
FROM INNOVATION TO NEGLIGENCE: THE PROMISE AND PERIL OF GENERATIVE ARTIFICIAL INTELLIGENCE IN LEGAL PRACTICE

Drishti Rao, B.A. LL.B. (Hons.), NLSIU, Bangalore

ABSTRACT

Generative Artificial Intelligence (Generative AI) is transforming legal practice, with its increasing integration across legal research, drafting, case analysis, strategy formulation, and client interaction reshaping how legal services are delivered. However, existing regulatory frameworks governing this integration remain largely suggestive in nature, creating a significant regulatory vacuum. This gap has given rise to serious ethical and professional concerns including fabricated citations, hallucinated legal references, lack of transparency, and the transfer of systemic AI biases into legal proceedings collectively undermining the integrity of legal institutions and public faith in the administration of justice. While several jurisdictions, including India, have begun developing guidelines, a comprehensive and enforceable regulatory framework remains largely absent. Through a comparative examination of international responses — including judicial standing orders and professional body guidelines from the United States and United Kingdom, and judicial guidelines from New Zealand, alongside an analysis of the Indian regulatory landscape, comprising Supreme Court directives and Bar Council Rules, this article identifies the collective inadequacy of existing frameworks. It further proposes mitigation measures spanning role delimitation, prompt engineering, human oversight, and client confidentiality safeguards — operable at both the individual and institutional level. The article ultimately argues that the responsible integration of Generative AI into legal practice demands clear, binding, and profession-specific regulatory standards.

Keywords: Generative Artificial Intelligence, Legal Practice, Regulatory Vacuum, Regulatory Framework, Hallucinated Legal References, Mitigation Measures

I. INTRODUCTION

The 21st century has witnessed rapid technological transformation, with Generative Artificial Intelligence (hereinafter 'Generative AI') emerging as one of its most significant developments. Once considered a domain requiring exclusively human judgment, the legal profession is increasingly encountering the integration of AI-driven tools across various aspects of practice. This integration progressed faster than the regulatory framework has been able to respond. Existing guidelines remain suggestive rather than conclusive, leaving a significant regulatory gap — this fast-paced adoption with fragile oversight has opened the door to a host of professional and judicial challenges that courts worldwide are only beginning to reckon with.

Generative AI refers to computational systems capable of producing new and meaningful content such as text, images, audio, or code by learning patterns from vast datasets on which they are trained.¹ Unlike traditional AI systems that primarily analyse or classify information, Generative AI models are designed to create novel outputs that resemble the characteristics of their training data. A major category of such systems is Large Language Models (LLMs), which are deep learning models trained on extensive corpora to generate human-like responses from prompts.² These models operate by predicting the most statistically probable sequence of words based on learned linguistic patterns. A commonly recognised limitation of these systems is the phenomenon of “AI hallucination,” where the model produces inaccurate responses based on noise or bias in training data rather than accurately reflecting the actual input provided.³ For legal practitioners, an unverified hallucination is not merely a technical error but a professional hazard, capable of misleading courts and exposing advocates to serious liability.

While AI shows promising capabilities in the field of legal practice, its increasing integration has also exposed several critical concerns. AI-powered tools are now being used in multiple areas of legal work, ranging from legal research and drafting to case analysis, preparation of legal strategies, and even aspects of client interaction, often saving considerable time and money. However, this advancement is accompanied by its unique drawbacks. The concerns that emerge are of two kinds — first, the problems associated with the outputs of AI-generated content, such as fabricated citations, false quotations, inaccurate content, and poor drafting

1 S. Feuerriegel et al., *Generative AI*, 66 BUS INF SYST ENG 111 (2024).

2 Bassel Almarie et al., *The Use of Large Language Models in Science: Opportunities and Challenges*, 9 PRINC. PRAC. CLIN. RSCH. 1 (2023).

3 Brydon T. Wang, *Prompts and Large Language Models: A New Tool for Drafting, Reviewing and Interpreting Contracts?*, 6 L. TECH. & HUMANS 88, 95 (2024).

lacking nuanced legal reasoning; and second, the deeper structural problems with how AI operates — including accountability, transparency, bias and discrimination in its generated content, and data privacy.

This article, therefore, examines the advancements and associated risks of the escalating use of unverified Generative AI content in legal practice, while also analysing real-life cases from Indian and foreign courts to understand how these issues have manifested in actual proceedings. It further maps the existing regulatory landscape by examining both international and Indian guidelines governing the use of AI within the legal profession, identifies the regulatory vacuum that enables unchecked AI reliance, and argues for the development of comprehensive and binding rules to ensure its responsible and accountable use, while also proposing mitigation steps that legal practitioners can adopt for the ethical and verified use of AI.

II. GENERATIVE AI IN LEGAL PRACTICE : PROMISE AND PERIL

A. Legal Research — Efficiency Gains and the Hallucination Problem

Generative AI LLMs are trained on vast datasets, thus, can make the process of legal research significantly quicker. Instead of going through a plethora of cases and manually filtering out useful ones, these tools can identify and retrieve pertinent cases with a single prompt outlining the required legal principle. Tasks that previously took hours can now be completed within minutes saving considerable time and effort.

For instance, by entering a simple prompt such as “List all relevant High Court or Supreme Court cases regarding violation of natural justice,” or a more specific query such as “List all cases where a court order was set aside because a party was not given an opportunity to be heard,” these tools can generate a list of potentially applicable precedents almost instantaneously.

Yet an important questions remain: **to what extent can these findings be relied upon?** Is conducting legal research truly as simple as entering a single prompt into an AI tool?

The first and most visible category of concerns relates to the outputs that these tools generate. As AI-based LLM tools such as ChatGPT, Gemini, and DeepSeek are increasingly used for legal research, a major issue that emerges is the risk of fabricated cases and fake quotations.

These cases when submitted to courts without proper verification, can potentially affect the course of proceedings, and such reliance on non-existent or inaccurate precedents can drive the court proceedings in a different direction. Moreover, it can lead to erroneous decisions that are based on some fictitious material. Although these models are capable of citing cases, the accuracy and reliability of such citations often remain questionable. While generating responses from extensive training data, AI models are not always capable of distinguishing between credible and non-credible sources, which can result **in inaccurate responses**. Unchecked reliance on such precedents can therefore mislead the court. Currently, courts in India as well as in other jurisdictions are facing similar concerns, **which not only delays** court procedures as additional time must be spent verifying the authenticity of such citations, but may also create a cloud of suspicion even over otherwise legitimate case references.

These concerns highlight the growing risks associated with unverified AI-assisted legal research, an issue that has already surfaced in several real-world judicial proceedings.

B. Drafting — Speed and the Accuracy Deficit

A number of tools such as Spellbook, Draftwise, Volody, and Luminance, now facilitate AI-assisted drafting, in addition to general tools such as ChatGPT and Claude AI. These platforms can generate AI-structured drafts with remarkable speed using prompts, significantly reducing time and costs compared to traditional methods that once took days. For example, Robin AI, a legal AI tool, claims that it is able to draft any common agreement, in an average time of 2m 37s.⁴

However, while drafting automation offers efficiency, it also introduces significant challenges. The risk of inaccurate content due to hallucinations persists not just in legal research but also in drafting. Despite their extensive training, AI tools cannot replicate the nuanced thinking and reasoning of lawyers. They may include relevant laws and details, but struggle to predict complex legal issues and unforeseeable problems that a human practitioner could identify. Relying solely on AI-drafted contracts, agreements, or notices may therefore lead to loopholes and complexities beyond AI's current capabilities.

Another problem that arises in the use of AI for drafting petitions relates to the **structure of legal documents**. There is a set format in which petitions are drafted and submitted before

4 Id. at 94.

courts to streamline the process and avoid unnecessary complications. Although AI tools can mimic the legalese of legal documents, each matter differs from the others. At times, changes to the draft are required for proper applicability and clearer understanding of the complex facts and rules involved in a case. AI tools may misidentify correct legal provisions and incorrectly refer to rules that do not apply, including citing an incorrect rule or attributing commentary to provisions where no such commentary exists.⁵ These structural and accuracy deficits highlight that AI-assisted drafting, while efficient, demands rigorous human oversight before any draft reaches a court or client.

C. Case Analysis and Legal Strategy — Capabilities and Accountability & Transparency Gaps

The use of AI tools in the legal profession has gradually extended beyond drafting and research to include case analysis and the formulation of legal strategies. By uploading a case file or contract, within minutes the tool can mark the important facts and information. Reviewing and analysing contracts, agreements, and files has become a more efficient and simpler process through these tools. For instance, Robin AI claims that its tool enables 80% faster review of agreements.⁶ AI tools can also help in generating legal strategies for a particular case based on the contents of the file, thus streamlining the entire process from research and drafting to strategy formulation. However, this integration has also brought forth some serious concerns — particularly around bias, transparency, and accountability.

One of the major concerns is the issue of **inherent bias and discrimination within AI systems**. These models are trained on specific datasets, and the information they generate may reflect biases present in the data on which they were trained. The datasets used for training may themselves be prejudiced, creating the possibility that AI systems may produce outputs that overlook or disadvantage certain groups. Non-diverse datasets may generate outputs that fail to consider different legal solutions for different populations or jurisdictions.⁷ This poses a substantial risk, particularly for minority groups and across diverse jurisdictions, as reliance on such AI-generated strategies drawn from limited and potentially non-representative databases

5 Herbert B. Dixon Jr., *My “Hallucinating” Experience with ChatGPT*, 62 JUDGES’ J. 37, 38 (2023).

6 Robin AI, <https://robinai.com> (last visited Mar. 16, 2026).

7 Keerthana R. & Parthasarathy S., *Ethical Challenges in AI-Powered Legal Drafting: Balancing Innovation and Responsibility*, in LEGAL ENGLISH AND THE LANGUAGE OF LAW: AN ANTHOLOGY 39, 42 (Jayini Bhaumik & Radha Ranjan eds., The Lawway with Lawyers Journal 2024).

may produce flawed legal advice and expose the practitioner to unwarranted liability.

Another problem that arises is the "black box" nature of AI systems, meaning only the input and output data are visible while what occurs in between remains unknown,⁸ making it impossible to understand how the algorithm reaches its conclusions.⁹ This opacity poses significant challenges particularly in legal and administrative contexts, where the reasoning behind a decision must be transparent and accountable. When suggesting a solution as part of a legal strategy in a contract dispute, it may be unclear why the system chose that specific strategy, what reasoning it applied, and on what data it relied to generate the output. Since such systems are trained on vast and complex datasets, identifying the precise information used in generating a particular outcome becomes impractical. This lack of transparency becomes particularly problematic when AI tools are used for legal analysis or strategic decision-making.

To illustrate the transparency issue associated with the "black box" nature of AI, experimental studies involving AI systems such as ChatGPT have attempted to evaluate the reasoning capabilities of these tools in simple legal scenarios. In one such experiment, the AI was asked to analyse a basic loan dispute involving an oral agreement and a witness statement acknowledging the existence of the loan. The system was able to summarise the facts and suggest that the claimant could pursue a civil action based on the evidentiary value of the witness testimony. While the output demonstrated that the system could identify relevant legal issues and provide a plausible legal assessment, the reasoning process through which the AI arrived at this conclusion remained unclear and inaccessible.¹⁰ Consequently, legal practitioners may find it difficult to evaluate the reliability of AI-generated legal strategies or recommendations when the underlying reasoning process often remains obscure to them.

This ambiguity and lack of transparency associated with AI systems also gives rise to an important issue of **accountability**. When an AI-generated legal strategy is used by practitioners and the outcome proves to be flawed, the question arises as to who should bear responsibility. It becomes unclear whether liability should fall on the developer of the AI system, the legal practitioner relying on the output, or the organisation deploying the technology. The

8 Martin Ebers, *Regulating Explainable AI in the European Union: An Overview of the Current Legal Framework(s)*, in NORDIC YEARBOOK OF LAW AND INFORMATICS 2020: LAW IN THE ERA OF ARTIFICIAL INTELLIGENCE 105, 107 (Liane Colonna & Stanley Greenstein eds., 2021).

9 Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147, 1159–1160 (2017).

10 Inesa Stolper, *Towards Automated Decision-Making at Court: The Use of Artificial Intelligence for Drafting and Rendering Court Decisions*, 130 TEISÉ 153, 159 (2024).

overlapping nature of these responsibilities complicates the issue of accountability and raises significant concerns regarding **professional liability and ethical oversight in legal practice**.¹¹

D. Client Interaction — Accessibility and Data Privacy Concerns

Apart from transforming how lawyers research and draft, **Generative AI has also begun reshaping how they communicate with clients**, making client interaction one of its most practically significant, yet relatively underexamined, applications in legal practice. **Many law firms are increasingly adopting AI-powered legal intake bots - automated tools designed to gather client information and facilitate initial interaction**. This increases response time efficiency and **results in higher conversion rates alongside providing 24/7 availability, multilingual support, and instantaneous real-time responses, which often surpass the limitations of human support**. Industry reported data that by using AI chat there is a 30% increase in conversion rates from visitors to leads and 60% of overall AI-generated chats turn into potential leads,¹² addressing a critical gap where nearly half of all law firms remain unreachable by potential clients.¹³ The bots can prioritise high-value prospects for follow-ups and automatically schedule meetings directly into lawyers' calendar.

However, with these transformations, significant concerns regarding data privacy remain. Clients may not place the same level of trust in sharing their personal information with a chatbot as they do with a human professional. In legal practice, client information shared with lawyers is protected by attorney-client privilege, but when the same information is shared with AI chatbots, this privilege may not apply, creating unanswered questions about accountability if such confidential information is leaked or compromised. In today's world, where data has become one of the most valuable digital assets, the susceptibility of such systems to breaches further exacerbates these concerns. Beyond data privacy concerns, these AI tools, although fluent in answering basic queries and frequently asked questions, may not always be able to account for the nuances and complexities of the law, as they may lack the same level of understanding and judgment as human lawyers when interpreting legal principles,¹⁴ and are

11 Keerthana R. & Parthasarathy S., *supra* note 7, at 44.

12 Scorpion, *Leveraging AI to Grow Your Legal Practice*, LAW TECHNOLOGY TODAY (Nov. 13, 2024), https://www.americanbar.org/groups/law_practice/resources/law-technology-today/2024/leveraging-ai-to-grow-your-legal-practice/ (last visited Mar. 16, 2026).

13 CLIO, *AI-Powered Legal Practices Surge: Clio's Latest Legal Trends Report Reveals Major Shift* (2024), <https://www.clio.com/about/press/clio-latest-legal-trends-report/> (last visited Mar. 16, 2026).

14 Andrew Perlman, *The Implications of ChatGPT for Legal Services and Society*, THE PRACTICE, Mar./Apr. 2023, <https://clp.law.harvard.edu/knowledge-hub/magazine/issues/generative-ai-in-the-legal-profession/the->

often unable to reliably navigate jurisdiction-specific legal questions.¹⁵ **Due to the unresolved gaps surrounding data breaches, accountability, and client trust, these chatbots, while showing considerable promise, cannot fully substitute human lawyers in the delivery of legal services.**

III. ILLUSTRATIVE CASE LAW ON AI-GENERATED LEGAL HALLUCINATIONS

Mata v. Avianca, Inc.

A widely cited example of AI-generated hallucination in legal filings arose in *Mata v. Avianca, Inc.*, where attorneys submitted a legal brief containing several fictitious judicial precedents. The counsel relied on legal research generated through ChatGPT, which produced case citations and quotations that appeared to originate from authoritative sources such as the Federal Reporter and Westlaw. However, many of these authorities were entirely fabricated. The lawyers failed to verify the accuracy of the AI-generated citations before submitting them to the court. Upon discovering the non-existent cases, the United States District Court for the Southern District of New York held the attorneys responsible for submitting false authorities and imposed sanctions of US\$5,000 on the lawyers and their law firm was also held jointly liable for the violations committed by their employees.¹⁶ The incident has since become a prominent illustration that reliance on unverified AI-generated content constitutes a serious professional violation, attracting liability not just for the individual lawyer but for the firm as a whole.

Gauthier v. Goodyear Tire & Rubber Co.

A similar issue arose in *Gauthier v. Goodyear Tire & Rubber Co.*, a wrongful termination dispute in which the defendant sought summary judgment. In opposing the motion, the plaintiff's counsel cited two non-existent judicial decisions: "Roca v. King's Creek Plantation, LLC, 500 Fed.Appx. 273, 276 (5th Cir. 2012) (unpublished)" and "Beets v. Texas Instruments, Inc., No. 94-10034, 1994 WL 714026, at *3 (5th Cir. Dec. 16, 1994) (unpublished)" and included fabricated quotations purportedly drawn from other cases. The defendant's counsel

implications-of-chatgpt-for-legal-services-and-society/ (last visited Mar. 16, 2026).

15 Harsh Gour, *The Many Difficulties with Legal AI Chatbots*, THE LEAFLET (Sept. 16, 2025), <https://theleaflet.in/digital-rights/law-and-technology/the-many-difficulties-with-legal-ai-chatbots> (last visited Mar. 16, 2026).

16 *Mata v. Avianca, Inc.*, 678 F.Supp.3d 443 (2023).

subsequently brought these inaccuracies to the court's attention, highlighting that the authorities relied upon by the plaintiff's attorney were inauthentic. Upon inquiry, it was found that the plaintiff's attorney had used "Claude," a generative AI tool, for drafting purposes and failed to verify these AI-generated citations before submitting them before the court. Upon discovering these non-existent cases and quotations, the United States District Court sanctioned a penalty of \$2,000 on the counsel.¹⁷ This incident demonstrates the risks associated with unverified reliance on Generative AI tools in legal research and drafting, highlighting the continuing duty of lawyers to verify the accuracy of authorities cited before presenting them before a court.

Deepak s/o Shivkumar Bahry v. Heart & Soul Entertainment Ltd.

A corresponding issue was highlighted by the **Bombay High Court**, where the court observed that written submissions filed by a party appeared to have been generated using an AI tool such as ChatGPT or similar platforms. The court noted several "give-away features" within the document, including the presence of green-box tick marks, bullet-point formatting, and repetitive submissions, which suggested the use of AI-generated content. The submissions also referred to an alleged case, "Jyoti w/o Dinesh Tulsiani v. Elegant Associates," without providing any citation or copy of the judgment. Despite attempts by the court and its law clerks, the case could not be located, resulting in a waste of judicial time. While acknowledging that AI tools may be used for research assistance, the court emphasised that advocates bear a responsibility to verify the authenticity and relevance of AI-generated material before submitting it in court, warning that failure to do so may attract costs or disciplinary consequences. The court consequently imposed a **penalty of ₹50,000** on the counsel.¹⁸ The episode illustrates the emerging challenges posed by generative AI in the Indian legal system and highlights the responsibility of legal practitioners to ensure the accuracy and authenticity of materials relied upon in court proceedings.

Supreme Court Observations on AI-Generated Legal Submissions

Concerns regarding the use of AI-generated legal content have also been acknowledged within the **Supreme Court of India**. During the hearing of a Public Interest Litigation before a Bench

¹⁷ Gauthier v. Goodyear Tire & Rubber Co., No. 1:23-CV-281, 2024 U.S. Dist. LEXIS 214029 (E.D. Tex. Nov. 25, 2024), <https://www.casemine.com/judgement/us/6746a063e8c044126187cfea> (last visited Mar. 16, 2026).

¹⁸ Deepak v. Heart & Soul Entertainment Ltd., 2026:BHC:AS:626.

comprising Chief Justice Surya Kant and Justices B.V. Nagarathna and Joymalya Bagchi, the judges expressed alarm over the increasing use of AI tools for drafting legal submissions. The Chief Justice observed that courts had recently encountered instances where lawyers cited precedents that did not exist, recalling a similar situation previously noted by Justice Dipankar Datta where a series of non-existent judgments had been cited. Justice B.V. Nagarathna referred to an example in which a case titled “Mercy v. Mankind” was cited before the court despite no such decision existing. She further noted that in some instances, even when the cited judgment is genuine, fabricated quotations are attributed to it, requiring judges to verify the authenticity of the quoted material.¹⁹ Justice Joymalya Bagchi also expressed concern that the quality of legal drafting has declined in recent years, observing that many pleadings have become compilations of extracted passages rather than carefully reasoned submissions.²⁰ The Supreme Court's unprompted intervention on this issue, during an otherwise unrelated matter, is itself an indicator of the seriousness of the growing menace of unverified AI-generated submissions.

These cases are not isolated incidents but reflect a growing pattern being encountered in courts across India and worldwide. With the rapid integration of AI into legal practice, fabricated citations and false quotations have become an increasing concern, opening the door to sanctions against lawyers, waste of judicial time, and a broader threat to the institutional integrity of the legal profession.

IV. THE REGULATORY AND ETHICAL VACUUM

The growing use of Generative AI in legal practice has prompted regulatory responses across multiple jurisdictions — from judicial standing orders in the United States to several bar association and judicial body guidelines in several countries and institutional guidelines in India. However, a closer examination of these responses reveals a fundamental inadequacy: existing frameworks are fragmented across jurisdictions, inconsistent in their requirements, and largely advisory in nature, leaving practitioners without clear, enforceable standards governing

19 LIVE MINT, *Supreme Court flags ‘alarming’ use of AI in drafting petitions, cites use of non-existent ‘Mercy vs Mankind’ case*, <https://www.livemint.com/news/india/supreme-court-flags-alarming-use-of-ai-in-drafting-petitions-cites-use-of-non-existent-mercy-vs-mankind-case-11771334601294.html> (last visited Mar. 16, 2026).

20 Vanshika Shahi, *“Mercy v. Mankind”? Supreme Court Shocked by Fake AI Citations in Petitions*, LAWFUL TALKS (Feb. 18, 2026), <https://www.lawfultalks.net/news/mercy-v-mankind-supreme-court-shocked-by-fake-ai-citations-in-petitions>.

their use of AI-generated content in legal proceedings.

A. Judicial Standing Orders: A Fragmented Framework

In the United States, individual federal judges have responded to the rise of Generative AI in legal practice by issuing standing orders primarily aimed at regulating the use of AI-generated content in court filings and proceedings. These orders, while limited in jurisdictional reach, broadly reflect judicial concern over the accuracy of court submissions and the professional responsibility of legal practitioners who employ such tools.

1. Judge Stanley Blumenfeld

In a standing order by Judge Stanley Blumenfeld, he highlighted the requirement to disclose AI-generated content and certify that the filer has verified the source material and content.²¹ This reflects a basic but critical principle that the use of AI does not transfer or diminish the lawyer's personal responsibility for the accuracy of court submissions. This becomes particularly significant in guarding against hallucinated citations and fabricated references from reaching courts, thereby preserving the integrity of judicial proceedings and the responsible use of judicial resources.

2. Judge Peter H. Kang

Additionally, in a standing order by Judge Peter H. Kang of the United States District Court, he acknowledged the recent developments relating to generative AI in the legal field. While recognising the potential utility of such tools, the court emphasised that lawyers and parties remain responsible for ensuring the accuracy and reliability of all materials submitted before the court. Building upon the disclosure framework, the order further clarifies that failure to verify the authenticity and accuracy of AI-generated content may attract sanctions under applicable procedural rules. It also requires disclosure when AI-generated materials are used as evidence and cautions lawyers about confidentiality risks associated with third-party AI platforms.²²

21 *Standing Order for Civil Cases Assigned to Judge Stanley Blumenfeld, Jr.*, § 5(c), United States District Court, Central District of California (updated Mar. 1, 2024).

22 *Standing Order for Civil Cases Before Magistrate Judge Peter H. Kang*, § VII(c), United States District Court, Northern District of California, San Francisco Division (effective July 16, 2025).

3. Judge Araceli Martínez-Olguín

Going beyond AI outputs to regulate the input, a standing order by District Judge Araceli Martínez-Olguín included the counsel's responsibility to maintain records of all prompts or inquiries submitted to any generative AI tool, apart from personally verifying the accuracy of the content.²³ This order addresses one of the fundamental problems associated with transparency in AI-generated content — the process by which AI reaches its output remains opaque to the user. Preserving such prompts may therefore enable a subsequent review of the lawyer's reasoning and approach while employing AI tools, introducing a degree of process-level accountability.

Although these judicial orders represent an important response to the escalating use of Generative AI in courtrooms, they remain largely limited in scope and jurisdiction. These orders, passed by individual judges, are confined to their respective courtrooms and therefore lack uniform applicability across courts. This highlights the absence of comprehensive and uniform guidelines that can be applied across the judiciary. Additionally, these orders primarily emphasise disclosure obligations and the duty to verify AI-generated content, rather than establishing broader standards governing the responsible use of AI in legal practice. Consequently, significant regulatory gaps remain with respect to professional accountability, verification protocols, and broader ethical obligations relating to the use of AI in the legal field.

B. Professional Body Guidelines

While judicial standing orders represent an important institutional response to the growing use of generative AI in litigation, professional bodies governing the legal profession have also issued formal guidelines to address the responsible use of generative AI in legal practice. Several bar associations, law societies, and judicial bodies across jurisdictions have issued formal guidelines in this regard. These include bodies from the United States, the United Kingdom, and New Zealand, whose approaches are examined in this section.

1. American Bar Association

The American Bar Association (hereinafter 'ABA') has addressed the use of Generative AI in

²³ *Standing Order for Civil Cases Before District Judge Araceli Martínez-Olguín*, § H(4), United States District Court for the Northern District of California (revised Nov. 14, 2025)

legal practice, emphasising the obligation of lawyers to provide competent representation to their clients, including an understanding of the benefits and risks associated with the technologies used in delivering legal services. The ABA recognises the risks of inaccurate content and biases that may lead to hallucinations in generative AI tools. The guidelines emphasise that while generative AI tools may assist lawyers in tasks such as generating legal analysis or drafting documents, lawyers bear full responsibility for all legal work and cannot delegate their professional judgment entirely to AI tools. Consistent with Model Rule 1.1, lawyers must exercise independent professional judgment and undertake appropriate verification of AI-generated outputs depending on the nature of the task involved.²⁴ The guidance further highlights the duty of confidentiality under Model Rule 1.6, requiring lawyers to safeguard all information relating to the representation of a client and to take reasonable measures to prevent unauthorized disclosure when using AI tools.²⁵ Additionally, the guidelines stress that managerial and supervisory lawyers must establish appropriate policies and training within law firms to ensure that both lawyers and non-lawyer staff use generative AI tools in compliance with their professional and ethical obligations.²⁶ Taken together, the ABA's guidance reflects a broader principle that the advent of generative AI does not diminish existing professional obligations but rather demands that lawyers actively uphold those duties within an evolving technological landscape.

2. United Kingdom Bar Council

Subsequently, the United Kingdom Bar Council, while highlighting the risks associated with generative AI and large language models, notes that irresponsible use of such technologies may lead to serious consequences, including claims for professional negligence, breach of confidence, defamation, and other legal liabilities. The guidelines expressly warn that sharing legally privileged or confidential client information with third-party AI platforms places such information at risk of exposure, as inputs may be used to train the system and subsequently shared with other users — constituting a breach of Core Duty 6 and rule rC15.5 of the Code of Conduct, undermining the trust fundamental to the attorney-client relationship, and potentially attracting disciplinary liability. It also highlights concerns regarding transparency arising from the 'black box' nature of AI platforms. The guidance further cautions that barristers must

24 American Bar Association Standing Committee on Ethics and Professional Responsibility, *Formal Opinion 512: Generative Artificial Intelligence Tools*, § II(A) (July 29, 2024).

25 *Id.* § II(B)

26 *Id.* § II(E)

critically assess AI-generated content for potential intellectual property violations, as LLMs are trained on vast datasets, raising a serious risk that outputs may reproduce content in violation of third-party copyright.²⁷ The UK Bar Council's guidance is primarily risk-oriented in nature — identifying risks and cautioning barristers to exercise due diligence, while stopping short of establishing uniform and enforceable standards governing the responsible use of generative AI in legal practice.

3. Courts of New Zealand

In a comparable vein, the Courts of New Zealand have also issued guidelines addressing the use of generative AI in court and tribunal proceedings. The guidelines emphasise that lawyers, as officers of the court, must not mislead the court and must take reasonable steps to ensure the accuracy of all information, including legal citations, presented before it. Lawyers remain responsible for verifying the accuracy of any information generated by AI tools before relying on it in proceedings. The guidance further cautions against entering private, confidential, suppressed, or legally privileged information into third-party AI platforms and emphasises the need to safeguard client confidentiality. While the use of generative AI tools is not prohibited, the guidelines clarify that such use must remain consistent with lawyers' professional obligations, and disclosure of AI use is not generally required unless specifically requested by the court or tribunal.²⁸ Notably, the New Zealand guidelines place the onus squarely on lawyers to verify the accuracy of AI-generated content, reflecting a self-regulatory approach that prioritises professional responsibility over prescriptive restriction.

These guidelines, issued by bar associations and judicial bodies, represent an important step towards addressing the ethical implications of generative AI in legal practice, but they do not establish clear and uniform regulatory standards governing the use of AI by legal practitioners. In particular, issues relating to the extent of verification required, disclosure of AI use to clients, professional accountability for AI-generated inaccuracies, and the protection of confidential legal information when interacting with third-party AI platforms remain insufficiently addressed. Consequently, despite the growing recognition of the risks associated with Generative AI in legal practice, a comprehensive and enforceable regulatory framework

27 The Bar Council Information Technology Panel, *Considerations when Using ChatGPT and Generative Artificial Intelligence Software Based on Large Language Models* (Jan. 30, 2024).

28 Courts of New Zealand, *Guidelines for Use of Generative Artificial Intelligence in Courts and Tribunals: Lawyers* 3–4 (Dec. 7, 2023).

governing its responsible use within the legal profession continues to be largely absent.

C. The Indian Position

The fragmentation and inadequacy that characterise the global regulatory response to AI in legal practice is equally if not more acutely reflected in the Indian context, where the integration of Generative AI into legal proceedings has outpaced the development of any comprehensive institutional framework to govern it. In India, certain guidance relevant to the use of AI in legal practice is issued by the Supreme Court of India, along with existing professional conduct rules governing advocates.

1. Report of the Supreme Court of India

A report was published by the Centre for Research and Planning of the Supreme Court of India, in November 2025, addressing the growing use of AI in Indian Judiciary and reflecting upon the drawbacks that come with this advancement. The report further suggests guidelines for lawyers for ethical use of AI in legal practices. The recommendations emphasise independent verification of all AI-generated citation and advise counsels to refrain from delegating legal search and strategy to the AI systems. It also advises lawyers to avoid revealing any personal and confidential information of the clients to the AI tools and conduct regular supervision and training for awareness regarding the capabilities and drawbacks of AI tools.²⁹

Despite being issued for lawyers, these guidelines are merely suggestive in nature. Although these instructions represent an initial step towards addressing AI-related problems in the legal field, the growing use of AI demands more robust and comprehensive regulations to effectively deal with such problems.

2. Bar Council of India Rules

Beyond judicial guidelines, the Bar Council of India Rules also contain provisions relevant to the use of AI in legal practice, albeit without explicit reference to it. These rules illustrate the professional conduct of the lawyers. Under Chapter II of Part VI of the Bar Council of India Rules, Rule 3 asserts that an advocate shall not influence the decision of a court by any illegal

²⁹ Centre for Research and Planning, Supreme Court of India, *White Paper on Artificial Intelligence and Judiciary* 79–80 (Nov. 2025).

or improper means.³⁰ While Rule 3 does not explicitly address the use of AI, it may reasonably be interpreted to include the submission of AI-generated fabricated citations or hallucinated quotations before a court, as such conduct could constitute an attempt to influence judicial decisions through false information, a clear violation of an advocate's duty to the court.

Under Rule 14 of the same chapter, it requires a lawyer to make full and frank disclosure to the client regarding any interest or circumstance that may affect the client's judgment in either engaging with the lawyer or continuing the engagement.³¹ Although this provision too was not formulated with AI in mind, it may reasonably be interpreted to mean that a lawyer should inform the client of any factor that could materially affect the manner in which legal services are being delivered. In this context, where a lawyer substantially relies on generative AI tools for research, drafting, or strategy formulation, the client may have a legitimate interest in being aware of such reliance, particularly given that the client engages and remunerates the lawyer precisely for the exercise of independent professional expertise.

Thus, the rules under the Bar Council of India regulations broadly aim to ensure the proper conduct of lawyers in court proceedings and emphasise their responsibilities towards clients, the court, and the administration of justice. However, these rules remain silent on the use of AI, leaving a significant gap in the regulatory framework governing advocates. In the present scenario, where AI tools are rapidly being integrated into the legal field, there is an increasing need for explicit guidelines and regulatory standards governing the use of such technologies by legal professionals. Clear frameworks would help ensure that the adoption of AI in legal practice remains consistent with the professional standards, ethical duties, and expected conduct of advocates.

India is yet to establish a comprehensive regulatory framework governing AI, though recent initiatives such as the Digital Personal Data Protection Act represent nascent steps in this direction. At the global level, the European Union has taken a significant step by enacting the EU AI Act, one of the first comprehensive regulatory frameworks governing the use of AI, while countries such as the United States, China, and Brazil are also attempting to enact or formulate their own AI regulatory mechanisms to accelerate the development of AI

30 Bar Council of India Rules, Chapter II (Part VI), Rule 3, made under the Advocates Act, Act No. 25, passed by Parliament in 1961 (India).

31 *Id.* r. 14.

technologies while also ensuring their ethical and responsible use.

However, beyond general AI governance frameworks, specific regulations are also required for particular professional fields such as the legal domain, as the integration of AI tools in legal practice raises a distinct set of challenges beyond the general risks associated with AI. The existing guidelines in this regard largely remain suggestive in nature and are neither comprehensive nor uniform in scope and jurisdiction, even within the same country. Judicial standing orders, professional body guidelines, and advisory frameworks issued by courts primarily emphasise disclosure obligations and the duty to verify AI-generated content, but they do not establish uniform standards governing the professional use of such technologies.

In addition, these frameworks do not clearly address issues such as the extent of verification required for AI-generated legal material, the obligation of lawyers to inform clients regarding the use of AI tools in legal work, or the potential risks relating to confidentiality and privileged information when interacting with third-party AI platforms. The absence of clear and binding regulatory standards also creates uncertainty regarding professional accountability and liability in situations where AI-generated inaccuracies, such as fabricated citations or incorrect legal analysis, are submitted before the court. Consequently, while several jurisdictions have begun recognising the implications of AI in legal practice, a comprehensive and enforceable regulatory framework governing the responsible use of AI in the legal profession remains largely absent.

V. STEPS TOWARD RISK MITIGATION

While waiting for a robust framework to regulate the use of AI in Legal profession, certain steps could be adopted by the practitioners to use AI ethically in practice. These steps are not the replacement of any regulatory framework rather a self-induced measures to ensure the moral and responsible use of AI that is required to maintain the integrity of the institution and trust of the people in legal practice. The use of AI must proceed from role limitation and disciplined prompting to verified outputs, structured prompt retention and strict confidentiality safeguards.

A. Individual-level Safeguards

1. Defining Role of AI

At the outset, the role of AI in legal practice must be clearly delimited as a supportive aid and

not as an autonomous or sole operator. With the increasing advancement in the field of AI, it becomes impractical to keep it completely separated from legal practice. AI, despite its inherent limitations, offers several practical and beneficial advantages. It may act as a bridge to reduce the gap between large and small firms, as well as between senior established lawyers and junior advocates, by providing greater efficiency, enhanced access to resources, and additional time to focus on substantive legal work beyond routine tasks.

AI tools must therefore be utilised for purposes of efficiency, such as brainstorming on legal strategy, summarising case files, structuring legal documents, and assisting in legal research. However, matters that inherently require human judgment must not be entrusted to AI, including the formulation of legal strategy, drafting of petitions and contracts, interpretation of law, and final decision-making. AI-generated outputs should inform and assist legal reasoning, but must not replace independent professional judgment.

This distinction is essential to ensure that accountability, ethical responsibility, and professional standards remain firmly vested in the human practitioner. Accordingly, AI must remain a facilitative tool, with all determinative functions retained by the legal professional.

2. Prompt Engineering Discipline

Following role delimitation, AI outputs must be refined through detailed and structured prompts. Avoidance of complete reliance on the very first result it produces is essential. Outputs generated after several methodological prompts elevate the overall quality of results. This process helps to clearly define important elements such as facts, jurisdiction, issues, and the required format. Providing necessary background and breaking down queries into parts further guides coherent and legally relevant outputs.

Through iterative refinement of prompts, improvements can be made in the depth, analysis, and accuracy of the generated responses. In this regard, prompt engineering operates as a critical control mechanism, shaping both the scope and reliability of AI-assisted legal work

3. Human Oversight

Subsequently, the use of AI necessitates human supervision in ensuring the accuracy of AI-generated content. This includes verifying the authenticity of citations and quotations by referring to the original judgments, checking for jurisdictional variations and applicability of

laws, and ensuring the correctness of facts, legal provisions, and terminology in documents drafted by AI.

It further involves assessing whether the reasoning, analysis, and conclusions are legally sound and contextually appropriate. Applying human critical thinking to AI-generated strategies and measures, as well as checking for any potential copyright violations in the generated content, is also essential. Outputs must be modified, corrected, and improved based on professional judgment before use. Thus, human oversight functions as the final safeguard, ensuring the accuracy, reliability, and accountability of AI-assisted legal work.

4. Maintain a Repository of Prompts

Further, maintaining a structured repository of prompts helps in establishing a clear understanding of the intentions of legal practitioners while using AI. Given the concerns of transparency, bias, and accountability in AI-generated responses, the documentation of prompts aids in delineating the extent of liability of the legal professional. Such documentation may also prove useful in anticipating disputes arising from ambiguity or negligence.³² Developing fixed and tested prompts for repetitive tasks can save time and ensure standardised and reliable responses. It also reduces variability in AI-generated outputs.

Additionally, a structured prompt repository functions as an audit mechanism, enabling traceability of how specific AI-generated outputs were derived. This is particularly relevant in contexts of professional accountability and dispute resolution, where the ability to reproduce or justify a particular output becomes crucial. It also facilitates knowledge transfer within legal practice, allowing junior practitioners to rely on tested prompt structures, thereby reducing inconsistencies arising from individual variations in prompt design. Thus, maintaining a structured repository of prompts, guided by clearly defined intentions, ensures consistency, efficiency, and goal-oriented AI-assisted legal work.

5. Client Confidentiality and Data Protection

Finally, the use of AI in legal practice must be governed by strict adherence to client confidentiality and data protection obligations. Information shared by clients with their legal practitioners is ordinarily protected under attorney–client privilege; however, when AI

³² Wang, *supra* note 3, at 96.

assistance is employed, the inadvertent disclosure of such information through prompts or generated responses may raise serious ethical concerns. Inputting sensitive client data into AI tools may lead to unintended storage, processing, or third-party access. Moreover, AI systems, being susceptible to hacking or security vulnerabilities, may further increase the risk of data breaches.

Certain measures can be adopted to mitigate these risks, such as anonymising names and sensitive information, using hypothetical data where possible, and obtaining client consent where AI tools are used in handling their data. Additionally, preferring secure or offline tools in confidentiality-sensitive matters, and understanding how AI systems store, use, or process data, can assist in navigating challenges of data privacy and breach. Thus, safeguarding client confidentiality must remain paramount, with AI use carefully calibrated to prevent any compromise of professional obligations.

B. Institutional (firm-level) Safeguards

Beyond individual-level safeguards, the responsible use of AI in legal practice necessitates the adoption of institutional mechanisms at the firm level to ensure proper utilisation and mitigate potential liabilities.

1. Supervision and Internal Guidelines

Firms must establish clear supervision and internal guidelines governing the responsible use of AI by lawyers and non-staff. Such frameworks should define permissible uses, set boundaries, and ensure accountability in AI-assisted legal work.

2. Secure Infrastructure and Tool Evaluation

Firms should prioritise the use of secure, enterprise-grade AI tools with controlled data handling. AI systems must be carefully evaluated before adoption, and internal safeguards should be implemented to prevent data leakage, misuse, or unauthorised access.

3. Institutionalised Human Oversight Mechanisms

Firms must mandate review mechanisms before AI-generated outputs are used externally, with clearly defined levels of review depending on the sensitivity and nature of the task. This ensures

reliability and reduces the risk of errors.

4. Auditing of AI-Based Client Interaction Tools

Where firms employ AI-based tools for client interaction, such as chatbots, regular auditing of such systems is essential to assess bias, accuracy, and operational capabilities, ensuring alignment with legal and ethical standards.

5. Training and Awareness Mechanisms

Regular workshops and training sessions should be conducted to build awareness regarding the ethical and effective use of AI. Such initiatives should address both the advantages and limitations of AI, ensuring informed usage by legal professionals.

6. Maintenance of a Shared Prompt Repository

Maintaining a repository of prompts within firms can function as a shared intellectual resource, enhancing responses and ensuring standardisation and reliability of AI-generated content through tested and refined prompts.

7. Client Transparency and Disclosure Policies

Firms should develop clear policies on whether and how clients are informed of AI use in their matters, ensuring compliance with ethical obligations and maintaining transparency in professional engagement.

8. Liability Allocation Framework

Firms must define how liability is to be allocated in cases where AI-assisted work leads to errors, thereby ensuring clarity in responsibility and mitigating potential legal and professional risks.

VI. CONCLUSION

Generative AI LLMs can be a game-changer, bridging knowledge and resource gaps among law practitioners and firms. Small firms, previously constrained by limited resources, can now compete with larger firms with abundant staff and infrastructure. With advancements in these

models, tasks such as research, drafting, and analysis can be performed efficiently, saving time and resources that can instead be invested in constructing arguments, handling more cases, and enhancing operational efficiency. However, alongside these advantages, unregulated AI use poses serious ethical and professional risks, including threats to data privacy, accountability, liability, bias, and accuracy. AI integration is inevitable, but it must remain structured and controlled to maximise benefits without undermining the integrity of the legal system.

AI adoption will transform routine legal work, accompanied by evolving roles for legal professionals and increased potential for efficiency, access to justice, and standardisation. Yet, AI can assist but cannot replace legal reasoning; critical human oversight remains central and cannot be delegated. The balance lies in determining the extent and manner of AI use while preserving duties that fundamentally remain with the lawyer. The future of legal practice will increasingly rely on AI-assisted tools, yet human judgment, ethical oversight, and professional responsibility must remain at the core. Ultimately, AI may assist legal practice, but it cannot replace the centrality of human judgment, accountability, and professional responsibility, ensuring that technology augments rather than governs the law. In the absence of clear and enforceable regulatory standards, the unchecked integration of AI into legal practice risks undermining not only professional accountability but also the institutional integrity of the justice system. The need for a coherent and binding regulatory framework is therefore no longer prospective, but immediate.