Anthropic), Gemini (Google), and LLama (Meta). Additionally, there are many AI-driven applications tailored for specific subject matters or writing tasks.

---

<sup>1</sup> New York State Commission on Judicial Conduct, “Delegation Authority,” [2019 Annual Report](#) pp 22-24.

Some Canadian judges have already embraced AI tools to improve their efficiency and accuracy, while others may be employing generative AI without realizing it.

The purpose of these Guidelines is to raise awareness of the risks of using any form of artificial intelligence in court administration and judicial decision-making and to prevent the delegation of decision-making authority while encouraging the safe, effective and appropriate uses of AI by the judiciary.

## INTRODUCTION

These Guidelines aim to provide Canadian judges with a principled framework for understanding the extent to which AI tools can be used appropriately to support or enhance the judicial role. The objective is twofold: to establish a rationale for a consistent approach to the utilization of AI in Canadian courts and to shed light on both the opportunities and risks inherent in AI's potential incursion into court administration and judicial decision-making. We hope this guidance can also serve as the basis for a more standardized approach to judicial education across all jurisdictions in order to foster continued competence as these systems evolve.

When it comes to information governance and technology, the current model of court administration in Canada exists in an ambiguous landscape. While courts maintain exclusive jurisdiction over their own records, necessary collaboration between the judiciary and the executive results in shared control over many aspects of technology and court information management. Although the Council has, for more than two decades, asserted that court information governance is a judicial function, the clarification and establishment of appropriate policies, responsibilities, and definitions remain contentious issues for some courts. To address these concerns, the Council has been issuing guidance on these matters since 2002, commencing with the [Computer Monitoring Guidelines](#), which were subsequently incorporated into the initial 2004 Blueprint for the Security of Judicial Information two years later.

Presently, we face a qualitative departure from the status quo, even as courts are moving registries to the cloud, and developing robust remote hearing capabilities. Generative AI, unlike programmed automation, holds profound implications for the adjudication process itself.

Formulating AI guidelines at this juncture necessitates a delicate balance between embracing innovation and acknowledging that many of the ramifications of this potent technology remain uncharted. Established values and fundamental principles must not be discarded solely due to the pressures of modernization. At the same time, AI shows promise in providing innovative solutions to longstanding practical challenges within the justice system.

Consequently, this modest endeavour proposes to adhere to three guiding principles: (1) Awareness and Education, (2) Caution and (3) Leadership.

## 1. AWARENESS AND EDUCATION

The adoption of AI cannot be a passive or reactive process. Some forms of AI are already embedded in everyday judicial applications for tasks such as translation, grammar checking, speech recognition and legal research. As generative AI becomes more prevalent, it becomes imperative that judges appreciate the implications, limitations, evolving risks, and mitigation strategies associated with its use.

## 2. CAUTION

AI can be entertaining, appealing and, at times, habit-forming. This impressive performance does not, however, mean that it is infallible. Combined with social and psychological factors, this appearance of infallibility can induce misplaced trust and undue reliance, especially as the forms of interacting with AI move from text to chat, to humanoid robot, and to forms as yet undiscovered. Judges must exercise due caution in proportion to the potential impact of AI, which as yet is not fully known.

Trust should coexist with a commitment to verification to the greatest extent possible – if at all possible.<sup>2</sup> Furthermore, beyond the surface appeal of AI lies a realm of potential legal intricacies. Even when AI output proves accurate and valuable, though, its use, particularly in the case of certain generative models, may inadvertently entangle judges in legal complexities such as copyright infringement. Beyond copyright, the extent to which AI has accessed and utilized specific training data – such as personal information used without consent – can significantly impact the legal considerations at play.

## 3. LEADERSHIP

Leadership entails proactive strategic planning, policy development, and the safeguarding of judicial independence. The requirement to maintain an independent judiciary must take precedence in any decisions related to AI procurement, deployment, configuration, and utilization, especially as this relates to the high-risk area of judicial decision-making.

As AI disrupts every corner of our society, it is imperative for judges to adapt, while maintaining their determined commitment to the principles of judicial independence, fairness and access to justice.

---

<sup>2</sup> As AI models evolve, human verification of content may not be achievable. Determining the legality of AI-generated content, such as respecting copyright and privacy laws, adds another layer of complexity to human verification.

## GUIDELINES

We fully expect these Guidelines will need to be continuously updated and refreshed as the technology changes and strategies for its deployment mature.

### GUIDELINE 1: PROTECT JUDICIAL INDEPENDENCE

Parliament cannot empower a state agency, comprised neither of officials from the legislative nor the judicial branch, to oversee the use of AI by and before constitutionally independent courts. Even as governments move forward with legislation governing the use of AI, judicial independence must be preserved.

Canadian legal and normative instruments such as the *Directive on Automated Decision-Making*,<sup>3</sup> the *Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems*<sup>4</sup> and the proposed *Artificial Intelligence and Data Act*<sup>5</sup> are intended to regulate the use of generative AI by government agencies and in administrative decision-making.

The planning, procurement, and deployment of generative AI in Canadian courts must firmly uphold the fundamental principle of judicial independence, encompassing its individual and institutional dimensions. This effort involves all stakeholders, including the executive branch of government, court administration and judicial leadership.

Numerous generative AI applications proposed for courts, including case management systems and alternative dispute resolution, harbour the potential to erode judicial agency and independence. For instance, the selection of training data sources for neural networks exerts a fundamental influence on the output they generate. AI systems honed for optimal efficiency or profitability may not always align with the broader public interest if design choices are made by commercial entities, or subject to substantive influence from public purchasing authorities.

Furthermore, placing too much reliance on any proprietary AI (whether commercial or publicly funded) could compromise judicial independence. This vulnerability highlights the critical need for judicial leadership and involvement in the process of AI adoption.

It must also be understood that the role of judges transcends the resolution of individual disputes. It encompasses the crucial task of interpreting the law, moving the law forward, and providing a stable, separate third branch of government. At the institutional level, a scenario could arise in which a government introduces an AI-powered dispute resolution platform in an effort to reduce case load backlogs. Despite the best of intentions, this could erode the jurisdiction of an

---

<sup>3</sup> [Government of Canada Directive on Automated Decision-Making.](#)

<sup>4</sup> [Innovation, Science and Economic Development Canada Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems.](#)

<sup>5</sup> [Bill C-27, Digital Charter Implementation Act, 2022](#) (which includes the *Artificial Intelligence and Data Act*).

independent judiciary and exacerbate a situation that already poses some risk to the role of public courts. Judges must be empowered to monitor developments in dispute resolution efforts established outside the confines of the justice system.

## GUIDELINE 2: USE AI CONSISTENTLY WITH CORE VALUES AND ETHICAL RULES

Any consideration of the use of assistive AI by judges should always be consistent with the core values of the court<sup>6</sup> and judicial ethics.<sup>7</sup> Together these include independence, integrity and respect; diligence and competence; equality and impartiality, fairness, transparency, accessibility, timeliness and certainty.

Following these precepts would ensure judges avoid introducing biases, which can inhabit AI outputs, promote AI tools that are accessible to all individuals, including those with disabilities, and that they do not inadvertently exclude any segments of the population or inadvertently perpetuate bias against anyone, including marginalized groups.

## GUIDELINE 3: HAVE REGARD TO THE LEGAL ASPECTS OF AI USE

Court administrators and Chief Justices must ensure that the integration of AI into any court process must consistently adhere to applicable laws, including those governing privacy, intellectual property, and criminal activities. Courts should be particularly attentive to the nature of the source material used to train proposed AI systems, aiming to strike an optimal balance between safety and accuracy.

Generative AI is not mystical; it derives insights from pre-existing content, often sensitive data entrusted to the courts by litigants for safekeeping, not for commercial purposes. Concerns over the legality of some practices involving generative AI have gained attention. For instance, allegations have arisen (and lawsuits filed) over the use of copyrighted material without proper authorization. Courts must exercise vigilance in avoiding inadvertent complicity in practices like data scraping, whether from privately-held copyrighted sources or otherwise protected public data repositories.

Moreover, the content used in training data for widely used generative AI models may have been procured under potentially unlawful circumstances, a concern that courts cannot overlook. For instance, content classified as hate speech, which might enjoy protection in some jurisdictions like the US, may contravene Canadian law. Other examples include the European Union General Data Protection Regulation (GDPR), and the *Artificial Intelligence Act*, which is likely to have

---

<sup>6</sup> Canadian Judicial Council, [Policy Framework to Accommodate the Digital Environment](#) (Discussion paper, 2013)

<sup>7</sup> Canadian Judicial Council, [Ethical Principles for Judges](#) (2021)

relevant extraterritorial impact.<sup>8</sup>

#### GUIDELINE 4: AI TOOLS MUST BE SUBJECT TO STRINGENT INFORMATION SECURITY STANDARDS (AND OUTPUT SAFEGUARDS)

For court administrators and Chief Justices, the integration of AI tools in a court brings unique information security challenges. For example, there is a risk of AI algorithms inadvertently exposing sensitive training data, for example information from sealed court files. There is always some risk of unauthorized tampering with AI algorithms or training data to influence outcomes or to inject malicious behavior. Access to external, third-party AI systems can also introduce vulnerabilities. Ensuring the security of these components is crucial to prevent supply chain attacks. The best defence is to have in place a robust information and cybersecurity program, and to give special consideration to addressing additional AI-specific threats.<sup>9</sup>

Individual judges must also recognize and endeavour to prevent the security and privacy risks associated with using generative AI. Uploading a draft judgment, or any sensitive or personal information to a free AI editing or translating website brings with it serious privacy implications. Videoconferencing platforms can now generate high-quality transcripts and summaries of confidential meetings. Judges must think about who has access to these, and how well secured they may be.

#### GUIDELINE 5: ANY AI TOOL USED IN COURT APPLICATIONS MUST BE ABLE TO PROVIDE UNDERSTANDABLE EXPLANATIONS FOR THEIR DECISION-MAKING OUTPUT

Backlogs in court scheduling represent a significant access to justice problem. Canadian courts are investigating and will be experimenting with the use of AI to improve efficiency in case management, alternative dispute resolution and other internal and public-facing areas.<sup>10</sup>

*Explainability* is the need for such AI tools to provide clear, understandable explanations for their output, making it easier for users (and for those affected) to interpret, trust, contest, or accept AI output in critical workflows. Explainability is akin to the requirement for judges to provide reasoned explanations for their decisions in law.

In a broader societal context, explainability fosters accountability. It enables external scrutiny of AI systems, whether by legal experts, ethics committees, or the public, to ensure adherence to

---

<sup>8</sup> The *Artificial Intelligence Act* (AI Act) is a European Union regulation concerning artificial intelligence (AI). It establishes a common regulatory and legal framework for AI within the European Union (EU). It came into force on 1 August 2024, with provisions coming into operation gradually over the following 6 to 36 months. [Wikipedia](#). (Footnotes and links omitted.)

<sup>9</sup> Courts are referred to the [Blueprint for the Security of Court Information \(7<sup>th</sup> edition, April 2024\)](#), in particular Policy 8.

<sup>10</sup> For example, see [Federal Court Strategic Plan 2020-2025](#): “At this point in time, A.I. is not being considered to assist with the adjudication of contested disputes. Rather, the Court is exploring how A.I. may assist it to streamline certain of its processes (e.g., the completion of online “smart forms”) and may be a potential aid in mediation and other types of alternative dispute resolution.”



ethical and legal standards. By making AI's functioning and impact explainable, the public can better navigate the ethical and legal complexities that AI raises, even as it enjoys improvements in access to justice.

As a side note, the explainability of AI systems being used in administrative decision-making has been subject to judicial review: an immigration decision made with input from generative AI was challenged in the Federal Court.<sup>11</sup>

#### GUIDELINE 6: COURTS MUST REGULARLY TRACK THE IMPACT OF AI DEPLOYMENTS

Before introducing AI into any court, administrators must perform a comprehensive, formal, and impartial assessment of its impact on judicial independence, workload, backlog reduction, privacy, security, access to justice, and the court's reputation. This assessment should also encompass an evaluation of how the court's information might be utilized by others if accessed by an AI provider. Best practices advocate commencing with a pilot project or establishing a controlled testing environment, known as a sandbox, to allow users to assess AI's capabilities without incurring the risks of a full-scale deployment. Impact assessments should be done on a continual basis to keep up with the dynamic development cycle of AI tools and include projected long-term impacts.

#### GUIDELINE 7: DEVELOP A PROGRAM OF EDUCATION AND PROVIDE USER SUPPORT

Given the complex nature of AI systems, the judiciary must possess advanced knowledge to identify potential concerns. Two indispensable measures come to the forefront: the training of judges and the provision of technical support for AI integration in court administration. Judicial education is a prerequisite for upholding and maintaining independence.

AI should not be employed without users undergoing a comprehensive educational process and understanding best practices for interacting with the technology, whether as a standalone service or an integral component of court software.

---

<sup>11</sup> [Haghshenas v. Canada \(Citizenship and Immigration\), 2023 FC 464](#). "Whether a decision is reasonable or unreasonable will determine if it is upheld or set aside, whether or not artificial intelligence was used. To hold otherwise would elevate process over substance." Per Brown J at para. 24.