Comments and Suggestions on the Draft Regulations for Use of Artificial Intelligence (AI) in Courts, 2026

To: The Member Secretary, Artificial Intelligence Committee, Supreme Court of India (office.regec@sci.nic.in).

Submitted by:

Study Group on Regulations for Use of AI in Courts, Law and Technology Research Group, Centre for the Study of Law and Governance, Jawaharlal Nehru University

Members of the Study Group: Prabhat Katyayan Mishra (lead author), Dr. Mandira Narain (lead author), Professor P. Puneeth, Dr. Nupur Chowdhury, Dishari Roy, Professor Amit Prakash

Email: lawandtech.cslg@gmail.com

We have divided this document into two parts. Part I provides general comments. Each of the general comments is numbered and includes recommendations. Part II provides provision wise comments.

Part I: General Comments

1. Overview

We welcome the public consultations initiated on the Draft Regulations for Use of Artificial Intelligence (AI) in Courts, 2026 (hereafter, Regulations).

The general comments that follow concern: the source of authority and reach of the Regulations; the missing annexures; the absence of a stated purpose and evaluation of AI systems already deployed; the absence of a judicial data policy on which any AI policy must rest; the tension between the instrument's precautionary and promotional commitments; the degree of centralisation, combined with institutional self-regulation; and pervasive drafting defects.

2. The source of authority and the reach of the instrument

The instrument styles itself “Regulations” and purports to apply to the Supreme Court, the High Courts, and all courts, tribunals and statutory commissions performing adjudicatory functions (Regulation 2). It nowhere identifies the constitutional or statutory power under which it is made. Article 145 of the Constitution of India authorises the Supreme Court, with the approval of the President and subject to law made by Parliament, to regulate its own practice and procedure. The High Courts derive independent authority over their own practice and procedure, and over the courts subordinate to them, from Articles 225, 227 and 235. Tribunals and statutory commissions owe their procedure to their parent statutes, and generally are subject to executive/administrative control of the government and the supervisory control of their respective High Courts under Article 227 of the Constitution. There is no evident legal basis on which a body constituted within the Supreme Court may prescribe “minimum mandatory standards” for “compliance by all Courts” (Regulation 24(1)(a)).

Commencement of the Guidelines under Regulation 1(2) relies on separate notification by the Chief Justice of each High Court while Regulations 24 (Powers of Apex Body) and 33 (AI Committees) simultaneously assume a hierarchy of binding national standards. The only domestic precedent, the Kerala High Court’s Policy of July 2025, confines itself to the district judiciary that the High Court superintends and controls respectively under Article 227 and Article 235.[1]

Recommendation. Regulation 1(2) as drafted permits selective and piecemeal enforcement that would undermine the integrity of the framework. State the source of authority expressly. If a binding instrument is intended, frame rules under Article 145 for the Supreme Court, accompanied by a model framework for adoption by each High Court under Articles 225, 227 and 235 and by tribunals under their own rule-making powers.[2]

---

[1] *Policy Regarding Use of Artificial Intelligence Tools in District Judiciary*, High Court of Kerala, Official Memorandum dated 19.07.2025, clause 2.

[2] See, as precedent, *Model Rules for Live-Streaming and Recording of Court Proceedings*, e-Committee, Supreme Court of India, available at https://ecommitteesci.gov.in/document/model-rules-for-live-streaming-and-recording-of-court-proceedings/

3. The missing annexures

Annexures I and II, on which the disclosure obligations in Regulation 43(3) depend, have not been published, and their content cannot be assessed.

Recommendation. Publish the complete draft with the missing annexures and allow a further round of comment on them.

4. The absent statement of purpose, and the unexamined record of systems already deployed

The draft contains no statement of objects and reasons. It is consequently unclear what specific problem or challenge the Regulations address. Reducing pendency, improving administrative efficiency, expanding access to justice, facilitating AI adoption, or something else entirely could be the intended object. In the present form, the Regulations are highly ambiguous. "Efficiency" is not a specific objective, and no rational nexus is demonstrated between the uses the draft permits and any intended outcome. Without clear and specific stated objectives, it is not possible to assess whether AI is an apt solution, whether the safeguards are proportionate, or what trade-offs, including fiscal and environmental costs, are acceptable.

Several provisions also cast the court as a developer and supplier of technological services, a posture removed from its constitutional function of authoritative interpretation and enforcement of laws of the land. A statement of purpose would also have to explicitly address this constitutional question of the court's role.

There is also the absence of any public evaluation of the judiciary's existing deployments, such as SUVAS, SUPACE, Saransh, and other third-party or co-developed tools already in use in the Supreme Court and across several High Courts. There is no published assessment of whether these tools, where significant public investment has been made, achieved their objectives, at what error rates, or with what incidents. Regulation 41 contemplates only a compliance review of existing systems, not an evaluation of outcomes.

Recommendation. Preface the Regulations with a clear statement of objects and reasons identifying the problems addressed, the outcomes sought, and the evidence relied upon. Quantify the link asserted between AI use and pendency reduction, if that is the objective sought to be achieved. Commission and publish an independent evaluation of all AI systems already deployed in courts before, or as a condition of, their continued use under Regulation 41.

5. A judicial data policy must precede an AI policy

Most of the hard questions the draft raises are questions about data. Judicial data are aggregates of personal and non-personal data. The judiciary currently has no comprehensive, public framework governing their collection, management, retention, sharing, reuse and protection. The draft proceeds as if custodial possession of court data resolves the question of permissible use. The use of litigants' data to train or refine AI systems is distinct from adjudication and the litigants' and witnesses' interests in such data are nowhere addressed. The recommendations of the Committee of Experts on Non-Personal Data Governance merit consideration in framing such a policy.[3]

Recommendation. Adopt, before or together with these Regulations, a comprehensive judicial data governance framework addressing ownership, access, retention, secondary use, anonymisation, privacy protection and accountability. Until such a framework is in force, the use of judicial data (whether or not anonymised) to train, fine-tune or refine any externally developed or vendor-operated model should be prohibited rather than made subject to approval under Regulation 20(1)(a). The extent of AI use permitted should further be pegged to the quality and completeness of the underlying data. Much judicial data remains partially digitised, fragmented and unrepresentative, and systems trained on it will inherit those defects.

6. Precaution, not promotion

Chapter II commits the framework to a precautionary posture (heightened scrutiny of opaque systems, safeguards rising with risk, mandatory human oversight). However, Regulations 16 and 17 state a "presumption … in favour of responsible adoption" and a general principle of "Innovation over Restraint" under which every court "shall actively seek opportunities to deploy" AI, with reasons in writing required for restriction but not for adoption. A normative presumption in favour of deployment, embedded among the governing principles of an adjudicatory instrument, inverts the burden of justification that ought to attach to introducing imperfectly understood systems into the exercise of judicial power, and shifts onto affected persons the burden of demonstrating harm.

By comparison, the "under user control" principle of the Council of Europe CEPEJ's *European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and their Environment*

[3] Committee of Experts on Non-Personal Data Governance Framework (Gopalakrishnan Committee), Report (2020), on the governance of aggregates of personal and non-personal data.

(2018) expressly precludes a prescriptive approach and requires that users of such systems remain informed actors in control of their choices.[4]

There is also an internal incoherence. Regulation 16 announces the presumption of adoption, yet its proviso lodges two substantive prohibitions: that AI shall not replace human decision-making, and that it shall not be deployed for dispute-outcome prediction. The first substantially duplicates Regulation 20(1)(b) and (c). The second appears nowhere in Regulation 20. Consequently, the bar on outcome prediction, though evidently intended to be absolute, enjoys no express protection against relaxation: Regulation 56 exempts from the Chief Justice's power of relaxation only the prohibitions specified in Regulation 20.

Recommendation. Adopt a precautionary approach in place of the presumption, with an express no-harm principle governing generative AI, in particular. Recast Regulations 16 and 17 so that encouragement of responsible modernisation operates as a policy orientation addressed to the institutional bodies, subordinate to the general principles, and not as an interpretive presumption bearing on deployment or adjudication. Relocate the proviso's prohibitions to Regulation 20, where they would fall under the non-derogability clause.

7. Centralisation, and regulators who are also developers

The draft contemplates a scale of centralisation for which we are aware of no comparator among judicial AI frameworks. An Apex Body at the Supreme Court setting binding national standards, approving systems, overseeing research, and reviewing its own regulations, its members nominated substantially by the Chief Justice of India without stated tenure protections or expertise requirements (Regulation 22(3) leaves terms of appointment and removal to the Chief Justice).

The architecture also conflates development with regulation. The Judicial Committee is to "innovate, develop and deploy" AI tools to assist adjudication (Regulation 26(2)), as is the Case and Data Management Committee (Regulation 29(2)), while the same institutional hierarchy approves, audits and reviews those tools. A body that builds the systems it regulates cannot supply independent oversight, and no provision subjects the Apex Body's own decisions to external audit or review.

---

[4]European Commission for the Efficiency of Justice (CEPEJ), *European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and their Environment* (2018), fifth principle ("under user control").

Recommendation. Separate the development function from the regulatory function institutionally. Recast the Apex Body's relationship with the High Courts as collaborative standard-setting. Prescribe member tenure, qualifications and conflict-of-interest rules in the instrument itself rather than leaving them to the appointing authority and provide for periodic independent review of the framework's operation.

8. Drafting quality

The draft deploys operative terms that remain undefined in the text. Some illustrative examples being "judicial function" (Regulation 2), "quasi-adjudicatory" function (Regulation 3(1)(q)), "high-risk" (Regulations 46 4(h) and7(3)), and the "emergency" that triggers Regulation 42.

Several definitions go beyond defining: the prohibition on filing undisclosed GenAI content sits inside the definition of GenAI (Regulation 3(1)(y)). Same for explainability standard inside Regulation 3(1)(w). Regulation 23(d) has the Apex Body striving "to realise the vision of the AI Committee", inverting the intended hierarchy.

Recommendation. Define every operative term once in Regulation 3 and use it consistently. Relocate operative content from the definitions to the substantive chapters.

Part II: Provision-wise Comments (selected)

1. **Regulation 3(1)(e), AI incident:** The definition of an 'AI incident', should be expanded to include legal terms like 'harm', 'loss', 'damage' and 'fault', in place of, or in addition to, 'malfunction' and 'erroneous output'. This is to comprehensively capture the range of technical failures an AI System may exhibit, rather than relying on an enumerated list of type of failure that may not be exhaustive.
2. **Regulation 3(1)(k), anonymisation:** The draft relies on anonymisation spelled in Regulation 3(1)(k), as a key safeguard for privacy and confidentiality. This complete reliance should be treated with caution since an existing body of technical literature shows that anonymised data can be re-identified through de-anonymisation techniques. We recommend that the regulations acknowledge this limitation and provide safeguards against re-identification, including legal consequences where de-

anonymisation occurs, instead of treating anonymisation as a sufficient privacy safeguard on its own.

3. **Regulation 3(1)(m) defines Artificial Intelligence:** The definition excludes ‘general-purpose software or digital tools,’ unless such software is specifically embedded with, augmented by, or functionally dependent upon artificial intelligence. This exclusion should be removed as nearly all digital tools now incorporate AI-enabled functionality. Therefore, making a clear line between ‘general-purpose software’ and an ‘AI System’ difficult to sustain in operation.
4. **Regulation 2, Judicial, adjudicatory, or administrative functions:** The regulation treats these as three distinct categories, which is a drafting error, as ‘judicial function’ ordinarily encompasses both adjudicatory and administrative functions and should not be listed as a separate third category.
5. **Regulation 3(1)(w), Explainability:** Regulation 3(1)(w) provides that an AI System shall be considered explainable, if such account may be provided upon request without requiring specialist technical knowledge. The definition of ‘explainable’ should not make the provision contingent upon a request. This Regulation 3(1)(w) should be revised as, if such account may be provided upon request and is understandable without requiring specialist technical knowledge. This will establish a default explainability requirement for all high-risk functions, be it administrative or adjudicatory.
6. **Regulation 3(1)(lb), Human-in-the-Loop as a sufficient safeguard:** Human-in-the-Loop (HITL), used in the operative provisions should not be treated as sufficient, as it will lead to default acceptance of AI-generated outputs rather than subjecting them to genuine independent scrutiny. In such circumstances, the human decision-maker’s role may serve to primarily legitimise the automated AI output rather than to provide meaningful oversight.
7. **Regulation 3(1)(zh), Definition of sensitive personal data or information:** Regulation 3(1)(zh) defines ‘sensitive judicial data’ by reference to ‘personal identifiable information of parties, witnesses, or legal representatives.’ This is a problem of terminological congruence as Section 43A of the Information Technology Act, 2000 already uses the established term ‘sensitive personal data or information.’

The draft should use this same phrase rather than introducing a different formulation for the same concept.

8. **Regulation 4, Human primacy and judicial independence:** Regulation 4, states that, the use of AI remain strictly subservient to human judgment and judicial authority. To give practical effect to the principles of human primacy, judicial independence, transparency and explainability, the regulation should expressly address scenarios involving the use of AI, rather than a single undifferentiated principle of 'subservience to human judgment.' For instance, in the primary scenario where an AI system in used, litigants should be provided with mandatory disclosure informing them of the extent of use. This disclosure should contain information such as, that AI was consulted, the nature of its output, and whether the judicial officer followed or departed from that output. Within this scenario, a further distinction must be drawn to preserve human primacy, as a judge should not be required to provide reasons solely because they departed from an AI recommendation. The validity of a judgment must continue to rest on its own merits and legal reasoning, irrespective of whether an AI system was used to assist the preparation.
9. **Regulations 18, 19 and 20, Permissible and prohibited uses not organised by the administrative/adjudicatory distinction:** Regulation 19(1) lists permissible uses without classifying them by the adjudicatory and administrative functions. Several listed uses include, citation verification, document summarisation, and aspects of legal research under Regulation 19(1)(d) which are adjudicatory in character. However, these stay undifferentiated alongside genuinely administrative uses, like case filing assistance and record management under Regulation 19(1)(e). Regulation 20 (prohibited uses) exhibits the same problem, where a prior classification, is unclear as to what standard of scrutiny, disclosure, or human oversight attaches to a given use.
10. **Regulation 19(1), Human oversight not universal across permissible uses:** Human oversight is required for some permissible uses. Regulation 19(1)(f), for instance, conditions conversational AI assistants on 'human oversight of their functioning' but is absent for others, including Regulation 19(1)(d), an adjudicatory function. There is no principled basis for confining oversight to only public-facing functions. It should extend and uniformly apply to every permissible use. Second, there is an absence of a

time limitation on human oversight. Even where oversight is required the regulation 19 prescribes no defined standard nor any time limit within which it must be exercised.

11. **Regulation 19(1)(j), Court performance assessment not confined to institutional use:** AI should not be used to assess individual judges or courts, which remains a human function. An undefined 'performance assessment,' term risks legitimising informal or opaque AI-based judge-profiling or evaluation. This should be confined to institutional use, such as pendency, while prohibiting an AI-based assessment of individual judges or courts.
12. **Regulation 20(1)(e), undisclosed, opaque, or unexplainable AI Systems:** While, the terms undisclosed and opaque AI system is welcomed, the regulation does not define what qualifies as 'explainable' for this purpose. The regulations should adopt a recognised international technical standard for assessing explainability before its use is authorised, with the person or body conducting such assessment.
13. **Regulation 20(1)(g), surveillance/continuous monitoring exception:** The exception permitting continuous monitoring of judicial officers, advocates, and litigants 'if authorised by law' should be deleted, on two grounds. First, any form of, indiscriminate surveillance breaches the fundamental right to privacy. Second, such monitoring fails to protect privilege communications under evidentiary law. This inevitably contradicts the draft regulations' own commitment privacy under Regulation 10(2). In addition, it may also affect independence of subordinate judiciary.
14. **Regulation 21, Reporting and remedial mechanism:** Regulation 21 does not designate who bears the reporting obligation. It also confines remedial measures to suspension of the AI System, with no provision for action against the individual who committed the violation.
15. **Regulation 22(1), The Apex Body's mandate to promote innovation:** A body constituted to regulate AI use in the judiciary should not be mandated to promote the technology it is meant to oversee. Its role should be confined to regulation aimed at protecting fundamental rights and ensure access to justice. Combining the two roles, creates a conflict of interest and risks judicial independence.
16. **Regulation 24, Subordination of AI Committees to the Apex Body:** Every AI Committee constituted under Regulation 33 is then required to comply with. This

regulation subordinates each High Court's AI Committee to a body whose own composition, rests entirely on nomination by the Chief Justice of India. This should be recasted and not be a binding subordination.

17. **Regulation 30(1), the Cyber Security Committee:** This regulation specifies, that Cyber Security Committee should include serving officers of CERT-In, MeitY and NIC *ex officio*. Taking such officers who are continuing representatives of ministries fails to preserve an institutional separation between the executive and the judiciary. These members should be inducted only through deputation, in which an officer is formally released from their cadre for the duration of their service

18. **Regulation 32, Centre of Research and Excellence on Artificial Intelligence (CoRe):** AI in the judiciary is not a self-contained technical subject. It intersects with disciplines such as political science, economics, sociology and others that have a bearing on law and governance. It should be ensured, that the research function remains institutionally independent of the judiciary and genuinely multidisciplinary in collaboration with independent research chairs or centres across universities in India.

19. **Regulations 34(4)(c), Functional similarity to an already approved by the Appropriate Authority:** The 'functionally similar' standard presumes a stable AI tool, as it resembles an already approved one. This disregards the nature of agentic and generative AI, which evolve simultaneously through continuous learning. This regulation assumes a static nature of the tool and should be revised to not ground 'expedited approval' by the same, basis the functionality presumption.

20. **Regulation 35(3)(c), Risks assessed under the Technical and Ethical Impact Assessment:** Risks of bias, error, hallucination and misuse in the relevant judicial context, omits reference to the resulting loss, harm, or damage which may be caused as consequences they may produce for persons or processes.

21. **Regulation 36(3), Controlled Environment Testing:** There is an internal contradiction in Regulation 36, which assures that the real-life testing will have no effect on court processes, whereas the purpose of real-life sandbox testing is to observe how a tool performs, under actual operating conditions, including, whether it causes errors or disruption. Therefore, if live testing is permitted, the regulation should explicitly acknowledge that some form of residual risk is unavoidable and provide a

clear accountability and liability framework for any harm caused during the testing period.

22. **Regulation 38(2), In-house and within Court premise audit:** Regulation 38(2) requires audits to be conducted ‘in-house’, and provides that source code, algorithms, or datasets or other architectural information should not be shared with any third party for audit outside Court premises. This rests on a spatial and physical conception of audit that marks a clear disjuncture from how actual AI systems are typically architected. Most AI systems rely on cloud infrastructure or third-party hosting which do not physically reside within Court premises. Therefore, internal audit is insufficient to ensure independent scrutiny. We recommend that AI Systems used by Courts be hosted on infrastructure under sovereign control, be it on premises or through a sovereign cloud arrangement; second, that external independent audit be expressly permitted that such external access be governed by conflict-of-interest principles and binding data confidentiality agreements.