
WHEN AI ENTERS THE JUDGMENT: JUDICIAL ACCOUNTABILITY AND THE DUTY TO GIVE REASONS IN INDIA

S. Mahesh Kumar, MKPM RV Institute of Legal Studies, Bengaluru

ABSTRACT

The increasing use of artificial intelligence in legal research and judicial administration raises the question of how best to reap the benefits of technology while upholding judicial integrity.

This dilemma took on a concrete form when it was revealed that the publicly available version of *Gayathri v. Guru Reddy Raju S.N.*, O.S. No. 4956/2011, contained apparent AI prompts and chatbot content. Beyond the issue of drafting mistakes, this incident raises questions about the verification of legal reasoning, liability for judicial documents and public trust in adjudication.

This article will explore the incident in the context of the constitutional principles of non-arbitrariness and fair procedure as well as the judicial mandate to provide reasoned decisions. It will also examine the Supreme Court's approach to fictitious legal authorities generated by AI in the case of *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.* In doing so, this article will argue that while assistance from AI may be appropriate, the responsibility of evaluating evidence, applying law and approving the final judgment must remain with the human adjudicator.

The article will propose a balanced framework based on human verification, pre-publication review, document integrity safeguards and context-sensitive remedies. The goal will be to develop a model of judicial AI governance that promotes efficiency without undermining accountability, procedural fairness or the rule of law.

1. Introduction

Artificial intelligence is increasingly being used to undertake tasks traditionally performed through human legal research, drafting, translation and document analysis. Within the judicial system, these applications promise considerable benefits. They may reduce repetitive work, assist in managing large volumes of case records and improve access to legal information. Yet, the incorporation of AI into adjudication presents a problem that extends beyond technological accuracy: who remains responsible for the reasoning expressed in an official judicial decision?

This question assumed particular significance following reports concerning *Gayathri v. Guru Reddy Raju S.N.*, O.S. No. 4956/2011, a partition dispute decided by a Bengaluru civil court. The publicly accessible copy of the judgment appears to contain chatbot-style prompts, AI-generated legal discussion and drafting language within the judicial document.¹ The material reportedly appears alongside a placeholder in the section dealing with reasons. The incident, considered cautiously and without presuming who generated or inserted the material, exposes a critical vulnerability in the production and publication of judicial decisions.

The concern is not that every use of AI in judicial work is improper. Rather, it is that a document formally presented as a judgment must communicate the court's own adjudicative conclusions and the reasons supporting them.² Where unreviewed AI-generated material enters that document, the distinction between assistance with drafting and the exercise of judicial reasoning becomes difficult to assess.

This article argues that India requires a proportionate framework for AI-assisted judicial drafting, grounded in the constitutional values of non-arbitrariness, fair procedure and judicial accountability. The central principle should be straightforward: technology may assist the judicial process, but responsibility for the decision, its reasoning and its final published form must remain human.

2. The Bengaluru Incident: More Than a Drafting Error

The reported incident in *Gayathri* is significant because it appears to concern the integrity of the final judicial document rather than merely an inaccurate citation in legal research. The

¹ *Gayathri v. Guru Reddy Raju S.N.*, O.S. No. 4956 of 2011 (XXX Addl. City Civil Judge, Bengaluru, Jan. 27, 2026)

² *S.N. Mukherjee v. Union of India*, (1990) 4 SCC 594.

publicly available text contains material resembling a conversation with a generative AI system, including a prompt concerning the legal consequences of a will, a response discussing the relevant property interests, a subsequent correction to the factual premise and a revised response.

Such material, if reproduced in an official judgment without appropriate review, raises three distinct concerns.

First, it creates uncertainty about the relationship between the court's reasoning and the text generated during the drafting process. An AI system may produce a plausible legal explanation, but plausibility does not establish that the explanation accurately reflects the evidence, the applicable law or the adjudicator's independent assessment.

Second, the incident raises questions about document verification. Judicial decisions pass through processes of preparation, correction, authentication and publication. A conspicuous prompt or an unfinished placeholder surviving into a publicly accessible judgment suggests that document-integrity checks may be inadequate, irrespective of how the material entered the document.

Third, the incident affects public confidence. A litigant ordinarily understands a judgment as the product of a judicial determination reached after considering the pleadings, evidence, submissions and governing law. The unexplained appearance of conversational AI material can undermine confidence in that process.

Nevertheless, an important distinction must be maintained. The presence of AI-generated text does not, by itself, prove that the adjudicator delegated the decision to AI, failed to consider the record or adopted the generated reasoning. Nor does the appearance of an editorial defect automatically establish that the operative order is legally invalid. The incident should therefore be treated as a serious institutional warning, not as proof of misconduct by any particular individual. The proper question is whether the judicial process contains safeguards capable of ensuring that the final document accurately records an independently reached judicial decision.

3. The Duty to Give Reasons: A Constitutional Foundation

The duty to provide reasons is an established feature of judicial and quasi-judicial decision-making in India. Reasons disclose the path through which the decision-maker has moved from

the relevant facts and legal principles to the conclusion. They enable the parties to understand why they succeeded or failed, facilitate appellate review and discipline the decision-making process itself.

In *Kranti Associates Pvt. Ltd. v. Masood Ahmed Khan*³, the Supreme Court emphasised the importance of reasons in judicial, quasi-judicial and administrative decisions. The principle is not merely a matter of literary form. A reasoned decision demonstrates that the relevant considerations have been evaluated and that the conclusion is not arbitrary or unexplained.

Similarly, *S.N. Mukherjee v. Union of India* recognised the importance of recording reasons in the exercise of adjudicative power, while explaining the legal principles governing that obligation. The precise content of the duty may vary with the nature of the proceeding and the applicable law, but the underlying purpose remains the same: to make the exercise of public power intelligible and reviewable.

Articles 14 and 21 of the Constitution provide a broader constitutional context.⁴ Article 14 protects against arbitrary State action, while Article 21 requires legally valid and fair procedures where life or personal liberty is implicated. The relevance of these provisions to judicial reasoning must, however, be expressed carefully. Not every defective judgment necessarily constitutes an independent constitutional violation, and Article 21 should not be invoked as a universal formula for every procedural irregularity. Rather, the constitutional significance depends on the nature of the decision, the applicable legal requirements and the effect of the defect on fairness and legality.

AI-assisted drafting complicates this established framework because the final text may be fluent and apparently coherent without necessarily reflecting a verified application of law to the record. A generative system can organise propositions, suggest arguments and produce language resembling judicial reasoning. It cannot, merely by generating that language, establish that the evidence has been correctly evaluated or that the legal conclusion is justified.

Accordingly, the duty to give reasons cannot be satisfied solely by the presence of polished legal prose. The reasons must be those adopted by the adjudicator after an independent

³ *Kranti Associates Pvt. Ltd. v. Masood Ahmed Khan*, (2010) 9 SCC 496

⁴ INDIA CONST. arts. 14, 21

assessment of the relevant material.⁵ Human review must therefore be substantive rather than ceremonial.

4. From AI Hallucinations to Judicial Accountability

The Supreme Court's decision in *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.*,⁶ provides an important contemporary point of comparison. In that case, the Court set aside tribunal decisions that relied on non-existent or fabricated legal authorities generated through AI. It emphasised the need for human control over adjudication and verification of AI-generated legal material.

The case concerns a problem distinct from the reported *Gayathri* incident. In *Pooja Ramesh Singh*, the central issue was reliance on fictitious precedents as legal authorities. In *Gayathri*, the concern arises from apparent chatbot prompts and generated drafting material appearing within the publicly accessible judicial document. The first involves the reliability of legal sources used in adjudication; the second raises questions about the integrity of the document that communicates the adjudication.

The distinction matters because the risks of AI in courts cannot be reduced to hallucinated citations alone. Even where every cited authority exists, AI-generated reasoning may misunderstand a statutory provision, omit a material factual distinction, misstate a party's submission or offer a conclusion unsupported by the evidence. Conversely, a drafting error may appear in a judgment without having affected the court's actual reasoning. The legal response must be sensitive to these differences.

The appropriate principle is therefore not a blanket prohibition on AI, but a requirement that the adjudicator verify and independently adopt every material proposition relied upon in the final decision. Verification must extend to factual premises, statutory interpretation, case citations and the relationship between the evidence and the conclusion.

Accountability also cannot be allocated solely by asking whether AI was used. A meaningful inquiry must consider who was responsible for reviewing the output, whether the final document was checked, what role the generated material played and whether the defect affected

⁵ Supreme Court of India, Centre for Research & Planning, White Paper on Artificial Intelligence and Judiciary 8, 65–66 (Nov. 2025).

⁶ *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.*, 2026 INSC 668

the determination. The objective is to preserve the human responsibility that makes judicial authority legitimate.

5. Efficiency, Human Oversight and the Risk of Automation Bias

Any proposal to regulate AI in courts must recognise the institutional context in which it operates. Indian courts face substantial caseloads, extensive records and demands for timely adjudication. AI-assisted transcription, translation, document classification and legal research may help address these difficulties.

A policy that treats all AI use as inherently suspect could discourage beneficial applications and preserve avoidable administrative burdens. The relevant distinction is between assistance with tasks that can be independently checked and reliance on AI for conclusions requiring judicial evaluation.

For example, AI may assist in arranging a chronology of pleadings or identifying authorities for further research. The judge or authorised legal researcher must still verify the source material and determine its relevance. Similarly, transcription and translation tools may improve efficiency, but errors affecting a witness's statement or a party's submission require correction before they influence the decision.

The deeper risk is automation bias: the tendency to place excessive confidence in a system's output because it appears authoritative, technically sophisticated or well written.⁷ In a judicial setting, this risk is especially serious because generated text may reproduce the formal language of legal reasoning. Its apparent precision can obscure factual omissions, doctrinal errors or unsupported assumptions.

Human oversight must therefore involve more than a final glance at a completed document. The responsible adjudicator must understand the relevant reasoning, verify material propositions and remain willing to reject the generated output. AI should assist the exercise of judicial judgment, not become a substitute for it.

Confidentiality also requires attention. Judicial records may contain personal information, financial details, family disputes and sensitive evidence. Entering such material into an

⁷ Raja Parasuraman & Victor Riley, *Humans and Automation: Use, Misuse, Disuse, Abuse*, 39 *Hum. Factors* 230, 230–53 (1997).

unauthorised external AI system may create risks concerning data security, confidentiality and control over information. Any institutional framework should accordingly distinguish approved, secure tools from systems that have not been evaluated for judicial use.

6. A Proportionate Framework for AI-Assisted Judicial Drafting

India should adopt a practical framework that preserves the benefits of technology while protecting the integrity of judicial decisions. Four safeguards are particularly important.

First, human approval must be substantive and traceable. The judicial officer responsible for the decision should retain responsibility for the final reasoning and operative order. AI-generated text should not be incorporated merely because it is fluent or appears legally plausible. Material findings and conclusions must be independently assessed against the record and the applicable law.

Second, legal and factual verification should be mandatory. Every authority, quotation, statutory provision and material factual proposition suggested by AI must be checked against a reliable source. Verification should include whether the cited case exists, whether the quoted passage appears in it, whether it remains good law where relevant and whether it actually supports the proposition for which it is cited. This safeguard responds directly to the problem identified in *Pooja Ramesh Singh* while extending beyond fictitious citations.

Third, final-document checks should be institutionalised. Before a judgment is authenticated or published, a proportionate review should identify obvious drafting artefacts, unfinished placeholders, duplicated passages, unexplained prompts and inconsistent formatting. Such checks should complement, rather than replace, substantive judicial review. A technically clean document is not necessarily a legally sound decision, but an avoidable publication defect should not be allowed to pass unchecked.

Fourth, correction and accountability procedures should be clear. If an AI-related defect is discovered after publication, the response should depend on its nature and effect. A typographical or editorial error that does not affect the reasoning may be capable of correction through an appropriate procedural mechanism. A defect that conceals missing reasoning, introduces an unverified factual premise or materially affects the determination may require more substantial judicial scrutiny. Any correction must follow the applicable procedural law

and respect the rights of the parties.

These safeguards should be accompanied by judicial training, secure technology infrastructure and institutional guidance on acceptable uses of AI. Training should address not only technical limitations but also the legal consequences of relying on unverified output. A clear protocol would protect judges and court staff by defining responsibilities rather than leaving them to navigate these risks without common standards.

The framework should be proportionate. Disclosure of every routine use of spelling assistance or transcription software may add little value. By contrast, the use of AI-generated substantive legal analysis may justify stronger verification and internal review. The governing principle should be risk-sensitive oversight, with stricter safeguards where the output may affect findings, legal conclusions or the rights of litigants.

7. Conclusion

AI presents Indian courts with both an opportunity and a responsibility. It can assist in managing legal information, processing documents and reducing repetitive work. Yet, the authority of a judgment does not arise from the sophistication of the technology used to prepare it. It arises from the lawful exercise of judicial power through independent evaluation, reasoned conclusions and fair procedure.

The reported incident in *Gayathri v. Guru Reddy Raju S.N.* demonstrates why judicial AI governance must address more than fabricated case citations. It must also protect the integrity of the final judicial document and ensure that drafting assistance does not obscure the source or adequacy of the court's reasoning. Read alongside *Kranti Associates, S.N. Mukherjee* and *Pooja Ramesh Singh*, the incident supports a framework based on verification, human responsibility, document-integrity checks and proportionate remedies.

The objective should not be to exclude technology from the courtroom, nor to treat every appearance of AI-generated text as proof of invalid adjudication. It should be to ensure that technological assistance remains subordinate to the judicial function. AI may help prepare the judgment, but it cannot assume responsibility for the judgment. That responsibility must remain with the human decision-maker, whose reasons must withstand scrutiny under the law.