



INTERNATIONAL JOURNAL OF HUMAN RIGHTS LAW REVIEW

An International Open Access Double Blind Peer Reviewed, Referred Journal

Volume 5 | Issue 5 | 2026

Art. 15

When Law Meets Generative AI: Preserving Judicial Integrity through Human Verification

Saksham Bharadwaj

*Law Student, BBA.LL.B., 4th Year,
Himachal Pradesh National Law University, Shimla*

Recommended Citation

Saksham Bhardwaj, *When Law Meets Generative AI: Preserving Judicial Integrity through Human Verification*, 5 IJHRLR 159-168 (2026).
Available at www.ijhrlr.in/current-issues/.

This Article is brought to you for free and open access by the International Journal of Human Rights Law Review by an authorized Lex Assisto & Co. administrator.

For more information,
please contact submission.ijhrlr@gmail.com

When Law Meets Generative AI: Preserving Judicial Integrity through Human Verification

Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd. & Anr.,
2026 INSC 668

1. Introduction

Artificial intelligence has moved from being a peripheral technological tool to an increasingly consequential component of professional decision-making. The legal system presents a particularly sensitive setting because judicial reasoning depends upon the accurate identification, interpretation and application of authoritative sources. The Supreme Court of India's decision in *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.* therefore marks an important early intervention in the relationship between artificial intelligence and adjudication.¹

The appeal arose from insolvency proceedings in which the National Company Law Tribunal, Mumbai (NCLT) relied upon several purported judicial precedents which were later found to be non-existent, incorrectly cited, or accompanied by paragraphs that did not appear in the cited judgments. The National Company Law Appellate Tribunal (NCLAT) subsequently affirmed the NCLT's conclusion without detecting the problem. The Supreme Court, on 2 July 2026, set aside both orders and directed a fresh consideration of the insolvency application. More significantly, the Court declared a zero-tolerance approach towards the production, citation or reliance upon AI-generated legal authorities without verification and directed the Bar Council of India to constitute a committee to formulate guiding principles and disciplinary consequences.

This commentary argues that the judgment is substantially correct in protecting adjudicative integrity, but that its strongest proposition—that even an “iota” of hallucinated material renders a decision legally unsustainable—requires careful future clarification. A principled regulatory framework should distinguish deliberate fabrication, negligent reliance, and an immaterial technological error. The judgment is therefore best understood not as an anti-AI decision, but as an early framework for accountable human use of AI in legal decision-making.

¹ *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd. & Anr.*, 2026 INSC 668 (2 July 2026).

2. *Facts of the Case*

Pooja Ramesh Singh was a suspended director of Essel Infraprojects Ltd. (EIL), which was the corporate debtor and corporate guarantor of Pan India Utilities Distribution Company Ltd. (PIUDCL). PIUDCL had obtained loan facilities from Jammu and Kashmir Bank Ltd., and EIL had executed a corporate guarantee securing those facilities. Following continuing defaults, the lender classified the accounts as non-performing assets and initiated proceedings under section 7 of the Insolvency and Bankruptcy Code, 2016 (IBC) before the NCLT, Mumbai.²

The NCLT admitted the section 7 application on 28 August 2024, appointed an Interim Resolution Professional and imposed the statutory moratorium under section 14 of the IBC.³ The appellant challenged the order before the NCLAT. Among other grounds, she contended that liabilities had been affected by a scheme of demerger and subsequent amalgamation, and that a renewed sanction letter did not contain the earlier guarantee. The NCLAT rejected the challenge on 11 September 2025 and held, inter alia, that the guarantee survived the internal restructuring because the guarantee deed itself protected it against absorption or amalgamation.⁴

The NCLAT referred to six purported Supreme Court decisions in support of the conclusion. Before the Supreme Court, senior counsel for the appellant pointed out that the authorities were either fictitious or contained passages that could not be traced to the reported judgments. The respondent's bank subsequently filed an affidavit stating that the disputed judgments had not been cited by its counsel and had instead been located through the NCLT's own research.

The Supreme Court independently examined the authorities. It found that some citations were entirely non-existent, while others corresponded to genuine cases but were accompanied by non-existent paragraphs. For example, the citation given as *State Bank of India v. M/s Shree Ram Urban Infrastructure Ltd.*, 2020 SCC OnLine SC 341 was an incorrect citation for an existing judgment⁵; *Everest Kento Cylinders Ltd. v. Union of India* had a correct citation but a non-existent paragraph; *ICICI Bank Ltd. v. Urban Infrastructure Real Estate Ltd.* and *vs. Dempo & Co. Ltd. v. Reliance Communications*

² The Insolvency and Bankruptcy Code, 2016, ss. 7, 14.

³ *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd. & Anr.*, 2026 INSC 668, paras 10–13.

⁴ *Ibid.*, paras 15–17.

⁵ *M. Subramaniam v. S. Janaki*, (2020) 16 SCC 728; 2020 SCC OnLine SC 341.

Ltd. were non-existent citations; and other cited authorities similarly contained fabricated passages.

The Supreme Court consequently held that the adjudicatory process had been tainted by reliance upon unreal material. It set aside the NCLT and NCLAT orders, restored the section 7 application to its original number, directed the NCLT to decide it afresh, and expressly clarified that it had not adjudicated the underlying merits.

3. Issues before the Court

The judgment raised interconnected questions concerning adjudication, professional responsibility and artificial intelligence

- Whether a judicial or quasi-judicial decision based upon non-existent or hallucinated legal authorities can be sustained in law.
- Whether reliance upon AI-generated fake precedents without verification constitutes professional misconduct or a serious judicial lapse.
- Whether the proper response is to invalidate the entire decision even where the fabricated authority may not have materially altered the ultimate outcome.
- What institutional safeguards are necessary to preserve the integrity of legal decision-making while permitting legitimate and useful applications of AI.
- Whether the Bar Council of India should develop specific professional standards governing advocates' use of generative AI in litigation and legal research.

These issues mattered beyond the immediate insolvency dispute because courts increasingly depend upon digital databases and technological research tools. If generative AI is permitted to generate plausible but false authorities without human verification, the reliability of precedent – the foundation upon which judicial consistency operates – is directly threatened.

4. Arguments of the Parties

The appellant challenged the underlying NCLT and NCLAT conclusions and, most importantly for the present appeal, demonstrated that the authorities relied upon by the tribunals could not be verified. The appellant submitted that even where a citation

appeared genuine, the passages attributed to the judgments were not traceable to the actual reports.

The respondents contested the appellant's submissions on the merits. Significantly, the bank's affidavit stated that the disputed authorities had not been cited by its counsel at the hearing and had instead been obtained through the adjudicating authority's own research. This fact was important because it shifted the immediate concern from advocate-generated false citation to the broader institutional problem of AI-assisted judicial research. The Supreme Court therefore considered the matter as a problem affecting both the Bar and the Bench.

5. *Judicial Reasoning*

The Supreme Court's reasoning proceeds from a distinction between the use of technology and the delegation of legal judgment.⁶ The Court recognised that AI can improve efficiency and assist professionals in dealing with increasing workloads. At the same time, it observed that the present generation of AI can produce non-existent or "hallucinated" material. The central danger arises when such material is presented as authoritative law without verification.

The Court emphasised that legal decision-making cannot be reduced to the mechanical retrieval of text. Judicial reasoning requires human assessment of facts, precedent, competing values and the consequences of competing legal interpretations. Technology can therefore remain an aid, but the responsibility for the correctness of the material used in adjudication must remain human.

The Court then drew a sharp line between legitimate AI use and fabricated authority. It clarified that its judgment does not prohibit the rightful use of AI. The concern is specifically with "presentation or reliance on fake or hallucinated material as if it were a court precedent." This is an important qualification because a blanket prohibition on AI would be technologically unrealistic and normatively unnecessary.

The Court's most consequential holding was that courts must adopt zero tolerance towards producing, citing or using AI-generated precedents without verification.⁷ It stated that an advocate who cites such judgments without verification commits misconduct and that a judge who relies upon such fabricated material commits a serious

⁶ *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd. & Anr.*, 2026 INSC 668, paras 7–9.

⁷ *Ibid.*, paras 18–20.

lapse. More fundamentally, the Court held that a decision based on fake or hallucinated material is “no decision in the eyes of the law” and must be set aside, irrespective of whether the material had a direct or indirect bearing on the decision-making process.

The Court linked this conclusion to the sanctity of adjudication and the rule of law. Judicial authority is legitimate not merely because an institution possesses jurisdiction, but because its decision is based upon ascertainable facts and law. A non-existent precedent cannot be tested, distinguished or applied by the opposing party. Its use therefore undermines procedural fairness as well as substantive correctness.

The Court also recognised that judicial prohibitions alone cannot solve the problem. It directed the Bar Council of India, as the apex statutory body, to constitute a committee to deliberate upon advocates submitting hallucinated AI material as precedent and to prescribe guiding principles and disciplinary consequences. The Court thereby moved from an individual dispute to institutional regulation.

Finally, the Court returned to the case itself. Because the NCLT and NCLAT decisions had been tainted by fabricated authorities, both were set aside. The section 7 application was restored for fresh determination, without expressing any opinion on its merits. The NCLT was directed to decide it expeditiously, preferably within two weeks, while the parties were directed to maintain status quo pending disposal.

6. Critical Analysis

The judgment’s principal strength is that it identifies the correct regulatory object. The problem is not artificial intelligence itself; the problem is unverified authority entering the adjudicatory record. Legal research has always involved tools – law reports, databases and electronic repositories. Generative AI is different because it can produce an apparently authoritative answer without guaranteeing that the underlying authority exists. The Court therefore correctly places verification at the centre of professional responsibility.

The decision also reinforces a foundational principle of adversarial adjudication. A party must have an opportunity to respond to the law relied upon against it. A fabricated authority deprives the opposing party of that opportunity because the supposed precedent cannot be located, scrutinised or distinguished. The problem is consequently not merely one of citation style; it affects fairness, transparency and

institutional legitimacy.

The Court's treatment of the NCLAT is equally significant. Appellate review is expected to provide a safeguard against errors at first instance. Here, the fabricated authorities escaped the first appellate tribunal's scrutiny. This demonstrates that technological error can propagate through hierarchical institutions when each level assumes that the previous decision has verified its sources. The judgment therefore rightly treats AI verification as a shared Bar-Bench responsibility.

However, the breadth of the "iota" principle deserves critical scrutiny. The Court states that a decision should be set aside even if an iota of fake or hallucinated material enters the decision-making process. Such a rule provides a powerful incentive for verification, but it may also create disproportionate consequences. A fabricated authority that supplies the decisive proposition is materially different from a mistaken citation appearing in background reasoning where the conclusion independently rests upon verified statutes and binding precedent.

A rigid invalidation rule could create strategic incentives. Parties might search lengthy judgments for minor citation errors and argue that any inaccurate AI-generated reference automatically nullifies the entire decision. This could increase remands, delay finality and burden courts. The present case involved multiple false authorities affecting the reasoning, so the result was justified on its facts. Future cases should clarify whether "iota" is categorical or whether materiality and the role of the false authority remain relevant.

The judgment also raises an institutional question concerning judicial research. The respondents' affidavit established that the disputed authorities were not cited by counsel but were apparently found through the NCLT's own research. Courts cannot simply impose verification obligations upon advocates while treating judicially generated research as inherently reliable. Judges and tribunal members must have access to verified databases and training in the limitations of generative AI.

The judgment is persuasive in insisting that the human being remains responsible. The "human in the loop" approach is preferable to both extremes: uncritical technological adoption and complete rejection of AI. Human review should require verification against an authoritative judgment, assessment of the proposition for which the case is cited, and confirmation that the quoted passage actually exists and remains good law.

At the same time, the judgment leaves unanswered questions concerning confidentiality, data protection, privilege and bias. A lawyer who uploads confidential pleadings or client information into an external generative AI system may create risks independent of hallucinated citations. AI-assisted legal research may also reproduce biases present in training data. These issues were not necessary to decide the appeal, but they are likely to become central to future regulation.

The judgment is strongest when read as an integrity decision rather than an AI-ban decision. It protects the chain of authority on which judicial reasoning depends. Its weakness is that it leaves the contours of proportionality and culpability to future development. That is understandable in a first-generation AI case, but those questions will require clarification as AI use becomes routine.

7. Comparative Perspective

The judgment's comparative dimension demonstrates that legal regulators are beginning to treat AI as a professional-risk issue rather than merely a technological innovation.⁸ The Supreme Court referred to the United Kingdom's regulatory experience in approving an AI-driven legal service subject to safeguards concerning confidentiality, conflicts, quality control, supervision and AI hallucinations.⁹ The underlying principle is similar to the Indian Court's approach: technological assistance may be permitted, but regulated professionals remain accountable for the output.

The comparative lesson is therefore not that India should copy another jurisdiction's rules. Instead, it is that accountability should remain attached to a human professional or judicial officer. AI systems should not become independent repositories of legal authority. Their outputs must remain traceable to authoritative sources and subject to human review.

In India, this approach also fits the existing statutory architecture. Section 35 of the Advocates Act, 1961 provides a mechanism for disciplinary proceedings where a State Bar Council has reason to believe that an advocate has been guilty of professional or other misconduct.¹⁰ The Supreme Court's direction to the Bar Council of India can therefore develop a specialised set of standards within the broader framework of professional responsibility.

⁸ Artificial Intelligence Act, Regulation (EU) 2024/1689.

⁹ Solicitors Regulation Authority, 'SRA approves first AI-driven law firm', 2025.

¹⁰ The Advocates Act, 1961, s. 35.

8. *Policy Implications*

The judgment has immediate implications for courts, tribunals, advocates, law firms and legal-technology providers. For courts and tribunals, verified legal databases should be preferred for authoritative research. If AI tools are used, the final judgment should rely only on authorities independently verified by the judicial officer or research staff. Training should explain hallucination risks and source verification.

For advocates and law firms, AI-use policies should require every generated citation, quotation and statutory proposition to be checked against an authoritative source before filing. Supervising advocates should remain accountable for filings.

For regulators, the Bar Council of India should develop a graded framework addressing verification, confidentiality, supervision and disciplinary consequences. The framework should distinguish intentional fabrication from negligent reliance and consider the materiality of the false authority.

For legal education, law schools should teach AI literacy as part of legal research methodology. Students should learn that a fluent AI response is not equivalent to an authoritative legal source.

For future litigation, the judgment is likely to generate applications seeking review or setting aside of orders where fabricated citations are discovered. Courts will have to determine how materially the false authority affected the reasoning. The Supreme Court's ruling provides the foundation, but future cases will refine its limits.

9. *Recommendations*

First, the Bar Council of India should issue an AI-specific professional practice guideline requiring verification of all AI-generated legal authorities before filing or oral reliance.

Second, disciplinary consequences should be graduated. Deliberate fabrication should attract the strongest sanction, while an isolated and promptly corrected negligent error should be treated differently, subject to the circumstances.

Third, courts and tribunals should adopt source-verification protocols. Where AI is used for research, every authority relied upon in the final order should be traceable to an official or recognised legal database.

Fourth, judicial training institutions should provide continuing

education on generative AI, including hallucination, bias, confidentiality and automation risks.

Fifth, legal-technology systems used by institutions should prioritize retrieval from verified legal databases and provide source links rather than generating unsupported authorities.

Sixth, future judicial decisions should clarify the materiality standard. The law should preserve the strong principle that fabricated precedent cannot be treated as law while avoiding automatic invalidation for immaterial typographical or citation errors that did not influence the decision.

Finally, confidentiality and data protection safeguards should accompany AI adoption. Lawyers and judges should not assume that every publicly available AI system is appropriate for confidential legal material.

10. Conclusion

Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd. is an important early judicial response to the arrival of generative AI in legal practice and adjudication. Its central proposition is both simple and constitutionally significant: technology may assist the administration of justice, but responsibility for the authenticity of law and the reasoning founded upon it cannot be delegated to a machine.

The Supreme Court was justified in setting aside the NCLT and NCLAT decisions because the process relied upon fabricated authorities. Its direction to the Bar Council of India recognises that the problem requires institutional regulation rather than isolated judicial warnings.

The judgment should not, however, be interpreted as a rejection of AI. Its more sustainable contribution is the establishment of a human-accountability model. AI may accelerate research, but it cannot transform an unverified output into precedent. The next stage of Indian jurisprudence should therefore focus on proportionality, culpability, materiality, confidentiality and institutional safeguards.

Ultimately, the case presents a broader question about the future of law. The value of judicial reasoning lies not in the speed with which an answer is produced, but in the reliability of the process by which truth is separated from error. In an age in which machines can generate convincing legal fiction, verification is not a technical formality. It is an essential condition of the rule of law.