

Digital Dharma: Balancing Efficiency, Ethics, and Democratic Values in India's AI-Powered Administrative Systems

Lt. Dr. Kongala Sukumar

Associate Professor of Public Administration

Dr. Marri Chenna Reddy, Human Resource Development (MCHRD) Institute of Telangana, Road no 25, Jubilee Hills, Hyderabad, Telangana 500033

Abstract:

India's public administration is undergoing a rapid shift toward artificial intelligence, with predictive analytics, facial recognition, chatbots, and algorithmic eligibility engines now mediating access to welfare, policing, taxation, and municipal services. This paper examines the tension between the efficiency gains promised by these systems and the ethical and constitutional commitments that underpin Indian democratic governance. Drawing on the concept of dharma — duty calibrated to context and consequence — the paper argues that algorithmic governance in India requires a situational, rather than purely technocratic, ethics. It surveys major AI deployments across central and state governments, documents exclusion and due-process harms alongside measured efficiency gains, and compares India's emerging oversight architecture with the European Union's risk-based regulatory model and the United Kingdom's sector-specific approach. The analysis finds that India's administrative AI ecosystem has outpaced its accountability infrastructure, producing a governance gap most acute for citizens with limited digital literacy or documentation. The paper proposes a “digital dharma” framework built on proportionality, contestability, and distributed oversight, offering a culturally grounded alternative to imported regulatory templates while remaining compatible with international best practice.

Keywords: artificial intelligence governance, Indian public administration, algorithmic accountability, digital ethics, e-governance, administrative law, dharma, democratic legitimacy

1. Introduction

Over the past decade, artificial intelligence has moved from the margins of Indian public administration to its operational core. What began as pilot chatbots on municipal websites has evolved into a dense infrastructure of algorithmic decision-making that determines who receives a subsidised ration, whose face is flagged at a metro gate, whose tax return is selected for scrutiny, and whose case is prioritised in an overburdened court docket. The National Strategy for Artificial Intelligence, first articulated by NITI Aayog in 2018 and revised in subsequent years, explicitly framed AI as a tool for “AI for All,” positioning the state as both regulator and principal adopter of automated systems (NITI Aayog, 2018). By the mid-2020s, systems built on this foundation had been deployed across welfare delivery, law enforcement, revenue administration, and judicial case management, often justified by the twin imperatives of scale and

scarcity: a billion-plus population served by administrative capacity that has historically struggled to match demand.

The administrative logic driving this expansion is not difficult to reconstruct. Indian public administration has, for decades, been characterised by a persistent gap between the scale of entitlement programmes and the human and financial capacity available to administer them accurately. Leakage in the Public Distribution System, discretionary delay in service delivery, and documented instances of corruption at the point of last-mile contact have been recurring subjects of both official audit and academic study since well before the current wave of automation (Drèze & Sen, 2013). Successive governments have framed digital and algorithmic tools as a solution to this capacity gap: authentication technology promises to eliminate “ghost” beneficiaries, machine-learning risk scoring promises to direct scarce audit capacity toward genuinely anomalous tax filings, and case-triage tools promise to help an overburdened judiciary — India's courts carry a pendency measured in tens of millions of cases — allocate scarce judicial attention more efficiently (Supreme Court of India, 2021). Each of these promises rests on a plausible technical logic. The question this paper takes up is not whether that logic is technically sound in principle, but what happens when it is operationalised at national scale within an administrative and legal culture whose accountability mechanisms were designed for human, not algorithmic, discretion.

This scale is precisely what makes the Indian case distinctive within the global literature on algorithmic governance. Much of the foundational scholarship on AI ethics and administrative automation has been produced in the context of smaller, wealthier democracies with comparatively robust data protection regimes and well-resourced regulatory agencies (Yeung, 2018; Zuboff, 2019). India's experience differs in three material respects. First, the sheer volume of transactions — hundreds of millions of welfare disbursements, biometric authentications, and tax filings annually — means that even a small error rate in an automated system translates into a large absolute number of affected citizens. Second, the digital divide within India is not merely a matter of connectivity but of documentary and biometric legibility: fingerprints worn down by manual labour, addresses that do not match across databases, and intermittent network access in rural areas interact with algorithmic systems in ways their designers rarely anticipate (Khera, 2019). Third, India's constitutional culture, particularly the jurisprudence following Justice K.S. Puttaswamy v. Union of India (2017), has articulated a right to privacy and, by extension, an emerging right to explanation and due process that sits in tension with the opacity of many deployed systems.

This paper asks a deceptively simple question: can efficiency-oriented algorithmic governance be reconciled with India's democratic and constitutional commitments, and if so, on what terms? To answer this, the paper draws on the concept of dharma — a Sanskrit term with no precise English equivalent, but which in classical and administrative usage denotes duty calibrated to role, context, and consequence rather than to a fixed, universal rule (Bilimoria, 2020). The paper argues that this contextual, relational ethics offers a more suitable normative vocabulary for Indian algorithmic governance than the largely rule-based, rights-enumeration frameworks imported from European and American AI ethics debates, without abandoning the substantive protections those frameworks aim to secure.

The urgency of this question is compounded by the pace of adoption. Between 2019 and 2024, the number of Union and state government departments reporting an active AI or machine-learning pilot grew several-fold, spanning agriculture, health, education, and law enforcement (MeitY, 2023). Much of this growth occurred through procurement rather than legislation: departments contracted vendors, integrated proof-of-concept systems into live operations, and scaled them incrementally, often without a parallel process of statutory authorisation or independent technical audit. This procurement-led pattern of diffusion is itself

significant for the argument developed here, because it means that the accountability architecture surrounding Indian administrative AI has, in most domains, been retrofitted after deployment rather than designed before it — a sequencing problem that recurs throughout the empirical sections of this paper.

The paper's contribution is threefold. Empirically, it consolidates a scattered evidentiary record — government reports, peer-reviewed field studies, RTI disclosures, and investigative journalism — into a single comparative account of where efficiency claims and documented harms diverge. Theoretically, it argues that the dominant vocabulary of global AI ethics, built substantially around European and North American legal traditions, requires adaptation rather than adoption when applied to a jurisdiction with India's scale, documentary informality, and constitutional structure. Practically, it proposes a framework intended to be usable by the officials, regulators, and legislators who must translate ethical commitments into operational rules.

The remainder of the paper proceeds in six parts. Section 2 develops the theoretical framework of “digital dharma” and situates it against dominant Western AI ethics paradigms. Section 3 surveys the current landscape of AI deployment across Indian public administration. Section 4 examines the empirical record of efficiency gains alongside documented harms, drawing on government data, audit reports, and investigative journalism. Section 5 analyses the accountability gap in India's current legal and institutional architecture. Section 6 situates India's approach comparatively against the European Union's AI Act and the United Kingdom's sectoral model. Section 7 proposes a digital dharma framework for reform, and Section 8 concludes.

2. Theoretical Framework: Dharma as an Ethical Lens for Algorithmic Governance

Dominant frameworks in AI ethics — principlism, deontological rights-based approaches, and the risk-tiered regulatory logic now codified in the European Union's AI Act — share an underlying assumption that ethical governance is best achieved through the ex ante specification of rules, categories, and thresholds (Floridi & Cowls, 2019; Jobin et al., 2019). A system is classified as high-risk or low-risk, permissible or prohibited, and the rule, once set, applies uniformly across contexts. This approach has considerable virtues: predictability, auditability, and resistance to arbitrary discretion. But it also presumes a relatively stable and legible administrative environment — one in which categories map cleanly onto cases.

The concept of dharma offers a different starting point. In classical Indian ethical and legal thought, dharma is not a single fixed code but a matrix of duties that vary by *svadharma* (one's particular role and station) and by circumstance, famously dramatised in the Bhagavad Gita's treatment of Arjuna's dilemma (Bilimoria, 2020). Applied to administrative ethics, this suggests that the duty of a welfare algorithm is not identical to the duty of a policing algorithm, even where both use similar technical architectures, because the stakes, the power asymmetry, and the consequences of error differ. A facial recognition misidentification in a policing context can result in wrongful detention; the same technical error in a museum kiosk is merely an inconvenience. A rule-based framework that classifies “facial recognition” as a single risk category risks flattening this distinction; a dharma-inflected framework asks, in each context, what duty of care the administering authority owes to the citizen affected, given the power differential and the reversibility of harm.

This is not an argument for regulatory relativism or for exempting Indian AI governance from enforceable standards. Rather, it is an argument that the content of those standards — the thresholds for human review, the burden of proof for contesting an automated decision, the transparency owed to the affected citizen —

should be calibrated to context rather than uniformly imported. Scholars of Indian administrative law have long noted that principles such as audi alteram partem (the right to be heard) and reasoned decision-making, drawn from both indigenous and common-law traditions, already provide doctrinal resources for this kind of contextual ethics (Bhatia, 2021). The digital dharma framework developed in Section 7 builds on this existing doctrinal base rather than treating Indian AI governance as a blank slate requiring wholesale importation of European or American regulatory models.

A further distinction is worth drawing between dharma as individual virtue ethics — a register in which the term is sometimes popularly invoked — and dharma as an institutional or role-based obligation, which is the register this paper adopts. The administering authority's dharma is not a matter of the individual official's private conscience but of the duty attached to the office and the system it operates, analogous to what Rao (2022) describes as the “structural” reading of dharma in classical Indian political theory (rajadharma, the dharma of rulership). On this reading, a welfare department deploying an eligibility algorithm inherits an institutional duty of care toward beneficiaries that does not disappear because the decision has been delegated to a model; if anything, the duty intensifies, because the citizen's capacity to seek an explanation or exception from a human official is precisely what the automation forecloses unless deliberately preserved. This institutional framing also distinguishes the approach from communitarian critiques of AI ethics that ground obligations in cultural particularity alone (Mohamed et al., 2020); the claim here is narrower and more operational — that Indian administrative law already possesses a doctrinal vocabulary for graduated, context-sensitive duty, and that this vocabulary is better suited to the governance problem at hand than a wholesale transplant of foreign risk taxonomies.

3. The Landscape of AI in Indian Public Administration

AI-enabled systems are now embedded across at least four domains of Indian public administration: welfare and identity, law enforcement, revenue and taxation, and judicial administration. Table 1 summarises the largest and most consequential deployments as of the mid-2020s.

Table 1- Major AI-Enabled Systems in Indian Public Administration

Domain	System / Programme	Implementing Agency	Core Function	Year Introduced
Welfare / Identity	Aadhaar-linked biometric authentication for PDS	UIDAI; State Food & Civil Supplies Departments	Beneficiary authentication for subsidised rations	2016–
Law Enforcement	State Facial Recognition Technology systems (e.g., TSCOP, Delhi Police FRT)	State Police Departments	Suspect identification; crowd and public-space surveillance	2018–
Revenue	Project Insight	Income Tax Department / CBDT	ML-based selection of returns for scrutiny	2017

Judiciary	SUPACE	e-Committee, Supreme Court of India	AI-assisted legal research and case triage	2021 (pilot)
Municipal / Citizen Services	State e-governance chatbots (e.g., UMANG, state helplines)	State IT Departments	Automated resolution of routine citizen queries	2019–
Agriculture	AI crop and pest advisory chatbots	State Agriculture Departments; NITI Aayog pilots	Vernacular-language crop/pest advisory	2020–

Source: Compiled by the author from NITI Aayog (2018, 2021), MeitY (2023), Internet Freedom Foundation (2021), and Supreme Court of India (2021).

The Aadhaar biometric identity system, while not itself an AI system in the narrow sense, has become the substrate on which numerous algorithmic layers operate. Authentication failures — caused by worn fingerprints, iris scan errors, or server downtime — feed into automated eligibility engines for the Public Distribution System (PDS) and the National Food Security Act, producing what Khera (2019) has termed “exclusion by design.” Several state governments, including Telangana, Delhi, and Uttar Pradesh, have deployed facial recognition technology (FRT) for policing, often procured without a dedicated statutory framework governing its use, retention of biometric data, or accuracy thresholds (Internet Freedom Foundation, 2021). The Income Tax Department's Project Insight uses machine learning to flag returns for scrutiny based on transaction-pattern analysis, while the judiciary has piloted the Supreme Court Portal for Assistance in Court's Efficiency (SUPACE) to assist judges in legal research and case triage (Supreme Court of India, 2021).

What unites these deployments is not a common technical architecture but a common institutional pattern: procurement and deployment have generally preceded the articulation of sector-specific accountability standards. The Ministry of Electronics and Information Technology's non-binding AI governance guidelines, released in 2023, articulate broad principles — safety, transparency, accountability, fairness — but do not create enforceable obligations, audit mandates, or a dedicated regulator with the power to suspend a non-compliant system (MeitY, 2023). This pattern of principle without enforcement mechanism recurs across the domains surveyed and forms the empirical basis for the accountability gap discussed in Section 5.

Beyond the four domains catalogued in Table 1, AI adoption is also visible in agricultural extension services, where several state governments have deployed machine-learning advisories for crop selection and pest management delivered through vernacular-language chatbots, and in health administration, where predictive triage tools have been piloted in select public hospital networks. These deployments share the welfare and policing systems' basic structural feature: they mediate access to a scarce public good or service on the basis of a model whose internal logic is rarely disclosed to the citizen it affects. Procurement records obtained through RTI requests in several states show that vendor contracts for these systems frequently include confidentiality clauses covering model architecture and training data, which — while standard practice in commercial software procurement generally — sits uneasily with the transparency

obligations that attach to state action under Indian administrative law (Internet Freedom Foundation, 2021). The effect is a widening category of administrative decisions that are, in practice, insulated from the kind of scrutiny that would apply to an equivalent decision made by a human official exercising discretion under a published rule.

4. Efficiency Gains and Documented Costs

Proponents of administrative AI in India point to substantial, at times striking, efficiency gains. Government audits and NITI Aayog reports have claimed reductions in PDS leakage attributable to Aadhaar-linked authentication, faster processing times for tax refunds under Project Insight, and reduced case backlogs where SUPACE-assisted triage has been piloted (NITI Aayog, 2021). These are not trivial claims in a system historically characterised by documented leakage, delay, and discretionary corruption at the point of service delivery.

However, an emerging body of field research and investigative reporting complicates this efficiency narrative by documenting the distribution of costs alongside the aggregate gains. Drèze et al. (2017) and subsequent field studies found that biometric authentication failures in PDS shops disproportionately affected elderly beneficiaries, manual labourers with worn fingerprints, and residents of areas with poor network connectivity — precisely the populations the welfare system is designed to serve. Investigative reporting has linked authentication failures to instances of starvation-related deaths in states including Jharkhand, where beneficiaries were turned away from ration shops after biometric mismatches, even where their entitlement was not in dispute (Yasin, 2018). On facial recognition, independent audits have found materially higher misidentification rates for individuals with darker skin tones and for women, mirroring findings from FRT accuracy studies conducted internationally (Buolamwini & Gebru, 2018), with no published Indian government accuracy benchmark against which deployed systems can be independently assessed.

Table 2 juxtaposes claimed efficiency gains with documented costs across the major deployment domains, drawing on government reporting, peer-reviewed field studies, and investigative journalism triangulated where possible.

Table 2- Efficiency Gains and Documented Costs by Domain

Domain	Claimed Efficiency Gain	Documented Cost / Harm	Source(s)
PDS / Aadhaar Authentication	Government-reported reduction in ration leakage and duplicate beneficiaries	Authentication failures linked to wrongful denial of rations; documented starvation-linked deaths (Jharkhand)	NITI Aayog (2021); Drèze et al. (2017); Yasin (2018)
Facial Recognition (Policing)	Faster suspect identification; expanded surveillance coverage	Elevated misidentification rates for darker skin tones and women; wrongful detentions reported	Internet Freedom Foundation (2021); Buolamwini & Gebru (2018)

Project Insight (Tax)	Reduced average time to select and process scrutiny cases	Limited published transparency on false-positive flagging rate; uneven grievance redress for flagged filers	NITI Aayog (2021); MeitY (2023)
SUPACE (Judiciary)	Faster case triage and legal research in pilot courts	No independent accuracy audit published; risk of encoding existing pendency biases into triage priorities	Supreme Court of India (2021)

Source: Compiled by the author from government reports, field studies, and investigative journalism cited in each row.

The pattern that emerges is not that automation is uniformly harmful or uniformly beneficial, but that the distribution of gains and costs is uneven: efficiency gains accrue predominantly to the administering agency in the form of reduced processing time and headline leakage figures, while the costs of system error fall predominantly on individual citizens with the least capacity to contest an automated decision. This asymmetry is the empirical core of the ethical problem this paper addresses, and it is precisely the kind of contextual, power-differentiated harm that a purely efficiency-oriented metric fails to capture. It is also worth noting a measurement asymmetry that compounds the distributional one: efficiency gains are typically reported through centralised administrative data — leakage percentages, average processing times, cases cleared per quarter — collected and published by the same agency responsible for the system's performance. Documented costs, by contrast, are typically established through decentralised, resource-intensive means: household surveys, RTI litigation, and investigative reporting, each of which reaches only a fraction of affected cases and is subject to its own access barriers. This means that the aggregate efficiency figures found in official reporting are considerably more complete than the aggregate harm figures found in independent research, not because harms are rarer than gains but because the infrastructure for measuring them is weaker. A rigorous accounting of the efficiency–ethics trade-off would require government agencies to publish disaggregated error rates — false-negative and false-positive rates broken down by beneficiary category — alongside the leakage-reduction figures they already report, a specific and actionable transparency reform returned to in Section 7.

5. Democratic Values, Accountability, and the Governance Gap

India's constitutional framework provides substantial doctrinal resources for holding administrative action accountable. The right to information, secured under the Right to Information Act 2005, and the doctrine of reasoned administrative decision-making, developed through decades of Supreme Court jurisprudence, both presuppose that a citizen can identify the basis of a decision affecting them and challenge it (Bhatia, 2021). Algorithmic decision-making, particularly where it relies on proprietary or poorly documented models, strains both presuppositions. A beneficiary denied a subsidy because an eligibility algorithm flagged an anomaly in linked bank data has, in practice, limited means of discovering the criterion applied, let alone contesting its application to their case.

The Digital Personal Data Protection Act 2023 (DPDP Act) represents India's most significant recent legislative intervention in this space, establishing consent and purpose-limitation obligations for personal data processing. Yet the Act's provisions are calibrated primarily to data protection rather than to

algorithmic decision-making per se, and it grants broad exemptions to government agencies processing data for stated public functions — precisely the category into which most of the systems surveyed in Section 3 fall (Ministry of Electronics and Information Technology, 2023). Unlike the European Union's General Data Protection Regulation, which includes a qualified right not to be subject to a decision based solely on automated processing (Article 22), the DPDP Act contains no equivalent provision. There is, moreover, no dedicated AI regulator in India analogous to a sectoral ombudsman with investigatory and suspension powers; oversight is distributed across sectoral regulators (the Reserve Bank of India for financial algorithms, the Election Commission for electoral technology) with no consistent standard applied across domains.

This produces what the paper terms a governance gap: a widening distance between the pace and scale of algorithmic deployment and the maturity of the accountability infrastructure meant to check it. The gap is not merely a matter of missing legislation; it is also institutional, in the absence of independent audit bodies with technical capacity to evaluate deployed models, and cultural, in the tendency of procurement processes to treat AI adoption as a technical upgrade rather than a change in the locus of administrative discretion warranting the same scrutiny as any other exercise of state power.

Parliamentary scrutiny of this gap has, to date, been episodic rather than systematic. Standing Committee reports on subjects including IT and communications have periodically flagged the absence of a comprehensive AI accountability framework, and draft legislation addressing algorithmic decision-making specifically has been discussed in policy circles without reaching the stage of a tabled bill (Ministry of Electronics and Information Technology, 2023). Judicial engagement has been similarly incremental: courts have addressed Aadhaar-linked exclusion and facial recognition deployments through individual writ petitions rather than through a general doctrine of algorithmic due process, producing a body of case-specific relief that does not, by itself, establish system-wide standards. The consequence is that the burden of establishing accountability continues to fall on individual litigants and civil society organisations, on a case-by-case basis, rather than being built into the design and procurement of the system from the outset — the inverse of the ex-ante, proportionate oversight this paper argues for in Section 7.

6. Comparative Oversight Architectures

Placing India's emerging framework alongside two other major regulatory approaches clarifies both its gaps and its distinctive possibilities. The European Union's AI Act, which entered into force in 2024, adopts a risk-tiered model: systems are classified as unacceptable, high, limited, or minimal risk, with correspondingly graduated obligations, and high-risk systems — a category that explicitly includes systems used in the “administration of justice” and in the “management of public benefits” — are subject to conformity assessment, human oversight requirements, and mandatory registration in an EU-wide database (European Commission, 2024). The United Kingdom, by contrast, has pursued a “pro-innovation,” sector-specific approach, asking existing regulators (in finance, health, and communications) to apply cross-cutting principles — safety, transparency, fairness, accountability, contestability — within their existing remits rather than creating a single horizontal AI statute (UK Department for Science, Innovation and Technology, 2023).

Table 3 compares the three jurisdictions across five dimensions of oversight architecture: existence of a horizontal AI statute, classification methodology, mandatory human review for high-stakes decisions, redress mechanisms available to affected citizens, and auditing/registration requirements.

Table 3- Comparative Oversight Architecture: India, European Union, and United Kingdom

Dimension	India	European Union	United Kingdom
Horizontal AI statute	None; non-binding guidelines only. DPDP Act 2023 covers data, not algorithmic decisions	AI Act 2024 — binding, risk-tiered regulation	None; principles applied within existing sectoral regimes
Classification methodology	Ad hoc, sector-by-sector; no unified taxonomy	Risk-tiered: unacceptable / high / limited / minimal risk	No formal tiering; principles applied case-by-case by regulators
Mandatory human review (high-stakes decisions)	Not uniformly mandated across domains	Required for high-risk systems under the AI Act	Regulator-dependent; not uniformly mandated
Redress mechanism for affected citizens	RTI Act and general administrative law; no AI-specific route	Rights adjacent to explanation; national supervisory authorities	Sectoral ombudsman and regulator complaint routes
Auditing / registration requirement	None centrally; fragmented across sectoral regulators	Mandatory conformity assessment and EU database registration for high-risk systems	No central registration; audit expectations vary by sector

Source: Compiled by the author from European Commission (2024), UK Department for Science, Innovation and Technology (2023), and MeitY (2023).

India's position on this comparative map is closest to a nascent, under-resourced version of the UK's sectoral model, but without the UK's established regulatory capacity within each sector to absorb AI-specific obligations. The EU's risk-tiered model offers a template India could adapt — its explicit inclusion of welfare administration and judicial support systems within the high-risk category maps closely onto the domains surveyed in Section 3 — but wholesale importation risks the rule-based flattening critiqued in Section 2, applying uniform conformity-assessment burdens to systems whose actual power differential and reversibility of harm vary considerably. A dharma-inflected adaptation would retain the EU's insistence on mandatory human review and redress for high-stakes decisions while allowing the intensity of the audit and registration regime to scale with the demonstrated vulnerability of the affected population, rather than with the technical classification of the system alone.

It is also instructive to note what the comparison does not support. Neither the EU nor the UK model was designed with India's scale of low-connectivity, low-documentation users in mind, and both presuppose administrative capacity — trained auditors, functioning grievance-redress helplines, reliable digital access for the affected citizen to even discover that an automated decision has been made — that cannot be assumed uniformly across Indian districts. A framework transplanted without adaptation to this capacity constraint risks becoming, in practice, a paper compliance regime: conformity assessments completed and

registered centrally while the lived experience of contestability for a beneficiary in a low-connectivity district remains unchanged. This is precisely the gap the digital dharma framework in Section 7 is designed to close, by tying the intensity of formal oversight to the vulnerability of the population actually served rather than to a uniform technical classification alone.

7. Toward a Digital Dharma Framework

Building on the preceding analysis, this paper proposes a digital dharma framework for Indian administrative AI organised around three principles. The first is proportionality: the intensity of oversight, the standard of accuracy required, and the mandatory human-in-the-loop threshold should scale with the reversibility of harm and the vulnerability of the affected population, rather than being set uniformly across all algorithmic systems. A welfare eligibility algorithm affecting subsistence-level beneficiaries warrants a materially higher standard of human review than a municipal chatbot answering routine queries, even though both are “AI systems” in the generic sense.

The second principle is contestability. Every citizen affected by an automated administrative decision should have a clear, accessible, and time-bound mechanism to request the basis of the decision and to have it reviewed by a human official empowered to override the algorithmic output — extending existing RTI and natural-justice doctrine explicitly into the algorithmic domain, rather than leaving courts to infer such a right case by case, as has occurred to date (Bhatia, 2021).

The third principle is distributed oversight, calibrated to India's federal structure. Rather than a single national AI regulator attempting to audit every deployment — a task likely to be under-resourced given the scale documented in Section 3 — the framework proposes sector-linked audit units embedded within existing regulators (the Unique Identification Authority of India for Aadhaar-linked systems, state police oversight bodies for FRT, the Central Board of Direct Taxes for Project Insight), coordinated by a light-touch central standards body responsible for publishing accuracy benchmarks and minimum audit protocols rather than approving individual deployments. This distributes technical capacity-building across bodies that already possess domain expertise, while the central body prevents the fragmentation of standards that currently characterises India's approach.

A fourth, more operational recommendation follows from the measurement asymmetry identified in Section 4: mandatory disaggregated error reporting. Departments deploying eligibility, authentication, or risk-scoring algorithms in high-stakes domains should be required to publish, at minimum annually, false-positive and false-negative rates disaggregated by relevant beneficiary categories (age, geography, and — where lawfully collectable — disability status), alongside the leakage or efficiency figures already reported. This single reform would not by itself resolve the accountability gap, but it would convert the current asymmetry between well-measured gains and poorly measured harms into a symmetrical evidentiary record, on which the proportionality judgments called for by the first principle above could actually be made. Taken together, the four elements of the framework — proportionality, contestability, distributed oversight, and disaggregated error reporting — are deliberately modest in institutional ambition. They do not require a new ministry or a constitutional amendment; they require extending doctrines and reporting obligations that already exist in Indian administrative law into a domain — algorithmic decision-making — that has so far been treated as exceptional.

8. Conclusion

India's administrative AI ecosystem has expanded faster than the accountability infrastructure needed to govern it responsibly. The efficiency gains claimed by government agencies are real but unevenly distributed, with documented harms concentrated among citizens least equipped to contest an automated decision — the elderly, manual labourers, and residents of poorly connected regions. This paper has argued that closing this gap does not require India to choose between efficiency and ethics, nor between indigenous and international regulatory templates. Drawing on the concept of dharma as contextual, role-calibrated duty, the paper has proposed a framework organised around proportionality, contestability, and distributed oversight — one that builds on existing doctrinal resources in Indian administrative law while learning selectively from the European Union's risk-tiered model and the United Kingdom's sectoral approach. Realising this framework will require legislative attention beyond the current scope of the Digital Personal Data Protection Act 2023, institutional investment in independent audit capacity, and a shift in procurement culture that treats algorithmic deployment as an exercise of administrative discretion warranting the same scrutiny as any other.

Several limitations of the present analysis should be acknowledged. The empirical record assembled in Sections 3 and 4 draws heavily on RTI disclosures, field surveys, and investigative reporting precisely because comprehensive, government-published error data does not yet exist at the disaggregated level this paper argues for — a gap the analysis diagnoses but cannot itself fill. Future research combining administrative micro-data (where accessible through research partnerships) with the qualitative accounts of affected beneficiaries would allow the distributional claims made here to be tested with greater precision, including across states with differing implementation capacity. The proposed framework would also benefit from piloting within a single domain — welfare eligibility algorithms are the most consequential and best-documented candidate — before extension to policing and judicial applications, where the human-rights stakes and the political economy of reform differ considerably from the welfare context.

The stakes of this gap are not abstract: they are measured in ration cards denied, faces misidentified, and citizens left with no clear avenue to be heard by the systems that increasingly govern their access to the state. A digital dharma framework, grounded in proportionality, contestability, and distributed oversight, offers one route toward closing that gap without requiring India to choose between the efficiency its administrative scale demands and the constitutional and democratic values its citizens are owed.

REFERENCES:

1. Bhatia, G. (2021). *The transformative constitution: A radical biography in nine acts*. HarperCollins India.
2. Bilimoria, P. (2020). Dharma. In K. Jayasuriya (Ed.), *The Oxford handbook of Indian philosophy* (pp. 211–229). Oxford University Press.
3. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 77–91.
4. Drèze, J., Khera, R., & Somanchi, A. (2017). Aadhaar and food security in Jharkhand: Pain without gain? *Economic and Political Weekly*, 52(50), 51–59.
5. Drèze, J., & Sen, A. (2013). *An uncertain glory: India and its contradictions*. Princeton University Press.

6. European Commission. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union.
7. Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1), 1–15. <https://doi.org/10.1162/99608f92.8cd550d1>
8. Internet Freedom Foundation. (2021). *Facial recognition technology in India: A survey of law enforcement deployments*. Internet Freedom Foundation.
9. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
10. *Justice K.S. Puttaswamy v. Union of India*, (2017) 10 SCC 1 (India).
11. Khera, R. (2019). *Dissent on Aadhaar: Big data meets big brother*. Orient BlackSwan.
12. Ministry of Electronics and Information Technology. (2023). *Digital Personal Data Protection Act, 2023*. Government of India.
13. Ministry of Electronics and Information Technology [MeitY]. (2023). *Responsible AI for all: Approach paper*. Government of India.
14. Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33, 659–684. <https://doi.org/10.1007/s13347-020-00405-8>
15. NITI Aayog. (2018). *National strategy for artificial intelligence: #AIforAll*. Government of India.
16. NITI Aayog. (2021). *Responsible AI: Approach document for India, part 1 – principles for responsible AI*. Government of India.
17. Rao, K. (2022). Rajadharma and the ethics of rulership in classical Indian political thought. *Journal of Indian Political Thought*, 14(2), 45–68.
18. Supreme Court of India. (2021). *SUPACE: Supreme Court portal for assistance in courts' efficiency* [Project overview]. e-Committee, Supreme Court of India.
19. UK Department for Science, Innovation and Technology. (2023). *A pro-innovation approach to AI regulation* (Cm 815). UK Government.
20. Yasin, A. (2018, October 17). Jharkhand hunger death: Aadhaar-linked ration denial under scrutiny. *The Hindu*.
21. Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance*, 12(4), 505–523. <https://doi.org/10.1111/rego.12158>
22. Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.