
THE UNREGULATED LAWYER: USE OF ARTIFICIAL INTELLIGENCE BY LEGAL PROFESSIONALS IN INDIA AND THE ABSENCE OF AN ETHICAL FRAMEWORK - A COMPARISON WITH THE AMERICAN BAR ASSOCIATION (ABA), UK BAR COUNCIL, AND SINGAPORE LAW SOCIETY

Aneeka Himayat, Rizvi Law College

ABSTRACT

Generative artificial intelligence has entered Indian legal practice quietly and without rules. Advocates draft pleadings with ChatGPT, summarize briefs through Gemini, and run client documents through cloud-hosted assistants whose data retention policies they have not read. The Bar Council of India has, so far, said nothing. Other jurisdictions have done a great deal. The American Bar Association issued Formal Opinion 512 in July 2024, mapping the ethical duties that govern generative use of artificial intelligence. The Bar Council of England and Wales updated its barristers' guidance in November 2025 in the light of fresh High Court rulings on artificial intelligence hallucinations. Singapore's Ministry of Law, working with the Law Society, launched a comprehensive Guide on the use of generative AI in the legal sector on 6th March 2026. India's silence is not benign. This paper argues that the BCI's inaction has produced a regulatory vacuum that exposes clients to harm, distorts billing, and increases the risk of hallucinated authorities being placed before the courts. Drawing on the American Bar Association (ABA), UK Bar Council, and Singapore frameworks, the paper proposes a five-point reform agenda for the BCI.

Keywords: Artificial Intelligence; Legal Ethics; Bar Council of India; Generative AI; Professional Responsibility; ABA Formal Opinion 512; DPDP Act, 2023; Comparative Regulation.

I. INTRODUCTION

In May 2023, Steven Schwartz, a New York lawyer with three decades at the Bar, filed a brief opposing a motion to dismiss in a routine personal injury suit. The brief cited cases that did not exist: *Varghese v. China Southern Airlines*, *Martinez v. Delta Airlines*, *Zicherman v. Korean Air Lines*. Their captions looked real. Their reasoning sounded plausible. Their citations even bore the conventional reporter formatting. None of them, it turned out, had ever been decided by any court. When opposing counsel pointed this out, Schwartz produced what he claimed were the full texts of the missing decisions. Those too had been generated by ChatGPT, which the lawyer had asked to “double-check” its earlier output. Judge P. Kevin Castel sanctioned Schwartz and his colleague Peter LoDuca jointly under Rule 11, imposing a fine of US\$ 5,000 and ordering them to write to each of the real judges whose names had been falsely affixed to the fictitious opinions.¹

Mata v. Avianca has since become a fixture in every conversation about lawyers and generative AI. But the more interesting question is not what happened to Schwartz. It is what would have happened to him in India.

There is no straightforward answer. The Bar Council of India has issued no opinion, no circular, and no practice direction on the use of generative artificial intelligence by advocates. The Advocates Act, 1961, drafted in an era when “research” meant turning pages, contains no obvious hook on which to hang a sanction for citing a hallucinated authority. The professional-misconduct provisions could conceivably reach the conduct, but no advocate has yet been disciplined under them for artificial intelligence generated work, and the BCI’s Standards of Professional Conduct and Etiquette are entirely silent on the subject.

The silence is not for want of facts. A bench of the Delhi High Court in September 2025 allowed a group of homebuyers to withdraw their petition after it emerged that several of the case citations and quoted paragraphs in the pleading had been hallucinated by ChatGPT — including invented passages attributed to landmark constitutional judgments.² The Bombay High Court, in January 2026, imposed costs of ₹50,000 on a litigant whose written submissions

¹*Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443, 448–49 (S.D.N.Y. 2023). The sanction order, authored by Judge P. Kevin Castel, runs to 34 pages and imposes a joint penalty of US\$ 5,000 under Rule 11.

²Tarini Kalra, *Phantom Precedents: The Rise of AI-Generated Case Law in Indian Courts*, Livelaw (Mar. 17, 2026), <https://www.livelaw.in/articles/phantom-precedents-ai-generated-case-law-indian-courts-526665>.

carried the unmistakable green-tick formatting and repetitive phrasing of raw chatbot output.³ In March 2025, the then Chief Justice of India, B.R. Gavai, cautioned that “platforms like ChatGPT have generated fake case citations and fabricated legal facts”.⁴ Earlier still, in *Jaswinder Singh v. State of Punjab*, Chitkara, J. of the Punjab and Haryana High Court took the unusual step of consulting ChatGPT on bail jurisprudence in a case involving an assault with cruelty — a use limited, the order made clear, to providing a broader picture and not as an opinion on the merits, but a use, nonetheless, that no Indian code of professional ethics has anything to say about.⁵

The scale of the underlying problem is not local. By late 2025, an open database maintained by Damien Charlotin of HEC Paris had catalogued nearly 800 documented instances across at least twenty-five jurisdictions in which AI-generated material had found its way into court filings; 128 of those incidents involved lawyers rather than self-represented litigants.⁶ A 2024 study by researchers at the Stanford RegLab tested the leading commercial legal research products and found hallucination rates between 17 and 33 per cent even for tools using retrieval-augmented generation — that is, tools specifically built to mitigate the problem.⁷ General-purpose chatbots performed considerably worse. The Stanford team’s earlier work, in fact, suggested that for legal questions, off-the-shelf large language models hallucinated answers in approximately three out of every four queries.

The temptation, faced with this, is to say that the existing duties of competence and candor to the court are sufficient. They are not — at least not in the form they take in Indian rules. The BCI Rules speak of competence in the language of an earlier century: an advocate is to act with dignity, to maintain the standards of the profession, and to do her best for the client. None of these formulations tells an advocate whether she may type a confidential client document into a foreign-hosted chatbot, whether she must disclose to the court that her brief was drafted with

³*Id.*

⁴Justice B.R. Gavai, Address at the International Conference on Artificial Intelligence and Law (Mar. 2025), reported in MEDIANAMA, *Risks of AI in Indian Courts: Justice Gavai Flags Fake Citations* (Mar. 11, 2025).

⁵*Jaswinder Singh @ Jassi v. State of Punjab*, CRM-M-22496-2022 (P&H HC Mar. 27, 2023) (order) (per Chitkara, J.). For commentary, see *ChatGPT-fication of the Judicial Process*, SCC ONLINE BLOG (Nov. 14, 2023).

⁶Damien Charlotin, *AI Hallucination Cases Database* (Nov. 2025), <https://www.damiencharlotin.com/hallucinations/>. The database, maintained by a researcher at HEC Paris, catalogues nearly 800 documented instances across at least 25 jurisdictions.

⁷Varun Magesh et al., *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools* 9–10 (Stanford Law Sch., RegLab Working Paper, June 6, 2024). An earlier iteration of the same study found that general-purpose large language models hallucinated legal answers in approximately 75 per cent of queries. See Matthew Dahl et al., *Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models* (Stanford Law Sch., Jan. 2024).

the assistance of an AI assistant, or how she ought to bill a client for thirty minutes of prompt engineering that produced the work of ten hours of associate time. The ABA Model Rules, by amendment in 2012, expressly tied competence to “the benefits and risks associated with relevant technology” (Comment [8] to Rule 1.1, at p. 1). The BCI Rules, framed in 1975, were never updated.

The thesis of this paper is plainly stated. The Bar Council of India’s continuing silence on the use of generative artificial intelligence by Indian advocates is not regulatory neutrality. It is regulatory failure. The cost of that failure falls in the first instance on clients — whose confidential information is uploaded to systems whose data-retention policies the advocate has not read, and whose matters may be argued on the basis of fabricated authorities. But it also falls on courts asked to police the integrity of pleadings they cannot verify in real time, and on advocates themselves, who must currently negotiate a fast-moving technology without any guidance from the body that licenses them to practice.

This paper asks three questions. First, what ethical obligations should govern an Indian advocate’s use of generative artificial intelligence tools? Second, how should those obligations be operationalized across the recognized heads of professional duty — competence, confidentiality, disclosure, billing, and liability for error? Third, what regulatory architecture is best suited to fill the existing gap? On each question the paper draws comparatively on three frameworks: the ABA’s Formal Opinion 512 of July 2024, the Bar Council of England and Wales’ updated guidance of November 2025, and the Singapore Ministry of Law’s Guide for Using Generative Artificial Intelligence in the Legal Sector of March 2026, read alongside the Singapore Supreme Court Registrar’s Circular No. 1 of 2024. The conclusion sets out a five-point reform agenda for the Bar Council of India.

II. THE CURRENT STATE OF THE USE OF ARTIFICIAL INTELLIGENCE IN INDIAN LEGAL PRACTICE

The Indian legal market has, in the last three years, absorbed generative artificial intelligence faster than any regulator has been able to write about it. The adoption has happened in three roughly distinct layers: the institutional, the commercial, and the individual.

At the institutional level, the Supreme Court has been measured but visible. In April 2021, the then acting Chief Justice of India, S.A. Bobde, launched the Supreme Court Portal for

Assistance in Court Efficiency (SUPACE), an artificial intelligence portal designed to assist judges in factual extraction and precedent retrieval, not in decision-making.⁸ The Vidhik Anuvaad Software (SUVAS), an AI-based translation tool, has translated more than 31,000 of the Court's 36,000 reported judgments into Indian languages on the e-SCR portal. A Legal Research Analysis Assistant (LegRAA), trained on Supreme Court judgments to draft issue summaries and structured briefs, has been piloted by the National Informatics Centre in conjunction with the Court's Centre for Research and Planning. In December 2025, the Centre for Research and Planning released a White Paper on Artificial Intelligence and the Judiciary in which the Court reiterated, with some emphasis, that "no AI and ML based tools are being used by the Supreme Court of India in the decision making process".⁹ These initiatives illustrate institutional thinking; they do not regulate the Bar.¹⁰

The commercial layer is harder to map. LexisNexis, Manupatra, and SCC Online have each rolled out AI-enabled features ranging from semantic search to summarization. Manupatra's AI assistant, integrated into its case database, is now in active use across litigation chambers. Indian-built tools include Casemine's research products and a growing number of drafting assistants aimed at first- and second-year associates. International tools used in Indian firms include Harvey, Co-Counsel (Thomson Reuters), and Lexis+ AI. These are typically deployed inside enterprise contracts that specify data residency and confidentiality terms; what those terms actually say, however, is rarely circulated to the junior associates who use the tools daily.

The third and most consequential layer is the use of free, general-purpose chatbots by individual lawyers — and, increasingly, by interns and paralegals who do the bulk of first-stage research. ChatGPT, Gemini, and Microsoft Copilot have no terms tailored to legal practice. Their default settings, in several cases, retain user inputs for the purposes of model improvement. Lawyers who paste excerpts from confidential client correspondence into the prompt window of a free chatbot have, without realising it, made a disclosure to a foreign service provider that is, in some respects, irreversible. There is no industry survey of how widespread the practice is among Indian lawyers, but anyone who has spent time in a litigation

⁸Press Information Bureau, Government of India, *Use of AI in Supreme Court Case Management*, PIB Release ID 2113224 (Mar. 2025); see also Supreme Court of India, *Press Note on Launch of SUPACE* (Apr. 2021).

⁹Supreme Court of India, Centre for Research and Planning, *White Paper on Artificial Intelligence and the Judiciary* (Dec. 2025).

¹⁰*Id.* at 4. The Court emphasised that AI tools remain in an experimental stage and are confined to administrative and research-support functions.

chamber in the last two years will have witnessed it.

This three-tier landscape has emerged inside a regulatory regime that, on its face, predates the very category of “generative artificial intelligence”. The Advocates Act, 1961 is the foundational statute. It vests the Bar Council of India with the power to lay down standards of professional conduct and etiquette under Section 49(1)(c).¹¹ Pursuant to that section, Part VI of the BCI Rules, framed in 1975, prescribes those standards under several heads. Rule 49 prohibits an advocate, while in practice, from being a full-time salaried employee. Rule 33 imposes a duty on the advocate not to act for opposing parties in the same matter. Other rules require fearless presentation of the client’s case, the protection of the client’s interest, and the maintenance of confidence. Across all of them, two things are conspicuously absent: any concept of technological competence as a component of professional skill, and any concept of supervision of non-human assistance.

The omissions are not surprising. The BCI Rules were drafted at a time when an advocate’s tools were her clerk, her library, and a manual typewriter. The rules were never amended to keep pace with electronic legal research; they were certainly not amended to keep pace with generative artificial intelligence. The Supreme Court has, on more than one occasion, urged the Bar to update its understanding of professional duty in the light of new technology. In *State of Maharashtra v. Praful Desai*, the Court endorsed video-conferencing for witness examination in criminal trials.¹² In *Justice K.S. Puttaswamy v. Union of India*, the nine-Judge Bench’s articulation of informational privacy as a fundamental right reshaped the entire backdrop against which any professional duty of confidentiality must now be read.¹³ Neither development has produced an amendment to the BCI Rules.

The Digital Personal Data Protection Act, 2023, the implementing Rules under which were notified on 13 November 2025, sharpens the issue from another direction.¹⁴ According to Section 2(i) of the Act, the “Data Fiduciary” is the one who decides on the purpose and means of the processing of personal data, while Section 8 states that a Data Fiduciary has a duty to comply, irrespective of contractual provisions to the contrary. A law firm that an advocate

¹¹The Advocates Act, No. 25 of 1961, § 49(1)(c) (India). The corresponding rules are contained in pt. VI, ch. II of the Bar Council of India Rules, 1975.

¹²*State of Maharashtra v. Praful Desai*, (2003) 4 SCC 601.

¹³*Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2017) 10 SCC 1 (nine-judge bench).

¹⁴The Digital Personal Data Protection Act, No. 22 of 2023 (India); Digital Personal Data Protection Rules, 2025, notified by MeitY (Nov. 13, 2025). Phased compliance is required by May 13, 2027.

belongs to and that uploads a client's personal information, such as confidential medical records in a medical negligence case, onto an artificial-intelligence system that is stored in the cloud is obviously processing that information. That processing is lawful hinges on the Act's consent and purpose limitation aspects. The Bar Council of India has failed to provide any guidelines for its members to adapt their use of AI within the framework of the DPDP regime.

A more positive note has been struck at the policy level. The IndiaAI Mission, run out of the Ministry of Electronics and Information Technology (MeitY), has documented the early use of AI-driven legal aid platforms — chatbots that route queries about tenancy, labour, and welfare entitlements to the appropriate forums — in semi-urban and rural settings.¹⁵ The potential gain in access to justice is real. So is the risk of unsupervised legal advice. The MeitY's India Artificial Intelligence Governance Guidelines, released on 5th November 2025, sketch a sectoral, non-prescriptive framework that envisages existing regulators tailoring guidance to their own constituencies.¹⁶ The Guidelines explicitly anticipate sectoral bodies stepping in. The Bar Council has not.

III. THE CORE ETHICAL OBLIGATIONS: A FRAMEWORK ANALYSIS

The ethical questions thrown up by generative artificial intelligence map, with surprising neatness, onto five long-established heads of professional responsibility: competence, confidentiality, transparency and disclosure, billing, and liability for error. The first four are duties owed by the advocate before something goes wrong; the fifth governs what happens afterwards. Each is addressed in turn.

A. Competence

ABA Model Rule 1.1 was amended in 2012 to require that a lawyer's competence include keeping abreast of "the benefits and risks associated with relevant technology" (Comment [8] to Rule 1.1, at p.1). Formal Opinion 512 elaborates: a lawyer using a generative artificial intelligence tool must understand the tool's capabilities and limitations to a degree that allows

¹⁵IndiaAI Mission, Ministry of Electronics and Information Technology, *India's AI-Driven Legal Future: Opportunities and Emerging Trends in 2025* (2025), <https://indiaai.gov.in>.

¹⁶Ministry of Electronics and Information Technology, *India AI Governance Guidelines* (Nov. 5, 2025). The Guidelines are anchored in seven *sutras* adapted from the Reserve Bank of India's August 2025 FREE-AI Committee Report.

informed evaluation of its output.¹⁷ The lawyer need not be a data scientist, but she must know enough to spot a hallucinated authority, to estimate the likelihood that a summarised brief has omitted a controlling holding, and to ask the right questions of an artificial intelligence vendor before agreeing to use its product.

The Bar Council of England and Wales' updated guidance of November 2025 puts the same point with more urgency. The Chair of the Bar Council, Ms. Barbara Mills KC, in launching the updated note, warned of "the dangers of the misuse by lawyers of artificial intelligence ... and its serious implications for public confidence in the administration of justice".¹⁸ The guidance — formally not "guidance" under the BSB Handbook but presented as principles that "reflect current professional expectations" (at p.5) — emphasises that the "ultimate responsibility for all legal work remains with the barrister" (at p.5), regardless of how much of that work was generated by software.¹⁹

The Singapore Ministry of Law's Guide of 6th March 2026 organises its discussion of competence around what it calls the "professional ethics" pillar, of which professional accountability is the cornerstone.²⁰ The Guide makes the obvious but important point that "AI tools do not bear professional liability. They are not officers of the court. They are not regulated by the Law Society. They do not carry professional indemnity insurance. The signing partner remains responsible" (at p.1).²¹

There is no comparable time of baseline in India. There is no place in the BCI Rules that specifies an advocate's need to be competent in the technologies she is using. In India, continuing legal education is at best sporadic and at best at best voluntary. No module on legal technology is available in the All India Bar Examination which is partly a competency test. What this means is that advocates are employing generative AI tools at a practical and sophisticated level, from the "I copy-paste and trust the answer" user to the sophisticated power

¹⁷American Bar Association, Standing Committee on Ethics and Professional Responsibility, *Formal Opinion 512: Generative Artificial Intelligence Tools* 1, 7–9 (July 29, 2024) [hereinafter **ABA Formal Opinion 512**].

¹⁸The Bar Council of England and Wales, *Considerations When Using ChatGPT and Generative Artificial Intelligence Software Based on Large Language Models* 3 (Nov. 25, 2025), <https://www.barcouncilethics.co.uk> [hereinafter **UK Bar Council Guidance (2025)**].

¹⁹UK Bar Council Guidance (2025), *supra* note 18, at 5.

²⁰Ministry of Law, Republic of Singapore, *Guide for Using Generative Artificial Intelligence in the Legal Sector* (Mar. 6, 2026) [hereinafter **MinLaw Singapore Guide (2026)**].

²¹Sundaresh Menon & Aurill Kam, *Using Generative AI in Legal Work (Part 2)*, THE SINGAPORE LAW GAZETTE, Mar. 25, 2026, at 1.

user. Variance is a competence issue in itself.

A proposed concrete action is presented. Rule 33 of the BCI Rules, or, more aptly in this context, a new Rule in Part VI of the BCI Rules, should be modified to include, in substance, that an advocate's duty of competence requires him to be knowledgeable of the advantages and disadvantages of the relevant technology used in the performance of the legal profession, including artificial intelligence tools, and that the advocate must be personally responsible for the accuracy and propriety of any work product created in connection with the use of such tools.

B. Confidentiality

Confidentiality is the most acute risk and the least appreciated one. The mechanics are simple. An advocate, working under time pressure, pastes a confidential client document — a draft settlement agreement, an internal corporate memo, a sensitive medical record — into the prompt window of a free chatbot. The chatbot returns a useful summary. The advocate has, in that moment, transmitted privileged material to a third-party service provider in a jurisdiction over which she has no control, governed by terms of service she has not read, and may have authorised that provider to retain or use the input for purposes she has not considered.

ABA Formal Opinion 512 frames the duty bluntly: client data processed through a generative AI tool is entitled to the same protection as data handled in conventional practice. The lawyer must obtain informed consent from the client and must conduct meaningful due diligence on the AI vendor's data practices.²²

The UK Bar Council's November 2025 note is more prescriptive on this point than its January 2024 predecessor. Barristers are warned that they "should never input any personal data in response to prompts from the system" (at p. 7).²³ The note refers to the Italian Data Protection Authority's December 2024 enforcement action against OpenAI, and to parallel investigations in France and Spain, as concrete illustrations of the data-protection risk.²⁴

Singapore's Guide treats confidentiality as one of its three structural pillars and frames the

²²ABA Formal Opinion 512, *supra* note 17, at 11–13.

²³UK Bar Council Guidance (2025), *supra* note 18, at 7.

²⁴Garante per la Protezione dei Dati Personali, Prov. n. 825 (Dec. 20, 2024) (imposing a fine of EUR 15 million on OpenAI); UK Bar Council Guidance (2025), *supra* note 18, at 7.

obligation in terms of platform choice. Cloud-based generative artificial intelligence solutions must be assessed against the Law Society's Guidance Note 3.4.1 on Cloud Computing; client data stored on overseas servers requires particular attention to applicable data-protection law and to the practical recoverability of the data if the vendor relationship ends.²⁵

For India, the analysis must run through the DPDP Act, 2023. Section 4 requires that personal data be processed only for a lawful purpose for which consent has been given. Section 6 prescribes the quality of the consent — it must be free, specific, informed, unconditional, and given through a clear affirmative action. Section 8 makes the Data Fiduciary responsible regardless of contractual arrangements with processors. A law firm that uploads a client's personal data into a generative artificial intelligence tool without first obtaining the client's specific consent for that use, and without ensuring that the artificial intelligence vendor's data-handling practices satisfy DPDP requirements, is on its face exposed.

There is also a prior legal-privilege question. The Bharatiya Sakshya Adhiniyam, 2023, retains the substance of the older Section 126 of the Evidence Act, 1872, protecting professional communications between an advocate and her client from disclosure without consent.²⁶ The doctrine has not been tested in the context of inputs to a third-party artificial intelligence tool. The conservative view is that uploading a privileged communication to a service whose terms permit retention or training-use of the input is itself a waiver, or at the very least a serious risk of waiver. Until the position is judicially clarified, the conservative view should govern conduct.

C. Transparency and Disclosure

Should an advocate disclose to a court that her brief was prepared with the assistance of generative artificial intelligence? Should she disclose to her client?

The first question is currently answered differently across jurisdictions. In the United States, several federal district courts and at least one State — most prominently Pennsylvania, with effect from August 2024 — require explicit disclosure of AI-generated material in court filings,

²⁵MinLaw Singapore Guide (2026), *supra* note 20, § 3.2; Law Society of Singapore, *Guidance Note 3.4.1: Cloud Computing*.

²⁶The Bharatiya Sakshya Adhiniyam, 2023, § 132 (India) (corresponding to the former Indian Evidence Act, 1872, § 126).

accompanied by certifications of verification.²⁷ The Court of King's Bench of Manitoba issued a Practice Direction in June 2023 to similar effect. The United Kingdom has so far stopped short of mandating disclosure across the board, but the High Court has, in three rulings between mid-2024 and late 2025, repeatedly stressed the duty of verification and warned of contempt consequences for unverified citations.²⁸ Singapore's Supreme Court Registrar's Circular No. 1 of 2024 takes a permissive but cautionary approach: court users are told that the use of generative artificial intelligence is not prohibited, but that they bear ultimate responsibility for the accuracy of their submissions.²⁹

India has no analogous rule. The closest analogue is the Supreme Court Rules, 2013, and the High Court Rules of various State High Courts, which require that pleadings be signed by an advocate and verified by the party. Section 151 of the Code of Civil Procedure preserves the inherent power of the court to make orders necessary for the ends of justice. There is no obstacle, in principle, to either the Supreme Court or a High Court framing a practice direction requiring the disclosure of AI-generated material in pleadings — only the absence, so far, of any such direction.

Disclosure to the client is a separate question, with a clearer answer. Where artificial intelligence is used in a material way on a client's matter, the client has a legitimate interest in knowing. The ABA Opinion 512 takes the position that informed consent is required where artificial intelligence use is material to the engagement. Singapore's March 2026 Guide makes provision for client opt-outs and requires the law firm to explain the implications of opting out — for instance, slower turnaround times and potentially higher costs.³⁰ Both frameworks contemplate that use of artificial intelligence disclosure be addressed in the engagement letter at the outset of the retainer. An Indian practice direction could readily adopt the same approach: a standard clause in the vakalatnama or engagement letter recording the firm's policy on the use of generative artificial intelligence tools, together with the client's option to require or

²⁷See Supreme Court of Pennsylvania, *Disciplinary Board Notice on Use of AI in Court Submissions* (Aug. 2024); United States District Court for the Northern District of Texas, *Standing Order on Artificial Intelligence* (May 30, 2023) (Brantley Starr, J.).

²⁸*Ayinde v. The London Borough of Haringey* [2025] EWHC 1383 (Admin) (sanctioning a barrister for citing twelve fabricated authorities); Court of King's Bench of Manitoba, *Practice Direction Re: Use of Artificial Intelligence in Court Submissions* (June 23, 2023).

²⁹Supreme Court of the Republic of Singapore, *Registrar's Circular No. 1 of 2024 — Guide on the Use of Generative Artificial Intelligence Tools by Court Users* § 5.4 (Oct. 23, 2024).

³⁰MinLaw Singapore Guide (2026), *supra* note 20, § 3.3.

prohibit such use.

D. Billing

Generative artificial intelligence changes the economics of legal work in ways that the existing fee structures of the Indian Bar are ill-equipped to absorb. A piece of legal research that took an associate ten hours of book-work and database queries may now be done, with appropriate verification, in thirty minutes of prompt engineering plus an hour of source-checking. Hourly billing for ten hours is no longer defensible; hourly billing for ninety minutes may, in turn, understate the genuine value delivered.

ABA Formal Opinion 512 is explicit: the efficiencies generated by artificial intelligence must be reflected in fair and transparent billing, and a lawyer may not charge the client for time that the artificial intelligence has saved. Where the lawyer charges the client for the use of an artificial intelligence tool itself — for instance, where a per-query API charge is incurred — the charge must be disclosed in advance and must reflect actual cost.³¹

UK guidance is less detailed on billing but consistent in principle: the cab-rank rule and the Bar Standards Board's expectations of fair fees apply to AI-assisted work as much as to traditional work, and chambers' leadership is expected to set clear policies.³²

In India, fee arrangements between advocates and clients are largely unregulated beyond the broad professional-misconduct framework of Section 35 of the Advocates Act.³³ The absence of any concrete billing guidance leaves Indian clients in a weak position. A client who is billed for forty hours of “research” on a matter that was substantially completed in a chatbot session has no easy means of knowing this, and no easy means of contesting it.

A reasonable starting position would be the adoption, by way of a BCI circular, of two basic principles. First, an advocate may not bill the client for time saved by the use of artificial intelligence. Second, any AI-related charges (platform-licensing fees or per-query costs) passed on to the client must be itemised, disclosed in advance, and limited to actual cost.

³¹ABA Formal Opinion 512, *supra* note 17, at 14–15.

³²UK Bar Council Guidance (2025), *supra* note 18, at 9.

³³The Advocates Act, No. 25 of 1961, § 35 (read with §§ 36–37) (India).

E. Hallucination Liability

The *Mata v. Avianca* scenario is no longer hypothetical in India. The Delhi High Court's September 2025 episode of phantom precedents, and the Bombay High Court's January 2026 imposition of costs, both demonstrate that Indian courts are now encountering hallucinated authorities. The question is what happens to the lawyer responsible.

There are three potential heads of liability. The first is professional misconduct under the Advocates Act, 1961, read with the BCI Rules. An advocate who knowingly cites a fabricated authority is acting in breach of her duty of candour to the court; an advocate who cites a fabricated authority without verifying it is, at the least, negligent. The remedy lies in disciplinary proceedings before the State Bar Council, with an appeal to the BCI under Section 37 of the Act.

The second head is contempt. The Delhi High Court in its September 2025 order warned that the misuse of AI-generated material could attract contempt proceedings or, in particular cases, perjury liability. The substantive law of contempt is wide enough to reach the conduct, though no Indian court has yet imposed contempt sanctions specifically for AI-generated hallucinations.

The third head is civil liability for professional negligence. The relevant standard remains that articulated by the Supreme Court in *Jacob Mathew v. State of Punjab*.³⁴ Although the case concerned medical negligence, the test it endorsed — the Bolam standard, requiring that the professional act with the degree of skill and care that an ordinarily competent member of the profession would exercise — applies, *mutatis mutandis*, to other learned professions, including the law.³⁵ An advocate who places before a court an authority she has not verified, when verification was practicable and would have revealed the authority to be fabricated, has fallen below the Bolam threshold. The harm caused — a wasted hearing, an adverse costs order against the client, the loss of credibility before the court — is recoverable in principle in a civil suit for professional negligence.

The picture, then, is not that India has no law to govern artificial intelligence hallucinations in

³⁴*Jacob Mathew v. State of Punjab*, (2005) 6 SCC 1.

³⁵*Bolam v. Friern Hosp. Mgmt. Comm.*, [1957] 1 WLR 582; applied in *Jacob Mathew*, *supra* note 34, ¶¶ 18, 48. The Court treated the Bolam test as the governing standard for negligence by a professional in India.

pleadings. It is that the existing law, fashioned for an analog era, will reach the conduct only awkwardly and only after significant damage has been done. A clear *ex ante* rule — that the advocate's signature on a pleading carries with it a personal certification that all cited authorities have been verified against the original — would reduce the incidence of the problem before it arrives at the court's door.

IV. COMPARATIVE REGULATORY MODELS

Each of the three jurisdictions surveyed below has approached the problem differently, and each has something to teach the Indian regulator.

A. United States: American Bar Associations' formal opinion 512 and the limits of guidance without enforcement

The American Bar Association's Formal Opinion 512, released on 29th July 2024, is the most comprehensive treatment of artificial intelligence ethics in the legal profession so far produced.³⁶ It organises its discussion around six existing Model Rules — competence (Rule 1.1), confidentiality (Rule 1.6), client communication (Rule 1.4), candour to the tribunal (Rules 3.1 and 3.3), supervisory responsibility (Rules 5.1 and 5.3), and reasonable fees (Rule 1.5) — and applies each to generative artificial intelligence use without attempting to invent new rules.³⁷

This is the Opinion's principal strength and its principal limit. Its strength is that it does not pretend artificial intelligence is so unprecedented that wholly new ethical categories are required. The duties of competence, confidentiality, and candour have always applied to whatever technology the lawyer uses; the Opinion's contribution is to map the application to a new tool. Its limit is that the ABA's Model Rules are not directly enforceable. They are guidance for State bar regulators, who adopt — or do not adopt — the Model Rules into binding State rules of professional conduct. As of early 2026, several State bars including California, Florida, New Jersey, New York, and Pennsylvania have either issued their own opinions or adopted rules tracking Formal Opinion 512. Others have not. The result is uneven.

³⁶ABA Formal Opinion 512, *supra* note 17.

³⁷*Id.* at 2–3 (mapping the discussion to ABA Model Rules 1.1 (Competence), 1.4 (Communication), 1.5 (Reasonable Fees), 1.6 (Confidentiality), 3.1, 3.3 (Candor Toward the Tribunal), 5.1, and 5.3 (Supervisory Responsibilities)).

What has driven actual change in lawyer behaviour in the United States has not been the ABA Opinion alone. It has been judicial enforcement through sanctions in cases like *Mata v. Avianca*, *Park v. Kim*, *Gauthier v. Goodyear*, and most recently *Shelton v. Parkland Health*.³⁸ In *Shelton*, decided in November 2025, the Northern District of Texas held that a lawyer's good faith was not enough to insulate her from Rule 11 sanctions for citing AI-fabricated authorities.³⁹ The combination of sanctions and the ABA Opinion has produced a feedback loop in which firms have rapidly developed internal AI policies and training programmes — driven less by aspiration than by liability risk.

The Indian takeaway is twofold. First, an ethics opinion alone is rarely sufficient to change behaviour; it must be coupled with credible enforcement. Second, the Indian advantage is that the BCI sits at the apex of a unified professional regulatory regime in a way that the ABA does not. A circular issued by the BCI would have direct application across the country in a way that the ABA Opinion does not.

B. United Kingdom: The Bar Council's principles, recalibrated by High Court rulings

The Bar Council of England and Wales took a different route. Its January 2024 guidance, *Considerations when using ChatGPT and Generative AI Software based on Large Language Models*, was framed not as binding “guidance” under the Bar Standards Board Handbook, but as practical assistance from the Council's Information Technology Panel.⁴⁰ The November 2025 update was triggered by a series of High Court rulings, most prominently *Ayinde*, in which a barrister was sanctioned for citing twelve fabricated authorities in submissions.⁴¹ The update preserves the structure of the 2024 version but recalibrates the language: the principles, the revised note states, “reflect current professional expectations” (at p. 2) in the light of the case law, and a “heightened imperative” exists to use AI responsibly.⁴²

The UK approach has three features worth noting. First, it deals concretely with the risks of generative artificial intelligence rather than abstractly. The 2025 note expressly names

³⁸*Mata v. Avianca*, *supra* note 1; *Park v. Kim*, 91 F.4th 610 (2d Cir. 2024); *Gauthier v. Goodyear Tire & Rubber Co.*, 2024 WL 4882651 (E.D. Tex. Nov. 25, 2024); *Shelton v. Parkland Health*, No. 3:24-cv-2190, 2025 WL 3012828 (N.D. Tex. Nov. 10, 2025).

³⁹*Shelton v. Parkland Health*, *supra* note 38, holding that a lawyer's good faith does not, by itself, insulate her from Rule 11 sanctions for citing AI-fabricated authorities.

⁴⁰The Bar Council of England and Wales, *Considerations When Using ChatGPT and Generative AI Software Based on Large Language Models* (Jan. 30, 2024).

⁴¹*Ayinde*, *supra* note 28; see also *Bandla v. Solicitors Regulation Authority* [2024] EWHC 1230 (Admin).

⁴²UK Bar Council Guidance (2025), *supra* note 18, at 2.

ChatGPT, Gemini, Perplexity, Copilot, Lexis+ AI, Clio Duo, and Thomson Reuters Co-Counsel as the products in scope, and addresses each according to the risks specific to it.⁴³ Second, it integrates data-protection considerations directly into ethics guidance; the EU GDPR continues to apply post-Brexit through the UK GDPR, and the note repeatedly cross-references this regime. Third, it makes a point — emphasised by Barbara Mills KC at the launch — that responsibility cannot be delegated to chambers' staff or junior counsel.⁴⁴

For India, the lesson is that ethical guidance can be issued in light, non-statutory form so long as it is anchored in concrete examples and supported by enforcement at the level of the court. The UK note is not a statute. It is a working document, updated as the technology and the case law develop. The BCI could plausibly produce a comparable working document for Indian advocates, anchored in Indian case law and Indian data-protection statute.

C. Singapore: A whole-of-sector framework

Singapore's approach is the most systematically architected of the three. The Ministry of Law's Guide for Using Generative Artificial Intelligence in the Legal Sector, launched on 6th March 2026 after a public consultation in September 2025, takes the form of a comprehensive sectoral framework.⁴⁵ It works in tandem with the Supreme Court Registrar's Circular No. 1 of 2024 (which governs court users) and with the Law Society's Guidance Note 3.4.1 on Cloud Computing.⁴⁶

There are three aspects in the Singapore framework that warrant careful consideration. First, the Guide is divided along three axis: professional ethics, confidentiality and transparency, and not abstract exhortations, but operational checklists. Second, it includes client opt-outs in the engagement letter and makes it clear what happens when a client opts out and the fee and turnaround time. Third, it adds the element of law-firm governance to the mix: law firms are expected to have written Artificial Intelligence policies in place, designate a senior lawyer to oversee compliance with Artificial Intelligence, and keep track of how Artificial Intelligence

⁴³*Id.* at 4 (expressly naming ChatGPT, Gemini, Perplexity, Copilot, Lexis+ AI, Clio Duo, and Thomson Reuters Co-Counsel as the products in scope).

⁴⁴Statement of Barbara Mills KC, Chair of the Bar Council, at the Launch of the Updated Guidance (Nov. 25, 2025), <https://www.barcouncil.org.uk/resource/updated-guidance-on-generative-ai-for-the-bar.html>.

⁴⁵MinLaw Singapore Guide (2026), *supra* note 20.

⁴⁶Supreme Court of Singapore, Registrar's Circular No. 1 of 2024, *supra* note 29; Law Society of Singapore, *Guidance Note 3.4.1: Cloud Computing*.

has been used on client matters and maintain records that can be audited.⁴⁷

Singapore also gains from a regulator (the Ministry of Law working with the Law Society) that has historically operated in close partnership with the profession. The Singapore Academy of Law's collaboration with Microsoft on a Prompt Engineering for Lawyers guide is a concrete example of regulatory-private co-production of guidance materials.⁴⁸

The Indian analogue would be a BCI working group, including representatives of the Supreme Court Bar Association, the SCC Online and Manupatra editorial teams, MeitY, and senior litigators, tasked with producing an Indian version of the Singapore Guide. The architecture is replicable. The political will is, so far, absent.

D. A composite Indian model

Taken together, the three jurisdictions suggest a composite Indian model with four ingredients. First, an interim ethics advisory issued by the BCI under Section 49(1)(c) of the Advocates Act, on the model of ABA Formal Opinion 512, mapping each existing duty of professional conduct to use artificial intelligence. Second, a Supreme Court practice direction — which could be issued under Article 145 of the Constitution and the Supreme Court Rules, 2013 — requiring disclosure of AI-generated material in pleadings and certification of verification. Third, integration with the DPDP Act through a memorandum of understanding between the BCI and MeitY, addressing the position of law firms as Data Fiduciaries. Fourth, a mandatory CLE module on artificial intelligence competence as a condition of continuing enrolment, tied to the All India Bar Examination architecture. Australia's Law Council and the Law Society of New South Wales have, in 2024 and 2025, issued similar guidance documents and may also be drawn on, though their regulatory shape differs from the Indian.⁴⁹

V. CONCLUSION

The Bar Council of India's absence in the matter of using generative artificial intelligence by Indian advocates has now become a policy to be followed instead of a lag. It's not a wise policy decision. At every stage, from the chambers of the senior counsel to the district court

⁴⁷MinLaw Singapore Guide (2026), *supra* note 20, §§ 4–5 (governance and implementation framework).

⁴⁸Singapore Academy of Law & Microsoft Corp., *Prompt Engineering for Lawyers* (2024).

⁴⁹Law Council of Australia, *Generative AI in Legal Practice: Guidance for Australian Legal Practitioners* (2024); Law Society of New South Wales, *Lawyers Adopting AI: Practical Guidance* (2024).

youngsters, the Bar in India is utilising the tools of artificial intelligence (AI) and the rules for their use have yet to be drafted.

This paper has suggested that, the BCI's ongoing inaction results in four tangible harms. Clients are put at risk when the data they upload to a system that has not been vetted for the advocate's retention practices. Courtrooms are open, and cases are being presented, where the citations the courts are receiving cannot be easily checked at the moment. The advocates are exposed because they are using tools without any information about what is and isn't allowed. The integrity of the entire system of justice is being undermined, one bogus citation at a time. All these are not hypothetical, but are represented in the case law and contemporary record.

The remedies are also clear, and they are within the BCI's existing statutory authority. Five steps would suffice as a starting framework.

First, the BCI ought to issue an Interim Ethics Advisory under Section 49(1)(c) of the Advocates Act addressing the use of generative artificial intelligence by advocates. The Advisory should track the five-head framework set out in Part III of this paper: competence, confidentiality, transparency, billing, and liability. It should be issued as a "best practices" document in the first instance, with a defined timeline for transition into binding amendments to the BCI Rules.

Second, artificial intelligence skills should be required to be taught as a mandatory course in Continuing Legal Education. The All India Bar Examination apparatus is already in place, and there is an easy additional module to administer. It is not about making everyone an AI engineer; it's about making sure that nobody signs a pleading that has AI generated content they don't understand.

Third, the BCI should be a Technology Ethics Committee comprising of senior members of the Bar, the academy and the technology sector who would have an ongoing role in monitoring developments and in updating guidance. The rate of AI development in the field won't slow down and the regulator needs to be able to keep up.

Fourth, the BCI should formally engage with MeitY with respect to the application of the DPDP Act to the application of law firms. Sectoral clarifications on legal correspondence and pleadings under Section 8, the cross-border processing and the legal position of cloud-based

legal AI tools are all in need of clarification. An MoU between BCI and MeitY as is happening with the RBI and IRDAI is the right instrument.

Fifth, the BCI should create a disclosure clause for AI, similar to the disclosure clause in Singapore, to be incorporated in the vakalatnamas and engagement letters to ensure that the AI usage is discussed from the outset of the retainer.

Administratively, each of these steps is achievable! All need no statutory changes. These can be all achieved under the current structure of the Advocates Act 1961.

The deeper question is whether the BCI sees the regulation of AI use as part of its mandate at all. The argument for doing so is not, in the end, about technology. It is about the foundational character of the legal profession as one whose members owe duties of competence, confidentiality, candor, and fidelity to their clients and to the courts they appear before. Those duties do not lapse when the lawyer reaches for a new tool. They are recast, instead, against the affordances and the risks of that tool. The Bar Council of England and Wales, the American Bar Association, and the Singapore Ministry of Law have each begun this work. The Bar Council of India has not. The longer the delay, the harder the catch-up; and the higher the cost, in the meantime, for the clients and the courts whose protection is the only justification for the existence of a regulated profession.