

PUCAR

(Public Collective for Avoidance and Resolution of Disputes)

Comments on the **Draft Regulations for Use of Artificial Intelligence in Courts, 2026**

Submitted by PUCAR

July 2026

INTRODUCTION

We welcome the Supreme Court's initiative to develop a comprehensive governance framework for the responsible adoption of Artificial Intelligence (AI) in the judiciary. Given the transformative potential of AI for improving access to justice, judicial administration and court efficiency, the effort to establish clear principles and institutional safeguards is both timely and commendable. By inviting public consultation, the Supreme Court has adopted an inclusive approach that will strengthen both the quality and legitimacy of the eventual regulatory framework.

PUCAR (Public Collective for Avoidance and Resolution of Disputes) is a non-profit-collaborative public mission to transform the dispute resolution experience of every person.

To ensure that these submissions reflect a broad range of expertise and perspectives, PUCAR conducted a closed stakeholder consultation as part of the process to arrive at these comments and recommendations. The consultation included contributions from the following community members of PUCAR:

1. Aishani Rai, Aapti Institute
2. Mousmi Panda, Aapti Institute
3. Priyanka Angadi, AnP Advocates
4. Leah Verghese
5. Smita Mutt
6. Manish Srivastava
7. Nehaa Chaudhari, Ikigai Law
8. K K Prahalad, PACTA
9. Nivedita Krishna, PACTA
10. Deepika Kinhal, TERES and CORD
11. Ganesh Gopal, Tinkerhub
12. Nikhil Narendran, Trilegal
13. Pavithra Manivannan, XKDR Forum and PUCAR
14. Supriya Sankaran, Agami and PUCAR
15. Varun Hemachandran, Agami and PUCAR
16. Anshumath Rao, PUCAR

SCOPE AND FOCUS OF COMMENTS

The Draft Regulations are broad in scope, spanning principles, technical standards, processes, and governance. Rather than responding section by section, we have focused our comments on two foundational areas that are difficult to change once a framework is in place, and on which everything else will depend.

1. Institutional Architecture — We examine the institutional framework proposed in Chapter IV in light of the full range of functions assigned across Chapters V to IX. We identify four structural concerns: fragmentation of responsibility and overlapping mandates; limited institutional continuity; sub-optimal use of judicial time and lack of functional specialisation; and high transaction costs for innovation. Underlying all four is a more fundamental point: institutional structures must not be built around individual technologies. Computerisation, digitisation, video conferencing, and AI are not separate challenges — they are successive tools in the ongoing transformation of court administration. Policies for each technology will legitimately differ — an AI governance policy will look different from a video conferencing policy. But the institutional machinery that plans, procures, implements, and evaluates technology must be unified. A new body for each new technology produces fragmentation, not governance. This makes institutional design the most consequential choice in any technology framework. Unlike substantive regulations, which can be revised as technology and needs evolve, institutional structures are significantly harder to reverse once established. Getting this right now matters more than getting anything else right later.

2. Standards for AI Systems and Vendors — The Draft Regulations establish who approves AI systems and how, but not what any system must actually demonstrate to qualify for approval. This is a significant gap. Across sectors, approval/ entry mechanisms are supported by clearly articulated standards. Whether approving a bridge, a pharmaceutical product or an aircraft, regulators rely on predefined standards/ criteria relating to the product, such as safety (eg. ISO 9001, FDA certifications), or relating to the firm, such as reliability (prior experience, financial soundness, net worth requirements). Without a defined benchmark, each High Court must independently determine what constitutes an acceptable AI system — duplicating effort across the system and, given the wide variation in technical capacity across India's judiciary, arriving at different answers. Defined standards would benefit all stakeholders. For courts, they reduce the burden of evaluation and create a consistent, well-documented approval process. For AI developers and vendors, they provide clarity and predictability on what their systems must satisfy. And for the public, they make the approval process more transparent and consistent, inspiring greater confidence in the judiciary's adoption of AI.

RECOMMENDATIONS AT A GLANCE

Recommendation 1: Establish a permanent specialised organisation for judicial administration

We recommend the establishment of a permanent Indian Courts Service — dedicated to judicial technology management, operating under judicial oversight. Rather than fragmenting governance across successive technology-specific bodies, a single institution capable of absorbing the functions currently held by the

e-Committee alongside the new AI governance mandate is more durable, more accountable, and more effective. In the short term, a dedicated office or team can be put in place.

Recommendation 2: Mandate minimum standards for judicial AI — and designate CoRE-AI to own them

We recommend that the Court mandate the development minimum standards — including on explainability, accuracy, data governance, human oversight, interoperability and training. These standards are a prerequisite to any AI system being sanctioned for court use. CoRE-AI, as a neutral expert body with no role in procurement or service delivery, is best placed to develop, maintain and periodically revise them. Designating CoRE-AI for this purpose also ensures the standards remain independent of vendor interest and institutional pressure.

SECTION 1: COMMENTS AND RECOMMENDATIONS ON INSTITUTIONAL ARCHITECTURE

A. Issues with the Proposed Design

(i) Fragmentation of responsibility and overlapping mandates

The Draft Regulations establish an Apex Body, five standing committees, an AI Committee in every High Court, an AI Secretariat at every High Court and District/Subordinate Court, and the Centre of Research and Excellence on AI (CoRE-AI). These institutions collectively perform functions ranging from policy formulation and standard-setting to procurement, technology evaluation, research, cybersecurity, audits, infrastructure planning and implementation.

The breadth of this framework reflects the importance that the Draft Regulations place on responsible AI adoption. At the same time, the institutional design raises an important governance question:

How should these functions and responsibilities relate to the judiciary's existing institutions for technology and reforms?

Technology-enabled reform in the judiciary is not a standalone activity. AI systems depend upon digital records, case management systems, cybersecurity infrastructure, procurement processes, budgeting, judicial administration and continuous training. These are functions that are already undertaken, at least in part, through existing institutions such as the Supreme Court e-Committee/ High Court Computer Committees, Committee for Digitisation and sometimes Committees for Administration or Infrastructure.

As AI becomes embedded into existing court processes, these functions will inevitably intersect. For instance, an AI-assisted scheduling system is simultaneously a case management project, a digitisation project, an infrastructure project and an AI project. Similarly, deploying AI for filing scrutiny or legal research requires coordination across procurement, cybersecurity, data management and judicial administration.

This creates two institutional risks:

- Accountability may become fragmented across multiple bodies, making it more difficult to identify which institution is ultimately responsible for initiating, planning, and delivering technology-enabled reform.
- Overlapping mandates may slow decision-making, particularly where multiple committees are required to examine related aspects of the same project before implementation can begin.

(ii) Limited institutional continuity

Technology systems require continuous evaluation, maintenance, upgrade, cybersecurity review, vendor management and periodic replacement. The Draft Regulations establish several committees and expert bodies to perform these functions. However, committees are, by design, deliberative institutions rather than permanent delivery organisations. Their

composition changes over time, members rotate, and institutional knowledge is not always retained.

Institutional experience suggests that organisations responsible for continuous execution perform better when constituted as permanent entities rather than committees. A permanent organisation allows operational autonomy based on clearly delegated authority. This enables quicker decision-making. The long tenure builds institutional memory by allowing expertise to accumulate over time. It also creates clear ownership over implementation, making it easier to attribute responsibility and create the feedback mechanisms for operational successes and failures.

Three examples from India illustrate this principle:

- **TNMSC:** The Tamil Nadu Medical Services Corporation (TNMSC) was established as a specialised statutory corporation to professionalise procurement and supply-chain management of medicines. By building permanent expertise in procurement, logistics and quality assurance, TNMSC transformed public health procurement in the State and served as a model for several others.
- **UIDAI:** The Unique Identification Authority of India (UIDAI) was created as a permanent institution to oversee the design, deployment, maintenance and continuous evolution of the Aadhaar ecosystem, including technology standards, cybersecurity, vendor management and ecosystem governance.
- **NHAI:** The National Highways Authority of India (NHAI) provides a dedicated institutional capability for planning, procuring, implementing and maintaining the national highway network. This execution is separated from policy formulation, which remains with the Ministry of Road, Transport and Highways.

AI adoption within the judiciary would benefit from having a permanent professional organisation capable of building technical expertise over time, with sufficient delegated authority that simultaneously enables timely operational decisions and ownership and accountability over those decisions.

(iii) Sub-optimal use of judicial time for operational functions

The Draft Regulations envisage judges occupying leadership positions across the Apex Body and AI Committees, and require committee oversight over approvals, audits, procurement, policy formulation, audits, evaluation and implementation.

Judicial leadership in AI adoption is both necessary and appropriate. Questions concerning judicial independence, fairness, due process and the appropriate use of AI ultimately require judicial oversight. However, not all functions in an AI governance framework require the same degree of judicial attention.

At present, operational responsibilities are combined with governance/ approval responsibilities within the same institutional body. For instance,

1. Regulation 46 makes the HC/ SC AI Committee as the final authority for **approval of an AI system** to be deployed. However, the approval process involves operational elements such as: evaluation of the AI system and procurement of the system which includes floating tenders, drafting contracts, and contract management post-award of contract.

2. The mandates of the five committees proposed in Regulations 26 to 30 such as innovating, developing and deploying AI tools, monitoring advances in AI research, establishment of infrastructure related to AI does not require the attention of judges. Whereas the financial aspects such as budgetary approval, cost-benefit evaluations and the audit of expenditure should be handled by the SC/ HC AI Committee.
3. Regulation 50 requires the SC/ HC AI Committee to maintain a living repository of best practices, case studies, lessons drawn from AI Incidents and guidance notes. This function does not necessarily require the time and attention of the judges.

Judicial time is among the judiciary's scarcest resources. Every additional governance structure entails meetings, agenda-setting, approvals, review of proposals and continuous oversight. At a time when reducing delays and improving adjudicatory capacity remain central institutional priorities, it is worth asking whether operational management of technology programmes should rest primarily on judge-led committees.

A widely accepted principle of organisational design is that responsibilities should be allocated according to comparative expertise — enabling specialist resources to focus on functions where they create the greatest value. Judges possess expertise that is uniquely valuable in adjudication, interpreting laws and constitutional principles, and policy-making. Technology evaluation, procurement, cybersecurity, programme management, vendor oversight and systems implementation require different forms of professional expertise that can be effectively exercised by permanent technical and administrative teams operating under judicial supervision.

(iv) Transaction costs for innovation

The Draft Regulations prioritise innovation over restraint. However, the proposed approval structure may inadvertently raise the cost of participating in the judicial technology market, particularly for smaller and newer entrants.

As the ecosystem matures, technology developers, start-ups and research institutions will require a predictable institutional interface through which they can engage with the judiciary. If approvals, evaluations and pilot deployments are undertaken independently across 25 High Courts and multiple tribunals, vendors may face repeated evaluation processes for substantially similar technologies. From the judiciary's perspective, repeated evaluations across jurisdictions also require significant institutional effort and may unintentionally favour well-established vendors over newer entrants that lack the resources to navigate multiple approval processes.

Many regulated sectors address similar concerns through nationally recognised qualification frameworks, empanelment mechanisms or common technical standards that permit decentralised procurement while avoiding repetitive evaluation. Exploring comparable approaches for judicial AI could reduce transaction costs, encourage innovation and promote greater competition, while preserving each court's autonomy over procurement decisions.

B. Recommendation: Establishment of a specialised organisation for judicial administration

The concerns outlined above point to a clear institutional need: the judiciary requires a permanent professional organisation capable of planning, executing and continuously managing technology-enabled reform, while preserving judicial leadership over policy, standards and strategic direction.

This distinction between governance and execution is well recognised across both public institutions and private organisations. Governing bodies determine policies and strategy, approve priorities, and exercise oversight. Permanent executive organisations translate those strategic decisions into operational outcomes through specialised expertise, institutional continuity and professional management. AI adoption presents an opportunity to adopt a similar institutional model within the judiciary.

Rather than fragmenting governance based on technologies, creating a single institution — capable of absorbing the functions currently held by the e-Committee alongside the new AI governance mandate, and whatever comes next — is more durable, more accountable, and more effective.

We respectfully recommend the establishment of a permanent institution Indian Courts Service — a specialised administrative and technical organisation dedicated to judicial administration and technology management. To preserve the federal nature of the judiciary, the services of this institution could be shared across multiple High Courts and Tribunals while retaining the flexibility for individual High Courts or Tribunals to establish their own dedicated institution where the scale of operations or local requirements so warrant. This body would support courts and judges by undertaking operational and technical functions, while judicial leadership would continue to determine policy, approve standards, set institutional priorities, and exercise oversight. In the short term a dedicated office can be created that performs the same functions and role.

This model is consistent with proposals for a National Judicial Technology Council under the eCourts Phase-III mission document, which envisaged a permanent institutional capability to support technology adoption across the judiciary. Rather than creating separate institutional arrangements for each technological reform initiative, a single organisation could support present and future technology initiatives through a common institutional platform.

This proposed institutional structure does not require a substantial redistribution of functions under the Draft Regulations. It merely requires a reallocation of existing functions according to their character.

This alternative structure is guided by three principles.

First, optimal use of judicial time. Decisions relating to the adoption of AI within courts including policy, governance, approval of AI systems and strategic priorities should remain under judicial leadership. At the same time, judges should devote their time to functions that uniquely require judicial judgment, while operational, technical and administrative functions are undertaken by permanent professional institutions possessing the requisite expertise.

Second, functional separation. The Draft Regulations presently combine decision-making, implementation and research functions within the same institutional

bodies. These functions require different capabilities and need not be performed by the same institution or actors. Separating strategy from day-to-day execution, and execution from research, improves accountability, avoids conflicts between implementation and oversight, and allows each institution to develop specialised expertise.

Third, separation of powers. Where the same institution both sets standards and procures against them, or both implements technology and audits its performance, the result is a conflict of interest rather than a system of accountability. Separating these functions — so that those who set standards do not procure, those who procure do not audit, and those who audit are independent of both — creates the checks and balances that responsible AI governance requires.

Based on these principles, the alternative institutional structure could comprise three distinct bodies with the following roles and responsibilities:

1. SC and HC AI Committee:

- a. Purpose: Ensuring that the adoption and use of AI in courts is safe to deploy and use consistent with constitutional principles, the governing principles in the Regulations and the broader objectives of the judiciary.
- b. Role: Primary governing body for AI in judiciary. Owns the decisions on what systems are deployed.
- c. Responsibilities: Providing strategic direction for AI projects, policy formulation, approval of AI systems and exercising oversight over the two bodies mentioned below.

2. The Indian Courts Service:

- a. Purpose: Translate judicial policy and decisions into sustained operational outcomes in a timely and cost-effective manner.
- b. Role: Judiciary's permanent execution and implementation arm. Owns the management and implementation of a project for deploying an AI system in court.
- c. Responsibilities: Providing the professional, technical and operational capability necessary to plan, procure, evaluate, deploy, maintain, operate and continuously improve technology and AI systems. Further, it can play the following roles:
 - i. Maintain an empanelled list of qualified vendors and approved tools — similar to the role performed by the Government e-Marketplace (GeM) — while leaving the ultimate procurement decision with the respective High Courts.
 - ii. Build and maintain the infrastructure for the controlled testing environment envisaged in the Draft Regulations.
 - iii. Create a permanent innovation platform — open to researchers, start-ups and academic institutions — through which AI solutions for judicial use can be developed and tested under controlled conditions. This model has precedent: SEBI's Regulatory Sandbox and NPCI's shared digital infrastructure both enable responsible innovation within defined boundaries.

- iv. Draw on empirical data on caseloads, staffing, infrastructure and technology adoption to ensure that planning and procurement are evidence-driven rather than ad hoc.
- v. Given that AI-related complaints often require technical investigation, examination of system logs, interaction with vendors and assessment of model behaviour, the Indian Courts Service would act as the first point of contact for such complaints. It would conduct the necessary fact-finding and prepare a technical report for consideration by the appropriate court, which would then adjudicate the dispute/ address the grievance.

3. Core-AI

- a. Purpose: To be the intellectual backbone of judicial AI governance, undertaking research, developing and maintaining the independent standards, benchmarks and evidence base against which all AI systems used in courts are evaluated.
- b. Role: Judiciary's neutral body for research, innovation and standard-setting. Owns the standards and benchmarks against which AI systems in courts are evaluated and audited.
- c. Responsibilities: Developing technical benchmarks and methodologies, evaluating emerging technologies, and fostering continuous learning.

Such a structure would improve accountability, build permanent technical capability, provide the required checks and balances for AI governance, and allow judges to focus on those functions that uniquely require judicial oversight.

SECTION 2: COMMENTS AND RECOMMENDATIONS ON STANDARDS FOR AI SYSTEMS AND VENDORS

Regulation 35 requires a Technical and Ethical Impact Assessment before any AI system is approved for use by the Appropriate Authority, and sets out the criteria that assessment must cover — purpose and architecture, training data quality, risks of bias and hallucination, cybersecurity, explainability, and incident reporting.

However, it does not provide the standard against which the answers will be judged. Knowing that an AI system must be evaluated for bias is necessary but not sufficient. The framework must also specify what level of bias is acceptable, for which functions, and how it is to be measured. Without this, the assessment becomes a documentation exercise rather than a genuine qualification threshold. The same gap applies across each criterion.

The absence of such standards creates a significant risk. Without a defined benchmark, each High Court must independently determine what constitutes an acceptable AI system — a burden that will be discharged differently across courts, given the wide variation in technical capacity and resources across India's judiciary. The result is likely to be inconsistent approvals.

We rely on minimum standards to govern almost every system that affects public safety and rights. They exist not to constrain innovation, but to ensure that innovation serves people reliably and consistently. Four familiar examples illustrate the point:

- **ISO 9001** — Before a pharmaceutical company can sell a drug in India, its entire manufacturing process must meet internationally defined quality standards. The product is not simply tested once and approved; the process by which it is made must consistently meet standards.
- **BIS Certification** — Before an electrical appliance can be sold in India, it must carry a Bureau of Indian Standards certification defining what the product must do and how it must perform under stress, before it reaches a single household.
- **Building Codes** — Before a commercial unit is constructed, it must meet defined structural, accessibility and fire safety standards. The architect or developer does not decide what "safe" means; the standard does.
- **PDF/A** — Courts across India already require judgments and filings to be stored in PDF/A format — a specific technical standard for long-term document preservation. This is a specification: it defines not just the file type but what the file must be capable of doing over time.

Standards serve two distinct functions. First, they enable courts to make objective, evidence-based approval decisions and set a minimum benchmark for approval. Second, they signal to the market the capabilities the judiciary seeks to encourage — providing AI developers and vendors with the clarity and predictability needed to design, test and invest accordingly.

We respectfully submit that articulating minimum specifications and evaluation criteria is not merely desirable — it is foundational to the integrity of the approval mechanism itself.

What must be specified?

In the context of the adoption of AI by the judiciary, the following are some of the critical standards that need to be defined:

- **Explainability** — The standard must require that AI systems provide reasons and evidence for every output, that those explanations are understandable to the user receiving them, and that the explanation reflects the system's actual decision-making process rather than a post-hoc justification.
- **Accuracy and error thresholds** — The standard must specify the maximum acceptable rate of error for each category of judicial function, the methodology by which accuracy is tested, and the consequences — including suspension of use — when a system falls below the defined threshold.
- **Data governance and provenance** — The standard must require that training data be documented, lawfully sourced and representative of the diversity of courts and litigants the system will serve, with clear records of when and how the data was collected and updated.
- **Human oversight and override** — The standard must define, for each category of AI function, the point at which a human must review or approve an output, and must guarantee that any AI-assisted output can be overridden without penalty or procedural barrier.
- **Data and API standards** — The standard must require AI systems to conform to open data formats and API specifications, ensuring that courts can migrate between providers without loss of data or functionality. Without this, the approval of one system becomes, in effect, a permanent procurement decision and the judiciary will be locked in with the vendor.

This list is not exhaustive. As AI technologies evolve and courts gain operational experience, additional specifications and standards will be needed. Specifications and standards will also evolve with technology. To facilitate evaluation against such standards, it could be accompanied by the following:

- **Common benchmark datasets.** For categories of AI applications used in courts, such as OCR, translation, anonymisation, or case classification, the judiciary could develop shared reference datasets against which competing systems must be evaluated. Vendors seeking approval would report performance on standardised metrics against these common datasets. This enables objective comparison, reduces duplication of evaluation effort across courts, and gives vendors a clear target to meet.
- **Codification of administrative benchmarks.** AI also creates an opportunity to articulate what good performance looks like for judicial administration itself. Many administrative functions, such as scheduling, workload allocation, or cause-list management, are today guided by convention rather than documented standards. Before any AI system can be evaluated on such tasks, the judiciary must first define the objective: is a scheduling system meant to reduce adjournments, balance judicial workload, or minimise litigant waiting time? Developing these benchmarks would not only enable AI evaluation, it would improve transparency and consistency in judicial administration more broadly.

Who should create these specifications and standards?

Setting standards and procuring technology must be kept institutionally separate. The body that evaluates and selects AI vendors cannot also be the body that defines what those vendors must demonstrate — that would be akin to asking a contractor to write the building code their own construction must meet. The value of a standards body lies precisely in its neutrality: it has no stake in which vendor wins, no procurement budget to defend, and no implementation timeline to protect.

This is the model that made India's payments revolution possible. NPCI did not build PhonePe or Google Pay — it wrote the specification and stepped back. The standard was shared; the market built everything else.

For AI in courts, the Indian Courts Service proposed under Recommendation 1 will be well positioned to plan, procure and manage technology. But the standards against which it procures — what any AI system must demonstrate before approval — must be set by a separate, neutral institution.

CoRE-AI, an expert institution envisioned in the Draft Regulations, is well placed to play such a role. Its role should remain focused on neutral roles such as research, setting specifications and standards, organising conferences etc.

CoRE-AI's neutrality is its most important asset. The Draft Regulations should protect it by ensuring CoRE-AI has no role in procurement or service delivery — and should mandate it to develop and maintain the minimum specifications that the approval mechanism currently lacks.

CONCLUSION

PUCAR respectfully submits that the two issues addressed in this submission — institutional architecture and standards — are foundational to the success of the Draft Regulations. Getting them right at this stage will determine whether the governance framework enables responsible, effective and equitable AI adoption across India's judiciary, or inadvertently creates new forms of fragmentation, inconsistency and institutional dependency.

We respectfully urge the Honourable Court to consider the following specific steps:

First, establish a permanent Indian Courts Service — a specialised organisation dedicated to judicial administration and technology management, operating under judicial oversight. Rather than creating a new institutional body for each new technology, a single permanent organisation capable of absorbing the functions currently held by the e-Committee alongside the new AI governance mandate is more durable, more accountable, and more effective. This structure should be guided by three principles: optimal use of judicial time, functional separation between governance and execution, and separation of powers to ensure that those who set standards do not procure, and those who procure do not audit.

Second, mandate the development of minimum specifications and standards against which AI systems must be evaluated before approval — including for explainability, accuracy, data governance, human oversight, interoperability and training. Entrust CoRE-AI, as a neutral

expert body with no role in procurement or service delivery, to develop, maintain and periodically revise these standards.

Both recommendations are designed to be durable — built to evolve as technology changes and as courts gain experience, rather than to be replaced by the next wave of reform. PUCAR would be honoured to support the Honourable Court in developing either of them further.