Canadian Judicial Council
Conseil canadien de la magistrature

Guidelines for the Use of Artificial Intelligence in Canadian Courts

September 2026

Second Edition

Prepared for the Canadian Judicial Council by Martin Felsky, PhD, JD

TABLE OF CONTENTS

Introduction **3**

Section 1: Principles **4**

Principle 1: Judicial Decision-Making Authority Cannot Be Delegated **4**

Principle 2: Judicial Independence Must Be Preserved **4**

Principle 3: Judicial Accountability and Public Confidence **5**

Transparency of Usage 6

Explainability and Contestability 7

Verification of Output 7

Principle 4: Judges Must Develop AI Literacy **8**

Broad Scope of Judicial Education 9

Principle 5: Responsible AI Adoption Requires Judicial Leadership **10**

Section 2: AI in Court Administration **11**

Judicial Governance and Independence 11

Assessment Before Deployment 11

Systems Affecting Court Users 12

Cybersecurity and Operational Reliability 13

Continuing Monitoring and Accountability 13

Education and Support 13

AI Governance Frameworks **14**

Privacy and Data Sovereignty 14

Protecting Sensitive Information 15

Automated Language Technologies 15

Conclusion: Looking Ahead **16**

Appendix: Judicial Uses of AI **18**

Examples of Appropriate Uses 18

Examples of High-Risk Borderline Uses 19

Examples of Uses That Are Inappropriate 20

"*Perhaps in a future cyber-court, in session somewhere on Amazon's cloud, a robo-lawyer will beat the speeding ticket that RoboCop issued to your driverless car, all while you go to the beach, and Leibniz's dream of reducing all argument to calculation will finally have come true*."

Pedro Domingos, 2015[1]

INTRODUCTION

These Guidelines are intended to provide Canadian judges with a principled framework for the responsible use of generative artificial intelligence (AI) in the performance of judicial duties. Their objective is twofold: to encourage judges to take advantage of AI where it can appropriately support the administration of justice, while ensuring that its use remains consistent with the constitutional role of the judiciary and the enduring principles that govern judicial office.

This document builds on and replaces the Council's *2024 Guidelines for the Use of Artificial Intelligence in Canadian Courts*. It provides clearer guardrails for distinguishing permissible, borderline and impermissible uses of AI by judges, reinforces the importance of judicial oversight, and provides a direction for the enhanced governance of AI in court administration.

Generative AI systems are now capable of producing sophisticated legal analysis, summarizing lengthy documents, translating text, generating software code, creating realistic images, audio and video, and engaging in what appears to be increasingly complex reasoning. Many of these capabilities are becoming embedded in everyday applications used for personal productivity, such as legal research, document drafting, translation and note-taking.

These developments create significant opportunities. Properly governed AI tools may also assist judges with judicial administration, including by reducing routine workloads and timelines. Used appropriately, AI has the potential to allow judges to devote more time to those aspects of their work that require careful legal analysis, deliberation, discretion and the exercise of independent judgment.

At the same time, AI introduces risks that differ fundamentally from those associated with earlier technologies. Unlike conventional software, modern AI systems generate new content rather than simply retrieving existing information or automating routine tasks. Their outputs may be incomplete, inaccurate, biased, or completely fabricated. Even highly capable systems may confidently propose incorrect legal authorities, misunderstand factual context or omit important

[1] Domingos, Pedro, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, 2015.

considerations. Their output must invariably be approached with informed caution and careful verification.

The challenge for judges today is to harness the power of AI effectively, without compromising their independence, fairness, integrity or sound decision-making.

SECTION 1: PRINCIPLES

PRINCIPLE 1: JUDICIAL DECISION-MAKING AUTHORITY CANNOT BE DELEGATED

Judges hold personal responsibility for every exercise of judicial authority. That responsibility cannot be delegated to human or synthetic intelligence and always remains exclusively with the judge. The Council's *Ethical Principles for Judges*, including independence, integrity and respect, diligence and competence, equality and impartiality, apply equally when judges make use of artificial intelligence. Technology may change how judicial work is performed, but not the ethical standards that govern it.

Judges have always relied upon appropriate forms of assistance in carrying out their work. Consulting with law clerks, directing legal researchers, proofreading draft judgments, using speech recognition software, translation tools, spelling and grammar checkers, or other technologies designed to improve efficiency are all examples of assistance that may appropriately support the judicial role. Such support does not constitute judicial decision-making because the judge is expected to select the appropriate tool, evaluate the work performed, exercise independent judgment and make every substantive decision.

Although AI represents the latest and most sophisticated generation of these supporting technologies, the legitimacy of judicial decisions does not depend solely upon the correctness of the outcome. It depends upon the fact that decisions are reached by an independent judge who has personally considered the evidence, applied the law, exercised judicial discretion and accepted responsibility for the result. Artificial intelligence, regardless of its sophistication, cannot assume these constitutional responsibilities.

PRINCIPLE 2: JUDICIAL INDEPENDENCE MUST BE PRESERVED

Judicial independence is one of the foundations of the Canadian justice system. It protects both the institutional independence of the courts and the individual independence of judges in the exercise of their judicial functions. Decisions concerning the adoption and use of artificial intelligence within Canadian courts must preserve both dimensions of that independence.

Judicial independence requires more than freedom from external interference. It also requires intellectual independence. Judges must guard against becoming unduly influenced by AI-

generated analyses, recommendations or conclusions simply because they are expressed confidently, appear persuasive or are generated quickly. Independent judicial reasoning requires active engagement with the record, the applicable law and the reasoning process itself.

Judges must avoid allowing AI systems to frame legal issues, narrow the scope of inquiry or shape substantive conclusions without careful and independent consideration. AI may identify possibilities for consideration, but it should never determine the direction or outcome of judicial reasoning.

Likewise, courts should retain an appropriate level of institutional authority over AI technologies deployed within the justice system. The judiciary's approach to the selection, implementation and oversight of AI systems should safeguard judicial independence from inappropriate governmental, commercial or other external influences. Canadian legal and policy instruments governing the use of AI by government departments or agencies and private organizations may provide useful guidance in areas such as risk management, transparency and accountability. They do not, however, govern the exercise of judicial authority. Courts occupy a constitutionally distinct position within Canada's system of government, and the adoption of AI for judicial administration and adjudication must remain subject to judicial direction and oversight.

Judicial leadership should play a central role in evaluating proposed AI technologies before they are introduced into the adjudicative process. Considerations should include not only efficiency, cost and functionality, but also their potential impact upon judicial independence, procedural fairness, confidentiality, public confidence and the proper administration of justice.

Artificial intelligence is becoming increasingly integrated into case management, including scheduling, document management, transcription, translation and public-facing registry services. It has also become embedded in the network and cloud infrastructure on which courts depend, supporting essential functions such as system logging, incident management and the detection of operational or security anomalies.

While many of these systems and applications can improve efficiency and access to justice, their cumulative effect should be monitored carefully.

PRINCIPLE 3: JUDICIAL ACCOUNTABILITY AND PUBLIC CONFIDENCE

Public confidence in the administration of justice depends upon judges remaining accountable for the decisions they make. That accountability is achieved through published reasons, independent appellate review, and established mechanisms governing judicial conduct and education. The use of AI should not limit a judge's ability to explain and justify a decision through reasons that reflect the judge's own analysis and conclusions.

Impartiality and Procedural Fairness

Judges should ensure that their use of AI neither creates nor appears to create unfairness. Artificial intelligence systems may reflect biases present in their training data or may perform unevenly across different populations or legal contexts. Judges should therefore consider whether AI-assisted processes could affect the fairness of proceedings, the equality of treatment between parties, or public confidence in the impartial administration of justice.

One of the less visible challenges posed by generative AI is that each tool embodies an implicit "world model": a set of assumptions, values, priorities, and statistical patterns learned from its training data and reinforced through its design and post-training alignment. These embedded perspectives may influence how information is framed, summarized, or prioritized, even when responses appear neutral and authoritative. Such biases are not limited to political or social viewpoints; they may also reflect cultural norms, linguistic conventions, even the political leanings of their developers.

Judges should therefore recognize that AI-generated outputs are not objective representations of reality. They are shaped, often invisibly, by the models that produce them, including the data used to train those models and the choices made in their subsequent refinement and configuration. Different AI systems may produce materially different responses to the same prompt because they embody different world models and have been trained and refined to behave in different ways.

Transparency of Usage

Transparency promotes public confidence while respecting judicial independence, deliberative secrecy and legitimate confidentiality interests.

The use of AI by an individual judge does not ordinarily require disclosure any more than the use of other accepted judicial support tools, such as law clerks, legal researchers, translation software, spelling and grammar checkers or other editorial technologies. Where AI is used only to provide appropriate administrative, organizational or editorial assistance, responsibility for the resulting work remains entirely that of the judge.

Transparency is, however, an important institutional responsibility. Courts should be appropriately open about the AI systems they deploy and the purposes for which they are used. Subject to legitimate security, confidentiality and operational considerations, courts should be prepared to explain how AI systems have been evaluated, approved, governed and secured; the safeguards that apply to their use; the types of functions they perform; and the oversight mechanisms established to ensure that their use remains consistent with judicial independence, procedural fairness and the proper administration of justice.

Explainability and Contestability

Artificial intelligence systems differ significantly in the extent to which they explain how particular outputs are produced. Some systems provide little or no meaningful insight into the reasoning, sources or data underlying their responses. Even where an explanation is provided, it may be incomplete, inaccurate or generated after the fact rather than reflecting the system's actual internal processes. Another relevant factor is the "sycophancy" syndrome exhibited by some AI platforms, whereby the system tends to reinforce or validate a user's assumptions and stated views rather than critically assessing their accuracy. This behavior may produce responses that reinforce a user's existing views or assumptions rather than critically examining them or presenting relevant contrary perspectives.

Judges must avoid uncritical reliance on AI outputs. They should instead exercise independent judgment, critically evaluating the reliability and appropriateness of any output before using it. While transparency and explainability are useful, they are not always possible; as AI systems grow more complex, even their developers may not fully understand how outputs are generated. Decreased transparency requires increased judicial scrutiny and independent verification. If an AI system cannot explain or support a material factual assertion, legal proposition, or recommendation, that output should be treated with caution, if not outright suspicion, and the ultimate responsibility for it resides solely with the judge who issues it in their reasons, endorsement or other format.

Judges should also consider whether AI-assisted analyses are capable of meaningful contestability.[2] Parties should, where fairness requires, have a reasonable opportunity to understand, challenge and respond to AI-assisted output that may affect the outcome of proceedings to their disadvantage. AI must not be used in a manner that prevents effective scrutiny or impedes the court's ability to explain the basis for its decisions.

The reasons for judgment must remain capable of being understood and defended through conventional legal reasoning. The use of AI must never reduce the transparency, accountability or contestability that are fundamental to the judicial process.

Verification of Output

Judges and the legal profession as a whole have been preoccupied with the astonishing number of court submissions around the world containing "hallucinated" propositions of law, fabricated citations and propositions, now reaching more than 1,900.[3] Less attention is paid to the same issue on the bench: judges who have delegated legal research to a law clerk, researcher, or AI,

---

[2] Readers are encouraged to refer to the UNESCO Guidelines for the use of AI systems in courts and tribunals (2025), which expands on this issue and many others touched on in these Guidelines.

[3] As of August 17, 2026. Worldwide cases of AI Hallucination are tracked and reported here: AI Hallucination Cases Database – Damien Charlotin/.

without sufficient verification. The AI Hallucination Cases Database is tracking 29 such cases so far, including judges, judicial officers and arbitrators. This embarrassing situation has caused some decisions to be withdrawn, others to be appealed or relitigated entirely. In one case, the flawed decision was sent to a judicial disciplinary commission.[4]

The extent of verification required will generally increase with the significance of the AI's contribution. Where an AI system is used to generate substantive legal research or analysis, the time and effort required to verify its accuracy, completeness and reasoning may substantially reduce, or even eliminate, any efficiency gains. This practical reality reinforces the principle that AI is best suited to supporting, rather than replacing, the exercise of judicial reasoning.

Verification may include, where possible:

- reviewing the primary source material
- confirming legal authorities from authoritative sources
- checking quotations and citations
- ensuring factual summaries accurately reflect the record
- considering whether significant issues or authorities have been omitted
- cross-checking output with another AI system, particularly to identify possible errors, omissions, alternative interpretations or issues requiring further investigation
- exercising independent legal judgment throughout the process

A second AI system can provide a useful additional check but should not itself be treated as an authoritative or independent source of verification. AI systems may make similar errors or reproduce the same inaccurate information. Verification should ultimately be based on the original record, primary legal authorities or other reliable sources. Verification becomes increasingly important as AI systems become more persuasive: greater fluency should never be mistaken for greater reliability.

PRINCIPLE 4: JUDGES MUST DEVELOP AI LITERACY

Judges who choose not to use AI should nevertheless have sufficient awareness of its capabilities, limitations and implications to understand AI-related issues that may arise in proceedings, court administration and the work of other participants in the justice system. Besides, it is becoming increasingly difficult to separate AI tools from the suite of programs normally used by judges in their daily work, as courts across Canada are implementing Microsoft

[4] *Fiscalía General de la Nación v. Luis Alberto Contreras Fandiño*, Tribunal Superior de Bogotá, Colombia. "Applicant alleges the lower court's judgment contains non-existent or 'made-up' scientific authorities and academic citations relied on in reasoning (not produced at trial). Applicant alleges the lower court added or relied upon legal/precedential authorities or case citations in the judgment that were not part of the trial record and may be non-existent."

365 Copilot on their integrated platforms. In addition to a familiar interactive chat format, Microsoft's embedded AI tools can edit, rewrite and compose text directly within applications such as Word. Legal research platforms are also increasingly incorporating AI, including Lexis+ AI with Protégé and CanLII's AI-supported research features. Similarly, Thomson Reuters has integrated CoCounsel and generative-AI-assisted research into Westlaw Edge Canada.

Judges who choose to use AI in their judicial work should develop a practical understanding sufficient to use the particular systems available to them responsibly and effectively. This does not require technical expertise or mastery of how AI systems operate internally. Rather, it involves understanding their principal capabilities and limitations, recognizing common risks and failure modes, knowing when outputs require verification, and understanding the applicable requirements concerning confidentiality, security and appropriate judicial use. Practical experience and skills training can also help judges develop more effective ways of interacting with AI systems and better recognize the tasks for which they are most and least reliable. The level of knowledge reasonably required will necessarily depend on the nature and extent of a judge's use of AI and should evolve as new capabilities emerge.

Developing this capacity should not be regarded solely as an individual responsibility. Courts and judicial education organizations play an important role in providing judges with appropriate training, practical guidance and opportunities for hands-on experience. Judicial education can help judges develop broader AI literacy while also providing the more specific knowledge needed by those who choose to incorporate AI into their work.

The objective is therefore not universal proficiency in the use of AI, but an appropriate level of AI literacy across the judiciary, supported by more practical education for those who wish to use these technologies.

Broad Scope of Judicial Education

AI literacy includes topics such as the capabilities and limitations of AI systems, effective prompting and verification techniques, confidentiality and information security. Developments in AI technology and judicial practice, emerging legal issues, and applicable court policies governing AI use are also important. Practical demonstrations and hands-on experience may be particularly valuable. They can help judges understand how different AI systems behave in practice, the variability of their outputs, and the circumstances in which their use is most helpful.

All members of a court, including judicial officers, should also understand confidentiality and cybersecurity risks (including data sovereignty), the implications of vendor terms of service, and the legal and ethical responsibilities associated with AI-assisted judicial work. Continuing education will be essential as these technologies continue to evolve.

PRINCIPLE 5: RESPONSIBLE AI ADOPTION REQUIRES JUDICIAL LEADERSHIP

Artificial intelligence presents the judiciary with an unusual challenge. Unlike previous technological advances, it has the potential not only to improve the efficiency of judicial work, but also to perform tasks that increasingly resemble human reasoning. As these capabilities continue to develop, courts will face growing pressure to expand the role of AI on the bench and in the registry.

Some of that pressure will be justified. Well-designed AI systems may improve access to justice, reduce delay, assist self-represented litigants, and help courts manage increasing workloads despite limited public resources. Governments may understandably ask whether routine or high-volume proceedings could one day be resolved, in whole or in part, through automated decision-making.

Those questions should not be answered by technological capability alone. They require careful consideration of constitutional principles, judicial independence, procedural fairness, transparency, accountability and public confidence. The question is not simply whether artificial intelligence can perform a judicial function, but whether society should entrust that function to a AI, and if so, under what safeguards and with what continuing human oversight.

Accountability for answering those questions belongs first and foremost to the judiciary. Courts have a long history of adapting to technological change while preserving the essential characteristics of judicial decision-making. Judges should neither resist innovation because it is unfamiliar nor embrace it because it is novel. Rather, they should lead its responsible adoption by ensuring that technological advancement remains consistent with the core values of justice in the public interest.

The pace of AI development also calls for humility. Artificial intelligence is advancing rapidly, and systems with capabilities far beyond today's models can no longer be dismissed as merely speculative. Whether those developments prove beneficial, disruptive or otherwise transformative, they underscore a central principle: the legitimacy of the judicial process rests not on the sophistication of the technology used, but on the independence, integrity and accountability of the human judges entrusted with judicial responsibilities. Whatever advances may come, those constitutional responsibilities remain irreplaceable.

SECTION 2: AI IN COURT ADMINISTRATION

As discussed, AI has potential applications throughout court administration including many public-facing functions. These uses are distinct from the use of AI by individual judges in carrying out judicial work. They nevertheless require careful governance because administrative systems can affect access to justice, the treatment of court users, the management of judicial information and, in some circumstances, the institutional independence of the courts.

Judicial Governance and Independence

The adoption of AI in court administration should occur with appropriate judicial leadership and oversight. Decisions concerning the procurement, configuration and deployment of AI systems may involve questions of judicial independence and control over court information even where the systems are intended primarily for administrative purposes.

Particular attention should be given to systems procured or operated through government-wide technology arrangements, including environments shared with external government ministries or agencies. The fact that a system is technically secure or available through an approved government platform does not, by itself, resolve questions concerning judicial independence, confidentiality, access to information or institutional control.

AI systems should also be assessed against the objectives and values of the justice system. Systems designed primarily to maximize efficiency, reduce costs or increase throughput may not necessarily reflect other important considerations, including fairness, accessibility, transparency, procedural rights and the quality of justice. Courts should therefore have an appropriate role in determining the objectives against which AI systems used in court operations are designed and evaluated.

Assessment Before Deployment

The introduction of a significant AI system into court operations should be preceded by an assessment proportionate to the nature of the proposed use and the risks it presents. Depending on the application, relevant considerations include those that have been raised elsewhere, including judicial independence, privacy and confidentiality, information security, procedural fairness, accessibility, equality, workload, access to justice and public confidence.

The assessment should also consider how the system operates, the information to which it will have access, where that information will be processed and stored, who may have access to it, whether information will be retained or used for model training, and the extent to which the court will retain effective control over the system and its data. Appropriate legal, privacy, security and information-governance review should form part of this process.

Where feasible, significant or higher-risk applications should ordinarily be tested through pilot projects or controlled environments before broader deployment. Testing should evaluate the system under realistic conditions and should seek to identify not only whether it performs its intended function, but also foreseeable errors, biases, security vulnerabilities and unintended consequences.

Courts should establish appropriate policies governing the acquisition, approval and use of AI systems, including requirements respecting:

- information classification
- privacy and security assessments
- acceptable use
- contractual protections
- audit and oversight mechanisms
- incident reporting
- periodic review as technologies evolve

Systems Affecting Court Users

Particular care is required where AI systems interact directly with litigants or the public, or influence how individuals access and navigate court processes. Systems used to provide information, assist with forms, schedule proceedings, triage matters or manage cases may appear administrative but can have substantive consequences, particularly where they influence which matters receive attention, how a case proceeds or what options are presented. The significance of a system should therefore be assessed by its practical effects, rather than by whether it is formally characterized as administrative or decision-making.

Safeguards should be proportionate to those effects. Where the operation of an AI system may materially affect an individual's interests or access to justice, appropriate measures may include transparency about the role of AI, meaningful human review, mechanisms to identify and correct consequential errors, and the ability to challenge outcomes influenced by the system.

Where AI has materially assisted in significant administrative decisions, courts may wish to consider whether an appropriate record of the interaction should be preserved for governance, audit or accountability purposes. Depending on the circumstances, this may include relevant prompts, model information and significant generated outputs. The purpose of preserving such records is to support institutional accountability and, where necessary, subsequent review of the appropriate use of AI.

As courts increasingly receive documents electronically, information flows not only from the courts to AI systems, but also from external parties into AI-enabled court processes. AI systems used to review, summarize, classify or otherwise analyze documents submitted by litigants or the

public should therefore be designed and deployed with appropriate safeguards against prompt injection and similar attacks.[5]

Cybersecurity and Operational Reliability

AI-specific risks should be incorporated into existing court cybersecurity, information-governance and risk-management programs, in accordance with the applicable model policies and guidelines from Council.[6]

Continuing Monitoring and Accountability

Approval for deployment is not the end of the governance process. AI systems, including their characteristic models and capabilities, tend to mature over time, and their behaviour could change in an evolving technological environment Courts should therefore monitor significant AI deployments throughout their lifecycle to determine whether they continue to perform as intended and remain consistent with applicable legal requirements, court policies and institutional values.

Monitoring should be proportionate to the nature and risk of the application and may include consideration of accuracy, bias, security incidents, unintended consequences, user experience and the system's actual effects on workload, efficiency and access to justice.[7] Courts should also establish clear responsibility for oversight and for responding when an AI system produces material errors or otherwise fails to operate as intended.

Where appropriate, courts should modify, suspend or discontinue AI deployments that no longer satisfy these expectations.

Education and Support

The introduction of AI into court administration should be accompanied by appropriate education, guidance and technical support. Judges and court personnel should understand, at a level appropriate to their roles, the purpose of the system, its principal capabilities and limitations, the information it uses, and the safeguards and policies governing its operation.

---

[5] Prompt injection occurs when a document contains hidden or embedded instructions intended to influence the behavior of an AI system rather than communicate information to a human reader. Such instructions may be concealed in metadata, formatting, embedded objects or other machine-readable content and may not be readily apparent during ordinary human review. If not properly mitigated, prompt injection may cause an AI system to ignore its intended task, generate misleading outputs, disclose information or otherwise behave in unintended ways. Courts should recognize prompt injection as a cybersecurity and information-governance risk whenever AI systems process documents originating outside the court's trusted environment.

[6] Council's technology related publications are available on the CJC website.

[7] Any monitoring of judicial officers must be compliant with the Council's *Computer Monitoring Guidelines*, 2002.

Institutional support is particularly important because responsible AI adoption cannot depend solely upon individual users identifying risks or developing appropriate practices on their own. Courts should provide clear guidance concerning approved uses, appropriate safeguards and available sources of assistance, and should review that guidance as technologies and institutional experience evolve.

AI GOVERNANCE FRAMEWORKS

Privacy and Data Sovereignty

Judges routinely work with some of the most sensitive information entrusted to any public institution. Court records may contain personal information, privileged communications, commercially confidential material, information affecting public safety, and information protected by statute, court order or the common law.

Courts may increasingly provide judges with access to AI systems operating within secure enterprise, court-specific or government-managed technology environments. Such environments may offer substantially greater protection than publicly accessible consumer AI services, including contractual restrictions on the use of customer data, stronger access controls, audit capabilities and defined data-retention and processing arrangements. They may also be configured so that information remains within a Canadian cloud region or an institutionally controlled environment. These protections can materially reduce confidentiality, cybersecurity and data sovereignty risks.

However, labels such as “enterprise,” “private” or “sandboxed” do not describe a uniform level of protection. The architecture, contractual terms, configuration and governance of each implementation must be understood. In particular, courts should consider whether the environment is controlled exclusively by the court or is shared with executive branch ministries or agencies. Shared environments may raise additional questions concerning access to judicial information, confidentiality, privilege, information governance and judicial independence, even where the supporting technology is considered secure.

The use of confidential judicial information for model training is of particular concern because it represents a secondary use of information beyond the immediate judicial purpose for which it was submitted. Depending on the system and its safeguards, such use may also create a risk that information derived from submitted material could influence model behaviour or be reproduced or indirectly exposed in later outputs.

Court-approved systems used for non-public information should therefore ordinarily provide clear and enforceable assurances that submitted information will not be used to train the available models and will be handled in accordance with judicially appropriate requirements for confidentiality, control and retention.

Administrators and users should have a reasonable understanding of:

- where information will be processed and stored
- how information is protected against unauthorized access, including through encryption
- whether prompts or uploaded documents are retained
- who may access that information
- whether court information is used to train or improve the provider's models
- what contractual, technical and organizational safeguards are in place
- whether the information may become subject to foreign legal regimes
- whether the proposed use complies with applicable court policies and information governance requirements

Judges should not enter non-public information into publicly accessible AI services unless the court has determined that the service provides appropriate safeguards for the information involved. Where court-approved AI systems are available, judges should ordinarily use those systems for judicial work involving non-public information.

Protecting Sensitive Information

Particular caution should be exercised before entering into an AI system:

- draft reasons for judgment
- sealed or confidential court records
- privileged communications
- personal information
- information protected by publication bans or statutory confidentiality provisions
- information relating to judicial deliberations or internal court processes
- any other material that should not be disclosed outside the court's secure environment

Even where an AI provider offers contractual assurances respecting confidentiality, judges remain responsible for ensuring that the proposed use is consistent with their continuing obligations to safeguard court information. Considerations of confidentiality, security and data sovereignty should inform that assessment. The absence of an identified security risk does not necessarily resolve broader governance concerns arising from the processing of Canadian court information outside the judiciary's control.

Automated Language Technologies

Artificial intelligence is increasingly capable of producing highly accurate transcripts, translations and real-time interpretation. These technologies may improve access to justice, reduce costs and increase efficiency, particularly in multilingual proceedings.

Nevertheless, automated language technologies remain susceptible to errors arising from context, legal terminology, dialects, accents and cultural nuance. Appropriate verification remains necessary, in particular where the precise evidence or words spoken are material to the issues before the court.

Judges should distinguish between AI-generated language tools used for administrative convenience and those relied upon as authoritative records of court proceedings.

CONCLUSION: LOOKING AHEAD

Artificial intelligence will continue to evolve, and these Guidelines must evolve with it. Particular technologies, capabilities and risks will change, sometimes unexpectedly. Judicial independence, impartiality, fairness, accountability, transparency and respect for the rule of law should remain the foundation against which new applications of AI are assessed.

The next phase of AI development may, however, present challenges materially different from those addressed by today's conversational systems. The prevailing model of generative AI remains largely based on interaction: a user asks a question, the system produces an output, and the user decides what to do with it. Increasingly, AI development is moving toward more agentic systems capable of planning and carrying out multi-step tasks, interacting with other software, retrieving and processing information, making choices autonomously and taking actions with progressively less human involvement. The interface may therefore shift from asking an AI system to produce a draft or answer toward assigning it an objective and allowing it to determine how that objective should be accomplished.

For courts, this development will make the distinction between assistance and delegation more difficult to draw. An AI agent might eventually manage elements of a case workflow, communicate with parties, organize and classify filings, conduct research, identify procedural issues, prepare materials for hearings or initiate administrative actions. Such systems could provide substantial benefits, but they also raise important questions about authorization, supervision and accountability. Courts will need to consider what decisions an AI agent may make in pursuing an assigned objective; what information and systems it may access; what actions it may take without prior human approval; how its activities can be monitored and reconstructed, and who bears responsibility when an autonomous or semi-autonomous system makes an error or takes an inappropriate action. Governance frameworks designed primarily around reviewing AI-generated content may therefore need to evolve toward governing the authority, access and permissible actions of AI systems themselves.

These developments also broaden the concept of technological sovereignty. Data sovereignty concerns where information is held, processed and subject to legal control. AI sovereignty raises a related but potentially more consequential issue: dependence on the systems that increasingly interpret, organize and act upon that information. Even where judicial information remains

securely within Canada, the AI models used to process it may be developed and controlled elsewhere, with their design, capabilities and continued availability determined by external providers. For courts, responsible adoption of AI therefore requires attention not only to the control of judicial information, but also to the degree of dependence created by the technologies through which that information is used.

No set of guidelines adopted today can anticipate every development that may follow. The task is therefore to establish principles sufficiently durable to guide courts through technological change while remaining prepared to revisit the practical rules governing particular applications. As the capabilities and autonomy of artificial intelligence increase, the central question will remain not simply what the technology can do, but what role it should be permitted to assume within a justice system whose legitimacy ultimately depends upon judicial independence, accountability and public confidence in the administration of justice.

APPENDIX: JUDICIAL USES OF AI

This section provides examples of judicial uses of artificial intelligence organized by level of risk. The examples are intended to assist judges in applying the principles set out in these Guidelines. They are illustrative rather than exhaustive, and the appropriate use of AI will always depend upon the circumstances, the sensitivity of the information involved, and the nature of judicial work. In practice, the examples below are moving targets requiring adaptation and periodic review in each court.

Before relying upon AI in performing a judicial task, judges should consider:

1. What judicial function is being performed?
2. Does the task require the exercise of legal judgment, discretion or independent reasoning?
3. Will I personally review and verify the AI-generated output?
4. Could I fully explain and justify the resulting work through my own reasoning?
5. If the AI output were inaccurate, incomplete or biased, would I be able to detect the problem?
6. Would a reasonable and informed member of the public regard the work as the product of an impartial judge?

If any of these questions gives rise to doubt, judges should reconsider whether the proposed use of AI is appropriate.

EXAMPLES OF APPROPRIATE USES

With court-provided secure AI tools, judges can safely use AI to assist with administrative tasks, including:

- preparing speaking notes, educational presentations or judicial training materials
- drafting and translating administrative correspondence, notices or internal memoranda
- assisting with scheduling meetings and the preparation of agendas
- note-taking, transcription of recorded meetings[8], and preparation of minutes

[8] Judges considering the use of AI-enabled tools that record or transcribe court proceedings should first confirm that the creation and use of a separate recording is permitted by applicable law and court policy. Any AI-generated transcription should ordinarily be treated as a personal note-taking aid, not as an official transcript or substitute for the official court record, and its accuracy should not be assumed where the precise evidence or words spoken are material to the decision.

AI may also be used to assist judges with tasks that support adjudication. These uses generally involve organizing, locating, formatting or communicating information. The judge remains responsible for independently reviewing the output and determining its accuracy, completeness and appropriateness.

Examples include:

- organizing electronic exhibits, transcripts or disclosure into searchable categories
- generating a chronology from an agreed statement of facts or other undisputed evidence
- preparing a procedural history of the litigation from the court record
- summarizing lengthy affidavits, expert reports, transcripts or written submissions for later verification against the originals
- identifying potentially relevant legal authorities or secondary sources for further legal research[9]
- comparing draft reasons or successive versions of documents to identify changes
- improving grammar, spelling, clarity or readability of one's own draft text
- converting material into tables, timelines, checklists or other organizational formats

EXAMPLES OF HIGH-RISK BORDERLINE USES

Some uses of artificial intelligence to support adjudication present significantly greater risks because they involve legal analysis, factual assessment or discretionary judgment. These uses require particular caution, careful verification and, in some circumstances, should be avoided altogether.[10]

Examples include:

- identifying inconsistencies in evidence for later verification against the original
- identifying the legal tests
- identifying sentence ranges and damages award
- translating draft reasons for judgment for later verification

[9] This use should be limited to AI-assisted legal research tools that are designed to retrieve and reference existing legal authorities from authoritative sources. General-purpose AI systems have demonstrated a propensity to generate fictitious or inaccurate case citations ("hallucinations") and should therefore not be relied upon to identify legal authorities unless their output can be independently verified against authoritative legal sources. In every case, the judge should review the original authorities before relying upon them.

[10] As models become more capable, individual examples may migrate between the first and second categories.

These uses do not necessarily preclude AI assistance, but the greater the system's involvement in legal reasoning or judicial discretion, the greater the need for independent judicial scrutiny and restraint.

EXAMPLES OF USES THAT ARE INAPPROPRIATE

Certain functions are so closely connected to the judicial role that they must not be delegated to artificial intelligence.

Examples include:

- entering court documents and information in non-court provided tools (ex., ChatGPT, Claude, Grok, Notability, NotebookLM, etc.)
- entering non-public information into publicly accessible AI services that does not meet the appropriate level of security
- resolving disputed questions of law or fact
- assessing credibility or reliability
- deciding whether statutory or common-law tests have been satisfied
- exercising judicial discretion
- editing, paraphrasing or otherwise altering the wording of testimony, witness statements or other quoted text that forms part of the evidentiary record
- drafting substantive portions of reasons for judgment without thorough independent review
- determining the appropriate disposition or remedy[11]
- deciding guilt, liability or sentence
- researching legal authorities using systems that do not reliably identify authentic sources
- weighing competing evidence
- drawing factual inferences from the evidentiary record
- generating draft findings of fact from documentary evidence and transcripts

[11] These Guidelines do not prohibit judges from using established decision-support tools that perform calculations based on mathematical models or recognized Canadian legal formulas, including tools commonly used in family law. The fact that a software product also incorporates generative AI functionality does not make its non-AI functions inappropriate, provided that the judge's use is confined to those functions.