
AI AS A TOOL, NOT A JUDGE: EXAMINING THE BOUNDARY BETWEEN MECHANICAL TASKS AND HUMAN JUDICIAL INTELLIGENCE IN THE INDIAN LEGAL SYSTEM

Ms. Soumya Mittal, Independent Legal Researcher, Bilaspur, Chhattisgarh, India

ABSTRACT

Artificial Intelligence (AI) is increasingly present in India's legal landscape, from translating judgments and searching databases to assisting with contract drafting. Yet a critical question remains unanswered in most public discussions: where exactly does AI's competence end and human intelligence begin? This paper argues that AI is genuinely effective when it performs mechanical, repetitive, volume-based tasks that require no judgment - what the paper calls 'backend labor.' However, AI fails, and fails dangerously, when it is asked to perform tasks that require contextual understanding, legal intuition, ethical reasoning, or human judgment. The paper draws on actual judicial practice in India, including the Supreme Court's SUVAS translation system, the SUPACE case management tool, the growing fake citation crisis, and the ongoing debate around AI-assisted adjudication, to demonstrate that the line between what AI can and cannot do is not a matter of technology alone - it is a matter of jurisprudence, constitutional principle, and the very meaning of justice.

Keywords: Artificial Intelligence; Human Intelligence; Indian Judiciary; Judicial Decision-Making; Legal Technology.

I. Introduction

In 2026, Artificial Intelligence has arrived inside India's courts - not in the dramatic way science fiction imagined, not as a robot judge sitting on a bench, but quietly, in the background, doing work that nobody used to notice and everybody used to spend enormous time on. A junior advocate once spent two full days searching for a precedent. A court clerk used to type out every word of a deposition by hand. A translator used to work for weeks converting a Supreme Court judgment into Hindi. Today, AI can do each of these tasks in minutes. This is remarkable. But it is also being seriously misunderstood.

The problem is that when people see AI perform these impressive background tasks efficiently, they assume it can do everything else too - that it can identify which case is truly relevant to a dispute, that it can draft a contract that genuinely protects a client's specific interests under Indian law, that it can predict the correct outcome of a case, and perhaps one day, pronounce judgment. This assumption is not just wrong. It is dangerous.

This paper makes one central argument: AI is a powerful tool for mechanical, repetitive, volume-heavy tasks that require no original judgment - what this paper calls 'backend legal labor.' But the moment a task requires contextual understanding, ethical reasoning, application of competing legal principles to specific human facts, or any form of what we ordinarily call wisdom, AI is not merely inadequate - it is structurally incapable, because it is not a thinking being.

The Indian legal system is an especially important context for this discussion. With over 5.4 crore cases pending across Indian courts, the pressure to use AI to speed things up is real and urgent.¹ But India also has a rich constitutional tradition, a deeply developed jurisprudence around fundamental rights, and a judiciary whose legitimacy rests on the perception that a human mind - trained, experienced, and accountable - has weighed the facts and the law and made a reasoned choice. That legitimacy cannot survive automation.

This paper proceeds in six parts. Part II explains what is meant by 'backend labor' versus 'applied intelligence' in the legal context. Parts III through V examine three concrete areas - translation, legal research, and contract drafting - where this distinction plays out in practice.

¹. National Judicial Data Grid (NJDG), *Pendency Statistics (July 2025)*, Sup. Ct. of India, <https://njdg.ecourts.gov.in>

Part VI addresses the most urgent and sensitive question: the use of AI in judicial decision-making. Part VII concludes.

II. The Core Distinction: Backend Labor vs. Applied Intelligence

A. What is 'Backend Labor' in Law?

Every legal task can be broken down into at least two layers. The first layer is what this paper calls 'backend labor' - work that is necessary, time-consuming, and requires accuracy, but does not require original judgment or the application of law to specific human facts. Consider some examples:

Translating a judgment from English to Tamil is backend labor. The meaning is fixed. The task is to convert it into another language accurately. No one is deciding what the judgment means. The judgment already decided that. The translator is just making it readable to more people.

Searching a database of five million judgments for every case that mentions 'Section 138 of the Negotiable Instruments Act' is backend labor. The task is retrieval. The computer is looking for a string of text across an enormous archive. No judgment is being applied about whether those cases are legally relevant to today's dispute.

Generating a first draft of a standard Non-Disclosure Agreement based on a template is backend labour. The basic clauses of an NDA are well-established. The structure is known. The AI is filling in blanks based on patterns it has learned from thousands of similar documents.

What makes all of these tasks 'labor' rather than 'intelligence' is that they can be fully defined in advance. You can write a complete specification of what success looks like before the task begins, and that specification does not depend on judgment calls that require understanding the specific human situation at hand.

B. What is 'Applied Legal Intelligence'?

The second layer is what this paper calls 'applied legal intelligence' - work that requires a trained human mind to evaluate, weigh, and reason about the specific facts of a specific situation, in light of applicable law, and arrive at a judgment that reflects that reasoning.

From the same examples: once you have retrieved all cases mentioning Section 138, a lawyer must read them and determine which ones are actually relevant to this particular dispute between this particular creditor and this particular debtor - given the specific history of their transaction, the specific defences available, the specific bench preferences, and dozens of other contextual factors. That determination requires applied intelligence.

Once you have a first draft NDA, a lawyer must decide whether its specific clauses actually protect this client in this business relationship, in light of how Indian courts interpret these clauses, the power dynamics between the parties, the specific risks involved in this industry, and what modifications are needed. That requires applied intelligence.

The critical insight this paper proposes is this: AI can perform backend labor effectively. AI cannot perform applied legal intelligence, because applied legal intelligence is not a more complicated version of pattern-matching - it is something qualitatively different.

A machine can find every needle in a haystack. Only a lawyer can tell you which needle is the right one for this patient.

C. Why This Distinction Matters Especially in India

India's legal system is unusually complex even by global standards. It operates across 29 states and 8 union territories, with state-specific laws, different High Court jurisdictions, multiple languages, overlapping regulatory regimes, colonial-era statutes still in force alongside modern legislation, and a constitutional jurisprudence that is among the most developed in the world. Applied legal intelligence in India must navigate all of this simultaneously.

India also has a massive access-to-justice gap. Most citizens who need legal help cannot afford experienced lawyers. Into this gap, AI-based tools are being introduced - sometimes to do backend labour helpfully, and sometimes to substitute for the applied intelligence that clients urgently need but cannot access. The consequences of that substitution failing are not abstract. They land on real people in real disputes.

III. Translation: Where AI Genuinely Works

A. The Problem It Solves

India is a country of twenty-two officially recognized languages and hundreds of dialects. Yet

under Article 348(1)(a) of the Constitution, all proceedings before the Supreme Court and High Courts must be conducted in English² - a language that the overwhelming majority of India's 1.4 billion people either do not speak or do not speak well enough to understand complex legal reasoning.

The result has been a profound access-to-justice problem. A farmer from rural Andhra Pradesh who wins a landmark Supreme Court judgment about land rights cannot read that judgment. He must rely entirely on what his lawyer tells him it says. He has no independent ability to verify, understand, or explain to his family what was decided and why. This is not a small problem. It goes to the heart of whether justice is truly accessible in a democratic republic.

B. SUVAS: The Tool and Its Achievements

The Supreme Court of India developed SUVAS - the Supreme Court Vidhik Anuvaad Software - specifically to address this problem. SUVAS is not a general-purpose translation tool. It is an AI system trained on legal texts, designed to understand judicial language, preserve legal terminology, and translate Supreme Court judgments between English and Indian languages with contextual accuracy.³

As of 2026, SUVAS supports translation between English and 19 Indian languages. By 2025, it had translated over 36,000 Supreme Court judgments - more than 22,000 into Hindi alone, with thousands more in Punjabi, Kannada, Tamil, Gujarati, and other languages. These translations are publicly available on the e-SCR portal.⁴

This is a genuine achievement. The scale is significant. Before SUVAS, the idea of translating tens of thousands of Supreme Court judgments into regional languages would have required decades of human translator time and enormous financial resources. SUVAS has accomplished this in a fraction of the time and cost.

² INDIA CONST. art. 348, cl. 2 (permitting a State Governor, with Presidential consent, to authorize use of Hindi or another state language in that State's High Court). *See also* INDIA CONST. sch. VIII (listing 22 official languages).

³ *Supreme Court Vidhik Anuvaad Software (SUVAS)*, Sup. Ct. of India, <https://main.sci.gov.in> (launched Nov. 26, 2019)

⁴ *SUVAS Translation Statistics (2025)*, e-SCR Portal, Sup. Ct. of India, <https://judgments.ecourts.gov.in> (last visited June 29, 2026) (recording 36,271 total translated judgments, including 22,396 in Hindi, 3,572 in Punjabi, 1,899 in Kannada, 1,172 in Tamil, 1,112 in Gujarati, and across 13 additional languages).

C. Why Translation Is 'Backend Labor'

Translation - at least in the context of judicial documents where the meaning is already fixed - is a paradigmatic example of backend labor. The Supreme Court has already decided what a judgment means. The words of the judgment express that decision. The translation task is to render those words accurately in another language. No new meaning is being created. No judgment is being made about what the law should be.

A skilled AI system trained on legal language can do this task well. It does not need to understand why the Court decided the way it did. It does not need to evaluate whether the decision was correct. It simply needs to find the most accurate equivalents in the target language for what the original text says. This is a sophisticated pattern-matching task, and sophisticated AI systems are good at sophisticated pattern-matching.

D. The Limits That Remain

SUVAS's achievements come with important caveats that the Supreme Court itself acknowledges. Translated judgments carry a disclaimer that the Registry is not responsible for inaccuracies.⁵ This matters enormously in practice.

Legal language contains terms that do not have direct equivalents in regional languages. The word 'injunction,' for example, carries a very specific legal meaning that has no perfect equivalent in many Indian languages. When SUVAS translates it, it must choose a word or phrase that is close - but 'close' in legal matters can create real confusion. A litigant reading the translation might understand something slightly different from what the Court actually held.

Additionally, translated versions are informational only. Under the constitutional framework, they cannot be cited in proceedings. High Courts, except in certain Hindi-speaking states, still conduct proceedings in English only. So while SUVAS has made judgments more accessible to citizens for understanding, it has not changed the fundamental linguistic barrier that prevents full participation in the legal system.

The lesson from SUVAS, therefore, is precise: AI is effective at the mechanical task of

⁵*Disclaimer*, Supreme Court Vidhik Anuvaad Software (SUVAS), Sup. Ct. of India, <https://main.sci.gov.in> (last visited June 29, 2026) (stating that translated versions are issued for informational purposes only and the Registry accepts no responsibility for inaccuracies). *See also* INDIA CONST. art. 348, cl. 1(a) (establishing English originals as the sole authoritative text).

converting language. It is not a substitute for the deeper reform of making Indian courts genuinely multilingual - which is a policy, constitutional, and human challenge, not a technological one.

IV. Legal Research: Where AI Helps and Where It Harms

A. The Scale of the Problem

India has one of the most voluminous bodies of case law in the world. The Supreme Court alone has delivered over 35,000 judgments, with High Courts across the country adding millions more. When a lawyer needs to research a point of law, the universe of potentially relevant material is staggering. A junior associate at a large firm might spend two or three full working days simply retrieving and reading cases before being able to begin the actual analysis.

This is where AI-assisted legal research has been most widely adopted in Indian legal practice. Platforms like Manupatra, SCC Online, CaseMine, and BharatLaw.AI now use machine learning and natural language processing⁶ to search enormous legal databases, retrieve relevant cases, generate summaries, and present research in organized formats. Tools like CaseMine's AMICUS module can visually map how cases relate to each other - which cases cite which, how a legal principle has evolved over decades.

B. What AI Does Well: The Retrieval Layer

The retrieval layer of legal research - finding cases that mention a specific legal provision, contain a specific phrase, or deal with a specific subject - is genuine backend labour. It requires no judgment about the law. It is a sophisticated search and classification task, and AI is genuinely excellent at it.

A lawyer can now ask a tool like Manupatra AI to find all High Court judgments in the last five years that discuss limitation periods in landlord-tenant disputes under the Specific Relief Act, and receive a comprehensive list in seconds. Before AI, this would have taken days. This time saving is real, measurable, and professionally valuable.

⁶ See Deepak Kapoor, *Why Native Legal AI is Required for India*, ARTIFICIAL LAWYER (Apr. 2, 2026), <https://www.artificiallawyer.com/2026/04/02/why-native-legal-ai-is-required-for-india/> (discussing the Manupatra Legal Tech Suite (manupatra.ai), SCC Online, CaseMine AMICUS AI module, and BharatLaw.AI NyaI stack).

Case summarization is another backend labor task AI handles reasonably well. A long judgment of a hundred pages can be summarized by AI into the key findings, the court's reasoning, and the ultimate order - giving a lawyer a quick orientation before deciding whether to read the full text. This does not require the AI to understand the law. It requires the AI to identify and extract the parts of the text that perform certain functions - findings, reasoning, orders - which is a structural classification task.

C. What AI Cannot Do: The Relevance Layer

Here is where the critical limitation begins. Retrieving cases is backend labour. Determining which retrieved cases are actually relevant to the specific dispute at hand is applied legal intelligence - and AI cannot reliably do this.

Relevance in Indian law is not determined by keyword similarity. A case about a bank guarantee may be directly relevant to a dispute about a construction contract because of how the Supreme Court interpreted the 'unconditional' nature of obligations in commercial instruments - a principle that cuts across the specific subject matter. A case from 1963 about the interpretation of a colonial-era statute may be the controlling authority for a 2025 dispute about digital contracts. No keyword search will surface these connections. Only a lawyer who understands the underlying legal principles, the hierarchy of courts, the evolution of the doctrine, and the specific facts of the present dispute can make these connections.

This is not a deficiency that better AI will eventually overcome. It reflects something fundamental: identifying legal relevance requires understanding the law as a system of principles applied to human situations - not as a database of text. AI processes text. Law is meaning, not text.

D. The Fake Citation Crisis: When AI Becomes Actively Harmful

The most serious practical problem to emerge from AI in Indian legal research is not that AI fails to find relevant cases. It is that AI invents cases that do not exist - and presents them with complete confidence.

This phenomenon, known as 'hallucination,' has created a crisis in Indian courts. General-purpose AI tools like ChatGPT, when asked to provide supporting case law for a legal position, will sometimes generate entirely fictitious judgments - with realistic-sounding case names,

realistic citation numbers, and even fabricated quotes from real judges - that do not exist in any judicial database. Lawyers who copy these outputs without verification have submitted these invented cases to courts as genuine authority.

The consequences have been severe. In February 2026, the Supreme Court bench led by Chief Justice Surya Kant raised alarm when a petition cited a case called 'Mercy versus Mankind'⁷ - a case that does not exist in any judicial record, generated entirely by an AI tool. In a 2025 incident, a trial court in Andhra Pradesh relied on four AI-generated fictitious precedents to dismiss objections in a case. The affected parties had to approach the Supreme Court, which described the situation as a grave failure of professional responsibility.

In September 2025, a petition before the Delhi High Court cited specific paragraphs from the constitutional landmark *Raj Narain v. Indira Nehru Gandhi* - paragraphs that the AI had invented. The actual judgment has only 27 paragraphs.⁸ The AI cited paragraphs 73 and 74.

In late 2024, a bench of the Income Tax Appellate Tribunal in Bengaluru itself - not just a lawyer appearing before it - cited three non-existent Supreme Court judgments and one fabricated Madras High Court ruling in its own order. The order had to be recalled within a week.⁹

These are not isolated mistakes. The Bombay High Court has imposed costs on lawyers for relying on fake AI citations.¹⁰ The Supreme Court's November 2025 White Paper on Artificial Intelligence and Judiciary formally identified hallucination of case law as a primary risk requiring mandatory independent verification.¹¹

The deeper lesson here is not simply that lawyers must verify AI output - though they absolutely

⁷ See 'Mercy vs Mankind' Never Existed: Supreme Court Slams AI-Generated Fake Judgments, LAWBEAT (Feb. 17, 2026), <https://lawbeat.in/top-stories/mercy-vs-mankind-never-existed-supreme-court-slams-ai-generated-fake-judgments-1566075> (reporting statements by a Supreme Court of India bench led by Chief Justice Surya Kant reprimanding counsel for submitting entirely hallucinated AI citations and warning against systemic unverified generative AI usage).

⁸ *Ipleaders.in*, *supra* note 13 (reporting that in Sept. 2025 a Delhi High Court bench observed that a petition cited paragraphs 73 and 74 of *Raj Narain v. Indira Nehru Gandhi*, AIR 1975 SC 865—a judgment containing only 27 paragraphs—generated entirely by an AI tool).

⁹ *In re Buckeye Trust*, ITA No. 1051/Bang/2024, Income Tax Appellate Tribunal, Bengaluru Bench (late 2024) (order recalling a prior mandate within one week after verification revealed all four supporting citations—three non-existent Supreme Court judgments and one fabricated Madras High Court ruling—were AI-generated).

¹⁰ *Artificial Intelligence in Courts: Promise, Peril, and Path Ahead*, LIVELAW (Mar. 31, 2026), <https://www.livelaw.in/articles/artificial-intelligence-courts-528424>

¹¹ *SC White Paper*, *supra* note 4, at § 3.2 (formally identifying fabrication of case citations as a primary risk and mandating independent verification of all AI-generated legal information under threat of disciplinary action).

must. The deeper lesson is that AI has no concept of truth. It has no access to the real world of court judgments. It generates text that looks like a citation because it has learned the pattern of what citations look like. It does not know - it cannot know - whether the case it is generating actually happened. Legal research requires dealing with reality. AI, fundamentally, deals with patterns in text.

An AI does not know what it does not know. It fills gaps in its knowledge with plausible-sounding text. In medicine, this is dangerous. In law, where a fictional precedent can send a person to prison or deprive them of their property, it is a catastrophe.

E. The Institutional Response

Indian courts and bar associations have begun responding to the fake citation crisis with a combination of warnings, penalties, and formal regulation. The Supreme Court's draft Regulations for Use of Artificial Intelligence in Courts, 2026, released in June 2026¹², requires mandatory disclosure when AI tools are used in preparing any court filing, and creates a 'filing provenance record' requiring lawyers to document which tools were used, what sources were consulted, and what human verification was performed.¹³

The Kerala High Court issued India's first binding AI policy for the district judiciary, explicitly prohibiting AI tools from being used to 'arrive at any findings, reliefs, order or judgment,'¹⁴ and requiring that AI transcription and translation outputs be treated as drafts requiring human verification.

The Bombay Bar Association issued advisory guidelines in July 2025 emphasizing verification, professional responsibility, and client confidentiality obligations when using AI tools.¹⁵

These responses confirm the paper's central argument. Retrieval - backend labor - is permitted,

¹² *Draft Regulations for Use of Artificial Intelligence in Courts, 2026*, Sup. Ct. of India, AI Comm. (issued June 3, 2026), <https://main.sci.gov.in> [hereinafter *Draft AI Regulations*] (open for public comment until June 20, 2026; prepared under supervision of the Supreme Court's AI Committee).

¹³ *Draft AI Regulations*, *supra* note 19, reg. 5 (mandatory disclosure obligation), reg. 7 (establishing a 'filing provenance record' requiring lawyers to document AI tools used, original sources consulted, and human verification steps performed).

¹⁴ *Policy Regarding Use of Artificial Intelligence Tools in District Judiciary*, Kerala H.C. (issued July 19, 2025) [hereinafter *Kerala AI Policy*] (distinguishing officially 'Approved AI Tools' from general generative systems and explicitly cautioning that tools such as ChatGPT pose severe confidentiality risks).

¹⁵ Bombay Bar Association, *Advisory Guidelines on Use of Artificial Intelligence for Members* (July 2025), <https://bombaybar.com> (emphasizing strict verification, transparency with clients, and confidentiality of client material while remaining advisory rather than enforceable).

even encouraged. Judgment - applied intelligence - remains entirely the domain of the human lawyer.

V. Contract Drafting: The Illusion of Competence

A. What AI Does in Contract Work

Contract drafting is the area where AI tools have been most aggressively marketed to Indian legal professionals, and where the illusion of AI competence is most seductive. Tools like SpotDraft, Legistify, VIDUR AI, and globally, Kira Systems and Harvey AI, can now produce complete contract drafts in minutes.¹⁶¹⁷ The drafts look professional. They contain recognizable clauses. They follow standard structures. They can be generated for a Non-Disclosure Agreement, a joint venture agreement, an employment contract, a lease deed, or a share purchase agreement.

Proponents cite efficiency gains of up to 90 percent in contract review time for large-volume, standardized work. For a corporate legal team reviewing hundreds of similar vendor contracts, AI can flag unusual clauses, identify missing standard terms, and highlight deviations from the firm's playbook - all in the time it previously took a paralegal to review one contract.

This is genuine value, and it is the backend labour value: pattern recognition applied to standardised, repetitive work.

B. The Three Layers of Indian Contract Complexity

But Indian contract law operates in three layers of complexity that AI tools, particularly those trained primarily on international contract databases, systematically fail to handle well.

The first layer is statutory complexity. Indian contracts are governed by the Indian Contract Act, 1872 - a colonial-era statute still in force, interpreted through 150 years of Indian jurisprudence¹⁸ that has developed distinctive interpretations in areas like consideration,

¹⁶ *AI Tools for Lawyers in India: 2026 Definitive Guide*, IPLEADERS (May 5, 2026), <https://blog.ipleaders.in/ai-tools-for-lawyers-in-india-the-2026-definitive-guide/>

¹⁷ See *SpotDraft Contract Lifecycle Management*, SpotDraft, <https://www.spotdraft.com> (last visited June 29, 2026); *Legistify Litigation & Contract Tracking Stack*, Legistify, <https://www.legistify.com> (last visited June 29, 2026); *Kira Systems Clause Extraction Engine*, Kira Systems, <https://kirasystems.com> (last visited June 29, 2026) (with vendor performance logs reporting up to a 90% reduction in contract review timelines for high-volume standardized instruments).

¹⁸ Indian Contract Act, 1872, §§ 10, 56, 73; Specific Relief Act, 1963, § 10 (as amended by the Specific Relief

privity, impossibility, and damages. The Companies Act, 2013, the Foreign Exchange Management Act, the Arbitration and Conciliation Act (as amended in 2019 and 2021), the Specific Relief Act (as amended in 2018), sector-specific regulations from SEBI, RBI, TRAI, and dozens of other regulators - all of these shape what is and is not enforceable in an Indian contract.

An AI tool trained primarily on American or English contracts does not understand this framework. It may produce a clause about specific performance that is perfectly standard under English law but will be interpreted differently by Indian courts under Section 10 of the Specific Relief Act. It may omit a stamp duty consideration that is legally required under the Indian Stamp Act for a contract to be admissible as evidence. It may include an arbitration clause that inadvertently conflicts with the mandatory provisions of the Arbitration and Conciliation Act, 1996.

The second layer is jurisdictional complexity. India is not one legal market. A contract involving a party in Maharashtra needs to account for Maharashtra's stamp duty schedules. A tenancy agreement in Karnataka operates under the Karnataka Rent Control Act, which has specific provisions that cannot be contracted away. A service agreement with a party in Telangana may be governed by Telangana-specific shops and establishments legislation. An AI tool with no knowledge of this jurisdictional texture will produce a contract that looks complete but may be legally defective in the specific jurisdiction where it matters.

The third layer is commercial and relational complexity - the most important and the one AI is most completely unable to address. What does this specific client actually need from this specific contract? What is the power dynamic between the parties? What risks are most important to protect against in this specific industry at this specific moment? What does this client's litigation history suggest about which clauses need to be particularly robust? What is the other party's reputation for honoring or contesting obligations?

These questions require the lawyer to understand the client, the relationship, the industry, and the legal environment simultaneously - and to make judgment calls that no amount of training data can replicate. The AI has no client. The AI has no understanding of the specific situation. It has only patterns learned from previous contracts. It will produce a contract that reflects what

(Amendment) Act, 2018); Arbitration and Conciliation Act, 1996 (as amended by the Arbitration and Conciliation (Amendment) Acts of 2019 and 2021); Indian Stamp Act, 1899.

similar contracts have looked like, not what this contract needs to be.

C. The Difference between a First Draft and a Right Contract

The distinction that Indian legal professionals must internalize is the difference between a first draft and the right contract. AI can produce a first draft - a structural starting point that saves time on formatting, standard clauses, and basic organization. This is backend labor value.

The right contract for a specific client requires a lawyer to apply their mind. It requires asking: what is the client trying to achieve? What could go wrong? How would an Indian court interpret this clause if the relationship breaks down? What does the client need protection against that they have not thought to ask about? Is this clause enforceable in the jurisdiction where disputes are likely to arise? Does this agreement interact correctly with the tax implications of the transaction?

No AI tool - including the most advanced global tools currently available - can perform this analysis. It requires what the Indian legal tradition calls the 'application of mind' - a phrase used in judicial decisions precisely to distinguish mechanical processing from genuine intellectual engagement with the specific facts.

When a contract produced by AI without adequate human review fails - when a clause is found to be unenforceable, when a critical risk was not addressed, when a jurisdictional requirement was missed - the damage falls on the client, not on the AI tool. The client suffers a business loss. They may lose a dispute that a well-drafted contract would have protected them from. They pay the price for the illusion that AI competence in producing first drafts translates into competence in producing the right contract.

VI. AI and Judicial Decision-Making: The Most Critical Question

A. The Background of the Debate

The debate about using AI in judicial decision-making is no longer hypothetical. Around the world, AI systems are already influencing - and in some cases making - decisions that affect people's lives. In the United States, the COMPAS algorithm has been used to assist judges in bail and sentencing decisions. In China, 'cyber courts' have adjudicated e-commerce and intellectual property disputes using AI. Estonia has piloted an AI judge for small claims. The

question for India is not whether this debate will arrive - it has arrived. The question is how India should respond.

Within India's judicial system, there are advocates for using AI to help clear the enormous backlog of pending cases. With 5.4 crore matters pending, the pressure is intense. If AI could handle some portion of routine cases - straightforward recovery suits, minor traffic violations, simple landlord-tenant disputes - courts could focus human judicial capacity on complex matters. This argument has surface appeal. It deserves serious engagement and serious rebuttal.

B. What AI Can Legitimately Do in Judicial Processes

Let us begin where AI is legitimate. There are genuine backend labor tasks within judicial processes that AI can and should assist with.

Case scheduling and administrative management - determining which cases are due for hearing, flagging cases that have been pending beyond a certain threshold, sending notices, generating cause lists - is pure administrative labor. It requires no legal judgment. AI can do this better, faster, and more consistently than any manual system. SUPACE, the Supreme Court's case management tool, is appropriately focused on this kind of work.¹⁹ It helps judges find the relevant documents and precedents in a case so they can focus on reasoning, not retrieval.

Transcription of court proceedings - converting what lawyers and judges say during hearings into accurate written records - is backend labor. Tools like Adalat AI are doing this in real time for Constitution Bench hearings at the Supreme Court.²⁰ This is valuable because India's courts have historically had poor or no real-time transcription, meaning there was often no reliable record of what was argued. AI transcription improves this.

Summarization of pleadings - giving a judge a quick, accurate summary of what each party has argued before oral arguments begin - is also backend labor, provided the summary is accurate

¹⁹ *Supreme Court Portal for Assistance in Court Efficiency (SUPACE)*, Sup. Ct. of India, <https://main.sci.gov.in> (launched Apr. 2021) (built by ManCorp Innovation Labs; noting that as of Jan. 2026, the system remains deployed exclusively for select criminal matters before judges of the Bombay High Court and Delhi High Court).

²⁰ *Adalat AI Real-Time Transcription Stack*, Adalat AI, <https://www.adalat.ai> (last visited June 29, 2026) (documenting deployment for live Constitution Bench transcriptions at the Supreme Court of India). *See also Technology Enabled Resolution (TERES) Transcription and Recording System*, TERES, <https://teres.ai> (operating in parallel for official record-keeping workflows under Phase III of the e-Courts Project).

and the judge treats it as a starting point, not a conclusion.

C. Why AI Cannot Give Judgment

None of the above - scheduling, transcription, summarization - is judgment. Judgment is something categorically different, and the case against AI judgment-giving in India is not merely practical. It is constitutional, jurisprudential, and philosophical.

Judgment in the legal sense means this: a human decision-maker, who has the authority of the state, reads the specific facts of a dispute, applies the law to those facts, weighs competing considerations, and makes a reasoned choice about what justice requires in this situation - a choice for which that decision-maker is personally and institutionally accountable. Every element of this definition is violated by AI adjudication.

Consider the meaning of 'applying the law to the facts.' The law is not a lookup table. 'Section 302 of the IPC provides for punishment for murder' - but what is murder? Was this specific killing done with the intention the statute requires? Was this specific person in a position of mental duress that reduces culpability? Were there specific mitigating circumstances that require the application of a lesser charge? These questions require a judge to understand what it means to be human in specific circumstances - to understand fear, coercion, desperation, and intent as human experiences, not as text classifications.

AI has never felt fear. It has never been under coercion. It has no understanding of human emotion, human motivation, or human circumstance from the inside. It can classify descriptions of these things based on patterns in text. But classifying descriptions of fear is not understanding fear. When a judge decides whether a confession was voluntary or coerced, they draw on their own humanity - their own experience of what it means to make a free choice versus a forced one. AI has no such experience to draw on.

D. The Constitutional Foundation of Human Adjudication in India

India's Constitution provides the deepest basis for insisting on human adjudication. Article 21 guarantees that no person shall be deprived of their life or personal liberty except according to procedure established by law.²¹ The Supreme Court has interpreted this to require not just the

²¹ INDIA CONST. art. 21 (guaranteeing that no person shall be deprived of life or personal liberty except according to procedure established by law, interpreted to require a fair, just, and reasonable procedure). *See*

existence of a legal procedure, but a procedure that is fair, just, and reasonable. A procedure in which an algorithm - which cannot be cross-examined, which cannot explain its reasoning in terms a human can engage with, and which cannot be held accountable - makes the decision about a person's liberty is arguably inconsistent with the right to a fair hearing.

Article 14 guarantees equality before the law. AI systems trained on historical data replicate historical patterns - including historical biases.²² An AI trained on decades of Indian criminal sentencing data will learn, and reproduce, whatever systemic biases existed in that data - biases related to caste, religion, economic class, and geography that the Indian constitutional framework is specifically designed to counteract. Using AI in adjudication would risk encoding these biases in an opaque, unaccountable system.

The doctrine of natural justice - audi alteram partem (hear the other side) and nemo iudex in causa sua (no one shall be a judge in their own cause) - requires that parties have a genuine opportunity to be heard by a decision-maker who can understand their arguments and who has no pre-programmed bias. An AI judge has, in effect, already formed its 'view' based on its training data before the parties even appear before it. Its 'hearing' of the parties is an input to a process whose direction is already largely determined. This is not the hearing that natural justice requires.

E. The Jurisprudential Argument: What Judgment Is For

Beyond the constitutional arguments, there is a deeper jurisprudential point. Judgment - genuine judicial judgment - is not primarily about reaching the correct answer. It is about legitimately resolving a dispute in a way that the parties and society can accept as just, even if not everyone agrees with the outcome.

This legitimacy rests on several foundations. It rests on the perception that the decision-maker has genuinely engaged with the specific situation. It rests on the availability of reasons - a judgment must explain why this outcome follows from the law and the facts, in terms that a trained lawyer and ultimately a citizen can engage with and, if necessary, challenge. It rests on accountability - the judge whose decision is wrong can be identified, their reasoning can be

Maneka Gandhi v. Union of India, (1978) 1 S.C.C. 248.

²² INDIA CONST. art. 14 (guaranteeing equality before the law and equal protection of the laws). *For the risk of AI replicating historical bias in adjudication, see generally* BERNARD E. HARCOURT, *AGAINST PREDICTION: PROFILING, POLICING, AND PUNISHING IN AN ACTUARIAL AGE* (2007).

criticized, their decision can be appealed, and in extreme cases of misconduct they can be removed.

An AI judgment destroys all of these foundations simultaneously. It cannot genuinely engage with a specific situation - it processes the situation as a data point. Its 'reasons' are statistical patterns, not logical chains that a human can meaningfully evaluate. It cannot be held accountable - when an AI makes an error that destroys a person's life, who is responsible? The company that trained the model? The court that deployed it? The programmer who wrote the algorithm? The answer to this question has no satisfactory form in current Indian law, or indeed in any legal system.

There is also the question of evolution. Indian law develops through the accumulation of judgments over time - each case adds something, refines something, extends a principle or limits it in response to new factual situations the legislature did not anticipate. The common law tradition, which India inherited and developed, depends on this organic, case-by-case growth of legal understanding. AI, trained on past data, can at best reproduce what has already been decided. It cannot develop the law. It cannot exercise the creative function that is essential to a living legal system.

The common law is not a database. It is a conversation between generations of judges, conducted through the medium of specific cases, about the meaning of justice in an evolving society. AI can read the transcripts of that conversation. It cannot participate in it.

F. The Specific Risks to Indian Jurisprudence

India's jurisprudence has developed several distinctive doctrines - the basic structure doctrine, the doctrine of proportionality in fundamental rights cases, the expansive interpretation of Article 21 - that are products of bold, creative judicial reasoning applied to specific situations that the framers of the Constitution did not foresee. The Kesavananda Bharati judgment, which established that Parliament cannot amend the basic structure of the Constitution, required judges to reason about the limits of constituent power in a way that had no direct precedent.²³ The Vishakha judgment, which created sexual harassment guidelines in the absence of

²³ Kesavananda Bharati v. State of Kerala, (1973) 4 SCC 225

legislation, required judges to read constitutional principles into the employment context.²⁴

None of these landmark decisions could have been produced by an AI system trained on existing law. They were produced by human judges who were willing to ask: what does justice require here, even if the existing law does not clearly provide for it? This is the creative, generative function of judicial reasoning that makes a legal system capable of responding to a changing society. AI cannot perform this function. If Indian adjudication is handed to AI, this function will be lost, and Indian jurisprudence will stagnate - frozen at whatever the state of the law was when the AI was trained.

G. India's Current Position: Cautious and Correct

India's judiciary has, to its credit, taken a clear and principled position on this question. The Supreme Court's AI Committee has established that AI assists with administrative tasks but cannot draft judgments or make outcome predictions. The Kerala High Court's binding AI policy explicitly prohibits AI from being used 'to arrive at any findings, reliefs, order or judgment.'²⁵ The draft Regulations for Use of Artificial Intelligence in Courts, 2026, confirm that AI is permitted as an assistive tool - for research, summarization, translation, transcription - but that the decision-making function is exclusively human.

This position is correct, and this paper supports it. The caution is not technophobia. It is a principled recognition of what adjudication is - and what makes it legitimate.

VII. Conclusion

This paper has argued that the future of AI in the Indian legal system depends on correctly identifying the boundary between two kinds of tasks: backend labor, which AI can perform effectively; and applied legal intelligence, which AI cannot perform and should not be asked to perform.

In translation, AI has demonstrated genuine, substantial value. SUVAS has made over 36,000 Supreme Court judgments accessible in regional languages, meaningfully improving access to information about the law. The limitations - translation inaccuracies, the informational rather

²⁴ Vishakha v. State of Rajasthan, (1997) 6 SCC 241

²⁵ Kerala AI Policy, *supra* note 21, § 4(b) (prohibiting AI software from being used "to arrive at any findings, reliefs, order or judgment" and requiring that AI outputs—including transcription and translation—be treated as working drafts subject to human review).

than authoritative status of translations - are real but manageable, and the overall contribution is positive. This is backend labor that AI does well.

In legal research, AI is effective at the retrieval layer - finding cases, summarizing judgments, mapping citation networks. It is not effective, and is actively dangerous, at the relevance layer - determining which cases matter for this specific dispute. The fake citation crisis, which has now produced multiple incidents before the Supreme Court and High Courts, demonstrates that using AI for applied legal intelligence in research does not merely fail to help. It actively introduces false information into judicial proceedings, wastes judicial time, and harms clients. This is a structural problem with AI that better tools will mitigate but not eliminate.

In contract drafting, AI can produce first drafts that save time on formatting and standard clauses. It cannot produce the right contract for a specific client, in a specific commercial relationship, governed by the specific texture of Indian law across its multiple jurisdictions and regulatory regimes. The right contract requires applying a legal mind to a specific human situation. AI has no legal mind. It has a pattern-matching engine.

In judicial decision-making, the case against AI is not just practical but constitutional and jurisprudential. Adjudication in India rests on the legitimacy of a human decision-maker who genuinely engages with the specific situation, reasons through competing considerations, explains their reasoning in terms that can be evaluated and challenged, and is accountable for their decisions. AI satisfies none of these requirements. Its use in adjudication would undermine Article 21's guarantee of fair procedure, encode historical biases in violation of Article 14, violate the principles of natural justice, and foreclose the creative, generative function of judicial reasoning on which the development of Indian jurisprudence depends.

The right framework for AI in the Indian legal system is clear: use it to do what it does well - handle volume, remove administrative friction, democratize access to legal information. Do not use it for what it cannot do - apply judgment, understand human situations, reason about what justice requires, or make decisions that affect people's lives and rights.

AI is a remarkable tool. A hammer is also a remarkable tool. Neither should be asked to do a surgeon's work.