

**SUBMISSION OF
COMMENTS/SUGGESTIONS**

on the

Draft 'Regulations for Use of Artificial Intelligence (AI) in Courts, 2026'

published by the
Artificial Intelligence Committee,
Supreme Court of India
vide Notice dated 03.06.2026



SUBMITTED BY:
The Dialogue[1]



July 2026

[1] with expert inputs from Ms. Saraswathy Vaidhyamanthan, Domain Expert, Assistant Professor, School of Law, BML Manjal University, Gurgaon

SUBMISSION OF COMMENTS/SUGGESTIONS ON THE DRAFT 'REGULATIONS FOR USE OF ARTIFICIAL INTELLIGENCE (AI) IN COURTS, 2026'

Authors: Artificial Intelligence Policy Team

The Dialogue is a public policy think tank with a vision to drive a progressive narrative in India's policy discourse. Founded in 2017, we believe in facilitating well-researched policy debates at various levels to help develop a more informed citizenry, on areas around technology and development issues. The Dialogue has been ranked as the world's Top 10 think tanks to watch out for, by the Think Tank and Civil Societies Programme (TTCSP), University of Pennsylvania in their 2020 and 2021 rankings.

For more information visit
thedialogue.org.in

Publication date: July 6, 2026

Catalogue no.: TD/ET/WC/0726/11

Suggested Citation: Artificial Intelligence Policy Team. (2026). Submission of Comments on the Draft Regulations for Use of Artificial Intelligence (AI) in Courts, 2026. The Dialogue.

Disclaimer- This submission has been prepared by the Artificial Intelligence Team at The Dialogue. The contents may be reproduced, in whole or in part, provided appropriate acknowledgement and attribution are given to the Artificial Intelligence Team and The Dialogue.[®]

CONTENTS

I. PREFATORY REMARKS	1
II. ON THE FORM AND CHARACTER OF THE INSTRUMENT	1
A. The Choice of ‘Regulations’ as the Legislative Vehicle	1
III. ON THE INSTITUTIONAL ARCHITECTURE	2
A. Bureaucratisation of Governance A Note of Caution.....	2
B. Capacity Building and Training The Most Critical Gap.....	2
IV. ON THE HUMAN-IN-THE-LOOP REQUIREMENT – SUBSTANCE OVER FORM.....	3
V. ON THE AI INCIDENT DATABASE THE CASE FOR OPENNESS	4
VI. ON GRIEVANCE REDRESSAL THE GAP BETWEEN ENTITLEMENT AND ACCESS	4
VII. ON THE DIGITAL DIVIDE AND INCLUSIVITY	5
VIII. ON THE DISTRICT COURT STRUCTURE THE MISSING CENTRE	6
IX. OTHER SPECIFIC OBSERVATIONS	6
A. Prohibition on Risk Scoring (Regulation 20(1)(d)).....	6
B. AI Content Verification Authority (Regulation 44)	6
C. Controlled Environment Testing (Regulation 36)	7
D. Bridging the Capacity Gap in “In-House” Audits (Regulation 38).....	7
E. Tightening Private Sector Controls and Securing Data Independence (Regulation 46)	7
F. Defining Verification Standards for Legal Professionals (Regulation 43)	7
X. CONCLUSION AND SUMMARY OF RECOMMENDATIONS	7

To,

The Member Secretary,
Artificial Intelligence Committee,
Supreme Court of India,
New Delhi
Email: office.regcc@sci.nic.in
Date: 19.06.2026

Subject: Submission of Comments and Suggestions on the Draft ‘Regulations for Use of Artificial Intelligence (AI) in Courts, 2026’.

Reference: Notice issued by the Member Secretary, AI Committee, Supreme Court of India, dated 03.06.2026.

I. PREFATORY REMARKS

The publication of the draft *Regulations for Use of Artificial Intelligence in Courts, 2026* (“Draft Regulations”) by the Artificial Intelligence Committee of the Supreme Court of India marks a watershed moment in the evolution of India’s justice delivery system. The initiative reflects commendable institutional foresight and a willingness to engage proactively with the transformative potential of Artificial Intelligence (“AI”) in the judicial sphere. The underlying aspiration to harness AI as a tool for reducing pendency, widening access to justice, and improving administrative efficiency is not only welcome, but overdue.

These comments are submitted in the spirit of constructive engagement, with the sincere hope that the final regulations will be calibrated to reflect not merely what is desirable in principle, but what is implementable in practice across the vast and structurally diverse landscape of India’s courts; from the Supreme Court of India to the remotest District Courts and Tribunals. The suggestions herein draw upon a reading of the Draft Regulations alongside the practical realities of the Indian judicial system, including questions of institutional capacity, existing infrastructure, the digital divide, and the psycho-social dimensions of human-AI interaction in adjudicatory settings.

The comments are organised thematically and, where possible, are cross-referenced to specific regulations in the Draft.

II. ON THE FORM AND CHARACTER OF THE INSTRUMENT

A. The Choice of ‘Regulations’ as the Legislative Vehicle

At the outset, it is necessary to flag a foundational question about the form of the instrument. The Draft has been titled ‘Regulations’ and is expressed to be made under the aegis of the AI Committee of the Supreme Court of India. While the judiciary unquestionably possesses the authority to regulate its own internal functioning, including the adoption of administrative tools and technologies, the scope of the Draft Regulations is remarkably expansive. It creates institutional bodies with quasi-regulatory powers (the Apex Body, AI Committees, CoRE-AI), prescribes obligations upon private entities, mandates disclosures and audits, and establishes grievance redressal mechanisms; a sweep that is ordinarily the province of primary legislation enacted by Parliament or delegated legislation made under statutory authority.

The creation of a Centre of Research and Excellence on Artificial Intelligence (“CoRE-AI”) under Regulation 32, with powers to conduct original research, evaluate AI tools, and maintain a centralised record of AI systems available to public institutions, is, with respect, a function more appropriately vested by an Act of Parliament than by a court-issued regulatory instrument. Similarly, the obligations imposed upon private entities under Chapter VI, including mandatory indemnity clauses, sovereign cloud deployment requirements, and prohibition on retraining of models using court data, derive their legitimacy from contract and procurement policy rather than from any express statutory delegation.

It is respectfully submitted that the Supreme Court and High Courts, exercising their inherent powers under Articles 142 and 145 of the Constitution, and under the respective High Court Acts and Letters Patent, have ample authority to regulate the internal use of AI within court processes. However, the provisions that extend beyond court administration, particularly those governing the conduct of private entities, advocates, and litigants in submitting AI-generated content, would benefit from either a statutory anchor or from being incorporated into the Rules of Court, which carry express procedural authority.

It is recommended that the AI Committee examine whether certain provisions ought to be: (i) incorporated as amendments to the Rules of Court of the Supreme Court and respective High Courts; (ii) recommended to Parliament for enactment as part of the Digital India framework; or (iii) introduced as Practice Directions with a sunset clause pending comprehensive legislation. A regulation without clear statutory backing risks contestation and inconsistent application across jurisdictions.

III. ON THE INSTITUTIONAL ARCHITECTURE

A. Bureaucratisation of Governance | A Note of Caution

Chapter IV of the Draft creates a multi-tiered institutional structure: the Apex Body, five Committees of the Apex Body (Regulations 25–30), the CoRE-AI (Regulation 32), AI Committees in each court (Regulation 33), and dedicated AI Secretariats (Regulation 34). While institutional rigour is necessary, there is a real risk that the aggregation of so many overlapping bodies, each with its own mandate, composition, procedures, and reporting obligations, will result in governance paralysis rather than governance efficiency.

Consider the approval pathway for a single AI Tool: the AI Secretariat may grant expedited approval for administrative tools under Regulation 34(4); the AI Committee approves AI Systems under Regulation 33(3)(a) following Technical and Ethical Impact Assessments under Regulation 35; the Apex Body sets minimum mandatory standards under Regulation 24(a) and approves systems under Regulation 24(b); and the CoRE-AI evaluates tools for use in court systems under Regulation 32(3)(b). It is not always clear where one authority's power ends and another begins. The Draft would benefit from a streamlined decision matrix that maps the exact pathway for different categories of AI Tools, to avoid unnecessary duplication and delay.

It is also observed that the Apex Body, as constituted under Regulation 22(2), is predominantly judicial in composition. While this is appropriate for judicial governance, the effective regulation of AI, a rapidly evolving technology will require meaningful and continuous input from technologists, ethicists, and domain experts, who may be co-opted only on a 'case-to-case or standing basis' with the permission of the Chief Justice of India (proviso to Regulation 22(2)). This structure may inadvertently limit the Apex Body's access to the specialised expertise that will be essential for sound AI governance.

B. Capacity Building and Training | The Most Critical Gap

Chapter VIII, which addresses capacity building and training, is perhaps the most consequential chapter of the Draft Regulations and, unfortunately, the least developed. Regulation 49 prescribes that all Judges, advocates, and court staff "shall receive regular, structured training" and that training "shall be accessible to all such persons, including those in district courts" and "shall be offered in a manner that accounts for linguistic diversity." These obligations are welcome but are, at present, aspirational rather than operational.

The challenge of training across the Indian judicial system cannot be overstated. There are over 25 million pending cases across approximately 24,000 courts. District and subordinate courts operate under severe resource constraints, with sanctioned judicial strength consistently falling short of actual strength. The judiciary comprises over 20,000 judicial officers at the district and subordinate court level, many of whom operate in single-court jurisdictions without dedicated support staff. The proposition that all of these officers will receive "regular, structured training" on "the technical, legal and ethical dimensions of AI" within a timeframe and budget that is not specified in the Draft Regulations, with training programmes to be "reviewed at least once in every two years" (Regulation 51), underestimates the scale of the task.

It is submitted that the regulations should prescribe:

- i. A minimum number of training hours per year for judicial officers at each tier of the system.
- ii. A dedicated allocation within court budgets for AI literacy programmes, separate from general judicial training budgets.
- iii. A certification or competency assessment mechanism before a judicial officer is permitted to supervise or verify AI outputs in adjudicatory contexts.
- iv. Regional language modules for training, developed in collaboration with State Judicial Academies, given that many subordinate court officers may not have working proficiency in English.

Without a credible capacity-building framework, the regulations risk creating a formal structure in which AI tools are adopted, but not understood; a scenario that is potentially more dangerous than no adoption at all.

IV. ON THE HUMAN-IN-THE-LOOP REQUIREMENT – SUBSTANCE OVER FORM

The Draft Regulations make frequent and commendable reference to the ‘Human-in-the-Loop’ (“HITL”) requirement, defined under Regulation 3(1)(zb) as a governance process in which AI outputs are subject to “mandatory human review, supervision and verification, with final decision-making authority, responsibility, and accountability resting at all times with a human.” HITL is also mandated under Regulation 12(2) for high-risk applications affecting personal liberty, and under Regulation 20(1)(c) for adjudicative and sentencing functions.

The regulatory intent is sound, but the design of the HITL requirement raises a concern that is not merely procedural – it is psychological. Research in cognitive science and behavioural law and economics has consistently demonstrated the phenomenon of *automation bias*: the tendency of human reviewers to defer uncritically to algorithmic or AI-generated outputs, particularly when those outputs are presented with apparent authority, numerical precision, or institutional validation. In a judicial setting, where case backlogs exert enormous time pressure, the risk that HITL verification becomes a rubber-stamping exercise, perfunctory, formulaic, and devoid of independent scrutiny, is significant and deserves explicit regulatory attention.

Closely related is the risk of *confirmation bias*: a judicial officer who has already formed a tentative view on a matter may unconsciously interpret AI-generated outputs, whether case summaries, precedent retrieval, or document analysis, in a manner that confirms rather than challenges that view. The AI system, trained on historical data that may itself encode existing patterns of adjudication, could thus create a feedback loop that entrenches rather than corrects systemic tendencies in judicial decision-making.

The Draft Regulations are silent on these psycho-social dimensions of AI use. It is submitted that the regulations should, at a minimum:

- i. Include in mandatory training curricula (Regulation 49) a specific module on automation bias, confirmation bias, and the conditions under which AI outputs should be treated with heightened scepticism.
- ii. Mandate that AI Systems used in legal research or precedent retrieval present outputs in a format that actively alerts the judicial officer to the limitations of the output, including potential gaps in coverage, dataset bias, or hallucination risk, rather than presenting conclusions in a manner designed to project false confidence.
- iii. Institute mandatory interface controls that enforce active human deliberation during the verification process. To prevent judicial staff from blindly trusting algorithmic outputs, the software should block one-click approvals for AI-generated summaries or scrutiny reports. Users must be compelled to manually enter their reasoning or confirm the cross-checking of citations, ensuring the oversight mandate translates into genuine, independent scrutiny rather than a superficial administrative step.

Regulation 8(3) states that AI-generated output “shall be treated as advisory in nature and reasonable care shall be taken to verify the accuracy of such output before the same is utilised”, with a proviso permitting the officer to dispense with verification for reasons recorded in writing. The dispensation proviso, while operationally necessary, must not become a routine escape from the verification obligation. The regulations should clarify that dispensation with verification is an exception and not a norm, and that it shall be subject to review during audits.

V. ON THE AI INCIDENT DATABASE | THE CASE FOR OPENNESS

Regulation 39 requires every AI Secretariat to maintain an AI Incident Database, recording “all AI Incidents, including their type, cause, manner of occurrence, consequences and the remedial measures taken.” This is a commendable provision. However, the Draft Regulations do not address the question of public accessibility of the Incident Database, and the only transparency mechanism is the Annual Transparency Report under Regulation 45.

It is submitted that the AI Incident Database should be made publicly accessible, subject to appropriate redactions for data protection, security, and confidentiality. There are three principal reasons for this:

- i. First, judicial accountability. AI systems deployed in courts affect the rights and liberties of litigants. Litigants and their counsel have a legitimate interest in knowing whether an AI system that participated in their matter has a documented record of malfunction, bias, or error. Restricting incident data to internal circulation defeats the purpose of maintaining the database.
- ii. Second, systemic learning. India’s 25 High Courts and thousands of subordinate courts will, over time, generate a rich body of AI incident experience. This knowledge is a public good. Restricting it to internal channels, even with the mechanism of communication between AI Secretariats under Regulation 39(2), will slow the pace of learning and correction across the system.
- iii. Third, public trust. The legitimacy of AI adoption in courts will depend, in large measure, on public confidence in the system’s ability to acknowledge and correct its own errors. A publicly accessible, anonymised incident log would demonstrate institutional integrity and strengthen trust.

It is accordingly recommended that Regulation 39 be amended to provide for the publication of a curated, anonymised version of the AI Incident Database on the official websites of the respective High Courts and the Apex Body, at intervals not exceeding six months, with appropriate redactions as may be determined by the AI Secretariat. The Annual Transparency Report under Regulation 45 should also include a consolidated summary of incident trends, patterns, and systemic recommendations.

VI. ON GRIEVANCE REDRESSAL | THE GAP BETWEEN ENTITLEMENT AND ACCESS

Chapter IX of the Draft Regulations, which addresses grievance redressal, is, with respect, the weakest part of an otherwise thoughtful document. Regulation 52 provides that any party harmed by the prohibited use of AI under Regulation 20 “may file an application at the earliest opportunity to the Court in which the relevant AI System was or is being used.” The High Court may prescribe principles and procedures for adjudication of such grievances and may prescribe formats for complaints “so as to ensure accessibility and ease of use for all stakeholders, including persons with limited legal literacy.”

While the intention is laudable, the mechanism as designed raises several concerns:

- i. **Identification of the harm.** How will a litigant, particularly one without legal representation, know that an AI system was used in their matter, let alone that it was used in a prohibited manner? The disclosure obligations under Regulation 43 address court-initiated use but do not always guarantee that AI involvement will be communicated in a manner comprehensible to the affected party. Without awareness, the right to grieve is purely theoretical.
- ii. **No time-bound resolution.** The Draft specifies no timeline within which grievances must be decided. In a system already burdened with pendency, the risk that AI-related grievances will themselves become long-pending matters is real.
- iii. **Remedy is limited to prohibited uses.** Regulation 52(1) restricts the grievance mechanism to harm caused by *prohibited* uses under Regulation 20. But harm may equally arise from the *permitted* use of AI, for instance, a transcription error in automated court proceedings, a faulty case summary that leads to a misconducted hearing,

or a biased document classification. The restriction to prohibited uses creates a significant gap in the remedial framework.

iv. Impracticality of fault-based claims for algorithmic harm. The proposed redressal framework assumes litigants can effectively prove negligence or pinpoint exactly how a system failed. However, the inherent opacity of machine learning models and the complex web of developers and operators make establishing direct causation virtually impossible for an ordinary citizen. Looking back at the Supreme Court's own historical rulings on inherently hazardous activities and state accountability for rights violations (such as the principles established in *M.C.Mehta and Nilabati Behera*), the rules should impose strict, non-delegable accountability on the state. The judicial system must absorb the liability for any technological misfires or hallucinations, preventing the state or its tech contractors from escaping liability when a litigant's rights are compromised.

It is recommended that:

- i.** A dedicated AI Grievance Officer be designated in each District Court and High Court, who is independent of the registry and is accessible to unrepresented litigants.
- ii.** Grievances be capable of being filed with the AI Secretariat as an alternative forum to the court in which the AI system was used, to address concerns of institutional bias.
- iii.** Time-bound resolution, not exceeding 60 days, be prescribed, with a right of appeal to the AI Committee.
- iv.** The scope of grievances be expanded beyond prohibited uses to cover any AI-related harm to a party in a court process, with the AI Committee having jurisdiction over systemic complaints.
- v.** Establish a clear standard of strict institutional accountability for harms resulting from court-sanctioned AI, guaranteeing that the judiciary bears ultimate and vicarious responsibility for rectifying and compensating algorithmic errors.

VII. ON THE DIGITAL DIVIDE AND INCLUSIVITY

Clause 13 of the Draft Regulations enunciates the principle of 'Inclusivity and Accessibility' and directs that "specific attention shall be given to ensuring that the deployment of AI does not create or widen digital divides." This is among the most important principles in the instrument. Unfortunately, the Draft does not translate this principle into concrete operational directives, leaving it as an aspiration rather than an obligation.

The Indian judicial system serves a population of 1.4 billion people, the majority of whom reside in rural areas, lack reliable internet connectivity, and have limited digital literacy. The adoption of AI-driven court services like chatbots, digital filing systems, AI-assisted case management will, in the short to medium term, be more readily accessible to English-speaking, urban, technologically literate litigants and their counsel than to first-generation litigants appearing in the vernacular before a Tehsil-level Court. The risk is not hypothetical; it is structural.

It is submitted that the regulations should incorporate the following specific safeguards:

- i.** Litigant-facing AI systems, chatbots, grievance portals, case status tools must be made available in all 22 Scheduled languages, with voice-enabled interfaces for those with limited literacy. The Digital Infrastructure standards should be made a prerequisite for AI deployment, not an afterthought.
- ii.** The Annual Transparency Report under Regulation 45 should be required to include a specific section on digital equity, measuring the extent to which AI-assisted services have been accessed by rural, economically disadvantaged, and linguistically diverse litigants, so that the system can track whether AI adoption is widening or narrowing the access-to-justice gap.
- iii.** Court Legal Services Committees and District Legal Services Authorities should be formally brought within the ambit of AI adoption, so that free legal aid beneficiaries, who are among the most vulnerable users of the justice system, can benefit from AI-assisted services rather than being left behind.

VIII. ON THE DISTRICT COURT STRUCTURE | THE MISSING CENTRE

The most conspicuous structural gap in the Draft Regulations is the limited attention paid to the district court system. District and subordinate courts handle over 80% of the total litigation in India and are the primary point of contact between the justice system and the ordinary citizen. Yet the institutional architecture created by Chapter IV is almost entirely Supreme Court and High Court-centric.

The Apex Body sits at the Supreme Court. AI Committees are constituted in the Supreme Court and High Courts. AI Secretariats are headed by officers of the cadre of District Judge, but they operate at the High Court level. District courts appear in the Draft only as recipients of regulations and obligations, not as active participants in governance.

This is concerning for several reasons:

i. Implementation distance. A regulation framed at the Supreme Court level may not percolate meaningfully to a Judicial Magistrate in a Tier-3 city who handles 60 to 80 matters per day and has no dedicated technology support. The Draft requires this officer to supervise AI outputs, verify AI-generated transcriptions, report AI incidents, and comply with data protection obligations, without any corresponding investment in the support structure at that level.

ii. No district-level AI Committee or focal point. The Draft does not contemplate a District-level AI focal point, nodal officer, or committee. The district judiciary, comprising Principal District Judges and their courts, has no defined role in the AI governance architecture, despite being the system's largest operating unit.

iii. Infrastructure deficit. Many District Courts lack reliable power supply, high-speed internet, and modern hardware. Before AI systems can be deployed at these courts, there must be a concurrent plan for baseline technological infrastructure. The Committee on Infrastructure and Finance under Regulation 28 is mandated to “oversee the establishment of infrastructure related to AI” but the Draft is silent on timelines, standards, or the role of State Governments in funding and deploying this infrastructure at the district level.

It is strongly recommended that the final regulations include a dedicated chapter or schedule specifically addressing the implementation of AI in district and subordinate courts, including phased deployment timelines, minimum infrastructure prerequisites, and the designation of a District-level AI Nodal Officer in each district judiciary.

IX. OTHER SPECIFIC OBSERVATIONS

A. Prohibition on Risk Scoring (Regulation 20(1)(d))

The absolute prohibition on the use of AI systems for Risk Scoring, covering flight risk, recidivism, bail eligibility, and credibility assessments, is strongly endorsed. This prohibition is consistent with global best practices and with the constitutional right to personal liberty under Article 21. However, the Draft does not address the possibility that risk-scoring tools may already be informally in use, or may be introduced indirectly through tools ostensibly used for other purposes. The review mechanism under Regulation 41 and the AI audit mechanism under Regulation 38 should specifically mandate a check for covert or indirect risk-scoring functionality in any deployed AI system.

B. AI Content Verification Authority (Regulation 44)

The creation of an AI Content Verification Authority is a forward-looking provision. However, the Draft leaves the composition, powers, and procedures of this authority entirely to the Appropriate Authority to determine. Given the significance of this body, which will oversee the verification of GenAI-generated content in court proceedings, the draft should provide at least a skeletal framework specifying its independence, access to technical expertise, and the standards it will apply.

C. Controlled Environment Testing (Regulation 36)

The provision for Controlled Environment Testing (or ‘sandboxing’) is valuable. However, Clause 36(1) states that the Appropriate Authority *may* direct such testing, making it discretionary rather than mandatory. Given the sensitivity of court environments, it is suggested that Controlled Environment Testing be made mandatory for all AI systems intended for use in adjudicatory functions, before deployment.

D. Bridging the Capacity Gap in “In-House” Audits (Regulation 38)

Regulation 38(2) strictly restricts all algorithmic audits to the Court premises, forbidding external third parties from accessing source codes or datasets. While safeguarding data sovereignty is crucial, this inward-looking approach ignores the reality that the judiciary lacks the highly specialised technical bandwidth required to rigorously audit deep learning models. To balance security with competence, the regulations should permit “Controlled Independent Audits.” By allowing fully vetted yet independent external technologists to conduct evaluations with the secure “Controlled Environment Testing” spaces as authorised under Regulation 36, the Court can maintain sovereign control without compromising world-class scrutiny.

E. Tightening Private Sector Controls and Securing Data Independence (Regulation 46)

Chapter VI initiates boundaries for private sector involvement, and Regulation 46(9) helpfully restricts private firms from monopolising intellectual property (IP) built on court resources. However, the safeguards need strengthening. Relying heavily on external tech companies for public governance frequently weakens democratic oversight and increases the vulnerability of sensitive citizen data to commercial exploitation. We urge the Committee to make comprehensive algorithmic openness and robust financial indemnity provisions non-negotiable in vendor contracts. Commercial entities should not be allowed to use IP or trade secrecy as an excuse to hinder the inner working of their models, especially when these models directly impact legal rights of litigants or case outcomes.

F. Defining Verification Standards for Legal Professionals (Regulation 43)

Regulation 43(6) makes an important intervention by holding lawyers completely responsible for filing fabricated “hallucinations” effectively preventing them from blaming the technology for their poor submissions. Nevertheless, the draft lacks a defined threshold for what constitutes proper vetting. The mandated disclosure format (Annexure I) under Regulation 43(3) should be modified to make lawyers explicitly detail the steps they took to corroborate the AI’s findings. By forcing the legal professionals to confirm that they manually cross-checked AI citations against authoritative and traditional legal reporters, the regulations will establish an objective standard for measuring professional diligence thus moving beyond mere threatening of punitive action after the fact.

X. CONCLUSION AND SUMMARY OF RECOMMENDATIONS

The draft Regulations for Use of Artificial Intelligence in Courts, 2026 represent a serious and substantive attempt to bring regulatory coherence to the adoption of AI in India’s judicial system. The foundational principles of human primacy, judicial independence, transparency, fairness, and accountability are well-chosen and thoughtfully articulated. The absolute prohibitions in Regulation 20 reflect a principled understanding of the boundaries that must not be crossed in the application of AI to adjudication.

The suggestions offered in this submission are not directed at the vision that animates the Draft, which is wholesome and necessary. They are directed at the imperative of ensuring that the final regulations are not merely aspirational but operationally credible. The central submission is this: the adoption of AI in courts is only as good as the capacity of the human beings who use, supervise, and govern it. The regulations that accompany that adoption must, accordingly, invest as much in people and processes as they do in platforms and policies.

The following are the principal recommendations in summary:

i. On form: Examine whether certain provisions require a statutory anchor or incorporation into Rules of Court to ensure legal validity and consistent enforceability.

- ii. On institutional architecture:** Streamline the approval and governance pathway to prevent duplication and delay; ensure meaningful expert participation in the Apex Body on a regular basis.
- iii. On capacity building:** Prescribe mandatory minimum training hours, regional language modules, competency assessments, and dedicated AI literacy budgets, with a realistic implementation timeline.
- iv. On Human-in-the-Loop:** Address automation bias and confirmation bias in training curricula; require substantive documentation of HITL verification; audit the quality, not merely the fact of human review.
- v. On AI Incident Database:** Mandate public accessibility of an anonymised Incident Database at regular intervals, as a matter of institutional transparency and public trust.
- vi. On grievance redressal:** Expand the scope of grievable harms, provide an independent forum, prescribe time-bound resolution, and designate AI Grievance Officers at the district level.
- vii. On digital inclusion:** Mandate parallel manual processes; require multilingual voice-enabled interfaces; include digital equity metrics in Annual Transparency Reports; integrate legal aid bodies into AI adoption planning.
- viii. On district courts:** Develop a dedicated implementation framework for district and subordinate courts, including minimum infrastructure prerequisites, phased timelines, and district-level nodal officers.
- ix. On Controlled Environment Testing:** Make sandboxing mandatory, not merely discretionary, for AI systems intended for adjudicatory use.
- x. On Liability and Redressal:** Elevate grievance mechanisms to a standard of absolute and vicarious liability for state-deployed AI, removing the burden of proving algorithmic faulty from the litigant.
- xi. On Oversight and Private Engagement:** Permit “Controlled Independent Audits” by vetted external technologists within secure environments to bridge internal capacity deficits and enforce strict algorithmic transparency to prevent private vendors from hiding behind trade secret protections.
- xii. On Professional Standards:** Amend disclosure formats to require legal professionals to certify their specific methodology for verifying AI-generated outputs against primary legal databases, establishing an objective standard of care.
- xiii. On User Interface and Verification:** Mandate interface controls that block one-click approvals and require active, documented human deliberation to actively combat automation bias.

These comments are submitted with the utmost respect for the institution and the initiative, and in the hope that the final regulations will reflect the aspirations of both the institution that drafts them and the hundreds of millions of citizens who will ultimately be governed by them.

Respectfully submitted,



Jameela Sahiba,
Associate Director, AI and Public Affairs,
[The Dialogue](#)



thedialogue.org.in



LinkedIn | The Dialogue



X | The Dialogue



Whatsapp | The Dialogue



Instagram | The Dialogue