

# Artificial Intelligence in Law and Legal Research: Capability, Practice, And the Limits of Delegation

ABHINAV BISWAL  
*Institute MS Law University*

## I. INTRODUCTION

Every profession that trades in information eventually has to reckon with a machine that reads faster than its practitioners. Law is having that reckoning now. For a discipline built on precedent, citation and the patient accumulation of authority, the arrival of large language models capable of drafting a memorandum, summarising a five-hundred-page case file, or translating a judgment into Odia in seconds is not a minor efficiency gain. It is a change in the unit of labour itself. This report takes stock of where that change actually stands in mid-2026 — not the speculative version of AI-in-law that circulated a few years ago, but the version now sitting inside law firms, legal research platforms, and, notably, the Indian judiciary itself.

The honest starting point is that the technology has proved genuinely useful and genuinely unreliable in the same breath, often in the same document. That tension — between AI as a serious research accelerant and AI as a source of fabricated authority — is the organising thread of what follows.

## II. FROM BOOLEAN SEARCH TO GENERATIVE RESEARCH

Legal research has been digitised for three decades, but the digitisation was, until recently, a matter of indexing. Westlaw, Manupatra, SCC Online and Lexis built enormous, searchable libraries of case law and statute, and the researcher's skill lay in constructing the right keyword string and reading judiciously through the results. What has changed since roughly 2023 is that the platforms themselves now read for you. Tools such as Lexis+ AI, Westlaw's CoCounsel, and Harvey generate a synthesised answer to a legal question, complete with citations, drawn from the underlying database using

retrieval-augmented generation — a method that grounds the model's output in retrieved source documents rather than in whatever the model happens to remember from training.

The appeal is obvious to anyone who has spent an afternoon shepardising authorities by hand. A first-pass memorandum that once took a junior associate a day can now take an hour of drafting and several hours of verification. Due diligence review, contract redlining, and the identification of conflicting precedent are increasingly AI-assisted first, human-reviewed second. What has not changed — and this is the point courts keep having to restate — is that verification remains a human, non-delegable duty. The tool changes the starting point of the work, not the standard of care owed at the end of it.

## III. THE INDIAN JUDICIARY'S OWN EXPERIMENT

India's courts have not been passive observers of this shift; they have built their own infrastructure for it, and the pace of that build-out over the past two years is worth setting out in some detail, since it rarely gets the attention that commercial legal-tech does.

### 3.1 SUPACE and the research layer

SUPACE — the Supreme Court Portal for Assistance in Court Efficiency — was conceived as a research aid for judges rather than a decision-making system. In its 2026 form it performs factual-matrix extraction from voluminous case files and precedent mapping through retrieval-augmented generation, surfacing relevant case law rather than recommending an outcome. The distinction is deliberate and has been preserved carefully in every official description of the tool: SUPACE identifies precedent, it does not decide anything. Judges retain the reasoning.

### 3.2 SUVAS and linguistic access

SUVAS, the Supreme Court Vidhik Anuvaad Software, addresses a different and in some ways more consequential problem — the fact that Indian court proceedings and judgments have historically been locked in English in a country where the overwhelming majority of litigants are not comfortable reading legal English. Launched on Constitution Day 2019, SUVAS is a domain-specific neural machine translation system trained on judicial text rather than general prose, and by 2026 it supports bidirectional translation across nineteen Indian languages, having translated tens of thousands of Supreme Court judgments into Hindi and other regional languages. High Courts including Allahabad, Delhi and Kerala now publish judgments through it. For a litigant in a district town who has never seen the inside of a courtroom in Cuttack, let alone New Delhi, that is not a convenience feature — it is closer to a precondition for actually understanding the order that governs their case.

### 3.3 TERES, LegRAA, and the wider e-Courts build-out

Sitting alongside SUPACE and SUVAS are TERES, a real-time transcription system used in Constitution Bench matters, and LegRAA, a newer generative-AI research tool. All of this sits under Phase III of the e-Courts Mission Mode Project (2023–2027), which has been explicitly framed around AI, machine learning, OCR and natural language processing for case management, research and translation, with a substantial allocation of public funds behind it. The institutional instinct throughout has been compartmentalisation rather than consolidation — keeping translation, transcription and research as separate, narrowly scoped tools instead of building one opaque system that does everything. That is a sensible design choice for an institution whose legitimacy depends on being able to explain its own reasoning.

### 3.4 The 2026 Draft Regulations

The most significant recent development is regulatory rather than technical. On 3 June 2026, the Supreme Court's AI Committee published the Draft Regulations for Use of Artificial Intelligence in Courts, 2026 — India's first binding, national

framework for judicial AI, applicable across the Supreme Court, every High Court, and tribunals performing adjudicatory functions. Its governing principles are worth stating plainly because they read as a template for how any institution ought to think about this technology: AI must remain strictly assistive and subordinate to judicial authority; final determination on law and fact rests exclusively with judicial officers; opaque or unexplainable systems attract heightened scrutiny; and accountability for AI-assisted output cannot be outsourced to the software. AI must actively avoid discrimination; opaque and unexplainable systems face heightened scrutiny; and accountability for AI use remains with the judicial officer. — Draft Regulations for Use of AI in Courts, 2026 — summarised principles

In March 2026, the Supreme Court went further and characterised the filing of AI-generated citations that turn out to be fabricated as professional misconduct — a signal that the judiciary intends to treat this not as a novel technical problem requiring a novel excuse, but as an old professional duty (verify what you sign) applied to a new tool.

## IV. THE OTHER SIDE OF THE LEDGER: HALLUCINATION AND SANCTION

It would be a disservice to write this report without dwelling on the failure mode, because the failure mode is now extensively documented and shows no sign of tapering off. The now-canonical case is *Mata v. Avianca* (2023), in which a New York attorney submitted a brief citing six cases that ChatGPT had invented outright, complete with fabricated quotations and case numbers; the presiding judge imposed a five-thousand-dollar sanction that, at the time, felt severe. Three years on, it reads as the opening entry in a very long list.

A public tracker maintained by the researcher Damien Charlotin now records roughly one and a half thousand court decisions worldwide — more than a thousand of them in the United States alone — in which a party relied on AI-hallucinated material and a court had to respond. The penalties have escalated well past that first five-thousand-dollar figure. In February 2026, an Omaha attorney filed a

divorce appeal brief in which fifty-seven of sixty-three citations were defective, twenty of them wholly invented; when asked directly whether he had used AI, he first denied it and later admitted it, and the Nebraska Supreme Court suspended him pending a disciplinary investigation. In May 2026, a federal judge in Oregon imposed a one-hundred-and-ten-thousand-dollar sanction — the largest AI-hallucination penalty in American legal history — after two lawyers submitted twenty-three fabricated citations and eight invented quotations. The Sixth Circuit and a California appellate panel have separately converged on almost identical language: no filing in any court should contain a citation the responsible attorney has not personally read and verified.

Two details from this body of cases are worth pulling out because they cut against comfortable assumptions. First, the tool responsible is usually a general-purpose consumer chatbot rather than a purpose-built legal research product — but purpose-built tools are not immune either; in one tracked case counsel using established legal-research platforms still filed sixteen fabricated quotations. Second, courts consistently punish the cover-up more severely than the original error. Lawyers who admitted the mistake candidly and corrected the record have generally fared better than those who denied using AI and were caught, or who repeated the error after being warned. The lesson embedded in that pattern is not really about artificial intelligence at all — it is about candour with the court, which has always been the more fundamental professional obligation.

#### V. ACCESS TO JUSTICE AND THE DIGITAL DIVIDE

A purely doctrinal account of AI in law risks missing the sociological question underneath it, which is who actually benefits from this shift and who is left further behind by it. The optimistic case for tools like SUVAS is genuine: a litigant who could previously not read the order governing their own case can now do so in their own language, and that is a real expansion of access to justice, not a cosmetic one. But the same digitisation drive assumes a baseline of connectivity that a large share of the country does not

have. Estimates cited in recent research on the e-Courts project suggest roughly seventy percent of the population still lacks reliable digital connectivity — which means the population most in need of translated, simplified access to the legal system is disproportionately the population least able to reach the interface delivering it.

This is the familiar pattern of technology-led reform in unequal societies: the tool is neutral, but the distribution of who can use it is not. Any serious evaluation of AI's contribution to access to justice in India has to hold both facts at once — that the tools genuinely help the litigants who can reach them, and that reach itself is unevenly distributed along lines that map onto existing patterns of rural-urban and class disadvantage. A legal-technology reform that is not paired with a connectivity and digital-literacy reform will tend to widen the gap between litigants who are already well served by the system and those who are not.

#### VI. WHERE THIS LEAVES LEGAL PRACTICE AND LEGAL EDUCATION

For practising lawyers, the emerging professional norm is reasonably settled even if compliance with it is not: AI may generate a first draft, a research summary, or a list of candidate authorities, but every citation, quotation and factual claim that reaches a filed document must be independently verified by the lawyer whose name is on it. That is not a new ethical standard invented for the AI era — it restates the existing duty of competence and candour to the tribunal — but the tools have made it easier to breach that duty at scale, which is precisely why courts have begun treating breaches of it as misconduct rather than as excusable technical error.

For legal education and for early-career researchers, the more interesting shift is in what skill is actually scarce. When a machine can produce a serviceable first draft of a research memorandum in minutes, the differentiating skill stops being the ability to locate a case and becomes the ability to evaluate whether the case actually says what it is claimed to say, to spot the gap between a plausible-sounding proposition and a supportable one, and to exercise the kind of

judgment that a retrieval system cannot substitute for. That is, if anything, a stronger argument for rigorous doctrinal training than existed before these tools arrived — not a weaker one.

For an institution like the judiciary, the emerging model — visible in India's compartmentalised, assistive, subordinate-to-the-judge design — is probably the right template for other public institutions considering similar adoption: narrow the tool's function, keep a human decision-maker unambiguously in the loop, and build the accountability framework before the scandal that would otherwise force one into existence.

## VII. CONCLUSION

The technology underlying AI-assisted legal research is not going away, and pretending otherwise serves no one. What the past three years have demonstrated is that the technology is best understood as a very capable, very confident research assistant that has never once been to law school and does not know the difference between a real case and a plausible one. Used that way — as a starting point subject to the same verification any human junior's work would receive — it has already made legal research faster and, through tools like SUVAS, meaningfully more accessible. Used as a substitute for that verification, it has produced a genuinely embarrassing catalogue of sanctions, suspensions, and dismissed cases across jurisdictions.

The institutions handling this transition most credibly — India's judiciary among them — are the ones treating AI as an instrument that extends judicial and professional capacity without ever being permitted to exercise judgment on its own behalf. That, more than any particular tool, is the principle worth carrying forward.