

Between Surveillance and Service: Reconciling AI-Driven Governance with Democratic Legitimacy in the Indian State

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:

The rapid diffusion of artificial intelligence (AI) into Indian public administration has produced a state that is simultaneously more capable of delivering welfare and more capable of watching its citizens. This paper examines the structural tension between AI-driven governance as an instrument of service delivery and as an instrument of surveillance, and asks whether India's constitutional and administrative architecture can reconcile the two without eroding democratic legitimacy. Drawing on the National Strategy for Artificial Intelligence, the Aadhaar biometric identity system, facial recognition deployments by state police forces, and the Digital Personal Data Protection Act, 2023, the paper develops a three-part legitimacy framework, procedural, substantive, and relational, to evaluate algorithmic governance practices. It argues that the same infrastructural capacities that enable direct benefit transfer and grievance redressal also enable predictive policing and biometric tracking, and that legitimacy depends less on the technology itself than on the accountability architecture surrounding it. The paper compares Indian arrangements with the European Union's risk-based regulatory model and identifies gaps in independent oversight, algorithmic audit capacity, and judicial review. It concludes with a policy framework for embedding proportionality, purpose limitation, and contestability into AI-enabled state functions.

Keywords: Artificial intelligence governance; democratic legitimacy; surveillance state; Aadhaar; algorithmic accountability; Digital Personal Data Protection Act; India

1. Introduction

Over the past decade, the Indian state has undergone what might be described as an infrastructural transformation, in which artificial intelligence and allied digital technologies have moved from pilot projects into the routine machinery of governance. The world's largest biometric identity system, Aadhaar, now underwrites subsidy disbursement, banking access, and welfare eligibility for well over a billion residents. Direct Benefit Transfer (DBT) schemes route enormous sums annually through algorithmically verified bank accounts. State police departments have deployed facial recognition systems at railway stations, protests, and public gatherings. Tax authorities use predictive analytics to flag returns for scrutiny, and municipal corporations experiment with AI-driven chatbots for grievance redressal. Each of these initiatives sits at the intersection of two governance logics that are rarely examined together: the logic of

service, in which AI is deployed to widen access, reduce leakage, and improve responsiveness, and the logic of surveillance, in which the same computational infrastructure is repurposed to monitor, classify, and, at times, discipline citizens.

This duality is not accidental. The technical architecture that allows a welfare department to verify a beneficiary's identity in real time is structurally similar to the architecture that allows a police department to match a protester's face against a watchlist. Both depend on centralized biometric or behavioural data, both rely on probabilistic matching rather than deterministic proof, and both are typically deployed with limited independent oversight (Singh, 2019). The question this paper addresses is not whether AI should be used in Indian governance, a question effectively settled by the scale of existing deployment, but whether the institutional arrangements surrounding that use are adequate to sustain democratic legitimacy. Democratic legitimacy, as used here, refers to more than an electoral mandate. Following administrative law scholarship, legitimacy is treated as resting on three related grounds: procedural legitimacy, which asks whether decisions are made through transparent, contestable, and rule-bound processes; substantive legitimacy, which asks whether outcomes are proportionate, non-discriminatory, and aligned with constitutional values; and relational legitimacy, which asks whether the relationship between citizen and state produced by the technology preserves dignity, consent, and a meaningful right of appeal (Citron & Pasquale, 2014; Yeung, 2018). AI-driven governance in India, this paper argues, has advanced furthest on capability and slowest on all three legitimacy dimensions.

The paper proceeds as follows. Section 2 situates the discussion within the comparative literature on algorithmic governance and the specific trajectory of Indian digital state-building. Section 3 maps the landscape of AI-enabled government functions, distinguishing service-oriented from surveillance-oriented deployments. Section 4 examines the surveillance dimension in detail, focusing on Aadhaar-linked authentication, facial recognition, and predictive policing. Section 5 turns to the service dimension, focusing on direct benefit transfer, grievance redressal, and administrative decision support. Section 6 evaluates these deployments against the procedural, substantive, and relational legitimacy framework, drawing comparisons with the European Union's risk-tiered regulatory approach. Section 7 proposes a reconciliation framework built around proportionality, purpose limitation, and contestability. Section 8 concludes with recommendations for legislators, courts, and administrators.

2. Theoretical Framework and Literature Review

2.1 Algorithmic governance and the administrative state

The turn toward automated decision-making in public administration has generated a substantial literature on what scholars term algorithmic governance: the use of data-driven, often opaque, computational systems to perform functions historically reserved for human bureaucratic discretion (Yeung, 2018). Citron and Pasquale (2014) argue that automated scoring systems, whatever their domain, tend to migrate consequential decisions away from due-process-bound procedures and toward statistically derived predictions that are difficult for affected individuals to contest. Eubanks (2018) extends this argument to welfare administration specifically, showing how automated eligibility and risk-scoring systems in the United States disproportionately burden low-income populations who are least equipped to challenge erroneous determinations. Zuboff (2019) situates this trend within a broader account of surveillance capitalism, in which behavioural and biometric data become a form of extractable value, whether the extracting entity is a private firm or a state agency. These strands of literature converge on a common

insight: automation does not simply digitize existing bureaucratic processes; it restructures the relationship between the citizen and the institution, often in ways that concentrate informational power in the hands of the administrator while diffusing the burden of error onto the individual.

2.2 The Indian digital state trajectory

India's experience departs from the Euro-American literature in scale and in the centrality of identity infrastructure. Singh's (2017, 2019) ethnographic studies of Aadhaar describe the project as an attempt to build a foundational, platform-like layer of identity verification beneath the country's welfare bureaucracy, an ambition captured in the Hindi word *aadhaar* itself, meaning foundation. Because Aadhaar authentication became a near-universal gateway to subsidised food grains, cooking gas, pensions, and banking access, its failures, such as biometric mismatches for manual labourers or elderly residents, translate directly into denial of essential services. This is the seam at which the service and surveillance logics meet: the same authentication event that grants access to a ration also generates a durable, centrally logged record of the citizen's transaction, location, and biometric profile. NITI Aayog's National Strategy for Artificial Intelligence (2018) and the subsequent Responsible AI approach papers (NITI Aayog, 2021) frame AI primarily as a growth and inclusion strategy, identifying healthcare, agriculture, education, smart cities, and mobility as priority sectors, but they devote comparatively little attention to the surveillance applications that have grown alongside these service-oriented use cases.

2.3 Legitimacy as a multidimensional construct

This paper adopts a three-part legitimacy framework rather than treating legitimacy as a single, binary property. Procedural legitimacy concerns the transparency and contestability of the process by which an algorithmic decision is reached: is the logic of the system disclosed, is there a right to explanation, and can an adverse decision be appealed to a human authority? Substantive legitimacy concerns the fairness of outcomes: does the system produce proportionate, non-discriminatory results, and is its use proportionate to the stated governmental purpose? Relational legitimacy concerns the quality of the citizen-state relationship the technology produces: does it preserve dignity and consent, or does it normalise continuous monitoring as a condition of accessing basic entitlements? The Puttaswamy judgment of 2017, in which a nine-judge bench of the Supreme Court of India held that privacy is a fundamental right protected under Articles 14, 19, and 21 of the Constitution, supplies the constitutional anchor for all three dimensions, since the Court explicitly required that any state incursion into informational privacy satisfy a standard of legality, necessity, and proportionality. The remainder of this paper uses that tripartite proportionality standard as the benchmark against which specific AI deployments are assessed.

2.4 Federalism and the layered nature of Indian digital governance

A further complication specific to the Indian case is federalism. Policing, and therefore most facial recognition and predictive law-enforcement deployment, falls constitutionally within the domain of state governments, while identity infrastructure, data protection legislation, and civil aviation policy are largely central subjects. This layering means that a coherent national legitimacy standard cannot simply be legislated once at the centre; it must be capable of binding state police departments that procure surveillance technology through local budgets and vendor contracts largely insulated from central oversight. Comparative federal systems, including the United States, face a structurally similar problem, where state and municipal law-enforcement agencies adopt AI tools well ahead of federal guidance, producing the patchwork pattern noted in Table 3 below. The Indian case is distinguished, however, by

the near-universal reach of a single central identity system, Aadhaar, that state-level surveillance systems can draw upon even where the surveillance application itself is a state, not a central, initiative. This vertical integration of identity data across a federal structure is precisely the seam at which the service and surveillance logics of the state most easily blur into one another.

3. Mapping AI-Driven Governance in India

AI-enabled government functions in India can be arranged along a spectrum running from purely service-oriented applications, whose principal purpose is to widen access or improve efficiency, to purely surveillance-oriented applications, whose principal purpose is monitoring, verification, or law enforcement. In practice, most large deployments combine elements of both, since the infrastructure built for one purpose is frequently repurposed for the other. Table 1 summarises a representative set of AI-enabled initiatives across central ministries and select states, indicating the nominal governance function and the extent to which each also carries a monitoring or verification dimension.

Table 1. Selected AI-Enabled Government Initiatives in India

Initiative	Implementing Agency	Primary Function	Governance Dimension	Year
Aadhaar Authentication	UIDAI	Identity verification for welfare and banking access	Service + Surveillance	2010
Direct Benefit Transfer (DBT)	Ministry of Finance / DBT Mission	Algorithmically verified subsidy disbursement	Service	2013
TSCOP Mobile Application	Telangana Police	Field policing support, facial matching	Surveillance	2016
Automated Facial Recognition System (AFRS)	National Crime Records Bureau	Criminal identification	Surveillance	2020
Digi Yatra	Ministry of Civil Aviation	Facial-recognition airport boarding	Service + Surveillance	2022
AI Grievance Redressal Chatbots	Various state e-governance departments	Citizen grievance handling	Service	2019 onward
Integrated Command and Control Centre	Telangana / Hyderabad Police	City-wide camera-based surveillance	Surveillance	2021

Source: Compiled by the author from NITI Aayog (2018), Internet Freedom Foundation (2023), and Software Freedom Law Centre (2023).

Three patterns emerge from this landscape. First, identity infrastructure is the common substrate: Aadhaar-linked authentication underlies not only DBT and banking access but also, indirectly, several facial recognition and law-enforcement systems that draw on government-issued identity databases for training data or verification. Second, the institutional home of a given system, whether a welfare ministry, a police department, or a municipal corporation, shapes its legal basis and oversight far more than the underlying technology does; a facial matching algorithm used by an airport to speed boarding and one used by a police department to identify protesters may be functionally similar but sit under entirely different accountability regimes. Third, service-oriented systems have generally been introduced through executive notification and scheme guidelines rather than dedicated primary legislation, while surveillance-oriented systems have often been procured through state police budgets with minimal public consultation, a pattern documented by digital rights organisations tracking facial recognition procurement (Software Freedom Law Centre, 2023).

There is also significant variation across states in the pace and transparency of adoption. Telangana has emerged as the state with the most extensive documented footprint of AI-enabled policing infrastructure, including the TSCOP application and the Hyderabad Integrated Command and Control Centre, while several other states remain at earlier tender or pilot stages. This unevenness matters for a legitimacy analysis because it means that the practical exposure of an Indian citizen to AI-enabled surveillance depends heavily on state of residence, a variation that has no equivalent on the service side of the ledger, where DBT and Aadhaar-linked welfare payments apply with substantially greater national uniformity. A legitimacy framework calibrated only at the national level therefore risks understating the intensity of surveillance exposure in the states where deployment has advanced furthest.

4. The Surveillance Dimension: Identity, Recognition, and Prediction

4.1 Aadhaar and biometric authentication

Aadhaar remains the foundational surveillance-capable infrastructure in Indian governance, not because it was conceived as a policing tool, but because its architecture centralises biometric and demographic data at a scale that makes secondary use difficult to prevent. Authentication failures, whether from worn fingerprints, network outages in rural areas, or server downtime, have been documented to result in denial of rations and pensions to eligible beneficiaries, a problem Singh (2017) describes as the system achieving inclusion only unevenly, at the seams between interlocking bureaucracies. The Puttaswamy litigation that ultimately established privacy as a fundamental right originated precisely from a challenge to Aadhaar's mandatory linkage requirements, underscoring how service delivery and rights-based objection have been intertwined from the project's inception.

4.2 Facial recognition technology

Facial recognition technology (FRT) has expanded rapidly across Indian policing since the late 2010s. Tracking by the Internet Freedom Foundation's Project Panoptic identified well over one hundred FRT systems in various stages of procurement, tender, or operation across Indian states, with Telangana emerging as the state with the largest concentration of deployed projects, including the TSCOP mobile policing application and integration with the Crime and Criminal Tracking Network and Systems. Delhi Police has used facial recognition to screen individuals at protests and, in one widely reported instance, to identify more than a thousand people allegedly involved in communal violence in the national capital by matching images against government identity databases. These deployments have typically proceeded

without dedicated statutory authorisation, relying instead on general policing powers, and civil society groups have repeatedly called for a moratorium pending the introduction of purpose-specific legislation and independent accuracy audits (Software Freedom Law Centre, 2023).

4.3 Data security incidents and the limits of self-regulation

The surveillance dimension of AI governance is further complicated by data security failures within the very databases these systems depend on. Reported breaches involving state police data repositories, including an incident affecting Telangana Police systems, have raised concerns that biometric and behavioural data collected for one stated purpose can circulate well beyond the control of the collecting agency once it is compromised or informally shared between departments. Because biometric identifiers, unlike passwords, cannot be reset once exposed, such incidents carry a different order of risk than conventional data breaches, and they underscore why relational legitimacy, the durability of citizen trust in the state's custodianship of sensitive data, cannot be assessed independently of the state's demonstrated data security practices. The Digital Personal Data Protection Act, 2023 introduces breach-notification obligations for data fiduciaries generally, but its exemptions for government processing undertaken in the interest of state security limit its practical reach over precisely the systems most implicated in these incidents (Ministry of Electronics and Information Technology, 2023).

4.4 Predictive and automated decision systems

Beyond identity verification, Indian tax and law-enforcement agencies have adopted predictive analytics for tasks such as flagging returns for scrutiny and, in some states, forecasting crime hotspots. These systems share the due-process concerns identified in the broader algorithmic accountability literature: the statistical models are generally proprietary or otherwise undisclosed, the criteria for flagging an individual or location are not published, and there is no standard mechanism for an affected party to obtain an explanation or contest the classification (Citron & Pasquale, 2014). Table 2 summarises the legal basis, deploying authority, and disclosed oversight mechanism for a representative sample of surveillance-oriented AI systems.

Table 2. Legal Basis and Oversight of Selected Surveillance-Oriented AI Systems

System	Deploying Authority	Stated Purpose	Legal Basis	Disclosed Oversight
Aadhaar Authentication Logs	UIDAI	Welfare identity verification	Aadhaar Act, 2016	UIDAI regulations; limited external audit
Delhi Police FRT	Delhi Police	Missing-persons tracing; public order	General policing powers (CrPC)	None publicly disclosed
TSCOP / CCTNS (Telangana)	Telangana Police	Field intelligence, criminal tracking	State police acts; MHA scheme guidelines	Internal departmental review

System	Deploying Authority	Stated Purpose	Legal Basis	Disclosed Oversight
Hyderabad Command and Control Centre	Telangana / Hyderabad Police	City-wide surveillance	State police acts	None publicly disclosed
Digi Yatra Facial Boarding	Airport operators (PPP)	Passenger facilitation	Civil aviation policy; user consent	Airline / airport operator terms of use

Source: Compiled by the author from Internet Freedom Foundation, Project Panoptic (2023), and Software Freedom Law Centre (2023).

The pattern visible in Table 2 is a consistent gap between technological capability and statutory authorisation. Where service-oriented systems generally operate under scheme guidelines issued by the concerned ministry, surveillance-oriented systems most often operate under the general investigative powers of the Code of Criminal Procedure or state police acts, instruments that were not drafted with algorithmic identification in mind and that provide no bespoke mechanism for audit, redress, or sunset review.

5. The Service Dimension: Delivery, Redressal, and Administrative Support

The service-oriented face of AI-driven governance is most visible in Direct Benefit Transfer, which routes subsidies for food, fuel, and cash entitlements directly into Aadhaar-linked bank accounts, substantially reducing the leakage associated with paper-based ration and subsidy systems. Municipal and state governments have also introduced AI-driven chatbots and virtual assistants for grievance redressal, intended to reduce the time and effort citizens spend navigating bureaucratic counters, and several states have piloted machine-translation and speech-recognition tools to extend digital government services to residents who are not literate in English or Hindi script. In healthcare, AI-assisted diagnostic support has been piloted in public hospitals for radiology and pathology triage, an application NITI Aayog (2018) specifically identified as high-value given the shortage of specialist personnel relative to India's population.

These service-oriented applications generally enjoy stronger public legitimacy than their surveillance counterparts, both because their stated purpose maps more directly onto a recognisable public good and because failure modes, a denied subsidy or a mistranslated form, are more visible and more readily contested through existing grievance mechanisms than the failure modes of a facial recognition mismatch. Even so, service-oriented systems are not free of the legitimacy concerns raised in Section 2. Automated eligibility determination for welfare schemes can replicate the same opacity problems documented in comparable systems abroad (Eubanks, 2018): a beneficiary removed from a scheme because an algorithm flagged a data mismatch frequently has no accessible channel to learn why the flag was raised or to correct the underlying record. The distinction between service and surveillance, in other words, is a distinction of stated purpose and public perception rather than of underlying legitimacy risk; both categories of system

depend on the same categories of data, and both are capable of producing the same due-process failures if left without adequate procedural safeguards.

A further consideration on the service side is the digital divide itself. AI-mediated grievance redressal, chatbot interfaces, and app-based service delivery presuppose a baseline of smartphone access, data connectivity, and digital literacy that remains unevenly distributed across rural and urban India, and across age, gender, and income groups. Citizens who cannot navigate a chatbot or complete an app-based authentication step are frequently pushed toward the same physical counters the technology was meant to reduce reliance on, or, worse, toward informal intermediaries who complete the digital step on their behalf in exchange for a fee, reproducing the very access barriers digitisation was intended to remove. NITI Aayog's own framing of AI as an inclusion strategy (NITI Aayog, 2018) is therefore only partially self-fulfilling: the service dimension of AI governance improves aggregate efficiency, but its distributional benefits depend on complementary investments in connectivity, device access, and assisted digital literacy that lie outside the AI system itself.

6. Evaluating Legitimacy: A Comparative Assessment

6.1 Procedural legitimacy

On procedural grounds, Indian AI governance exhibits substantial variation by sector. Welfare and taxation systems generally operate within published scheme guidelines and are subject to the Right to Information Act, 2005, which provides at least a formal channel for citizens to request information about how a decision was reached, even where practical responsiveness is limited. Surveillance systems, by contrast, are frequently procured and deployed without prior public notification, and their underlying matching algorithms are treated as vendor intellectual property, unavailable for independent audit. The Digital Personal Data Protection Act, 2023, passed by Parliament in August 2023, introduces a Data Protection Board and a consent-based processing framework, but it also permits government agencies to process personal data without consent for functions including state security and public order, a carve-out that directly affects the oversight of law-enforcement AI systems (Ministry of Electronics and Information Technology, 2023; PRS Legislative Research, 2023).

6.2 Substantive legitimacy

On substantive grounds, the central question is proportionality: does the scale and intrusiveness of a given AI deployment match its stated governmental purpose? Facial recognition deployed to identify missing children, one of the original justifications offered for Delhi Police's system, occupies different proportionality territory than the same system's later use to screen individuals at political gatherings. India currently lacks a general statutory proportionality test comparable to the necessity and proportionality standard articulated in *Puttaswamy*, applied consistently across agencies; instead, proportionality review occurs largely through post hoc litigation, an inherently reactive mechanism that offers little guidance to agencies before a system is deployed.

6.3 Relational legitimacy and comparative context

On relational grounds, the cumulative effect of biometric authentication, facial recognition, and predictive scoring is to normalise a state-citizen relationship premised on continuous verification, in which access to basic entitlements is conditioned on successful biometric matching. This dynamic contrasts with the risk-tiered approach adopted in the European Union's Artificial Intelligence Act, which classifies biometric identification and law-enforcement AI systems as high-risk or, in specific circumstances, prohibited, and

imposes mandatory conformity assessment, human oversight, and registration requirements before such systems may be deployed (European Commission, 2021). Table 3 compares the Indian and European Union accountability architectures across the three legitimacy dimensions, together with a stylised assessment of the United States' largely sector-specific approach.

Table 3. Comparative Accountability Architecture for AI Governance

Dimension	India	European Union	United States
Procedural legitimacy	RTI Act provides general disclosure; no AI-specific transparency mandate	Mandatory technical documentation and registration for high-risk AI (AI Act)	Sector-specific disclosure rules; no general federal mandate
Substantive legitimacy	Proportionality standard from Puttaswamy applied mainly via post hoc litigation	Ex-ante conformity assessment required for high-risk systems	Varies by sector and state; patchwork of laws
Relational legitimacy / oversight body	Data Protection Board (DPDP Act, 2023), with security and public-order exemptions	European Artificial Intelligence Board; national market-surveillance authorities	No single federal AI oversight body

Source: Author's synthesis based on European Commission (2021), Ministry of Electronics and Information Technology (2023), and PRS Legislative Research (2023).

The comparison in Table 3 suggests that India's regulatory gap is least pronounced in the service domain, where scheme guidelines and the Right to Information Act provide at least a baseline of procedural accountability, and most pronounced in the surveillance domain, where neither a risk-tiering framework comparable to the EU AI Act nor a dedicated facial recognition statute currently exists.

6.4 The accountability gap in practice

The practical consequence of this uneven regulatory landscape is that the burden of identifying and challenging AI-related harm currently falls almost entirely on civil society organisations and individual litigants rather than on any standing state institution charged with proactive oversight. Project Panoptic's tracking of facial recognition procurement, and the Software Freedom Law Centre's compilation of deployment records, function today as the closest approximation to a public registry of surveillance-oriented AI systems in India, a function that in the European Union is performed by statutory market-surveillance authorities operating under the AI Act. This reliance on non-governmental tracking is informative in itself: it indicates that transparency in Indian AI governance has so far been supplied largely through the resourcefulness of external monitors rather than through a designed feature of the accountability architecture, and it means that systems procured by less-scrutinised departments or smaller states may escape public documentation altogether.

7. Toward Reconciliation: A Proportionality-Based Framework

Reconciling the service and surveillance faces of AI-driven governance does not require India to choose between technological capability and democratic legitimacy; it requires embedding the constitutional proportionality standard articulated in Puttaswamy into the ordinary administrative lifecycle of AI systems, from procurement through deployment to review. Three design principles follow from the analysis above.

First, purpose limitation should be enforced technically as well as legally. A biometric authentication event generated for a subsidy transaction should not, by default, be queryable by a law-enforcement agency without a separate, judicially authorised process; the current pattern, in which identity infrastructure built for welfare delivery becomes a de facto resource for policing, is precisely what the necessity and proportionality standard in Puttaswamy was intended to constrain. Second, contestability should be built into every consequential automated decision, whether it denies a subsidy, flags a tax return, or matches a face against a watchlist. This means, at minimum, a published account of the categories of data and logic used, a channel for an affected individual to request human review, and a defined timeline for that review, obligations broadly consistent with the due-process safeguards Citron and Pasquale (2014) proposed for automated scoring generally. Third, independent audit capacity needs to be built into the state's own institutional architecture rather than left to litigation or civil society tracking projects. The Data Protection Board established under the Digital Personal Data Protection Act, 2023 could, if its mandate were extended and its independence strengthened, serve as the nucleus of such a capacity, but its current design is oriented primarily toward private-sector data fiduciaries rather than toward government surveillance systems, and the Act's own security and public-order exemptions currently place many of the highest-risk systems outside its practical reach.

A proportionality-based framework of this kind would not treat all AI-enabled governance uniformly. Service-oriented systems with well-defined, narrow purposes and existing grievance-redressal channels would face comparatively light-touch obligations, principally around explanation and correction rights. Surveillance-oriented systems, and particularly biometric identification systems used for law enforcement, would face the heavier obligations associated with high-risk classification: mandatory registration, independent accuracy testing disaggregated by demographic group, sunset clauses requiring periodic reauthorisation, and a statutory requirement that deployment be authorised by a specific, publicly disclosed legal instrument rather than general policing powers. This tiering mirrors, while adapting to Indian constitutional and administrative conditions, the risk-based structure of the European Union's AI Act (European Commission, 2021), without requiring wholesale adoption of a foreign regulatory model.

8. Conclusion and Policy Recommendations

The Indian state's embrace of artificial intelligence has produced genuine gains in the reach and efficiency of public service delivery, alongside a rapid and largely unregulated expansion of biometric and algorithmic surveillance capacity. Because both trajectories draw on the same underlying identity and data infrastructure, they cannot be evaluated in isolation from one another; a policy framework that strengthens service delivery while leaving surveillance unregulated will, over time, allow the surveillance function to expand through the same channels built for service, as the history of Aadhaar and its facial recognition

adjacencies illustrates. Democratic legitimacy, this paper has argued, is not a property of the technology itself but of the accountability architecture, procedural, substantive, and relational, that surrounds its use. Three recommendations follow directly from the analysis. First, Parliament should consider dedicated legislation governing biometric identification and facial recognition technology, comparable in ambition to the Digital Personal Data Protection Act but specifically addressed to law-enforcement and security use cases, including mandatory pre-deployment risk assessment and periodic independent audit. Second, the mandate and institutional independence of the Data Protection Board should be strengthened and clarified to include oversight of government surveillance systems, narrowing the current security and public-order exemptions that place many high-risk deployments outside its reach. Third, courts and administrative tribunals should be equipped, through technical capacity-building and standing expert panels, to apply the necessity and proportionality standard set out in *Puttaswamy* at the point of system design and procurement rather than only after harm has occurred and litigation has commenced. None of these steps requires India to slow its adoption of AI in governance; each is intended to ensure that the resulting state, whatever its computational capability, remains recognisably answerable to the citizens it serves and, at times, watches.

REFERENCES:

1. Citron, D. K., & Pasquale, F. A. (2014). The scored society: Due process for automated predictions. *Washington Law Review*, 89(1), 1–33.
2. Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
3. European Commission. (2021). Proposal for a regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) (COM/2021/206 final). European Union.
4. Internet Freedom Foundation. (2023). *Project Panoptic: Tracking facial recognition technology projects in India*. Internet Freedom Foundation.
5. Justice K.S. Puttaswamy (Retd.) & Anr. v. Union of India & Ors., (2017) 10 SCC 1 (Supreme Court of India).
6. Ministry of Electronics and Information Technology. (2023). *The Digital Personal Data Protection Act, 2023 (No. 22 of 2023)*. Government of India.
7. NITI Aayog. (2018). *National strategy for artificial intelligence: #AIforAll*. Government of India.
8. NITI Aayog. (2021). *Responsible AI: Approach document for India, part 1 – principles for responsible AI*. Government of India.
9. PRS Legislative Research. (2023). *The Digital Personal Data Protection Bill, 2023: Bill summary*. PRS India.
10. Singh, R. (2017). From margins to seams: Imbrication, inclusion, and torque in the Aadhaar identification project. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 4776–4824.
11. Singh, R. (2019). Give me a database and I will raise the nation-state: Insights from building Aadhaar. *South Asia: Journal of South Asian Studies*, 42(3), 501–518.
12. Software Freedom Law Centre. (2023). *Deployment of facial recognition technology for state surveillance and monitoring*. SFLC.in.



13. Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance*, 12(4), 505–523.
14. Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.