

Guidelines on the Use of Generative Artificial Intelligence by Judges

PREAMBLE

The rapid emergence of generative artificial intelligence (‘generative AI’) is transforming many sectors, including the legal field. Quebec courts are no exception. Indeed, generative AI tools are now available and are increasingly being used by various stakeholders in the judicial system, including litigants and their representatives.

Recognising this reality, the Court of Appeal of Quebec, the Superior Court of Québec, the Court of Québec and the Municipal Courts of Québec reaffirm that the act of judging remains a fundamentally human function, exercised exclusively by independent, responsible and accountable judges.

Generative AI, however sophisticated it may be, possesses neither judgment nor conscience, and is not subject to any form of accountability. It cannot understand the human, social and legal context specific to each case. It cannot, therefore, under any circumstances, replace judicial reasoning, the assessment of evidence or deliberation.

The act of judging cannot be reduced to a technical exercise. It is an intrinsically human function, based on deliberation, consideration of context, respect for dignity, the requirement of fairness and the exercise of moral responsibility. No algorithm can fully replicate its nature. Artificial intelligence systems are not intelligent in the human sense of the term. They merely imitate certain forms of human reasoning.

Furthermore, these tools rely on algorithms whose inner workings remain largely opaque. This opacity alone justifies the cautious approach called for by these guidelines.

In the current context, characterized in particular by the absence, to date, of institutional generative AI tools that are explicitly authorized, secure and deployed within the environment of each court, it is necessary to adopt a cautious, responsible and consistent approach, based on the fundamental values of justice, the protection of judicial independence and the maintenance of public confidence.

These guidelines establish a common framework for the four courts for the potential use of generative AI by judges. They do not encourage its adoption but place strict limits on permitted uses and clearly identify the lines that must not be crossed. They constitute recommendations aimed at: (i) helping judges to make appropriate use of generative AI, in accordance with the courts’ mission and the judicial role; and

(ii) enabling the public to better understand the legitimate uses judges may make of this technology in the performance of their duties.

They form part of an evolving approach, which will adapt to technological and institutional developments, while always preserving the primacy of human judgment and the personal responsibility of the judge.

No judge may delegate, in whole or in part, their decision-making function to a generative AI system, however sophisticated it may be.

Judging is an exclusively human act.

GUIDELINES

1. Purpose and scope

These guidelines set out the framework for the use of generative AI by judges of the Court of Appeal of Quebec, the Superior Court of Québec, the Court of Québec and the Municipal Courts of Québec.

They are aimed at:

- upholding the fundamental values of justice, in particular judicial independence, impartiality, fairness, equality, the protection of human rights and access to justice;
- reaffirming the principle of judicial accountability;
- maintaining public confidence in the administration of justice.

These guidelines govern the use of AI tools, not the judge's working methods. Any information submitted to an AI tool, regardless of its form or the method by which it was collected, is subject to the same requirements of caution: the necessary precautions depend on the sensitivity of the content and the safeguards offered by the tool used, and not on the form, nature or method of collection of the information.

2. Nature and role of generative artificial intelligence

Generative AI is a technological tool capable of producing content based on statistical models derived from training data.

This category encompasses large language models designed to generate natural language, as well as certain transcription or machine translation tools which, in some cases, may be based on specialized generative models. The tools actually used or currently being evaluated within the courts may be listed, for information purposes and subject to change, in a separate accompanying document.

Generative AI:

- does not understand the law;
- exercises no judgment;
- assumes no responsibility;
- cannot independently explain its reasoning in legal terms.

Specialized tools may eventually be used, developed and authorized specifically to support the activities of the courts, provided that they comply with recognized institutional standards. However, to date, no such tool is available, which calls for a cautious approach. It is essential that any tool intended to assist judicial functions be rigorously evaluated to ensure compliance with institutional standards.

In any event, generative AI may only be used as a secondary support tool. It cannot constitute a source of authority, legal reasoning, assessment of evidence or adjudication. In accordance with these guidelines, however, it may be used for certain ancillary, descriptive, organizational or linguistic tasks.

3. Responsibility and accountability

The judge remains fully responsible for:

- legal reasoning;
- the content of the judgments handed down;
- the accuracy of the facts established;
- the accuracy of the references, quotations and authorities cited;
- and the consequences of any use of a generative AI tool.

Judges cannot shift responsibility for an error to the tool.

The use of a generative AI tool does not diminish the judge's duty to verify the facts or their full responsibility for the final result.

4. Non-recommended uses

4.1 General principle

Judges should not use generative AI for any activity directly affecting the exercise of judicial decision-making power.

4.2 In particular, the following are not recommended:

- substantive legal reasoning;
- analysis and characterization of the facts;
- assessment of evidence and credibility;
- determination of the outcome of a dispute;
- the drafting, in whole or in part, of the analytical or substantive reasons for a decision, including legal analysis, the assessment of credibility, the weighing of relevant factors, and the determination or direction of the outcome of the dispute;
- the use of generative AI to search for or obtain extrajudicial information or analyses that could affect the decision, without such information having been the subject of adversarial debate.

4.3 Bright-line test

If the use of AI has the effect or purpose of:

- determining the legal outcome;
 - assessing credibility or interpreting disputed facts;
 - assessing the relevance of the facts;
 - or articulating the reasons of a decision,
- it must not be used.

5. Uses that may be contemplated subject to strict conditions

5.1 General principle

Within the limits set out above, certain ancillary uses may be contemplated, subject to the strict conditions set out below.

The use of generative AI may be contemplated, with caution, solely as a secondary support tool for tasks ancillary to the performance of judicial functions, and provided that it does not undermine any fundamental values of justice or any ethical obligations.

5.2 Contemplated uses

Subject to these guidelines, generative AI could be contemplated, *inter alia*, for:

- the purely linguistic correction, revision, rewriting or translation of texts already written by the judge, on the understanding that this does not include draft judgments;

- certain ancillary administrative tasks, unrelated to judicial reasoning;
- the preparation of a factual timeline based on elements already identified by the judge;
- the production of a strictly descriptive summary of a witness statement, a document or a factual sequence, without any assessment, interpretation, normative selection or conclusion;
- the search for information in specific, closed sources (i.e. a specific document submitted to the tool, as opposed to an open web search), subject to systematic verification at source.

Content generated for strictly descriptive purposes may be used as working material or in the descriptive or contextual sections of a text alone, provided that it is fully verified, reproduced, amended where necessary and fully endorsed by the judge.

For the purposes of this section, descriptive selection — rendering what has been said or written, its order and its author, in a manner that can be verified at source — is distinct from normative selection, which consists of determining what is relevant, determinative or credible. Only the former may be assisted by the tool; only the judge may make the latter kind of selection.

5.3 Mandatory conditions

Any contemplated use is subject to the following conditions:

- full and continuous human oversight;
- independent verification of all generated content, which must include, as appropriate, reference to the judge's notes, the recording, the transcript, the source document or the original decision;
- the absence of confidential or identifiable data;
- an understanding of the tool's limitations;
- the ability to explain and take full responsibility for the final result.

Where in doubt, it is recommended to refrain from using the tool.

6. Confidentiality, security and data protection

The obligations of confidentiality and secrecy of deliberations apply in full to any use of generative AI.

Consequently, judges should not provide an AI tool with:

- draft judgments;
- notes relating to a pending case;
- any information protected by law or by a court order.

The judge must also ensure, in the current context where no institutional generative AI tool is provided, that the settings of the tool used do not permit the data entered to be used for the purposes of training or improving the system. In the event of uncertainty regarding these settings, it is recommended to refrain from using the tool.

In the current context, no transcription or recording processing tool has undergone an institutional assessment; the submission of such content to a generative AI tool is not recommended until such an assessment has been carried out.

Documents provided by the parties or by third persons must not be entered as-is into a generative AI tool without first carrying out an appropriate check of their content, origin and format. These documents may contain hidden information or instructions, which could unduly influence the tool's behaviour or compromise its reliability.

In any permitted use, the judge must ensure that they have control over the nature, scope and reliability of the information submitted to the tool.

7. Training, awareness-raising and development

7.1 Technological competence

The responsible use of generative AI requires the judge to develop adequate technological competence.

While not requiring specialized technical expertise, this competence means that the judge must be able to:

- understand the general capabilities and structural limitations of generative AI tools;
- recognize the risks of hallucinations, bias, factual errors and cognitive anchoring;
- distinguish clearly between technical, descriptive, organizational or linguistic assistance and judicial reasoning, which remains exclusively human;
- critically evaluate all generated content;
- recognize the risks associated with the use of documents or content from third persons, particularly where these may influence, distort or compromise the tool's output;
- understand that apparent efficiency gains never remove the need for full verification and may, in some cases, be negated by control requirements.

7.2 Ongoing training and awareness-raising

This competence must be supported by an institutional commitment, with the courts ensuring that the following are provided for this purpose:

- regular training and awareness-raising activities;
- the dissemination of up-to-date information on risks, limitations and best practices;
- institutional support tailored to changes in tools and contexts of use.

These activities are aimed at:

- strengthening judges' caution, autonomy and accountability;
- preventing inappropriate or uncontrolled use;
- supporting judges' professional judgment and critical assessment of these tools, without promoting their adoption.

7.3 The evolving nature of the guidelines

These guidelines constitute an evolving framework.

They are intended to be reviewed periodically to take account of:

- rapid technological developments;
- experience gained in the courts;
- the potential availability of secure institutional environments;
- best judicial practices in Canada and in comparable jurisdictions.

August 31, 2026

Original signed

Geneviève Cotnam, Chief Justice of Quebec and Chief Justice of the Court of Appeal of Quebec

Signed original

Marie-Anne Paquette, Chief Justice of the Superior Court of Québec

Signed original

Henri Richard, Chief Judge of the Court of Québec

Signed original

Nathalie Duchesne, Chief Municipal Judge of the Municipal Courts of Québec

APPENDIX A

Preliminary checklist for assessing a proposed use of generative AI

(To assist judges with self-assessment)

This framework is an optional guidance tool, separate from the guidelines and without any normative value. It helps the judge to structure their own thinking and may be used, adapted or disregarded as appropriate; it is neither a test nor a condition of compliance.

STEP 1 — Nature of the task (see section 4)

1. Does the task relate directly or indirectly to:

- legal reasoning?
- the assessment of evidence or credibility?
- the determination of the outcome of the dispute?
- the articulation of the analytical reasons of a decision?

☐ Yes / Unsure: use not recommended

☐ No: proceed to step 2

STEP 2 — The role of AI (see section 2)

2. Is AI used solely as a secondary support tool, rather than as a source of authority, reasoning or adjudication?

☐ Yes: proceed to step 3

☐ No / Unsure: use not recommended

STEP 3 — Nature of the requested content (see section 5.2)

3. Is the intended use limited to:

- purely linguistic correction, revision, rewriting or translation?
- a factual timeline based on elements already identified?
- a strictly descriptive summary, without any interpretation or conclusion?
- a search in specific, closed sources?

☐ Yes: proceed to step 4

☐ No / Unsure: use not recommended

STEP 4 — No interpretation (see section 4.3)

4. Are the judge's instructions to the AI tool intended to:

- interpreting the facts?
 - selecting relevant elements according to a normative logic?
 - drawing conclusions?
 - assessing credibility or consistency?
 - guiding judicial reasoning?
- ☐ No: proceed to step 5
- ☐ Yes / Unsure: use not recommended

STEP 5 — Data used (see section 6)

5.1 Does the data submitted to AI consist of:

- draft judgments?
- notes relating to a pending case?
- information protected by law or by a court order?

In practice, very little data relating to deliberations falls outside these categories; if in doubt, it must be treated as protected.

- ☐ No: proceed to step 5.2
- ☐ Yes / Unsure: use not recommended

5.2 Where the data comes from documents provided by parties or third persons, have I:

- checked their content, source and format?
 - exercised reasonable care in assessing their source, format and plausibility?
- ☐ Yes: proceed to step 6
- ☐ No / Unsure: use not recommended

STEP 6 — Verification and human oversight (see section 5.3)

6. Am I able to:

- verify the generated content in full?
- verify any quotations, references, authoritative sources or key factual information directly at the source?

- explain how the result has been used?
- take full responsibility for the final content?
- discard any suggestion without justification?

☐ Yes: proceed to step 7

☐ No: use not recommended

STEP 7 — Cognitive anchoring and independence (see section 5.3)

7. Have I put the following guidelines in place?

- I have articulated my own analysis before consulting the tool.
- I have not used the generated content as a starting point for the analysis or reasoning, except for the strictly descriptive material referred to in section 5.2.
- I have kept a record of my preliminary work.

☐ Yes: proceed to step 8

☐ No / Unsure: use not recommended — refrain from using the tool

STEP 8 — Appearance and public confidence (see section 3)

8. Would the intended use stand up to:

- an explanation to a colleague?
- an institutional review?
- hypothetical disclosure without undermining public trust?

☐ Yes: use can be contemplated, subject to strict conditions

☐ No / Uncertain: use not recommended

Final reminder to exercise caution: if in doubt, it is recommended that you refrain from using generative AI.