

When Algorithms Meet Adjudication: The Emerging Jurisprudence of AI in Indian Courts



AASHWYN SINGH
Senior Associate



ANIKET KANHAUI
Associate

On a routine day in October 2025, the Bombay High Court confronted a problem that would have seemed absurd just a few years earlier: a tax assessment officer had cited three judicial precedents to justify adding Rs. 22.66 crores to a company's taxable income. The difficulty was simple, none of those judgments existed. They were phantoms and fabricated citations presented with the confidence of authoritative law. Whether the error arose from careless drafting or unverified AI-assisted research, the consequence was the same, the court had to set aside the assessment, not merely because it was wrong, but because it exposed a deeper risk, the creeping substitution of human judgment with machine output in contexts where lives, livelihoods, and the rule of law itself hang in balance.

This incident, documented in *KMG Wires Private Limited v. National Faceless Assessment Centre*¹, exemplifies the challenges confronting India's legal system as it navigates the promises and perils of artificial intelligence. The question is no longer whether AI will transform legal practice, it already has. The question is how the profession can harness its capabilities while preserving the human element in justice.

The Allure of Efficiency

The appeal of AI to an overburdened legal system is understandable and may even feel inevitable. India's judiciary faces millions of pending cases across various levels of courts. The promise that technology can accelerate research, automate routine documentation, and identify relevant precedents from vast databases offers hope for a system straining under its own weight. During the COVID-19 pandemic, virtual hearings proved that technology could expand access to justice, allowing litigants from remote areas to participate without the expense and hardship of travel.

Justice Surya Kant, speaking at the 29th National Law Conference in Kandy, Sri Lanka, acknowledged these benefits while observing²: "We are not replacing the lawyer or the judge; we are simply augmenting their reach and refining their capacity to serve." The distinction between augmentation and replacement is not semantic. It marks the difference between technology that strengthens human judgment and technology that displaces it.

Consider India's successful digital initiatives, the e-Supreme Court Reports portal providing free access to judgments in thirteen languages³, live streaming of constitutional court proceedings that enhances transparency

¹ *KMG Wires Private Limited v. National Faceless Assessment Centre*, WP(L) No. 24366 of 2025 (Bombay High Court, 6 October 2025) Available at https://www.livelaw.in/pdf_upload/kmg-wires-private-limited-627647.pdf.

² Justice Surya Kant, Keynote Address, "Technology in the Aid of the Legal Profession – A Global Perspective," 29th National Law Conference, Bar Association of Sri Lanka, Kandy (October 2025) Available at <https://www.livelaw.in/top-stories/ai-can-aid-lawyers-judges-cant-replace-them-justice-will-remain-a-human-enterprise-justice-surya-kant-307812>.

³ *Ibid.*

and e-filing systems that reduce the barriers of geography. These innovations empower rather than replace. They make information more accessible, proceedings more transparent, and justice more democratic. They reflect technology deployed in service of human values, not as a substitute for human judgment.

The Manipur Judge and the Google Search

Yet the boundary between assistance and abdication can blur with startling ease. The case of *Md. Zakir Hussain v. State of Manipur*⁴ illustrates this tension in unexpected ways. A Village Defence Force member challenged his dismissal, but the respondent state failed to provide adequate information about VDF service conditions despite specific court directions. Justice A. Guneshwar Sharma found himself in an uncomfortable position, recording: "In the circumstances, this Court is compelled to do extra research through Google and ChatGPT 3.5."⁵

What followed was remarkable. The judge used these tools to locate a crucial government order dated 18.10.2022 establishing VDF service conditions⁶, a document the State should have produced but did not. The order mandated show cause notices before adverse action, precisely the procedural protection the petitioner had been denied. On that basis, the court held that the dismissal violated principles of natural justice, leading to reinstatement with 50% back wages.

This case cuts both ways. On one hand, it demonstrates technology's capacity to uncover relevant information even when parties fail to present it, helping ensure that adjudication proceeds on a complete record. On the other hand, it raises hard questions about judicial function. Should judges have to conduct such electronic research to compensate for inadequate advocacy or administrative failures? And what happens when that research yields inaccurate or fabricated material, as occurred in *KMG Wires*?

Justice Sharma was careful to note that the Office Memorandum "is not mentioned in the counter-affidavit of the respondent in spite of specific direction of this Court." His resort to Google and ChatGPT, albeit successful in this instance, ventures into territory that deserves scrutiny. Had the tool produced fabricated information about VDF service conditions, the consequences could have been severe.

The Anatomy of AI Failure

The *KMG Wires* judgment provides one of the most detailed Indian judicial analyses of AI-related failures in legal contexts. Justices B.P. Colabawalla and Amit S. Jamsandekar confronted an assessment order that made two primary additions to the petitioner's income⁷. First, it disallowed purchases

⁴ *Md. Zakir Hussain v. State of Manipur*, WP(C) No. 70 of 2023 (Manipur High Court, 23 May 2024) Available at https://www.livelaw.in/pdf_upload/asds-541351.pdf.

⁵ *Ibid.*, at para 7.

⁶ *Ibid.*, at para 9. The Office Memorandum was issued by the Home Department, Government of Manipur.

⁷ *KMG Wires* (n 1).

worth Rs. 2.15 crores from Dhanlaxmi Metal Industries on the basis that the supplier never responded to notices. In fact, the supplier had filed a detailed hundred-page response with supporting documents on 08.03.2025, well before the assessment order. The assessing officer simply failed to consider material already on the file.

The second addition was more troubling. The officer added Rs. 22.66 crores as the peak balance in directors' loans, relying on three judicial precedents for the proposition that opening balances should be included. But those precedents did not exist. The court's warning deserves quotation: "In this era of Artificial Intelligence ('AI'), one tends to place much reliance on the results thrown open by the system. However, when one is exercising quasi-judicial functions, it goes without saying that such results [which are thrown open by AI] are not to be blindly relied upon, but the same should be duly cross-verified before using them."

The court did not merely criticize, it identified the mechanism of failure. AI systems can generate convincing-looking citations, complete with case names, court references, and legal propositions that may be entirely fabricated, yet presented with an authority that invites reliance. This phenomenon, often described as "hallucination", is documented even in legal-domain systems. A Stanford HAI study found hallucinations in roughly one out of six (or more) benchmark legal research queries⁸. Beyond controlled benchmarks, public trackers and court reporting show a sharp rise in judicial decisions addressing fabricated citations and AI-generated misinformation across jurisdictions⁹. The KMG Wires episode therefore sits within a wider pattern, when unverified AI output enters a legal record, it can distort the very inputs on which adjudication depends.

What makes KMG Wires particularly instructive is the court's nuanced response. Rather than condemning AI technology wholesale, the judges distinguished between different failures. The overlooked supplier response represented human negligence, not AI error. The fabricated precedents reflected unchecked reliance on generated output. The court set aside the assessment not as punishment but to ensure natural justice, remanding the matter with specific directions i.e., clear show cause notices, sufficient response time, mandatory personal hearings, and seven days' notice before relying on judicial decisions.

Why "Hallucinations" Happen

Large language models (LLMs) do not conduct legal research in the conventional sense. They generate responses by predicting likely text, rather than by locating and confirming primary legal material. As a result, when asked for precedent, they may output a citation that appears correct in form but has no real judgment behind it, precisely the risk that surfaced in KMG Wires.

A further difficulty is that these tools can reproduce the style of legal writing

⁸ Varun Magesh et al., "AI on Trial: Legal Models Hallucinate in 1 out of 6 (or More) Benchmarking Queries," Stanford HAI (23 May 2024) Available at <https://hai.stanford.edu/news/ai-trial-legal-models-hallucinate-1-out-6-or-more-benchmarking-queries>.

⁹ See Damien Charlotin, "Legal decisions involving AI hallucinations" (database), available at <https://www.damiencharlotin.com/hallucinations/>.

with ease. They can mimic judicial tone, bench and party descriptions, and citation formats in a manner that makes the output appear reliable even when it is not. The risk is higher when the prompt itself assumes that supporting authority exists (for instance, "find judgments holding X"), because the model may supply a plausible answer to match the assumption. For these reasons, the only safe approach is strict verification. Any AI-generated citation or proposition must be checked against primary sources before it is relied upon.

The Principle Behind the Practice

These cases illuminate broader principles about AI's proper role in legal systems. The Supreme Court's observation in *Annaya Kocha Shetty v. Laxmibai Narayan Satose*¹⁰ addresses a related problem: computer-generated pleadings that prioritise length over substance. Justice S.V.N. Bhatti noted that courts increasingly face "AI-generated or computer-generated statements" characterised by unnecessary verbosity. The judgment invoked Abraham Lincoln's sardonic description of a lawyer who "can compress the most words into the smallest ideas of any man I ever met"; that nineteenth-century critique applies with renewed force to AI-generated legal documents.

The deeper point concerns the relationship between form and substance in legal reasoning. Traditional legal training emphasises concision, precision, and relevance. Every word should advance the argument; every citation should genuinely support a proposition. AI tools, by contrast, can generate text that appears sophisticated but lacks the discipline of human legal writing, producing thirty pages where three would suffice.

Justice Bhatti observed that "lengthy pleadings and avoidable evidence are well within the scrutiny of trial courts, and, at the right stage, must be regulated within the four corners of the law." This is more than stylistic preference. Excessive length burdens courts, delays justice, and can hinder understanding. The judgment calls for invoking Order VI Rule 16 of the Code of Civil Procedure to strike unnecessary pleadings, a remedy that may need more frequent use in the age of AI-generated documents.

Wisdom from the Bench in Chandigarh

The Punjab and Haryana High Court has addressed whether lawyers should use AI tools in the courtroom itself. In *RXXXXXX v. State of Haryana (CRM-M-31392-2025)*¹¹, Justice Sanjay Vashisth expressed concern about members of the Bar using mobile phones during hearings to search for information and answer judicial queries, material that ought to have been prepared in advance.

Justice Vashisth noted that in *Ravneet Singh Sandhu v. UT of Chandigarh*¹², he

¹⁰ *Annaya Kocha Shetty v. Laxmibai Narayan Satose*, Civil Appeal No. 84 of 2019 (Supreme Court of India, 8 April 2025), 2025 INSC 466 Available at https://www.livelaw.in/pdf_upload/38628201813150160771judgement08-apr-2025-595054.pdf.

¹¹ *RXXXXXX v. State of Haryana*, CRM-M-31392-2025 (Punjab & Haryana High Court, interim order dated 30 September 2025) Available at https://courtbook.in/pdf_upload/crm-m31392202530092025interimorder-624098.pdf.

¹² *Ibid.*

had seized a mobile phone being used during the hearing. The concern was professional: "Even sometimes, proceedings are to be stalled, awaiting the answer, which would come only after retrieving information from such mobile phone." The order warned lawyers against compelling the court "to pass any harsh order" on account of repeated mobile-phone use to update themselves through artificial intelligence/online platforms/Google information."

This reflects deeper concerns about preparation, professionalism, and the attorney-client relationship. A lawyer relying on real-time online searches during argument has not mastered the case. More troubling is the risk that, if an AI tool supplies fabricated information, the consequences extend beyond embarrassment to potential miscarriage of justice.

Guidelines from Kerala: A Template for Responsible Use

Recognising these challenges, the Kerala High Court issued comprehensive guidelines in July 2025 titled "Policy Regarding Use of Artificial Intelligence Tools in District Judiciary"¹³. This document represents one of the most systematic Indian judicial responses to AI adoption.

The policy identifies several core principles. First, transparency, fairness, accountability, and confidentiality remain "integral aspects of judicial administration, which shall not be compromised by the use of AI tools." Second, cloud-based services such as ChatGPT and DeepSeek pose confidentiality risks because user inputs may be accessed by service providers, requiring avoidance except for specifically approved applications.

Most critically, the Kerala guidelines mandate that "AI tools shall not be used to arrive at any findings, reliefs, orders or judgment under any circumstances¹⁴, as the responsibility for the content and integrity of the judicial order, judgment or any part thereof lies fully with the Judges." This squarely frames AI as assistive, never decisional.

The policy also insists on process safeguards: verification of AI-generated outputs (especially citations); qualified translator verification for AI-translated texts; human supervision for administrative tasks; auditability of AI tool usage; and mandatory training on AI's ethical, legal, and technical dimensions. It treats inappropriate AI use as a matter of professional responsibility, with disciplinary consequences.

Justice Surya Kant's Framework

Justice Surya Kant's Kandy address synthesised these concerns into an overarching philosophy: "Technology may be a powerful ally, but justice will always remain a profoundly human enterprise¹⁵." Technology should enhance rather than displace human capacity: "Let technology be the guide, and the human govern." Certain aspects cannot be automated: "The judge's

¹³ Kerala High Court, "Policy Regarding Use of Artificial Intelligence Tools in District Judiciary" (July 2025) Available at https://images.assettype.com/theleaflet/2025-07-22/mt4bw6n7/Kerala_HC_AI_Guidelines.pdf.

¹⁴ *Ibid.*

¹⁵ Justice Surya Kant (n 2).

discernment, the advocate's reasoning, the litigant's dignity, and the empathy that animates every fair trial, these are the living fibres of justice that no machine can replicate."

Data may inform but must never dictate decisions. Justice Kant warned against over-reliance on algorithmic outputs in matters requiring moral judgment or contextual understanding. He emphasised that AI is prone to inaccuracy and bias, making human oversight non-negotiable, particularly in light of errors or "hallucinations" that can have serious consequences.

Finally, he identified the digital divide as a critical challenge. Technology's benefits must not accrue only to elite practitioners while those serving marginalised communities lack access or training. This requires sustained capacity building and carefully designed partnerships that expand access without compromising integrity.

International Parallels

India's judicial concerns echo those worldwide. Bar bodies and courts in several jurisdictions have issued ethics guidance on generative AI, emphasising competence, confidentiality, supervision, and verification. For instance, the American Bar Association's Formal Opinion 512 links generative AI use to core professional duties¹⁶, including confidentiality and the duty of competence. International reporting has also chronicled a growing pattern of sanctions and judicial frustration where lawyers cite hallucinated or fabricated authorities.

Regulators have also responded. The European Union's AI Act treats certain uses of AI in the administration of justice as¹⁷ high-risk, triggering obligations around transparency, risk management, and human oversight. At the technical level, the error profile is real: legal-domain benchmarking has found hallucinations in roughly one in six queries¹⁸, while studies of general-purpose chatbots on legal tasks have reported substantially higher error rates in some settings. These figures do not argue against adoption; they argue for guardrails.

Emerging Indian Regulatory Direction

In India, the approach to AI in courts is taking shape largely through judicial guidance and court-led policy, guided by the need to protect independence, confidentiality and procedural fairness. A clear articulation is found in the Kerala High Court's 2025 Policy Regarding Use of Artificial Intelligence Tools in District Judiciary-controlled adoption of technology, one that may improve access and efficiency, but keeps adjudication and accountability firmly with

¹⁶ American Bar Association, *Formal Opinion 512* (29 July 2024) Available at <https://www.americanbar.org/news/abanews/aba-news-archives/2024/07/aba-issues-first-ethics-guidance-ai-tools/>.

¹⁷ European Commission, *EU AI Act; see Annex III on High-Risk AI Systems (Administration of Justice)* Available at <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>; see also <https://ai-act-service-desk.ec.europa.eu/en/ai-act/annex-3>.

¹⁸ Stanford HAI (n 8); see also Varun Magesh et al., "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," *Journal of Empirical Legal Studies* (2025).

the human judge¹⁹, which proceeds on a simple principle: AI may assist, but it cannot decide. The policy cautions against using AI output for findings, reasons or final orders, stresses verification of any AI-generated material (particularly citations), flags confidentiality risks in cloud-based tools, and requires human oversight even for ancillary tasks such as translations and administrative support. High Courts have also begun signaling expectations for responsible courtroom use, including warnings against real-time AI searching during oral arguments in a manner that substitutes preparation with unverified assistance. Taken together, these developments point to a regulated, court-

Key Recommendations for Courts

Courts may consider adopting a few basic safeguards to ensure that AI improves efficiency without compromising accuracy or fairness.

- First, a clear “verify-before-use” rule should be institutionalized: any AI-generated citation, quotation or proposition must be checked against the primary source before it is relied upon.
- Second, AI should be expressly treated as assistive and never decisional, while it may support drafting, translation or administrative tasks, findings, reasons and final orders must remain the independent work of the judge.
- Third, courts may maintain an “approved tools” protocol, favouring systems with clear data-handling practices, audit trails and reliable citation support, while discouraging use of general cloud chatbots for confidential casework.
- Fourth, training and simple bench checklists can help judges, registries and the Bar identify recurring risks such as fabricated authorities, mistranslations and data leakage.
- Fifth, practice directions may regulate real-time AI use during oral hearings so that advocacy is not reduced to unverified, on-the-spot outputs.
- Finally, to protect the integrity of the record, parties may be required to file authenticated copies of all authorities relied upon, and courts may discourage annexures that contain AI-generated extracts without verifiable sources.

Policy priorities for the next phase

Looking ahead, the immediate need is a coherent court-led framework that builds on the Kerala High Court’s policy and can be adopted, with suitable local adaptations, across High Courts and the district judiciary. In parallel, courts should move towards secure, court-controlled research and translation systems that rely on authenticated databases of judgments and statutes, so that confidentiality is protected and fabricated authorities are avoided. Finally, a simple incident-reporting mechanism for AI-related errors, such as false citations or mistranslations, can help identify patterns,

¹⁹ Kerala High Court (n 13).

issue corrective guidance, and refine safeguards over time.

The Path Forward

Several practical principles emerge from Indian cases and international experience.

- Verification is non-negotiable: Every AI-generated citation must be checked against primary sources. The KMG Wires court was unequivocal: results "should be duly cross-verified before using them."
- Transparency should be encouraged: Where AI tools materially assist with drafting or research, courts may consider requiring appropriate disclosure, at least in sensitive matters. International practice on disclosure varies, but the direction of travel is toward clearer norms.
- Professional responsibility rules should be clarified: Bar associations can make explicit that uncritical reliance on unverified AI output breaches duties of competence and diligence. The Kerala guidelines' treatment of AI misuse as a disciplinary matter offers a workable model.
- Judicial training must expand: As the Md. Zakir Hussain case shows, judges themselves may turn to AI tools. Structured training can help courts understand both capabilities and limits and reduce the risk of inadvertent reliance on flawed outputs.
- Approved tool ecosystems matter: Following Kerala's model, courts should identify and vet specific AI applications, and specify what is prohibited. The distinction drawn between approved tools and general cloud-based services is a useful starting point.
- Confidentiality must be paramount: Justice Kant noted that "the confidentiality, privilege, data integrity, and cybersecurity of their clients must remain sacrosanct²⁰." Institutional policies should reflect this, particularly for cloud-based AI systems.
- Human judgment must remain supreme: The Kerala policy's prohibition on using AI to arrive at "any findings, reliefs, orders or judgment" establishes the proper hierarchy: AI assists; humans decide.

Conclusion

The cases emerging from Indian courts in 2025 reveal a judiciary neither reflexively hostile to technology nor naively embracing it. KMG Wires states the point crisply: when exercising quasi-judicial functions, blind reliance on generated output is incompatible with the responsibility that adjudication demands. Justice Surya Kant's metaphor remains apt: technology as guide, humans to govern.

²⁰ Justice Surya Kant (n 2).

In *Annaya Kocha Shetty*, Justice Bhatti reminded the profession that "the effort of pleading and evidence should be to be concise to the cause and must not confuse the cause." The principle applies with equal force to AI-generated content.

The Manipur case demonstrates both possibility and peril. Justice Sharma's use of digital tools to find a government order served justice, but the same approach could mislead if the tool hallucinated a non-existent regulation. The Chandigarh court's concern about lawyers using mobile devices during hearings stems from the same root: professional standards can erode when convenience substitutes for preparation and verification.

Justice Surya Kant captured the essential insight: "Artificial intelligence may assist in researching authorities, generating drafts, or highlighting inconsistencies, but it cannot perceive the tremor in a witness's voice, the anguish behind a petition, or the moral weight of a decision." Justice ultimately depends on distinctively human capacities for empathy, moral reasoning, and contextual judgment. That recognition should guide every decision about AI deployment in legal systems.

Technology can make legal research faster, court access broader, and information more transparent. But it cannot, and should not, make the judgments that define justice itself. That responsibility remains, as it must, human.