



Research Article

AI-Driven Governance in India: Balancing Administrative Efficiency, Transparency, Accountability and Data Protection

Dr. E. Priyanka

Assistant Professor (Sr. Gr.), Department of Politics and Public Administration, Crescent School of Law,
B.S. Abdur Rahman Institute of Science and Technology, Vandalur, Tamil Nadu, India

Corresponding Author: * Dr. E. Priyanka

DOI: <https://doi.org/10.5281/zenodo.23317790>

Abstract

India is embedding artificial intelligence (AI) in public administration at the very moment its legal architecture for data protection, transparency and AI oversight is being rewritten. Between November 2025 and August 2026, the government issued the India AI Governance Guidelines, notified the Digital Personal Data Protection (DPDP) Rules 2025, hosted the AI Impact Summit, and began defending in the Supreme Court a constitutional challenge to the DPDP Act's amendment of the Right to Information (RTI) Act. This paper asks whether administrative efficiency, transparency, accountability and data protection can be reconciled in AI-driven governance in India. Using doctrinal analysis, policy analysis and a case study of Telangana's Samagra Vedika welfare-profiling system, it argues that the four values are not inherently opposed. The tensions arise from design choices, thin institutional capacity and broad exemptions for the State. The paper identifies four structural gaps: voluntary AI guidance without enforceable duties; consent-free State processing and wide exemptions under the DPDP Act; the narrowing of the RTI public-interest test; and the absence of a right to human review and explanation. It proposes a proportionality-based, risk-tiered framework combining statutory authorisation, impact assessments, algorithm registers, meaningful human review, independent audit and accessible redress. Legal positions are stated as at 28 September 2026.

Manuscript Information

- **ISSN No:** 2583-7397
- **Received:** 11-08-2026
- **Accepted:** 07-10-2026
- **Published:** 10-10-2026
- **IJCRM:** 5(5); 2026: 1152-1161
- **©2026, All Rights Reserved**
- **Plagiarism Checked:** Yes
- **Peer Review Process:** Yes

How to Cite this Article

E. Priyanka. AI-Driven Governance in India: Balancing Administrative Efficiency, Transparency, Accountability and Data Protection. Int J Contemp Res Multidiscip. 2026;5(5):1152-1161.

Access this Article Online



www.multiarticlesjournal.com

KEYWORDS: algorithmic governance; e-governance; transparency; accountability; data protection; DPDP Act; proportionality; India.

1. INTRODUCTION

Governance in India has long been shaped by a gap between what law promises and what the state can deliver at scale. With more than 1.4 billion people, a large informal economy and uneven administrative capacity across States, the Indian state has repeatedly turned to technology to extend its reach. Aadhaar, the Unified Payments Interface, DigiLocker and the wider "India Stack" made digital public infrastructure central to state-building. AI is the next layer: systems that classify, predict, recommend and, increasingly, decide (Wirtz et al., 2019; Kuziemski & Misuraca, 2020).

The pace of adoption is notable. The IndiaAI Mission was approved in March 2024 with an outlay of about ₹10,372 crore, and by early 2026 more than 38,000 GPUs had been provisioned, with 20,000 more announced at the India AI Impact Summit in February 2026 (PIB, 2026). Courts are experimenting with AI-assisted transcription, research and filing scrutiny (MediaNama, 2025). States have built data-linkage platforms that profile welfare beneficiaries (Amnesty International, 2024). At the same time, the legal framework is in motion. MeitY published the India AI Governance Guidelines on 5 November 2025 and confirmed that India will not, for now, enact an umbrella AI statute (MeitY, 2025; Chambers and Partners, 2025). The DPDP Rules 2025 followed on 13 November 2025, with obligations phased in until May 2027.

Governance, however, is not only a matter of efficient output. In a constitutional democracy, administrative action must be lawful, reasoned, open to scrutiny and respectful of rights. AI systems can be opaque, can encode historical bias, can disperse responsibility away from identifiable officials, and depend on large-scale collection and linkage of personal data. The central question of this paper is therefore: **how can India secure the efficiency gains of AI-driven governance without eroding transparency, accountability and data protection?**

The paper makes three contributions. First, it updates the Indian debate to September 2026, incorporating the 2025 AI Governance Guidelines, the DPDP Rules 2025, the pending Supreme Court challenge to the DPDP Act, and recent judicial

AI policies. Second, it uses the Samagra Vedika case to show how the four values interact in practice. Third, it proposes a risk-tiered, proportionality-based framework tied to concrete legal and institutional reforms. The method is doctrinal and policy-analytic, supported by one documented case study; sources include statutes, court records, government documents, peer-reviewed scholarship and investigative reports. Because litigation and implementation are ongoing, the analysis flags where positions may change.

2. CONCEPTUAL FRAMEWORK

2.1 Algorithmic discretion

Algorithmic governance refers to the use of computational systems to inform or make decisions traditionally made by public officials. Bullock (2019) and Young et al. (2019) show that AI does not remove discretion but relocates it, from street-level officials to system designers who choose training data, features, thresholds and tolerable error rates. These choices are seldom visible to the citizen affected, which is why the values below matter.

2.2 Four values and their tensions

Administrative efficiency means achieving public objectives with fewer resources, less delay and greater consistency. *Transparency* is the availability of information about how and why decisions are made; Ananny and Crawford (2018) warn that disclosure of code alone can amount to "seeing without knowing", so what counts is intelligibility to affected persons, overseers and courts. *Accountability*, following Bovens (2007), is a relationship in which an actor must explain and justify conduct to a forum that can pose questions and impose consequences; AI complicates it by spreading agency across developers, vendors, procuring departments and officials (Busuioc, 2021). *Data protection* concerns lawful, fair and limited processing of personal data, grounded in the right to privacy recognised in *Justice K.S. Puttaswamy (Retd.) v. Union of India* (2017) 10 SCC 1. Table 1 summarises what each value demands, how it typically fails, and what reconciles it with the others.

Table 1. The four values: requirements, failure modes and reconciling design

Value	What it requires of AI in government	Typical failure mode	Reconciling design
Administrative efficiency	Faster, cheaper, more consistent decisions; reach in low-capacity settings	Error costs shifted to those least able to contest them; flawed data automated at scale	Count error and appeal costs; pilot and publish accuracy before scale-up
Transparency	Disclosure that AI is used; intelligible reasons; auditor access	Trade-secret claims; undocumented procurement; RTI exemptions	Public registers; reasoned notices; tiered disclosure; contract clauses
Accountability	Named responsibility; ability to contest and obtain remedy	Diffuse responsibility; automation bias; vendor insulation	Human review with real authority; audit trails; independent audit; redress
Data protection	Lawful, minimal, purpose-limited and secure processing	Function creep through data linkage; broad State exemptions	Impact assessments; privacy-enhancing technologies; statutory limits on linkage

2.3 Proportionality as the reconciling standard

The values are often framed as zero-sum: transparency may enable gaming of fraud-detection systems, minimisation may reduce model accuracy, and human review may slow processing. This paper takes a more nuanced position. Some

tensions are real but manageable through calibrated design, and several values are complementary: explainable systems are easier to audit, audited systems earn trust, and data minimisation reduces breach exposure. The reconciling standard in Indian public law is proportionality, set out in

Puttaswamy (2017) and applied in the Aadhaar judgment, Puttaswamy (Aadhaar) (2019) 1 SCC 1, and in Anuradha Bhasin v. Union of India (2020) 3 SCC 637. State action that limits rights must rest on law, pursue a legitimate aim, be

rationally connected to it, be necessary in that no less intrusive means would serve, and strike a fair balance between benefit and infringement. Figure 1 places this test at the centre of the four values.




Figure 1. The four values and the proportionality standard that reconciles them (author's framework)

3. India's Policy and Legal Landscape

3.1 Policy: principles first, statute later

India's approach has been principles-based and innovation-oriented. NITI Aayog's 2018 national strategy and 2021 *Responsible AI* documents set non-binding principles. The MeitY *India AI Governance Guidelines* (2025) build on seven "sutras": trust as the foundation, people first, innovation over restraint, fairness and equity, accountability, understandable by design, and safety, resilience and sustainability. They organise recommendations under six pillars across three domains (enablement, regulation and oversight), propose an AI Governance Group, a Technology and Policy Expert Committee and an AI Safety Institute, and rely on voluntary measures, existing law and sectoral regulators rather than a new statute (MeitY, 2025; Chambers and Partners, 2025). The Reserve Bank of India's FREE-AI framework (2025) applies similar principles to finance through 26 recommendations.

The judiciary has moved faster on binding internal rules. The Kerala High Court's July 2025 policy confines AI to an assistive role and prohibits its use to arrive at findings, reliefs, orders or judgments (Jhalnia, 2025). The Supreme Court's white paper on AI and the judiciary (late 2025) stresses human verification of every AI output, and the Chief Justice reconstituted the Court's AI Committee in December 2025 (MediaNama, 2025; Prasar Bharati News Services, 2025).

3.2 Law: existing instruments and their limits

No single statute governs AI in government. Table 2 lists the instruments that do most of the work and the limits of each, and Figure 2 places them on a timeline alongside the European Union's AI Act as a comparator.

Table 2. Principal instruments governing AI in Indian public administration (as at September 2026)

Instrument	Status	Relevance	Principal limitation
Constitution: Arts 14, 19, 21; Puttaswamy (2017, 2019)	Binding	Legality, non-arbitrariness, proportionality	Enforced after the fact, through litigation
RTI Act 2005	In force; s 8(1)(j) amended by DPDP Act s 44(3)	Access to records; proactive disclosure (s 4)	Exemptions; unclear reach to vendor-held information; amendment under challenge
Information Technology Act	In force	Security duties; cyber	Not designed for algorithmic

2000		offences	decision-making
DPDP Act 2023 and Rules 2025	Board constituted Nov 2025; consent-manager rules from Nov 2026; core duties from May 2027 (computed)	Purpose limitation, security, breach notification, data-principal rights	Consent-free State processing (s 7(b)); s 17(2) exemptions; no automated-decision right; under challenge
MeitY India AI Governance Guidelines (Nov 2025)	Voluntary	Seven sutras; proposed institutions; risk-mitigation guidance	Non-binding; no enforcement or individual remedy
Sectoral and judicial instruments (RBI FREE-AI 2025; Kerala HC policy 2025; SC white paper 2025)	Varies	Sector- or court-specific guardrails	Fragmented; uneven coverage across regulators and High Courts




Figure 2. Timeline of key Indian and EU regulatory milestones, 2023-2028. Sources: MeitY (2025); Shardul Amarchand Mangaldas (2025); PIB (2026); Cloud Security Alliance (2026). Dates for DPDP Rules provisions are computed from the notification of 13 November 2025.

Two features of the timeline matter. First, the DPDP Rules bring into force only some provisions immediately (including the constitution of the Data Protection Board); the rule on State processing (Rule 5) and the government's power to call for information (Rule 23) take effect eighteen months after notification, on a computed date of 13 May 2027 (Shardul Amarchand Mangaldas, 2025; ddpdprules.org, 2026). Second, section 44(3), which amends the RTI Act, is already operative even though its constitutionality is under challenge. State AI deployments are therefore expanding during a period in which much of the data protection regime is not yet enforceable.

4. Administrative Efficiency: Promise and Limits

4.1 The efficiency case

The efficiency argument is strong in India. Pendency in courts, delay in benefit delivery and manual scrutiny of filings impose real costs, and these fall hardest on the poor. AI can triage grievances, flag anomalous transactions, translate documents across India's many languages and answer routine queries at scale. In the judiciary, live speech-to-text is used in Constitution Bench hearings, and a filing-scrutiny tool piloted with IIT Madras detects defects such as missing annexures and formatting errors, reducing registry workload (MediaNama, 2025). Where AI reduces discretionary encounters between citizens and officials, it may also reduce petty corruption.

Telangana's government made exactly this argument in defence of its welfare system, saying that earlier in-person verification was discretionary, opaque and prone to misuse (Al Jazeera & Pulitzer Center, 2024).

4.2 Case study: Samagra Vedika

Samagra Vedika, built for the Telangana government by Posidex Technologies, draws on more than thirty State databases to build "360-degree" profiles of residents and uses entity resolution, a machine-learning technique that decides whether records refer to the same person, to verify welfare eligibility and flag suspected fraud (Oxford Human Rights Hub, 2026; Amnesty International, 2024). Officials are required to consult it before approving applications. A 2024 investigation by Al Jazeera and the Pulitzer Center's AI Accountability Network reported that thousands of families lost subsidised food entitlements after the system attributed assets, such as cars, to the wrong people (Al Jazeera & Pulitzer Center, 2024). Amnesty International's independent technical analysis identified how such record-matching errors arise and raised concerns about social security, privacy and redress, noting that opacity and the regulatory vacuum made independent investigation very difficult (Amnesty International, 2024). Comparable platforms are reported to be in use or development in other States (Oxford Human Rights Hub, 2026).

The case shows that efficiency is not neutral. A system tuned to catch fraud will accept some false positives, and these fall on those least able to contest them, a pattern documented in other welfare automation regimes (Eubanks, 2018) and in the Aadhaar authentication literature (Khera, 2019). Amnesty's researchers add that the business case for such tools presumes a large volume of fraud, which in practice can increase incorrect matches. Three further limits follow. Output quality depends on data quality, and administrative records often contain errors in names, addresses and linkages. Historical bias can be embedded, since models trained on past enforcement or welfare records reproduce patterns shaped by caste, religion, gender or region. And measured efficiency often ignores downstream costs: time saved at the point of decision can be offset by appeals, litigation and lost trust when decisions are wrong. A credible efficiency claim should therefore count error and redress costs, and should be supported by published accuracy data, which are largely absent in India.

5. Transparency: From Disclosure to Intelligibility

5.1 The transparency deficit

Transparency in AI governance has three layers: disclosure that an automated system is being used; disclosure of its purpose, data sources and general logic; and an individualised explanation of a specific decision. Indian practice is thin on all three. Procurement of AI tools is often not publicly documented, vendors may assert trade secrecy, and citizens are frequently unaware that an algorithm influenced a decision about them. The MeitY Guidelines' sutra of "understandable by design" recognises the problem but is voluntary. Wachter et al. (2017) show that even the GDPR does not clearly confer a legally robust right to explanation, which suggests that India needs express statutory duties rather than reliance on general principles.

5.2 The RTI Act under strain

The RTI Act is the main transparency instrument and has been central to algorithmic accountability work in India; the Samagra Vedika investigation, for example, relied on RTI filings (Rebello, 2025). Its application to algorithmic systems is nonetheless uncertain. Public authorities can invoke section 8(1) exemptions, including commercial confidence and impediment to investigation, and questions arise where systems are built or run by private vendors.

Section 44(3) of the DPDP Act has added a new strain by rewriting section 8(1)(j) of the RTI Act. Before the amendment, personal information was exempt only if it had no relationship to public activity or interest or would cause an unwarranted invasion of privacy, and disclosure could be ordered where larger public interest justified it. Petitioners argue that the amended clause removes the public-interest override and in effect exempts personal information wholesale (LiveLaw, 2026). On 16 February 2026 the Supreme Court issued notice in *Venkatesh Nayak v. Union of India* (W.P.(C) 177/2026) and connected petitions but declined an interim stay, and on 7 August 2026 it gave the Centre two weeks to respond (Supreme Court Observer, 2026; Business Standard, 2026). The petitioners rely on *CPIO, Supreme Court of India v. Subhash*

Chandra Agarwal (2020) 5 SCC 481, in which the Constitution Bench treated both privacy and information as fundamental and found the original balance appropriate. For AI governance the stakes are practical: the identities of officials who approved a system, and data about beneficiaries needed for audit, may both be characterised as "personal information".

5.3 Toward meaningful transparency

Transparency should be structured rather than absolute. Four measures follow. *Public algorithm registers*, as adopted in Amsterdam and Helsinki, should list the automated systems used by public bodies with their purpose, data sources and responsible officers (Ada Lovelace Institute et al., 2021). *Reasoned decisions*, consistent with the principle that administrative orders must give reasons, should state the factors relied on and be provided in the citizen's language. *Tiered disclosure* would give independent auditors and courts technical detail under confidentiality while publishing plain-language summaries, which addresses gaming concerns without accepting blanket secrecy. *Contractual transparency* should require vendors to give documentation, audit access and test data, and should bar trade-secret claims from limiting oversight of systems that perform public functions.

6. Accountability: Locating Responsibility

Busuioc (2021) argues that public-sector AI creates an accountability deficit because of opacity, technical complexity and diffuse responsibility. In India, three features aggravate this. Development is often outsourced, so the designers of a system sit outside public-law duties. Human oversight is frequently nominal. Alon-Barkat and Busuioc (2023) find that public officials tend to follow algorithmic advice selectively, and Green (2022) argues that policies requiring human oversight of government algorithms rest on assumptions about human capacity that the evidence does not support. Reports on Samagra Vedika are consistent with this: officials formally retain authority to override the system but are said in practice to defer to it (Oxford Human Rights Hub, 2026). Finally, the institutions proposed by the 2025 Guidelines appear designed to coordinate and advise rather than to receive and decide complaints, so an excluded citizen may find no identifiable person able to explain or reverse the decision.

Indian public law offers established avenues: judicial review under Articles 32 and 226, statutory appeals, grievance systems such as CPGRAMS, parliamentary oversight and audit by the Comptroller and Auditor General. Courts have scrutinised technology-mediated State action, most notably in the Aadhaar judgment. But judicial review is slow, costly and reactive, and for high-volume automated decisions affecting poor households it cannot be the primary remedy.

Accountability can instead be built into systems and procedures (Kroll et al., 2017; Raji et al., 2020). The Kerala High Court policy is instructive: by prohibiting AI from producing findings or orders and leaving responsibility for the content of judicial orders with the judge, it allocates accountability by design (Jhalnia, 2025). Extending this logic to executive decision-making suggests five elements. Human review must have real authority: the reviewer needs information, time and institutional

protection to depart from the system's output. Every high-risk system needs a named accountable officer within the deploying department, so that responsibility cannot be deflected to a vendor or an algorithm. Algorithmic impact assessments should precede deployment and be published in redacted form; Metcalf et al. (2021) caution that such assessments only work when the affected communities' concerns shape what counts as an "impact". Independent audit, potentially by the CAG or the proposed AI Safety Institute, should test accuracy and disparate impact periodically. And citizens need a simple, time-bound right to human review, with restitution for wrongful exclusion, without having to litigate.

7. Data Protection: Rights, Exemptions and Risks

7.1 The constitutional baseline

Puttaswamy placed informational privacy within Article 21 and set the proportionality test for State intrusions. AI-driven governance implicates it directly because of the scale of aggregation, profiling and inference involved. The concern is

not only collection but linkage: combining datasets from many departments creates comprehensive profiles no single dataset would reveal and invites function creep. Samagra Vedika's integration of over thirty databases is a clear example. In the Aadhaar judgment the Court accepted the scheme only with limits, a reminder that scale and convenience do not displace necessity.

7.2 The DPDP Act and Rules: strong on paper, weaker against the State

The DPDP Act 2023 imposes duties on data fiduciaries, including purpose limitation, accuracy, security safeguards, breach notification and erasure, and gives data principals rights of access, correction, erasure and grievance redress. The Rules (2025) add operational detail, a phased timeline, and a Data Protection Board, constituted in November 2025 with four members (Shardul Amarchand Mangaldas, 2025). Its significance for AI governance lies as much in its limits as in its provisions, summarised in Table 3.

Table 3. DPDP framework features that shape AI-driven governance, and reform options

Feature	Provision	Effect on AI-driven governance	Reform option
Consent-free State processing	s 7(b); Rule 5; Second Schedule	Welfare and licensing AI may process personal data without consent; safeguards depend on Schedule standards	Read Schedule standards through necessity; require impact assessment for profiling and linkage
Exemption of State instrumentalities	s 17(2)(a)	Notified agencies exempt from most obligations on security and public-order grounds	Publish notifications with reasons; add sunset and independent review
Government calls for information	Rule 23	Confidential State access to data held by fiduciaries	Prior authorisation; periodic transparency reporting
Board independence	Rules 17(1)-(2)	Executive-dominated selection of a quasi-judicial body; challenged	Independent selection committee; security of tenure
Automated decisions	No specific provision	No right to human review, reasons or information on logic	Add right to human review and reasoned explanation for significant decisions
Interface with RTI	s 44(3) amending RTI s 8(1)(j)	Public-interest override removed; pending before Supreme Court	Restore override or define "personal information" narrowly

Section 7(b) allows processing without consent to provide State subsidies, benefits, services, certificates, licences and permits, and Rule 5 requires such processing to follow the standards in the Second Schedule. Because welfare AI is the paradigm case, this is where the regime matters most, and the protective content depends on those standards being interpreted through a necessity lens. Section 17(2)(a) allows the central government to exempt notified State instrumentalities on grounds such as sovereignty, security and public order. Rule 23 preserves the government's power to call for information from fiduciaries and to direct non-disclosure of having done so (IKIGAI Law, 2025). The Board's appointment mechanism under Rules 17(1) and 17(2) has been challenged as executive-dominated, which matters when the Board must adjudicate against the State (The Wire, 2026). Unlike Article 22 of the GDPR, the Act contains no right to human review of solely automated decisions or to information about their logic. These provisions are before the Supreme Court, which has not stayed them (Internet Freedom Foundation, 2026).

7.3 Reconciling data protection with AI utility

Some tension between data protection and AI performance is genuine: minimisation restricts dataset reuse, and erasure duties

conflict with the long retention that model training favours. It can be reduced by privacy-enhancing technologies such as differential privacy and federated learning, strict access controls and logging, purpose-bound data-sharing frameworks with independent oversight, and data protection impact assessments integrated with algorithmic impact assessments. The proportionality test should decide whether large-scale linkage is justified at all: if eligibility can be verified with a narrower, less intrusive check, the broader profile fails the necessity limb.

8. Comparative Lessons

The European Union offers the nearest binding comparator. Its AI Act (Regulation (EU) 2024/1689) classifies AI systems by risk, and systems used by public authorities to determine eligibility for essential public benefits fall in the high-risk category, attracting duties on risk management, data quality, logging, human oversight and transparency. The GDPR adds safeguards around solely automated decisions. Yet even this binding regime had difficulty operationalising its rules: because harmonised standards and supervisory infrastructure were not ready, Regulation (EU) 2026/1744, the Digital Omnibus on AI, entered into force on 27 July 2026 and deferred the high-risk obligations for stand-alone systems to 2 December 2027 (Cloud

Security Alliance, 2026). The lesson is not that binding rules fail but that obligations need sequencing, standards and resourced regulators.

Europe also supplies cautionary cases from the digital welfare state. In the Dutch childcare-benefits scandal, algorithmic risk profiling contributed to wrongful fraud accusations against thousands of families (Amnesty International, 2021), and the

Dutch SyRI system was held unlawful by a Hague court for failing the proportionality and transparency requirements of the right to private life (Rachovitsa & Johann, 2022). These episodes resemble the concerns raised about Samagra Vedika, and they show that courts can enforce proportionality against welfare algorithms where a statutory hook exists. Table 4 compares the two systems.

Table 4. India and the European Union compared (as at September 2026)

Dimension	India	European Union	Design lesson for India
Legal form	Voluntary guidelines plus existing law; sectoral regulators	Binding AI Act plus GDPR	Statute for high-risk State uses; keep guidance for low-risk uses
Risk classification	Risk-mitigation guidance; no legal tiers	Four-tier legal classification; benefit-eligibility systems are high-risk	Adopt tiers tied to enforceable duties
Rights on automated decisions	None specific in DPDP Act	GDPR Art. 22 safeguards	Add right to human review and reasons
State exemptions	Broad power to exempt instrumentalities (s 17(2))	Security carve-outs subject to conditions	Necessity test and independent review
Oversight	Proposed coordinating and advisory bodies; Data Protection Board	National authorities and EU AI Office	Independent auditor with investigative powers
Timing	Core DPDP duties May 2027; no AI statute planned	High-risk duties from 2 Dec 2027 after deferral	Sequence duties and build standards and capacity early

9. An Integrated Framework for Responsible AI Governance

The analysis suggests that the four values are best reconciled by an institutional framework rather than a single rule. This paper proposes five interlocking components, followed by a sequenced roadmap.

Risk-tiered regulation. AI uses in government should be classified by impact on rights and entitlements (Figure 3). Low-

risk uses such as internal drafting or translation need only light guidance. High-risk uses, including benefit eligibility, policing, tax enforcement and any system with legal or similarly significant effects, trigger the full safeguards below. Some uses, such as untargeted biometric surveillance in public places, should be prohibited or permitted only under a specific statute with independent authorisation. Tiering preserves efficiency where risk is low and concentrates scrutiny where harm is greatest.




Figure 3. Proposed risk-tiered framework for AI in government (author's proposal)

Statutory basis for high-risk State use. In line with *Puttaswamy*'s requirement of legality, intrusive AI uses should rest on a law or binding rules that define purpose, scope, retention limits and oversight. Executive guidelines are inadequate for measures that limit fundamental rights.

Pre-deployment and continuing assessment. Algorithmic and data protection impact assessments should precede deployment of high-risk systems, be revisited when systems change materially, and be published in redacted form. Pilots should

report error rates and disparate impact, by caste, gender, region and language, before scale-up.

Transparency and contestability infrastructure. This includes public algorithm registers, reasoned decision notices, a right to human review, and low-cost appeals with defined timelines and default restoration of benefits during review. The RTI regime should be clarified to cover privately built tools used for public functions, and the public-interest override in section 8(1)(j) should be restored or the term "personal information" narrowed, depending on the Supreme Court's decision.

Independent oversight and capacity. An independent body,

whether a dedicated regulator, an expanded mandate for an existing institution or a network of sectoral auditors, should have powers to audit systems, receive complaints and publish findings. The Data Protection Board's independence should be strengthened and State exemptions under section 17(2) made subject to necessity review. Officials need training to interrogate outputs, and governments need in-house technical expertise to avoid dependence on vendors. Because welfare delivery is mostly State-run, national standards must be paired with support for State capacity, and non-digital and assisted channels must remain available for those with limited connectivity, literacy or documentation. Table 5 sets out a sequenced roadmap.

Table 5. Proposed reform roadmap

Action	Instrument	Lead body	Indicative timing
Public register of high-risk government AI systems	Executive order or rules under MeitY guidelines	MeitY with States	0-12 months
Mandatory algorithmic and data protection impact assessments for high-risk uses	Binding rules; procurement conditions	MeitY; line ministries	0-12 months
Right to human review and reasoned notice for significant automated decisions	Statutory amendment or sectoral rules	Parliament; welfare departments	12-24 months
Narrow s 17(2) exemptions; secure Board independence	Amendment to DPDP Act and Rules	Parliament; MeitY	12-24 months
Independent technical audit of high-risk systems	Statutory mandate; audit standards	CAG or AI Safety Institute	12-36 months
Capacity-building and assisted non-digital channels	Funding; State-level programmes	MeitY; State governments	Ongoing

10. Anticipating Objections and Limitations

Three objections merit response. First, regulation may slow innovation and delay benefits. But the costs of unregulated deployment are also borne by the poorest, through wrongful exclusion and loss of trust, and tiering deliberately limits burdens on low-risk uses. Second, transparency may compromise security and enable fraud. Tiered disclosure and auditor access preserve confidentiality of sensitive parameters while maintaining oversight. Third, existing law may be sufficient. The constitutional standards are robust, but enforcement is reactive, the DPDP Act's exemptions reduce its force against the State, the RTI public-interest test is under strain, and no statute imposes a duty of algorithmic explanation. Limitations should be acknowledged. The paper is doctrinal and conceptual, and evidence on Indian deployments comes largely from journalistic, civil-society and non-governmental technical investigations rather than systematic public evaluation, partly because government does not publish performance data, itself a symptom of the transparency problem. Several central questions, notably the Supreme Court's ruling on the DPDP Act, are pending, and the DPDP Rules' main provisions are not yet in force. Empirical research on error rates, exclusion and outcomes across States should be a priority.

11. CONCLUSION

AI offers the Indian state real opportunities to improve service delivery, reduce delay and extend reach. But the legitimacy of governance depends on how decisions are made as well as how quickly. This paper has argued that efficiency, transparency, accountability and data protection can be reconciled, but not by

default. Reconciliation requires legal grounding for intrusive State uses, risk-proportionate regulation, structured transparency, human oversight with genuine authority, independent audit, accessible redress and rigorous data protection standards. It also requires closing four gaps in the current architecture: a voluntary AI framework without enforceable duties, broad State exemptions in the DPDP Act, the strain on the RTI public-interest test, and the absence of a right to human review and explanation.

India has a strong constitutional tradition of proportionality and non-arbitrariness, an active transparency movement, and an opportunity to shape its rules while deployment is still expanding. The task is to turn these strengths into enforceable institutional arrangements before automated systems become entrenched. If AI-driven governance is to honour India's constitutional promise, efficiency must serve, and be constrained by, the rights and dignity of the citizens it is meant to benefit.

REFERENCES

1. Ada Lovelace Institute, AI Now Institute, Open Government Partnership. Algorithmic accountability for the public sector. 2021. Available from: <https://www.adalovelaceinstitute.org>
2. Al Jazeera, Pulitzer Center AI Accountability Network. How an algorithm denied food to thousands of poor in India's Telangana. Pulitzer Center. 2024 Jan. Available from: <https://pulitzercenter.org/stories/how-algorithm-denied-food-thousands-poor-indias-telangana>

3. Alon-Barkat S, Busuioc M. Human-AI interactions in public sector decision making: “automation bias” and “selective adherence” to algorithmic advice. *J Public Adm Res Theory*. 2023;33(1):153–169.
4. Amnesty International. Xenophobic machines: Discrimination through unregulated use of algorithms in the Dutch childcare benefits scandal. Amnesty International; 2021.
5. Amnesty International. Use of entity resolution in India: Shining a light on how new forms of automation can deny people access to welfare. 2024 Apr. Available from: <https://www.amnesty.org/en/latest/research/2024/04/entity-resolution-in-indias-welfare-digitalization/>
6. Ananny M, Crawford K. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media Soc*. 2018;20(3):973–989.
7. Anuradha Bhasin v. Union of India. (2020) 3 SCC 637.
8. Bovens M. Analysing and assessing accountability: A conceptual framework. *Eur Law J*. 2007;13(4):447–468.
9. Bullock JB. Artificial intelligence, discretion, and bureaucracy. *Am Rev Public Adm*. 2019;49(7):751–761.
10. Busuioc M. Accountable artificial intelligence: Holding algorithms to account. *Public Adm Rev*. 2021;81(5):825–836.
11. Business Standard. SC seeks Centre's response on DPDP Act amendment to right to information. 2026 Aug 7. Available from: https://www.business-standard.com/india-news/sc-seeks-centre-s-response-on-dpdp-act-amendment-to-right-to-information-126080701888_1.html
12. Central Public Information Officer, Supreme Court of India v. Subhash Chandra Agarwal. (2020) 5 SCC 481.
13. Chambers and Partners. MeitY unveils India's approach towards regulating artificial intelligence. 2025 Nov 17. Available from: <https://chambers.com/articles/meity-unveils-india-s-approach-towards-regulating-artificial-intelligence>
14. Cloud Security Alliance. EU AI Act's high-risk deadline: Deferred, not cancelled. Research note on Regulation (EU) 2026/1744. 2026 Aug 1. Available from: <https://labs.cloudsecurityalliance.org/research/csa-research-note-eu-ai-act-high-risk-deadline-omnibus-20260/>
15. India. Digital Personal Data Protection Act, 2023.
16. India. Digital Personal Data Protection Rules, 2025. G.S.R. 846(E). Published 2025 Nov 13.
17. dpdprules.org. DPDP FAQs: Questions and answers. 2026. Available from: <https://dpdprules.org/questions>
18. Eubanks V. Automating inequality: How high-tech tools profile, police, and punish the poor. St. Martin's Press; 2018.
19. European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). 2024.
20. European Union. Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI). Off J Eur Union. 2026 Jul 24.
21. Green B. The flaws of policies requiring human oversight of government algorithms. *Comput Law Secur Rev*. 2022; 45:105681.
22. IKIGAI Law. A closer look at the DPDP Rules 2025. 2025 Nov 14. Available from: <https://www.ikigailaw.com/article/647/a-closer-look-at-the-dpdp-rules-2025>
23. Internet Freedom Foundation. Supreme Court issues notice on constitutional challenge to the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025. 2026 Feb. Available from: <https://internetfreedom.in/supreme-court-issues-notice-on-constitutional-challenge-to-the-digital-personal-data-protection-act-2023-and-the-digital-personal-data-protection-rules-2025/>
24. Jhalaria S. Bench and bot: The Kerala HC's AI guidelines and the bigger judicial puzzle. *SpicyIP*. 2025 Aug 18. Available from: <https://spicyip.com/2025/08/bench-and-bot-the-kerala-hcs-ai-guidelines-and-the-bigger-judicial-puzzle.html>
25. Justice K.S. Puttaswamy (Retd.) v. Union of India. (2017) 10 SCC 1.
26. Justice K.S. Puttaswamy (Retd.) v. Union of India (Aadhaar). (2019) 1 SCC 1.
27. Khera R, editor. Dissent on Aadhaar: Big data meets big brother. Orient BlackSwan; 2019.
28. Kroll JA, Huey J, Barocas S, Felten EW, Reidenberg JR, Robinson DG, et al. Accountable algorithms. *Univ Pa Law Rev*. 2017;165(3):633–705.
29. Kuziemski M, Misuraca G. AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. *Telecomm Policy*. 2020;44(6):101976.
30. LiveLaw. “Interesting question is what's public data & personal data?”: Supreme Court on challenge to DPDP Act. 2026 Mar 12. Available from: <https://www.livelaw.in/top-stories/whats-public-data-personal-data-supreme-court-on-challenge-to-dpdp-act-526194>
31. MediaNama. Explained: How Indian courts are using AI and why the SC wants humans in the loop. 2025 Dec 4. Available from: <https://www.medianama.com/2025/12/223-supreme-court-ai-white-paper-indian-judiciary/>
32. Metcalf J, Moss E, Watkins EA, Singh R, Elish MC. Algorithmic impact assessments and accountability: The co-construction of impacts. In: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. ACM; 2021. p. 735–746.
33. Ministry of Electronics and Information Technology. India AI governance guidelines. Government of India; 2025. Available from: <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2228315>
34. NITI Aayog. National strategy for artificial intelligence: #AIforAll. Government of India; 2018.
35. NITI Aayog. Responsible AI #AIforAll: Approach document for India. Parts 1 and 2. Government of India; 2021.

36. Oxford Human Rights Hub. Algorithmic welfare exclusion and the right to food in India: Lessons from Samagra Vedika. 2026 Feb 19. Available from: <https://ohrh.law.ox.ac.uk/algorithmic-welfare-exclusion-and-the-right-to-food-in-india-lessons-from-samagra-vedika/>
37. Press Information Bureau. India AI Impact Summit 2026: Landmark global declaration and major AI investment commitments. Government of India; 2026 Feb. Available from: <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2234343>
38. Prasar Bharati News Services. Supreme Court reconstitutes AI Committee to oversee adoption and deployment of AI tools. 2025 Dec 11. Available from: <https://www.newsonair.gov.in/supreme-court-reconstitutes-ai-committee-to-oversee-adoption-and-deployment-of-ai-tools>
39. Rachovitsa A, Johann N. The human rights implications of the use of AI in the digital welfare state: Lessons learned from the Dutch SyRI case. Hum Rights Law Rev. 2022;22(2): ngac010.
40. Raji ID, Smart A, White RN, Mitchell M, Gebru T, Hutchinson B, et al. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. ACM; 2020. p. 33–44.
41. Rebelo K. How to hold algorithms accountable in India. Journalist fellowship paper. Oxford: Reuters Institute for the Study of Journalism, University of Oxford; 2025.
42. Reserve Bank of India. Framework for responsible and ethical enablement of artificial intelligence (FREE-AI): Report of the Committee. Reserve Bank of India; 2025.
43. India. Right to Information Act, 2005.
44. Shardul Amarchand Mangaldas & Co. Enforcement of the DPDP Act and notification of the DPDP Rules. 2025 Nov 27. Available from: <https://www.amsshardul.com/insight/enforcement-of-the-dpdp-act-and-notification-of-the-dpdp-rules/>
45. Supreme Court Observer. Challenge to DPDP Act: Bench refuses to stay amendment, issues notice. 2026. Available from: <https://www.scobserver.in/reports/challenge-to-dpdp-amendment-of-rti-bench-refuses-to-stay-amendment-issues-notice/>
46. The Wire. SC to consider plea challenging parts of DPDP Act and Rules, including amendment to RTI. 2026 Feb 15. Available from: <https://m.thewire.in/article/government/sc-to-consider-plea-seeking-striking-down-of-dpdp-act-and-rules-including-amendment-to-rti>
47. Wachter S, Mittelstadt B, Floridi L. Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. Int Data Priv Law. 2017;7(2):76–99.
48. Wirtz BW, Weyerer JC, Geyer C. Artificial intelligence and the public sector: Applications and challenges. Int J Public Adm. 2019;42(7):596–615.
49. Young MM, Bullock JB, Leczy JD. Artificial discretion as a tool of governance: A framework for understanding the impact of artificial intelligence on public administration. Perspect Public Manage Gov. 2019;2(4):301–313.

Creative Commons (CC) License

This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution–Non-Commercial–No Derivatives 4.0 International (CC BY-NC-ND 4.0) license. This license permits sharing and redistribution of the article in any medium or format for non-commercial purposes only, provided that appropriate credit is given to the original author(s) and source. No modifications, adaptations, or derivative works are permitted under this license.

About the Author



Dr. E. Priyanka is an Assistant Professor (Senior Grade) in the Department of Politics and Public Administration at Crescent School of Law, B.S. Abdur Rahman Institute of Science and Technology, Vandalur, Tamil Nadu, India. Her academic interests include political science, public administration, legal studies, and interdisciplinary research.