

AI's Positive Role in Indian Law and Technology

Naveen Yadav¹, Bindu Thakral²

²School of Engineering & Technology, Sushant University, Gurugram, India

¹School of Law, Sushant University, Gurugram, India

binduthakral@sushantuniversity.edu.in

naveen.2501lb008@sushantuniversity.edu.in

Abstract - Artificial intelligence offers practical assistance with legal research, translation and court records, but its value depends on whether institutions can detect and correct consequential errors. This paper examines how regulation could support those uses in India while preserving judicial independence, procedural fairness and access to justice. It uses doctrinal analysis of constitutional provisions, legislation, policy documents and the Supreme Court's draft Regulations for Use of Artificial Intelligence in Courts, 2026, with selected commentary and a limited comparison with the European Union and Japan. The analysis finds support for assistive applications alongside unresolved questions about verification exceptions, implementation capacity and correction of errors within permitted uses. It proposes clearer approval conditions, source-linked outputs, supported human review and procurement terms that allow institutions to test and monitor performance. A future evaluation is outlined for judgment summarisation and citation verification, including comparison with manual work; no original performance results are claimed. The paper concludes that regulation can support useful development by making the conditions of adoption predictable and the responsibility for error enforceable, although the benefits of that approach require empirical evaluation.

Key Words: artificial intelligence, Indian law, responsible innovation, judicial independence, access to justice, data protection.

1. INTRODUCTION

Consider a litigant trying to understand an order written in an unfamiliar language. A preliminary translation could help that person identify the next procedural step, but a mistranslated date or condition could also cause a missed hearing. A lawyer using a judgment summary faces a related problem: the summary may locate the relevant issue yet omit an exception that changes the result. These hypothetical examples show why the positive potential of artificial intelligence (AI) in legal work must be examined at the level of particular tasks. Translation, document search and filing assistance are among the uses contemplated by the Supreme Court's 2026 draft regulations. [21, reg 19] Their value depends on how reliably they help people complete those tasks and how quickly mistakes can be corrected.

Regulation influences the decisions made before a legal AI service reaches its users. A developer needs to know which records may be processed, what testing an institution expects and which functions require approval. The purchasing institution must decide who checks outputs and who handles complaints. Unclear answers can delay a useful product; an elaborate approval process can also consume resources without detecting its principal defects. This paper therefore treats regulatory design as part of the conditions for innovation. Its argument is that requirements should address identifiable failures in a proposed use. For example, a translation tool should be assessed for errors in legal meaning, while a citation tool should be assessed for both the existence and relevance of its authorities.

The Indian constitutional context supplies the standards for evaluating these choices. Article 14 guarantees equality before the law and equal protection of the laws; Article 21 protects life and personal liberty. Article 39A, a Directive Principle of State Policy, directs the State towards equal justice and free legal aid. [6] Their different constitutional positions should be recognised when connecting them to technology policy. The argument advanced here is that an AI-assisted procedure should improve access while preserving fair opportunities to understand and contest the material used in a proceeding. A service that helps a litigant read an order may advance that objective. A procedure that leaves the litigant unable to challenge a consequential error would require a different assessment, however convenient it appears to its operators.

The Supreme Court's consultation draft of 3 June 2026 provides a detailed setting in which to examine this argument. Regulations 16 and 17 encourage responsible adoption and innovation, while regulations 4 and 20 reserve judicial authority to human decision-makers and restrict specified applications. [21] These provisions bring encouragement and restraint into the same instrument. The instrument remains a draft for the purposes of this analysis: in *Vijay Ghanshyam Gadiya v Union of India & Anr*, decided on 2 September 2026, the Court stated that the regulations had not yet been finalised. [22, para 6] References below to approvals, committees and disclosure duties describe proposed arrangements. They should not be read as a claim that the draft itself already imposes binding obligations across all courts.

Three research questions guide the discussion. Which uses of AI offer defensible public benefits in Indian law and related technology? How far do the selected Indian legal and policy materials support those uses while addressing foreseeable harm? What changes would make the emerging framework more effective without imposing the same burden on every application? The paper advances a conditional argument: regulation can encourage innovation when its requirements are tied to the purpose, consequences, and verifiability of a system. Its contribution is a practical reading of the court draft alongside wider governance materials, with particular attention to verification, procurement, institutional competence, and redress. It does not claim that AI has already reduced judicial delay or that the proposed framework has been empirically proven to do so.

2. LITERATURE REVIEW

Indian commentary raises a question about the institution that should oversee AI. Tulika Avni Sinha discusses alternative regulatory approaches, while Adya Madhavan argues for a dedicated independent authority. [12, 13] These contributions draw attention to the choice between coordinating existing institutions and creating a new regulator. Each approach leaves a practical question unanswered unless its responsibilities are specified. A new authority would still need expertise, staff and procedures for dealing with sector-specific problems. Coordination would need a way to resolve

overlapping mandates. For legal services, the proposed regulator's name matters less than the allocation of identifiable tasks: approving an application, investigating an incident and securing correction for an affected person. This paper uses those tasks to assess institutional design.

NITI Aayog's two responsible AI documents, published in 2021, connect principles such as safety, equality, privacy, transparency and accountability with questions of implementation. [2, 3] The Ministry of Electronics and Information Technology (MeitY), in its November 2025 announcement of the India AI Governance Guidelines, also places innovation and risk mitigation within a common policy framework. [16] Together, these materials support an account of trust that extends beyond a model's technical performance. They do not, by their publication alone, establish that a deployment meets the stated principles. For a court translation service, for instance, accountability would require an assigned reviewer and an available correction procedure. The analysis below follows this movement from a policy principle to a specific institutional duty, and asks what records would allow compliance with that duty to be assessed.

Yuthika Bhargava's 2020 report discusses industry apprehension that premature restrictions could obstruct development, including interest in regulatory sandboxes. [1] Aroon Deep's January 2025 account highlights the case for coordinated governance across ministries. [11] Read together, these sources suggest that uncertainty and fragmented requirements may burden innovation as much as strict substantive rules. However, their dates matter: the January report cannot substitute for the later November 2025 guidelines. A further question follows from their arguments. Can a small developer discover the applicable requirements before investing in a product, and can a public institution explain its approval decision? Predictable procedures are therefore part of the innovation question, rather than merely administrative detail.

Empirical scholarship supplies a necessary caution about legal reliability. Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E. Ho document legal hallucinations in tested language models answering questions about American case law. [7] Their study demonstrates why confident language and recognisable legal terminology cannot establish accuracy. It also shows why users should question a model's acceptance of a false premise. The study concerns particular models, tasks, and a foreign legal setting; its numerical findings should not be presented as measured error rates for Indian courts or current AI products. Its contribution here is narrower but valuable: legal AI evaluation should test whether an authority exists, whether the stated proposition matches it, and whether the system recognises an unsupported question.

Comparative materials complicate the assumption that enforceable safeguards necessarily oppose innovation. The European Union (EU) AI Act includes specified justice-related systems within its high-risk framework and provides for regulatory sandboxes. [8, Annex III 8(a), art 57] Japan's 2025 AI legislation, by contrast, gives prominence to research, development, utilisation, information gathering, and international cooperation. [15, arts 16–17] Sanhita Chauriha's article interprets the Japanese approach as particularly supportive of innovation. [14] These models offer different institutional emphases, not a controlled experiment establishing which country has found the best policy. India can learn from both without copying either wholesale. The appropriate comparison concerns particular functions—testing, documentation, support, and supervision—rather than broad labels such as strict or permissive regulation.

A further strand places AI within inclusion and international cooperation. The Global Digital Compact, adopted with the Pact for the Future in September 2024, supplies a framework for digital cooperation. [9] The commentary by Neethu Rajam and Krishna Ravi Srinivas explains its shared-goal character rather than treating it as domestic legislation. [10] Mansukh Mandaviya and Gilbert F. Hounbo emphasise social justice and inclusive work in their February 2026 article. [20] These perspectives widen the meaning of positive potential. A service cannot be considered successful solely because well-resourced users find it convenient. Its distributional effects also matter. In the legal sector, that means examining language barriers, disability access, support for less experienced users, and opportunities for workers to acquire new skills.

The selected literature establishes why reliability, institutional coordination and inclusion matter, but these aims still require decisions about everyday legal work. The Supreme Court draft permits a focused examination of those decisions: who verifies an output, which uses need approval and how an affected person obtains correction. [21] This paper assesses those arrangements against the policy arguments and reliability research reviewed above. Its contribution lies in connecting particular provisions to their implementation requirements. The recommendations therefore concern approval records, verification duties, procurement and remedies; they do not depend on claiming that responsible AI is an unexplored subject.

3. RESEARCH METHODOLOGY

3.1 Research design and scope

The study adopts a doctrinal method: it reads relevant legal provisions together, identifies their scope and status, and examines points requiring interpretation or reform. Its principal subject is AI used for legal information, professional assistance and court administration in India. Data protection and technology policy enter the analysis where they affect these functions. A limited comparison with the European Union and Japan helps identify alternative ways to organise testing and public oversight. The comparison concerns selected provisions and policy choices; it does not measure national innovation performance. Sources were checked up to 5 September 2026. The resulting conclusions concern the legal framework and its design, with practical effectiveness left for future empirical evaluation.

3.2 Materials and selection

Materials were selected for their relevance to legal reliability, institutional responsibility, access to justice or the conditions for developing AI services. They include the Constitution, statutory texts, the Supreme Court's consultation draft and subsequent decision in *Vijay Ghanshyam Gadiya*, official policy publications, selected newspaper extracts and published research on legal hallucinations. The separate commencement notification is read alongside the data-protection Act and Rules. [2, 3, 4, 5, 6, 7, 12, 13, 16, 17, 18, 21, 22] Primary texts support statements about legal provisions; newspaper extracts identify attributed arguments in the public debate. The empirical article supplies evidence about the systems and tasks it tested. This purposive selection permits close examination of the chosen legal issues, but does not constitute an exhaustive review of AI governance literature.

3.3 Source checking and legal status

Source checking distinguishes a document's legal status from its policy argument. The court draft was compared with the official publication, and its status checked against the subsequent Supreme Court decision. [21, 22] The comparative

discussion uses the enacted EU instrument and the Japanese government's English translation of its AI legislation. [8, 15] For data protection, the Act, Rules and separate commencement notification are read together; publication of the Rules is not treated as evidence that every provision of the Act is already operative. [4, 17, 18] Earlier commentary is assessed at its publication date. The discussion identifies proposed reforms separately from requirements stated in the legal materials.

3.4 Analytical framework

The materials are compared through five questions: what task the system performs; what benefit is claimed; whose interests may be affected; which safeguard addresses the relevant failure; and who remains answerable. Particular attention is given to whether an output can be checked against an authoritative source before it affects a person. A searchable index, a preliminary translation, and a proposed assessment of witness credibility therefore require different analysis. The findings section reports the relationships and tensions identified in the texts. The discussion then develops recommendations from those findings. Where the paper proposes an approval route, metric, or technical practice beyond the sources, it presents that choice as a recommendation rather than as a binding requirement.

3.5 Integrity and limitations

No interviews, surveys, field observations, or controlled comparisons of AI systems were conducted for this paper. There is consequently no original dataset from which to estimate accuracy, cost savings, or changes in case disposal. Illustrations describe possible workflows; they are not reports of deployments observed by the author. Published evidence from other jurisdictions is used with its contextual limits intact. The selected newspaper extracts are incomplete, so arguments are attributed only to material actually available. These limits make a doctrinal approach appropriate, but restrict the strength of causal conclusions. The paper can explain why a safeguard is legally or institutionally desirable; it cannot establish how much benefit that safeguard will produce in practice without subsequent testing.

4. RESULTS: FINDINGS FROM THE LEGAL AND POLICY ANALYSIS

4.1 Useful assistance has a concrete but conditional place

The court draft identifies a substantial range of potentially permissible activities, including translation, transcription, research, citation verification, document summarisation, filing assistance, and accessibility services. These uses are subject to approval and relevant supervision or verification conditions. [21, reg 19] An official February 2026 account describes Supreme Court Vidhik Anuvaad Software (SUVAS) translation work and states that the Supreme Court Portal for Assistance in Court Efficiency (SUPACE) remained experimental rather than in regular judicial deployment. [19] The finding is therefore a supported scope for assistance, not proof of system-wide performance gains. These applications share an important feature: their output can often be compared with an existing document, recording, or procedural requirement. Nevertheless, verifiability varies by task. A translated order requires attention to legal meaning, whereas the insertion of a case number into a notice may require a simpler comparison with a verified record.

Figure 1 summarises the proposed assistive architecture. Figure 2 sets out a verification and correction workflow developed from these findings; both are conceptual proposals, not evaluated systems.




Fig -1: Proposed architecture for source-linked legal AI assistance. Authors' synthesis of Sections 4–5; not a deployed system.

4.2 The draft favours adoption but imposes substantive limits

Regulations 16 and 17 of the court draft support responsible adoption, while regulation 20 prohibits specified applications, including risk scoring and prediction or profiling of future behaviour in court processes. Regulation 16 separately excludes dispute-outcome prediction. [21, regs 16–17, 20] The requirement for human involvement does not override those specific prohibitions. Regulation 20(1)(c), concerning adjudication and sentencing without mandatory human involvement, should consequently be read with the reservation of judicial authority in regulation 4 and the limits in regulation 16. Reading the provisions together prevents a general reference to supervision from becoming permission for an otherwise prohibited function.

Verification presents a separate interpretive problem. Regulation 8(3) first requires reasonable care to check an advisory output. Its first proviso allows the responsible officer to dispense with verification by recording reasons; its second provides verification on a class basis for exclusively administrative tools whose reliability the AI Secretariat certifies. Regulation 19 expressly requires accuracy review and certification for transcription, human verification for translation, and human review before action on administrative authenticity or fraud-detection outputs. [21, regs 8(3), 19(1)(b)–(c), (h)] The draft does not expressly resolve the relationship between the general dispensation and these particular safeguards. Clarifying that relationship would help officers know which checks remain necessary and developers know which review features their products must support.

4.3 Data protection and evidence require separate legal inquiries

Data protection and evidentiary proof concern different questions. In its enacted form, section 17(1)(b) of the Digital Personal Data Protection Act, 2023 exempts processing necessary for specified judicial, quasi-judicial, regulatory or supervisory functions from designated provisions. The exemption retains section 8(1), concerning the Data Fiduciary's responsibility for compliance, and section 8(5), concerning reasonable security safeguards. [4, ss 8(1), 8(5), 17(1)] Commencement must be considered separately. Notification GSR 843(E), dated 13 November 2025, places sections 8 and 17 in the group commencing eighteen months after Gazette publication. Those provisions had therefore not commenced by this paper's cut-off of 5 September 2026. [17,

cl (c)] The 2025 Rules also provide phased commencement. [18, r 1] The discussion of the judicial-function exemption accordingly explains the enacted framework rather than an exemption already operative on that date.

Under the Bharatiya Sakshya Adhiniyam, 2023, sections 61–63 address electronic records and their proof, including the conditions and certification mechanism in section 63. [5] An AI summary of a recording does not, by itself, establish that the underlying event occurred as described. The court must still examine the relevant record and applicable proof requirements. Likewise, the authority to process a record does not settle its authenticity, admissibility or weight. Keeping these questions distinct helps identify the limited contribution of an AI tool: it may assist with locating, organising or describing material, while the legal assessment of that material remains necessary.

4.4 Accountability is becoming more specific, but delivery remains unresolved

The court draft assigns institutional work through impact assessments, controlled testing, registers, audits, incident reporting, procurement and training provisions. [21, regs 35–51] The importance of that work is illustrated by Vijay Ghanshyam Gadiya. The Supreme Court identified problems including non-existent cases and incorrect propositions attributed to real authorities. It set aside the customs penalty order and the High Court order confirming it, and directed fresh determination by another officer of the same rank. [22, paras 4–8] The immediate lesson for deployment is that a plausible legal output requires checking before it enters official reasoning. Assigning that responsibility to an officer is only the starting point. Institutions must also provide access to source material, enough time to review it and a procedure for stopping use when the output cannot be substantiated.

Redress requires attention even where the initial use is permitted. Regulation 52 addresses harm from a use prohibited by regulation 20, while regulation 53 preserves other legal remedies. [21] Suppose an approved transcription tool records an admission where the speaker actually denied a fact. The litigant needs a practical means of comparing the transcript with the recording and having the record corrected. The availability of another legal remedy does not by itself explain how that routine correction should proceed. The draft would benefit from an express procedure covering material errors in permitted applications, identifying where to complain, who reviews the source and how the corrected record reaches those relying on it. This proposal concerns the accessibility of correction; it does not suggest that existing remedies disappear.

5. DISCUSSION

5.1 Regulation should follow consequences and verifiability

The findings support regulation tied to the actual use of a system. An application that identifies missing attachments should not automatically face the same approval burden as one analysing material that could influence a person's liberty. Conversely, the label “administrative” should not become an exemption from scrutiny. Scheduling and defect scrutiny may affect whether an urgent matter is heard or a filing proceeds. The proposed distinction should therefore consider consequences, personal data exposure, reversibility, and opportunities for correction. Draft regulations 34(4) and 46(7) already envisage expedited approval for specified administrative tools that do not involve parties' personal data, do not affect adjudicatory functions, and resemble approved tools. [21, regs 34(4), 46(7)] Their conditions should be retained when designing a simpler route.

An approval record should make the permitted function clear to both the institution and its supplier. It could identify the task, data sources, intended users, verification steps and officer responsible for incidents. Consider a tool approved to explain filing requirements. If a later update begins predicting whether an individual claim will succeed, its function has changed in a way that calls for fresh scrutiny and engages the draft's restriction on outcome prediction. [21, reg 16] A change to display colours would ordinarily raise no comparable issue. Reassessment should therefore turn on a material change in purpose, data use or performance risk. This would give institutions a basis for controlling extensions of use while allowing routine improvements to proceed.

This approach also faces a substantial objection: approval and documentation duties may favour established suppliers and delay a small developer with a useful product. The answer cannot be to assume that every safeguard pays for itself. Institutions should publish clear submission requirements, reuse technical assessments where appropriate and explain refusals by reference to the proposed function and its risks. The draft's expedited route supplies a starting point for eligible administrative tools, while its conditions limit the cases in which that route is available. [21, regs 34(4), 46(7)] A pilot should therefore record approval time and compliance effort as well as output quality. If a requirement repeatedly consumes resources without addressing a relevant failure, the institution should review its design.

5.2 Verification must be a supported professional task

Verification should be designed around the legal mistake a reviewer needs to detect. A research answer that links only to a long judgment leaves the reviewer to search for its supposed support; identifying the relevant paragraph would make the check more direct. A summary should distinguish a party's submission from the court's holding, since confusing the two can reverse the apparent meaning of a decision. A translation should remain attached to its original, with important dates, obligations and legal terms checked by an appropriate reviewer. These are proposed features for reviewable assistance. Where the available record does not support an answer, the system should indicate that limitation and refer the unresolved question for human attention.

The first proviso to regulation 8(3) should be narrowed so that a recorded dispensation cannot remove the specific review requirements for uses governed by regulation 19. The second proviso already provides a more defined route for reliable administrative tools certified by the AI Secretariat. [21] That route could accommodate routine automation, subject to monitoring for recurring mistakes. Material legal propositions, the existence and meaning of authorities, and translations capable of affecting rights should receive direct verification. A court should also establish who performs each check and how an unresolved discrepancy reaches a supervisor. Requiring verification without allocating time and access to the original material risks producing a formal sign-off with little substantive examination. The proposed distinction would concentrate scarce review capacity on errors that can alter a proceeding while permitting appropriate sampling of routine administrative outputs.




Fig -2: Proposed verification and correction workflow. Authors' synthesis of Sections 4.4, 5.2 and 5.4; not a binding procedure.

5.3 Technical development, procurement, and access must work together

A legal research service could begin with an institutionally approved collection of judgments and statutes. Its answers should identify the passages on which they rely, and its interface should make an unsupported citation visible as an unresolved item. Restricting the source collection would still leave possible failures: an omitted later decision, an incomplete search or a mistaken reading of the retrieved text. Evaluation should separate these errors. For each legal proposition, a reviewer could record whether the cited authority exists, whether the passage supports the proposition and whether a material qualification has been omitted. The time needed to complete those checks should also be measured. This proposed evaluation would test the usefulness of the complete service, including the work it leaves for the lawyer.

Purchasing arrangements should allow that evaluation to continue after installation. Regulation 46 of the draft addresses matters including documentation, data use, liability and restrictions on retraining with court data. [21] A contract could give those provisions practical form through agreed retention periods, access controls, advance notice of material updates and cooperation when an incident is investigated. It should also explain how the institution retrieves its records when the service ends. These requirements have implications for competition: requiring one supplier to provide every judicial function may make entry harder for a developer with a useful specialist tool. Separate purchases for translation, search or accessibility deserve consideration where common security and compatibility requirements can be met. Whether that approach lowers costs would need to be tested in procurement practice.

The people least familiar with legal procedure should be included when assessing a proposed legal-aid assistant. Can a user understand its explanation in the relevant language, recognise its limits and reach a human adviser when the problem remains unresolved? Evaluation should also include disability access and assistance for people unable to use the digital service independently. Mandaviya and Hounbo's emphasis on social justice supports this attention to the distribution of benefits. [20] A pilot should record the extra work created by corrections, referrals and staff support

alongside any time saved. Keeping an assisted or non-digital route available would also permit comparison with the new service. Improved access should be demonstrated through these users' ability to complete the task, rather than inferred from the availability of an application.

5.4 Testing and remedies should build evidence for further adoption

A future pilot could evaluate judgment summarisation and citation verification using a fixed pack of official judgments and statutory provisions. Each system should receive the same task and source-access conditions; model versions, dates, prompts, settings and outputs should be retained. A legally qualified reviewer should prepare the expected answers from the original materials before assessing the outputs. The scoring criteria should distinguish an invented citation, a real authority that does not support the proposition, and a summary that omits a qualification material to the result. Questions with misleading premises or insufficient source support should be included. A second reviewer could examine disputed scores, with disagreements recorded. This is a proposed study; no such experiment or independent scoring has been completed for this paper.

A manual baseline is necessary to assess the claimed benefit. The pilot should compare a lawyer's work using ordinary source access with a comparable task completed using AI assistance followed by verification. Where practicable, reviewers should not know which system produced the answer they score, and task order should be varied to reduce learning effects. Total completion time should include checking and correcting outputs. Results should report the number and seriousness of errors alongside time, since faster drafting may simply move work into review. Failure on consequential assertions should be investigated before any expansion of use. For court-supervised controlled testing, draft regulation 36(3) would prohibit test outputs from influencing actual judicial decisions. [21] These design choices would allow the proposed benefits to be examined without exposing litigants to untested results.

Correction procedures should be settled before wider adoption. Regulation 43 of the draft already proposes timely, accessible information for parties when a tool materially assists case management, document analysis or judicial administration in a way that may affect their proceedings. [21, reg 43(1)–(2)] An operational procedure should then explain how a party flags an inaccurate record, who compares it with the source, and how the correction is communicated to those relying on it. Relevant versions and verification records should be preserved, with urgent hearing or filing errors receiving prompt human attention. Anonymised incident records could reveal repeated failures requiring retraining, restricted use or withdrawal, subject to controls on sensitive information. This would connect disclosure and correction to decisions about continued deployment, giving an institution evidence on whether its service is becoming dependable.

6. CONCLUSION

AI's positive role in Indian law should be assessed through the legal work it helps people complete and the responsibility retained when it fails. The materials examined support the development of translation, research, accessibility and record-management services whose outputs can be checked against reliable sources. They also show why permission to use a system requires more than confidence in its general capabilities. The Supreme Court's June 2026 draft connects encouragement of adoption with restrictions and oversight, but its verification exceptions need clarification and its correction arrangements would benefit from more explicit treatment of mistakes within permitted uses. A workable

framework must also account for the staff, expertise and time required to carry out those duties. Predictable approval procedures and testable procurement conditions could help developers invest in useful applications, while common requirements could reduce repeated compliance work. Such advantages remain propositions to be evaluated, not measured outcomes of this study. The proposed pilot would examine them by comparing legal accuracy and total task time with a manual baseline. The case for regulating to innovate is therefore conditional: regulation should make beneficial assistance easier to develop, its performance easier to examine and its errors easier to remedy. The standard of success is an improvement in legal service that remains accountable to the people using it.

ACKNOWLEDGEMENT

I express my sincere gratitude to Dr. Bindu for her guidance, supervision and constant support throughout the preparation of this research paper. As my mentor, she helped me develop my ideas and approach the research with greater clarity. Her suggestions, encouragement and willingness to help at every stage have been invaluable. I am deeply thankful for the time and effort she devoted to supporting my work.

REFERENCES

Online sources accessed 5 September 2026. Newspaper commentary is based on supplied extracts.

- [1] Y. Bhargava, "Are regulations a boon or bane for tech?," The Hindu, updated Jan. 27, 2020. Supplied extract.
- [2] NITI Aayog, "Responsible AI #AIforAll: Approach document for India, Part 1—Principles for responsible AI," Feb. 2021. <https://www.niti.gov.in/sites/default/files/2021-02/Responsible-AI-22022021.pdf>
- [3] NITI Aayog, "Responsible AI #AIforAll: Approach document for India, Part 2—Operationalizing principles for responsible AI," Aug. 2021. <https://www.niti.gov.in/sites/default/files/2021-08/Part2-Responsible-AI-12082021.pdf>
- [4] Government of India, Digital Personal Data Protection Act, 2023, Act No. 22 of 2023, secs. 8 and 17(1), 2023. <https://www.meity.gov.in/static/uploads/2024/06/2bf1f0e9f04e6fb4f8fef35e82c42aa5.pdf>
- [5] Government of India, Bharatiya Sakshya Adhiniyam, 2023, Act No. 47 of 2023, secs. 61–63 and Schedule, 2023. https://www.mha.gov.in/sites/default/files/2024-04/250882_english_01042024_0.pdf
- [6] Legislative Department, Government of India, Constitution of India, 2024 ed., arts. 14, 21, and 39A, 2024. <https://cdnbbsr.s3waas.gov.in/s380537a945c7aaa788ccfcd1b99b5d8f/uploads/2024/07/20240716890312078.pdf>
- [7] M. Dahl, V. Magesh, M. Suzgun, and D. E. Ho, "Large legal fictions: Profiling legal hallucinations in large language models," Journal of Legal Analysis, vol. 16, pp. 64–93, 2024. <https://academic.oup.com/jla/article/16/1/64/7699227>
- [8] European Parliament and Council of the European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), Annex III, point 8(a), and art. 57, 2024. <https://eur-lex.europa.eu/eli/reg/2024/1689>
- [9] United Nations, "Global Digital Compact," Annex I to the Pact for the Future, adopted Sep. 22, 2024. <https://www.un.org/digital-emerging-technologies/global-digital-compact>
- [10] N. Rajam and K. R. Srinivas, "Global Digital Compact: Advancing digital innovation in a sustainable fashion," The Hindu, updated Oct. 13, 2024. Supplied extract.
- [11] A. Deep, "Inter-ministerial committee needed to strengthen AI governance, says Centre's IndiaAI Mission," The Hindu, updated Jan. 8, 2025. Supplied extract.
- [12] A. Madhavan, "India's AI investment, the budgetary push, the need for regulations," The Hindu, Feb. 25, 2025. Supplied extract.
- [13] T. A. Sinha, "The approach to regulating AI in India," The Hindu, Apr. 16, 2025. Supplied extract.
- [14] S. Chauriha, "How Japan's new AI Act fosters an innovation-first ecosystem," The Hindu, June 3, 2025. Supplied extract.
- [15] Japan, Act on Promotion of Research and Development, and Utilization of Artificial Intelligence-related Technology, Act No. 53, June 4, 2025, arts. 16–17. English translation, Japanese Law Translation. <https://www.japaneselawtranslation.go.jp/en/laws/view/5066/en>
- [16] Ministry of Electronics and Information Technology, "MeitY unveils India AI Governance Guidelines under IndiaAI Mission to ensure safe, inclusive, and responsible adoption of artificial intelligence across sectors," Press Information Bureau, Release ID 2186639, Nov. 5, 2025. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2186639>
- [17] Ministry of Electronics and Information Technology, "Notification GSR 843(E)," Nov. 13, 2025, issued under sec. 1(2) of the Digital Personal Data Protection Act, 2023, cl. (c), Gazette identifier CG-DL-E-14112025-267647. <https://www.meity.gov.in/static/uploads/2025/11/c56cae6c383460ca69577428d36828b.pdf>
- [18] Ministry of Electronics and Information Technology, Digital Personal Data Protection Rules, 2025, GSR 846(E), Nov. 13, 2025, r. 1(2)–(4). <https://www.meity.gov.in/static/uploads/2025/11/53450e6e5dc0bfa85ebd78686cadad39.pdf>
- [19] Press Information Bureau, "From digitisation to intelligence: How AI is enhancing access to justice in India," Press Information Bureau, Release ID 2226283, Feb. 11, 2026. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2226283&lang=1®=3>
- [20] M. Mandaviya and G. F. Hounbo, "AI for people, applying technology for social good," The Hindu, updated Feb. 19, 2026. Supplied extract.
- [21] Supreme Court of India, "Notice and draft Regulations for Use of Artificial Intelligence in Courts, 2026," June 3, 2026, regs. 4, 8, 16–20, 34–36, and 43–53. Consultation draft. <https://cdnbbsr.s3waas.gov.in/s3ec0490f1f4972d133619a60c30f3559e/uploads/2026/06/2026060342.pdf>
- [22] Supreme Court of India, Vijay Ghanshyam Gadiya v. Union of India & Anr., 2026 INSC 947, Sep. 2, 2026, paras. 4–8, arising out of SLP (Civil) No. 15605 of 2026. https://www.sci.gov.in/sci-pdf/?diary_no=219912026&from=latest_judgements_order&order_date=2026-09-02&type=j