

Bond University

Bond Law Review

Volume 38

Issue 2

2026

Disciplined Integration: AI and the Legal Profession's Existing Duties

Pouya Lajevardi

Bond University

Follow this and additional works at: <https://blr.scholasticahq.com/>



© Copyright the authors.

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivative 4.0 International Licence](https://creativecommons.org/licenses/by-nc-nd/4.0/).

DISCIPLINED INTEGRATION: AI AND THE LEGAL PROFESSION'S EXISTING DUTIES

AUTHOR: POUYA LAJEVARDI*

Abstract

This article argues that the legal profession's existing duties of competence, supervision, candour, confidentiality, and procedural fairness provide the necessary and sufficient framework for integrating generative AI into legal practice, without requiring a wholly new regulatory paradigm. Narrowing its focus to transformer-based large language models, the article examines the 'fake law' phenomenon, where models generate plausible but non-existent authorities, and traces its implications for miscarriages of justice, particularly where adversarial safeguards are thinnest: self-represented litigants, resource-constrained parties, and off-record settlements. It then considers AI in adjudication, arguing that transparency, contestability, and human-authored reasons remain non-negotiable for judicial legitimacy. The central recommendation is disciplined integration: augmentation, not automation. AI should operate within the profession's existing load-bearing beams through auditable workflows, clear disclosure norms, and continuous education. Used with discipline, AI can strengthen practice; used without it, it risks eroding the trust and outcomes the profession exists to protect.

I Introduction

Lawyers owe specific professional duties and obligations, including to the administration of justice and to the courts. Lawyers may be characterized as “agents of the rule of law”.¹ In an era of rapid technological change, the role of lawyers, judges, and court officers in safeguarding the integrity of the rule of law has never been more important.

This article argues that the legal profession's existing duties – competence, supervision, candour to the tribunal, confidentiality, and procedural fairness – provide the necessary and sufficient framework

* JD (Bond University); BSc (University of Toronto). Research Associate, Meshkini Legal Services PC. The author would like to thank Wendy Bonython for their guidance.

¹ Richard Moorhead, Steven Vaughan and Kenta Tsuda, *What Does It Mean for Lawyers to Uphold the Rule of Law?* (Report, Legal Services Board, October 2023) 2–3, 16 (‘Rule of Law Report’).

for the disciplined integration of generative artificial intelligence into legal practice. The challenge posed by AI is not that it demands an entirely new regulatory paradigm, but that it tests the rigour with which the profession applies the duties it already has. To call these duties sufficient is not to call them static. The argument is that, read in a manner appropriate to the conditions of generative AI, they already reach its characteristic risks, so that what the moment demands is not a new body of rules but the disciplined, and continually renewed, application of the duties the profession already owes, of which a working literacy in AI is now an inseparable part. The profession's response should be proactive and by design: not to reject AI, but to make it work within the load-bearing beams of professional responsibility, through auditable workflows, clear disclosure norms, and continuous education.

“Disciplined integration”, as this article uses the term, denotes the incorporation of generative AI into legal work through, and subordinate to, the profession's existing duties, rather than through a new, AI-specific rulebook. The discipline lies in the rigour with which those duties are applied: in treating AI outputs as unverified leads until they are corroborated against primary sources, in maintaining auditable checkpoints and records of verification, in disclosing material AI assistance where candour requires it, and in sustaining the education that keeps competence current. So understood, disciplined integration is a narrower and more exacting idea than the familiar language of “responsible” or “ethical” AI, which tends to state high-level principles addressed to those who design or deploy these systems; its claim is that the profession requires no new normative paradigm, because the duties it already owes are equal to the task. It is also more than “augmentation, not automation”: augmentation names the mode of use this article defends, in which AI assists rather than displaces the human decision-maker, while disciplined integration supplies the architecture of responsibility, verification, disclosure and continuing education that keeps augmentation honest across practice, adjudication and access to justice alike.

These duties are not the invention of any single legal system. Competence, supervision, candour to the tribunal and confidentiality are the common currency of the professional-conduct rules that govern lawyers across the common-law world, expressed in the model codes and solicitors' conduct rules examined below, and procedural fairness draws on the lawyer's duty to the court and the requirements of natural justice those systems share. Their precise formulation varies from one jurisdiction to another, and this article does not attempt a comparative survey. Rather, it takes its examples predominantly from common-law systems – Canada, Australia, the United Kingdom and the United States – and pitches its argument at the level of the obligations those systems

hold in common. That shared foundation is what allows the analysis to travel: the particular rules differ, but the underlying duties, and the pressures generative AI places upon them, do not.

The paper's method is correspondingly doctrinal and normative rather than empirical. It establishes the content of these duties from professional-conduct rules, case law and regulatory guidance, and argues for how they should be understood and applied to the conditions generative AI creates. The claims that follow are therefore about the proper interpretation of existing obligations, not about the measured behaviour of particular tools or practitioners.

Because "artificial intelligence" is an overloaded term, Part II narrows the scope of this article to transformer-based large language models accessed through consumer interfaces or APIs. Part III examines the 'fake law' phenomenon – fabricated cases and citations generated by these models – and the emerging regulatory responses from courts and professional bodies. Part IV analyses the consequent risk of miscarriages of justice, particularly where adversarial safeguards are weakest. Part V turns to AI in adjudication, considering judicial impartiality, the debate over AI judges versus AI-assisted judging, and the Canadian administrative law context. Part VI makes the case for augmentation rather than automation, drawing together the risks identified in the preceding Parts. Part VII examines how these risks bear most heavily on self-represented litigants and the broader access-to-justice implications. Part VIII argues for a "back to basics" approach, anchoring AI integration in technological competence, ethical judgment, continuous education, and the primacy of professional duty. Part IX concludes.

II What is AI?

"Artificial intelligence" references a long, shifting research programme rather than a single stable technology.² From the symbolic reasoning systems of the 1950s and 1960s, through the expert systems and knowledge-representation work of the 1980s, to the statistical pattern-learning that followed the revival of neural networks and

² Alan M Turing, "Computing Machinery and Intelligence" (1950) 59(236) *Mind* 433 ('Computing Machinery and Intelligence').

backpropagation, the term's referent has repeatedly shifted.³ For legal and policy purposes, 'AI' is best understood as a moving target.⁴

Today's dominant paradigm centres on what Bommasani et al term 'foundation models': large-scale, generative systems pre-trained on vast datasets whose general-purpose capabilities underpin many downstream applications, including in legal and governmental settings.⁵ The form of AI with which this article is principally concerned is the contemporary family of transformer-based large language models ('LLMs'), such as the GPT series.⁶ These models generate fluent, contextually coherent text by predicting the next token in a sequence, having been trained on internet-scale text corpora.⁷ Crucially, a single pre-trained model can perform diverse tasks from natural-language prompts alone, without task-specific retraining, a capability that has made LLMs rapidly accessible to lawyers and litigants alike through consumer interfaces such as ChatGPT.⁸

The capabilities of these systems are considerable. Empirical evaluations of GPT-4 document human-level or near-human-level performance on a range of professional benchmarks, including a simulated bar examination.⁹ Yet the same evaluations reveal persistent and poorly understood failure modes, most notably 'hallucination' — the generation of fluent, confident text that is factually incorrect or entirely fabricated.¹⁰ It is this combination of persuasive output and unreliable accuracy that makes LLMs both a powerful tool and a professional hazard, and that animates the analysis in the Parts that follow.

³ John McCarthy et al, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence" (Web Document, 31 August 1955) ('Dartmouth AI Proposal'); Marvin Minsky, "Steps Toward Artificial Intelligence" (1961) 49(1) *Proceedings of the IRE* 8 ('Steps Toward AI'); Marvin Minsky and Seymour Papert, *Perceptrons* (MIT Press, 1969) chs 1, 11 ('Perceptrons'); David E Rumelhart, Geoffrey E Hinton and Ronald J Williams, "Learning Representations by Back-Propagating Errors" (1986) 323 *Nature* 533, 533–5 ("Learning Representations"); see also Marvin Minsky, *The Society of Mind* (Simon & Schuster, 1986) ('Society of Mind').

⁴ See Yann LeCun, Yoshua Bengio and Geoffrey Hinton, 'Deep Learning' (2015) 521 *Nature* 436; See also David Silver et al, 'Mastering the Game of Go with Deep Neural Networks and Tree Search' (2016) 529 *Nature* 484; Jia Deng et al, 'ImageNet: A Large-Scale Hierarchical Image Database' (Conference Paper, IEEE Conference on Computer Vision and Pattern Recognition, 2009).

⁵ Rishi Bommasani et al, 'On the Opportunities and Risks of Foundation Models' (Preprint, arXiv:2108.07258, 16 August 2021).

⁶ See generally Tom B Brown et al, 'Language Models are Few-Shot Learners' (Preprint, arXiv:2005.14165, 28 May 2020).

⁷ Ashish Vaswani et al, 'Attention Is All You Need' (Preprint, arXiv:1706.03762, 12 June 2017).

⁸ Brown et al (n 6).

⁹ OpenAI, 'GPT-4 Technical Report' (Preprint, arXiv:2303.08774, 15 March 2023); Sébastien Bubeck et al, 'Sparks of Artificial General Intelligence: Early Experiments with GPT-4' (Preprint, arXiv:2303.12712, 22 March 2023).

¹⁰ Bubeck et al (n 9); OpenAI (n 9); See also Varun Magesh et al, 'Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools' (2025) 22 *Journal of Empirical Legal Studies* 216.

Unless otherwise indicated, references to 'AI' in this article denote general-purpose language models accessed through consumer interfaces or APIs, and their purpose-built legal counterparts, sometimes with retrieval-augmented generation or domain-specific fine-tuning layered on top.¹¹

III Fake Law

Since the advent of LLMs, courts have repeatedly confronted “fake law”, fabricated cases and quotations placed into filings from generative tools. The emblematic U.S. episode is *Mata v. Avianca*,¹² where counsel submitted non-existent decisions and stood by them even after judicial orders, prompting sanctions and a detailed admonition that, while AI tools are not inherently improper, existing rules require lawyers to exercise judgment and independently verify outputs.¹³ In Colorado disciplinary proceedings, *People v. Zachariah C. Crabill*, a lawyer cited ChatGPT-sourced cases he had not read or verified, then failed to correct the record and misattributed the errors, conduct recounted in the disciplinary findings.¹⁴ In *Park v. Kim*, sanctions were issued after counsel relied on a fake case located via ChatGPT; the Second Circuit emphasized that Rule 11 of *Federal Rules of Civil Procedure* (‘*FRCP*’) obliges attorneys to read and confirm the existence and validity of the authorities on which they rely.¹⁵ Massachusetts’ *Smith v. Farwell* similarly imposed a monetary sanction after interns and an associate used an AI system to prepare filings with fabricated citations; although the court found no intent to mislead, it sanctioned counsel to deter recurrence.¹⁶

Outside the U.S., Canadian courts have also addressed “fake law”: in *Zhang v. Chen*, two AI-fabricated citations led to personal cost consequences (no special costs absent intent), with the court characterizing the practice as an abuse of process that, if unchecked, risks miscarriage of justice.¹⁷ Other jurisdictions have flagged similar issues. Some such issues were identified in Australia’s legal spheres in the PWC report, which highlighted the friction between AI “hallucinations” and existing legal regimes. Practical examples of these

¹¹ This includes systems such as Lexis+ AI and CoCounsel that layer retrieval-augmented generation or fine-tuning on top of general-purpose LLMs: see, eg, Thomson Reuters, ‘CoCounsel: The Industry-Leading AI Assistant for Professionals’ (Product Page, 2025); LexisNexis, ‘Lexis+ AI Security Information’ (Information Sheet, June 2025).

¹² *Mata v Avianca Inc*, 678 F Supp 3d 443 (SD NY, 2023) (‘*Mata*’).

¹³ William A Ryan, Allen Garrett and Brad Sears, “Practical Lessons from the Attorney AI Missteps in *Mata v. Avianca*” (Web Page, 8 August 2023) <<https://www.acc.com/resource-library/practical-lessons-attorney-ai-missteps-mata-v-avianca>> (‘Practical Lessons from AI Missteps’).

¹⁴ *People v Crabill* (Colo OPDJ, 2023 WL 8111898, 22 November 2023) (‘*Crabill*’).

¹⁵ *Park v Kim*, 91 F 4th 610 (2nd Cir, 2024) [9]–[10] (‘*Park*’).

¹⁶ *Smith v Farwell* (Mass Super Ct, No 2284CV01197, 12 February 2024) (‘*Smith*’).

¹⁷ *Zhang v Chen* [2024] BCSC 285, [29] (‘*Zhang*’).

risks have already manifested: a government review conducted by Deloitte Access Economics was found to contain hallucinated case citations, and a submission to a parliamentary inquiry by Australian academics was retracted after it was revealed to include AI-generated case studies containing false allegations against consulting firms.¹⁸

In the reasons for judgment in *Luck v Secretary, Services Australia*, the court redacted a false case citation to avoid further propagation by the AI system. The Court said:

The case referred to in the first paragraph of [a quoted extract] does not exist. The judgment with the medium neutral citation referred to is a completely different matter ... We apprehend that the reference may be a product of hallucination by a large language model. We have therefore redacted the case name and citation so that the false information is not propagated further by artificial intelligence systems having access to these reasons.¹⁹

Collectively, these matters trace a clear through-line: courts will sanction or impose costs when AI-generated authorities go unverified, and they are increasingly issuing guidance that AI may assist drafting but is never itself an authority requiring anything less than full, independent verification by counsel.

This article argues that the issue is, and always has been, fundamentally one of professional responsibility, now heightened by the emergence of AI tools such as ChatGPT. Most glaring mistakes and hallucinations tended to occur at the early stages of the technology itself. As with every technology, there is a curve from birth to maturity. At the low end of that curve, the risks are high and problems many, with many unaware of the rules of the brave new world.²⁰

None of this is to dismiss the concerns surrounding the fake law discussion. While these are primarily professionalism concerns, technological safeguards remain relevant. It is also important to address the concern head-on with thorough and rigorous education within the profession that has been happening for a while now.²¹

¹⁸ PwC Australia, “Real” Laws for Artificial Intelligence (Report, February 2024) 5; see also Deloitte, Targeted Compliance Framework Assurance Review: Final Report (Report, 4 July 2024) 47–8; James Guthrie et al, Submission No 5 to Senate Finance and Public Administration References Committee, Parliament of Australia, Inquiry into the Management and Assurance of Integrity by Consulting Services (2023).

¹⁹ *Luck v Secretary, Services Australia* [2025] FCAFC 26, [14] (“*Luck*”).

²⁰ See generally David Collingridge, *The Social Control of Technology* (St Martin’s Press, 1980), describing the “dilemma of control” wherein the social consequences of a technology cannot be predicted early in its life, rendering regulation difficult until the technology is entrenched. See also Everett M Rogers, *Diffusion of Innovations* (Free Press, 3rd ed, 1983).

²¹ Canadian Bar Association, “Ethics of Artificial Intelligence for the Legal Practitioner” (Practice Toolkit, November 2024) <<https://www.cba.org/resources/practice-tools/ethics-of-artificial-intelligence-for-the-legal-practitioner>> (“CBA AI Ethics Toolkit”); Canadian Bar Association, “AI Technologies in Practice” (Webinar Series, June 2024) <<https://www.cba.org/our-impact/initiatives/legal-futures/ai-technologies-in-practice>> (“CBA AI Webinar”); Law Society of British Columbia, “Guidance on

Another arm of the fake law concern is the risk that client confidential data may enter LLM training datasets.²² Although this risk is receding in medium to large firms with safeguards put in place for such issues by the developers themselves, this can remain a concern for smaller firms and solo practitioners who may not be giving attention to or have the resources to discern which tools and versions preserve client confidentiality and are compliant.²³

Accordingly, the way forward is not to exceptionalize these episodes but to treat generative outputs as unverified leads until corroborated by primary source;²⁴ institutionalize simple, auditable checkpoints, including attaching the sources relied on, maintaining citation-validation logs, and adding a brief read-the-case certification to the file; adopt matter-level protocols that confine AI use to defined tasks such as issue spotting, drafting alternatives, summarization, with human verification at points of reliance; require vendors to disclose training limits and known failure modes;²⁵ and embed ongoing education so that

Professional Responsibility and Generative AI" (Practice Resource, 2023) <<https://www.lawsociety.bc.ca>> ('LSBC AI Guidance'); see also "AI and the Practice of Law" (Lawcast BC, Law Society of British Columbia, 5 June 2024) ('Lawcast BC'); Law Society of Alberta, "Generative AI Playbook" (Guidance, January 2024) ('LSA Playbook'), summarised in Canadian Lawyers Insurance Association, "2025 Update: Generative AI Guidance" (Web Page, 2 September 2025) <<https://www.clia.ca>> ('CLIA 2025 Update'); Law Society of New South Wales, "A Solicitor's Guide to Responsible Use of Artificial Intelligence" (Guideline, November 2023, updated October 2024) <<https://www.lawsociety.com.au>> ('LSNSW AI Guide'); New South Wales Bar Association, "Issues Arising from the Use of AI Language Models (including ChatGPT) in Legal Practice" (Guideline, 12 July 2023) <<https://inbrief.nswbar.asn.au>> ('NSW Bar AI Guidelines'); Law Institute of Victoria, "Ethical and Responsible Use of Artificial Intelligence" (Ethics Guideline, 2025) <<https://www.liv.asn.au>> ('LIV AI Guidelines').

²² Law Society of Alberta, LSA Playbook (n 21) 'Confidentiality and Security'; Professional Support Unit, "A Solicitor's Guide to Responsible Use of Artificial Intelligence" (14 November 2023) Law Society Journal 3–4 ('LSJ AI Article').

²³ OpenAI, "Business Data Privacy, Security and Compliance" (Web Page, 2025) <<https://openai.com/business-data>> ('OpenAI Business Data'); OpenAI, "Enterprise Privacy" (Web Page, 4 June 2025) <<https://openai.com/enterprise-privacy>> ('OpenAI Enterprise Privacy'); OpenAI, "Introducing ChatGPT Enterprise" (Blog Post, 28 August 2023) <<https://openai.com/index/introducing-chatgpt-enterprise>> ('ChatGPT Enterprise Blog'); LexisNexis, "Lexis+ AI Security Information" (Information Sheet, June 2025) 1–3 <<https://www.lexisnexis.ca/pdf/lexis-plus-ai-security-info.pdf>> ('Lexis+ AI Security'); Jeff Jenkins, "7 Key Facts about Legal AI Security and Privacy with LexisNexis®" (Thought Leadership, LexisNexis Legal & Professional, 28 July 2025) <<https://www.lexisnexis.com/community/insights/legal/b/thought-leadership/posts/ai-privacy-security-the-lexisnexis-commitment>> ('LexisNexis Privacy Facts'); Thomson Reuters, "CoCounsel: The Industry-Leading AI Assistant for Professionals" (Product Page, 2025) <<https://www.thomsonreuters.com/en/cocounsel>> ('CoCounsel Product Page'); ABA Standing Committee on Ethics and Professional Responsibility, Generative Artificial Intelligence Tools (Formal Opinion 512, 29 July 2024) 5–7 ('ABA Opinion 512'); Canadian Bar Association, 'CBA AI Ethics Toolkit' (n 21) pt 4 'Choosing AI Tools'; Law Society of New South Wales, 'LSNSW AI Guide' (n 21) 4–5.

²⁴ See Model Rules of Professional Conduct r 1.1 cmt 8 (duty of technological competence, including understanding the "benefits and risks associated with relevant technology") ('Model Rules'); ABA Opinion 512 (n 23) 5–7 (lawyers must independently review and verify AI-assisted work product).

²⁵ See 'ABA Opinion 512' (n 23) 10–11.

interns through partners can recognize telltale artifacts of hallucination and understand when model confidence is not justified.²⁶ In line with transparency to the court, it would be ideal to include brief, standardized disclosures of when and how an AI system assisted. However, conversely, some judges have taken the other view that in complex litigation, this may create a significant administrative hurdle.²⁷

Courts and regulators are now creating guardrails to target the precise failure modes exposed by hallucinated citations. In the United States, standing orders such as Judge Brantley Starr's require a certificate that any AI-assisted text has been independently verified, expressly to prevent fabricated quotations and authorities from entering the record.²⁸ The decision in *Park v. Kim* drives the point home under Rule 11 FRCP: a "fake opinion is not "existing law,"" and counsel must read and confirm the authorities they cite.²⁹

In Canada, the Federal Court's Notice requires a declaration stating when and how AI was used in preparing filed materials, while the Court's AI policy commits that it will not use AI to make judgments or orders absent public consultation.³⁰ The Court of Appeal of Quebec's Notice likewise urges caution, reliability and human involvement when AI is used for research or drafting.³¹ Canadian judicial administration is moving in step: the Canadian Judicial Council's 2024 *Guidelines for the Use of AI in Canadian Courts* caution that judges cannot delegate decision-making to AI, and the Ontario Superior Court of Justice has added a dedicated "Use of AI" section to its Consolidated Civil Provincial Practice Direction. This trend carries in other jurisdictions such as Queensland, Australia as well.³²

²⁶ Ibid 6, 10–12.

²⁷ Ibid 2; Victorian Law Reform Commission, Artificial Intelligence in Victoria's Courts and Tribunals (Report, November 2025) [7.116]–[7.117] ('VLRC AI Report').

²⁸ "A Sampling of AI Court Orders" (Web Document) <<https://www.textenrls.org/wp-content/uploads/2024/07/A-sampling-of-AI-Court-Orders.pdf>> ('AI Court Orders').

²⁹ *Park* (n 15) [9]; Federal Rules of Civil Procedure r 11.

³⁰ Federal Court of Canada, Updated Notice to the Parties and the Profession: The Use of Artificial Intelligence (AI) Tools in Court Proceedings, 22 October 2024 <<https://www.fct-cf.ca/Content/assets/pdf/base/FC-Updated-AI-Notice-EN.pdf>> ('Federal Court AI Notice'); Federal Court of Canada, "Artificial Intelligence" (Web Page, 2024) <<https://www.fct-cf.ca/en/pages/law-and-practice/artificial-intelligence>> ('Federal Court AI Page').

³¹ Court of Appeal of Quebec, The Use of Artificial Intelligence (AI) in Proceedings Before the Court (Notice, 2023) <https://courdappellduquebec.ca/fileadmin/dossiers_civils/avis_et_formulaires/eng/avis_utilisation_intelligence_artificielle_ENG.pdf> ('QCA AI Notice').

³² Canadian Judicial Council, Use of Artificial Intelligence (AI) by Judges – Guidelines (September 2024) <<https://cjc-ccm.ca/sites/default/files/documents/2024/AI%20Guidelines%20-%20FINAL%20-%202024-09%20-%20EN.pdf>> ('CJC AI Guidelines'); Ontario Superior Court of Justice, Consolidated Provincial Practice Direction for Civil Proceedings at the Superior Court of Justice, 2024 <<https://www.ontariocourts.ca/scj/filing-procedures/provincial/consolidated-civil-provincial-practice-direction/>> ('Ontario Civil Practice Direction'); See also Queensland Courts, The Use of Generative AI: Guidelines for Judicial Officers, 15 September

Professional regulators are converging on the same operational baseline, treating LLM outputs as unverified leads until a lawyer has exercised competence, supervision and candour. In July 2024, the American Bar Association issued Formal Opinion 512 to clarify the ethical implications of using generative AI in legal practice. The opinion establishes that existing obligations apply fully to the use of AI tools. Furthermore, the guidance explicitly warns practitioners about the technological risks associated with these tools, including the potential for ‘hallucinations’ and the inadvertent exposure of confidential client data.³³ In Canada, the Federation of Law Societies’ Model Code embeds technological competence and confidentiality obligations that apply equally to digital tools, and provincial law society resources direct licensees to understand the benefits and risks of the technologies they deploy.³⁴ Read together with the court notices above, this body of authority supports simple, auditable controls so that AI assistance never dilutes the lawyer’s non-delegable duty to verify the record.³⁵

The result is not romanticizing the new tools nor rejecting them, but encouraging disciplined integration in which competence, supervision, and candour to the tribunal remain the load-bearing beams and AI is made to work within those beams rather than around them.

IV Miscarriage of Justice

All of this discussion leads to concerns over the miscarriage of justice. What if there have been fake cases presented to the court that the court and opposing counsel have not checked?

Courts themselves warn that undiscovered “fake law” can jeopardize the integrity of outcomes. In *Zhang v. Chen*, the BCSC called the practice “an abuse of process... tantamount to making a false statement to the court” and cautioned that, “unchecked, it can lead to a miscarriage of justice.”³⁶ Similar concerns have arisen abroad: following multiple incidents of fabricated citations, including one matter in which 18 of 45 authorities were fictitious, the President of the King’s Bench Division

2025 <https://www.courts.qld.gov.au/_data/assets/pdf_file/0009/879714/the-use-of-generative-ai-guidelines-for-judicial-officers.pdf> (‘Queensland Courts AI Guidelines’).

³³ ABA Opinion 512 (n 23); Mike Scarcella, “Lawyers Using AI Must Heed Ethics Rules, ABA Says in First Formal Guidance”, Reuters (online, 29 July 2024) <<https://www.reuters.com/legal/legalindustry/lawyers-using-ai-must-heed-ethics-rules-aba-says-first-formal-guidance-2024-07-29/>> (‘Reuters ABA Article’).

³⁴ Federation of Law Societies of Canada, Model Code of Professional Conduct (at October 2019) r 3.1–2 (‘FLSC Model Code’); Law Society of Ontario, “Using Technology” (Web Page, 2024) <<https://lso.ca/lawyers/technology-resource-centre/practice-resources-and-supports/using-technology>> (‘LSO Technology Guide’).

³⁵ Federal Court AI Notice (n 30); Park (n 15).

³⁶ Zhang (n 17) [26]–[29].

issued an urgent warning to the profession that misuse of AI threatens the justice system, with sanctions to follow.³⁷

The structural vulnerability to false negatives lies in the DNA of the adversarial system. Courts function as neutral arbiters rather than active investigators.³⁸ As the U.S. Supreme Court emphasized in *United States v. Sineneng-Smith*, the system is designed around the premise that parties represent their own interests appropriately and competently, meaning courts rely on counsel to frame the issues and vet the authorities.³⁹ The Supreme Court of Canada has similarly cautioned that appellate courts must not step into the arena to raise new issues without strict procedural safeguards, lest they compromise their neutrality.⁴⁰

This structural reliance creates a blind spot: if the adversary fails to challenge a fabrication, the court may default to accepting it. In *Mata v. Avianca*, the court initially relied on the plaintiff's submissions, and the fraud only surfaced because opposing counsel, not the court's internal vetting, flagged the anomalies.⁴¹ As recent scholarship notes, while ethical rules like Rule 11 FRCP theoretically forbid this, they are "ill-equipped" for the scale of generative error if the adversarial check fails.⁴²

At the same time, several systemic "fail-safes" reduce the probability that fabricated law will silently contaminate outcomes. First, the routine use of citators serves as a baseline verification mechanism, confirming the existence and treatment of authorities, a practice that law libraries categorize as an essential component of professional competence.⁴³ Second, courts are rapidly implementing procedural guardrails specific to generative AI. The Federal Court of Canada, for instance, now mandates that parties disclose AI-generated content in their filings, a requirement expressly designed to alert the Court and opposing counsel to the nature of the content and restore contestability.⁴⁴ Third, professional regulators emphasize that AI

³⁷ Robert Booth, "High Court Tells UK Lawyers to Urgently Stop Misuse of AI in Legal Work", The Guardian (online, 6 June 2025) <<https://www.theguardian.com/technology/2025/jun/06/high-court-tells-uk-lawyers-to-urgently-stop-misuse-of-ai-in-legal-work>> ('Guardian AI Warning'); *Al-Haroun v. Qatar National Bank QPSC* [2025] EWHC 1383 (Admin) ('Al-Haroun').

³⁸ *United States v. Sineneng-Smith*, 140 S Ct 1575, 1579 (2020) ('*Sineneng-Smith*'); *Greenlaw v. United States*, 554 US 237, 243 (2008) ('*Greenlaw*').

³⁹ *Sineneng-Smith* (n 38); *Greenlaw* (n 38).

⁴⁰ *R v. Mian* [2014] 2 SCR 689, [38]–[41] ('*Mian*').

⁴¹ *Mata* (n 12) 443–50.

⁴² Jessica R. Gunder, "Rule 11 Is No Match for Generative AI" (2024) 27 Stanford Technology Law Review 308, 334–6 ('Rule 11 Is No Match').

⁴³ David Hamer et al, *Laying Down the Law* (LexisNexis Butterworths, 12th ed, 2024) 541–5 ('*Laying Down the Law*').

⁴⁴ Federal Court of Canada, Notice to the Parties and the Profession: The Use of Artificial Intelligence in Court Proceedings, 7 May 2024 <<https://www.fct-cf.ca/Content/assets/pdf/base/FC-Updated-AI-Notice-EN.pdf>> ('Federal Court AI Notice May 2024'); see also Zhang (n 17) [29].

outputs do not dilute core duties of competence and candour. The American Bar Association's *Formal Opinion 512* clarifies that lawyers must independently verify AI outputs,⁴⁵ while UK regulators have issued similar warnings regarding "hallucination risks."⁴⁶ Finally, even if a fabricated authority slips through to judgment, post-judgment remedies exist to correct the record. In Ontario, *Rule 59.06* permits the Court to set aside or vary an order obtained by fraud or misrepresentation, ensuring a procedural remedy remains available.⁴⁷

Consequently, the miscarriage-of-justice risk concentrates where these safeguards are thinnest: self-represented litigants (SRLs) and off-record settlement negotiations. Empirical literature documents that SRLs face profound informational asymmetries that impair their ability to detect erroneous legal claims. The National Self-Represented Litigants Project highlights that generative tools, while promising for access to justice, risk "leading astray" users who lack the legal training to verify confident but inaccurate outputs.⁴⁸ Those same asymmetries can manifest in negotiation. As Owen Fiss's seminal critique 'Against Settlement' warned, weaker litigants are structurally susceptible to pressured resolutions because they lack the resources to gather and analyze legal information, a deficit that is weaponized if an opponent relies on fabricated authorities.⁴⁹ Canadian family-law experts report that this dynamic is particularly acute for survivors of violence, who may be nudged or coerced into disadvantageous settlements to avoid the "legal abuse" of prolonged litigation.⁵⁰ Professional conduct rules squarely forbid such misrepresentation, meaning the ethical apparatus already treats "fake law" used to pressure settlement as sanctionable conduct, even outside the courtroom.⁵¹

⁴⁵ ABA Opinion 512 (n 23) 2–11.

⁴⁶ Solicitors Regulation Authority, Risk Outlook Report: The Use of Artificial Intelligence in the Legal Market (Report, 20 November 2023) ('SRA Risk Outlook').

⁴⁷ Rules of Civil Procedure, RRO 1990, Reg 194, r 59.06(2)(a).

⁴⁸ Julie Macfarlane, The National Self-Represented Litigants Project: Identifying and Meeting the Needs of Self-Represented Litigants (Final Report, May 2013) 50–5, 63–7, 108–10 ('NSRLP Report'); Jennifer Leitch, "The Inevitability of AI in Court: What Does It Mean for Self-Represented Litigants?" (Web Page, National Self-Represented Litigants Project, 16 April 2024) ('NSRLP AI Article'); 'CJC AI Guidelines' (n 32) 5; see also Canadian Judicial Council, Statement of Principles on Self-represented Litigants and Accused Persons (September 2006) 6–7 ('CJC SRL Principles').

⁴⁹ Owen M Fiss, "Against Settlement" (1984) 93(6) Yale Law Journal 1073, 1076–7 ('Against Settlement'); Marc Galanter, "Why the 'Haves' Come out Ahead: Speculations on the Limits of Legal Change" (1974) 9(1) Law and Society Review 95, 97–100, 149–51 ('Why the "Haves" Come out Ahead').

⁵⁰ Luke's Place, "Tech Safety Toolkit: Identify, Manage and Document Tech Abuse" (Web Page, 2024) <<https://lukesplace.ca/resources/tech-abuse/>> ('Tech Safety Toolkit') (detailing how abusers use technology to manipulate legal processes); Peter G Jaffe et al, Making Appropriate Parenting Arrangements in Family Violence Cases: Applying the Literature to Identify Promising Practices (Report, February 2023) pt 5 ('Parenting Arrangements in Family Violence Cases').

⁵¹ Law Society of Ontario, Rules of Professional Conduct (at 1 October 2024) rr 5.1-1, 7.2-1 ('LSO Rules'); see also Model Rules (n 24) r 4.1; Australian Solicitors' Conduct Rules

This article argues that the risk of “false negatives” or undetected legal hallucinations is mitigated, though not eliminated, by the justice system’s existing redundancies. In fully contested litigation, the adversarial hierarchy offers multiple interception points: scrutiny by opposing counsel, vetting by law clerks and judges, and the corrective mechanism of appellate review. Consequently, the true danger may not lie in the courtroom, but in the shadows of the process where these safeguards are absent or diluted. The risk concentrates in two specific areas of vulnerability: first, matters involving self-represented or resource-constrained litigants who lack the capacity to police their opponents’ citations; and second, off-record settlement negotiations, where fabricated authority, intentional or otherwise, can be deployed to coerce or inadvertently coerce resolutions without ever facing judicial scrutiny.

V Judicial Impartiality and AI Judges

A further challenge concerns the present and future role of AI in decision-making and AI judges. If impartiality is indeed “...the supreme judicial virtue”,⁵² then it is warranted to be concerned and have a more thoughtful discussion about the subject.

Judicial impartiality has two dimensions: institutional independence from the executive and legislature, and individual objectivity in interpreting and applying the law. It is the second dimension that AI most directly tests.⁵³ Individual judges being impartial means that they are not influenced improperly and are free from it; they hold an open mind and adhere to these duties.⁵⁴ The Supreme Court of Canada has held that the same open-mindedness requirement also applies to any decision maker, even administrative decision makers. The same line of thinking also appears in Australian jurisprudence.⁵⁵

AI is increasingly being deployed in contexts where a human judicial officer would traditionally exercise discretion, from bail

2012 (Qld) rr 4.1.2, 22.1 (‘ASCR’). See also Queensland Law Society, Guidance Statement No 31: Courtesy in Communications (2025) [2]–[3.2] (‘QLS Guidance No 31’).

⁵² Lord Devlin, “Judges and Lawmakers” (1976) 39(1) *Modern Law Review* 1, 4 (‘Judges and Lawmakers’).

⁵³ Felicity Bell and Michael Legg, “Judicial Impartiality: AI in Courts” (Research Paper No 4, University of New South Wales, 2025) 1 (‘Judicial Impartiality’); William Blackstone, *Blackstone’s Commentaries*, ed Henry Winthrop Ballantine (Blackstone Institute, 1915) 366, 471 (‘Blackstone’s Commentaries’); Tom Bingham, *The Rule of Law* (Penguin Books, 2011) 51 (‘The Rule of Law’); Brian Z. Tamanaha, *On the Rule of Law: History, Politics, Theory* (Cambridge University Press, 2004) 122–6 (‘On the Rule of Law’); Raphael Franck, “Judicial Impartiality in Politically Charged Cases” (2018) 29(2) *Constitutional Political Economy* 193, 194 (‘Judicial Impartiality in Politically Charged Cases’).

⁵⁴ *Valente v R* [1985] 2 SCR 673 (‘*Valente*’).

⁵⁵ *Valente* (n 54); *Committee for Justice and Liberty v National Energy Board* [1978] 1 SCR 369; *Livesey v New South Wales Bar Association* (1983) 151 CLR 288; *Ebner v Official Trustee in Bankruptcy* (2000) 205 CLR 337; *Minister for Immigration and Multicultural Affairs; Ex parte Jia* (2001) 205 CLR 507.

assessments to benefits triage. This might not be the final decision yet, but an interim one that leads to a final decision.⁵⁶ Conversely, there could come a day that an advanced enough AI system finds its way to a more core exercise of judicial power.⁵⁷

On the other hand, AI systems can also serve to identify judicial bias, especially in a systematic way and help assess it. In this way, they may assist in enhancing judicial impartiality.⁵⁸ AI should be understood as a tool, one that can both introduce and detect bias. As with any tool, if used incorrectly and without proper safety measures, there may be negative consequences. The more powerful the tool, however, the more severe the consequences. AI is such a powerful tool.

A *AI Judges*

The emerging debate about “AI judges” versus AI-assisted judging is best understood as a struggle to reconcile promised gains in accuracy, consistency and efficiency with deep concerns about fairness, transparency and legitimacy.⁵⁹ Early experiments with algorithmic risk assessment in criminal justice⁶⁰ illustrate the core tension: models trained on historical data can embed and reproduce structural inequalities. Even where an AI tool is framed as merely advisory, its opacity and the difficulty of interrogating its design, training data and error profile create acute due-process problems.⁶¹ If neither judges nor

⁵⁶ Gradient Institute and Minderoo Foundation, *De-Risking Automated Decisions* (Report, 16 March 2022) 6–7 (‘De-Risking Automated Decisions’); Canadian Bar Association, *Law, Technology, and Accountability: Reimagining Canadian Immigration for the 21st Century* (Submission to Parliament, 2025) [IV.4]–[IV.5] (‘CBA Immigration Submission’); Institute on Governance, *Towards a Considered Use of AI Technologies in Government* (Report, 28 September 2023) 8–10 (‘IOG AI Report’).

⁵⁷ See, eg, Vera Lúcia Raposo, “The Digital ‘To Kill A Mockingbird’: Artificial Intelligence Biases in Courts” (2024) 54 *California Western International Law Journal* 459, 461 (‘Digital Mockingbird’); Christoph K Winter, *The Challenges of Artificial Judicial Decision-Making for Liberal Democracy* in Piotr Bystranowski, Bartosz Janik and Maciej Próchnicki (eds), *Judicial Decision-Making: Integrating Empirical and Theoretical Perspectives* (Springer, 2022) 179, 180 (‘Artificial Judicial Decision-Making’); Benjamin Minhao Chen, Alexander Stremtizer and Kevin Tobia, “Having Your Day in Robot Court” (2022) 36 *Harvard Journal of Law and Technology* 127, 128 (‘Robot Court’).

⁵⁸ Bell and Legg, “Judicial Impartiality” (n 53) 5–6.

⁵⁹ *Ibid* 6–11.

⁶⁰ Julia Angwin et al, “Machine Bias: There’s Software Used Across the Country to Predict Future Criminals. And It’s Biased Against Blacks”, ProPublica (online, 23 May 2016) <<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>> (‘Machine Bias’).

⁶¹ Danielle Keats Citron, “Technological Due Process” (2008) 85(6) *Washington University Law Review* 1249 (‘Technological Due Process’); Finale Doshi-Velez et al, “Accountability of AI Under the Law: The Role of Explanation” (Research Paper, Berkman Klein Center for Internet & Society, 2017) (‘Accountability of AI’); Brandon L Garrett and Cynthia Rudin, “The Right to a Glass Box: Rethinking the Use of Artificial Intelligence in Criminal Justice” (2024) 109(5) *Cornell Law Review* 1421 (‘Right to a Glass Box’); Cynthia Rudin, “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead” (2019) 1 *Nature Machine Intelligence* 206 (‘Stop Explaining Black Box Models’).

parties can meaningfully understand or challenge how a risk score or recommendation was produced, the resulting decision may be perceived as tainted by hidden bias.⁶² This is why instruments such as the EU AI Act⁶³ move to prohibit certain forms of predictive profiling in criminal justice, even as they leave a narrow lane for AI to “augment” human assessments based on verifiable facts.⁶⁴

At the same time, simply contrasting “biased humans” with “biased machines” is too crude. Psychological research on heuristics, noise and unconscious bias suggests that human judging is often inconsistent and context-sensitive, and that structured statistical tools can, in some circumstances, dampen those distortions.⁶⁵ Conversely, it has been argued that well-designed systems, case-outcome predictors, sentencing recommendation engines, triage tools for appellate courts, might deliver more uniform treatment of like cases, filter out irrelevant considerations and reduce backlogs.⁶⁶ Yet these same systems create new vulnerabilities: judges may defer too readily to an apparently precise algorithmic output (automation bias, anchoring), and the use of generative AI to draft reasons may erode confidence that judgments genuinely reflect the judge’s own deliberation.⁶⁷ The more decisional weight is shifted to AI, the more important traditional safeguards become: an articulated duty to give reasons in the judge’s own voice, opportunities for parties to contest the data and logic underlying AI tools, and an effective appellate structure capable of scrutinizing how

⁶² Han-Wei Liu, Ching-Fu Lin and Yu-Jie Chen, “Beyond State v Loomis: Artificial Intelligence, Government Algorithmization and Accountability” (2019) 27(2) *International Journal of Law and Information Technology* 122 (‘Beyond State v Loomis’); Alicia Solow-Niederman, “Algorithmic Grey Holes” (2023) 5(1) *Journal of Law and Innovation* 116 (‘Algorithmic Grey Holes’).

⁶³ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) [2024] OJ L 2024/1689, art 5 (‘EU AI Act’).

⁶⁴ Bell and Legg, “Judicial Impartiality” (n 53) 5–6.

⁶⁵ Eugene Volokh, “Chief Justice Robots” (2019) 68(6) *Duke Law Journal* 1135, 1169 (‘Chief Justice Robots’); Chris Guthrie, Jeffrey J Rachlinski and Andrew J Wistrich, “Blinking on the Bench: How Judges Decide Cases” (2007) 93(1) *Cornell Law Review* 1, 13–19 (‘Blinking on the Bench’); (judges are as susceptible as others to relying on intuition rather than deliberative judgment); Daniel Kahneman et al, “Reducing Noise in Decision Making” (2016) 94(12) *Harvard Business Review* 18 (‘Reducing Noise’).

⁶⁶ Frank Fagan, “Big Data Legal Scholarship: Toward a Research Program and Practitioner’s Guide” (2016) 20(1) *Virginia Journal of Law and Technology* 1, 2 (‘Big Data Legal Scholarship’); Arthur Rizer and Caleb Watney, “Artificial Intelligence Can Make Our Jail System More Efficient, Equitable, and Just” (2018) 23(1) *Texas Review of Law and Politics* 181 (‘AI in Jail Systems’).

⁶⁷ Jane Dalton, “Judge Admits Using “Jolly Useful” ChatGPT to Write Court Ruling”, *The Independent* (online, 15 September 2023) <<https://www.independent.co.uk/news/uk/home-news/chatgpt-ai-judge-chatbot-rulingb2412378.html>> (‘Judge Admits Using ChatGPT’); Chief Justice Andrew Bell, “Truth Decay and Its Implications for the Judiciary: An Australian Perspective” (Speech, Durham University, 23–26 April 2024) (‘Truth Decay’).

those tools shape outcomes.⁶⁸ In this light, the key normative question is not whether AI can replace judges, but under what conditions AI can be integrated into judicial reasoning without hollowing out the human responsibility, contestability and public justification on which impartial adjudication ultimately depends.

B *AI Decision Makers – Viewed from the Canadian Context*

Transposing these concerns into the administrative state, the same worries about opacity, bias and automation bias apply with at least as much force to administrative decision-makers as to judges, given the volume and often fragile circumstances of the people affected. Canadian law already insists that any public authority making an administrative decision that affects rights, privileges, or interests owes a duty of procedural fairness, with its twin pillars of the right to be heard and the right to an unbiased decision-maker, as articulated in *Cardinal* and *Knight* and *Baker*.⁶⁹

If an AI system is used to triage benefit claims, flag immigration applications as “high risk”, or generate template reasons, it will inevitably shape the process and outcome of those decisions. On orthodox doctrine, these procedures should engage obligations of notice, disclosure and participation (e.g. the right to know the “case to meet” and to respond to adverse material), as developed in *Suresh*, *Khela* and *Charkaoui*.⁷⁰ A black-box score or recommendation that cannot be explained in human-intelligible terms risks both actual unfairness and a reasonable apprehension of bias. These risks are amplified if decision-makers rely uncritically on the vendor’s assurances about accuracy. AI-assisted administration, therefore, sits uneasily with the requirement that affected persons be able to understand and, if necessary, challenge the basis of the decision that has altered their legal position.

On review, Vavilov recasts reasonableness as part of a “culture of justification”, in which administrative actors must offer transparent, intelligible reasons that show a rational chain of analysis and demonstrate fidelity to the statutory scheme, the evidence, past practice and the stakes for the individual.⁷¹ In an environment where AI tools

⁶⁸ Bell, ‘Truth Decay’ (n 67); Chief Justice Gibbs Salika, ‘Ethics: The Construct of Impartiality and Unconscious Bias’ (Paper, Pacific Judicial Integrity Program Judicial Fraud & Anti-Corruption Workshop, 28 November 2022) 10 (‘Ethics, Impartiality and Bias’).

⁶⁹ *Cardinal v Director of Kent Institution* [1985] 2 SCR 643 (‘Cardinal’); *Knight v Indian Head School Division No 19* [1990] 1 SCR 653 (‘Knight’); *Baker v Canada (Minister of Citizenship and Immigration)* [1999] 2 SCR 817 (‘Baker’); Bell and Legg, *Judicial Impartiality* (n 53) 8; Salika, ‘Ethics, Impartiality and Bias’ (n 68) 10.

⁷⁰ *Suresh v Canada (Minister of Citizenship and Immigration)* [2002] 1 SCR 3 (‘Suresh’); *Mission Institution v Khela* [2014] 1 SCR 502 (‘Khela’); *Charkaoui v Canada (Citizenship and Immigration)* [2007] 1 SCR 350 (‘Charkaoui I’).

⁷¹ *Canada (Minister of Citizenship and Immigration) v Vavilov* [2019] 4 SCR 653, [14] (‘Vavilov’).

are deployed to structure discretion, a bare assertion that “the model classified this claimant as ineligible” is unlikely to satisfy that culture of justification. The decision-maker must still explain why the output is being accepted in light of the statute, the record and the submissions before them, and avoid fettering their discretion by treating an algorithmic score as binding.

This tension is difficult to escape. If current generative systems are demonstrably unreliable even at the relatively lower-stakes task of producing case citations, then it is hard to see how they could be trusted with the much higher-stakes task of resolving rights and obligations. The fake law episodes show that, absent rigorous human verification, LLM outputs are not epistemically grounded in the legal materials they purport to summarize. That epistemic fragility sits uneasily alongside proposals for AI judges or AI-driven administrative triage, which presuppose that the same technologies can be treated as authoritative decision makers or decision inputs.

C *Augmentation, Not Automation*

Concerns raised by AI judges and administrative decision-makers point toward a particular way of framing legal AI: as a tool for augmentation, not automation. The Sedona Canada Primer on Artificial Intelligence and the Practice of Law explicitly distinguishes “automation,” where AI systems are deployed to perform tasks with minimal or no human involvement, from “augmentation,” where AI is used to enhance human capabilities while leaving lawyers responsible for judgment, supervision and accountability.⁷² The Primer stresses that although AI may “automate certain legal tasks,” its use must not compromise quality or displace “competent legal advice from a human lawyer,” instead requiring ongoing human oversight, technological competence, and protection of confidentiality and privilege.⁷³

Soft-law frameworks in the public sector adopt the same logic. The federal *Directive on Automated Decision-Making* applies to systems that “fully or partially automate an administrative decision,” but its guidance emphasizes transparency, recourse, meaningful explanations and ensuring human involvement in higher-impact decisions as necessary to preserve procedural fairness and accountability.⁷⁴

In the EU, the AI Act similarly classifies AI systems used in the “administration of justice and democratic processes” as “high-risk”,

⁷² The Sedona Conference, “The Sedona Canada Primer on Artificial Intelligence and the Practice of Law” (2025) 26 *Sedona Conference Journal* 103, 105–7 (“Sedona Canada AI Primer”).

⁷³ *Ibid* 159–60.

⁷⁴ Treasury Board of Canada Secretariat, Guide on the Scope of the Directive on Automated Decision Making (24 June 2025) <https://publications.gc.ca/collections/collection_2024/sct-tbs/BT48-46-2024-eng.pdf> (“ADM Scope Guide”).

subjecting them to stringent risk-management, transparency and human-oversight obligations precisely because they materially influence fundamental rights.⁷⁵ Taken together, these instruments assume that AI is properly situated as a decision-support layer within an essentially human process.

The access-to-justice scholarship on LLMs reinforces this “augmentation not automation” stance, particularly for self-represented litigants. Ogunde’s analysis of generative AI and SRLs in Canada concludes that LLMs can help users draft documents, distil issues and better understand procedure, potentially narrowing long-standing information gaps, but she also catalogues core limitations, including hallucinated or misleading outputs, jurisdictional mismatches, bias embedded in training data and oversimplified explanations, and emphasises that these tools cannot match the research and analytical capacity of competent counsel.⁷⁶ Cesta’s study of large language models in Australian community legal centres makes a similar argument. LLM-based systems can streamline intake and assist lawyers with research and drafting, but client-facing chatbots providing legal information may do more harm than good absent careful design, testing and expert oversight, given risks of hallucination, opacity and privacy breach.⁷⁷

If AI is deployed as a replacement for human assistance in low-income and SRL matters, while remaining merely supplementary for well-resourced parties, there is a risk of entrenching a two-tier system in which vulnerable litigants receive a thinner, more error-prone tier of automated justice.⁷⁸ On that view, the responsible use of legal AI is to extend and streamline human help, through more efficient intake, triage and document assistance, without treating LLMs as substitutes for legal advice or adjudication.

The same augmentation-only norm emerges in debates on AI judging and automated public decision-making. Bell and Legg’s work on judicial impartiality stresses that because both actual impartiality and the appearance of impartiality are foundational to the legitimacy of the courts, the deployment of opaque or proprietary AI in core adjudicative functions poses a structural threat. They argue that such systems threaten the judicial function because their training data, error profiles,

⁷⁵ EU AI Act (n 63) ch 3, annex III.

⁷⁶ Fife Ogunde, “Generative AI and Access to Justice in Canada: The Case of Self Represented Litigants [SRLs]” (2024) 40 *Windsor Yearbook of Access to Justice* 211, 215–18 (‘Generative AI and Access to Justice’).

⁷⁷ Will Cesta, “Large Language Models and Community Legal Centres: Could Chatbots Help Reduce Australia’s Justice Gap?” (2024) 49(3) *Alternative Law Journal* 181, 183–6 (‘LLMs and CLCs’); Joshua Kelsall et al, “A Rapid Evidence Review of Evaluation Techniques for Large Language Models in Legal Use Cases: Trends, Gaps, and Recommendations for Future Research” (2025) *AI and Society* (advance) (‘Rapid Evidence Review of LLMs’).

⁷⁸ Drew Simshaw, “Access to AI Justice: Avoiding an Inequitable Two-Tiered System of Legal Services” (2022) 24 *Yale Journal of Law and Technology* 150, 171–6 (‘Access to AI Justice’).

and value-laden design choices cannot be readily scrutinized by parties, making it difficult for judges to provide intelligible, human-authored reasons that justify the outcome.⁷⁹

Reflecting these concerns, judicial and court-administration guidance in multiple jurisdictions now explicitly frames GenAI as a secondary tool. The Courts of New Zealand's *Guidelines for Use of Generative Artificial Intelligence* (2023) allow use for tasks such as summarization and translation but insist that judges remain fully responsible for decisions, must verify AI-assisted content, and must guard against confidentiality breaches and automation bias.⁸⁰ Similarly, the Canadian Judicial Council and the Federal Court of Canada's 2024 *Notice to the Parties* state that judges cannot delegate decision-making functions to AI and that counsel must affirmatively disclose if and how AI has been used in their materials.⁸¹ In the administrative sphere, academic commentary on the federal *Directive on Automated Decision Making* emphasizes that meaningful human involvement and recourse mechanisms are essential to ensure that algorithmic recommendations do not hollow out administrative law principles of transparency, fairness, and legality.⁸²

Empirical work on legal-domain LLMs underlines why, on present evidence, they are convincing tools of augmentation but poor candidates for automation. Ogunde's separate comment on "legal LLMs" notes that while law-specific systems are increasingly marketed as hallucination-free research assistants, they remain prone to significant error.⁸³ This caution is grounded in a recent Stanford study, which demonstrated that leading AI legal research tools (including Lexis+ AI and Westlaw AI-Assisted Research) still hallucinated authorities in roughly 17–33% of tested prompts, even under controlled conditions.⁸⁴ Related evaluations of general-purpose LLMs and legal benchmarks report strong performance on some classification and

⁷⁹ Bell and Legg, "Judicial Impartiality" (n 53).

⁸⁰ Courts of New Zealand, *Guidelines for the Use of Generative Artificial Intelligence*, 7 December 2023 <<https://www.courtsofnz.govt.nz/assets/6-Going-to-Court/practice-directions/practice-guidelines/all-benches/20231207-GenAI-Guidelines-Judicial.pdf>> ('NZ Courts AI Guidelines').

⁸¹ Canadian Judicial Council, 'CJC AI Guidelines' (n 32) 3; Federal Court of Canada, "Federal Court AI Notice May 2024" (n 44).

⁸² Teresa Scassa, "Administrative Law and the Governance of Automated Decision Making: A Critical Look at Canada's Directive on Automated Decision Making" (2021) 54(1) *University of British Columbia Law Review* 251, 260–2, 274, 287, 297–8 ('Canada's Directive on ADM').

⁸³ Fife Ogunde, "Legal Large Language Models (LLMs): Legal Dynamos or "Fancifully Packaged ChatGPT"?" (2025) 5 *Discover Artificial Intelligence* art 2 ('Legal LLMs').

⁸⁴ Varun Magesh et al, "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools" (2025) 22 *Journal of Empirical Legal Studies* 216, 218–20, 224–6 ('Hallucination-Free?').

reasoning tasks, but also persistent hallucinations, brittleness to prompt phrasing, and sensitivity to subtle contextual changes.⁸⁵

The fake-law decisions, the miscarriage-of-justice risks for SRLs and less affluent parties, and impartiality concerns around AI-mediated judging reinforce the central claim of this article: generative AI may be a powerful collaborator in research, drafting and triage, but the non-delegable tasks of determining the law, applying it to facts, safeguarding fairness and preventing miscarriages of justice must remain firmly in human hands.

This is not to glamorize the status quo or to deny AI's promise. Rather, it is to insist that the integration of LLMs into legal practice and adjudication be disciplined. AI should be made to work within the profession's existing load-bearing beams, competence, supervision, candour to the tribunal, procedural fairness and a culture of justification, rather than around them.

VI Self-represented Litigants and Access to Justice

The risks traced in the fake-law and miscarriage-of-justice discussions bear most heavily on those who already sit at the margins of the system. Canadian data suggest that self-representation is no longer an aberration but a structural feature of the landscape. Justice Canada's report on family law estimates that between 40 and 57 per cent of parties in family cases appear without counsel, depending on jurisdiction and stage of proceedings, while more recent Statistics Canada data indicate that in 2019–2020, a majority of family case litigants in active cases (around 58 per cent) were self-represented.⁸⁶ National Self-Represented Litigants Project ('NSRLP') intake studies and follow-up reports similarly show that a large share of SRLs have annual incomes below \$30,000, with roughly two-thirds earning under \$50,000, and that cost, lack of legal aid and scarcity of affordable services are the predominant reasons for going without a lawyer.⁸⁷ Ogunde's article on *Generative AI and Access to Justice in Canada: The Case of Self-Represented*

⁸⁵ Matthew Dahl et al, "Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models" (Preprint, arXiv:2401.01301, 2 January 2024) 3–5 <<https://arxiv.org/abs/2401.01301>> ('Large Legal Fictions'); Neel Guha et al, "LegalBench: A Collaborative Benchmark for Legal Reasoning in Large Language Models" (Preprint, arXiv:2308.11462, 22 August 2023) 4–6, 8–9 <<https://arxiv.org/abs/2308.11462>> ('LegalBench'); Harry Surden, "ChatGPT, AI Large Language Models, and Law" (2024) 92 *Fordham Law Review* 1941, 1965–8, 1972 ('ChatGPT and Law').

⁸⁶ Department of Justice Canada, "Self-Represented Litigants in Family Law" (Web Page, 2022) <<https://www.justice.gc.ca/eng/rp-pr/fl-lf/divorce/jf-pf/srl-pnr.html>> ('DOJ SRL Family Law'); Statistics Canada, Profile of Family Law Cases in Canada, 2019/2020 (Catalogue No 85-002-X, 16 March 2021) ('StatsCan Family Law Profile').

⁸⁷ Keerthi Chintapalli et al, Tracking the Trends of the Self-Represented Litigant Phenomenon: Data from the National Self-Represented Litigants Project, 2021–23 (Report, 2 February 2024) ('NSRLP Intake Report 2021–23'); Macfarlane, 'NSRLP Report' (n 48).

Litigants focuses on SRLs, rather than lawyers. She argues that LLMs can “play a significant role in enabling SRLs to present decent cases” and participate more effectively in proceedings, for example, by helping them draft pleadings, organize arguments and understand procedure. At the same time, she stresses that the “inherent deficiencies” of LLMs, hallucinations, gaps or distortions in training data and uncertainty about jurisdiction may mean they “do more harm than good” for SRLs, especially those with no legal training, and that the capacity to harness these tools depends more on SRLs’ literacy and understanding than on the mere availability of the technology.⁸⁸ Her survey of the SRL literature underscores that this group is heterogeneous. Generative AI arrives not at a level playing field, but at an environment in which many SRLs already lack both legal knowledge and access to institutional support.

Experiments in adjacent access to justice spaces reinforce both the promise and the limits of LLMs for unrepresented users. Casta’s study of Australian community legal centres concludes that LLM-based chatbots are unlikely to reduce the “justice gap” by streamlining intake, they may do so by offering client-facing legal information but emphasises that this potential is bounded by substantial barriers: hallucinations that may mislead clients, difficulties aligning outputs with the *Legal Profession Uniform Law*, and the need to clarify whether autonomous systems are themselves engaged in legal practice.⁸⁹

Dilag’s analysis of predictive-analytics tools marketed on an access-to-justice narrative similarly finds that while vendors promise time savings, clarity and certainty, the concrete access to justice benefits are constrained by the economics of the legal services market: sophisticated tools are typically sold to institutional repeat players rather than to those who most need help.⁹⁰ In the Canadian doctrinal context, Castets-Renard warns that litigants without substantial legal resources may disproportionately rely on AI tools, and that existing frameworks such as Bill C-27 and the EU AI Act, while a step forward, still leave gaps around effective human control and accountability for AI-driven outcomes in legal contexts.⁹¹

For self-represented litigants, those structural and technological dynamics interact in a particularly sharp way. Most will not have access to enterprise-grade “legal LLMs” integrated into research platforms or

⁸⁸ Ogunde, ‘Generative AI and Access to Justice’ (n 76) 211–13, 219–22.

⁸⁹ Cesta, ‘LLMs and CLCs’ (n 77) 181.

⁹⁰ Matthew Dilag, “The Impact of Artificial Intelligence on Access to Justice: Predictive Analytics and the Legal Services Market” (2025) 48(1) *Dalhousie Law Journal* 237, 237–43, 259–60 (‘AI and Access to Justice’).

⁹¹ Céline Castets-Renard, “Do the Canadian Practical and Legal Frameworks on AI Adequately Address the Uses of LLMs in Courts?” (2025) 58(2) *University of British Columbia Law Review* 321, 324–7, 335–6, 342–3, 368 (‘Canadian Frameworks on AI’); Digital Charter Implementation Act 2022 (Canada) (Bill C-27, 1st Sess, 44th Parl, second reading 24 April 2023) (‘Bill C-27’); EU AI Act (n 63) art 5.

wrapped in institutional safeguards; instead, they are more likely to interact with general-purpose systems through consumer interfaces that neither guarantee jurisdictional accuracy nor provide robust guardrails against hallucination.

From an access-to-justice perspective, the lesson is not that self-represented litigants must be shielded from AI altogether. Indeed, Castets-Renard, Ogunde and Cesta suggest that carefully designed tools, embedded in legal aid, community clinics or court-based supports could help level informational imbalances, but that unmediated reliance on general-purpose LLMs by unrepresented parties is fundamentally and practically fraught. In this light, the “augmentation, not automation” stance developed in the previous Part has a specific access to justice inflection. The goal is to use AI to augment human assistance for SRLs, through better triage, clearer plain-language materials, and more efficient back-office support, rather than to automate legal advice or decision-making for those least equipped to contest errors. Embedding this model of supervised augmentation is not merely a preference but a structural necessity that must be integrated into the justice system’s architecture by design. The alternative is a two-tier system in which affluent parties enjoy both counsel and curated, supervised AI, while self-represented litigants are left to navigate high-variance, hallucination-prone tools on their own.

VII Back to Basics

Going back to basics is a good guide against getting lost in the AI forest. The solution to the challenges posed by AI is a renaissance of professionalism, grounded in the duty of competence, the necessity of human judgment, and the discipline of continuous education.

It is worth stating plainly why this is a claim about sufficiency and not merely relevance. The duties canvassed above are not fixed inventories of permitted conduct. They are standards, and standards move with their circumstances. The duty of competence has never been confined to a static body of knowledge, and it expands, as it has with every earlier shift in legal technology, to take in the tools that now shape practice; candour to the tribunal already governs what must be disclosed about how a submission was produced; confidentiality already governs what may be placed into a system whose handling of data is unknown; and the duty to the court already forbids the filing of authorities one has not verified. Read in this way, the duties do not merely bear upon the risks generative AI presents, they are adequate to them, and the disclosure norms and verification protocols now emerging from courts and regulators are best understood not as new obligations but as the working-out of existing ones. Such gaps as remain lie not in the profession's duties but in the separate domain of product and system

regulation, the lane occupied by instruments such as the EU AI Act, which govern how these tools are built and deployed rather than how lawyers must conduct themselves; those regimes complement the profession's duties without displacing them. The one thing the analysis genuinely demands is therefore not a new duty but a new competency: because competence is dynamic, a working literacy in AI, its capabilities, its failure modes and its ethical hazards must become a standing element of the profession's continuing education rather than an optional specialism.

A *Technological Competence as a Core Duty*

The first basic principle to reaffirm is that competence, as noted, is dynamic. Just as a lawyer cannot claim competence while ignoring major legislative changes, they can no longer ignore the tools that shape modern practice. The Federation of Law Societies of Canada, along with provincial regulators, have explicitly recognized that the ethical duty of competence includes “technological competence.”⁹² This does not require every lawyer to become a computer scientist. Rather, it mandates an understanding of the benefits and risks associated with relevant technology, recognizing that while AI can automate tasks, it cannot replace the professional obligations of the lawyer.⁹³

In the context of AI, this back-to-basics approach means applying the standard duty of verification to AI outputs. As demonstrated in *Zhang v Chen*⁹⁴ and *Mata v Avianca, Inc.*,⁹⁵ the failure was not merely technological but a failure of professional diligence, the abandonment of the gatekeeping role that lawyers hold to ensure the accuracy of their filings.⁹⁶ Guidance from the Law Society of British Columbia and the American Bar Association’s *Formal Opinion 512* reinforces that lawyers must critically review and verify AI-generated content to ensure it is accurate and supports the client's interests.⁹⁷ The technology may change, but the obligation to read the cases one cites remains immutable.

B *The “Better than a Bot” Standard: Ethical Judgment*

As AI automates routine tasks, drafting, summarizing, and categorizing, the value of the human lawyer shifts to what machines cannot do. Michael Legg argues that the future of the profession depends on

⁹² ‘FLSC Model Code’ (n 34) r 3.1-2 cmt [4A]–[4B].

⁹³ The Sedona Conference, ‘Sedona Canada AI Primer’ (n 72) 159.

⁹⁴ Zhang (n 17) [29].

⁹⁵ Mata (n 12) 448.

⁹⁶ Ibid.

⁹⁷ Law Society of British Columbia, “Guidance on Professional Responsibility and Generative AI” (Practice Resource, July 2024) 3 (‘LSBC AI Guidance (July 2024)’); ABA Opinion 512 (n 23) 6.

cultivating skills that allow a lawyer to be “Better than a Bot”.⁹⁸ At the heart of this is “ethical judgement.” While AI can predict the most statistically probable next word, it cannot exercise wisdom. Ethical judgement is multi-faceted; it combines expertise, experience, and emotional intelligence with a normative commitment to the administration of justice.⁹⁹

Going back to basics means prioritizing this human element. A lawyer's role is not merely to provide a technical answer but to offer holistic advice that accounts for the client's broader commercial, reputational, and moral interests.¹⁰⁰ As the *Sedona Canada Primer on Artificial Intelligence and the Practice of Law* emphasizes, AI should be viewed through the lens of “augmentation” rather than “automation.”¹⁰¹ The AI may act as a powerful engine for processing information, but the lawyer must remain the pilot, applying intellectual capacity, judgment and deliberation to every function.¹⁰²

C Continuous Education, Reflective Practice and Proactive Engagement

To maintain this human in the loop supremacy, the profession must embrace a more proactive approach to education. It is insufficient to learn the law once; lawyers must engage in “reflective practice”, a disciplined process of analyzing one's experiences to build professional wisdom.¹⁰³

In the age of AI, where junior lawyers may be deprived of the drudgery that traditionally formed the training ground for legal skills (e.g., document review, basic drafting), intentional education becomes paramount. Law firms and legal departments must ensure that the use of AI does not hollow out the learning opportunities for junior counsel.¹⁰⁴ The profession must return to the basic principle of mentorship and supervision, ensuring that those delegating work to AI systems, or to juniors using AI systems, maintain complete professional responsibility for the work product.¹⁰⁵

Going back to basics means reclaiming agency. The profession should not be a passive recipient of Silicon Valley's outputs. Being proactive means engaging with the technology to understand its specific application to the justice sector. As Ogunde notes, the capacity to harness these tools depends more on the user's literacy and

⁹⁸ Michael Legg, “Better than a Bot – Instilling Ethical Judgement into the Lawyers of the Future in the Age of AI” (2024) 33(3) *Griffith Law Review* 273, 273 (“Better than a Bot”).

⁹⁹ *Ibid* 282–3.

¹⁰⁰ *Ibid*.

¹⁰¹ *Sedona Canada AI Primer* (n 72) 105–6.

¹⁰² *Ibid* 103, 136, 144, 160.

¹⁰³ Legg, *Better than a Bot* (n 98) 283–4.

¹⁰⁴ *Ibid* 286.

¹⁰⁵ *Sedona Canada AI Primer* (n 72) 159–60.

understanding than on the mere availability of the technology.¹⁰⁶ This aligns with guidance from the Canadian Judicial Council and various courts, which urge legal professionals to understand not just that AI is being used, but how it functions, its potential for bias, and its limitations.¹⁰⁷

The “basics” are clear: the lawyer is an agent of the rule of law, not a button-pusher for an algorithm. By anchoring the practice of law in technological competence, ethical judgment, and rigorous supervision, the profession can harness the power of AI without losing its soul. As Legg notes, while AI will result in significant change, lawyers who possess these key professional skills will remain essential to the provision of legal services.¹⁰⁸ The future of the profession is not found in competing with machines on speed, but in excelling in the “humanity and ethics” that machines cannot replicate.¹⁰⁹

D *Ethical and Professional Obligation Above Technological Concerns*

A final “basic” is the hierarchy between technology and professional duty. However sophisticated AI tools have become, they remain subordinate to the lawyer’s obligations to client, court and the public. Some such obligations and duties are competence, supervision, confidentiality and candour to the tribunal as non-delegable duties. A lawyer cannot justify defective work, misuse of confidential data or abdication of judgment by pointing to a system’s limitations or the client’s insistence on using it.¹¹⁰ Contemporaneous guidance on generative AI is consistent: regulators stress that where efficiency gains conflict with accuracy, fairness or respect for legal processes, professional standards must prevail.¹¹¹ Courts and commentators similarly emphasize that AI may augment, but cannot displace, the human responsibilities that constitute the administration of justice.¹¹² In

¹⁰⁶ Ogunde, *Generative AI and Access to Justice* (n 76) 227.

¹⁰⁷ Canadian Judicial Council, ‘CJC AI Guidelines’ (n 32) 3; Federal Court of Canada, ‘Federal Court AI Notice May 2024’ (n 44).

¹⁰⁸ Legg, *Better than a Bot* (n 98) 273.

¹⁰⁹ *Ibid.*

¹¹⁰ FLSC Model Code (n 34) r 3.1-2 cmt [4A]–[4B]; Legal Profession Uniform Law Australian Solicitors’ Conduct Rules 2015 (NSW) rr 3, 4.1, 9 (‘Solicitors’ Conduct Rules’).

¹¹¹ Law Society of British Columbia, ‘LSBC AI Guidance (July 2024)’ (n 97); ABA Standing Committee on Ethics and Professional Responsibility, ‘ABA Opinion 512’ (n 23) 2–11, 14–15; Law Society of New South Wales, ‘LSNSW AI Guide’ (n 21); Law Society of New South Wales, Legal Practice Board of Western Australia and Victorian Legal Services Board and Commissioner, Statement on the Use of Artificial Intelligence in Australian Legal Practice (Joint Statement, 26 March 2025) (‘Joint Statement on AI’).

¹¹² ‘Sedona Canada AI Primer’ (n 72) 159–60; ‘CJC AI Guidelines’ (n 32) 3; ‘Federal Court AI Notice May 2024’ (n 44); Supreme Court of Victoria, Guidelines for Litigants: Responsible Use of Artificial Intelligence in Litigation (Guidelines, 6 May 2024) [4], [8]–[10] (‘SCV AI Guidelines’); New South Wales Bar Association, ‘NSW Bar AI Guidelines’ (n 21) 2–3; Law Society of Western Australia, ‘Submission to Supreme Court of Western Australia – Artificial Intelligence’ (Submission, 11 April 2025) 9–11 (‘LSWA Submission’).

practical terms, this means that decisions about whether, when and how to use AI must be made through the lens of ethical obligation rather than technological possibility or market pressure.¹¹³

VIII Conclusion

Generative AI has arrived in legal practice the way email and e-discovery once did: first as a convenience, then as an expectation, and finally as part of the profession's infrastructure. This article has argued that the profession's existing duties, competence, supervision, candour to the tribunal, confidentiality, and procedural fairness provide the necessary and sufficient framework for responding to that arrival. The challenge is not novelty but rigour: AI does not demand a new regulatory paradigm so much as it tests the discipline with which the profession applies the one it already has.

Courts and regulators have responded accordingly, converging on disclosure obligations and certification requirements that operationalise existing duties rather than inventing new ones.

The more subtle risk, as Part IV demonstrated, is not the mistake that is caught but the one that is not. Undiscovered hallucinations and confident misstatements threaten the integrity of outcomes precisely where the adversarial system's ordinary checks are weakest. As Part VII argued, those risks concentrate most acutely on self-represented litigants and resource-constrained parties, who are least equipped to detect error and most likely to interact with consumer-grade tools that lack institutional safeguards. If AI is permitted to function as a replacement for human assistance in low-income matters while remaining merely supplementary for well-resourced parties, the profession risks entrenching a two-tier system of justice.

The same augmentation-over-automation logic applies to adjudication and administration. As Part V argued, AI can support research, triage, and consistency checks, and it may even help identify patterns of bias. But legitimacy in judging depends on public reasons, meaningful contestability, and a human decision-maker who can be held to account. Automation of that function is not, on present evidence, compatible with those requirements.

The path forward, as Part VIII contended, is disciplined augmentation anchored in existing duties: build auditable checkpoints, verify primary sources, protect client data, disclose material AI assistance where required, and ensure that continuous education equips lawyers to understand both the promise and the failure modes of the tools they use. None of this assumes the technology will stand still; the pace of change may accelerate or be constrained in ways that are difficult to predict. But a framework built on durable professional

¹¹³ Legg, *Better than a Bot* (n 98) 283–4.

obligations, rather than calibrated to any single technological forecast, is able to address whatever comes next.