

System 1 Lawyering and the Governance of AI in Legal Practice:

A Practitioner's Framework for Supervised Intelligence

Shivam Shukla

Advocate, High Court of Judicature at Allahabad

Co-Founder, Bharat Litigation Law | shivam@advshivamshukla.in

March 2026

Abstract

The integration of generative artificial intelligence into legal practice has outpaced the development of professional standards governing its use. In February 2026, the Supreme Court of India took cognizance of a trial court judge's reliance on four AI-fabricated judgments in Gummadi Usha Rani v. Sure Mallikarjuna Rao (SLP(C) 7575/2026, order dated 27.02.2026), declared it misconduct, and issued notice to the Attorney General, the Solicitor General, and the Bar Council of India. This paper introduces the first practitioner-driven, cognitively grounded governance framework for artificial intelligence in legal practice. It identifies the cognitive failure mode that AI introduces into legal reasoning — a regression from deliberate analytical thinking (System 2) to automatic pattern-acceptance (System 1), drawing on Kahneman's dual-process theory — and proposes a comprehensive framework comprising five professional doctrines (Cognitive Sovereignty, Epistemic Diligence, Fiduciary Confidentiality, Adversarial Anticipation, and Intellectual Amplification), a five-stage Supervised Intelligence Model (SIM), and a four-step Accountability Review Test (ART). The paper grounds this framework in the existing statutory architecture of the Advocates Act, 1961 and the Bar Council of India Rules on Professional Conduct, and situates it within the emerging global jurisprudence on AI misuse in legal proceedings. It presents the full text of a Declaration on Responsible AI Use in Legal Practice, designed for voluntary adoption by individual advocates, courts, and bar regulatory bodies. The Declaration is platform-independent and jurisdiction-adaptable, grounded in professional obligations that pre-date and will survive every technological development.

Keywords: artificial intelligence in law, legal ethics, responsible AI, supervised intelligence, cognitive sovereignty, legal profession, advocate obligations, AI governance, legal technology, professional standards, dual-process theory, professional misconduct, AI hallucination, Bar Council regulation, verification protocols

I. Introduction

The legal profession globally confronts a governance deficit: generative artificial intelligence tools are being used by advocates, judges, and court staff in the preparation and adjudication of legal matters, yet no jurisdiction has established a comprehensive, binding professional standard that governs the cognitive, ethical, and procedural dimensions of this use. The consequences of this gap are no longer theoretical.

This paper makes three contributions: (i) identification of a distinct cognitive failure mode in AI-assisted legal reasoning ('System 1 Lawyering'), grounded in Kahneman's dual-process theory; (ii) formulation of a five-doctrine professional framework anchored in existing statutory obligations under the Advocates Act, 1961; and (iii) operationalisation of that framework through the Supervised Intelligence Model (SIM) and the Accountability Review Test (ART).

In India, as of March 2026, no Bar Council has issued a binding professional standard, continuing legal education module, or standardised verification protocol governing the use of artificial intelligence in legal practice. Courts are addressing AI-related professional failures through sanctions and censure on a case-by-case basis. The most significant development occurred on 27 February 2026, when the Supreme Court of India, in *Gummadi Usha Rani v. Sure Mallikarjuna Rao* (SLP(C) 7575/2026), took cognizance of a trial court judge's reliance on four AI-fabricated judgments that did not exist in any law reporter. The Court declared this misconduct, issued notice to the Attorney General, the Solicitor General, and the Bar Council of India, and appointed Shyam Divan, Senior Advocate, as *amicus curiae*.²

This paper does not propose regulation. It proposes a voluntary professional standard — a Declaration — that any advocate may adopt unilaterally, and that any court or bar regulatory body may incorporate into its rules. The framework is grounded not in the features of any specific AI platform, but in the enduring professional obligations of competence, diligence, confidentiality, and independent judgment that define the advocate's function.

II. The Legal Foundation: Existing Professional Obligations

The framework proposed in this paper does not create new obligations. It operationalises obligations that already exist under the Advocates Act, 1961 and the Bar Council of India Rules on Standards of Professional Conduct and Etiquette framed under Section 49(1)(c) of the Act. The advocate's duty to use AI responsibly is not a novel regulatory imposition — it is a necessary implication of duties that have governed the profession since 1961.

A. The Advocates Act, 1961

Section 7(1)(b) of the Advocates Act, 1961 empowers the Bar Council of India to "lay down standards of professional conduct and etiquette for advocates." Section 49(1)(c) confers the corresponding rule-making power: the BCI may make rules prescribing "the standard of professional conduct and etiquette to be observed by advocates," subject to the approval of the Chief Justice of India. These provisions establish the statutory basis for any professional standard governing AI use. The BCI has the existing authority — and, this paper argues, the existing

obligation — to issue such standards.

Section 35 of the Act provides the enforcement mechanism. Where a State Bar Council has reason to believe that an advocate has been guilty of “professional or other misconduct,” it shall refer the case to its disciplinary committee, which may reprimand the advocate, suspend the advocate from practice, or remove the advocate’s name from the State roll. The disciplinary framework already encompasses every form of professional failure that AI misuse may produce: filing fabricated authorities, submitting unverified citations, breaching client confidentiality through AI tool inputs, and abdicating independent professional judgment.

Section 49(1)(ah) specifies that the BCI may prescribe “the conditions subject to which an advocate shall have the right to practise.” Read together with Sections 7(1)(b) and 35, this creates a complete statutory architecture: the BCI defines the professional standard, the advocate’s right to practise is conditioned upon compliance, and breach constitutes misconduct subject to disciplinary proceedings. No new legislation is required. The statutory infrastructure exists.

B. Bar Council of India Rules on Professional Conduct

Part VI, Chapter II of the BCI Rules sets out the Standards of Professional Conduct and Etiquette.³ Section I prescribes the advocate’s duty to the court: the advocate shall not knowingly refer to any decided case which has been overruled, or to any statute which has been repealed; shall not mislead the court; and shall maintain towards the court a respectful attitude bearing in mind that the dignity of the judicial office is essential for the survival of a free community. Section II prescribes the duty to the client, including the duty of full and frank disclosure and the fiduciary obligation of confidentiality.

Each of these existing duties acquires heightened significance in an AI-assisted environment. The duty not to cite overruled authorities becomes, in the AI context, the duty not to cite AI-fabricated authorities — a more severe failure, because the authority may never have existed at all. The duty not to mislead the court is violated not only by affirmative misrepresentation but by the passive adoption of AI output without verification. The fiduciary duty of confidentiality is directly engaged when client information is entered into third-party AI systems with opaque data handling practices. The Five Doctrines proposed in this paper are not supplements to these duties. They are their operational translation for the AI-assisted practice environment.

III. The Cognitive Failure Mode: System 1 Lawyering

Daniel Kahneman’s dual-process theory, articulated in *Thinking, Fast and Slow* (2011),¹ distinguishes between two modes of cognition. System 1 operates automatically, quickly, and with little conscious effort — it is the mode of pattern recognition, intuitive judgment, and heuristic shortcuts. System 2 is deliberate, effortful, and analytical — it is the mode of legal reasoning, statutory construction, and doctrinal analysis.

Generative AI introduces a specific professional hazard: it produces output that is linguistically fluent, structurally coherent, and formatted with the surface markers of legal authority

— case names, citation formats, statutory references, and analytical structure — regardless of whether the underlying content is accurate, fabricated, or misleading. This fluency activates System 1 processing in the receiving advocate or judge: the output *looks* right, and the professional’s cognitive system defaults to acceptance rather than verification.

The author terms this failure mode **System 1 Lawyering**: the regression from deliberate legal analysis to automatic pattern-acceptance when confronted with AI-generated output that mimics the form of professional work product. The failure in *Gummadi Usha Rani* was not technological — it was cognitive. The judge’s System 2 was never engaged.

The governance response, therefore, must be cognitive and professional, not merely technological. No AI platform’s built-in verification feature can substitute for the advocate’s independent engagement of System 2 reasoning. The standard must be platform-independent and anchored in the advocate’s own professional discipline.

IV. The Five Doctrines

The framework proposes five professional doctrines as binding principles governing the advocate’s conduct whenever artificial intelligence is part of the work. The doctrines follow a progressive logic: protect → verify → guard → defend → amplify.

Doctrine 1 — Cognitive Sovereignty

The advocate’s independent legal judgment is non-delegable and sovereign. Artificial intelligence may assist research, analysis, and drafting; it may never substitute for the advocate’s authorship or responsibility. The machine produces; the advocate judges. Sovereignty is not a rule to follow — it is a practice to inhabit, rebuilt daily in every matter, in every decision to frame the legal question before opening the tool.

Doctrine 2 — Epistemic Diligence

Every AI-generated fact, citation, inference, or conclusion is presumptively unreliable until independently verified against primary sources. The advocate must possess a working understanding of the system’s architecture, biases, and failure modes. Partial diligence is professional negligence. The advocate who files a citation they have not read is not merely risking professional sanction — they are misrepresenting their knowledge to the court.

Doctrine 3 — Fiduciary Confidentiality

Client confidences are held under absolute fiduciary duty. No client data, prompt, document, or strategy may be submitted to any AI system unless the advocate retains complete control, auditability, and legal accountability over the handling of that information. The generative AI tool is a convenience; the client’s confidence is a duty.

Doctrine 4 — Adversarial Anticipation

The advocate must assume, as a baseline of professional competence, that opposing counsel, corporations, regulators, and the State are already deploying artificial intelligence at scale. This requires the development of both offensive mastery and forensic defensive skills to detect, challenge,

and neutralise AI-generated evidence, arguments, decisions, and pleadings.

Doctrine 5 — Intellectual Amplification

When governed by the first four doctrines, artificial intelligence becomes a legitimate and powerful force multiplier for rigorous legal reasoning. Properly supervised AI expands research capacity, accelerates drafting, democratises access to legal knowledge, and strengthens the administration of justice. The advocate carries an affirmative professional duty to harness this amplification responsibly.

V. The Supervised Intelligence Model (SIM)

The Supervised Intelligence Model operationalises the Five Doctrines into a sequential, stage-gated methodology. It comprises five stages: three human-only and two AI-assisted under supervision.

Stage 1 — Legal Framing (Human Only)

The advocate identifies and frames the legal question independently, before any AI engagement. Governed by Doctrine 1 (Cognitive Sovereignty). The legal question must originate in the advocate's own analysis.

Stage 2 — Pattern Expansion (AI-Assisted)

The advocate uses AI to expand the range of authorities, arguments, and perspectives relevant to the framed question. Governed by Doctrines 2 and 3. AI is deployed as a research accelerator, not as a reasoning engine.

Stage 3 — Doctrinal Reconstruction (Human Dominant)

The advocate reconstructs the legal argument from first principles, using AI-expanded material as raw input — not as a draft to be edited. The critical distinction is between editing within the AI's structure (System 1) and reconstructing the argument independently (System 2). Governed by Doctrines 1 and 2.

Stage 4 — Verification (Human Only)

Every citation, statutory provision, and factual assertion is independently verified against primary sources. Verification means: the authority exists, the text is accurate, the provision is in force, and it supports the proposition for which it is advanced. Governed by Doctrine 2.

Stage 5 — Strategic Judgment (Human Only)

All final strategic, tactical, and doctrinal decisions are made by the advocate alone. AI cannot assess the bench, the opposing counsel, the client's priorities, or the political and institutional dynamics of the matter. Governed by Doctrines 1, 4, and 5.

VI. The Accountability Review Test (ART)

The Accountability Review Test is a four-step self-assessment that every AI-assisted work product must pass before submission. It may also be applied forensically to the opponent’s work product.

Step 1 — Framing Diligence

“Was the legal framing the advocate’s own, or was it adopted from AI output?”

Step 2 — Supervisory Adequacy

“Was AI output subjected to meaningful, substantive oversight at every material stage?”

Step 3 — Verification Completeness

“Was every citation, fact, and legal proposition independently verified against primary sources?”

Step 4 — Judgment Independence

“Was the final strategic judgment exercised by the advocate without delegation to any system?”

A negative answer to any step requires the advocate to return to the relevant SIM stage before the work product may be filed or relied upon.

VII. Comparative Context

The governance challenge identified in this paper is not unique to India. Jurisdictions across the common law world have confronted the consequences of AI misuse in legal proceedings, though none has yet produced a comprehensive, cognitively grounded professional standard of the kind proposed here.

In the United States, the seminal case is *Mata v. Avianca, Inc.* (No. 22-cv-1461, S.D.N.Y., June 2023),⁴ in which counsel submitted a brief containing six fabricated case citations generated by ChatGPT. The court imposed monetary sanctions and referred the attorneys for disciplinary proceedings. The case prompted the American Bar Association to issue Formal Opinion 512 (July 2024),⁵ which confirmed that existing duties of competence (Model Rule 1.1), diligence (Model Rule 1.3), and candour to the tribunal (Model Rule 3.3) apply to AI-assisted legal work. Multiple federal courts — including the Northern District of Texas, the Eastern District of Pennsylvania, and the Fifth Circuit — have since adopted standing orders or local rules requiring disclosure of AI use and certification that AI-generated content has been independently verified.

In the United Kingdom, the Judiciary of England and Wales issued guidance in December 2023 cautioning legal professionals against reliance on AI-generated legal research without independent verification. The Bar Standards Board has indicated that existing professional obligations — including the duty not to mislead the court — encompass AI-assisted work product, though no comprehensive framework has been issued.

The common thread across these jurisdictions is the recognition that existing professional obligations already govern AI use — the gap is not in the law but in its operationalisation. This is precisely the gap that the Declaration and its supporting frameworks are designed to fill. The Five

Doctrines translate abstract duties into concrete professional practices; the SIM provides a stage-gated methodology; and the ART provides a verifiable self-assessment. No comparable integrated framework has been proposed in any jurisdiction to date.

VIII. Declaration on Responsible AI Use in Legal Practice

The following Declaration is presented as a voluntary professional standard. It may be adopted in whole or in part by individual advocates, incorporated into court rules or practice directions, or adapted by bar regulatory bodies as the basis for formal guidelines.

DECLARATION ON RESPONSIBLE AI USE IN LEGAL PRACTICE

Version 1 — March 2026

Declaration of Principles

- 1. Supervised Intelligence as the Professional Standard.** The use of artificial intelligence in legal practice is governed by the principle of supervised intelligence: the advocate retains complete authority over every stage of legal reasoning, and AI operates exclusively as an instrument subject to the advocate's professional judgment.
- 2. Non-Delegable Authorship Responsibility.** The advocate bears complete and non-delegable authorship responsibility for every submission, argument, and citation filed or relied upon in any proceeding.
- 3. Mandatory Independent Verification.** Verification of citations, statutory provisions, and legal analysis constitutes a core professional duty that cannot be outsourced to any technological system or platform.
- 4. Platform Independence.** Professional standards governing the use of artificial intelligence apply uniformly across all tools and models and are not tied to the features or verification mechanisms of any specific platform.
- 5. Preservation of Professional Judgment.** All strategic, doctrinal, and tactical judgments must remain the product of the advocate's independent analysis. AI systems may assist in research or drafting but cannot substitute for the advocate's cognitive engagement.
- 6. Disclosure and Structured Supervision.** Where the use of artificial intelligence is material to the preparation of a submission, the advocate may disclose such use and confirm compliance with a structured supervisory protocol ensuring independent verification and judgment at every material stage.

These principles do not restrict the use of artificial intelligence. They govern it in a manner consistent with long-standing professional obligations that pre-date and will survive every technological development.

This Declaration may be adopted, adapted, or incorporated by courts, Bar Councils, and individual advocates.

IX. Sample AI Usage Declaration Clause

The following clause is proposed as a standard form that may be appended to any AI-assisted filing or submission:

“The undersigned advocate certifies that (a) artificial intelligence was used in the research and drafting stage of this submission, (b) all AI-generated content has been subjected to independent verification in accordance with the Declaration on Responsible AI Use in Legal Practice dated March 2026, and (c) the advocate accepts full authorship responsibility for the final content.”

X. Application and Adoption

The Declaration and its supporting frameworks are designed for three levels of adoption. First, individual advocates may adopt the Declaration voluntarily as a personal professional standard, applying the SIM methodology and ART self-assessment to their daily practice. Second, courts may incorporate the Declaration’s principles into practice directions or standing orders, requiring AI usage disclosure clauses on filings where AI tools were used in preparation. Third, bar regulatory bodies may use the Five Doctrines as the conceptual foundation for formal rules governing AI use in legal practice, and the SIM as the basis for continuing legal education modules.

The framework is jurisdiction-adaptable. While this paper draws its illustrative case law from the Indian experience — *Gummadi Usha Rani* being among the most significant judicial interventions on AI misuse in any common law jurisdiction to date — the underlying professional obligations of competence, diligence, confidentiality, and independent judgment are universal to the legal profession. The Five Doctrines, the SIM, and the ART are platform-independent and do not depend on the features, limitations, or governance commitments of any specific AI provider.

XI. Conclusion

The profession needs a proactive standard — not a reactive response to judicial censure. The failure mode that AI introduces into legal reasoning is cognitive, not technological. The governance response must therefore be cognitive and professional: a framework that restores the advocate’s System 2 engagement at every stage where professional judgment is required.

The statutory architecture already exists. The Advocates Act, 1961 empowers the Bar Council of India to prescribe standards of professional conduct. The BCI Rules already impose duties of competence, candour, and confidentiality. The Five Doctrines, the Supervised Intelligence Model, and the Accountability Review Test translate these existing obligations into an operational framework for the AI-assisted practice environment. The Declaration on Responsible AI Use in Legal Practice offers that framework. It asks nothing of technology. It asks everything of the advocate.

Notes

¹ Daniel Kahneman, *Thinking, Fast and Slow* (Farrar, Straus and Giroux, 2011). Kahneman distinguishes System 1 (fast, automatic, heuristic) from System 2 (slow, deliberate, analytical). The application to legal reasoning in AI-assisted environments is the author's.

² *Gummadi Usha Rani v. Sure Mallikarjuna Rao*, SLP(C) No. 7575/2026, order dated 27 February 2026 (Supreme Court of India). The Court took cognizance of a trial court judge's reliance on four AI-fabricated judgments, declared it misconduct, and issued notice to the Attorney General, the Solicitor General, and the Bar Council of India. Shyam Divan, Senior Advocate, was appointed *amicus curiae*.

³ Bar Council of India Rules framed under Section 49(1)(c) of the Advocates Act, 1961, Part VI, Chapter II: "Standards of Professional Conduct and Etiquette." Section I: Duty to the Court. Section II: Duty to the Client. Section IV: Duty to Colleagues. These Rules require the approval of the Chief Justice of India under the proviso to Section 49(1).

⁴ *Mata v. Avianca, Inc.*, No. 22-cv-1461 (LJL), 2023 WL 4114965 (S.D.N.Y. June 22, 2023) (Castel, J.). The court imposed monetary sanctions on counsel who submitted a brief containing six fabricated case citations generated by ChatGPT.

⁵ American Bar Association, Formal Opinion 512: "Generative Artificial Intelligence Tools" (July 29, 2024). The Opinion confirmed that Model Rules 1.1 (Competence), 1.3 (Diligence), 1.6 (Confidentiality of Information), and 3.3 (Candor Toward the Tribunal) apply to AI-assisted legal work.

About the Author

Shivam Shukla is an Advocate practising at the High Court of Judicature at Allahabad, enrolled with the Bar Council of Uttar Pradesh. He is the author of *AI for Indian Advocates: The Practitioner's Standard for Supervised Intelligence* (forthcoming, Vinod Law Publications, 2026) and co-founder of Bharat Litigation Law. His practice spans constitutional and administrative law, revenue and property law, criminal law, and family and matrimonial law. The Declaration and the supporting White Paper ("System 1 Lawyering") are available at advshivamshukla.in/standard.

Contact: shivam@advshivamshukla.in

SSRN Author ID: 10920908