



# Artificial Intelligence and the Future of Legal Practice in India: Evidence from an Empirical Study of Legal Professionals

Prashant Panwar<sup>1</sup>, Prof. (Dr.) Shobhna Jeet<sup>2</sup>

<sup>1</sup>School of Legal Studies, K.R. Mangalam University, Gurugram

<sup>2</sup>School of Legal Studies, K.R. Mangalam University, Gurugram

## Abstract

India's legal profession has adopted artificial intelligence tools at a striking pace, yet the statutory framework governing that adoption remains anchored in legislation that predates modern AI by decades. This paper reports primary empirical findings from a survey of 229 Indian legal professionals carried out in April 2026, constituting the empirical chapter of a doctoral dissertation on progressive and protective approaches to AI regulation in the European Union and India. The survey reveals that 94.4 per cent of respondents already use AI tools professionally, while 76.4 per cent regard India's existing regulatory framework as inadequate or significantly inadequate. Professional demand for governance is specific and strong: 93 per cent support mandatory disclosure of AI use to courts and clients, 90.4 per cent favour Bar Council of India ethics guidelines, 84.3 per cent back compulsory AI ethics training, and 74.7 per cent endorse mandatory auditing of AI systems deployed in court proceedings. Against these findings, a mean AI-preparedness score of only 2.67 on a five-point scale points to an urgent professional-readiness deficit. The data provide direct empirical support for the thesis's first hypothesis — that India's reliance on pre-AI statutes creates structural governance gaps — and ground the original Responsible Acceleration Model proposed as a context-sensitive governance framework for the Indian legal profession.

**Keywords:** Artificial Intelligence governance, Indian legal profession, empirical survey, Responsible Acceleration Model, EU AI Act, Bar Council of India, professional ethics, digital regulation

## I. Introduction

Few developments in recent legal history have moved as quickly, or with as little accompanying governance, as the penetration of artificial intelligence into everyday legal work. Within a span of roughly three years, tools capable of drafting pleadings, retrieving case law, summarising judgments, and generating contract clauses have shifted from experimental curiosities to routine professional instruments. In India, this shift has taken place beneath a regulatory canopy that was never designed to hold it. The Information Technology Act 2000, the Digital Personal Data Protection Act 2023 (DPDPA), and the Copyright Act 1957 — the three statutes most commonly invoked in AI-related disputes — address, respectively, electronic transactions and cyber offences, personal data processing, and literary property. None of them speaks to the distinctive risks that AI systems introduce: the generation of plausible but fabricated legal authority, the processing of confidential client communications through commercial cloud services, or the amplification of historical biases embedded in training data [1].

Comparative legal scholarship has devoted considerable energy to examining the European Union's response to these challenges. The EU AI Act 2024, which entered into force on 1 August 2024, is the first horizontal, risk-stratified statutory instrument specifically designed to govern AI across all sectors of an economy [2]. For India, the Act matters not merely as a regulatory achievement to be admired from a distance but as a potential architectural reference, given that India's own governance evolution has repeatedly drawn on European regulatory models in domains ranging from data protection to competition law. Yet the question of whether the EU framework offers India genuinely transferable governance content — and which parts of it require contextual retuning — cannot be answered through doctrinal analysis alone. What India's AI governance debate lacks, and what existing scholarship has not supplied, is primary empirical evidence from the professional community at the heart of the problem.

This paper addresses that gap. The survey on which it reports was administered to 229 Indian legal professionals in April 2026 and constitutes the primary empirical chapter of a doctoral dissertation examining progressive and protective approaches to AI regulation in the EU and India. The study pursued two overarching research hypotheses: that India's existing statutory architecture, designed for pre-AI contexts, creates structural governance challenges for AI in legal practice (H1); and that the EU's rights-centred, risk-tiered governance model offers India a credible normative and institutional reference point, subject to adaptation for India's developmental and constitutional context (H2). The survey findings engage both hypotheses directly, providing quantified professional assessments of governance adequacy, documented harm experiences, and explicit regulatory preferences from across the legal profession.

The paper proceeds as follows. Section II describes the survey design, sampling strategy, and analytical approach. Section III presents the demographic profile of respondents. Sections IV through VII report substantive findings on AI use patterns, governance perceptions, professional risk assessments, and regulatory preferences. Section VIII draws together the governance implications of the aggregate data. Section IX outlines the Responsible Acceleration Model that these findings ground and explains how specific survey results map onto specific governance provisions. Section X reflects on what the data mean for India's near-term regulatory trajectory.

## II. Research Methodology

### A. Instrument Design

The survey instrument was structured in four parts. Part A gathered background information about each respondent: professional category, years of practice or teaching experience, primary area of specialisation, the High Court jurisdiction in which they primarily worked, and whether they had previously received any formal training in AI governance or ethics. Part B explored current AI use, asking respondents to identify which tools they used, for what professional purposes, and how often. Part C assessed governance awareness and perceptions — how familiar respondents were with existing regulatory frameworks, how adequate they considered those frameworks, and what governance failures or AI-related professional difficulties they had personally encountered. Part D recorded regulatory preferences, asking respondents to express views on a range of governance mechanisms from professional conduct rules to statutory authority and legislative models.

The instrument combined Likert-scale items, multiple-choice questions, and open-ended fields for qualitative elaboration. A pilot run with fifteen practitioners led to revision of two question stems where respondents had interpreted the original wording differently from what was intended. The revised instrument was then distributed in April 2026 through a combination of online dissemination via legal professional networks and direct outreach to bar associations affiliated with the five targeted High Courts.

### B. Sampling

Four professional categories were targeted: advocates and solicitors in active practice before High Courts; judicial officers at district court level and above; government law officers in central and state legal departments; and legal academics at recognised law schools. Geographic coverage spanned five High Court jurisdictions — Delhi, Bombay, Madras, Calcutta, and Kerala — chosen to capture variation in institutional culture, technology infrastructure, and the character of legal practice.

Sampling was stratified and purposive, not random. The decision to focus on upper-tier practitioners was deliberate: district-court-level AI adoption was, at the time of the survey, considerably lower, and the governance questions that animated the research — disclosure obligations, professional accountability for AI errors, data protection in client matters — were most immediately live at the High Court level. The study does not claim to represent the full breadth of India's legal profession, and the implications of this boundary are revisited in Section X. After removing incomplete or inconsistent responses, 229 valid responses were retained for analysis.

### C. Analysis

Quantitative responses were analysed using descriptive statistics: frequency counts, percentage distributions, and mean scores for Likert items. No inferential significance testing was performed, consistent with the research's governance-analytical rather than hypothesis-testing statistical orientation. Open-ended responses were thematically coded; selected quotations appear in the discussion where they add texture to the numerical findings. All data were handled in accordance with standard research ethics protocols, and respondents were assured of anonymity before participation.

## III. Respondent Profile

Advocates and solicitors in active practice formed the single largest group, accounting for 42.4 per cent of the sample ( $n = 97$ ). Legal academics followed at 28.8 per cent ( $n = 66$ ), then government law officers at 17.9 per cent ( $n = 41$ ), and judicial officers at 10.9 per cent ( $n = 25$ ). The judicial officer sub-group, though the smallest, carries disproportionate analytical weight: these are professionals making decisions in proceedings where AI is already being used, and their governance perceptions carry institutional authority that other respondents' views do not.

**Table 1. Respondent Professional Category Distribution**

Professional Category	n	Percentage
Advocates / Solicitors in Active Practice	97	42.4%
Legal Academics	66	28.8%
Government Law Officers	41	17.9%
Judicial Officers	25	10.9%
<b>Total</b>	<b>229</b>	<b>100%</b>

Experience levels skewed toward the senior end of the profession: 41.5 per cent of respondents had practised or taught for more than fifteen years, and a further 28.4 per cent reported between eight and fifteen years of experience. This is partly an artefact of the sampling frame — High Court practitioners tend to be more experienced than their district-court counterparts — but it also means that the survey captures the assessments of people with deep enough professional knowledge to evaluate AI's practical effects rather than offering speculative views.

One demographic finding deserves early attention. Only 23.1 per cent of respondents had received any formal AI governance or ethics training, and a mere 8.3 per cent had received training that focused specifically on AI in legal practice. Given that these numbers come from a sample already biased toward senior urban practitioners, the likely figures for the wider profession are sobering. This training gap is not just a footnote to the demographics: it runs through the rest of the survey data as a structuring condition.

#### IV. AI Adoption and Usage Patterns

##### A. How Widely is AI Already in Use?

The short answer is: almost universally, at least within this sample. Of the 229 respondents, 216 (94.4 per cent) reported using at least one AI tool in their professional or academic legal work. That figure is arresting by any measure, but especially so in a context where professional conduct rules have not been updated to address AI, where no mandatory disclosure obligations exist, and where the statutory framework governing AI was never designed for it. The governance challenge, in other words, is not anticipatory: it is already happening.

The thirteen non-users (5.6 per cent) were asked why they had not adopted AI tools. Accuracy and reliability concerns were cited by 69.2 per cent of non-users, followed by data confidentiality worries (53.8 per cent) and unfamiliarity with available tools (38.5 per cent). What is striking is how closely these non-user concerns mirror the risk perceptions of users (reported in Section VI). The profession, whether using AI or not, has arrived at similar conclusions about what is dangerous about it — a convergence that points toward a shared professional understanding awaiting formal codification.

##### B. Frequency of Use

Daily use was reported by 38.2 per cent of AI users ( $n = 82$ ), and a further 32.4 per cent used AI tools several times a week. Adding the 18.1 per cent reporting weekly use, roughly 88 per cent of AI users engage with these tools at least once a week. Only 11.3 per cent described their use as occasional or infrequent.

**Table 2. Frequency of AI Tool Use Among AI-Using Respondents ( $n = 216$ )**

Frequency	n	Percentage
Daily	82	38.2%
Several times per week	70	32.4%
Weekly	39	18.1%
Occasional / Infrequent	25	11.3%

The governance relevance of high-frequency use is often overlooked in regulatory debates that treat AI as an episodic tool. A practitioner who interacts with AI daily is not in a position to apply careful ad hoc scrutiny each time; the default behaviour of daily use is routinisation. Governance mechanisms designed for episodic, deliberate AI use — such as post-hoc review or voluntary disclosure — are structurally mismatched to a practice environment where AI outputs are generated and relied upon as naturally as a Google search. What the frequency data suggest is that governance rules need to be embedded in professional practice defaults, not bolted on as occasional checks.

##### C. What AI is Being Used For

Respondents could select multiple purposes. Legal research and case law retrieval was the most frequently cited use (79.2 per cent), well ahead of document drafting and review (63.1 per cent) and legal writing and argument structuring (58.1 per cent). Statutory and regulatory analysis was cited by 47.6 per cent. Predictive analytics and case outcome modelling came last at 24.9 per cent, a figure that likely reflects a combination of tool availability constraints and professional discomfort with algorithmic prediction in adjudicatory contexts.

Taken together, the usage pattern is one of tool-augmentation rather than decision-delegation. Practitioners are using AI to do existing research and drafting tasks faster, not (yet) to let AI make substantive calls. This matters for governance design. The EU AI Act's treatment of AI systems in "the administration of justice" as high-risk presupposes a system involved in decision-making; the Indian legal profession's AI use sits, for the most part, in an assistive register below that threshold. The governance risks in this assistive register are real — hallucination, data leakage, attribution error — but they call for different regulatory tools than those designed for decision-supporting systems.

#### V. Perceptions of India's Governance Framework

##### A. Is the Existing Framework Adequate?

The clearest single finding from the survey is also the most consequential for policy: 76.4 per cent of respondents described India's existing regulatory framework for AI in legal practice as "inadequate" or "significantly inadequate." Only 8.7 per cent considered it adequate or better. The remaining 14.9 per cent chose "partially adequate." These numbers are cross-demographically consistent. Advocacy practitioners expressed the sharpest dissatisfaction (81.4 per cent finding the framework inadequate), but the figure among legal academics was 74.2 per cent, and even among judicial officers — the group most institutionally invested in the existing framework's continuity — it reached 72.0 per cent.

The inadequacy finding is not merely a complaint about a regulatory gap in the abstract. When respondents were asked which specific aspects of the framework were most deficient, the responses pointed to concrete, nameable failures: 68.4 per cent identified the absence of any duty to verify AI outputs before professional reliance; 66.8 per cent identified the lack of mandatory disclosure requirements for AI use in court proceedings; 63.3 per cent pointed to the absence of data protection provisions specifically applicable to legal professional AI use; and 58.1 per cent cited the non-existence of any professional misconduct framework for AI-related failures. Each of these is a specific, drafts-ready governance gap. Each corresponds, with varying degrees of precision, to provisions that the EU AI Act includes: Article 14's human oversight requirements, Articles 13 and 50's transparency obligations, Article 10's data governance requirements, and the enforcement architecture of Chapter VIII respectively.

## B. Efficiency, Risk, and Preparedness: Likert Data

Respondents were invited to rate a set of statements about AI on a scale from 1 (strongly disagree) to 5 (strongly agree). On the efficiency side, the findings were moderately positive: time savings in legal research scored a mean of 3.82, improved access to legal information 3.74, and assistance in comparative research 3.69. These scores reflect genuine professional appreciation for what AI can do. They are not, however, the most important numbers in the Likert analysis.

**Table 3. Likert Scale Responses: AI Efficiency and Risk (1 = Strongly Disagree; 5 = Strongly Agree)**

Statement	Mean Score
AI tools save significant time in legal research	3.82
AI improves access to legal information	3.74
AI assists in comparative and cross-jurisdictional research	3.69
AI poses significant data privacy risks in legal practice	3.94
AI outputs require thorough verification before reliance	4.41
I feel professionally prepared to use AI responsibly	2.67
AI will significantly improve access to justice for marginalised communities	2.91

Two scores demand separate discussion. The statement that AI outputs require thorough verification before reliance produced the highest mean in the entire survey (4.41), with 89.6 per cent of respondents agreeing or strongly agreeing. This near-unanimity is remarkable because it exists entirely without a formal legal duty to support it. The professional norm is there; the enforceable obligation is not. Translating this professional value into a codified duty, with clear professional misconduct consequences for breach, is arguably the single most straightforward governance step available to the Bar Council of India or the Supreme Court at this moment.

The lowest score in the table — 2.67 for the statement "I feel professionally prepared to use AI responsibly" — is, if anything, more significant still. A profession using AI at a 94.4 per cent rate while placing its own preparedness below the midpoint of a five-point scale is not a profession with a confidence problem: it is a profession with a training problem. These two figures together — 94.4 per cent adoption, 2.67 preparedness — describe the essential governance emergency. The profession has run ahead of itself, and governance must catch up.

The equity-sceptical finding also warrants reflection. Only 34.2 per cent of respondents agreed that AI would significantly improve access to justice for marginalised communities (mean 2.91). This is not technophobia. These are lawyers who use AI daily and value its efficiency benefits. Their scepticism about AI's justice-equity effects reflects, rather, professional familiarity with the conditions that make AI benefits unlikely to reach those who need legal help most: the digital divide, the concentration of AI tools in urban English-medium practice environments, and the accumulated evidence that automated systems trained on historical legal data tend to replicate rather than correct existing patterns of disadvantage.

## VI. Professional Risk Perceptions

### A. What Concerns Practitioners Most?

Respondents selected, from a provided list, all risks they considered significant in the context of AI use in legal practice. The concern appearing most frequently — accuracy and hallucination risk — was selected by 84.7 per cent. Client data confidentiality followed at 79.3 per cent, AI-generated evidence fabrication at 72.5 per cent, algorithmic bias and discriminatory outputs at 69.4 per cent, and liability for professional errors attributable to AI at 67.2 per cent.

**Table 4. Risks Identified as Significant by Respondents (Multiple-Response, n = 229)**

Risk Category	% Selecting
Accuracy and hallucination risk (fabricated legal authority)	84.7%
Client data confidentiality and privacy	79.3%
AI-generated evidence fabrication	72.5%
Algorithmic bias and discriminatory outputs	69.4%
Professional liability for AI errors	67.2%
Erosion of attorney–client privilege	61.1%
Displacement of junior legal professionals	54.6%

The hallucination concern reflects concrete professional experience, not abstract anxiety. Open-ended survey responses documented twenty-three instances in which respondents had personally encountered, or discovered in opposing submissions, AI-generated citations that did not exist. These ranged from cases with invented parties and fictional bench compositions to statutory provisions that bore no relation to any enacted legislation. The encounters span multiple High Court jurisdictions and appear across all four professional categories. Taken

together, they constitute a primary evidentiary record of the governance failure that the ITAT Bengaluru’s recall of the Buckeye Trust order [6] and the Delhi High Court’s Greenopolis observations [7] have already begun to address through reactive judicial governance. The question the survey data raise is whether judicial reaction is the appropriate primary governance mechanism when the underlying professional behaviour is occurring at this scale. The data confidentiality concern (79.3 per cent) is particularly sharp in the Indian legal context. Most AI tools used by Indian practitioners are cloud-based services hosted outside India, typically by US or EU providers. When a practitioner submits a client’s facts to a commercial large language model to generate a draft pleading, that data is processed on foreign servers under foreign data protection law. The DPDPA 2023 applies to the Data Fiduciary (the practitioner), but it does not address the data protection implications of AI training on professionally generated content, nor does it impose the purpose-limitation constraints that would prevent a commercial AI provider from using submitted client data to improve its models. The 63.3 per cent of respondents who specifically identified the absence of legal professional AI data protection provisions as a governance gap are pointing, with reasonable precision, at a lacuna in the DPDPA.

### B. Who Bears Responsibility When AI Gets It Wrong?

Accountability attribution is one of the areas where professional norms are most fluid and where governance clarity would be most valuable. Asked to identify who should bear primary responsibility when AI use in legal practice causes professional or client harm, respondents placed the largest share on the practitioner who failed to verify the output (47.2 per cent). AI developers or providers were assigned primary responsibility by 28.8 per cent, shared liability between user and developer by 16.2 per cent, and the client by 7.8 per cent.

**Table 5. Attribution of Responsibility for Harmful AI Outputs (n = 229)**

Party Identified as Primarily Responsible	% Respondents
Advocate / practitioner (for failing to verify output)	47.2%
AI developer / service provider	28.8%
Shared (practitioner and developer jointly)	16.2%
Client (for accepting AI-assisted legal advice)	7.8%

The practitioner-responsibility plurality is convergent with judicial trends in other common-law systems. The US courts’ progression from *Mata v Avianca* [4] to *Lexis v Overstock* [5] consistently locates AI accountability in the duty of competence: attorneys are responsible for what they file, regardless of whether the error originated with a human or a machine. The Illinois Supreme Court’s 2025 AI Use Policy and Arizona’s binding disclosure rule both premise their requirements on the same principle. Indian courts arriving at analogous conclusions through the Greenopolis observations and the ITAT recall order are not isolated innovations; they are part of a convergent judicial movement whose regulatory logic points, in all jurisdictions, toward professional conduct codification.

The 28.8 per cent assigning primary responsibility to developers, and the 16.2 per cent favouring dual accountability, express a professional instinct that the EU AI Act formalises: obligations on both AI system providers (Article 16) and deployers (Article 26) of high-risk systems. A governance framework that places all accountability on the practitioner without any obligation on the developer of a system prone to hallucination would be, professionally speaking, unsatisfying — and the survey data suggest that a substantial portion of the profession would view it as inadequate.

## VII. Regulatory Preferences

### A. Specific Governance Mechanisms

Asked whether they supported a series of named governance mechanisms, respondents expressed strikingly high agreement levels across the board. The mechanism drawing the strongest support was mandatory disclosure of AI use to courts and clients (93.0 per cent). The near-absence of opposition to disclosure — only 7.0 per cent opposing or neutral — across all professional sub-groups establishes it as a genuinely consensual starting point for governance design. Any regulatory framework that omits mandatory disclosure as its baseline would be rejecting the express preference of 93 per cent of the practitioner community.

**Table 6. Support for Specific AI Governance Mechanisms (n = 229)**

Governance Mechanism	Support	Oppose / Neutral
Mandatory disclosure of AI use to courts and clients	93.0%	7.0%
Bar Council of India mandatory AI ethics guidelines	90.4%	9.6%
Mandatory AI ethics training for practitioners	84.3%	15.7%
Mandatory audit of AI systems used in courts	74.7%	25.3%
Independent regulatory authority (NARALP-type)	71.2%	28.8%
Liability provisions targeting AI developers	68.1%	31.9%
AI certification scheme for legal tools (LACS-type)	62.4%	37.6%

BCI ethics guidelines drew 90.4 per cent support, mandatory AI ethics training 84.3 per cent, mandatory auditing of court-deployed AI systems 74.7 per cent, and a NARALP-type independent statutory authority 71.2 per cent. Even the least popular item in the list — an AI certification scheme at 62.4 per cent — commands a clear majority. The survey does not reveal a profession reluctant to be regulated; it reveals a profession that has already decided regulation is necessary and is waiting for someone to implement it.

### **B. Preferred Legislative Model**

Respondents were asked to identify their preferred overall governance model. Comprehensive legislation through Parliament and a dedicated authority attracted the single largest plurality (41.5 per cent), followed by a co-regulatory model combining BCI authority with statutory backing (34.5 per cent). Pure self-regulation through BCI alone was preferred by 16.2 per cent, and sector-specific ministerial guidelines by 7.8 per cent.

The combined statutory preference — those choosing either comprehensive legislation or a co-regulatory model requiring statutory foundation — amounts to 76.0 per cent. This figure mirrors, rather precisely, the 76.4 per cent who found the existing framework inadequate. The two data points reinforce each other: the same professional community that experiences the governance gap is demanding statutory closure. The preference cannot be dismissed as an abstract ideal untethered from felt professional experience.

### **C. Data Protection Preferences**

Client data privacy was an active concern for 77.8 per cent of respondents. When asked to identify preferred data protection mechanisms specifically for legal AI use, the most popular response was a prohibition on using client data to train commercial AI models without explicit and informed consent (79.1 per cent). Mandatory data residency requirements for AI tools processing legal case data followed at 68.1 per cent, and enhanced DPDPA protections tailored to legal professional AI use at 64.2 per cent. The consent-for-training prohibition is particularly significant because it addresses a gap in the DPDPA 2023 that current statutory interpretation cannot easily bridge: the Act's purpose-limitation provisions do not clearly extend to the use of professionally generated content for AI model improvement, leaving a regulatory vacuum that practitioners have identified with considerable specificity.

## **VIII. What the Data Mean for Governance**

Several governance implications emerge from the aggregate findings that are worth drawing out explicitly rather than leaving to the reader to assemble from individual sections.

The profession is not a bystander waiting to receive governance from outside. With 94.4 per cent adoption and 76.0 per cent statutory governance preference, India's legal profession is simultaneously the primary AI user community in the legal sector and its most insistent governance advocate. Regulatory frameworks designed without professional input, or that treat practitioners as compliance objects rather than governance co-designers, are likely to produce rules that look competent on paper but have insufficient professional uptake. The BCI's authority under Section 49 of the Advocates Act 1961 to issue mandatory rules on professional conduct is already available; the survey data suggest that the political conditions for using it — professional consensus and demand — are unusually clear.

The readiness deficit is the most urgent challenge. The gap between 94.4 per cent AI adoption and a mean preparedness score of 2.67 is, put plainly, alarming. It means the profession is already operating at a level of AI use that substantially exceeds its own assessment of its capacity to use AI safely. Governance that begins with compliance obligations — disclosure forms, verification certifications, audit submissions — before building the professional competencies those obligations presuppose will generate cynical checkbox-ticking rather than genuine behavioural change. The training investment must be front-loaded, not retrofitted.

Hallucination is a documented professional harm, not a theoretical risk. Twenty-three specific instances of AI-generated fabricated legal authority were recorded in open-ended responses. These are real cases, real proceedings, real professional reputations affected. They span jurisdictions and professional categories. The appropriate response to documented, recurring professional harm is not continued reliance on ad hoc judicial correction — which by definition arrives after the harm has materialised in a court record — but a professional conduct rule making verification a non-delegable duty and treating failure to verify as a cognisable misconduct ground.

The equity concern cannot be finessed by general efficiency arguments. The finding that only 34.2 per cent of respondents believe AI will meaningfully improve access to justice for marginalised communities comes from practitioners with direct knowledge of how technology benefits and costs distribute across India's socio-economic landscape. AI governance frameworks that treat equity as an aspirational afterthought risk producing a two-tier legal profession: AI-augmented and well-resourced at the top, unassisted and under-resourced at the bottom, with the gap between them widening rather than narrowing.

International regulatory reference points are professionally legitimate. When 58.4 per cent of respondents are aware of the EU AI Act and 71.3 per cent of those consider it a relevant or highly relevant governance reference, the notion that EU-inspired governance principles represent an inappropriate foreign imposition — an argument sometimes advanced against comparative regulatory borrowing in India — loses purchase. The profession has looked at the EU framework and found it worth learning from. That professional endorsement matters for the political legitimacy of a governance reform process.

## **IX. The Responsible Acceleration Model: Evidence into Policy**

The Responsible Acceleration Model (RAM) is the governance framework proposed in the doctoral dissertation [8] to which this survey belongs as its empirical chapter. Its provisions derive directly from the survey findings;

the mapping between empirical problem and governance response is set out in Table 7 below. The following paragraphs offer a condensed account of the RAM's architecture sufficient to contextualise that mapping.

#### A. Constitutional Grounding and Core Principles

The RAM is built around eight principles, each independently grounded in Indian constitutional jurisprudence rather than simply borrowed from the EU. Human Primacy holds that all AI in legal practice must operate under binding human oversight, drawing on natural justice doctrine under Article 21 and the constitutional commitment to judicial independence under Article 50. Algorithmic Transparency requires that practitioners be able to explain AI-assisted outputs to courts and clients, rooted in the right to a reasoned decision that Article 21's procedural fairness guarantee has been interpreted to include. Verifiable Accuracy converts the profession's near-unanimous agreement (89.6 per cent) that AI outputs require verification into a codified professional duty under the BCI's professional conduct rules. Data Dignity addresses the 77.8 per cent data privacy concern through purpose-limitation and informed consent requirements grounded in Justice Puttaswamy [10]. Proportionate Risk-Governance calibrates regulatory burden to AI application risk rather than imposing uniform compliance costs. Equitable Access directly responds to the 34.2 per cent access-to-justice finding by embedding equity audit requirements at higher risk tiers. Professional Self-Governance recognises BCI's institutional authority while requiring statutory backing to give professional rules the legal force that 76.0 per cent of respondents demand. Responsible Acceleration is the organising principle: the framework enables beneficial adoption rather than simply constraining harm.

#### B. Institutions

The RAM proposes three new institutional bodies. The National AI Regulatory Authority for Legal Practice (NARALP) would be established as an independent statutory authority under Parliament's Union List powers, holding certification, audit, and enforcement functions. The Bar Council of India AI Ethics Committee (BCI-AEC) would be constituted under the BCI's existing authority under Section 49 of the Advocates Act 1961, developing mandatory professional conduct guidelines, a Legal AI Practice Guide, and an AI Supplementary Code to the BCI Rules. The Judicial AI Oversight Panel (JAOP) would be established under the Supreme Court's administrative authority under Article 142, developing disclosure standards for AI-assisted submissions and overseeing AI integration in court administration. The three bodies are designed to be complementary rather than overlapping: NARALP for statutory authority and certification, BCI-AEC for professional ethics, JAOP for the judicial sector specifically.

#### C. Risk Classification

The RAM classifies legal AI applications into four tiers calibrated to the actual risk profile of legal practice rather than directly transposing the EU Act's general categories. Tier 1 (Facilitative) covers legal database searches, document translation, and scheduling tools, subject only to mandatory disclosure. Tier 2 (Moderate Risk) covers AI-assisted research, drafting, and case summarisation, subject to certification under the Legal AI Certification Scheme (LACS), verification obligations, and periodic audit. Tier 3 (High Risk) covers predictive analytics, automated legal advice systems, and AI-assisted judicial decision support, subject to full conformity assessment, human override protocols, bias audits, and equity impact assessments. Tier 4 marks the boundary of the impermissible: autonomous adjudication systems, AI-generated evidence fabrication tools, and AI systems designed to manipulate legal processes fall entirely outside authorised use.

#### D. Survey Findings to RAM Provisions: The Evidence-Policy Bridge

**Table 7. Survey Finding to RAM Governance Provision Mapping (Selected)**

Survey Finding	Governance Gap Identified	RAM Provision
94.4% AI usage; no professional conduct rules cover it	No BCI rule on AI professional conduct	BCI-AEC Mandatory AI Ethics Guidelines under Advocates Act s. 49
76.4% find India's regulatory framework inadequate	No dedicated AI legislation	NARALP Act — Parliament to enact
Mean preparedness score 2.67 / 5.00	No mandatory AI literacy requirement	LACS Mandatory AI Ethics Training — Phase I rollout
93.0% support mandatory disclosure to courts and clients	No disclosure obligation in procedural court rules	AI Disclosure Statement — JAOP Court Rules amendment
77.8% concerned about client data privacy through AI tools	DPDPA 2023 silent on legal AI data processing	Data Protection and Confidentiality Framework (DPCF)
74.7% support mandatory AI audits in courts	No conformity assessment or audit requirement	AI Ethics Compliance Mechanism (AECM) — Tiers 2 and 3
Only 34.2% believe AI will improve access to justice equitably	No equity impact assessment requirement exists	Equity Audit — mandatory at Tier 3 before deployment

#### E. Phased Implementation

The RAM is designed for implementation in three phases, specifically to avoid the compliance overload risk that observers have identified in the EU AI Act's ambitious implementation schedule. Phase I (Years 1–2) establishes the governance infrastructure without heavy compliance burdens: BCI-AEC constitution, an Interim Professional AI Notification under the Advocates Act, mandatory AI disclosure in court submissions under JAOP guidelines,

and introduction of the National AI Governance Framework Bill in Parliament. Phase II (Years 3–4) operationalises NARALP, launches the LACS certification scheme, develops the Legal AI Explainability Standard (LAES), and rolls out mandatory AI ethics training. Phase III (Year 5 onwards) implements full Tier 3 conformity assessment requirements, institutionalises equity auditing, and establishes the India–Global South AI Governance Forum. This sequence is calibrated to build institutional capacity ahead of compliance demands, rather than imposing regulatory obligations that existing institutions cannot yet discharge.

## X. Conclusion

Surveys of legal professionals have a somewhat mixed reputation in the empirical legal literature: respondents may not know what they do not know, their stated preferences may not match their actual behaviour, and sample design inevitably introduces selection effects. These caveats apply here, and Section II acknowledges the most significant methodological boundary — the purposive focus on upper-tier practitioners. Nonetheless, the survey reported in this paper produces findings clear enough in their direction and consistent enough across professional sub-groups to carry genuine governance weight.

The core empirical picture is this: AI has entered India’s legal profession at scale (94.4 per cent adoption), without governance (76.4 per cent finding the framework inadequate), without adequate professional preparation (mean preparedness 2.67 out of 5), and with documented harm already occurring (twenty-three specific hallucination encounters recorded). Against this backdrop, the profession expresses strong, specific, and legally actionable governance preferences: disclosure (93 per cent), BCI ethics guidelines (90.4 per cent), mandatory training (84.3 per cent), mandatory auditing (74.7 per cent), and statutory governance models (76.0 per cent combined). The professional demand is not for some future governance: it is for governance now.

Relative to the thesis’s two hypotheses, the data are unambiguous. H1 — that India’s pre-AI statutory framework creates structural governance challenges — is confirmed not only by doctrinal analysis (which the other chapters supply) but by the direct experience of 76.4 per cent of the practitioners affected by that framework. H2 — that the EU’s governance framework can serve as a normative reference point for India — is supported by the convergence between the specific governance deficiencies that respondents identify and the provisions that the EU AI Act includes, and by the 71.3 per cent of EU-Act-aware respondents who consider it relevant or highly relevant to India’s governance design.

What the survey cannot do — and what the RAM in the doctoral thesis is designed to address — is move from empirical diagnosis to governance prescription. The survey documents what the profession thinks and experiences; the RAM translates those findings into specific, sequenced, constitutionally grounded governance provisions. Future research should extend the empirical base to district-court practitioners and the lower bar, where AI literacy and adoption patterns may differ significantly, and should revisit these questions in five years to measure whether any governance interventions that follow this research have moved the preparedness needle in the right direction.

## Acknowledgment

The author thanks the 229 legal professionals who gave their time to this survey. Doctoral supervision by Prof. (Dr.) Shobhna Jeet, School of Legal Studies, K.R. Mangalam University, is gratefully acknowledged. The research was conducted as part of a doctoral programme at K.R. Mangalam University, Gurugram.

## References

- [1] F. Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press, 2015.
- [2] M. Veale and F. Zuiderveen Borgesius, “Demystifying the Draft EU Artificial Intelligence Act,” *Computer Law Review International*, vol. 22, no. 4, pp. 97–112, 2021.
- [3] NITI Aayog, *National Strategy for Artificial Intelligence: #AIforAll*. New Delhi: Government of India, 2018.
- [4] *Mata v Avianca, Inc.*, No. 22-cv-1461 (PKC), S.D.N.Y., June 22, 2023 (sanctions for AI-fabricated citations).
- [5] *Lexos Media IP, LLC v Overstock.com, Inc.*, Case No. 22-2324-JAR, D. Kan., Feb. 2, 2026 (sanctions for AI-fabricated legal authority).
- [6] *ITAT Bengaluru, Buckeye Trust v PCIT*, ITA No. 1051/Bang/2024, December 2024 (recall of order citing non-existent cases).
- [7] *Delhi High Court, Greenopolis Welfare Association v Ireo Grace Realtech Pvt Ltd*, September 2025 (observations on impermissibility of AI-fabricated citations).
- [8] P. Panwar, “An Examination of Progressive and Protective Approaches in AI Regulation with Reference to the European Union and India,” Ph.D. dissertation, School of Legal Studies, K.R. Mangalam University, Gurugram, 2026.
- [9] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act), OJ L, 12.7.2024.
- [10] *Justice K.S. Puttaswamy (Retd.) & Anr. v. Union of India & Ors.*, (2017) 10 SCC 1 (Supreme Court of India, right to privacy as fundamental right).
- [11] T. Wischmeyer and T. Rademacher (eds.), *Regulating Artificial Intelligence*. Cham: Springer, 2020.
- [12] S. Sarangi and P. Sharma, *Artificial Intelligence: Evolution, Ethics and Public Policy*. Abingdon: Routledge India, 2018.
- [13] K.D. Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge: Cambridge University Press, 2017.

- [14] Kerala High Court, Policy Regarding Use of Artificial Intelligence Tools in District Judiciary, July 19, 2025.
- [15] Supreme Court of India, *Neha Lal v Abhishek Kumar*, 2026 INSC 73, January 2026.