

**INTERNATIONAL JOURNAL OF LAW
MANAGEMENT & HUMANITIES**
[ISSN 2581-5369]

Volume 9 | Issue 5

2026

© 2026 *International Journal of Law Management & Humanities*

Follow this and additional works at: <https://www.ijlmh.com/>

Under the aegis of VidhiAagaz – Inking Your Brain (<https://www.vidhiaagaz.com/>)

This article is brought to you for free and open access by the International Journal of Law Management & Humanities at VidhiAagaz. It has been accepted for inclusion in the International Journal of Law Management & Humanities after due review.

In case of **any suggestions or complaints**, kindly contact support@vidhiaagaz.com.

To submit your Manuscript for Publication in the **International Journal of Law Management & Humanities**, kindly email your Manuscript to submission@ijlmh.com.

AI-Generated Authorities and the Integrity of Judicial Proceedings: Reconstructing the Lawyer's Duty of Verification in India

SONIA ARORA*

ABSTRACT

Between December 2024 and July 2026, Indian tribunals and trial courts relied on judicial precedents that did not exist, generated by a large language model and accepted without verification; in one instance, a High Court that itself recognised such a fabrication on revision nonetheless declined to treat it as fatal to the outcome. The pattern culminated in Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd., where the Supreme Court set aside a National Company Law Tribunal order and its confirmation in appeal, both resting on fabricated precedents, and declared that citing AI-generated authorities without verification is professional misconduct. This paper argues that the Supreme Court's response, while doctrinally significant, has so far declared a duty of verification without defining its content: it does not specify what verification requires, against which sources, how it should be documented, or how responsibility should be apportioned between an Advocate-on-Record, briefing counsel, and the junior associates who typically conduct the underlying research. The paper traces the existing architecture of the advocate's professional obligations under Sections 35 and 36 of the Advocates Act, 1961 and the Bar Council of India's Standards of Professional Conduct and Etiquette, and shows that this architecture, framed before generative artificial intelligence existed, supplies only generalised duties of diligence and candour rather than a specific, source-anchored protocol. Drawing on the doctrinal lineage from Gummati Usha Rani v. Sure Mallikarjuna Rao through Pooja Ramesh Singh, on the Supreme Court's draft Regulations for Use of AI in Courts, 2026, and on comparative practice under Mata v. Avianca, Inc. in the United States, the paper proposes a tiered, source-anchored verification standard, keyed to source authoritativeness, procedural stage, and research-responsibility allocation within a legal team, capable of catching fabricated authorities before a judge relies on them rather than after an entire tier of appeal already has.

Keywords: AI Hallucination, Duty of Verification, Professional Misconduct, Advocates Act 1961, Judicial Integrity

* Author is a Student at Gitarattan International Business School, Guru Gobind Singh Indraprastha University, New Delhi, India.

I. INTRODUCTION AND RESEARCH PROBLEM

On 27 February 2026, a Bench of the Supreme Court of India, hearing a special leave petition arising from an injunction suit over property in Andhra Pradesh, recorded that a trial court had, by an order of 19 August 2025, dismissed objections to an advocate commissioner's report by relying on four purported Supreme Court judgments that did not exist.¹ The Andhra Pradesh High Court, on revision, had itself recognised that the citations were artificial-intelligence-generated fabrications, recorded a word of caution, and nonetheless proceeded to dismiss the revision petition on the merits, affirming the trial court's decision.² The Supreme Court was not satisfied that this was an adequate response. It observed that a decision founded on non-existent and fake judgments is not merely an error in decision-making; it declared, in an order that has since been widely reported, that such reliance "would be a misconduct and legal consequence shall follow," and issued notice to the Attorney General for India, the Solicitor General of India, and the Bar Council of India, while appointing Senior Advocate Shyam Divan to assist the Court.³

This was neither the first nor, as it turned out, the last such episode. An order of the Bengaluru Bench of the Income Tax Appellate Tribunal dated 30 December 2024, in a trust-taxation dispute involving approximately Rs 669 crore, had already been recalled by the Tribunal on 7 January 2025, after it emerged that the order relied on three Supreme Court judgments and one Madras High Court ruling which, as reported, did not exist.⁴ In *KMG Wires Private Ltd. v. National Faceless Assessment Centre*, a Division Bench of the Bombay High Court comprising Justice B.P. Colabawalla and Justice Amit S. Jamsandekar, by an order dated 6 October 2025, quashed an income tax assessment for Assessment Year 2023–24 which had determined the petitioner's total income at approximately Rs 27.91 crore, after finding that the Assessing

¹ *Gummadi Usha Rani v. Sure Mallikarjuna Rao*, Special Leave Petition (Civil) No. 7575 of 2026, 2026 SCC OnLine SC 341, ¶¶ 2–4 (India Feb. 27, 2026) (order), https://api.sci.gov.in/supremecourt/2026/6571/6571_2026_6_22_68896_Order_27-Feb-2026.pdf.

² *Gummadi Usha Rani v. Sure Mallikarjuna Rao*, Civil Revision Petition No. 2487 of 2025 (A.P. High Ct. Jan. 21, 2026) (India), as recorded in *Gummadi Usha Rani*, 2026 SCC OnLine SC 341, ¶ 4; see also Aakriti Bansal, *10 Cases That Show Indian Courts Have an AI Hallucination Problem*, MEDIANAMA (July 3, 2026), <https://www.medianama.com/2026/07/223-10-cases-ai-hallucination-cases-in-indian-courts/>.

³ *Gummadi Usha Rani*, 2026 SCC OnLine SC 341, ¶¶ 7–9; see also Ritu, *SC Raises Concern over AI-Generated Fake Precedents*, SCC ONLINE BLOG (Mar. 11, 2026), <https://www.sconline.com/blog/post/2026/03/11/sc-raises-concern-use-ai-generated-fake-precedents/>.

⁴ *Buckeye Trust v. Principal Comm'r of Income Tax*, ITA No. 1051/Bang/2024 (Income Tax App. Trib., Bangalore Dec. 30, 2024) (India), <https://indiankanoon.org/doc/169046995/>, recalled (Jan. 7, 2025); see *ITAT Bangalore Recalls Its Order Dated 30 December 2024 in Case of Buckeye Trust*, MONDAQ (Jan. 14, 2025), <https://www.mondaq.com/india/tax-authorities/1567478/itad-bangalore-recalls-its-order-dated-30-december-2024-in-case-of-buckeye-trust>; *Buckeye Trust v. Registrar, Income Tax Appellate Tribunal*, Writ Petition No. 25280 of 2025, 2025:KHC:37479 (Kar. High Ct. Sept. 18, 2025) (India) (noting the recall), <https://indiankanoon.org/doc/145492551/>; see also *AI-Hallucinated Case Law: How Fake Citations Are Getting Lawyers Sanctioned in India (and How to Avoid It)*, IPLEADERS (July 6, 2026).

Officer had relied on three non-existent judicial decisions while computing one of the impugned additions; the Court held that results thrown up by AI are not to be blindly relied upon by an authority exercising quasi-judicial functions but must be duly cross-verified before use.⁵ In March 2025, a single judge of the Karnataka High Court, allowing a civil revision petition, found that a City Civil Court judge in Bengaluru had relied on two Supreme Court decisions that were never rendered, in rejecting an application for return of the plaint, and directed that the order be placed before the Chief Justice for further action against the judge; this appears to be among the few instances in which a High Court has proceeded against a judge, rather than an advocate, over fictitious citations.⁶ And on 17 February 2026, during a hearing before a Bench of Chief Justice of India Surya Kant and Justices B.V. Nagarathna and Joymalya Bagchi, in the week of the India AI Impact Summit in New Delhi, the Court flagged what the Chief Justice described as an “alarming” trend: Justice Nagarathna referred to a petition that had placed before the Court a fictitious judgment styled “*Mercy v. Mankind*”, and the Chief Justice noted that a series of similarly fabricated judgments had been cited before Justice Dipankar Datta.⁷

The pattern reached its most consequential expression on 2 July 2026, when the Supreme Court decided *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.*⁸ The National Company Law Tribunal, Mumbai Bench, had admitted a Section 7 application under the Insolvency and Bankruptcy Code, 2016 against Essel Infraprojects Ltd. on 28 August 2024, and the National Company Law Appellate Tribunal had affirmed that order on 11 September 2025 without noticing that the NCLT’s reasoning rested, in material part, on non-existent, fake, and hallucinated judicial material. It was Ms. Madhavi Divan, senior counsel appearing for the appellant, who at the outset pointed out to the Supreme Court that the authorities relied upon by the NCLT, and left unexamined by the NCLAT, were fake and non-existent.⁹ The Supreme Court

⁵ *KMG Wires Pvt. Ltd. v. Nat’l Faceless Assessment Ctr.*, Writ Petition (L) No. 24366 of 2025, 2025:BHC-OS:19789-DB, ¶¶ 2, 9, 11 (Bom. High Ct. Oct. 6, 2025) (India), <https://indiankanoon.org/doc/181115899/> (quashing the assessment order dated Mar. 27, 2025); see also *AI ‘Hallucination’ in Tax Assessment: Bombay HC Sets Aside Rs27.91-Crore Order Passed on Non-existent Case Laws*, MONEYLIFE (Oct. 27, 2025); *Don’t Blindly Trust AI’: Bombay High Court Quashes Income Tax Assessment Passed on Unverified AI-Generated Case Laws*, LIVELAW (Oct. 27, 2025); *Bombay High Court Pulls Up Tax Officer for Citing Fake, AI-Generated Case Laws*, BAR & BENCH (Oct. 31, 2025).

⁶ *Sammaan Capital Ltd. v. Mantri Infrastructure Pvt. Ltd.*, Civil Revision Petition No. 49 of 2025, ¶¶ 10, 12 (Kar. High Ct. Mar. 24, 2025) (India) (R. Devdas, J.), <https://indiankanoon.org/doc/184681356/>; see also Bansal, *supra* note 2.

⁷ Manoj Rahul, *What the India AI Impact Summit Means for Indian Legal Practice*, BAR & BENCH (Feb. 27, 2026), <https://www.barandbench.com/columns/when-ai-agents-enter-the-courtroom-what-the-delhi-summit-means-for-indian-legal-practice>; see also Bansal, *supra* note 2.

⁸ *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.*, Civil Appeal No. 11950 of 2025, 2026 INSC 668 (India July 2, 2026) (Narasimha & Aradhe, JJ.), <https://indiankanoon.org/doc/113338666/>.

⁹ *Pooja Ramesh Singh*, 2026 INSC 668, ¶¶ 12–14; see also *SC Sets Aside Orders in Jammu and Kashmir Bank Case over AI-Hallucinated Judgments*, KASHMIR LIFE (July 2, 2026). Ms. Madhavi Divan, who appeared for the

set aside both orders, holding that a decision based even partly on such material is no decision at all, and describing the phenomenon in unusually vivid terms: that “the production of fake, non-existent, and hallucinated material and its utilisation as precedents in law, is like the release of methyl isocyanate in the province of law and justice: invisible, insidious, and catastrophic by the time anyone notices,” such that “[i]t not only contaminates but takes away the very lifeblood of judicial determination”.¹⁰ The Court went further, holding that it is misconduct on the part of an advocate to cite such judgments without verification, and that courts must adopt a zero-tolerance mode for producing, citing or using AI-generated precedents without verification. It directed the Bar Council of India, as the apex statutory body, to constitute a committee to deliberate on the issue and to prescribe both a guiding principle and the disciplinary action that should follow a violation.¹¹

This paper takes *Pooja Ramesh Singh* as its point of departure, but its research problem lies one step beyond the judgment. The Supreme Court has now said, with unmistakable clarity, that a duty of verification exists and that its breach is misconduct. What the Court has not said, because it was not, on the facts before it, required to say, is what verification actually consists of: against which source an advocate must check a citation before relying on it; how that check should be documented so that its performance, or its omission, can later be established; how responsibility should be distributed within a legal team in which the individual who conducted the underlying research, an Advocate-on-Record who settles and signs the pleading, and senior counsel who argues the matter may each have played a different role and possessed different degrees of actual knowledge; and how the standard should differ, if at all, between a habeas corpus petition drafted overnight and a considered written submission filed months into a commercial arbitration. The Bar Council of India’s Standards of Professional Conduct and Etiquette, framed under the Advocates Act, 1961 long before generative artificial intelligence existed, are silent on all these questions.¹² This paper is an attempt to reconstruct the content of that duty: not merely to note that it exists, which *Pooja Ramesh Singh* has already done, but to give it a shape precise enough to be taught, audited, and enforced.

appellant, is not to be confused with Mr. Shyam Divan, Senior Advocate, appointed to assist the Court in *Gummadi Usha Rani*.

¹⁰ *Pooja Ramesh Singh*, 2026 INSC 668, ¶ 6; see also *AI-Hallucinated Fake Judgments Are Like ‘Release of Methyl Isocyanate’ in Justice System: Supreme Court*, LAWBEAT (July 2, 2026).

¹¹ *Pooja Ramesh Singh*, 2026 INSC 668, ¶¶ 7, 9.

¹² Rahul, *supra* note 7 (observing that the BCI’s Standards of Professional Conduct and Etiquette are “entirely silent on AI delegation”).

II. RESEARCH OBJECTIVES

- To document, from verified judicial and tribunal records, the pattern of AI-generated fabricated authorities that has emerged in Indian adjudication between December 2024 and July 2026.
- To trace the doctrinal response of the Supreme Court of India, from its initial cognisance of the problem in *Gummadi Usha Rani v. Sure Mallikarjuna Rao* to its considered pronouncement in *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.*
- To examine whether the existing statutory and regulatory architecture governing advocates in India, namely the Advocates Act, 1961 and the Bar Council of India's Standards of Professional Conduct and Etiquette, supplies a workable standard of verification, or only a generalised duty of diligence incapable of being audited or enforced with precision.
- To assess the Supreme Court's draft Regulations for Use of Artificial Intelligence in Courts, 2026 as an institutional response, and to identify the gaps that remain even if the draft is notified in its current form.
- To propose a tiered, source-anchored standard of verification, together with procedural and institutional reforms, capable of detecting fabricated authorities before they are relied upon by a court rather than after an entire tier of appeal has already done so.

III. RESEARCH QUESTIONS

- What does the duty of verification articulated in *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.* require an advocate to do, and against which sources must that verification be conducted?
- How should responsibility for an unverified or fabricated citation be apportioned between an Advocate-on-Record, briefing or senior counsel, and the junior associates or paralegals who typically conduct legal research?
- Does the existing framework of professional misconduct under Sections 35 and 36 of the Advocates Act, 1961 provide an adequate basis for disciplining unverified reliance on AI-generated authorities, or does it require supplementation?
- Can the Supreme Court's draft Regulations for Use of Artificial Intelligence in Courts, 2026, as presently framed, meaningfully reduce the incidence of fabricated

authorities reaching a court, or do they address disclosure without addressing verification as such?

- What institutional and procedural safeguards, short of an outright prohibition on the use of generative AI in legal research, would allow fabricated authorities to be caught before, rather than after, they influence a judicial or quasi-judicial decision?

IV. RESEARCH HYPOTHESES

H1: The duty of verification recognised in *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.* is, at present, a duty declared in principle but undefined in content, and its practical enforceability depends on the Bar Council of India supplying the specificity that the judgment itself does not.

H2: The recurrence of AI-hallucinated authorities across trial courts, a High Court, a specialised tribunal, and its appellate counterpart between December 2024 and September 2025, most of them undetected until challenged by a party, demonstrates that reliance on after-the-fact judicial detection is structurally inadequate as the principal safeguard against fabricated citations.

H3: The existing professional conduct framework under the Advocates Act, 1961 and the Bar Council of India's Standards of Professional Conduct and Etiquette, being silent on generative artificial intelligence, cannot by itself supply the specific, source-anchored verification standard that a workable duty requires.

H4: A tiered standard of verification, calibrated to the authoritativeness of the source consulted, the stage of the proceeding, and the allocation of research responsibility within a legal team, is more likely to be both administrable and enforceable than a single, undifferentiated duty to "verify".

V. RESEARCH METHODOLOGY

This paper adopts a doctrinal, black-letter methodology,¹³ supplemented by a documentary case-study method in respect of the reported incidents discussed in Part VII.A, and follows the citation convention prescribed by the Bluebook (21st edition), adapted to incorporate the neutral citation format now used by Indian courts for reported judgments (e.g., 2026 INSC 668) alongside SCC OnLine citations (e.g., 2026 SCC OnLine SC 341).¹⁴ The primary legal sources examined are the Advocates Act, 1961, the Bar Council of India Rules (in particular the

¹³ See generally Terry Hutchinson & Nigel Duncan, *Defining and Describing What We Do: Doctrinal Legal Research*, 17 DEAKIN L. REV. 83 (2012), <https://doi.org/10.21153/dlr2012vol17no1art70>; LEGAL RESEARCH AND METHODOLOGY (S.K. Verma & M. Afzal Wani eds., Indian L. Inst., 2d ed. 2001).

¹⁴ See THE BLUEBOOK: A UNIFORM SYSTEM OF CITATION (Columbia L. Rev. Ass'n et al. eds., 21st ed. 2020).

Standards of Professional Conduct and Etiquette framed under Section 49(1)(c) of the Act), the Insolvency and Bankruptcy Code, 2016, the Income-tax Act, 1961, and the Supreme Court's draft Regulations for Use of Artificial Intelligence in Courts, 2026.¹⁵ Case law has been drawn from primary court records where available (the full text of *Pooja Ramesh Singh* has been consulted through Indian Kanoon and Casemine) and, where the primary record of an order was not independently accessible, as with certain tribunal orders discussed only in secondary reportage, the paper relies on contemporaneous reporting by specialised legal news outlets, cross-checked across at least two independent sources before inclusion. Comparative material is confined to the United States' experience with fabricated AI citations in litigation, most notably *Mata v. Avianca, Inc.*, used to test whether a jurisdiction with a more developed body of sanctions practice offers lessons transferable to the Indian professional conduct framework. Sources are treated as a hierarchy for this purpose: a primary court record (accessed directly, as with the full text of *Pooja Ramesh Singh*) ranks above reporting by a specialised legal news outlet, which in turn ranks above general news coverage. Where a factual claim rests on a single source below the level of a primary record, that source's standing is disclosed in the relevant footnote rather than presented with the same confidence as a directly verified fact, and a small number of such claims, identified as such in the text, have been retained on this qualified basis because they could not be corroborated within the scope of this paper, rather than being either omitted or silently upgraded in confidence. Primary judgment texts were retrieved through Indian Kanoon, Casemine, and Advocatekhaj; secondary reportage was located through general web search using terms including "AI hallucinated citations India," "fake case law tribunal," and the names of the specific matters once identified. The literature search underlying this paper was last conducted in August 2026, and developments after that date are not reflected.

¹⁵ SUPREME COURT OF INDIA, ARTIFICIAL INTELLIGENCE COMMITTEE, *Draft Regulations for Use of Artificial Intelligence (AI) in Courts, 2026* (June 3, 2026) (committee headed by P.S. Narasimha, J.); SUPREME COURT OF INDIA, Notice Seeking Views/Suggestions on Draft 'Regulations for Use of Artificial Intelligence (AI) in Courts, 2026' (June 3, 2026) (describing the draft as "grounded in the principles of human primacy, transparency, accountability, data protection, and judicial independence" and inviting comments by June 20, 2026), <https://cdnbbsr.s3waas.gov.in/s3ec0490f1f4972d133619a60c30f3559e/uploads/2026/06/2026060342.pdf>; SUPREME COURT OF INDIA, Notice Extending Time for Comments on Draft 'Regulations for Use of Artificial Intelligence (AI) in Courts, 2026' (June 16, 2026) (extending the deadline to July 15, 2026), https://allahabadhighcourt.in/Final_draft_with_Notice_v2.pdf. The draft proposes, among other measures, a permanent Apex Body at the Supreme Court, AI Committees in the High Courts, and a Centre of Research and Excellence on Artificial Intelligence (CoRE-AI) to test AI systems before judicial deployment. See Namrata Banerjee & Debjani M, *Order in the Digital Court: Artificial Intelligence Regulations, Supreme Court*, SUPREME COURT OBSERVER (July 13, 2026), <https://www.scobserver.in/journal/order-in-the-digital-court-artificial-intelligence-regulations-supreme-court/>; *Explained: The Supreme Court of India's Draft Regulations for Use of Artificial Intelligence in Courts, 2026*, THE LEAFLET (June 24, 2026); *India's Supreme Court Releases Draft Regulations for Use of AI in Courts, 2026 for Public Consultation*, ASIA IP (June 23, 2026); *Critical Analysis of India's Draft Regulations for Use of Artificial Intelligence in Courts, 2026*, MONDAQ (June 11, 2026).

Conflict of interest and funding: The author declares no conflict of interest and no funding source in connection with this paper, and has not submitted comments on the Supreme Court's draft Regulations for Use of Artificial Intelligence in Courts, 2026 discussed herein.

VI. LITERATURE REVIEW

Commentary on AI-hallucinated case law in India has grown rapidly over the period this paper examines, but it remains chronological and descriptive rather than doctrinally synthetic. A substantial body of writing, largely produced by legal news outlets and law-firm publications, has focused on cataloguing individual incidents as they emerged: the *Buckeye Trust* recall before the Income Tax Appellate Tribunal, the Bombay High Court's quashing of a faceless assessment order, the Karnataka High Court's direction for action against a trial judge, and the Andhra Pradesh proceedings that eventually reached the Supreme Court in *Gummadi Usha Rani*.¹⁶ This literature performs a valuable documentary function, and this paper draws on it extensively in Part VII.A, but it has generally stopped short of asking what, precisely, a reconstructed duty of verification should require as a matter of professional regulation, treating the problem primarily as a matter of individual carelessness to be remedied by judicial admonition rather than as a gap in the profession's regulatory architecture.

A second strand of commentary has focused on the institutional response, particularly the Supreme Court's own posture as both regulator of, and participant in, AI adoption within the judiciary. Writers have noted the apparent tension in a court that operates its own AI tools (the Supreme Court Portal for Assistance in Court Efficiency, the Supreme Court Vidhik Anuvaad Software for translation, and a generative legal research assistant) while simultaneously holding the Bar to account for unverified reliance on similar technology.¹⁷ This literature has usefully highlighted that the accountability gap runs in two directions, with litigants and advocates submitting unverified AI-generated material, and judges and their researchers doing the same in the course of preparing orders,¹⁸ but it has engaged only lightly with the specifically professional-conduct dimension of the problem: what changes, if anything, the reconstructed duty demands of the ordinary practice of legal research within a chamber or law firm.

A third and more analytically developed strand has begun to examine the Supreme Court's draft Regulations for Use of Artificial Intelligence in Courts, 2026 as a governance instrument. Commentary here has emphasised the draft's three organising principles, that AI may assist but

¹⁶ See generally Bansal, *supra* note 2; *AI-Hallucinated Case Law*, *supra* note 4.

¹⁷ Bansal, *supra* note 2 (describing the Supreme Court's SUPACE, SUVAS, TERES and LegRAA tools).

¹⁸ Aakriti Bansal, *Supreme Court Asks Bar Council to Form AI Expert Panel After Trial Court Cites Fake Judgments*, MEDIANAMA (May 7, 2026), <https://www.medianama.com/2026/05/223-supreme-court-seeks-ai-panel-over-fake-judgments/>.

never adjudicate, that its use in filings must be disclosed, and that a permanent apex body will approve and supervise the AI tools deployed across courts and tribunals,¹⁹ and has correctly identified disclosure, rather than verification, as the draft's central mechanism. This is, in this paper's assessment, precisely the gap the existing literature has not yet addressed in depth: a disclosure obligation tells a court that AI was used, but it does not by itself tell the court, or indeed the advocate submitting the material, whether the specific authorities relied upon are genuine. A Bar and Bench column written around the February 2026 India AI Impact Summit has observed, correctly in this paper's view, that the Bar Council of India's Standards of Professional Conduct and Etiquette, "framed long before large language models existed, are entirely silent on AI delegation",²⁰ a diagnosis this paper adopts as its starting premise but seeks to carry further, into a proposed reconstruction of the duty rather than a description of its absence.

Finally, a smaller body of comparative commentary has drawn attention to the United States' more developed experience with judicial sanctions for AI-fabricated citations, beginning with *Mata v. Avianca, Inc.* in 2023 and continuing through a series of escalating sanctions in subsequent years. In *Couvrette v. Wisnovsky*, for example, a federal magistrate judge in the District of Oregon imposed a \$15,500 monetary sanction on lead counsel, together with an award of the opposing party's attorneys' fees, after three briefs had cited fifteen non-existent cases and eight fabricated quotations; one litigation tracker reports that a later fee-apportionment order brought the combined amount against lead and local counsel to \$110,204.38.²¹ This literature has tended to treat the American experience as a cautionary parallel rather than as a source of transplantable regulatory technique, and it has not, so far as this paper has been able to establish, engaged closely with the specific question of how an India-specific, tiered verification standard might be constructed by reference to the country's own

¹⁹ *Supreme Court's AI Rules Explained: Why India Wants Artificial Intelligence in Courts, But Not in Judicial Decisions*, OUTLOOK INDIA (July 13, 2026).

²⁰ Rahul, *supra* note 7.

²¹ *Couvrette v. Wisnovsky*, No. 1:21-cv-00157-CL, 2025 WL 4109655 (D. Or. Dec. 12, 2025) (Clarke, Mag. J.) (imposing a \$15,500 monetary sanction on lead counsel, calculated at \$500 per non-existent case and \$1,000 per fabricated quotation, striking the offending briefs, dismissing the offending plaintiffs' claims with prejudice, awarding the defendants their attorneys' fees and directing notice to the Oregon State Bar). A follow-on fee-apportionment order of Mar. 23, 2026 (ECF No. 225) is reported by the Legal AI Governance case tracker to have brought the combined amount against lead counsel and local counsel to \$110,204.38. See Mark J. Fucile, *Parade of Horribles: Federal Court in Oregon Surveys Sanctions for AI Fake Citations*, NWSIDEBAR (Wash. State Bar Ass'n, Mar. 2, 2026), <https://nwsidebar.wsba.org/2026/03/02/parade-of-horribles-federal-court-in-oregon-surveys-sanctions-for-ai-fake-citations/>; Nadia Dahab, *Oregon's AI Sanctions Framework Is Growing, and So Are the Stakes*, SUGERMAN DAHAB (Apr. 11, 2026), <https://sugermandahab.com/uncategorized/oregons-ai-sanctions-framework-is-growing-and-so-are-the-stakes/>; *Couvrette v. Wisnovsky*, LEGAL AI GOVERNANCE (case tracker). Some secondary commentary, including at least one Indian-authored guide, has reported the number of fabricated citations as 23 rather than 15; 23 is the sum of the fifteen non-existent cases and the eight fabricated quotations, a split confirmed by the arithmetic of the base sanction ($\$500 \times 15 + \$1,000 \times 8 = \$15,500$).

citation ecosystem: official law reports, SCC Online, Manupatra, and the increasingly important e-SCR portal maintained by the Supreme Court itself. It is this gap, connecting the doctrinal recognition of a duty in *Pooja Ramesh Singh*, the disclosure-centred draft Regulations, and the profession's own citation practices, that this paper seeks to address.

VII. RESEARCH AND ANALYSIS

A. Anatomy of an Indian AI hallucination: a documentary survey

It is necessary to establish, with some precision, what happened in the reported Indian incidents before attempting to prescribe a remedy, because the incidents differ from one another in ways that matter for the reconstruction attempted in this paper. The earliest documented episode, the *Buckeye Trust* matter before the Income Tax Appellate Tribunal's Bengaluru Bench, arose in an order dated 30 December 2024 deciding a trust-taxation dispute of approximately Rs 669 crore; the order cited, in its own reasoning, three Supreme Court judgments and one Madras High Court ruling that, as reported, did not exist, and reporting at the time indicated that the Revenue's representative had used a generative AI tool without independently verifying its output.²² The Tribunal recalled its own order on 7 January 2025, citing "inadvertent errors", reportedly in exercise of its power under Section 254(2) of the Income-tax Act, 1961 to rectify a mistake apparent from the record; that mechanism, however, presupposes that the mistake is discovered at all.

The Andhra Pradesh episode that produced *Gummadi Usha Rani v. Sure Mallikarjuna Rao* illustrates a different and, in some respects, more troubling failure mode: here, the fabrication was not merely undetected but was detected and then tolerated. A trial court dismissed objections to an advocate commissioner's report on 19 August 2025 by relying on four non-existent Supreme Court judgments; on revision, the Andhra Pradesh High Court itself recognised that the citations were AI-generated fabrications, yet proceeded to affirm the trial court's decision on merits, reasoning, in substance, that the fabrication did not vitiate an outcome that was otherwise sustainable.²³ It was precisely this reasoning that the Supreme Court later found unacceptable in principle, holding in *Pooja Ramesh Singh* that a decision resting even partly on fabricated material is no decision in the eyes of the law, irrespective of whether such material had a direct or indirect bearing on the decision-making.²⁴

²² *Buckeye Trust*, ITA No. 1051/Bang/2024, ¶ 13; *AI-Hallucinated Case Law*, *supra* note 4; *ITAT Bangalore Recalls Its Order*, *supra* note 4.

²³ *Gummadi Usha Rani*, 2026 SCC OnLine SC 341, ¶¶ 2–4; Bansal, *supra* note 2.

²⁴ *Pooja Ramesh Singh*, 2026 INSC 668, ¶¶ 7, 17; see also *Pooja Ramesh Singh v. J&K Bank (2026): SC Tears Apart NCLT Order That Relied on AI-Hallucinated Judgments*, AASHAYEIN JUDICIARY (July 13, 2026), <https://aashayeinjudiciary.com/blog/pooja-ramesh-singh-v-jk-bank-2026>.

The *Pooja Ramesh Singh* matter itself supplies the most instructive failure mode for present purposes, because the fabrication survived not one but two tiers of adjudication. The National Company Law Tribunal admitted a Section 7 application under the Insolvency and Bankruptcy Code, 2016 on 28 August 2024, in a proceeding where the fabricated authorities apparently first appeared; the National Company Law Appellate Tribunal, on appeal, affirmed that order on 11 September 2025 without detecting the fabrication, more than a year after the original order and with the benefit of a full appellate hearing.²⁵ The Supreme Court's own judgment records that an independent examination of the six authorities relied upon by the NCLT, as reproduced by the NCLAT, disclosed fabrication of different kinds: three citations did not exist at all, two were correct citations attached to paragraphs that did not exist, and the citation attributed to *State Bank of India v. M/s Shree Ram Urban Infrastructure Ltd.*, 2020 SCC OnLine SC 341, was found to be a wrong citation of an existing reported judgment, in fact belonging to an entirely different case, accompanied by a non-existent paragraph.²⁶ Critically for the argument developed in Part VII.E, the Supreme Court recorded an affidavit filed by the respondent bank stating that the authorities had not been cited by its counsel at the bar and that the adjudicating authority had obtained them "through its own research", a strong indication that at least part of the accountability gap in this matter ran through the tribunal's own process rather than through a party's submissions.²⁷

Two further threads deserve mention because they show the pattern is not confined to any single tier or type of forum. The Karnataka High Court's direction in March 2025 for further action against a City Civil Court judge who had relied on two non-existent Supreme Court decisions in a commercial dispute (although that order does not itself attribute the citations to AI) is, on the evidence reviewed for this paper, among the few instances in which disciplinary or supervisory scrutiny has been directed at a judicial officer rather than at counsel,²⁸ while the remarks from the Bench of the Chief Justice of India in February 2026 about a petition citing a wholly fictitious case styled "*Mercy v. Mankind*", and about a series of similarly fabricated judgments cited before Justice Dipankar Datta, indicate that the problem has by now been observed at the apex of the system itself and is not confined to subordinate courts or specialised tribunals.²⁹ Taken together, these episodes span at least four distinct institutional settings,

²⁵ *Pooja Ramesh Singh*, 2026 INSC 668, ¶¶ 12–13.

²⁶ *Pooja Ramesh Singh*, 2026 INSC 668, ¶ 15 & n.6 (giving the correct cause title of 2020 SCC OnLine SC 341 as *M. Subramaniam v. S. Janaki*, (2020) 16 SCC 728); see also *Fake AI Cases 'Entered' an NCLT Insolvency Order: The Supreme Court Quashed It. But Who Is Accountable?*, THE WIRE (July 3, 2026).

²⁷ *Pooja Ramesh Singh*, 2026 INSC 668, ¶ 16; see also *Fake AI Cases 'Entered' an NCLT Insolvency Order*, *supra* note 26.

²⁸ *Sammaan Capital*, Civil Revision Petition No. 49 of 2025, ¶¶ 10, 12; Bansal, *supra* note 2.

²⁹ Rahul, *supra* note 7; Bansal, *supra* note 2.

namely a specialised tax tribunal, a civil trial court, a High Court exercising revisional jurisdiction, and the country's insolvency adjudicatory hierarchy, and at least three distinct failure points: fabrication going undetected entirely until challenged by a party; fabrication detected but treated as harmless because the outcome was otherwise defensible; and fabrication surviving a full tier of appellate review. Any reconstructed duty of verification must be able to address all three, not just the first.

B. The doctrinal response: from *Gummadi Usha Rani* to *Pooja Ramesh Singh*

The Supreme Court's engagement with this pattern proceeded in two distinct steps, and the difference between them is instructive. In *Gummadi Usha Rani*, the Court's order of 27 February 2026 was explicitly preliminary: it took cognisance of the trial court's reliance on fabricated material, declared that such reliance "would be a misconduct and legal consequence shall follow", but stopped short of laying down an operative standard, choosing instead to issue notice to the Attorney General, the Solicitor General, and the Bar Council of India, and to appoint senior counsel to assist the Court in a fuller examination of the issue.³⁰ The practical difficulty of enforcing any verification duty without a workable mechanism for discharging it surfaced later, when the same Bench, deciding *Pooja Ramesh Singh*, observed that courts and tribunals "implicitly trust lawyers when referring to precedents cited before them" and asked the reader to "[i]magine the hardship of a situation in which the Court must verify the authenticity of each judgment cited by an advocate".³¹

By the time the Court decided *Pooja Ramesh Singh* some four months later, it was prepared to state the underlying principle in unqualified terms. The judgment holds three things of direct relevance to this paper. First, a judicial or quasi-judicial decision resting even partly on fabricated precedent is not merely erroneous but is, in the Court's words, "no decision at all", a formulation that treats the defect as foundational rather than as an ordinary ground of appeal.³² Second, advocates carry an affirmative professional duty to verify the authenticity of legal authorities before placing them before a court or tribunal, and the failure to do so, not merely the act of deliberate fabrication but the citation of such judgments without verification as such, amounts to misconduct.³³ Third, the responsibility for enforcing this duty in the first instance

³⁰ *Gummadi Usha Rani*, 2026 SCC OnLine SC 341, ¶¶ 7–9; see also *Judges Citing Fake AI-Generated Case Laws Amounts to Misconduct: Supreme Court*, MONEYLIFE (Mar. 2, 2026).

³¹ *Pooja Ramesh Singh*, 2026 INSC 668, ¶ 16; see also 'AI Should Aid, Not Replace Human Reasoning': Supreme Court on Fake Case Law, LAWBEAT (July 30, 2026), <https://lawbeat.in/supreme-court-judgments/ai-should-aid-not-replace-human-reasoning-supreme-court-on-fake-case-law-1617273>.

³² *Pooja Ramesh Singh*, 2026 INSC 668, ¶¶ 7, 17.

³³ *Id.* ¶ 7; see also *SC Orders Strict Action Against Lawyers Citing Unverified AI-Generated Judgments*, STORYBOARD18 (July 2, 2026).

was handed to the Bar Council of India, which was directed to constitute a committee, take up the issue with what the Court called “utmost seriousness”, and prescribe both a guiding principle and the disciplinary action that will follow a violation.³⁴

What the judgment does not do, and this is the premise on which the remainder of this paper proceeds, is specify the content of the verification duty it recognises. It does not say whether an advocate satisfies the duty by confirming a citation against a single reported source, or whether cross-verification against more than one source is required for citations drawn from an unfamiliar or lesser-used database; it does not say whether the duty differs according to the seniority of the advocate concerned, or according to whether the advocate personally conducted the research or relied on a junior colleague’s output; and it does not say how a court, faced with a citation that later turns out to be fabricated, should go about establishing whether verification was in fact attempted, as distinct from simply not performed. These are not gaps of drafting; they are gaps inherent in deciding a case on its own facts, and they are precisely the gaps the Bar Council of India’s committee, and in the interim this paper, must attempt to fill.

C. Comparative perspective: the United States’ experience with fabricated AI citations

India’s confrontation with AI-hallucinated authorities is not unique, and it is useful, if only as a point of contrast, to note briefly how the problem first surfaced in the United States. The 2023 decision in *Mata v. Avianca, Inc.* (S.D.N.Y.), in which attorneys filed a submission citing several judicial decisions that a generative AI tool had entirely fabricated, complete with invented quotations and internal citations, is widely regarded as the first prominent instance of the phenomenon reaching a court. The presiding judge’s sanctions order stated a principle that this paper considers directly transferable to the Indian context: there is “nothing inherently improper about using a reliable artificial intelligence tool for assistance”, but existing rules “impose a gatekeeping role on attorneys to ensure the accuracy of their filings”, a responsibility commentators have described as non-delegable, whatever the tool used to prepare the filing.³⁵

The American experience since 2023 has been one of escalating and increasingly severe sanctions, including the Oregon sanction in *Couvrette v. Wisnovsky* discussed in Part VI above,³⁶ and it has developed largely through individual courts issuing their own standing orders requiring disclosure of AI use and certification of verification, rather than through a single nationwide regulatory instrument. This court-by-court approach has the advantage of flexibility,

³⁴ *Pooja Ramesh Singh*, 2026 INSC 668, ¶ 9.

³⁵ *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443, 448 (S.D.N.Y. 2023); see also *Phantom Precedents: The Rise of AI-Generated Case Law in Indian Courts*, LIVE LAW (Mar. 17, 2026).

³⁶ *Couvrette*, 2025 WL 4109655.

but the disadvantage of fragmentation: a lawyer practising across multiple federal districts may face different disclosure and verification requirements in each. India's institutional response, by contrast, has so far proceeded through the Supreme Court's own draft Regulations and the Bar Council of India's mandated committee, both operating at a national level; this centralised structure creates the possibility, discussed in Part VIII, of a single, consistently applied verification standard, provided the substantive content of that standard is worked out with the specificity that the American case-by-case approach has, out of necessity, supplied only incrementally.

D. The existing architecture of the advocate's duty: the Advocates Act, 1961 and the BCI Standards

Before proposing a reconstruction, it is necessary to establish what the advocate's duty of care and candour already comprises under Indian law, since any new verification standard must be anchored in this existing framework rather than floating free of it. Section 35 of the Advocates Act, 1961 requires a State Bar Council which has reason to believe that an advocate on its roll has been guilty of "professional or other misconduct" to refer the case to its disciplinary committee, which may reprimand, suspend or remove the advocate, without the Act itself defining that term; Section 36 confers a corresponding jurisdiction on the Bar Council of India over advocates whose names are not on any State roll, and allows its disciplinary committee, of its own motion or on a report or application, to withdraw to itself proceedings pending before a State Bar Council's disciplinary committee.³⁷ The undefined character of "misconduct" is deliberate: it allows disciplinary committees and, on appeal, the courts to develop the concept incrementally, and it is this open texture that *Pooja Ramesh Singh* has now used to bring unverified reliance on AI-generated authorities within its scope, without any need for Parliament to amend the Act itself.³⁸

The Bar Council of India's Standards of Professional Conduct and Etiquette, framed under the rule-making power conferred by Section 49(1)(c) of the Act, supply the more granular content of an advocate's duties. Those owed to the court require an advocate to conduct himself with dignity and self-respect, to maintain a respectful attitude towards the court, not to influence its decision by any illegal or improper means, and to restrain his client from sharp or unfair practices, refusing to act as a "mere mouth-piece" of the client; towards the client, the advocate

³⁷ The Advocates Act, 1961, No. 25, Acts of Parliament, 1961, §§ 35–36 (India); *see also Supreme Court Releases Draft 'Regulations for Use of Artificial Intelligence (AI) in Courts, 2026' | Here's What That Means*, INDIAN LEGAL TECH NETWORK (June 9, 2026) (noting that false citations and authorities expose advocates to disciplinary proceedings under these provisions).

³⁸ *Pooja Ramesh Singh*, 2026 INSC 668, ¶¶ 7–9.

must uphold the client's interests "by all fair and honourable means". Read together, these duties are capable of encompassing an obligation to verify legal authorities: a fabricated citation, once relied upon, is by definition a misrepresentation of the law to the court, whether or not the advocate who relied on it knew it to be false, and it is not a fair or honourable means of advancing a client's case. What the Standards do not supply, because they long predate the technology, is any provision specifically addressing generative artificial intelligence, or indeed any express rule on the verification of authorities, and it is this silence that the literature reviewed in Part VI has rightly identified as the central regulatory gap.³⁹ The general duties, in other words, already reach the conduct in question as a matter of principle; what is missing is a specific articulation of what diligence requires in the particular context of AI-assisted research, the equivalent, in a different professional register, of the difference between a general duty of care in tort and a specific, codified standard of practice in a regulated industry.

E. Gaps in the reconstructed duty: attribution, the AOR/junior divide, and the court's own tools

Three specific gaps in the current position deserve separate attention, because each requires a different kind of institutional response, and none is resolved by the general language of *Pooja Ramesh Singh* standing alone.

The first is a gap of attribution within a legal team. Indian litigation practice, particularly at the level of the Supreme Court and the High Courts, typically involves a division of labour between an Advocate-on-Record who settles and formally signs pleadings, senior or briefing counsel who argue the matter, and junior associates, paralegals, or, increasingly, AI-assisted research platforms that generate the first draft of the legal research underlying a submission. The *Pooja Ramesh Singh* judgment speaks of the duty resting on "an advocate" without differentiating between these roles, yet the practical reality of how a fabricated citation enters a filing will often depend heavily on which of these actors introduced it and which had a realistic opportunity to catch it before the document was filed. A verification standard that treats every advocate whose name appears on a pleading as equally culpable for a fabrication introduced by a junior researcher three tiers removed risks either over-deterrence, with senior counsel refusing to sign pleadings they have not personally re-verified line by line, which is both impractical and an inefficient use of scarce senior time, or, if enforcement in practice proves lenient toward senior counsel precisely because of this impracticality, under-deterrence at the point where the fabrication actually originates.

³⁹ Bar Council of India Rules, pt. VI, ch. II, § I, rr. 1–4 & § II, r. 15 (Standards of Professional Conduct and Etiquette, framed under the Advocates Act, 1961, § 49(1)(c)); *see also* Rahul, *supra* note 7.

The second gap concerns the court's own use of AI tools, a point the literature reviewed in Part VI has raised but which this paper considers under-examined as a matter of professional and institutional responsibility specifically. The Supreme Court itself operates a Portal for Assistance in Court Efficiency for research assistance, a translation tool, a transcription tool, and a generative legal research assistant,⁴⁰ and the *Pooja Ramesh Singh* record itself indicates, through the respondent bank's affidavit discussed in Part VII.A above, that the fabricated authorities in that matter originated in the Tribunal's research process rather than in either party's submissions. A duty of verification framed exclusively as an obligation running from advocate to court addresses only half of the accountability structure this record discloses; the other half, namely what verification obligations, if any, attach to a court's or tribunal's own research staff and AI-assisted tools before a citation is incorporated into a judgment, remains almost entirely unaddressed in the current doctrinal response, notwithstanding the Court's own acknowledgment, in *Pooja Ramesh Singh* itself, that it is "a serious lapse if a judge relies on such a fake or hallucinated AI-generated material as precedents".⁴¹

The third gap concerns the standard's temporal and forum-specific calibration. The incidents surveyed in Part VII.A span a trial-court order on an interlocutory objection, a considered appellate order of the NCLAT delivered more than a year after the original proceeding, and a specialised tax tribunal's final order in a matter involving several hundred crores of rupees. A verification standard that does not distinguish between these settings risks being either too onerous for genuinely time-pressed interlocutory work or too lax for high-value, considered submissions where the opportunity and the stakes both justify a materially higher degree of care. Part VII.G below attempts to address this by proposing a standard tiered along precisely this axis.

F. The draft Regulations for Use of Artificial Intelligence in Courts, 2026: an institutional first step

The Supreme Court's AI Committee released a draft titled the Regulations for Use of Artificial Intelligence (AI) in Courts, 2026 on 3 June 2026, with the comment period initially set to close on 20 June 2026 and subsequently extended to 15 July 2026.⁴² As of August 2026, the framework had not been finalised or notified, and readers should confirm its status, whether notified, revised, or still pending, before relying on this description. The draft is organised

⁴⁰ Bansal, *supra* note 2 (describing SUPACE, SUVAS, TERES and LegRAA).

⁴¹ *Pooja Ramesh Singh*, 2026 INSC 668, ¶ 7; see also 'AI Should Aid, Not Replace Human Reasoning', *supra* note 31.

⁴² *Supreme Court Draft AI Regulations 2026: New Rules for Artificial Intelligence in Indian Courts*, SANSA LEGAL (June 5, 2026); *Supreme Court's AI Rules Explained*, *supra* note 19.

around three principal commitments: that artificial intelligence may assist judicial and administrative functions but may never adjudicate; that any use of AI tools by lawyers or litigants in preparing pleadings, submissions, or evidence must be disclosed; and that a permanent apex body situated at the Supreme Court will approve and supervise the AI tools deployed across the judicial system, extending to the High Courts and subordinate courts as well.⁴³

This is, in this paper's assessment, a necessary but not sufficient response to the problem examined here. Mandatory disclosure is valuable because it converts a previously invisible practice, the use of a generative tool to prepare a submission, into a fact of record that a court or opposing counsel can, in principle, interrogate; commentary on the draft has rightly emphasised that this closes off the possibility of an advocate later disclaiming responsibility on the ground that "a machine produced it", since, as one commentary on the draft put it, the machine "neither signs the pleading nor owes duties to the court", and the advocate does.⁴⁴ But disclosure answers the question "was AI used?", not the question "was what it produced actually checked, and against what?", which is precisely the question *Pooja Ramesh Singh* leaves open. The draft does provide that, where AI-generated output is used in any court, "reasonable care shall be taken to verify the accuracy of such output before the same is utilised", and it allows a court to require disclosure of "the steps taken to verify the accuracy of any AI-generated content",⁴⁵ but it leaves that care as undefined as the judgment does. A litigant could, in principle, fully disclose that a submission was prepared with AI assistance and still have failed to verify a single citation within it; disclosure, without an accompanying and independently specified verification obligation, risks becoming a formality that is satisfied on the face of a filing while the underlying mischief, unverified and potentially fabricated authority reaching a court, continues unaddressed.

G. Synthesis: toward a tiered, source-anchored standard of verification

Bringing together the doctrinal recognition of a duty in *Pooja Ramesh Singh*, the existing but generalised professional-conduct architecture under the Advocates Act, the disclosure-centred draft Regulations, and the comparative experience under *Mata v. Avianca*, this paper suggests

⁴³ *Supreme Court Draft AI Regulations 2026*, *supra* note 42; Banerjee & Debjani M, *supra* note 15 (describing the declaration of AI use required by Regulation 43, the bar on autonomous adjudication and the Apex Body).

⁴⁴ *Supreme Court Releases Draft 'Regulations for Use of Artificial Intelligence (AI) in Courts, 2026'*, *supra* note 37.

⁴⁵ *Draft Regulations for Use of Artificial Intelligence (AI) in Courts, 2026*, *supra* note 15, regs. 8(3), 43(4), 43(6); *see also* Banerjee & Debjani M, *supra* note 15 (reporting that anyone using AI "has to exercise reasonable diligence in verifying the accuracy of the material generated").

that a workable, reconstructed duty of verification should rest on three linked propositions rather than a single undifferentiated command to “verify.”

First, verification should be source-anchored rather than merely asserted. An advocate should be taken to have discharged the duty only where a cited authority has been checked against an authoritative source, whether an official law report, a recognised commercial database such as SCC Online or Manupatra, or the Supreme Court’s own e-SCR portal, and not merely against the output of the generative tool that produced the citation in the first place. Verifying a citation generated by one model by asking the same or a different generative model to confirm it is not verification in any meaningful sense; it is repetition of the same unreliable process, and any reconstructed standard should say so explicitly.

Second, the standard should be tied to the stage and stakes of the proceeding, in recognition of the temporal gap identified in Part VII.E. A genuinely urgent interlocutory application prepared under severe time pressure cannot realistically be held to the same standard of exhaustive, multiply cross-checked verification as a considered written submission filed months into a commercial dispute or a final hearing before a specialised tribunal; what can reasonably be required in both cases, however, is that at least one authoritative-source check be performed for every citation relied upon, with a higher standard of cross-verification reserved for matters of higher value or complexity, mirroring the risk-based calibration the draft Regulations already contemplate for AI tools generally.⁴⁶

Third, responsibility should be apportioned within a legal team according to actual role and realistic opportunity to detect the defect, rather than falling uniformly on whichever advocate’s name happens to appear on the final pleading. An Advocate-on-Record who settles a pleading prepared by a junior colleague should bear a duty to confirm, at a minimum, that a stated verification process was in fact followed, a professional analogue to the certification requirements that have emerged in American practice following *Mata v. Avianca*, rather than a duty to personally re-verify every authority from first principles, provided a documented process for doing so exists within the chamber or firm. This calibration is not a dilution of the *Pooja Ramesh Singh* standard; it is what gives that standard the operational content needed for it to be taught to a fresh recruit, audited by a disciplinary committee, and enforced with the

⁴⁶ *Draft Regulations for Use of Artificial Intelligence (AI) in Courts, 2026*, *supra* note 15, reg. 12 (use of AI to be “proportionate to the nature, complexity and risk profile of the relevant task”, with higher-risk applications “subject to correspondingly heightened safeguards”); *see also Supreme Court Releases Draft ‘Regulations for Use of Artificial Intelligence (AI) in Courts, 2026’*, *supra* note 37 (describing a risk-based approach under which tasks involving higher risk attract heightened safeguards).

consistency the Bar Council of India's forthcoming committee has now been tasked with supplying.

VIII. SUGGESTIONS AND RECOMMENDATIONS

The Supreme Court has, through the direction issued in *Pooja Ramesh Singh*, already created the institutional occasion for reform: the Bar Council of India is now under a specific mandate to constitute a committee and prescribe guiding principles and disciplinary consequences. The following recommendations are offered as a contribution to that process and are addressed as much to the Bar Council of India and the Supreme Court's AI Committee as to individual practitioners.

A. Codify a source-anchored verification protocol in the BCI Standards

The Bar Council of India should amend the Standards of Professional Conduct and Etiquette to include an explicit provision requiring that any judicial or quasi-judicial authority cited in a pleading, written submission, or oral argument be confirmed against an authoritative source (an official law report, a recognised legal database, or the Supreme Court's e-SCR portal) before it is relied upon, and that verification against the output of the same or another generative AI tool does not, without more, satisfy this requirement. Codifying this as an express rule, rather than leaving it to be inferred from the general duty of diligence, gives disciplinary committees a specific yardstick against which conduct can be measured, and gives the profession fair notice of exactly what is expected.

B. Introduce a tiered standard calibrated to stage and stakes

Consistent with the analysis in Part VII.G, the Bar Council's committee should resist the temptation to prescribe a single, undifferentiated verification standard for every filing regardless of urgency or value. A workable framework would distinguish, at minimum, between urgent interlocutory filings (requiring at least one authoritative-source check per citation), considered written submissions in matters of ordinary value (requiring a documented single-source check with spot cross-verification), and submissions in high-value or high-complexity matters, including insolvency, taxation, and constitutional proceedings (requiring cross-verification against at least two independent authoritative sources for every citation relied upon).

C. Require a verification certificate for AI-assisted filings

Building on the disclosure obligation already contemplated by the draft Regulations for Use of Artificial Intelligence in Courts, 2026, this paper recommends that any pleading or written submission prepared with the assistance of a generative AI tool carry a short, signed certificate

confirming that every judicial authority cited therein has been independently verified against an authoritative source, identifying which Advocate-on-Record or counsel takes responsibility for that verification. This would give the disclosure obligation the substantive complement it currently lacks, converting “AI was used” into “AI was used, and here is who checked its output and how.”

D. Extend verification obligations to the court’s own research process

Given the indication, discussed in Part VII.E, that the fabricated authorities in the *Pooja Ramesh Singh* matter originated in a tribunal’s own research process rather than in either party’s submissions, this paper recommends that the same source-anchored verification principle be extended, through appropriate administrative circulars, to law clerks, research associates, and AI-assisted tools used by courts and tribunals in preparing draft orders and judgments. A verification duty that binds only the Bar while leaving the Bench’s own research process unaddressed responds to only part of the problem the reported incidents disclose.

E. Build citation-verification support into e-filing infrastructure

Finally, given that the incidents surveyed in Part VII.A demonstrate that after-the-fact judicial detection is an unreliable safeguard (a fabricated authority survived an entire tier of appeal in the very matter that produced *Pooja Ramesh Singh*), this paper recommends that the National Informatics Centre and the e-Committee of the Supreme Court explore integrating an automated citation cross-check into existing e-filing portals, flagging, at the point of filing, any cited authority that cannot be located in a recognised law report or database. Such a tool would not replace an advocate’s own duty of verification, but it would supply the kind of front-line, systemic safeguard capable of catching a fabricated citation before it reaches a judge, rather than relying solely on an opposing party’s diligence or a court’s own, necessarily variable, attentiveness.

IX. CONCLUSION

Between the Income Tax Appellate Tribunal’s order of December 2024 in the *Buckeye Trust* matter, recalled within days, and the Supreme Court’s decision in *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.* in July 2026, Indian adjudication confronted, in the space of some nineteen months, a recurring failure that touched a specialised tax tribunal, a civil trial court, a High Court exercising revisional jurisdiction, and the country’s insolvency adjudicatory hierarchy across two tiers. In each instance, a plausible-looking but entirely fictitious judicial authority was placed before, and in most instances relied upon by, the forum in question, without the safeguard that ought, in principle, to have caught it: an advocate’s or a court’s own

verification of the authority against the actual body of reported law. The Supreme Court's response in *Pooja Ramesh Singh*, declaring that such reliance amounts to professional misconduct, that a decision founded on fabricated precedent is no decision at all, and that the Bar Council of India must now prescribe both principles and consequences, marks the point at which Indian law formally recognised the problem as one of professional responsibility rather than mere technological mishap.

This paper has argued that recognition of the duty is different from its reconstruction. A duty to "verify" that does not specify against which sources verification must occur, how it should be documented, how it should be apportioned within a legal team, or how it should scale with the urgency and stakes of a given proceeding, risks becoming either unenforceable in practice or arbitrarily enforced against whichever advocate happens to be easiest to identify after the fact. The three-part standard proposed in this paper, source-anchored, tiered by stage and stakes, and apportioned by actual role, is offered as one way of giving the *Pooja Ramesh Singh* principle the operational precision it will need if it is to function as a genuine professional standard rather than a slogan invoked after the next fabrication is discovered. Whether the Bar Council of India's committee, the Supreme Court's AI Committee, or Parliament itself ultimately supplies that precision, the incidents surveyed in this paper make one thing clear: the cost of leaving the duty undefined is not borne only by the advocate who fails to discharge it, but by every litigant whose case is decided, even in part, on the strength of a precedent that was never anything more than a plausible arrangement of words.
