

Four Guiding Principles on the Use of Artificial Intelligence in Legal Scholarship

Yan Campagnolo, Adamo Crescenzi, Djouly Laube, Emma Pellegrini, Catherine Roy, and Fumi Shibutani

University of Ottawa, Faculty of Law

PUBLIÉ EN LIGNE | PUBLISHED ONLINE

Octobre 2026 | October 2026

VOLUME

57

NO

2

PAGES

329–354

CITATION RECOMMANDÉE • RECOMMENDED CITATION

ENGLISH

Yan Campagnolo et al, “Four Guiding Principles on the Use of Artificial Intelligence in Legal Scholarship” (2026) 57:2 Ottawa L Rev 329.

rdo-olr.org · lawrev@uottawa.ca · 57 Louis-Pasteur, Ottawa ON K1N 6N5

ÉTABLIE EN • EST. 1966

UNIVERSITÉ D'OTTAWA
FACULTÉ DE DROIT
SECTION DE COMMON LAW

UNIVERSITY OF OTTAWA
FACULTY OF LAW
COMMON LAW SECTION

CONTENTS

Four Guiding Principles on the Use of Artificial Intelligence in Legal Scholarship

Yan Campagnolo, Adamo Crescenzi, Djouly Laube, Emma Pellegrini, Catherine Roy, and Fumi Shibutani

- I. Introduction 331
- II. AI Is Not an Author 332
 - A. Accountability and Copyright 333
 - 1. Human Authors Must Remain Fully Accountable for Their Work 333
 - 2. Copyright: Legal and Public Policy Perspectives 335
 - B. AI in Scholarship 336
- III. AI Use Must Be Disclosed 338
 - A. Transparency 339
 - 1. Academic Integrity and Transparency Demand Disclosure 339
 - 2. Disclosure Must Occur at the Submission Stage 340
 - 3. Publication Eligibility: Limiting Acceptance to Categories 1 to 3 341
 - 4. Certification of Source Verification: A Necessary Safeguard 342
 - B. Sanctions for Breach 342
- IV. Rigorous Human Oversight Is Required to Ensure the Integrity of Publications 343
 - A. The Central Role of Human Oversight 344
 - 1. Accuracy, Hallucinations, and Verification 344
 - 2. Bias and Editorial Judgment 345
 - B. The OLR's Editorial Process and Reputation 345
 - 1. Independent and Multi-Layered Review 345
 - 2. Editorial Use of AI Tools 346
- V. Policies Governing the Use of AI Must Adapt to Change 347
 - A. Evolving Technological and Legal Landscapes 347
 - 1. Technological Developments 347
 - 2. Legal and Regulatory Uncertainty 347

B. Institutionalizing Adaptability	348
1. Creation of a Permanent AI Committee	348
2. Adaptability as a Complement to Oversight	349
VI. Conclusion	349
Appendix 1: General Conditions on the Use of Artificial Intelligence	350
Appendix 2: Disclosure and Certification Form	353

Four Guiding Principles on the Use of Artificial Intelligence in Legal Scholarship

*Yan Campagnolo, Adamo Crescenzi, Djoully Laube,
Emma Pellegrini, Catherine Roy, and Fumi Shibutani**

I. INTRODUCTION

In 2025, the *Ottawa Law Review* (“OLR”) revealed its insights on the use of artificial intelligence (“AI”) in legal scholarship in its book *Artificial Intelligence’s Impact on Legal Journals*. At the time of writing, legal journals had not yet caught up to the evolving and extensive use of AI in legal publications.¹

At the conclusion of the book, the OLR established its first policy on the use of generative AI. The policy stated that AI is allowed in the submission process but that “all authors must take full responsibility for the factual accuracy of their submission and the proper citation of sources.”² The policy also required authors to disclose whether they used generative AI to produce their work, along with the model and version used, the range of dates of that use, and the steps they took to verify the veracity of the AI-generated output.³ The policy prohibited the use of AI in the peer review process but allowed its use in the editing process for spelling and grammar purposes only.⁴ Lastly, the policy acknowledged that AI is fast evolving and

* Yan Campagnolo is a full professor at the University of Ottawa Faculty of Law and faculty advisor to the *Ottawa Law Review* (OLR). Adamo Crescenzi, Catherine Roy, Djoully Laube, Emma Pellegrini, and Fumi Shibutani are JD candidates at the same institution and senior OLR editors. For constructive comments on earlier versions, we are indebted to Wolfgang Alschner, Nachman Gutowski, Teresa Scassa, and Amy Salczyn. We also wish to acknowledge the excellent editing work performed by Desmond Fisher.

1 Yan Campagnolo, ed, *Artificial Intelligence’s Impact on Legal Journals* (Ottawa: University of Ottawa Press, 2025) at 5.

2 *Ibid* at 43.

3 *Ibid*.

4 *Ibid* at 44.

that the OLR's position was subject to change along with the evolution of copyright law and industry standards.

In response to quickly evolving AI technology, the OLR established an AI committee in September 2025. The committee was led by Professor Yan Campagnolo and included senior editors Adamo Crescenzi, Djoully Laube, Emma Pellegrini, Catherine Roy, and Fumi Shibutani. Through the committee's work, the OLR reaffirmed and enhanced its first policy and developed a new set of guiding principles to help the journal address the risks and opportunities associated with AI.⁵ The committee's definition of "AI" was limited to generative AI, that is, "AI that generates content—including text, image, audio, and video—based on the system's training data and in response to prompts that the user inputs."⁶

Four guiding principles were adopted. First, in order to ensure human accountability and remain consistent with Canadian copyright principles, the OLR will not recognize AI as an author. Second, the OLR will require AI use to be disclosed as a matter of academic integrity and transparency. Disclosure must occur at the submission stage, and the OLR will add sanctions for failing to engage in proper disclosure. Third, the OLR will reinforce human oversight during the assessment and editorial processes to ensure the integrity of publications. The presence of bias and hallucinations and the importance of editorial judgment call for independent and multi-level review by human editors in order to protect the OLR's reputation. Finally, the OLR will continue to monitor ongoing developments and adapt its AI policy to evolving technological and legal landscapes. The OLR will remain vigilant to ensure it takes a proactive rather than reactive approach.

II. AI IS NOT AN AUTHOR

The OLR does not and will not recognize AI as an author. Human authors must remain fully accountable for their work, from the accuracy of the basic facts in an article all the way to its eventual impact on society, rather than offloading responsibility to an AI tool.

5 The AI committee leveraged AI (ChatGPT and Gemini) to supplement its existing research and to locate additional academic sources. During the drafting and editing process, its AI use was restricted to correcting grammar and spelling and enhancing clarity. All outputs were edited, proofread, and reformulated by the committee. No part of the article or accompanying policy was written by AI. All sources were manually verified to ensure none were hallucinations.

6 Campagnolo, *supra* note 1 at 4.

From an intellectual property standpoint, the OLR's stance matches that of Canadian courts, which require a work to involve a human's exercise of skill and judgment to attract copyright protection.⁷ Human-centred authorship also aligns with the general public's views and the standards set by major scholarly publishers across multiple disciplines.

Short of authorship, AI continues to be a useful resource for ideation, brainstorming, and structural clarity. Rather than being recognized as an author, AI should be recognized as a tool for human authors to use to support the writing process.

A. Accountability and Copyright

1. *Human Authors Must Remain Fully Accountable for Their Work*

Kevin Frazier and Alan Rozenshtein have taken the position that AI should not be considered an author because authors have a responsibility to ensure the exactitude, accuracy, and clarity of submitted work.⁸ Specifically, they argue that “[l]isting AI as a co-author is inappropriate because it wrongly suggests shared responsibility, opening the door for human authors to deflect accountability”.⁹

Human accountability is particularly relevant in medical publications, where an author's claims and conclusions may have significant consequences. Indeed, scientific misconduct in medicine carries serious consequences for public health. Falsified data or fraudulent findings can mislead practitioners, policymakers, and the public. For example, Andrew Wakefield's 1998 article in *The Lancet* alleging a link between the MMR vaccine and autism continues to affect public confidence in vaccines despite its being fully retracted and exposed as fraudulent in 2010.¹⁰

In the legal context, human accountability and oversight helps mitigate some of the risks and possible harm caused by AI hallucinations. Research indicates that both general-purpose and legal-specific AI models hallucinate facts at alarming rates. Most domain-specific AI tools (that is, legal AI tools like Lexis+ AI) leverage retrieval-augmented generation (“RAG”).

7 *CCH Canadian Ltd v Law Society of Upper Canada*, 2004 SCC 13 at para 16.

8 In their article, the authors focus on the use of AI research tools such as Westlaw and LexisNexis rather than generative writing tools like ChatGPT (see Kevin T Frazier & Alan Z Rozenshtein, “Large Language Scholarship”, *Fla Intl U L Rev* [forthcoming in 2026]).

9 *Ibid* at 46.

10 TS Sathyanarayana Rao & Chittaranjan Andrade, “The MMR Vaccine and Autism: Sensation, Refutation, Retraction, and Fraud” (2011) 53:2 *Indian J Psychiatry* 95.

Although legal-specific RAG AI systems appeared to hallucinate less than general-purpose tools in earlier testing, they remained far from reliable. Faiz Surani and Daniel Ho's 2024 study, based on testing earlier models, discovered that Lexis+ AI hallucinated at a rate of 17 percent, while Westlaw AI hallucinated 34 percent of the time.¹¹

Legal AI tools hallucinate in one of two ways. First, an AI response may just be incorrect; it describes the law incorrectly or makes a factual error.¹² Second, an AI response may be misguided; the AI tool describes the law correctly, but cites a source that does not support the claims it makes.¹³ If an author relies on an AI tool that then cites sources that appear authoritative on the surface but are irrelevant or contradictory in reality, users could place undue trust in the tool's output, potentially leading to erroneous legal judgments and conclusions.¹⁴

According to a 2024 study by Matthew Dahl on four large language models, or LLMs, these tools also struggle to predict their own hallucinations and often uncritically accept users' incorrect legal assumptions.¹⁵ Importantly, hallucinations in LLMs increase when the complexity of the legal research task increases.¹⁶ As long as hallucinations persist, human oversight, verification, and responsibility is the only viable solution.

In addition to hallucination issues, AI models can replicate bias in their outputs when they are trained on biased data.¹⁷ Many AI models create their own algorithms based on training data to identify and recreate data patterns.¹⁸ Learning from past biased decisions and recognizing these outcomes as the norm results in "bias amplification".¹⁹ For example, when online retailer Amazon introduced an AI hiring tool to review job applications and filter unfit applicants, it trained its tool on biased data that did

11 Faiz Surani & Daniel E Ho, "AI on Trial: Legal Models Hallucinate in 1 out of 6 (or More) Benchmarking Queries" (23 May 2024), online: <hai.stanford.edu/news/ai-trial-legal-models-hallucinate-1-out-6-or-more-benchmarking-queries>.

12 *Ibid.*

13 *Ibid.*

14 *Ibid.*

15 Matthew Dahl et al, "Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models" (2024) 16:1 J Leg Analysis 64 at 64.

16 *Ibid* at 74.

17 Andrej Krištofik, "Bias in AI (Supported) Decision Making: Old Problems, New Technologies" (2025) 16:1 Intl J for Court Administration 1 at 8; Rory Pulvino, Dan Sutton & JJ Naddeo, "Hiding in Plain Sight: An Empirical Study of Prosecutorial Bias in AI Legal Analysis" (2026) 27:1 Colum Sci & Tech L Rev 1.

18 Krištofik, *supra* note 17 at 8–9.

19 *Ibid* at 8.

not include women in senior leadership roles.²⁰ This created a norm that being a woman disqualified a candidate from executive positions, resulting in female applicants' CVs being discarded during the hiring process.²¹

More recently, a 2026 empirical study in criminal law suggests that LLMs may adopt a pro-prosecution orientation not simply because of isolated bias, but because they are trained on legal corpora dominated by prosecutorial narratives and guilt-presuming language. Rory Pulvino, Dan Sutton, and Joseph Naddeo found that ChatGPT consistently favoured prosecution even when the facts revealed evidentiary weaknesses or possible constitutional violations, indicating that these systems can internalize and reproduce a broader "prosecutorial default bias." Their findings show how the data used for training may shape legal AI outputs in subtle but systematic ways, reinforcing existing punitive tendencies rather than neutrally assessing cases.²² When AI outputs cause tangible harm or discriminatory results, human authors who relied on such outputs should not evade responsibility by attributing harm to technology; rather, they should be held accountable for their published work. By rejecting AI authorship, the OLR can ensure that the negative impacts of a published article remain the responsibility of human authors, rather than being offloaded to AI.

2. Copyright: Legal and Public Policy Perspectives

While "author" is not defined in the *Copyright Act*, the Federal Court of Appeal has confirmed that an author "undoubtedly refers to a natural person".²³ Consequently, since AI is not a natural person, it cannot attract authorship and copyright protection on its own. The Federal Court of Appeal's decision aligns with the OLR's view that AI, on its own, cannot be an author. Recognizing AI authorship would be inconsistent with the legal framework, creating a disconnect between OLR practices and established legal doctrine.

However, Canadian courts have not decided whether AI can be a *co-author*. In December 2021, the Canadian Intellectual Property Office recognized RAGHAV (an AI painting tool) and a human as co-authors of an artwork titled SURYAST.²⁴ In response, the Samuelson-Glushko Canadian

²⁰ *Ibid* at 8–9.

²¹ *Ibid*.

²² Pulvino et al, *supra* note 17 at 10, 21.

²³ *PS Knight Co Ltd v Canadian Standards Association*, 2018 FCA 222 at para 147.

²⁴ "SURYAST" (painting) Ankit Sahni & RAGHAV Artificial Intelligence Painting App, Can 1188619 (1 December 2021) registered.

Internet Policy and Public Interest Clinic (“CIPPIC”) submitted a notice of application to the Federal Court seeking an order for the SURYAST registration to be struck from Canada’s copyright registry.²⁵ Some of CIPPIC’s reasons included authorship’s being restricted to humans and the artwork’s lacking originality.²⁶ Although a hearing date has yet to be set, the Federal Court will review whether RAGHAV can be recognized as a co-author when it hears *CIPPIC v Sahni*. The OLR will carefully consider the Federal Court’s decision when it is issued, but in the meantime, it will not recognize AI as a co-author for the same reasons it does not recognize AI as an author.

The Government of Canada’s consultation on the implications of AI for copyright enabled the public to participate and share their opinions on AI. The consultation revealed that many participants agree that human authorship is a requirement to attract copyright protection.²⁷ Some participants felt that human authors should be able to use AI as a tool in their creative process, while others proposed that AI-generated elements of a work should not attract copyright.²⁸ This general reluctance to grant copyright to AI-generated works aligns with the OLR’s decision that AI authorship will not be recognized.

B. AI in Scholarship

In the pure sciences, most journals have policies prohibiting AI authorship, reaffirming that accountability requires human oversight. The evidentiary thresholds and precision involved in STEM research mean that minor errors introduced by AI have significant ramifications on reproducing results and interpreting data. Although they do not recognize AI authorship, journals acknowledge that AI may be used in other ways, such as during the drafting or analysis process. Pure science journals also generally require disclosure of AI use.

25 *Samuelson-Glushko Canadian Internet Policy and Public Interest Clinic v Ankit Sahni* (8 July 2024), Ottawa T-1717-24 (FC) (Notice of Application), online (pdf): <cippic.ca/_files/ugd/391d3a_36e25294c78a490ebbe2870fo77c78db.pdf>.

26 *Ibid.*

27 Innovation, Science and Economic Development Canada, *What We Heard Report: Consultation on Copyright in the Age of Generative Artificial Intelligence* (Ottawa: Innovation, Science and Economic Development Canada, 2025) at 9.

28 *Ibid.*

In the social sciences, major journals, particularly those under publishing houses like Wiley or Elsevier, have implemented uniform AI guidelines. These guidelines generally prohibit AI authorship and co-authorship. Smaller independent journals tend to align with their publisher's standards. However, some of these journals authorize the use of AI to assist with content generation and analysis with adequate disclosure of such use. For example, some journals explicitly allow the use of AI for linguistic editing or data visualization, if that use does not alter the conceptual analysis or theoretical interpretation of the piece. Overall, social science journals are proactively developing AI policies.

In contrast, legal journals have been slower to adopt AI guidelines. In Canada, only 4 of 13 major law journals (less than 31 percent) have adopted such policies at the time of writing.²⁹ One explanation for this difference could be legal scholarship's focus on doctrine, compared with the empirical focus of research in other disciplines. Moreover, unlike non-legal journals, law journals typically establish their editorial standards independently, without oversight from a unifying professional body. Where policies do exist, they tend to mirror those of larger publishing bodies or include general ethical guidelines prohibiting AI authorship.

The clearest and most universally adopted principle across academic disciplines is the categorical exclusion of AI tools as authors. Every journal with a formal AI policy explicitly states that AI tools cannot meet the traditional criteria for authorship.³⁰ Under these standards, authorship requires intellectual responsibility for a study's conception, design, analysis, and interpretation of data, as well as accountability for errors. Since AI lacks consciousness, ethical agency, and the ability to consent or respond to peer review, it cannot bear this responsibility.³¹ Scholarly journals therefore reaffirm that ultimate accountability must remain with human authors.

Although AI should not be recognized as an author, AI should still be acknowledged as a valuable tool. Nachman Gutowski has compared the role of AI to that of collaborative support, like having a colleague review

29 The authors reviewed 13 Canadian law journal policies in the fall of 2025. At the time, the four journals that had adopted AI policies were *Osgoode Hall Law Journal*, *McGill Law Journal*, *University of Toronto Law Journal*, and *Dalhousie Law Journal*.

30 The authors reviewed the policies of 13 Canadian law journals, 24 pure science journals, and 10 social sciences journals in the fall of 2025.

31 David B Resnik & Mohammad Hosseini, "Disclosing Artificial Intelligence Use in Scientific Research and Publication: When Should Disclosure be Mandatory, Optional, or Unnecessary?" (2025) 33:2 *Accountability in Research* 1 at 2.

your work.³² Similarly, Spencer Nayar and Michael Cooper point out that AI's contributions do not reduce the value of an article. Indeed, AI can assist human authors in creating original contributions.³³ Short of authorship, AI can contribute to scholarship in various ways, such as by aiding with research tasks by summarizing legal precedents or suggesting additional relevant sources. AI may also be used to generate preliminary text or outlines that human authors refine and expand upon.³⁴ AI also serves as a valuable editing tool; beyond correcting grammar and misspellings, AI can be used to improve sentence structure and enhance clarity.

III. AI USE MUST BE DISCLOSED

The *OLR* requires full disclosure of AI use in legal scholarship. Human authors need to maintain complete responsibility for the accuracy of their work, the accurate citation of sources, and the reliability of their analysis rather than delegate these duties to an AI tool.

From a scholarly standpoint, the *OLR*'s approach aligns with contemporary journal policies that treat transparency as essential for evaluating the origin of claims and the soundness of methods. Disclosure enables editors, peer reviewers, and readers to assess the origins of an article's arguments, ensuring that the scholarly record remains trustworthy and verifiable. Disclosing AI use is analogous to authors acknowledging their research assistants and librarians, who contributed to an article from its inception to publication. Therefore, disclosure of AI is an extension of the commonly accepted and ethical practice of recognizing the contributions of others.

Despite being incapable of full authorship or primary analytical responsibility, AI remains a useful resource for drafting, linguistic editing, and preliminary research. Rather than replacing human responsibility, AI should be recognized as a tool that supports human authorship.

32 Nachman N Gutowski, "Disclosing the Machine: Trends, Policies, and Considerations of Artificial Intelligence Use in Law Review Authorship" *Jacksonville U L Rev* 1 at 4, 8 [forthcoming in 2026].

33 Spencer Nayar & Michael I Cooper, "The 'Why' & 'How' of Artificial Intelligence in Legal Scholarship" (2025) 11:4 *Texas A&M J Property L* 543 at 553–57.

34 *Ibid* at 561–62.

A. Transparency

1. *Academic Integrity and Transparency Demand Disclosure*

The use of AI in drafting and analysis must be disclosed. Disclosure is not simply an ethical preference; it is a condition for maintaining confidence in legal scholarship as a reliable record. Contemporary journal policies across disciplines increasingly require authors to disclose whether AI was used in drafting, analysis, image generation, or linguistic editing because transparency enables editors, reviewers, and readers to assess who made the claims and the integrity of the methods used to produce them.

In a law journal environment, transparency is especially important because doctrinal argument depends not only on persuasive reasoning but on demonstrable grounding in authority, which requires accurate citation and faithful characterization of sources. Where AI systems are involved, the core risk is the possibility that apparently well-formed legal propositions are paired with inaccurate, incomplete, or unsupported authority, thereby undermining the trust that scholarly legal writing assumes and requires.

Accountability and transparency are the key rationales underlying the OLR's existing posture toward AI. The OLR's current policy permits AI use at the submission stage but squarely assigns responsibility to authors by stating that "all authors must take full responsibility for the factual accuracy of their submission and the proper citation of sources."³⁵ The current policy also requires disclosure of any use of generative AI, the model and version used, the range of dates it was used, and the steps taken to verify the veracity of AI-generated output.³⁶ These requirements reflect a core publishing reality: because AI tools cannot be held accountable, the journal's integrity must be secured through human responsibility combined with a level of transparency that permits targeted verification.³⁷

The need for a robust transparency regime is sharpened by the persistence of hallucinations. Hallucinations occur in both general-purpose systems and legal-specific systems. For legal scholarship, a particularly damaging variant is *misgrounding*, where a statement of law may be plausible, or even correct in the abstract, yet the cited authority does not in fact support the proposition advanced. This failure mode is corrosive to

³⁵ Campagnolo, *supra* note 1 at 43.

³⁶ *Ibid.*

³⁷ Zafer Kocak, "Publication Ethics in the Era of Artificial Intelligence" (2024) 39:33 J Korean Medical Science 1 at 3; Resnik & Hosseini, *supra* note 31 at 2.

legal academic practice because it can produce an article that appears conventionally sourced while being substantively unreliable at the level that matters most in law: the relationship between proposition and authority. The practical implication is that verification cannot be treated as optional; human oversight and verification are the only viable responses where hallucination risk remains structurally present.

2. *Disclosure Must Occur at the Submission Stage*

Disclosure must occur at the submission stage through a standardized, mandatory form. The reason is simple: unstructured or purely narrative disclosures tend to be uneven, non-comparable, and difficult to operationalize. A structured disclosure form allows the journal to translate transparency into governance by enabling calibrated screening and verification. Such a form responds to a specific institutional concern: without a structured disclosure process, author declarations about AI use can become a pure trust exercise, in which journals give authors the benefit of the doubt while remaining exposed to the risk that citations do not support the propositions offered or, in some cases, are fabricated. Because a journal's credibility depends on citation integrity, the disclosure regime must be designed to reduce the extent to which the journal is forced to rely on trust alone, particularly when the underlying technology is known to generate plausible but unsupported claims or references.

A four-level disclosure regime provides a workable structure for this form because it maps the intensity of AI involvement to corresponding editorial expectations.³⁸ At Category 1, the author declares no AI use across drafting, analysis, research assistance, or editing. At Category 2, the author declares limited AI use. Examples include assistance with language correction or a literature review.³⁹ At this level of AI use, the author develops and drafts all substantive analysis and argumentation.⁴⁰ Category 2 also includes using AI to suggest edits or refine the author's draft for clarity and concision. At Category 3, the author declares moderate use, meaning AI was used to generate some results, text, or ideas such as preliminary outlines, draft passages later substantially rewritten, or suggested conceptual framings that the author adopted and refined.⁴¹ At Category 4, the author

38 Nayar & Cooper, *supra* note 33 at 561–63, 567–68.

39 Andrew M Perlman, *Generative AI and the Future of Legal Scholarship* (Research paper, Suffolk University Law School, 2025).

40 Nayar & Cooper, *supra* note 33 at 561.

41 *Ibid.*

declares high use, meaning AI was used to generate the article or a substantial part of it, with the author's role primarily consisting of selection, editing, rearrangement, or limited supplementation.⁴²

3. *Publication Eligibility: Limiting Acceptance to Categories 1 to 3*

A credible transparency framework should draw a principled boundary around what the journal will publish. The OLR will publish only submissions that involved AI use at Categories 1 to 3 and will not publish Category 4 submissions.⁴³ This is not a condemnation of AI; it is a risk-management posture that aligns the OLR's acceptance criteria with what it can responsibly stand behind as a verified scholarly product. Furthermore, rejecting Category 4 submissions affirms the OLR's guiding principle that AI is not an author and that publishing human-centric scholarship is the OLR's key priority. The persistence of hallucinations makes it difficult to treat AI generation of substantial portions of an article as merely another drafting tool without significantly expanding verification burdens and reputational exposure.⁴⁴

Within Categories 1 to 3, Category 3 requires special handling because it is the point at which AI contributions move beyond mechanical support into substantive generation. Accordingly, authors of Category 3 submissions will be required to provide specific information describing the nature of AI use, including the tool and version, the timeframe of use, the prompts used to generate AI outputs, and the steps taken to verify them. Most of these requirements are consistent with the OLR's existing disclosure expectations.

In addition, authors of Category 3 submissions will be required to cite AI-generated results, text, visual representations, and ideas where those outputs materially shaped the manuscript's analysis or argument.⁴⁵ Such citation is not intended to treat AI as an author; rather, it is a mechanism that enables readers to understand when a particular synthesis, framing, or passage originated as AI-generated content that was then adopted, adapted, or rewritten. An OLR-specific citation format must therefore be developed to ensure that AI citations are standardized, informative, and integrated into OLR citation practices. One solution could be to adopt the Canadian

⁴² Perlman, *supra* note 39; Nayar & Cooper, *supra* note 33 at 561.

⁴³ See Appendix 2. See also Nayar & Cooper, *supra* note 33 at 561.

⁴⁴ Nayar & Cooper, *supra* note 33 at 555–56, 565.

⁴⁵ *Ibid* at 561.

Open Access Legal citation rules for AI citations as part of the *OLR Style Guide* to add consistency in citation.⁴⁶

4. *Certification of Source Verification: A Necessary Safeguard*

Disclosure alone cannot protect the integrity of legal scholarship; verification must also be explicitly required. A separate certification, also incorporated into the new submission form, will require authors to attest that all sources cited were reviewed and verified and that the manuscript does not contain hallucinated facts or sources. The point of certification is not to create a formalistic checkbox; it is to anchor the author's non-delegable responsibility for factual accuracy and proper citation by imposing upon them a concrete procedural duty that is enforceable. This certification requirement is consistent with the underlying reality that hallucinations can involve fabricated citations or misgrounded references and that verification is the only viable response to mitigate these risks.

Certification must be explicit to ensure the integrity of publications. Even where authors disclose AI use, the OLR must still contend with the possibility that their citations are unreliable, and the appropriate editorial response is rigorous human oversight of supporting materials to ensure that each footnote adequately supports the text. Author certification and editorial oversight are therefore complementary; certification places primary responsibility on the author, while editorial oversight by the OLR provides an institutional backstop that protects the integrity of publications where the risk profile of current tools makes trust alone insufficient.

B. Sanctions for Breach

A transparency regime must be strictly enforced to be credible. The disclosure and certification obligations will therefore be paired with clearly stated sanctions for breach. These sanctions will be explicit in the submission materials and will include the article's rejection and a temporary publication ban. The justification for sanctions lies in the integrity function of the disclosure and certification regime itself: where declarations can otherwise become a trust exercise, the prospect of meaningful consequences is necessary to deter incomplete disclosure by authors and to discourage

46 COAL-RJAL Editorial Group, "Canadian Open Access Legal Citation Guide" (2025), online: <canlii.org/en/commentary/doc/2024CanLIIDocs830>; Ottawa Law Review, "OLR Style Guide" (2026), online (pdf): <rdo-olr.org/wp-content/uploads/2026/02/OLR-Style-Guide_Winter-2026.pdf>.

careless verification practices that risk introducing unsupported or fabricated authorities into the scholarly record. Sanctions are also institution-protective: they preserve the credibility of the journal by ensuring that authors understand disclosure and verification as substantive obligations rather than simple ethical preferences that can be ignored.

Transparency must also operate within the editorial and peer review workflow. To facilitate transparency, the *OLR* has developed new “General Conditions on the Use of Artificial Intelligence” (see Appendix 1) and a corresponding “Disclosure and Certification Form” (see Appendix 2). The completed form for each submission will be shared with *OLR* editors at the pre-screening stage and with expert reviewers at the peer review stage in an anonymized format. The “Disclosure and Certification Form” will equip decision-makers with process-relevant information that can support calibrated screening and verification. This approach also addresses confidentiality and integrity concerns that the *OLR* has already recognized in its governance posture, including the prohibition on using AI in peer review. The “Disclosure and Certification Form” therefore allows the journal to take AI use seriously as a workflow-relevant variable without undermining the normative commitments of peer review, including fairness, confidentiality, and reasoned evaluation of scholarly merit.

To sum up, a transparency framework for AI use in legal scholarship must be built around a non-negotiable disclosure principle and operationalized through a structured submission-stage form, a tiered classification of AI involvement, and an explicit certification that all sources have been reviewed and verified. The integrity rationale is straightforward: because hallucinated facts and misgrounded citations remain a persistent risk, journals cannot rely on trust alone and must require both disclosure and verification to protect the reliability of their articles. Limiting publication eligibility to submissions involving AI use at Categories 1 to 3, imposing heightened obligations at Category 3 (including a standardized AI citation format), and integrating sanctions transform transparency from a general ethical aspiration into a workable set of editorial controls to ensure that the articles published are accurate.

IV. RIGOROUS HUMAN OVERSIGHT IS REQUIRED TO ENSURE THE INTEGRITY OF PUBLICATIONS

The *OLR* believes that rigorous and consistent human oversight is a foundational requirement for maintaining the integrity, credibility, and scholarly

value of its publications. The principle of oversight applies regardless of whether AI is used for drafting, editing, or translating, and regardless of the degree of AI assistance involved.

AI systems, whether general-purpose or legal-specific, remain prone to hallucinations in their citations and to biased outputs.⁴⁷ These limitations make clear that AI tools cannot independently satisfy the standards of accuracy, accountability, and judgment expected in legal scholarship. Human oversight is therefore not merely precautionary; it is structurally necessary to ensure that each published article meets the *OLR*'s standards of rigour and reliability.

This principle is consistent with the dominant position taken across pure sciences, social sciences, and legal publishing, where AI governance frameworks emphasize that AI may assist scholars but cannot replace the need for human oversight and final judgment. It is also reinforced by reputational considerations: the credibility of the *OLR* depends not on whether AI is used, but on the demonstrable robustness of the review and editorial processes prior to publication.

A. The Central Role of Human Oversight

1. *Accuracy, Hallucinations, and Verification*

As mentioned in section II, authors such as Surani, Ho, and Dahl all consistently found that AI systems continue to hallucinate facts and sources at alarming rates.⁴⁸ This is true even for legal-specific tools that rely on RAG, such as Lexis+ AI and Westlaw AI.⁴⁹ While these systems hallucinate less frequently than general-purpose models, they still produce incorrect citations with sufficient regularity to pose serious risks to the integrity of legal scholarship.⁵⁰

Most alarming for legal scholarship are instances when an AI system correctly states a legal proposition while citing a source that does not support it, creating a false appearance of authority.⁵¹ These outputs are particularly dangerous in legal writing, where readers often rely heavily on citations as indicators of reliability.

⁴⁷ See Surani & Ho, *supra* note 11; Krištofik, *supra* note 17 at 7–8.

⁴⁸ Surani & Ho, *supra* note 11; Dahl et al, *supra* note 15 at 64–66.

⁴⁹ Dahl et al, *supra* note 15 at 64–74.

⁵⁰ Surani & Ho, *supra* note 11.

⁵¹ *Ibid.*

At present, no technological solutions can reliably eliminate these risks.⁵² AI-detection tools remain inaccurate and inconsistent.⁵³ As a result, the only effective safeguard is rigorous human verification of propositions, citations, and sources.⁵⁴ This requirement applies irrespective of whether an author has disclosed AI use, as hallucinations may arise even where AI involvement is minimal or indirect.

2. *Bias and Editorial Judgment*

In addition to hallucinations, Andrej Krištofik and Rahul Bharati have both emphasized the persistence of bias in AI systems, arising from biased training data and algorithmic design.⁵⁵

From an editorial perspective, this risk of bias raises serious concerns about the use of AI in any evaluative or pre-screening capacity. The persistence of bias reinforces the necessity of human-led review at all stages of the publication process, particularly during submission assessment and peer review. Human editors are uniquely positioned to contextualize arguments, recognize scholarly merit beyond formalistic patterns, and mitigate the risk of systemic bias that AI tools may inadvertently amplify.

B. The OLR's Editorial Process and Reputation

1. *Independent and Multi-Layered Review*

The OLR's existing editorial framework is a core institutional strength. The combination of article peer review by subject-matter experts, a multi-step editorial review, and manual citation verification aligns with best practices observed in disciplines that have more mature AI governance regimes.

⁵² *Ibid.*

⁵³ Ahmed M Elkhayat, Khaled Elsaid & Saeed Almeer, "Evaluating the Efficacy of AI Content Detection Tools in Differentiating Between Human and AI-generated Text" (2023) 19:1 Intl J Educational Integrity at 12–13.

⁵⁴ Surani & Ho, *supra* note 11. See the example used by the authors to illustrate the conclusion that access to AI tools still requires rigorous human oversight. The article mentions that Paul, Weiss, Rifkind, Wharton & Garrison LLP did over a year of testing on an AI product and was unable to develop reliable metrics because checking the AI system was so involved that it made any efficiency gains difficult to measure. See also Zach Warren, "GenAI Hallucinations are Still Pervasive in Legal Filings, but Better Lawyering is the Cure" (18 August 2025), online: <thomsonreuters.com/en-us/posts/technology/genai-hallucinations>.

⁵⁵ Krištofik, *supra* note 17 at 8; Dr Rahul Kailas Bharati, "Bias and Fairness in AI Algorithms: Legal Standards and Ethical Guidelines" (2025) 12:2 Intl J Novel Research Computer Science & Software Engineering 44.

What protects the *OLR*'s reputation is not the absence of AI but the presence of sound human judgment at every decisive stage. Transparent dissemination of the *OLR*'s process therefore serves a reputational function, reassuring readers, authors, and reviewers that published work has undergone meaningful scrutiny.

2. *Editorial Use of AI Tools*

Consistent with cross-disciplinary standards, AI tools can have a limited editorial use for grammar correction and translation, where such tools are widely accepted and where their outputs are always subject to human review. These uses are framed as assistive rather than determinative and do not implicate substantive legal analysis or editorial decision-making.

At the same time, the use of AI-detection software should be cautioned. Research surveyed across disciplines demonstrates that these tools are currently unreliable, are prone to false positives, and raise serious privacy and confidentiality concerns when applied to unpublished manuscripts.⁵⁶ Accordingly, in order to ensure procedural fairness and respect for submission confidentiality, the *OLR* does not use AI-detection tools.

Journals that have adopted AI policies consistently emphasize the importance of disclosing their own AI use as a means of preserving the trust of their stakeholders and remaining accountable.⁵⁷ In line with this trend, the *OLR* will disclose that it may use AI-powered tools for limited editorial purposes while clearly affirming that final editorial judgment always remains human. Peer reviewers will also be prohibited from using AI to evaluate submissions to ensure an independent human assessment process. These approaches align the *OLR* with leading pure science and social science journals.

56 Mike Perkins et al, "GenAI Detection Tools, Adversarial Techniques and Implication for Inclusivity in Higher Education" (March 2024) [unpublished, archived at Cornell University].

57 Dalhousie Law Journal, "Information for Contributors", online: <digitalcommons.schulichlaw.dal.ca/dlj/policies.html>; Lex Electronica, "Politique contre le plagiat (2025)", online: <lex-electronica.org/soumissions/politique-contre-le-plagiat>; University of Toronto Press, "Editorial Policies", online: <utppublishing.com/about/editorial-policies-for-ethics-and-publication-malpractice#_ai>; Osgoode Hall Law Journal, "Publication Ethics Statement", online: <digitalcommons.osgoode.yorku.ca/ohlj/publication_ethics.html>.

V. POLICIES GOVERNING THE USE OF AI MUST ADAPT TO CHANGE

AI governance in academic publishing cannot be static. AI tools evolve rapidly, legal frameworks remain unsettled, and scholarly norms continue to develop unevenly across disciplines. In this context, policies governing AI use must be understood as living instruments that can be updated as new information, technologies, and risks emerge. The OLR's initial AI policy was intentionally provisional, reflecting an awareness that technological capabilities and legal standards are still in development. This approach is correct; adaptability is a necessary condition for responsible governance.

A. Evolving Technological and Legal Landscapes

1. *Technological Developments*

Andrew Perlman has noted that AI tools, particularly in the fields of legal research and drafting, continue to advance.⁵⁸ At the same time, many authors have pointed out that these tools remain ill-suited for internal journal operations, especially given their training on predominantly American and English-language datasets. Since the OLR is a bilingual journal, AI tools' limited training on French materials is a significant concern.

Outside of research and drafting, AI tools could be useful for tasks such as automating citations, expediting translation, or editing syntax. However, progress on this front has been modest. Therefore, continuous monitoring is required to assess whether emerging tools can meaningfully alter the risk-benefit calculus for editorial adoption.

2. *Legal and Regulatory Uncertainty*

From a legal standpoint, developments in the fields of privacy, copyright, and AI regulation are ongoing.⁵⁹ Canadian authorities continue to investigate AI practices, while international initiatives such as the EU's

⁵⁸ Perlman, *supra* note 39 at Epilogue.

⁵⁹ Office of the Privacy Commissioner of Canada, "Joint Statement on Data Scraping and the Protection of Privacy" (24 August 2023), online: <priv.gc.ca/en/opc-news/speeches-and-statements/2023/js-dc_20230824>; Office of the Privacy Commissioner of Canada, "Statement on the Role of Data Protection Authorities in Fostering Trustworthy AI" (11 October 2024), online: <priv.gc.ca/en/opc-news/speeches-and-statements/2024/s-d_g7_20241011_ai>; Office of the Privacy Commissioner of Canada, "Statement on AI and Children" (11 October 2024), online: <priv.gc.ca/en/opc-news/speeches-and-statements/2024/s-d_g7_20241011_child-ai>.

Artificial Intelligence Act signal emerging expectations for transparency, bias mitigation, and accountability.⁶⁰ In 2026, the Information and Privacy Commissioner of Ontario and the Ontario Human Rights Commission jointly published a report titled *Principles for the Responsible Use of Artificial Intelligence*.⁶¹ The titular principles, along with Ontario's directive regarding the responsible use of AI, serve as important guidance for public institutions like the University of Ottawa. The four guiding principles on AI use presented herein are in line with the principles put forth in the report, especially the reliability, privacy, transparency, and accountability requirements.

Although these regimes do not directly regulate law reviews, they shape the broader normative environment in which academic publishing operates. The OLR must remain attentive to legal developments, including pending litigation on AI co-authorship and evolving copyright doctrine.

B. Institutionalizing Adaptability

1. *Creation of a Permanent AI Committee*

To operationalize adaptability, the OLR will establish a permanent AI committee. This committee will be responsible for monitoring technological developments, tracking legal and regulatory changes, reviewing cross-disciplinary best practices, and recommending timely policy updates. A permanent committee will ensure new developments in AI are caught and addressed in a timely manner.

The creation of a permanent AI committee mirrors governance structures observed in larger publishing institutions and reflects the OLR's conclusion that ongoing oversight and revision are essential to responsible AI integration. A permanent committee will ensure that AI governance does not depend on ad hoc responses or institutional memory, but on sustained and informed engagement.

60 EU, *Regulation 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations 300/2008/EC 167/2013/EU, 168/2013/EU, 2018/858/EU, 2018/1139/EU, 2019/2144/EU and Directives 2014/90/EU, 2016/797/EU and 2020/1828/EU*, [2024] OJ, L 2024/1689.

61 Privacy Commissioner of Ontario & Ontario Human Rights Commission, "Principles for the Responsible Use of Artificial Intelligence" (1 January 2026), online (pdf): <ohrc.on.ca/sites/default/files/2026-OHRC-IPC%20joint%20principles%20of%20responsible%20use%20of%20AI.pdf>.

2. *Adaptability as a Complement to Oversight*

Finally, adaptability and oversight are mutually reinforcing principles. Rigorous human oversight protects against current risks, while adaptive policy structures ensure that those safeguards remain effective as conditions change. Together, these principles position the OLR not merely as a cautious adopter of AI, but as an institution capable of leading thoughtful, principled discourse on AI governance in Canadian legal scholarship.

VI. CONCLUSION

By reviewing key developments and research on the use of AI in legal scholarship, the OLR has developed four guiding principles to steer its internal policy development and governance structures. First, the OLR does not and will not recognize AI as an author or co-author but acknowledges its utility as a research tool. Second, the OLR will require AI use to be disclosed as a matter of academic integrity and will sanction any breach of this disclosure requirement. Third, the OLR will reinforce human oversight to ensure the integrity of its publications and protect its reputation. Finally, the OLR will remain current with evolving AI technology and a changing legal landscape. These four principles will inform future changes to the OLR's policies and help ensure that the OLR remains proactive rather than reactive to AI going forward.

APPENDIX 1

General Conditions on the Use of Artificial Intelligence

This Policy outlines the general conditions that apply to the use of artificial intelligence (“AI”) in preparing submissions to the *Ottawa Law Review* (“OLR”).

1. Policy**1.1 Definitions**

In this policy,

- (1) “author” means any author or co-author of a submission,
- (2) “submission” means any text submitted to the OLR for potential publication, and
- (3) “AI” means any type of artificial intelligence that generates content based on the system’s training data and in response to prompts that the user inputs.

1.2 Compliance

Every author must read and comply with this policy and assume full responsibility for their submission’s compliance with it. Violation of this policy by an author may result in

- (1) immediate termination of the publication process for their submission, and
- (2) refusal to publish future submissions they author for a period of time.

2. AI Use in Submissions**2.1 AI as Author**

The OLR does not recognize AI as an author and takes the position that only human beings may be authors.

2.2 Acceptable Use of AI in Submissions

The OLR accepts the following categories of submissions:

- (1) No AI was used in creating the submission; all research, drafting, and revisions were carried out solely by the author (“Category 1”).
- (2) AI served as a research aid, by assisting with tasks such as summarizing legal decisions or identifying relevant sources. AI served as a linguistic editing aid, by suggesting edits or rephrasing for clarity. The submission itself was written by the author (“Category 2”).
- (3) AI contributed to the submission as a research aid and by generating preliminary text or an outline. The portions of text generated by AI were then substantially rewritten, refined, and expanded upon by the author (“Category 3”).

2.3 Unacceptable Use of AI in Submissions

The OLR does not accept submissions for which AI was used to generate the article or a substantial part of it, with the author’s role primarily consisting of selection, editing, rearrangement, or limited supplementation (“Category 4”).

3. Mandatory Disclosure of AI Use

3.1 Disclosure and Certification Form

All authors must include a duly completed copy of the “Disclosure and Certification Form” to disclose whether AI was used in their submission and, if so, to what extent it was used.

3.2 Responsibility

Any author making a submission to the OLR bears responsibility for any misrepresentation made in their “Disclosure and Certification Form” and is subject to the sanctions set out in section 1.2.

4. OLR Processes and AI

4.1 General

The OLR puts all submissions through a rigorous review process—including the preselection committee’s assessment, the peer review process, and the editing process—to ensure their accuracy, credibility, and reliability.

4.2 Peer Review

- (1) As part of peer review, an anonymized copy of the “Disclosure and Certification Form” will be shared with the peer reviewers.
- (2) To uphold the authenticity, confidentiality, and integrity of submissions, peer reviewers cannot use AI to assess submissions during the peer review process.

4.3 Editing

- (1) Subject to section 4.3(2), *OLR* editors are not permitted to upload or otherwise submit any portion of a submission to an AI system.
- (2) *OLR* editors may use AI to assist with linguistic editing, spelling, or grammar.
- (3) *OLR* editors will check whether the sources cited in each submission exist, are relevant, and are accurately cited.

4.4 Responsibility

The author retains final responsibility for ensuring the accuracy, credibility, and reliability of cited sources.

APPENDIX 2**Disclosure and Certification Form****1. Submission**

Name of Author

Email of Author

Title of Submission**2. Disclosure**

Please check the applicable category.

- ☐ **Category 1:** No artificial intelligence (“AI”) was used in creating the submission; all research, drafting, and revisions were carried out solely by the author.
- ☐ **Category 2:** AI served as a research aid, by assisting with tasks such as summarizing legal decisions or identifying relevant sources. AI served as a linguistic editing aid, by suggesting edits or rephrasing for clarity. The submission itself was written by the author.
- ☐ **Category 3:** AI contributed to the submission as a research aid and by generating preliminary text or an outline. The portions of text generated by AI were then substantially rewritten, refined, and expanded upon by the author.

Please note that the *Ottawa Law Review* does not accept Category 4 submissions, that is, submissions for which AI was used to generate the article or a substantial part of it, with the author’s role primarily consisting of selection, editing, rearrangement, or limited supplementation.

3. Additional Disclosure

If Category 3 was checked in the previous section,

- provide the name and version of the AI model used as well as the dates (or range of dates) when it was used, and
- describe how the AI model was used and the steps taken to confirm the veracity of the generated output.

4. Certification

I certify that

- I have read and understood the “General Conditions on the Use of Artificial Intelligence”,
- if applicable, I have fully disclosed all use of AI in my submission, and
- if applicable, I have reviewed and verified all content and sources provided by AI that are included in my submission.

Signature

Date