



**HUMAN IN THE LOOP: THE SUPREME COURT'S RESPONSE TO AI
HALLUCINATIONS IN JUDICIAL DECISION-MAKING VIS-À-VIS POOJA
RAMESH SINGH V JAMMU AND KASHMIR BANK LTD & ANR**

Abinash Mohanty*

INTRODUCTION

Artificial Intelligence ('AI') has transformed legal practice, allowing lawyers to generate case summaries, statutory interpretations, and arguments within seconds. Yet generative AI suffers from 'hallucination', the confident generation of fabricated facts, citations, or authorities that appear genuine but have no basis in reality. Such inaccuracies may merely inconvenience an ordinary user, but their entry into judicial reasoning threatens the legitimacy of the justice system itself.

The Supreme Court confronted this problem precisely in *Pooja Ramesh Singh v Jammu and Kashmir Bank Ltd & Anr*¹, where the National Company Law Tribunal ('NCLT') had relied on precedents that were either non-existent or contained AI-generated passages falsely attributed to genuine decisions, and the National Company Law Appellate Tribunal ('NCLAT') failed to catch the defect. The Court held that a decision founded on hallucinated authorities 'is no decision at all', and amounts to a 'subversion of the rule of law'¹, while clarifying that it was not condemning AI itself, only its uncritical use. It insisted on a 'human in the loop' at every stage of adjudication³ and declared zero tolerance for fake AI-generated precedents. The decision is among the first by the Supreme Court to address the intersection of AI and legal ethics and judicial accountability, and its implications extend well beyond the insolvency dispute in which it arose.

*BA LLB (HONS), FIRST YEAR, AMITY UNIVERSITY, BENGALURU.

¹ *Pooja Ramesh Singh v Jammu and Kashmir Bank Ltd* 2026 INSC 668

FACTS OF THE CASE

Pan India Utilities Distribution Company Limited ('PIUDCL') obtained credit facilities from Jammu and Kashmir Bank Limited, secured by a corporate guarantee from Essel Infraprojects Limited ('EIL'). After PIUDCL defaulted and the account was classified as a Non-Performing Asset, the Bank initiated insolvency proceedings against EIL as the guarantor under section 7 of the Insolvency and Bankruptcy Code, 2016.

The appellant, Ms Pooja Ramesh Singh, a suspended director of EIL, challenged the maintainability of the proceedings before the NCLT, which nonetheless admitted the application, relying extensively on precedents said to support the guarantor's continuing liability². The NCLAT affirmed this reasoning on appeal.

Before the Supreme Court, Senior Advocate Madhavi Divan demonstrated that several authorities relied upon by the NCLT either did not exist or contained passages that could not be found in the reported judgments; some genuine citations had been paired with fictitious extracts, while other decisions were fabricated altogether. On independent verification, the Court confirmed these findings, and Jammu and Kashmir Bank filed an affidavit stating that its own counsel had not cited the fabricated material; the Tribunal appeared to have introduced it through its own research, raising a broader concern about how AI-generated content can infiltrate judicial reasoning even absent any misconduct by advocates.

ISSUES BEFORE THE SUPREME COURT

Although the underlying dispute concerned admission of insolvency proceedings, the controversy before the Supreme Court centred on judicial integrity. The principal issues were:

1. Whether an order founded on hallucinated precedents can survive scrutiny.
2. Whether such reliance vitiates the decision-making process regardless of materiality.
3. Whether advocates must independently verify AI-generated research.
4. Whether courts bear an independent duty to verify authorities they rely upon.
5. Whether AI's growing role in legal practice requires institutional safeguards.

² Insolvency and Bankruptcy Code, 2016, s 7.

JUDGMENT OF THE COURT

Speaking through Narasimha J, the Court held that judicial legitimacy depends on authentic legal reasoning, and that a decision contaminated by unreal or hallucinated authorities 'is no decision at all', amounting to a 'subversion of the rule of law'. It set aside both the NCLT and NCLAT judgments and restored the section 7 application for fresh adjudication, deliberately expressing no view on the underlying merits.

The Court distinguished between using AI as a technological aid and blindly depending on its outputs: AI may assist efficiency, but adjudication must remain a fundamentally human function, with 'total and absolute control over adjudication, [and] a human in the loop at every stage'. It further held that citing hallucinated authorities without verification 'is a misconduct on the part of an advocate', that it is 'a serious lapse' for a judge to rely on such material, and that such a decision must be set aside 'even if an iota' of fabricated material enters the reasoning directing the Bar Council of India to constitute a committee on the responsible use of AI in legal practice, including possible disciplinary measures.

RATIO DECIDENDI

The central ratio is that decisions founded on fabricated authorities are fundamentally defective because they undermine the legitimacy of adjudication itself; the problem is not AI's use but its uncritical acceptance without verification. Three propositions underlie this holding: first, judicial legitimacy depends on the authenticity of the precedents relied upon, so fictitious authorities strike at the foundation of the rule of law; second, AI is an assistive technology, not an autonomous decision-maker, and the duty to verify authorities is personal and non-delegable for both lawyers and judges; and third, technological innovation cannot dilute the ethical obligations advocates already owe³.

CRITICAL ANALYSIS

Reinforcing Judicial Integrity: The judgment's greatest strength is its recognition that fabricated precedents threaten not just individual cases but public confidence in the justice system. Since the doctrine of precedent is the cornerstone of Indian law, unknowing reliance on AI-fabricated authorities risks destabilising stare decisis itself, particularly since a fabricated

³ Advocates Act 1961; Bar Council of India Rules, pt VI, ch II (Standards of Professional Conduct and Etiquette).

precedent creates an illusion of certainty that may be repeated by other courts. By declaring such reliance unacceptable in all circumstances, the Court reaffirms authenticity as the defining feature of judicial reasoning.

Human Accountability and Professional Ethics: Generative AI predicts plausible language rather than verifying legal truth, and its outputs often appear authoritative despite being wrong, making the 'human in the loop' principle an appropriate and internationally recognised safeguard.⁴ This reasoning extends naturally to legal ethics: advocates, as officers of the court, already owed a duty to verify authorities before relying on them, and the judgment confirms that this obligation is unchanged by the involvement of AI. Notably, the Court also placed responsibility on judges themselves, who increasingly use AI-assisted research tools, preventing this framework from becoming one-sided regulation aimed only at the Bar.

Institutional Implications: The direction to the Bar Council of India to frame guidelines is a useful first step, but broader reforms will likely be needed as AI becomes further embedded in judicial administration, including mandatory disclosure of AI-assisted drafting, verification protocols for citations, judicial training on AI reliability, and disciplinary mechanisms for repeated misuse. Such measures would complement the judgment by building a structured framework rather than relying solely on professional ethics.

Unresolved Tensions: The zero-tolerance standard is not without difficulty. The Court held that a decision must fall even if only 'an iota' of fabricated material entered the reasoning, regardless of its effect on the outcome, a rule that avoids materiality disputes but may invite parties to reopen otherwise sound decisions over a single peripheral citation. The judgment also leaves open how hallucination is to be proved in contested cases: here fabrication was conceded and independently verified by the Court, but future disputes over obscure or unreported citations may not be so straightforward, and no procedural mechanism such as requiring authenticated copies from the citing party has been laid down for resolving them.

⁴ OECD, 'Recommendation of the Council on Artificial Intelligence' (2019); UNESCO, 'Recommendation on the Ethics of Artificial Intelligence' (2021); European Commission, 'Ethics Guidelines for Trustworthy AI' (2019); NITI Aayog, 'Responsible AI for All' (2021).

COMPARATIVE PERSPECTIVE

India is not alone in confronting this problem. In *Mata v Avianca Inc*⁵, the US District Court for the Southern District of New York sanctioned lawyers who submitted a brief citing fictitious precedents generated by ChatGPT, holding that technological innovation does not excuse failures of professional competence; courts in Canada and Australia have issued similar practice directions, and English courts have likewise cautioned firms deploying AI research tools without adequate supervision. Where these foreign responses focused chiefly on sanctioning advocates, the Indian Supreme Court adopted a broader institutional view, recognising that fabricated authorities can enter judicial reasoning through multiple channels and distributing responsibility across lawyers, judges, and regulators, an approach that may prove more effective because it targets systemic risk rather than individual misconduct alone.

BROADER IMPLICATIONS FOR INDIAN JURISPRUDENCE

The significance of *Pooja Ramesh Singh* extends well beyond insolvency law. It lays down foundational principles for the ethical use of AI across legal practice, is likely to shape judicial training on emerging technologies, and may inform future legislative efforts on AI governance, particularly around transparency, accountability, and human oversight. At its core, the judgment reinforces a constitutional principle: technological advancement must strengthen, not weaken, the rule of law, and its benefits for access to justice and efficiency can be realised only if AI remains subordinate to human judgment.

CONCLUSION

The Supreme Court's decision in *Pooja Ramesh Singh v Jammu and Kashmir Bank Ltd* marks a defining moment in the relationship between AI and the Indian legal system, and is among the first authoritative pronouncements on the risks generative AI poses to judicial decision-making. Its central achievement is refusing to treat the choice as binary: rather than condemning AI or excluding it from practice, the Court preserved its legitimate utility while insisting that adjudication remain a fundamentally human function.

This carries direct consequences for professional ethics and judicial administration alike: advocates remain responsible for every authority they place before a court, and judges retain a non-delegable duty to verify what they rely upon. The direction to the Bar Council of India is

⁵ *Mata v Avianca Inc* No 22-cv-1461 (SDNY 2023).

a useful start, but durable regulation will likely need to go further, addressing disclosure, verification standards, and judicial use of AI tools directly. Ultimately, the judgment's enduring contribution is its reaffirmation of a simple proposition: technology may assist the administration of justice, but justice itself remains a human responsibility.