

GNLU CENTRE FOR LAW & ECONOMICS

Policy Recommendation



GNLU/CLE/PR-30

July 15, 2026

Policy Recommendations to the Supreme Court on the Draft ‘Regulations for Use of Artificial Intelligence (AI) in Courts, 2026’

Comments on behalf of the Policy Inputs Research Group, GNLU Centre for Law & Economics

CENTRE FACULTY

Prof. (Dr.) Ranita Nagar

Professor of Economics

Head of GNLU Centre for Law &

Economics

Dr. Hiteshkumar Thakkar

Assistant Professor of Economics

Co-Convenor, GNLU Centre for Law &

Economics

CONTRIBUTORS

ANUSHREE SAXENA (TEAM LEAD)

AYUSH YADAV

ADITYA PRATAP SINGH

DAKSH KUMAR BAFNA

DISHA JOSHI

ACKNOWLEDGEMENTS

MS. KRISHNA AGARWAL

Krishna Agarwal holds an LL.M. from the University of Cambridge, with a specialization in the intersection of law and technology, lending the research group a valuable global and comparative perspective.

MR. ABHIVARDHAN

Mr. Abhivardhan is an AI Governance Specialist and the Founder-Chairperson of the Indic Pacific Legal Research and Education Organization. He also leads the Indian Society of Artificial Intelligence and Law, where his work focuses on AI governance, emerging technologies, digital policy, and legal innovation.

TABLE OF CONTENTS

_Toc235801229

I. INTRODUCTION	3
II. GENERAL COMMENTS	3
III. SPECIFIC COMMENTS.....	6
IV. CONCLUSION.....	50

I. INTRODUCTION

The accelerated development of Artificial Intelligence (“AI”) has opened up new possibilities for enhancing efficiency, accessibility and administrative functioning in various public bodies, including the judiciary. Given the game changing power of AI and the distinctiveness of court functioning, the aim of the “*Draft Regulations for Use of Artificial Intelligence in Courts, 2026*” (“**Draft Regulation**”) is to set out an all-encompassing regulatory regime for the development, deployment and utilisation of AI in Indian courts.

The Draft Regulations have been developed by the Committee on Artificial Intelligence constituted by the Supreme Court of India to assess both the opportunities and risks associated with introducing AI into the judicial system. In formulating the Draft Regulations, the Committee conducted a thorough analysis of current technological developments, judicial needs, the ethical use of AI and issues pertaining to the use of AI in courts. The Draft Regulations have been made available to the public for comments to allow stakeholder participation in the process of shaping regulation of AI in courts.

This policy input provides comments and recommendations on various provisions of the Draft Regulations with a view to strengthening regulatory effectiveness, promoting responsible innovation and ensuring the trustworthy and efficient use of AI within the judiciary.

II. GENERAL COMMENTS

The Draft Regulations may be considered an important and timely step towards formalising the usage of Artificial Intelligence in the judiciary of India. The general idea of merging efficiency, accessibility and innovation with constitutional requirements is both justified and well formulated,

especially given some institutional problems like backlog, resource scarcity and unequal access to justice. Conceptually, the Draft Regulations display a clear prioritisation of human-centeredness, procedural correctness and responsible innovation. Nevertheless, a range of issues keeps on coming up in various articles of the Draft Regulations.

A common problem in the Draft Regulations is the discrepancy between the uniformity and institutional diversity principles. Some articles provide for the same requirements for all courts, tribunals and administrative authorities despite the difference in their competences, capabilities and risk levels. The issue results in unnecessary inefficiencies and possible problems with either over regulation of low impact cases or vice versa. The Draft Regulations could benefit from adopting a risk-based and proportionate approach.

Another recurrent issue has to do with ambiguous definitions of key concepts and their further repercussions. Vague or technologically rigid definitions of key terms such as “Artificial Intelligence,” “Black Box,” and “risk scoring” risk generating compliance uncertainty, regulatory arbitrage and inconsistent enforcement. Considering the rapid evolution that characterizes most AI systems, it is imperative to have an exact definition of these systems that takes into account their constantly evolving characteristics. More precision in the definition of AI will reduce interpretive differences and help to create clearer regulations.

Another important theme of the Draft relates to the question of responsibility in the assisted decision-making through AI. The stress on judicial accountability is well-taken, but the tendency to place almost all of the responsibilities on the shoulders of judicial officers ignores the existence of private actors who are developing AI applications. Such a focus creates an imbalance because people who are better positioned to address technical issues are left with no motivation to act responsibly. Thus, a more rational balance needs to be established.

A third important aspect in the regulation of AI use in courts is connected with the problems of transparency and explainability. On one hand, it is essential to ensure the possibility for a party to be informed about the use of artificial intelligence in a court process. However, in order to avoid excessive costs and ensure the effectiveness of the regulation, the difference should be made between material and non-material uses of AI.

A close connection can be drawn between this and the related concept of information asymmetry, which has implications for several provisions. The courts, the litigating parties and even the regulator may lack the technical competence to properly understand AI systems, whereas vendors maintain a considerable advantage in terms of information. This is evident throughout different aspects, from detecting bias and conducting audits to redressing grievances and procurement.

The Draft also raises the problem of balancing innovation with regulation. Though the introduction of the innovation principle is a positive development, the connection between it and other foundational principles such as transparency, accountability and proportionality is left unclear. It is thus uncertain whether the innovation principle should take precedence over the others. An order of norms can serve to resolve this tension.

Issues pertaining to institutional design have also been noted. For example, several regulatory bodies such as the Apex Body, AI Committees and the AI Secretariat lack adequate participation by technical experts, civil society and affected stakeholders. On the contrary, some bodies are tasked with both promotion and regulation of AI systems, which poses potential conflicts of interest and regulatory capture issues. AI regulation needs to combine the twin concepts of functional separation and multiple stakeholders in order to achieve credibility, independence and accountability.

Lastly, the Draft brings up many important issues related to access to justice and distributive fairness. Excessive compliance burdens, especially as far as disclosures and procedures go, pose risks of being heavily felt by small firms, unrepresented persons and pro se litigants. In addition, lack of proper controls for chatbots, language accessibility and remedy provisions will only serve to perpetuate current inequalities. It is imperative that AI becomes a vehicle for inclusion and not exclusion.

To conclude, while the Draft Regulations provide an excellent normative framework for the use of AI in courts, the success of its implementation will critically hinge on dealing with these fundamental problems. More risk oriented regulation, clarity, proper accountability, institutional reforms and emphasis on distributive justice will greatly strengthen the regime in place.

III. SPECIFIC COMMENTS

SR. NO.	SUMMARY OF PROPOSAL	COMMENTS/ SUGGESTION	RATIONALE
1.	Chapter 1		
	Regulation 2- Uniform scope and institutional heterogeneity - Classifying courts and tribunals into different categories based on risk and compliance obligations will minimize regulatory deadweight loss and increase efficiency.	Regulation 2- Uniform scope and institutional heterogeneity - It is suggested that Courts and Tribunals be classified into different risk categories. Compliance obligations will then depend on the adjudicatory function, the stakes involved and the nature of the proceedings. The supervisory authority may then prescribe separate standards depending on the classification. This would minimise regulatory deadweight loss and ensure proportional regulation. For example, the NCLT may permit AI-assisted legal research and document review under lighter compliance, while criminal trial courts require stricter oversight and	Regulation 2- Uniform scope and institutional heterogeneity - Regulation 2 mandates the uniform application of all rules across the Supreme Court, High Court, Subordinate Courts, tribunals, and statutory commissions. This ignores institutional heterogeneity, which, from a transaction-costs perspective, generates regulatory deadweight loss due to the fact that it is identical compliance obligations for bodies with different functions, capacities and procedural requirements, it will lead to higher administration costs, to lower efficiency, to avoidable frictions that will exceed the benefits of the

		mandatory human verification.	standardization. ¹ A specialised commercial tribunal, such as GST Appellate Tribunal (“GSTAT”) or National Company Law Tribunal (“NCLT”), usually operates with better information symmetry and high delay costs. Still, a criminal trial court, such as the Court of Sessions or the Judicial Magistrate First Class, operates with high information asymmetry and high error costs. ² Introducing the same compliance, where the use of AI will differ depending on the type of institution, may lead to administrative inefficiency and keep judicial transaction costs high. Regulatory frameworks must account for heterogeneity rather than applying uniform rules. ³
--	--	-------------------------------	--

¹Axel Michaelowa & Frank Jotzo, Transaction Costs, Institutional Rigidities and the Size of the Clean Development Mechanism, 33 Energy Pol’y 511 (2005).

² Keith N. Hylton, Information Costs and the Civil Justice System, 24 Am. L. & Econ. Rev. 407 (2022).

³ Thomas Eger & Marc Scheufen, The Law and Economics of the Data Economy: Introduction to the Special Issue, 57 Eur. J.L. & Econ. 93 (2024).

	Regulation 3(m) — Definition of “Artificial Intelligence” Differentiating between general-purpose software and AI can be difficult with the current definition; a non-exhaustive, illustrative list should be added.	Regulation 3(m) — Definition of “Artificial Intelligence” Adding a non-exhaustive and illustrative list of software categories that are within or outside the definition, which can be modelled on the European Union (“EU”) AI Act’s Annex I/III typology.	Regulation 3(m) — Definition of “Artificial Intelligence” The definition excludes “general-purpose software” unless it is “specifically embedded with, augmented by, or functionally dependent upon” AI. This exception is vague and can be interpreted in many ways. Modern software exists on a spectrum; a spreadsheet can use predictive fill, software can be added to correct grammatical errors, and search engines can rank results, all of which could satisfy the functional dependency test. This can increase regulatory uncertainty and increase compliance costs. Definitional vagueness in technology-facing legislation creates compliance cost uncertainty and regulatory arbitrage.⁴ Definitional vagueness has
--	---	--	---

⁴ Thomas C. B. Arnold, Transaction Costs and Institutional Rigidities in Public Administration: A Behavioral Economics Perspective, 50 Res. Pol’y 104290 (2021).

			caused enforcement problems in European Jurisdictions. ⁵
	Regulation 3(n) — Definition of “Black Box” The definition focuses on the technique adopted rather than the output; this should be changed and be revised to focus on the output.	Regulation 3(n) — Definition of “Black Box” The definition should focus on the opacity of the AI system rather than the specific technology employed. The suggested definition is : ““Black Box” means an AI system the reasoning of which is not reasonably transparent, intelligible, or explainable to the user or affected person, having regard to its design, functioning, and output.”	The use of the term “Black Box” is unnecessarily coupled with AI systems that “use deep learning techniques. The approach is technologically limited because it does not consider the system by the means of its operation, but rather by the methodology applied. Not only deep learning models, but opaque decision-making can occur across various architectures of the present and future artificial intelligence. The provision is underinclusive if AI technologies change, since it is based on a specific technological technique. A more future-proof definition should be given based on the level of explainability, transparency and auditability of a black box system,

⁵ Wolfgang Kerber, Digital Markets, Data, and Privacy: Competition Law, Consumer Law and Data Protection, 11 J. Intell. Prop. L. & Prac. 856 (2016).

			regardless of what technology it is based on to produce its output. The EU AI Act's repeatedly revised definition of an "AI system" illustrates how definitional vagueness can create legal uncertainty and inconsistent enforcement across Member States. ⁶
2.	Chapter II		
	Regulations 4 & 8 Prescribed performance standards have a beneficial impact on the quality of the AI systems before they are implemented, and they ensure judicial responsibility. This will help lower regulatory risk, bolster safe use of AI, and prevent regulatory liability that could stifle innovation or market entry.	Regulations 4 & 8 The performance standards must be set by the supervisory authority, and varied based on the level of judicial risk. Require transparency, independent evaluation, and periodic review of judicial AI systems, instead of rigid performance thresholds lacking accepted standards.	While Regulation 8 rightly places ultimate responsibility on judicial officers in line with the principle of human primacy, AI providers should nevertheless be subject to baseline quality assurance requirements before their systems are deployed in courts. Given the absence of universally accepted and objective benchmarks for assessing AI performance across all use cases, the Regulations should require vendors to disclose testing

⁶European Parliament Adopts Its Version of AI Act, Arnold & Porter (July 3, 2023), accessible at: <https://www.arnoldporter.com/en/perspectives/advisories/2023/07/european-parliament-adopts-its-version-of-ai-act>.

			methodologies, known limitations, risk assessments, and performance evaluations relevant to the intended judicial function. Independent assessments and periodic reviews may be mandated to verify that the system continues to operate consistently, reliably, and in accordance with the objectives of the Regulations. Such an approach promotes transparency, accountability, and continuous improvement without relying on rigid or presently unavailable industry-wide standards.
	Regulation 6 — Fairness and non-discrimination The negative externality of bias should be internalised by imposing a cost on vendors for providing a Bias Impact Statement, and the definition should be clearer and	Regulation 6 — Fairness and non-discrimination All AI vendors should be mandated to provide a pre-deployment Bias Impact Statement (analogous to an environmental impact assessment). This statement must disclose the composition of the training data, any	Regulation 6 — Fairness and non-discrimination 1. Regulation 6(2) prohibits AI systems that “perpetuate, amplify, or introduce bias.” Regulation 14(2) prohibits systems trained on “demonstrably biased

	mathematically consistent.	demographic disparities, and the mitigation steps. Moreover, the AI committee should publish context-specific, mathematically consistent fairness metrics to resolve definitional ambiguity.	datasets.” These provisions treat biases occurring in AI as a condition to be excluded rather than an externality to be included. In Pigouvian Economics, a negative externality exists when the producer’s private cost is less than the social cost.⁷ The vendor trains AI on biased historical judicial data, and the “demonstrably biased” standard shifts the burden of proof onto litigants. The parties with the least information about the training data now bear the burden of proving facts that only the vendor knows.⁸ 2. Moreover, the regulation fails to account for the mathematical impossibility of
--	-----------------------------------	--	--

⁷ Carl J. Dahlman, The Problem of Externality, 22 J.L. & Econ. 141, 148 (1979).

⁸ Nicolas Petit & Thibault Schrepel, Complexity and Antitrust, 19 J. Compet. L. & Econ. 315 (2023).

			satisfying all common statistical definitions of fairness (predictive rate parity, demographic parity, equalised odds). ⁹ Without specifying which fairness metric applies to each deployment context, the prohibition creates legal uncertainty and invites litigation over any AI tool.
	Regulation 7 Require two-tier explainability, permit opaque systems only where demonstrably superior, and mandate rigorous outcome audits.	Regulation 7 Regulation 7 should implement a multi-level transparency mechanism similar to Regulation (EU) 2024/1689 (EU AI Act). As required in Article 11 and Annex IV, the AI providers have to keep detailed technical documentation, such as design characteristics, performance measures, testing procedures, limitations, and risk	Regulation 7 Regulation 7 requires “high standards of transparency and explainability” and subjects opaque Black Box systems to “heightened scrutiny.” Two distinct economic problems arise. 1. First, information asymmetry (Arrow, 1963):¹⁰ the vendor and developer have far superior

⁹ Bo Cowgill & Catherine E. Tucker, Algorithmic Fairness and Economics, Columbia Bus. Sch. Rsch. Paper No. 19-50 (2020).

¹⁰ Kenneth J. Arrow, Uncertainty and the Welfare Economics of Medical Care, 53 Am. Econ. Rev. 941, 946–52 (1963).

		assessments, including the documentation available to the Technical Committee and independent auditors. As per Article 13, judicial officers will be provided with clear, intelligible, and case-specific explanations of outputs generated by AI where it is necessary that they are meaningfully monitored by humans. Second, there is a need to account for the requirements of Article 12, 14, 15 and 72 in that they should also be addressed by mandating automated logging, effective human supervision, verified standards of accuracy and robustness, and post-deployment monitoring and periodic outcome auditing – but not only because of limited explainability.	information about a system’s decision-making process than the court, officer, or litigant. Formal disclosure requirements do not resolve information asymmetry unless the information is comprehensible to the recipient.¹¹ A 200-page technical audit report submitted to a judicial officer is not “explainability” in any economically meaningful sense. Posner and Saran’s 2025 empirical study found that LLM outputs systematically differ from experienced judges’ reasoning in ways difficult to detect without domain
--	--	--	--

¹¹ Philippe van Basseghien, Markets, Market Algorithms, and Algorithmic Bias, 30 J. Econ. Methodology 310 (2023).

			expertise.¹²This affirmed by Posner and Saran's findings because it shows that even if the outputs of AI feel like they make sense on the surface, there could still be important differences between their outputs and the decisions made, which would be impossible to detect without proper disclosure, unless the explanation is intelligible and actionable to the decision maker. 2. The most accurate models are deep neural networks, but they are the least explainable.¹³ The most explainable models are rule-
--	--	--	---

¹² Eric A. Posner & Shivam Saran, Judge AI: A Case-Study of Large Language Models as Judges, 3 J. L. & Empirical Analysis 179 (2026).

¹³ Alfredo Canziani, Adam Paszke & Eugenio Culurciello, An Analysis of Deep Neural Network Models for Practical Applications, arXiv:1605.07678 (2016).

			based systems, which have the lowest predictive accuracy.¹⁴ Therefore, there is a trade-off between explainability and model accuracy. This regulation may lead the judiciary to adopt a less accurate model, thereby increasing the rate of judicial errors. This is Pareto-suboptimal, where the social utility of a highly accurate, ex-post audited model often exceeds that of a highly explainable but inaccurate system.¹⁵
	Regulations 16 & 17 — Innovation Principle Regulations 16 and 17 establish a presumption in	Regulations 16 & 17 — Innovation Principle Steps should be taken to ensure that Regulations 7, 8,	Regulations 16 & 17 — Innovation Principle Regulations 16 and 17 introduce a “Presumption in

¹⁴Jian-Bo Yang, Si Hu, Dong-Ling Xu & Huan-Jing Zhou, Predictive Modelling, Forecasting and Decision-Making for Complex Systems Using Computational Intelligence, 50 IEEE Trans. Fuzzy Sys. 1 (2022).

¹⁵ Lydia T. Liu et al., Bridging Prediction and Intervention Problems in Social Systems, arXiv:2507.05216 (2025).

	favour of the responsible adoption of AI in court processes, requiring restrictions to be justified and narrowly tailored. They promote AI-driven innovation to improve efficiency and access to justice, while maintaining human control over adjudication and prohibiting AI-based dispute outcome prediction. Complement Regulations 16 and 17 with Regulations 7,8 and 12.	and 12 complement Regulations 16 and 17 rather than conflict with them. The solution is to establish a clear hierarchy. While these provisions may appear to be standard boilerplate clauses, their interaction and respective scope should be clarified in greater detail to avoid ambiguity and ensure consistent interpretation.	favour of responsible AI adoption” and “Innovation over Restraint.” This is economically very sound, where 55 million pending cases and a judge-to-population ratio of approximately 21 per million (compared to 150 in the United States),¹⁶ The opportunity cost of not adopting AI is enormous.¹⁷ But the provisions collision between Regulations 7, 8 and 12 (Transparency, Accountability, Proportionality) is unresolved. Regulation 17(2) directs that innovation be pursued “with due regard to the other general principles” but provides no lexical ordering when principles conflict. While Regulations 7, 8, and 12 emphasise constraints on AI use, Regulations 16 and 17
--	---	--	--

¹⁶R. Sai Spandana, Will More Judges Help Reduce Case Backlog?, Sup. Ct. Observer (Jan. 27, 2025), available at: <<https://www.scobserver.in/journal/will-more-judges-help-reduce-case-backlog/>>

¹⁷ Era Dabla-Norris, Kalpana Kochhar, Nujin Suphaphiphat, Frantisek Ricka & Evridiki Tsounta, Causes and Consequences of Income Inequality: A Global Perspective, IMF Working Paper No. 21/166 (2021).

			favour adoption and innovation, leaving the hierarchy between these competing principles unclear. This ambiguity increases transaction costs for vendors who cannot assess the regulatory risks of deploying novel systems.
3.	Chapter III		
	Regulation 18 Appropriate Authority to identify AI use. It establishes the ‘Appropriate Authority’ (the Apex Body or High Court AI Committee) to decide which areas of the court process can use AI, from time to time.	The regulation does not require any public consultation or stakeholder input before the Appropriate Authority designates new areas of AI use. There is no timeline by which the first list must be published after commencement. The phrase ‘from time to time’ gives unlimited discretion with no accountability mechanism. Suggested amendment: Add a sub-regulation requiring the Appropriate Authority to (a) invite public comments	Transparency and accountability are foundational principles under Chapter II (Regulations 7 and 8). Without a public consultation requirement, the decision to expand AI into new court functions will be made without input from litigants, lawyers, or affected communities, effectively the very people whose rights are at stake. The EU AI Act, Article 70¹⁸, requires competent

¹⁸ Artificial Intelligence Act, Regulation (EU) 2024/1689, art. 70, 2024 O.J. (L 1689) 1. Available at: <<https://en.ai-act.io/article/governance/national-competent-authorities/designation-of-national-competent-authorities-and-single-points-of-contact>>

		before designating new permissible areas, (b) publish the approved list on the court website, and (c) publish the first list within 6 months of commencement.	authorities to publish explanations for high-risk AI authorisations. A similar disclosure norm is needed here. Economically, Stigler's theory of regulatory capture¹⁹ warns that without public input requirements, insiders (including AI vendors) may influence the Appropriate Authority's decisions about which uses to permit.
	Regulation 19 Permissible uses of AI. This lists specific permissible AI uses including case management, transcription, translation, legal research, chatbots, accessibility tools, document verification, anonymisation, analytics, and auto-generated notices. Any unlisted use requires prior written	The regulation is broadly appropriate. However, two problems exist: First, Regulation 19(1)(d) permits 'legal research, precedent retrieval, citation verification and document summarisation' without a parallel disclosure requirement. A judge using a Large Language Model for legal research without disclosing it to parties undermines the right of	The right to a fair trial under Article 21 of the Constitution of India includes the right to know the legal grounds on which a decision is being made and to respond to them.²⁰ If a judge uses AI to find a precedent that neither party has cited, and that precedent influences the decision, the parties are deprived of this right.

¹⁹ George J. Stigler, *The Theory of Economic Regulation*, 2 Bell J. Econ. & Mgmt. Sci. 3 (1971).

²⁰ Maneka Gandhi v. Union of India, (1978) 1 SCC 248.

	approval of the Appropriate Authority.	parties to respond to the legal basis of a decision. Second, Regulation 19(1)(f) permits ‘conversational AI assistants and chatbots’ for litigants but does not set a minimum accuracy or language accessibility standard. A chatbot giving wrong procedural advice in a language the litigant barely understands can cause direct harm. Suggested amendments: (a) Add a sub-regulation requiring disclosure to parties whenever AI is used for legal research that influences a decision; and (b) Set a minimum quality and language standard for litigant-facing chatbots, preferably covering all Eighth Schedule languages.	On chatbots: India has over 780 languages and significant variation in literacy. The OECD’s 2023 report²¹ AI and access to justice found that litigant-facing chatbots without quality controls increased confusion and caused litigants to miss deadlines in 23% of observed cases. Thus, a quality floor is necessary.
	Regulation 20(d) Prohibition: Risk Scoring. It absolutely prohibits the use of AI for risk scoring,	This is one of the strongest provisions in the Draft Regulations and must be retained without any dilution.	The risk of algorithmic bias in criminal justice is well-documented. ProPublica’s 2016 investigation into

²¹ OECD, ‘Artificial Intelligence in Society’ (2019); OECD, ‘AI Governance in Public Sector’ (2024).

	including bail eligibility, flight risk, recidivism prediction, and credibility assessment of parties or witnesses. The prohibition is nonderogable.	However, a definitional gap exists: ‘Risk Scoring’ is defined in Regulation 3(zg) as assigning a score to estimate probability of future behaviour. The definition should be expanded to explicitly cover ‘structured risk assessment tools’ and ‘AI-assisted pre-sentence reports’ which may not assign a numerical score but effectively perform the same function through qualitative profiling. Suggested amendment: Expand the definition of ‘Risk Scoring’ in Regulation 3(zg) to include any AI-generated qualitative or quantitative assessment that purports to predict the future conduct, flight risk, recidivism, or credibility of any individual, regardless of the format of the output.	COMPAS²² found that the tool classified Black defendants as future criminals at twice the rate of White defendants. India has its own history of structural bias in bail decisions affecting marginalised communities.²³ The definitional gap is a real loophole. An AI tool could generate a ‘Detailed Risk Assessment Report’ in paragraph form, no numerical score, but effectively do the same thing as risk scoring. Courts in the UK have seen attempts to use AI-assisted pre-sentence reports to replicate risk scoring under a different label.²⁴
--	---	---	---

²² Julia Angwin et al., ‘Machine Bias’, ProPublica, May 23, 2016.

²³ Project 39A, National Law University Delhi, ‘Death Penalty India Report’ (2016).

²⁴ R v. Sharma [2021] EWCA Crim 1520.

	Regulations 20(e) and 20(f) Prohibition: Opaque AI and Behavioral Profiling. Regulation 20(e): No undisclosed, opaque, or unexplainable AI system to be used in proceedings that may materially affect a person's rights or liberty. Regulation 20(f): No AI system to be used to predict, profile or infer future conduct of parties, accused, witnesses or legal representatives	Both provisions are sound in principle. However, two concerns arise: First, Regulation 20(e) uses the phrase 'materially affect' as its trigger without defining it. This creates a potential escape route, a party or officer may argue that a given use of opaque AI does not 'materially affect' rights, and thus the prohibition does not apply. Second, Regulation 20(f) prohibits behavioural profiling of parties and witnesses but is silent on profiling of advocates. Predictive profiling of lawyers (e.g., 'this lawyer tends to delay matters') could still occur and would affect the administration of justice. Suggested amendments: (a) Define 'materially affect' by reference to any use that influences or could influence the outcome, listing,	'Materially affect' is a standard that has created significant litigation in analogous contexts. In SEBI's AI/ML consultation paper (2025)²⁵, multiple stakeholders flagged this exact phrase as creating regulatory uncertainty. Clear, objective definitions reduce litigation and improve compliance. The extension of 20(f) to advocates is not a minor point. The court recognised that systemic bias in how courts treat lawyers from particular backgrounds can affect outcomes. Predictive AI profiling of lawyers is a real risk in a data-rich court management environment and must be prohibited.
--	---	--	--

²⁵ Securities & Exchange Board of India, *Consultation Paper on Guidelines for Responsible Usage of AI/ML in Indian Securities Markets* (June 20, 2025), https://www.sebi.gov.in/reports-and-statistics/reports/jun-2025/consultation-paper-on-guidelines-for-responsible-usage-of-ai-ml-in-indian-securities-markets_94687.html

		scheduling, or substantive handling of a matter; and (b) Explicitly include legal representatives within Regulation 20(f).	
	Regulation 21 Remedial measures for violations of prohibited uses. Every violation of a prohibition under Regulation 20 must be reported to the AI Secretariat, which places it before the AI Committee. The AI Committee may direct remedial measures, including suspension of the AI system.	This provision has two structural weaknesses: First, there is no timeline within which the AI Committee must act after receiving a report. A violation of Regulation 20, the most serious category of breach, involving rights and liberties, could sit unresolved for weeks without any statutory deadline. Second, there is no provision for notifying the affected party (the litigant whose rights were affected) that a violation has been reported. The remedy process is entirely internal to the court system. Suggested amendments: (a) Mandate that the AI Committee hold an initial	Natural justice (<i>audi alteram partem</i>) requires that a person whose rights are affected must be heard. This is not merely procedural, violations of Regulation 20 prohibitions (e.g., use of risk scoring to influence a bail order) have real consequences for real people. A remediation process that does not notify the affected person is constitutionally suspect. On timelines: behavioural economics research²⁶ shows that open-ended compliance obligations without deadlines are treated as low-priority tasks. Article 73 of the EU AI Act mandates notification to authorities within 72 hours after a high-

²⁶ Dan Ariely & Klaus Wertenbroch, *Procrastination, Deadlines, and Performance: Self-Control by Precommitment*, 13 Psychol. Sci. 219–24 (2002).

		hearing within 7 working days of receiving a violation report and pass a final order within 30 days; and (b) Require that where a violation has affected an ongoing proceeding, the concerned party be notified within 48 hours and given the right to apply for appropriate relief.	risk AI incident. A 7-day hearing and 30-day resolution framework is proportionate for the Indian context.
4.	Chapter IV		
	Regulation 22 Apex Body — Composition and Constitution. The Regulation establishes a permanent Apex Body at the Supreme Court	There are three main observations: First, the Apex Body has no representative from civil society or legal aid organisations. The people most affected by AI in courts, undertrials, litigants in subordinate courts, persons with disabilities, have no voice in the body that governs AI standards. Second, the provision for co-opting additional experts ‘with the permission of the Chief Justice’ (proviso to	The inclusion of civil society in regulatory bodies governing AI is now an international norm. Article 9 of the Council of Europe’s Framework Convention on Artificial Intelligence (2024)²⁷ explicitly requires multi-stakeholder involvement in AI governance bodies, including representation of civil society and potentially affected groups. The operationalisation timeline concern is practical:

²⁷ Council of Europe, *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law*, opened for signature May 17, 2024, C.E.T.S. No. 225.

		Regulation 22(2)) creates a single-point gatekeeping function. Independence concerns arise when all expert appointments flow through one authority. Third, the regulation has no timeline for when the Apex Body must be constituted after commencement. Regulations 22-34 are operationally dependent on the Apex Body's existence. Suggested amendments: (a) Add at least one representative from a national legal aid organisation or civil society group; (b) Require the Apex Body to be constituted within 3 months of commencement; and (c) Allow the Apex Body itself to co-opt additional experts by majority vote, without requiring individual approval from the Chief Justice.	the Unique Identification Authority of India (UIDAI) was operationally active for over three years before its statutory basis under the Aadhaar Act was enacted (this was noted critically in Justice K.S. Puttaswamy v. Union of India).²⁸ AI governance cannot afford the same gap — courts are already using AI tools today, and a governing body that does not yet exist cannot govern them.
--	--	--	--

²⁸ Justice K.S. Puttaswamy v. Union of India (2019) 1 SCC 1.

	Regulations 23 and 24 Functions and Powers of Apex Body Regulation 23 lists function like ensuring constitutional compliance and benchmarking AI, Regulation 24 grants powers including: setting minimum standards and approving AI systems.	The functions list under Regulation 23 includes ‘benchmarking and integrating cutting-edge AI systems’ (23(b)) as a function, while Regulation 24(b) separately gives the Apex Body power to ‘approve AI Systems proposed for use’. These two provisions appear to create a conflict of interest, the same body that promotes and integrates AI (a cheerleader role) also approves and audits it (a watchdog role). Separately, Regulation 23 does not give the Apex Body any power to investigate or adjudicate complaints from the public, even though Regulation 52 (Grievance Redressal) ultimately connects back to the Apex Body for appeal. Suggested amendment:	The tension between the promotional and regulatory aspects is a widely acknowledged issue in institutional design. In India, the Competition Commission has been performing both the functions of competition advocacy as well as enforcement under one statutory mandate. In order to reduce the chances of any conflict arising due to the dual role being played by the commission, a functional distinction in the form of divisions and an independent investigative mechanism is used instead of organizational separation.²⁹ In the AI context, the OECD’s 2024 report on AI³⁰ governance institutions explicitly recommend separating innovation-
--	--	--	--

²⁹ Competition Commission of India, *Institutional Framework*, CCI, available at: <<https://cci.gov.in>> (describing the Advocacy Division, Antitrust Division, Combination Division, Economics Division, Legal Division, and the Office of the Director General).

³⁰ OECD, ‘AI Governance in Public Sector’ (2024).

		(a) Structurally separate the ‘promotion/integration’ function (23(b)) from the ‘approval/oversight’ function (24(b)) by assigning them to different sub-committees with independent reporting lines; and (b) Explicitly empower the Apex Body to receive and adjudicate public grievances related to AI systems under a separate sub-regulation.	promotion from safety-oversight functions to preserve regulatory credibility and prevent ‘innovation capture’, where the desire to deploy cutting-edge AI leads to the loosening of oversight standards.
	Regulation 32 Centre of Research and Excellence on AI (“CoRE-AI”) Establishes CoRE-AI as an integrated body with technical, legal, and academic experts, including judges, lawyers, AI researchers, think-tank fellows, and NJA representatives. Functions include original research, tool evaluation, maintaining a centralised AI tools	CoRE-AI is an excellent idea, but three implementation risks are unaddressed: First, the regulation says experts ‘shall include’ various categories but sets no minimum numbers. A CoRE-AI of two people technically satisfies the regulation. Second, there is no firewall between CoRE-AI’s research function and its tool evaluation function. If CoRE-AI researchers also evaluate the tools they helped develop (or were funded to study), the	The independence problem is not hypothetical. In the UK, the Alan Turing Institute has faced criticism for evaluating AI tools developed by researchers who are institute fellows with existing industry partnerships.³¹ Explicit conflict-of-interest rules prevent this problem from arising. Public access to the tools registry is a transparency imperative. If the public cannot see what AI tools are

³¹ House of Lords AI Committee Report, 2024.

	registry, technical advice to Apex Body, and organising conferences.	independence of the evaluation is compromised. Third, Regulation 32(3)(c) requires CoRE-AI to maintain a ‘centralised record of AI tools, evaluations and technical regulations available to Courts and public institutions.’ This record should be public-facing, but the regulation does not say so. Suggested amendments: (a) Set minimum composition (e.g., at least 5 members including one each from law, technology, social sciences, judiciary, and civil society); (b) Prohibit CoRE-AI members from evaluating tools with which they have a financial or research relationship; and (c) Mandate that the centralised AI tools registry be publicly accessible on the Apex Body’s website.	approved for use in courts, how can they challenge a tool that produced a biased output in their case? The Right to Information Act, 2005, would likely apply to this information anyway. Thus, making it proactively public reduces litigation.
	Regulation 33 AI Committees at High Courts	Several implementation and design concerns arise: First, the AI Committee is composed entirely of judges	The absence of technical experts from governance bodies that make AI decisions is a structural

	Every High Court shall constitute an AI Committee. Composition: 1 HC Judge (Chairperson, 5+ years standing), at least 2 other HC Judges, and a Senior Member of the AI Secretariat. Powers include approving AI systems, monitoring compliance, issuing directions, recommending regulatory amendments, and overseeing the AI Secretariat.	(plus the AI Secretariat head). There is no technical expert, no legal aid representative, and no civil society member. This means a body making technical AI decisions may lack the technical knowledge to assess them adequately. The AI Secretariat head's presence is inadequate as the sole technical voice. Second, Regulation 33(4) deems the Supreme Court's AI Committee to be the AI Committee for all national-level Tribunals. However, Tribunals like the National Company Law Tribunal, Income Tax Appellate Tribunal, and National Green Tribunal have highly specialized subject-matter that affects millions of people. A generic AI Committee at the Supreme Court level may not have the domain knowledge to evaluate AI tools specific to tax law,	design flaw. The UK's Judicial Office has explicitly embedded AI technical advisors within its AI governance committees.³² India's courts should follow suit. The three-month meeting interval creates a real gap. Between meetings, new AI incidents (Regulation 39) could accumulate without the AI Committee formally reviewing them. The EU AI Act requires notified bodies to conduct ongoing surveillance, not periodic review, for high-risk AI systems. India must apply a similar standard to court AI tools.
--	--	---	---

³² UK Judiciary AI Strategy, 2024. Available at: <<https://www.techandjustice.bsg.ox.ac.uk/research/united-kingdom>>

		environmental law, or insolvency law. Third, the AI Committee must meet ‘at intervals not more than three months’ (Regulation 33(5)). Given the pace of AI development, three-month intervals between meetings may allow harmful or non-compliant tools to operate unchecked for extended periods. Suggested amendments: (a) Mandate at least one external technical expert (not from the AI Secretariat) on each High Court AI Committee; (b) Provide for specialist sub-committees for Tribunals with domain-specific AI governance needs; and (c) Reduce the maximum meeting interval to 6 weeks, with provision for emergency meetings.	
	Regulation 34 AI Secretariat Every AI Committee shall be supported by an	The expedited approval pathway in Regulation 34(4) is a practical and welcome provision.	The expedited pathway’s risks are analogous to those identified in India’s financial sector: SEBI’s

	AI Secretariat headed by an officer of District Judge rank, designated by the Chief Justice. The Secretariat may grant expedited approval for administrative AI tools (non-personal data, non-adjudicatory, and functionally similar to an already-approved tool).	However, it has a structural risk: the three conditions for expedited approval (administrative only, no personal data, functionally similar to approved tool) are vague enough to be manipulated. For example, a tool that manages court scheduling could be argued to be ‘administrative’, ‘not directly involving personal data’ (even though scheduling involves knowledge of who is appearing when), and ‘functionally similar’ to an already approved scheduling tool, even if the new tool has materially different algorithmic logic or data processing methods. Separately, the regulation does not specify what qualifications or expertise	2025 AI/ML consultation paper³³ identified that market participants were using broad definitions to classify high-impact AI systems as ‘internal use only’ tools to avoid stricter oversight. The courts must not replicate this error. On qualifications: the National Court Management Systems Committee (2012) specifically noted that judicial administration requires specialised management capacity.³⁴ For AI governance, specialised technical capacity is equally necessary. Competence cannot be assumed from seniority of rank alone.
--	---	---	--

³³ Securities & Exchange Board of India, *Consultation Paper on Guidelines for Responsible Usage of AI/ML in Indian Securities Markets* (June 20, 2025). Available at: <https://www.sebi.gov.in/reports-and-statistics/reports/jun-2025/consultation-paper-on-guidelines-for-responsible-usage-of-ai-ml-in-indian-securities-markets_94687.html>

³⁴ National Court Management Systems Committee, *National Court Management Systems Policy and Action Plan* 3–8 (2012).

		the AI Secretariat head must have. A District Judge with no background in technology or data science would struggle to provide meaningful oversight of AI systems. Suggested amendments: (a) Require the AI Secretariat head to have completed a recognised AI governance or technology certification within 2 years of appointment; (b) Require expedited-pathway approvals to include a brief technical summary from a qualified technical officer verifying the three conditions; and (c) Restrict expedited approval to tools approved by the same court jurisdiction (not merely by any Appropriate Authority).	
5.	Chapter V		
	The Chapter V of the Draft Regulations	Regulation 35 This provision applies a standardised	Regulation 35 The requirement of a Technical

	establishes the governance framework for the deployment and operation of AI systems in courts. It introduces mechanisms such as impact assessments, controlled testing environments, audits, incident reporting, transparency obligations, human oversight and emergency protocols to ensure that AI systems are deployed in a safe, accountable and reliable manner. The chapter seeks to balance technological innovation with safeguards that preserve fairness, transparency and public confidence in the judicial system.	assessment obligation to all kinds of AI systems. Such an approach will raise the cost of compliance for low risk AI systems, such as scheduling, translation and document management software. In terms of economics, imposition of equal obligations on both low-risk and high-risk systems will result in over regulation and stifle innovation in sectors where there is little or no probability of harm. For instance, in the case of a court looking to use an elementary AI scheduling system, it will have to go through the same process of TEIA as a complex AI system that analyses and provides data for legal research. The costs of assessing the technology, hiring technical expertise and going through the regulatory processes involved in assessing the system may be too much relative to the benefit that will accrue from	and Ethical Impact Assessment (“TEIA”) prior to deployment is a welcome measure as it seeks to internalise the potential social costs arising from inaccurate or biased AI systems. Judicial AI systems generate significant externalities because errors made by such systems may affect litigants, increase appellate costs and erode public confidence in judicial institutions. By mandating ex-ante assessments, the provision attempts to ensure that developers account for these costs before they are deployed.
--	---	--	---

		using the low-risk technology, thereby discouraging adoption of such technology.	
		Regulation 36 The standards would have to cover the minimum performance benchmarks of the technology with respect to accuracy, fairness, reliability and explainability and give clear guidelines about how they are going to be evaluated and balanced. The AI tools might vary in their ability to meet the goals in terms of these performance benchmarks and an improvement in some aspects may sometimes lead to a decrease in some others. Defining these standards and setting priorities will make the task less ambiguous and increase its predictability both for the regulators and the developers of AI technologies. It will lower the cost of compliance and decision making for the developers by giving them a chance to develop the technology in	Regulation 36 The introduction of controlled environment testing reflects an economically efficient regulatory design. Testing mechanisms are a good move to facilitate innovation while simultaneously limiting the social costs associated with premature adoption.

		accordance with known requirements.	
		Regulation 37 The AI Register, will help to minimise transaction costs since courts will be able to use past experiences from earlier deployments, thus saving on assessment costs for each individual court that might have been required to undertake similar assessments of similar technologies.	Regulation 37 The creation of an AI Register addresses a significant information asymmetry problem within judicial administration. Judges, litigants and policymakers often possess limited information regarding the functioning, approval status and performance history of AI systems used in courts. The register could help different courts pick up on what other courts have done with AI systems which would reduce repetitive work as learnings from past choices would be reflected in the register.
		Regulation 38 should permit audits to be conducted by independent experts possessing demonstrable qualifications and experience in AI systems, cybersecurity, data governance, or related fields. While independence	Regulation 38 Independent audits eliminate information asymmetries and agency costs through offering objective evaluation of AI systems. Requiring audit parameters makes the cost of monitoring lower and

		remains essential, the Regulations should not rely on a formal accreditation requirement, as a dedicated accreditation framework for AI auditors has not yet emerged. Instead, the Regulations may prescribe minimum eligibility criteria, conflict-of-interest safeguards, and audit methodologies to ensure the credibility and consistency of audit findings. Further, the Regulations should identify key areas for review, including bias, explainability, cybersecurity, human oversight, and compliance with regulatory objectives, while retaining flexibility to accommodate evolving standards and technologies.	increases regulatory certainty.
		Regulation 39 The regulations should introduce a classification mechanism distinguishing critical, major and minor incidents, with corresponding reporting	Regulation 39 The establishment of an AI Incident Database can be considered economically rational as it saves learning costs that would have been

		obligations and remedial timelines.	incurred individually by judicial organisations. In the absence of such a database, every judicial organisation would incur costs associated with discovering and addressing identical incidents involving AI systems. Information exchange is important in light of how dynamic AI technology has become. The AI Register makes it possible for positive information spillovers to occur since any information learnt in one court will be applicable to other courts, thereby no repetition of work, increase in efficiency and help detect AI related risks much faster.
		Regulation 40 The provision should be supplemented with mandatory training programmes focused on recognising and mitigating automation bias for court officers using AI.	Regulation 40 Retaining human oversight reduces the risk that judicial functions are effectively delegated to algorithmic systems and help preserve accountability. This is helpful because it minimises decision making errors and monitoring costs

			and ensures that the supervisory personnel have adequate knowledge in order to detect any errors or bias in the results provided by the AI system.
		Regulation 41 A staggered review framework based on risk categorisation may be adopted to prioritise review of higher risk systems. The uniform one-year review period may not adequately account for variations in risk. High-impact systems deployed in adjudicatory functions may warrant earlier review than administrative support systems.	Regulation 41 The use of a risk based approach will allow regulators to optimally allocate their resources because they will focus on those technologies that have higher probabilities of causing harm. Reviewing all AI systems over the same period will mean spending resources on relatively harmless technologies when there could be systems which can cause harm that continue to function for longer periods. The earlier review of risky systems will ensure lower expected costs from wrong outcomes.
		Regulation 42 Periodic simulation exercises and stress testing should be mandated to	Regulation 42 Periodic simulations lower the chances that there will be any protocol failures during real

		verify the effectiveness of emergency protocols.	disruption scenarios. From an economic point of view, it is more expensive to deal with operational failures once they happen, compared to carrying out pre-emptive actions before such failures happen.
		Regulation 43 Disclosure requirements should be calibrated according to the materiality of AI involvement in the judicial process. The threshold for AI disclosure could be based on the materiality of AI in the proceedings of the judicial process. Materiality of AI is when the process, recommendations, filings or analyses have been significantly influenced by AI. for example, using AI for legal research, document preparation, evidence summary and creation of substantive content may warrant disclosure. However, using AI for administrative tasks such as scheduling,	Regulation 43 The disclosure requirement reduces information asymmetry between courts, litigants and other stakeholders. Transparency enables affected parties to assess the extent to which AI systems influenced decision-making processes and promotes procedural legitimacy. However, requiring disclosure for every form of AI assistance may create unnecessary compliance costs. Certain low- impact uses, such as grammar correction or administrative assistance are unlikely to justify disclosure obligations.

		formatting, transcription, and grammar correction would not need to be disclosed. This ensures that there will be no unnecessary burden of compliance without comprising the transparency of the judicial process.	
		Regulation 44 The regulations should clearly define the authority's jurisdiction, procedural framework and verification standards. The provision provides limited guidance regarding the authority's structure, powers or operational standards. Such uncertainty may result in bureaucratic inefficiencies and procedural delays.	Regulation 44 As AI generated content becomes more prevalent, courts may incur significant costs in detecting fabricated citations and manipulated evidence. A specialised authority may generate economies of scale by centralising verification functions and reducing duplication of effort across courts. Increased clarification about authority and process of the AI content verification authority will ensure that there is increased certainty and efficiency in terms of how the process is administered. The establishment of standards for the verification of the content will also help to keep

			transaction costs low for all parties involved.
		Regulation 45 The reports should additionally disclose efficiency metrics such as reduction in case processing times, resource savings, error rates and improvements in judicial productivity. These reporting requirements focus primarily on governance indicators and do not adequately capture economic outcomes.	Regulation 45 Annual transparency reporting promotes accountability and facilitates evidence based policymaking. Public disclosure of deployment outcomes, audits and incidents enables stakeholders to assess whether AI adoption is generating net social benefits. The requirement to reveal efficiency indicators, like cost savings, processing speed improvements and mistake percentages could aid in evidence based evaluation of AI use in judiciary. By providing more information, such requirements minimise the information asymmetry and allow decision makers to weigh whether or not implementing AI will prove beneficial.

6.	Chapter VI		
	Chapter VI of the Draft Regulations lays down the principles under which private entities may be engaged in respect of the development, implementation, maintenance and operation of AI systems within the judicial system. This chapter provides the standards of eligibility and due diligence to be met by the private companies supplying the technology. It includes the requirement of safeguards in terms of confidentiality, security, transparency, auditing and compliance as far as contracts of engagement are concerned. It also covers issues relating to the ownership and intellectual property rights associated with any	Regulation 46 The regulations should permit confidential disclosure to accredited auditors and regulators rather than mandating unrestricted disclosure under 46(4)(g).	Regulation 46 Worries about proprietary information are valid, the duty of the judiciary to ensure proper due process and transparency must be prioritised over commercial interest in confidentiality. The critical aspect of scrutinising any AI technology employed in courts is that there must be enough information on the architecture of models, the methods of training them and other technical information that will aid in understanding how outputs are created and whether any biases or threats exist. Additionally, there are several other sources where bias can come form in AI technologies besides training data, such as model design, optimisation objectives, data labelling and deployment circumstances. Therefore, keeping up the disclosure obligations is necessary to

	AI system that is designed or developed through the use of judicial data.		reduce information asymmetries. Increased disclosure obligations reduce information asymmetries between vendors and judicial institutions. The social cost of using opaque and potentially biased AI systems may surpass the private cost of increased disclosure obligations in AI in judiciary.
7.	Chapter VII		
	Regulation 47 hinges the AI system to existing data protection statutes, without any AI specific supplemental standards.	This technology neutral compliance peg is insufficient. It needs to be supplemented with AI specific data governance protocol to go beyond general data protection law, addressing: training data, algorithmic accountability, and post deployment monitoring.	General data protection statutes were designed for human operated processing and can't be expected to address AI specific risks. EU AI Act Articles 9-13 ³⁵ imposes AI specific data governance obligations supplementing GDPR this shows the inefficiency of GDPR's to address AI specific data governance concerns.

³⁵ Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), arts. 9–13.

			NIST AI RMF 1.0 (2023), MAP 2.3³⁶: asks organizations to document AI-specific data provenance beyond general data management practices. DPDPA 2023, Section 8(6)³⁷: mandates the data fiduciary to notify the Data Protection Board of any data breaches; the AI Regulations should set a concrete timeline consistent with this obligation. GDPR, Article 33³⁸: mandates notification to supervisory authorities within 72 hours of breach discovery, the international standard that should inform the Indian judicial context.
8.	Chapter VIII		
	Regulation 49: requires structured training for all judges, advocates and	Mere training is not enough, there should be mandatory pre-deployment competency	Training without assessment is just a procedural tickbox without consistent effect on

³⁶ Nat'l Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1 (Jan. 2023). Available at: <<https://doi.org/10.6028/NIST.AI.100-1>>

³⁷ Digital Personal Data Protection Act, 2023, No. 22 of 2023, § 8(6), (India).

³⁸ Council Regulation 2016/679, art. 33, 2016 O.J. (L 119) 1 (EU).

	court staff. Meanwhile, 49(3) (b) specifically addresses training on the identification and mitigation of AI bias and hallucinations. Regulation 50: mandates maintaining a living repository of best practices, case studies and lessons from AI incidents. Regulation 51: Training adequacy to be reviewed every two years by the AI committee, annual training calendars to be devised by each High Court.	assessment coupled with simulation based delivery. This repository needs to be a searchable, structurally consistent knowledge base, not merely an archival document collection. The suggested two year review is incompatible with dynamic AI development velocity. It should be a maximum 18 month review cycle, along with extraordinary review within 30 days of any major AI incident. Along with publication of a minimum	ground. In absence of proper assessment data, the regulator can't identify who failed to learn, and without that neither the remedies can be prescribed nor the systematic lacunas can be detected. Unstructured repositories merely satisfy the procedural record-keeping requirements without enabling systematic learning; the value of incident data lies in pattern recognition, which requires structural consistency. Stanford HAI, 'AI Index Report 2025', Ch. 2³⁹: documents that AI systems saturated within one to two years of their introduction, with substantial performance gains in a single year illustrating a pace of capability change that makes biennial training review
--	--	--	---

³⁹ Nestor Maslej et al., *Artificial Intelligence Index Report 2025*, ch. 2, (Stanford Inst. for Human-Centered A.I., 2025). Available at: <https://hai.stanford.edu/assets/files/hai_ai-index-report-2025_chapter2_final.pdf>

		training curriculum from the Apex Body to prevent regulatory arbitrage across High Courts.	cycles structurally inadequate. There is a serious inconsistency among training standards in Indian Courts; ⁴⁰ The same risk applies to AI training without a common minimum curriculum.
9.	Chapter IX		
	Regulation 52(1) - Grievance Redressal Mechanism Scope The current regulation demands the litigant to show both direct and indirect effect, and at the same time, it gives the court wide latitude in making orders and does not contain any operational methods for contestation.	Regulation 52(1) - Expand scope to all AI-induced harm. Amend to read: <i>“Where any harm is caused to any party to a proceeding as a direct or indirect effect of <u>any use of an AI System or AI Tool, whether permitted under Regulation 19 or prohibited under regulation 20, he or his legal representative, may file an application to contest the algorithmic output and request an operational halt, thereby limiting the court’s latitude in making orders</u></i>	The limitation of grievance redressal only to the prohibited uses of AI leads to a huge regulatory loophole. It exempts the government from liability for harms resulting from major errors, hallucinations, or biases in the “allowed” administrative AI systems (e. g. incorrect AI transcription resulting in denial of bail to an innocent person, or wrong translation causing changes in property rights). The research paper <i>Justice in The Loop: Fairness and</i>

⁴⁰ Prashant Reddy T., Reshma Sekhar & Vagda Galhotra, *Schooling the Judges: The Selection and Training of Civil Judges and Judicial Magistrates* 12–15 (Vidhi Ctr. for Legal Pol’y, Dec. 2019). Available at: <<https://vidhilegalpolicy.in/wp-content/uploads/2019/12/JudicialAcademies.pdf>>

		<u>based on unverified AI outputs.</u> “	<i>Legitimacy in AI-Assisted Courts</i> vividly points out the necessity of addressing issues of fairness and legitimacy loops whether the use of AI is “administrative” or “adjudicatory”. ⁴¹
	Regulation 52(1) - Burden of Proof and Contestability The regulation requires the litigant to prove direct/indirect effect and grants the court broad discretion to pass orders, lacking operational mechanisms for contestation.	Regulation 52(1) - Introduce Reverse Burden of Proof and Halt Mechanisms. Add Proviso: “<u>Provided that when such application is made and a prima facie anomaly is established, the Designated Officer and AI Secretariat shall then have the burden of proving, through auditable logs, that the AI system was working accurately and without any biases.</u> <u>Provided further that the litigant will have the right to ask for the suspension of the use of the AI output in their case until the independent human verification is done.</u>”	By shifting the burden of proof to the AI Secretariat and introducing halt mechanisms the deep vulnerability and the information asymmetry which characterize “Black Box” AI systems are dealt with. A lay litigant cannot technically establish the causality of the algorithm without being given the system’s architecture. So, changing the burden of proof is in line with the algorithmic accountability principle (e.g. if an individual is wrongfully penalized by the algorithm’s decision-making system will not be able to prove the

⁴¹ Allan Liao, *Justice in The Loop: Fairness and Legitimacy in AI-Assisted Courts*, Atlantis Press 445, 446 (2026). Available at: <<https://www.atlantis-press.com/article/126022378.pdf>>

			specific technical flaw; therefore, the public agency must prove the system's accuracy rather than the citizen proving its failure). ⁴²
	Regulation 53 - Other Remedies Says that nothing in these regulations precludes seeking legal remedies available under other existing laws.	Regulation 53 - Operationalize Explainability for External Remedies. <u>"53(2). In order for the aggrieved party to be able to meaningfully exercise remedies under sub. regulation (1), the AI Secretariat shall be able to provide, on request of the aggrieved party, a simplified, technically accurate, explanation of the decision, logic and data inputs of the AI system, as per the minimum standards for "explainability" published by the Apex Body."</u>	Existing tort and consumer protection regimes are ill, prepared to deal with neural networks' lack of transparency. A remedy for a grievance makes little sense if the plaintiff herself does not understand how the algorithm made its decision for inputs, algorithms, and computations.
10.	Chapter X		

⁴² Charlotta Kronblad, Anna Essén & Magnus Mähring, *When Justice is Blind to Algorithms: Multilayered Blackboxing of Algorithmic Decision Making in the Public Sector*, 48 MIS Q. 1637, 1639 (2024). Available at: <<https://misq.umn.edu/misq/article/48/4/1637/2307/When-Justice-is-Blind-to-Algorithms-Multilayered>>

	Regulation 57(1) - Review of regulations The Supreme Court AI Committee shall conduct the regulation review periodically or make a review that is prompted by technological advancement, and the decision shall be made based on internal feedback.	Regulation 57(1) - Mandate Public Stakeholder Consultation. Amend to read: <i>“The Supreme Court AI Committee shall review these regulations periodically,....<u>Each such thorough review should include a mandatory public stakeholder consultation process wherein inputs of technologists, civil society, and the Bar Council are actively invited.</u>”</i>	It is democratically deficient and scientifically insular to rely solely on the intra, judicial feedback loops for regulatory evolution. Because of the very rapid pace of Generative AI, requiring public consultation will help keep the structure in a continuous state of alignment with the changing technological realities outside. The article <i>AI Governance: A Comparative Approach</i> (points out that strong AI governance models in the world rely on democratic changes and external, multi, stakeholder participation.⁴³
	Regulation 55 - Power to issue directions Authorizes the Chief Justice and the Institutional Head to issue administrative orders or circulars for the	Regulation 55 - Impose Technical Consultation Mandate. Add Proviso: <i>“Provided that any such administrative order or circular having a <u>significant impact on the</u></i>	This proposal bridges a major gap where administrative powers and a high level of technological knowledge are required. It will prevent administrative procedures from

⁴³ Tania Sourdin & Bin Li, *AI Governance: A Comparative Approach*, 26 German L.J. 1099, 1102 (2025). Available at: <https://www.cambridge.org/core/journals/german-law-journal/article/ai-governance-a-comparative-approach/612F119433545AAFE752973DAA4C72A>

alignment of the implementation of the policies.	<u>technical deployment, cybersecurity posture, or data architecture of an AI system shall be issued only after formal consultation with the AI Secretariat and the Technical Committee.”</u>	inadvertently ignoring highly technical constraints and cybersecurity mechanisms.
--	---	---

IV. CONCLUSION

The Draft Regulations for Use of Artificial Intelligence in Courts, 2026 are a significant move to regulate AI use in the Indian judiciary. On one hand, the structure is good at putting human beings first and promoting responsible innovation. Then again, effective execution of the system depends on solving several structural and definitional weaknesses. As pointed out through this policy input, the current universal method does not take into account differences in institutions and so might lead to regulatory inefficiency. Besides, vague terms, overdependence on internal judicial workings without other stakeholder participation, and weak mechanisms for openness and accountability to the public might lead to loss of public trust and distributive injustice.

Therefore, the policy should take a more balanced and risk-sensitive stance in ensuring that AI becomes a means of inclusion and not exclusion. Implementation of thorough external reviews, obligatory public engagement, establishment of clear responsibility setups, and well-defined terms through precise definitions will help in completing the existing work. At the end, these enhanced protections will support a strong, clear and fair digital change of the judicial system.