

Submission to Innovation, Science and Economic Development Canada

Consultation: Advancing AI Transparency in Canada

Submitted by: Tom Macintosh Zheng, co-founder, [Courtready.ca](https://courtready.ca)

Date: September 23, 2026

Via Email to: AIConsultations-ConsultationsIA@ised-isde.gc.ca

Summary

Unverified output from generative artificial intelligence (“AI”) has already had a detrimental effect on our justice system. As of September 21, 2026, at least 239 published decisions from 62 Canadian courts and tribunals had flagged AI-hallucinated content. Self-represented litigants were involved in 82% of them.

The discussion paper, “Enhancing trust in artificial intelligence through increased transparency”,¹ (the “**Discussion Paper**”) cites an earlier version of our research from March 2026, which flagged 132 decisions.² The count has since grown by 81%, and the problem has spread beyond fictitious case law to AI-generated evidence.

So far, institutional responses have placed the burden on users and adjudicators. In our view, transparency obligations should also reach the *developers* whose tools produce the output. This submission makes three recommendations:

1. Treat legal information as a high-stakes context requiring plain-language, point-of-use disclosure of AI limitations.
2. Recognize AI-generated material submitted as evidence in any legal proceeding as a distinct transparency problem that labelling alone cannot solve.
3. Include harms to the justice system in any framework for tracking AI incidents.

¹ Innovation, Science and Economic Development Canada, *Enhancing Trust in Artificial Intelligence Through Increased Transparency* (Ottawa: ISED, 2026), online: <https://ised-isde.canada.ca/site/ised/en/have-your-say-advancing-ai-transparency-canada/enhancing-trust-artificial-intelligence-through-increased-transparency>.

² *Ibid*, n 39, citing Tom Macintosh Zheng, “The Rise of AI-Hallucinated Case Law in Canadian Courts and Tribunals”, [2026 CanLIIDocs 738](#).

1. Our research on AI's impact on Canadian courts and tribunals

The dataset, which is hosted on Courtready.ca,³ covers published decisions from January 1, 2024, to September 21, 2026. Of the 239 decisions:

- 213 involved fictitious case citations, meaning cases cited as authority that do not exist;
- 18 involved invented legal rules, including real statutory provisions described as saying something they do not; and
- 8 involved fabricated quotations attributed to real decisions.

In 80% of decisions, the adjudicator found or presumed that AI generated the content.

The 239 figure is a conservative estimate. It captures only what an adjudicator detected and chose to address in written reasons. It also does not include private arbitrations or unpublished decisions.

2. Information about AI systems (Section 3.3)

The Discussion Paper notes that users may not realize AI responses come from patterns in training data rather than verified sources. The consequence of this lack of understanding, through no fault of the user, is that they inadvertently cite fictitious law as real.

General-purpose tools do not reliably check a legal database when asked for case law. They produce text that looks like a citation: a plausible party name, a real court abbreviation, a realistic year. The output is formatted and confident.⁴ As Professor Amy Salyzyn of the University of Ottawa has observed, “humans do sometimes produce bad

³ Tom Macintosh Zheng, *When AI Invents the Law: Canadian AI Hallucination Cases Database* (Toronto: Courtready.ca, last updated 21 September 2026), online: <https://courtready.ca/fictitious-citations-in-canadian-courts/>.

⁴ See e.g. *Giggs v. Michel (dba Soma MD Advanced Medical Aesthetics)*, 2026 BCCRT 1062 at para 8, online: <https://courtready.ca/wp-content/uploads/2026/07/Giggs-v.-Michel-dba-Soma-MD-Advanced-Medical-Aesthetics-2026-BCCRT-1062.pdf>, where the tribunal noted that the applicant cited “*Panchal v. Singh*, 2020 BCCRT 949, alleging it was about negligent eyebrow microblading. In reality, the correct citation is *Roofix Services Inc. v. Mike Stanfield (dba Kitchen Cabinets for Less)*, 2020 BCCRT 949, and relates to an agreement to install a kitchen exhaust roof vent.”

but well-presented work. LLMs always do. That makes it harder to spot the errors.”⁵ The only warning is typically a generic footer line that “AI can make mistakes.”

We respectfully submit that this warning is insufficient. Many self-represented litigants turn to AI because they cannot afford a lawyer,⁶ and they have no frame of reference to question what the tool produces. Court and tribunal rules⁷ that require parties to disclose AI use, or to certify that cited authorities are authentic, cannot solve this problem. A self-represented litigant who has no way of knowing that a case is fictitious will disclose and certify in good faith.

Recommendation: Developers of general-purpose AI systems should be responsible for clear disclosure at the point of use when a system produces legal authorities. This disclosure should state, in plain language, that the output has not been independently verified, that the cases or legal rules it cites may not exist, and that they should be checked against an authoritative source before anyone relies on them.

3. AI-generated content (Section 3.1)

The Discussion Paper frames AI-generated content around deepfakes, scams and misinformation, with watermarks, metadata and platform labels as solutions. The justice system faces a different problem: AI-generated *text* submitted as evidence.

Recent decisions include a party relying on material generated by two chatbots to support a claim,⁸ a party relying on a chatbot’s review of a veterinary clinic’s records in a claim

⁵ Amy Salyzyn, “Deceptive Dynamics of Generative AI: Beyond the ‘First-Year Associate’ Framing”, *Slaw* (28 October 2025), online: <https://www.slaw.ca/2025/10/28/deceptive-dynamics-of-generative-ai-beyond-the-first-year-associate-framing/>.

⁶ Jennifer Leitch, “The Inevitability of AI in Court: What Does It Mean for Self-Represented Litigants?”, *Slaw* (16 April 2025), online: <https://www.slaw.ca/2025/04/16/the-inevitability-of-ai-in-court-what-does-it-mean-for-self-represented-litigants/>.

⁷ Courtready.ca, “AI in Canadian Courts: Research, Tools and Guides”, online: <https://courtready.ca/ai-in-canadian-courts/>. As of 21 September 2026, 27 of the 249 Canadian courts and tribunals tracked require parties to disclose AI use.

⁸ *Lee v. AutoCanada MR Motors GP Inc.*, 2026 BCCRT 854 at para 8, online: <https://courtready.ca/wp-content/uploads/2026/06/Lee-v.-AutoCanada-MR-Motors-GP-Inc.-2026-BCCRT-854.pdf>.

against that clinic,⁹ and a worker filing a large body of AI-generated research to prove that a shoulder condition stemmed from a workplace injury.¹⁰

None of the tools in the Discussion Paper would reach these cases. The Discussion Paper acknowledges that text watermarks are brittle and easy to remove, and the labelling obligation in Bill C-34 would apply only to regulated social media services. More importantly, in several of these cases, the AI origin was never hidden. The litigant knew the material came from AI, and they believed, in good faith, that the material was reliable. A label tells a person that content is AI-generated, but it does not tell them whether it is true.

Verification also carries a cost to the already understaffed and underfunded justice system. Justice Browne of the Supreme Court of Newfoundland and Labrador has noted concerns that spouses in family proceedings may use deepfakes to attack credibility, and that authenticating such material requires expert evidence, which “only adds to issues like backlog and delay.”¹¹

British Columbia’s Civil Resolution Tribunal offers a working model. Since February 1, 2026, its rules have prohibited parties from providing false or fabricated evidence, including evidence created or altered by an AI tool.¹²

Recommendation: Any federal approach to AI-generated content should recognize legal proceedings as a distinct high-stakes context, and support adjudicative bodies in adopting a common standard: AI-generated material offered as evidence carries no weight unless the party can show where it came from and why it can be trusted.

⁹ *Ida v. Meadow Lane Equine Clinic Ltd.*, 2026 BCCRT 626 at paras 25–26, online:

<https://courtready.ca/wp-content/uploads/2026/04/Ida-v.-Meadow-Lane-Equine-Clinic-Ltd-2026-BCCRT-626.pdf>.

¹⁰ *Côté et Placements J.F. Desgagné inc.*, 2026 QCTAT 1311 at paras 21–25, online:

<https://courtready.ca/wp-content/uploads/2026/04/Cote-et-Placements-J.F.-Desgagne-inc.-2026-QCTAT-1311.pdf>.

¹¹ Anosha Khan, “Atlantic courts see rise in AI filings, cite accuracy and deepfake concerns”, *Law360 Canada* (21 September 2026), online: <https://www.law360.ca/ca/articles/2528111>.

¹² Civil Resolution Tribunal, *Standard Rules* (effective 1 February 2026), r 8.3(9), online:

<https://civilresolutionbc.ca/wp-content/uploads/CRT-Standard-Rules-effective-February-1-2026.pdf>.

4. AI interaction (Section 3.2)

We support the proposed rules in Bill C-34 addressing chatbots that deceptively pose as legal or other licensed professional. Legal matters should sit alongside health, finance and employment as a high-stakes context in any interaction disclosure framework.

5. AI incidents (Section 3.4)

Published decisions, including the ones in our database, are already an incident record. Each decision in our database identifies the harm, the setting and the consequence. To our knowledge, and as of the date of this submission, no public body collects this record systematically. The harm also extends beyond the parties. Justice Hamou of the Court of King's Bench of New Brunswick has observed that, with AI use, courts receive voluminous submissions and lengthy issues, leading to an “explosion of paperwork”.¹³

Recommendation: Any definition of a reportable AI incident should include harms to the administration of justice, and Canada should support independent, ongoing monitoring of this data.

Conclusion

The Canadian justice system has spent more than two years asking litigants, lawyers and adjudicators to catch what AI tools produce. The number of affected decisions has, nonetheless, continued to rise.

The underlying data is public, updated regularly and available at courtready.ca/fictitious-citations-in-canadian-courts. We would welcome the opportunity to provide updated figures or further information to support the Government's work.

About the Author

Tom Macintosh Zheng is the co-founder of Courtready.ca and is licensed to practise law in Ontario and in England and Wales. Courtready.ca builds practical tools for Canadians who litigate and conducts original research on access to justice issues in Canada. His research on the rise of AI-generated case law in Canadian courts and tribunals, which is available on Courtready.ca, has been cited by Innovation, Science and Economic Development Canada (ISED), the Court of Appeal of Quebec, Ontario's Law Society Tribunal, The Economist, CBC News, CTV News, Global News and the Toronto Star. Tom is this year's recipient of the Ontario Bar Association's OBA Foundation Award, which recognizes exceptional contributions to the improvement of the justice system. He can be reached at tom@courtready.ca.

¹³ Khan, *supra* note 11.