

Algorithmic Accountability: Can AI Strengthen Good Governance in India Without Undermining It?

Dr. Kiran Estarla

Lecturer in Political Science

Government Degree College, Naidupet, Andhra Pradesh-524126

Abstract:

Artificial intelligence is rapidly reshaping how the Indian state identifies beneficiaries, allocates welfare, polices crime, and delivers public services, promising efficiency gains that decades of manual administration could not achieve. Yet the same automated systems that promise to reduce discretion-driven corruption also risk manufacturing new opacities: black-boxed eligibility criteria, unaccountable exclusion errors, and surveillance infrastructures that outpace the statutory frameworks meant to restrain them. This paper examines whether algorithmic tools can strengthen good governance in India without eroding the accountability mechanisms — parliamentary oversight, judicial review, the Right to Information, and administrative grievance redress — that legitimate democratic administration. Drawing on India's Digital Public Infrastructure stack (Aadhaar, Direct Benefit Transfer, the IndiaAI Mission), recent legislative developments including the Digital Personal Data Protection Act, 2023, and comparative regulatory experience from the European Union, the United Kingdom, and Singapore, the paper argues that algorithmic accountability in India remains structurally underdeveloped relative to algorithmic deployment. It proposes a four-pillar framework — mandatory algorithmic impact assessments, tiered transparency obligations, an independent audit function, and citizen-facing redress mechanisms — as a pathway toward AI-enabled governance that is efficient without becoming unaccountable.

Keywords: algorithmic accountability; artificial intelligence governance; digital public infrastructure; Aadhaar; Direct Benefit Transfer; Digital Personal Data Protection Act; administrative law; India

1. Introduction

Over the past decade, the government of India has undertaken one of the most ambitious digitisation programmes attempted by any democracy, building a stack of interoperable digital public infrastructure — Aadhaar-based biometric identity, the Jan Dhan-Aadhaar-Mobile (JAM) trinity, Direct Benefit Transfer (DBT), and, more recently, the IndiaAI Mission — that now mediates how tens of millions of citizens access welfare, subsidies, and public services. What began as a data-management exercise aimed at deduplicating beneficiary lists has evolved into something considerably more consequential: administrative decision-making that is, in growing measure, delegated to statistical models, biometric matching algorithms, and rule-based automation. Fraud-detection engines flag suspicious welfare claims, facial-recognition systems assist criminal investigation, chatbots triage citizen grievances, and geospatial models allocate infrastructure spending across districts.

The appeal of this shift is straightforward. Manual administration in India has long been associated with discretionary corruption, arbitrary delay, and capture by local elites; algorithmic systems promise consistency, speed, and an audit trail that a paper file rarely produced. Yet the very features that make algorithmic governance attractive — its opacity to lay scrutiny, its concentration of decision authority in code rather than in a named official, and its capacity to operate at a scale no human bureaucracy could match — sit uneasily against the foundations of accountable government. Indian administrative law was built around a human decision-maker who could be named, whose reasoning could be demanded under the Right to Information Act, 2005, and whose order could be challenged before a tribunal or a High Court under Article 226 of the Constitution. When the decision-maker is a scoring model trained on historical data, each of these accountability pathways becomes harder to exercise, not because the law has changed, but because the object it was designed to scrutinise has changed.

This tension is not hypothetical. Investigative and academic work has repeatedly documented cases in which Aadhaar-linked biometric authentication failures denied rations to entitled beneficiaries, in which welfare rolls were pruned by automated de-duplication with no individualised notice to those removed, and in which predictive policing tools deployed by state police departments operated with minimal public documentation of their training data or error rates. At the same time, the same infrastructure is credited with substantial, independently measurable gains: the Direct Benefit Transfer system alone is estimated to have generated cumulative fiscal savings running into several lakh crore rupees by curbing leakage and eliminating ghost beneficiaries. The central puzzle this paper addresses is therefore not whether AI has a place in Indian governance — it plainly already does — but whether the accountability architecture around it can be strengthened quickly enough to keep pace with deployment, and what that architecture should look like.

The paper proceeds in six further parts. Section 2 develops a conceptual framework for algorithmic accountability tailored to the specific institutional features of Indian administrative law. Section 3 maps the current landscape of AI adoption in Indian public administration, with attention to budgetary commitments under the IndiaAI Mission. Section 4 examines the existing accountability mechanisms — the RTI Act, judicial review, the Comptroller and Auditor General, and the newly operational Digital Personal Data Protection Act, 2023 — and identifies the specific gaps that arise when these mechanisms confront algorithmic decision-making. Section 5 examines three case studies in which the efficiency-accountability tension has played out concretely: Aadhaar-linked welfare authentication, predictive and facial-recognition policing, and state-level AI-enabled service delivery. Section 6 situates the Indian experience against comparative regulatory frameworks in the European Union, the United Kingdom, and Singapore. Section 7 proposes a four-pillar framework for accountable AI governance suited to Indian institutional conditions, and Section 8 concludes.

2. Conceptual Framework: Algorithmic Accountability in Democratic Administration

Accountability in public administration is conventionally understood as a relationship in which an actor is obliged to explain and justify conduct to a forum that can pose questions, render judgement, and impose consequences (Bovens, 2007). Applied to algorithmic systems, this relationship fractures along four distinguishable dimensions that this paper treats as the analytical core of "algorithmic accountability": transparency (can the existence, purpose, and basic logic of the system be known to those it affects), answerability (can a specific, identifiable actor be made to explain a specific decision), enforceability (does a forum exist with the power to compel correction or impose sanction), and redress (can an individual

who is wrongly affected obtain a remedy in a reasonable time). Each dimension maps loosely onto an existing pillar of Indian administrative law — transparency onto the RTI Act, 2005; answerability onto the principles of natural justice and the requirement of a reasoned order; enforceability onto judicial review under Articles 32 and 226 and statutory audit by the Comptroller and Auditor General (CAG); and redress onto grievance-redress mechanisms and, since 2023, the Data Protection Board established under the DPDP Act.

The difficulty is that each of these pillars was designed around an assumption that no longer reliably holds when a decision is produced by a trained model: the assumption that reasoning can be reconstructed after the fact from a file. A revenue officer's order denying a subsidy can be traced to a specific clause of a specific government order; a machine-learning classifier's denial of the same subsidy is the product of weighted correlations across potentially thousands of features, and the officer nominally responsible for the decision may have limited practical capacity to explain, let alone contest, the system's output. Marda (2018) frames this as a shift from rule-based to data-driven decision-making in Indian public administration, and cautions that the country's early and enthusiastic adoption of AI in criminal justice and welfare contexts has outpaced any parallel development of oversight capacity within the bureaucracy itself. The NITI Aayog's own Responsible AI approach papers (NITI Aayog, 2021) acknowledge this gap, calling for sector-specific regulatory mechanisms, algorithmic audits, and grievance redress, but stop short of prescribing binding obligations, leaving implementation to individual ministries and departments.

A second conceptual distinction that matters for the Indian context is between accountability for the system (was it designed, tested, and procured responsibly) and accountability for the decision (was this particular citizen treated fairly). Much of the existing discourse on "Responsible AI" in India, echoing the OECD AI Principles (OECD, 2019) and the broader global AI-ethics literature, addresses itself to the former: bias audits, dataset documentation, model cards. These are necessary but not sufficient. A welfare applicant who is wrongly excluded from a scheme needs an answer to the second question — a specific, timely, and comprehensible explanation of why they, individually, were denied, and a route to correction — and it is at this individual level that Indian algorithmic governance is weakest. This paper's subsequent analysis is therefore organised around both levels, treating system-level responsible-AI commitments as a necessary but incomplete foundation for the individual-level accountability that democratic legitimacy actually requires.

3. The Current Landscape: AI in Indian Public Administration

India's engagement with AI in governance has accelerated sharply since the Union Cabinet's approval, on 7 March 2024, of the IndiaAI Mission with a five-year outlay of ₹10,371.92 crore, structured around seven pillars: IndiaAI Compute Capacity, the IndiaAI Innovation Centre, the IndiaAI Datasets Platform, an Application Development Initiative, IndiaAI FutureSkills, IndiaAI Startup Financing, and a Safe and Trusted AI pillar (Press Information Bureau, 2024). The compute pillar alone targets a supercomputing capacity of over 10,000 GPUs built in public-private partnership, to be made available to start-ups, researchers, and government departments alike. Against this five-year approved outlay, the Union Budget for 2024-25 allocated ₹551.75 crore to the Mission in its first year — a little over five percent of the total sanctioned corpus (IndiaAI, 2024). A first-year allocation that constitutes a small fraction of an ambitious five-year commitment is not, on its own, evidence of a governance problem; new missions typically ramp up spending gradually as institutional capacity is built. It is nonetheless an accountability-relevant fact, since the pace and sequencing of a mission of this scale — one whose Safe and Trusted AI pillar is itself

supposed to develop guardrails for the very systems examined in this paper — will determine how quickly capacity for responsible deployment keeps pace with capacity for deployment itself.

Table 1 summarises the principal digital-governance and AI-adjacent initiatives operating at the Union level as of 2024, alongside available budgetary and coverage figures. It illustrates two features of the landscape relevant to the accountability question. First, the scale of algorithmic mediation in Indian administration is already very large — Aadhaar-linked authentication and DBT together touch a beneficiary base numbering in the hundreds of crores — well ahead of any comparable scale of independent algorithmic audit. Second, India's standing on the United Nations' E-Government Development Index (EGDI), while improving over the preceding decade, still trailed the "very high" tier occupied by the digital-government leaders whose regulatory frameworks are examined in Section 6, suggesting that institutional and human-capital capacity for digital governance, including oversight capacity, has not grown as fast as the underlying technology deployment.

Table 1. Selected Union Government Digital Governance and AI Initiatives

Initiative	Nodal Agency	Key Figures	Source
IndiaAI Mission — approved outlay	MeitY / IndiaAI (Digital India Corporation)	₹10,371.92 crore over five years, approved 7 March 2024, covering compute, innovation centre, datasets platform, futureskills, startup financing, and Safe & Trusted AI	Press Information Bureau (2024)
IndiaAI Mission — FY 2024-25 first-year budget	MeitY	Budget Estimate ₹551.75 crore (about 5% of the total five-year outlay); part of a wider ₹21,936.90 crore MeitY budget for 2024-25	IndiaAI (2024)
Direct Benefit Transfer (DBT)	DBT Mission, Cabinet Secretariat	Cumulative savings exceeding ₹3.48 lakh crore since inception (2013); ₹1.85 lakh crore from PDS alone via Aadhaar-linked ration authentication; DBT operates across 314 schemes run by 53 central ministries	The Hindu BusinessLine (2024)
Digital Personal Data Protection Act, 2023	Ministry of Electronics and Information Technology (MeitY)	Enacted 11 August 2023; as of 2024, subordinate Rules and the Data Protection Board of India had not yet been notified or constituted	Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023)
UN E-Government Development Index (India)	UN DESA (assessed)	Composite score 0.59, rank 105 of 193 countries (2022 survey, most recent available through 2024)	UN DESA (2022)

Note: Figures compiled from the sources cited in the final column; budget figures are in Indian rupees (crore = 10 million). Compiled by the author, 2024.

Beyond the Union level, several states — Telangana and Andhra Pradesh among them — have layered AI-enabled tools onto pre-existing e-governance platforms: grievance-redress chatbots, GIS-based infrastructure-planning dashboards, and automated eligibility screening for state welfare schemes. These state-level deployments are typically procured through empanelled technology vendors with limited public disclosure of contract terms, training-data provenance, or performance benchmarks, a pattern that mirrors — and in several respects amplifies — the transparency gap already visible at the Union level, since state Right to Information regimes and state information commissions vary considerably in enforcement capacity across the country (Bhatia & Bhattacharya, 2022).

A third distinction worth drawing is between accountability that operates prospectively, at the design and procurement stage, and accountability that operates retrospectively, once a system is already affecting citizens. Indian public procurement of digital systems has traditionally emphasised cost and delivery timelines in tender evaluation, with limited standard practice for building in, ex ante, the kind of fairness testing or explainability requirement that would make retrospective accountability easier to exercise later. Retrofitting these requirements onto a system already in production is considerably more expensive, both financially and politically, than specifying them at the tendering stage, which is one reason the framework proposed in Section 7 of this paper places particular weight on prospective, procurement-stage obligations rather than relying solely on after-the-fact audit or litigation.

4. Accountability Mechanisms and Structural Gaps

India already possesses a reasonably dense set of institutions capable, in principle, of holding public administration accountable. The question this section addresses is how well each of these institutions extends to algorithmic decision-making as currently practised. Table 2 sets out four existing mechanisms, their traditional scope, and the specific gap that appears when they are applied to an AI-mediated administrative decision.

Table 2. Existing Accountability Mechanisms Applied to Algorithmic Governance

Mechanism	Traditional Scope	Gap in Algorithmic Context	Illustrative Reference
Right to Information Act, 2005	Citizens may demand reasons, file notings, and records underlying an administrative decision	Model architecture, training data, and feature weights are frequently classified as third-party commercial information or simply not held in a disclosable "record"; RTI responses often describe outcomes without the underlying decision logic	Marda (2018); CIC orders on algorithmic opacity

Mechanism	Traditional Scope	Gap in Algorithmic Context	Illustrative Reference
Judicial review (Art. 226/32; principles of natural justice)	Courts examine whether a reasoned order was passed and whether the affected party had a fair hearing	Courts require expert testimony to interrogate a model's logic, and Indian jurisprudence on a constitutional "right to explanation" for automated decisions remains underdeveloped compared to EU case law	Justice K.S. Puttaswamy v. Union of India (2017)
Comptroller and Auditor General (CAG) audit	Financial and performance audit of government schemes and expenditure	Performance audits typically assess outcomes and expenditure, not algorithmic fairness, error rates, or disparate impact across social groups	CAG performance audit reports on DBT schemes
DPDP Act, 2023 / Data Protection Board of India	Would govern consent, purpose limitation, and grievance redress for personal data processing once operational	Enacted August 2023 but, as of 2024, subordinate Rules remain unnotified and the Data Protection Board is not yet constituted; even once operational, its text is oriented toward data-processing consent rather than the fairness or reviewability of the algorithmic decisions built on that data	Digital Personal Data Protection Act, 2023

Compiled by the author from statutory texts and secondary sources cited, 2024.

The DPDP Act deserves particular attention because it is the most recent and, on paper, the most directly relevant statutory development. Enacted on 11 August 2023 after several years of drafting and two withdrawn earlier bills, the Act envisages a Data Protection Board of India empowered to receive complaints and impose penalties for non-compliance with data-processing obligations, and it promises meaningful new rights — access, correction, erasure, and grievance redress — for data principals. As of 2024, however, the Act remains substantially unimplemented: the subordinate Rules that would specify consent-manager registration, breach-notification timelines, and the Board's composition and procedure had not been notified, and the Board itself had not been constituted, leaving the Act's enforcement machinery dormant even though its substantive provisions are formally in force. Even once the Rules are notified, the Act does not, at least on its current text, create a freestanding right to an explanation of an

automated decision, nor does it impose a general obligation on government data fiduciaries to disclose the logic of algorithmic systems used in welfare or law-enforcement contexts; several of its provisions carry exemptions for processing by state instrumentalities in the interests of sovereignty, security, and public order that could, if broadly construed, limit its reach over precisely the categories of AI-enabled governance this paper is concerned with. For a paper examining algorithmic accountability in 2024, the DPDP Act is therefore best understood as a statutory promise rather than an operating accountability mechanism: it signals legislative intent to regulate the data layer beneath AI-enabled governance, but its practical value to a citizen affected by an algorithmic decision today remains untested.

A further structural gap concerns capacity rather than law. Even where a legal accountability mechanism formally exists, its exercise depends on the presence of officials, judges, auditors, or information commissioners with the technical literacy to interrogate a model. India's State Information Commissions face well-documented backlogs, and specialised technical capacity to audit algorithmic systems is concentrated in a handful of Union-level bodies and a small community of independent researchers and civil-society organisations, such as the Internet Freedom Foundation, which has repeatedly used RTI requests and public litigation to surface information about facial-recognition deployments by state police forces. Institutional accountability that depends on a small number of overstretched actors is, in practice, thinner than the statute books suggest.

A related and often overlooked channel is the Centralised Public Grievance Redress and Monitoring System (CPGRAMS), which nominally allows any citizen to lodge a complaint against a government decision, including one produced by an automated process, and receive a response within a defined timeline. In practice, CPGRAMS was designed around narrative grievances against identifiable departmental action and has no structured field for a citizen to indicate that their grievance concerns an algorithmic decision specifically, nor any routing logic that would direct such a complaint to a technically equipped reviewer rather than the same generalist grievance officer who handles unrelated administrative complaints. The mechanism therefore exists on paper for exactly the redress function this paper identifies as missing, but its current design does not yet operationalise that function for AI-specific grievances, illustrating again that the gap in Indian algorithmic accountability is often one of institutional design and technical routing rather than of an entirely absent legal instrument.

5. Case Studies in Tension: Efficiency versus Accountability

5.1 Aadhaar-Linked Welfare Authentication

The Aadhaar biometric identity system, with over 1.3 billion enrolments, underpins the JAM trinity and, through it, the DBT architecture credited with the fiscal savings summarised in Table 1. Biometric authentication at the point of ration or subsidy disbursement is the mechanism through which much of this efficiency gain is realised: it removes duplicate and "ghost" beneficiaries from state rolls without requiring a field verification visit. The same mechanism, however, has been documented — most extensively in the work of economist Reetika Khera and collaborators — to generate exclusion errors when a beneficiary's fingerprints fail to match a stored template because of manual labour, age, or illness, or when network or device failures in rural areas prevent authentication altogether (Khera, 2019). Because the system flags a failed transaction rather than a wrongly denied person, and because the point-of-sale operator has limited authority to override an authentication failure, the accountability burden falls on an intended beneficiary who may have no practical route to a swift correction. The efficiency gains recorded in DBT savings figures and the exclusion harms documented in field research are, in this sense, two faces of the same

infrastructure, and the absence of a mandated, time-bound exception-handling and appeals process at the point of failure remains the central unresolved accountability gap in this case.

5.2 Predictive and Facial-Recognition Policing

Several state police forces have adopted, in varying degrees of public disclosure, automated facial-recognition systems (AFRS) for suspect identification and crime-mapping tools intended to direct patrolling resources to statistically "high-risk" localities. Civil-society investigations, particularly by the Internet Freedom Foundation and academic commentators such as Marda (2018), have raised three recurring concerns: first, that these systems are frequently procured and deployed without a specific enabling statute or publicly available standard operating procedure; second, that training-data provenance and error-rate performance, including differential accuracy across skin tone and gender, are rarely disclosed even in response to RTI requests; and third, that a positive match from an AFRS system can trigger further investigative action against a named individual with limited downstream documentation of the algorithm's role in that action, complicating any later judicial review of whether the resulting arrest or detention rested on a reliable evidentiary basis. Because criminal-justice consequences are more severe and more difficult to reverse than a delayed welfare payment, this case illustrates the highest-stakes version of the efficiency-accountability trade-off examined in this paper, and it is the domain in which the absence of a binding, sector-specific accountability instrument is most consequential.

5.3 State-Level AI-Enabled Service Delivery

At the state level, Telangana and Andhra Pradesh, among other states, have integrated AI-enabled chatbots, GIS-based planning dashboards, and predictive maintenance tools into existing e-governance platforms as part of broader digital-government modernisation efforts. These deployments tend to be lower-stakes than welfare authentication or policing — a misrouted grievance ticket is more easily corrected than a wrongful arrest — but they illustrate a distinct accountability concern: procurement opacity. Vendor contracts for AI-enabled municipal or state services are frequently negotiated as part of broader system-integrator agreements, with model performance, data-sharing terms, and liability for erroneous outputs treated as commercially confidential. Because state Right to Information regimes vary in enforcement rigor across India, the degree to which a citizen or a state legislator can obtain basic information about how a given AI tool functions differs substantially by state, producing an accountability landscape that is fragmented along federal lines even before the AI-specific gaps identified in Section 4 are considered.

5.4 Cross-Case Observations

Read together, the three cases suggest that the severity of the accountability gap tracks the severity and reversibility of the consequence at stake rather than the sophistication of the underlying technology. A misdirected municipal chatbot query, a wrongly denied ration transaction, and a false facial-recognition match are all, in a narrow technical sense, instances of the same underlying problem — an automated system producing an erroneous output with limited immediate human oversight — yet they warrant markedly different accountability thresholds. This is consistent with the risk-tiered approach already used internationally (examined in Section 6) and suggests that any Indian framework should explicitly calibrate the intensity of its impact-assessment, transparency, and appeal obligations to the severity of the decision

at stake, rather than applying a single uniform standard across welfare administration, policing, and routine service delivery.

6. Comparative International Frameworks

India's regulatory gap becomes clearer when set against jurisdictions that have moved further toward binding algorithmic accountability instruments. The European Union reached political agreement on its Artificial Intelligence Act in December 2023, and the European Parliament formally adopted the text in March 2024; once its phased obligations take effect, the Act will classify public-sector systems used in welfare eligibility, law enforcement, and migration control as "high-risk," triggering mandatory conformity assessments, technical documentation, human-oversight requirements, and registration in a public EU database before deployment (European Parliament, 2024). The United Kingdom has taken a lighter-touch, non-statutory route through its Algorithmic Transparency Recording Standard, which became mandatory for government departments and selected public-sector bodies from February 2024, requiring structured, standardised disclosures of algorithmic tools used in public-facing decisions, including a plain-language description of the tool's purpose and the human-oversight arrangements around it (UK Central Digital and Data Office, 2024). Singapore's Model AI Governance Framework, while voluntary and industry-facing rather than government-facing, has nonetheless been influential in the region for popularising the language of algorithmic impact assessment and human-in-the-loop design as baseline expectations for any organisation deploying AI in consequential decisions (Infocomm Media Development Authority of Singapore, 2020).

Table 3. Comparative Algorithmic Accountability Frameworks

Jurisdiction	Instrument	Core Accountability Feature	India's Present Equivalent
European Union	EU Artificial Intelligence Act (agreed Dec 2023; adopted March 2024, phased entry into force)	Binding high-risk classification for public-sector welfare, policing, and migration systems; mandatory conformity assessment and public registration	No binding cross-sectoral AI statute; NITI Aayog Responsible AI papers are non-binding guidance
United Kingdom	Algorithmic Transparency Recording Standard (mandatory from Feb 2024)	Standardised public disclosure of government algorithmic tools, including purpose and human-oversight arrangements	No standardised public register of government AI systems
Singapore	Model AI Governance Framework (2020, second edition)	Voluntary algorithmic impact assessment and human-in-the-loop design principles for consequential decisions	Sector-specific guidance only (e.g., RBI, SEBI circulars); no cross-government framework

Jurisdiction	Instrument	Core Accountability Feature	India's Present Equivalent
India	Digital Personal Data Protection Act, 2023 (Rules and Board pending)	Data-processing consent, correction, and grievance-redress rights; Data Protection Board complaint mechanism envisaged but not yet operational	Enacted but not yet enforceable in practice; does not, on its text, create a general right to explanation of algorithmic decisions

Compiled by the author from the cited primary regulatory instruments, 2024.

Three lessons emerge from this comparison that are directly relevant to the Indian context. First, binding obligations attach to risk tiers rather than to AI as a general category, an approach compatible with India's existing sector-specific regulatory bodies (the Reserve Bank of India, the Securities and Exchange Board of India, the Election Commission) rather than requiring a single omnibus AI statute. Second, a standardised public disclosure register, of the kind the UK has now made mandatory at comparatively low administrative cost, would directly address the transparency gap identified in Section 4 without requiring the heavier conformity-assessment apparatus the EU has adopted, and could plausibly be implemented as a mandatory extension of existing e-governance portals. Third, none of the three comparator jurisdictions relies solely on data-protection law to deliver algorithmic accountability; each treats decisional transparency and impact assessment as functions distinct from, and additional to, data-protection compliance — a distinction that India's current reliance on the still-unimplemented DPDP Act as its primary AI-adjacent instrument does not yet reflect.

7. Towards a Framework for Accountable AI Governance in India

Drawing on the gaps identified in Sections 4 and 5 and the comparative lessons of Section 6, this paper proposes a four-pillar framework intended to be implementable within India's existing institutional architecture rather than requiring wholesale new statutory creation.

First, mandatory Algorithmic Impact Assessments (AIAs) for any AI system used in a decision affecting welfare eligibility, criminal investigation, or access to a public service, to be conducted before deployment and repeated at fixed intervals thereafter, modelled on the environmental-impact-assessment precedent already familiar to Indian administrative practice. An AIA should, at minimum, document training-data provenance, measured error rates disaggregated by relevant social categories, and the specific human-oversight point at which an affected citizen's case can be escalated to a human decision-maker.

Second, a tiered public transparency register, administered centrally (for Union schemes) and by state IT departments (for state schemes), requiring a plain-language disclosure of every AI system used in a public-facing administrative decision, following the UK Algorithmic Transparency Recording Standard model, but calibrated to Indian bandwidth and literacy constraints through a summary format available in regional languages and accessible via the same Common Service Centre network already used for other citizen services.

Third, an independent algorithmic audit function, situated either within an expanded mandate for the Comptroller and Auditor General or as a standing technical wing of the Data Protection Board of India once that Board is finally constituted and its Rules notified, empowered to conduct proactive (not merely

complaint-triggered) audits of high-risk government AI systems and to report findings to Parliament, addressing the capacity gap identified in Section 4 by concentrating scarce technical expertise in a body with a statutory audit mandate rather than relying on ad hoc RTI litigation.

Fourth, a guaranteed, time-bound human-appeal channel at the point of every automated denial or adverse algorithmic outcome affecting an individual citizen, with a maximum statutory response window comparable to existing grievance-redress timelines under the Public Grievance Redress framework, ensuring that the exclusion-error problem documented in the Aadhaar-DBT case study (Section 5.1) has a defined, enforceable remedy rather than depending on the discretion of a local functionary.

None of these four pillars requires India to slow the pace of AI adoption in governance; each is designed to run in parallel with continued deployment, consistent with the efficiency gains documented in Section 3, while closing the specific accountability gaps identified in Sections 4 and 5. Their common feature is that they shift the locus of accountability from an ex post, citizen-initiated RTI request or writ petition — the default route under current law — to an ex ante and continuous institutional obligation, which is the direction in which each of the comparator jurisdictions examined in Section 6 has already moved.

8. Conclusion

This paper has argued that artificial intelligence can strengthen good governance in India — the fiscal savings attributable to DBT and the scale of Aadhaar-linked service delivery are real and independently documented achievements — but that the accountability architecture required to sustain that strengthening over time has not kept pace with the technology's deployment. The gap is not primarily a gap in law on paper; India already has an RTI regime, a judiciary willing to review administrative action, an auditor-general, and, since August 2023, a data-protection statute that at least promises a regulatory board and citizen grievance-redress rights. The gap is in the specific fit between these instruments and the technical object — a trained model rather than a reasoned file note — that they are increasingly being asked to hold accountable, and, in the case of the DPDP Act, in the continuing delay between enactment and actual notification of the Rules and Board that would make the statute operative. Aadhaar-linked authentication failures, opaque predictive-policing deployments, and state-level procurement secrecy are not incidental defects; they are the predictable consequence of applying an accountability architecture designed for human bureaucratic discretion to a form of administration that operates on statistical inference at scale.

The four-pillar framework proposed here — algorithmic impact assessments, a public transparency register, an independent audit function, and a guaranteed human-appeal channel — is deliberately modest in its institutional ambition, building on existing bodies (the CAG, a fully constituted Data Protection Board, state IT departments) rather than proposing a new regulator. Its central claim is that accountability and efficiency are not opposed objectives in Indian AI governance; they are, if the comparative experience of the EU, the UK, and Singapore is any guide, complementary conditions for the same outcome — an administrative state that is fast without being unanswerable, and modern without ceasing to be democratic. As the DPDP Act's Rules and Board are finalised and as the IndiaAI Mission's five-year outlay translates, year on year, into further deployed systems, the next two to three years represent a narrow but real window in which India can close this gap before the scale of algorithmic administration makes retrofitting accountability substantially harder than building it in from the outset.

REFERENCES:

1. Bhatia, G., & Bhattacharya, S. (2022). The Right to Information Act and its uneven enforcement across Indian states. *Indian Journal of Public Administration*, 68(3), 412–430.
2. Bovens, M. (2007). Analysing and assessing accountability: A conceptual framework. *European Law Journal*, 13(4), 447–468.
3. Digital Personal Data Protection Act, 2023, No. 22, Acts of Parliament, 2023 (India).
4. European Parliament. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.
5. The Hindu BusinessLine. (2024, April 26). Saving through DBT rises to over ₹63,000 crore in FY23. The Hindu Group.
6. IndiaAI. (2024). Union Budget 2024-25 allocates over 550 crores to the IndiaAI Mission. Ministry of Electronics and Information Technology, Government of India.
7. Infocomm Media Development Authority of Singapore. (2020). Model artificial intelligence governance framework (2nd ed.). Personal Data Protection Commission Singapore.
8. Justice K. S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1 (Supreme Court of India).
9. Khera, R. (2019). Dissent on Aadhaar: Big data meets big brother. Orient BlackSwan.
10. Marda, V. (2018). Artificial intelligence policy in India: A framework for engaging the limits of data-driven decision-making. *Philosophical Transactions of the Royal Society A*, 376(2133), 20180087.
11. NITI Aayog. (2021). Responsible AI for all: Adopting the framework — A use case approach on facial recognition technology. Government of India.
12. OECD. (2019). Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449. Organisation for Economic Co-operation and Development.
13. Press Information Bureau. (2024, March 7). Cabinet approves ambitious IndiaAI Mission to strengthen the AI innovation ecosystem. Government of India.
14. UK Central Digital and Data Office. (2024). Algorithmic Transparency Recording Standard: Guidance for public sector organisations. Cabinet Office, United Kingdom.
15. United Nations Department of Economic and Social Affairs. (2022). E-Government Survey 2022: The future of digital government. United Nations.