



The Indian Journal for Research in Law and Management

Open Access Law Journal – Copyright © 2026

Editor-in-Chief – Dr. Muktai Deb Chavan; Publisher – Alden Vas; ISSN: 2583-9896

This is an Open Access article distributed under the terms of the Creative Commons Attribution-Non-Commercial-Share Alike 4.0 International (CC-BY-NC-SA 4.0) License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium provided the original work is properly cited.

ARTIFICIAL INTELLIGENCE IN THE INDIAN JUDICIARY: REVOLUTION OR RISK?

~*HITESH BHOOTRA*

INTRODUCTION:

Somewhere in a courtroom of Chandigarh asking a chatbot about bail jurisprudence and the Supreme Court of India comparing fake AI generated case laws to the Bhopal gas leak, Indian law kind of ended up in a story it never really asked to join. And no, that's not just me being dramatic it is, pretty much word for word, how the Supreme Court described the risk of AI hallucinations getting inside judicial reasoning: invisible, insidious and catastrophic, by the time anyone catches it.

As a law student I did not grow up thinking that courts were places where technology moves quickly, or moves fast at all. Courts are supposed to be slow, deliberate, almost stubbornly human. And yet, in the space of about three years, Indian courts have gone from curiously experimenting with ChatGPT in a single bail order to actively drafting a national regulatory framework for how artificial intelligence may and may not be used in the administration of justice. That arc, compressed and dramatic, is really the whole story of "Revolution or Risk." It has been both, often in the same case.

This piece tries to walk through that arc honestly as what AI is actually doing inside Indian courts today, what has already gone wrong, how the judiciary has responded, and what all of this might mean for someone like me, who will (hopefully) be practicing law in a system that is rewriting its own rules of engagement with machines while I am still a student.

WHERE AI ALREADY LIVES INSIDE INDIAN COURTS:

It helps to first separate the myth from the reality. AI in Indian courts is not, as pop culture might suggest, a robot judge deciding cases. What it actually looks like is far less dramatic and far more administrative.

The Supreme Court of India has developed and deployed a set of indigenous tools built specifically for judicial use: **SUPACE**, which assists in analyzing case records and surfacing relevant material for judges and their research teams; **SUVAS**, a translation tool that converts judgments between English and regional Indian languages; and **TERES**, used for real-time transcription of court proceedings. They are not really decision-making systems, more like very able research, and clerical helpers if you know what I mean. The scale is kind of wild honestly: by 2024, SUVAS-style translation work had turned tens of thousands of Supreme Court rulings into Hindi and a few other regional languages, which counts a lot in a country where most litigants simply do not read fluent legal English

In September 2024, the Supreme Court also released proper guidelines for designing, and running AI tools that deal with transcribing arguments and court proceedings. That felt like an early signal, like the institution was thinking about AI as something more than a gimmick, more like a basic service that still needs guardrails from the beginning. The Court also collaborated with IIT Madras to build AI and machine learning systems for spotting problems in electronic filings, and prototype access was already extended to roughly 200 Advocates-on-Record, which is pretty significant.

None of this is science fiction. It is, essentially, digitization with a layer of intelligence reducing the mountain of paperwork and translation backlog that has historically slowed Indian courts down. Given that Indian courts are dealing with a docket crisis running into tens of millions of pending cases, tools that speed up transcription, translation, and case-record review are not a luxury. They are closer to a necessity.

THE MOMENT EVERYONE NOTICED: CHATGPT AND A BAIL ORDER:

If there is, like one single moment that managed to pull AI in Indian courts out of the usual space of internal administration and into the public back-and-forth, it was March 2023. That's when the Punjab and Haryana High Court, in an order by Justice Anoop Chitkara, actually referred to a response generated by ChatGPT while dealing with a bail application in

“Jaswinder Singh @ Jassi v. State of Punjab¹”. The Court had put a pretty narrow question to the chatbot, essentially asking for wider global jurisprudence on bail when the facts involve cruelty and it came back with a generic, almost textbook type answer, on how courts typically weigh violent conduct and flight risk.

Credit where it's due, the order was also careful, to underline that the ChatGPT response was not being used as any view on the merits, of the particular case. and it further said that no subordinate court should treat the reference as binding, or even as persuasive, in any formal way. Bail was, in fact, denied primarily on the accused's prior criminal history and the severity of the assault grounds that had nothing to do with the chatbot.

Still, the reaction from the legal community was sharp, and reasonably so. Commentators pointed out that the order's own reasoning elsewhere relied on sources like an American criminal law textbook and a nineteenth-century case about capital punishment methods, without anyone in the process actually pausing to ask whether a generative AI tool trained on unverified internet text, with no access to Indian legal databases should be anywhere near an order that decides whether a person stays in jail or walks free. It was less a scandal than a warning shot: a demonstration that judges, like everyone else, were curious about these tools, and that curiosity without safeguards could quietly reshape judicial reasoning before anyone had agreed on the rules.

WHEN THE WARNING SHOT BECAME A CRISIS:

What happened three years later was no longer a warning shot. It was the real thing.

In July 2026, the Supreme Court delivered its judgment in **“Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.²”** a case that will likely be taught in law schools for years as the definitive Indian precedent on AI hallucination in judicial decision-making. The facts are almost unbelievable in their ordinariness, an insolvency dispute under **Section 7 of the Insolvency and Bankruptcy Code**, where the National Company Law Tribunal in Mumbai, back in August 2024 had admitted an insolvency application against Essel Infra projects Limited over an unpaid corporate guarantee. Then, on appeal, the National Company Law

¹ Jaswinder Singh @ Jassi v. State of Punjab 2026 INSC 23

² Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd. 2026 INSC 668

Appellate Tribunal upheld that order in September 2025, more or less without much fuss, or so it seemed.

But the trouble only really came into view when the matter reached the Supreme Court. On close scrutiny, several of the so called “precedents” the NCLT had leaned on to shut down the appellant’s points turned out to be entirely fictitious, cases that simply do not exist in any law report, and alongside them were “real” citations that had been attached to legal propositions they never actually said. Jammu and Kashmir Bank’s own lawyers, for instance, confirmed they had never cited these things at all. The Tribunal, the Supreme Court was told, appeared to have produced them through AI-assisted research of some sort, and nobody had checked before these items were woven into a binding order, an order that then went through appellate scrutiny at the NCLAT, and somehow the fabrication was not caught, not even once

The Supreme Court, sitting before a bench of Justices P.S. Narasimha and Alok Aradhe, did not treat this like some minor procedural stumble, you know. It said that fabricated precedents entering judicial reasoning are sort of like the release of methyl isocyanate, the gas connected to the Bhopal disaster, unseen at first, insidious, and then destructive by the time anybody even finally notices. That, honestly, is about as emphatic a warning as a court can put out there. The Bench set aside both the NCLT and NCLAT orders, restored the insolvency application for fresh consideration, and most significantly for practitioners held that courts must adopt a zero-tolerance approach toward AI-generated precedents that are not independently verified. Citing an unverified, AI-fabricated judgment, the Court said, amounts to professional misconduct on the part of a lawyer, and reliance on such material by a judge is an equally serious lapse. The ruling went further: even an "iota" of hallucinated material entering a judicial decision is enough to strike at its validity.

What makes this case analytically interesting, beyond its dramatic language, is what it reveals about the mechanics of hallucination. This was not a rogue lawyer trying to win a case through deception. The fabricated citations appear to have entered through the Tribunal's own AI-assisted research process meaning the institution meant to safeguard against error was itself the source of it. That should worry anyone who assumes AI risk in courts is primarily about dishonest advocates gaming the system. Sometimes the danger sits closer to the bench than to the bar. This was also not an isolated episode. Reports have surfaced of a trial court judge in Vijayawada dismissing a case in reliance on four Supreme Court judgments that, on later verification, simply did not exist. The pattern is consistent enough that the Supreme Court's

own AI Committee named "fabrication of cases and hallucination" as the primary risk associated with AI use in its White Paper on AI and Judiciary, released in November 2025.

THE INSTITUTIONAL RESPONSE: REGULATION CATCHING UP TO REALITY:

To its credit, the Indian judiciary has not simply absorbed these incidents and moved on. It has responded with a genuine, if still evolving, regulatory architecture.

The Kerala High Court got to be the first in the country, to put out a proper AI use policy for judicial officers and staff, back in July 2025. And yeah, it's pretty tough minded: it directly blocks the use of cloud based generative AI tools like ChatGPT or DeepSeek for anything case related research, with the stated worry being confidentiality. Basically, the court says if sensitive case material gets uploaded to outside servers, the chances of a leak are real, and not just theoretical. The policy then allows only tools that are pre-approved and vetted, it also insists that every response has to be checked and confirmed by a human independently, before anyone leans on it. If somebody violates it, they're looking at disciplinary action under existing Kerala civil service rules, so the fallout can go from a simple formal warning, up to suspension or worse... depending on how serious it gets.

Now at the wider national level, the Supreme Court's AI Committee released Draft Regulations for the Use of AI in Courts 2026, and these are out for public consultation. This is, in a way, India's first run at a broad judiciary wide set of rules covering how AI may be used across the Supreme Court, High Courts and also the subordinate courts. The draft is built around the idea that AI may help with court administration and legal research, but it cannot ever be permitted to decide anything, or to influence findings in a way that shifts judgment making. The draft keeps repeating that this responsibility, the judge's responsibility, stays solely with the judge, and it can't be delegated, not even indirectly, or "handed over", as the framework frames it.

What is kind of striking, reading through this timeline, is the tone of the Supreme Court itself when it speaks publicly about all this. In December 2025, Chief Justice of India Surya Kant, while addressing a petition seeking regulation of AI in the judicial process, was pretty unambiguous: AI would never be permitted to override judicial decision making, and the independence and discretion of a human judge stays firmly nonnegotiable. That doesn't sound like the sort of language an institution uses when it's rushing to hand over its core function to a machine. If anything, it reads more like the language of an institution that has watched something go wrong pretty close-up, and then adjusted its posture accordingly.

REVOLUTION OR RISK? READING THE TWO SIDES TOGETHER:

It would be intellectually lazy to pick one side of this debate and defend it uncritically, so let me try to hold both honestly.

“The case for revolution” is not abstract. Indian courts are kind of drowning in pending cases and a meaningful share of that backlog is not about complex legal reasoning at all. It is translation delays, transcription bottlenecks and the sheer administrative weight of processing paper, in a system serving over a billion people across more than twenty officially recognized languages. Tools like SUVAS and SUPACE attack exactly that problem. They do not replace judicial reasoning; they just clear the undergrowth so judges can spend their limited time on reasoning itself. For a litigant in a small-town waiting year for a judgment to be translated into a language they can actually read, that kind of AI is not some philosophical question. It is access to justice, in the most literal sense.

“The case for risk” is equally concrete, and the Jammu and Kashmir Bank case is the clearest possible illustration of it. Generated AI tools really don’t know the law the way a trained lawyer knows, AI tools just only predict text that sounds smooth, but smooth text can be and also convincingly wrong. When that wrongness enters a judicial order, it does not stay contained. It gets cited by other lawyers, relied upon by other tribunals and as happened here can survive an entire appellate process without detection. The damage is not to one case; it is to the reliability of the citation system that the entire common law method depends on. Precedent only works as a discipline if the precedents are real.

There is a subtler risk too, one the Supreme Court's own White Paper gestures toward algorithmic bias and a kind of lack of transparency. Even AI tools that never hallucinate outright can, very quietly encode the biases of their training data or the biases of their designers, and then produce outputs that look neutral, but are not. Unlike a judge’s reasoning, which can be interrogated, appealed, and laid out, an AI system's internal logic is often opaque even to the people using it. In a legal system built on the notion that justice must not only be done but be seen done, opacity is not a small problem it is rather a problem in plain view.

WHAT THIS MEANS FOR SOMEONE STUDYING LAW RIGHT NOW:

I think the honest answer, as a student watching this unfold in real time, is that the "revolution versus risk" framing is a little misleading as if we get to pick one. What the last three years actually show is that the revolution and the risk are the same event, viewed from two different

distances. Close up, in a single case, AI hallucination looks like a catastrophic institutional failure. Zoomed out, across the system, tools like SUVAS and TERES look like exactly the kind of quiet, unglamorous efficiency gain that a backlog-choked judiciary desperately needs.

What seems to be settling into place, at least in principle, is a fairly sensible division of labour AI as research assistant, translator, and administrative aide never as decision-maker. The Kerala policy's insistence on mandatory human verification, and the Supreme Court's own language about "human in the loop, at every stage" in the Jammu and Kashmir Bank judgment, both point toward the same underlying idea. Machines can gather and organize; humans must still decide and remain accountable for that decision.

For law students specifically, I think this moment carries a fairly direct professional lesson, one that has nothing to do with technology and everything to do with basic legal discipline: verify your citations. It sounds almost too obvious to write down, but the NCLT case shows exactly how a lapse in that most basic habit checking whether a case actually exists before relying on it can unravel an entire judicial order and, with it, the reputations of everyone who touched the file. AI has not changed that duty. It has just made it more urgent, because the tools that help us research faster are the same tools that can fail silently, without any obvious signal that something has gone wrong.

CONCLUSION:

Indian courts aren't exactly choosing between "embracing AI" and "rejecting it", because that choice has been made, sort of quietly, through translation software, transcription tools, and case-management systems that most litigants will never even see. The real decision being made right now, the one hiding in plain sight, via documents like the 2026 Draft Regulations and judgments like *Pooja Ramesh Singh*³, is really about where the line sits between assistance and abdication, not about whether AI exists in the room.

The Bhopal metaphor the Supreme Court reached for was extreme, sure, but it was also, in its own peculiar way, precise. Some risks don't look dramatic until the moment they become catastrophic, and by then, the damage has already slipped through the system, in ways that are hard to trace back and undo.

³ supra

If Indian courts can hold that line with AI doing the translate, transcribe, and organize work, while making sure every finding, every citation, and every judgment stays answerable to a human being who can be asked questions and held accountable, then yes, this can really be a revolution in access to justice. Especially for the millions of litigants for whom the system has always felt too slow, too distant, and too much in a foreign language, in more ways than one.