




BLOG | 13 APR 2026 | 11 MINS READ

Navigating the Crisis of “Phantom Precedents” and the Erosion of Professional Duty

When AI Hallucinates the Law



****Ojasvi**

The Crisis of the “Fake Law”

The Indian legal system is moving away from its old paper and ink tradition toward a digital future. For decades, a physical copy of a Supreme Court judgment was seen as the absolute truth, the unshakeable foundation of justice. But as we move through 2026, this foundation is starting to crack.

The problem is the sudden arrival of Large Language Models (LLMs) like ChatGPT or Gemini. These tools promise to clear India’s massive backlog of cases by doing research in seconds. However, they have introduced a dangerous new guest into our courtrooms, namely, “Synthetic judgements.”

The legal system relies on a doctrine called stare decisis, which is a legal principle that directs courts to adhere to previous judgments or judgments of a higher court or tribunal, as it has persuasive and binding authority while resolving a case with comparable facts. If the records are filled with fake cases that look and sound real but only exist in a machine’s imagination, the entire system of trust and reliability collapses. The core of the problem is two-fold:

1. **A Breach of Trust:** When a lawyer submits an unverified AI result, they are failing their professional duty to be honest with the court.
2. **A Constitutional Crisis:** Under [Article 21](#), every Indian has the [right to a Fair Trial](#). You cannot have a fair trial if the law being used against you is a digital fiction that you cannot find, read, or challenge.

To understand why this is so dangerous, we have to look at the judicial alarms that have been triggered across Indian courts in early 2026.

Evaluating Recent Judicial Alarms

The initial honeymoon phase between the Indian legal profession and Generative AI came to an abrupt and jarring end in early 2026. What was once seen as a [revolutionary tool for democratization of legal research](#) has instead introduced [synthetic judgments](#) into the stream of justice. This shift from curiosity to crackdown was necessitated by high-profile incidents where the boundary between research and fabrication was catastrophically blurred, leading to many fictional cases being cited in litigations presented to the judicial body.

A watershed moment occurred in March 2026 with the case of [Gummadi Usha Rani v Sure Mallikarjuna Rao \(2026\)](#). In this startling instance, the Supreme Court discovered that a lower court in Andhra Pradesh had inadvertently based an entire order on fictional precedents fabricated by an AI tool used by the petitioner's counsel. [The Court](#) noted that this is not merely an error in the decision making, rather it may amount to misconduct and entail legal consequences. The Court asserted that the matter required a deeper examination.

Simultaneously, [the Bombay High Court](#) emerged as an accidental pioneer in "AI Forensics" in January 2026. During a routine proceeding, a presiding judge identified a lawyer's submission as AI-generated by scrutinizing specific visual "tells." The court observed that raw LLM output often carries distinct, robotic markers that a seasoned legal mind should have caught. These markers included:

- **Residual Artifacts:** The presence of green tick-marks or unconventional symbols left over from the AI's user interface.
- **Syntactic Monotony:** Repetitive bullet points and a robotic flow of prose that mimics the style of a judgment without the substance of a legal mind.

The Bombay High Court's decision to impose [a fine of ₹50,000](#) on the offending party was a significant policy signal. This penalty was more than a mere fine; it served as a clear message of zero-tolerance for wasting judicial time with unverified drafts. For trial court judges across India, who may lack technical sophistication, this case provides a blueprint for identifying and penalizing the use of AI. This judicial pushback highlights a growing consensus: submitting AI-generated fictions is not an innovative use of technology, but a question of the court's dignity.

Deconstructing the Hallucination

For the modern legal practitioner, understanding the technical mechanics of AI is no longer a peripheral skill; it is a mandatory requirement for professional survival. We must demystify what a Large Language Model (LLM) is to understand why it "hallucinates." At its core, an LLM is a system designed to [predict the next word](#) in a sequence based on statistical probability derived from vast amounts of training data. It does not "search" for law in the way a human researcher queries a database; it constructs sentences that sound plausible based on linguistic patterns.

The danger in legal practice arises from the primary phenomena of "Mirror Effect." The "[Mirror Effect](#)" occurs when an AI, programmed to be helpful, creates exactly what the lawyer is asking for, regardless of reality.

The Anatomy of a Hallucination

A lawyer prompts an AI: "Find a Supreme Court case from 2018 regarding digital asset property rights."

The AI, wanting to satisfy the user, predicts a statistically likely response: *State of Maharashtra v. K.L. Sharma (2018)*. It may even provide a plausible-sounding citation like (2018) 4 SCC 210 and a 500-word excerpt that mimics the tone of a Supreme Court Justice. However, this case does not exist. It is a digital fiction born of probability, not a precedent born of the law.

The predictive nature of AI is fundamentally at odds with the verification-heavy nature of legal practice. Authorized databases like SCC Online and Manupatra are built on the principles of curation and accuracy, whereas LLMs are built on the principles of fluency and plausibility. When a lawyer, under the pressure of a heavy backlog, fails to verify an LLM's output against these authorized sources, they are not practicing law; they are engaging in a dangerous form of digital hearsay.

[Cookie settings](#) [privacy policy](#) [Accept](#)

The reliance on unverified AI is not merely an administrative error or a lapse in research; it represents a significant shift toward the violation of our fundamental right under [Article 21](#), where every citizen is guaranteed the Right to a Fair Trial.

While the use of unverified AI is often dismissed as mere professional negligence or ‘lazy lawyering,’ its systemic integration into the judicial record triggers a profound Constitutional conflict. Article 21 is the guarantee of ‘procedure established by law,’ however, when a court’s reasoning is built upon citations that exist only as statistical probabilities, the procedure ceases to be ‘law’ and becomes a fiction. This compromises the Rule of Law, which [demands](#) that the law be certain, discoverable, and equal for all. If the judicial record is polluted by digital fabrications, the doctrine of stare decisis collapses.

If the reasoning behind a legal submission or worse, a judicial order is generated by an algorithm whose internal logic is based on statistics and behaviour readings, the transparency of the law vanishes. [The “Right to Reasoned Decisions”](#) is a cornerstone of Indian administrative law; an unverified output is the antithesis of a reasoned decision. A defendant cannot effectively challenge a decision or cross-examine a logic that originated from a system that cannot explain why it reached a specific conclusion. This lack of auditability violates the principle of natural justice.

The Ethical Erosion: A Crisis of Professional Identity

The reliance on unverified AI is far more than a technical shortcut; it is a violation of the lawyer’s role as an “Officer of the Court.” The legal profession is built on the bedrock of a promise that the human advocate has applied their mind to the client’s cause. When a lawyer dumps AI-generated content into a petition without rigorous verification, they effectively delegate their professional conscience to a probability-based algorithm. This is not just bad lawyering, rather it is a breach of the [Section 35](#) of the Advocates Act regarding professional misconduct, as it misleads the court and gambles with the client’s legal standing.

Furthermore, this trend threatens to trigger a systemic collapse of public trust. During proceedings arguments may vary, but there is a belief that the facts and laws cited to support or contradict those arguments are real. If the judiciary becomes a factory for “synthetic law,” the public will cease to view lawyers as learned professionals and start viewing them as mere prompt-engineers. Once the client realizes that their legal advice, arguments to be presented in the court were generated by a chatbot while they were charged premium fees the moral authority of the advocate evaporates. If lawyers treat the job of drafting legal arguments as a routine task to hand over to machines, the profession risks losing its depth and turning into a data-processing business.

Comparative law provides essential insights for developing local AI standards. The global legal community has been grappling with these issues since the [Mata v Avianca Inc.](#) incident in New York, where lawyers were sanctioned for citing six non-existent cases. While the American response focused heavily on punitive measures for lawyer dishonesty, the United Kingdom has developed a more nuanced, structural model by 2026.

The British threshold, as established by the [UK Courts and Tribunals Judiciary](#), creates a clear taxonomy of tasks, distinguishing between the mechanical and the intellectual:

1. **Administrative Tasks (Allowed):** Using AI for grammar checks, summarizing undisputed facts, or organizing large volumes of discovery documents.
2. **Substantive Legal Reasoning (Prohibited):** Outsourcing the core intellectual work of interpreting statutes, finding case law, or constructing the ‘ratio decidendi’ (reason for the decision) to an LLM.

The UK model reinforces the idea that AI should function as a sophisticated secretary rather than a research associate. The difference lies in the delegation of intellectual agency. As a ‘secretary,’ AI is utilized for administrative tasks such as composing, summarising and prioritising emails, transcribing and summarising meetings, and composing memoranda. Conversely, treating AI as a ‘research associate’ involves delegating the core tasks of legal discovery and analysis. The UK Judicial Guidance warns that because LLMs function on probability rather than precedent, they are fundamentally unfit for this associate role. India should synthesize these global lessons by focusing on structural guidance that prevents the error before it reaches the judge’s desk, ensuring that the “British Threshold” is adapted for the unique complexities of the Indian district courts.

Strategic Recommendations

The solution to the judicial backlog and the use of AI in the legal sphere is not merely faster automation, but the strengthening of infrastructure and human accountability. To restore institutional legitimacy and protect the constitutional rights of litigants, I propose a three-tiered strategy:

1. **Judicial Sanctions as a Deterrent:** Courts must move beyond simple monetary costs. For repeat offenders who knowingly submit fake judgments to gain an advantage, the court should charge the counsel with [contempt of court](#). Misleading the court is a deliberate act that falls

2. **Certified Legal AI:** The Ministry of Law and Justice should collaborate with the Digital Courts 2.0 project to provide a State-Certified AI Tool. This tool would draw information exclusively from the National Judicial Data Grid (NJDG). To address the inherent risk of hallucinations where models may still misattribute party names or citations even within a closed dataset this tool must utilize [Retrieval-Augmented Generation \(RAG\)](#). Unlike standard AI that guesses information from memory, a RAG-based system allows generative AI models to access additional external knowledge bases, such as internal organizational data, scholarly journals and specialized datasets. By integrating relevant information into the generation process, [chatbots](#) and other [natural language processing \(NLP\)](#) tools can create more accurate domain-specific content. By incorporating mandatory hyperlinking to the original NJDG source files, the system creates a verification loop. This ensures that rather than offering a complete delegation of research, this framework reinforces professional responsibility by requiring the advocate to perform a final verification of the primary text.

These recommendations are designed to ensure that while our tools evolve, the accountability of the practitioner remains the primary safeguard of the law. Automating the *process* of justice is a virtue; automating the *judgment* of justice is a vice.

Conclusion

The law is not a math problem for an algorithm to solve. It is a human effort that requires empathy, context, and a sense of right and wrong. While an AI can replicate the linguistic patterns of a judicial opinion, it lacks the capacity for legal reasoning. The recent judicial crackdown on “Synthetic judgements” is a necessary and urgent correction to a growing culture of blind reliance on AI.

As we move through 2026, the legal profession must adopt a use but verify culture. We must remember that AI is a powerful mirror, it reflects what it has seen, but it is not a source of truth. The future of Indian justice depends on our ability to ensure that while our tools may be artificial, our intelligence, our ethics, and our accountability remain strictly human. The dignity of our courts and the rights of our citizens are too precious to be left in the hands of a machine which answers based on guess and probability.

****Ojasvi is a First year student [BA LLB (Hons.)] at Jindal Global Law School.**

[privacy policy](#)