JOURNAL OF LAW, INFORMATION AND SCIENCE
ISSN 0729-1485• Vol. 4, Issue 2 (2026) • https://jlisjournal.com/

REVIEW ARTICLE Article ID: JLIS/2026/4/2/010

When the Algorithm Says No: Artificial Intelligence, Welfare Delivery, and Accountability in the Global South

Satyam Tomar [1]

[1] *Lecturer, TMCLLS, Teerthanker Mahaveer University, Moradabad, Uttar Pradesh – 244001, India.*

Received: 22 June 2026 **Revised:** 29 July 2026 **Accepted:** 25 August 2026 **Published:** 1 Oct 2026

ABSTRACT

Artificial intelligence has entered development practice across most of the Global South through one channel above all; the automation of welfare delivery. Biometric authentication, reduplication algorithms and eligibility scoring now sit on top of what is called digital public infrastructure, or DPI. This article draws on doctrinal legal analysis and published field evidence from India to examine what happens when these systems come between the poorest citizens and the entitlements owed to them. India's regulatory framework, including the "Digital Personal Data Protection Act 2023", governs data flows. Five safeguards are proposed here, each implementable without new legislation: a human fallback guarantee, exclusion audits, an explainability floor, grievance parity, and procurement transparency. As India exports its DPI model abroad, its accountability gaps travel with it. The question, in the end, is not innovation versus regulation. It is who bears the burden when an algorithm gets it wrong.

Keywords: *artificial intelligence; digital public infrastructure; social protection; algorithmic accountability; governance and public policy.*

© 2026. The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

Introduction

Much of the existing writing on artificial intelligence and development tends to present AI primarily as a promise, whether in the form of precision agriculture for smallholders, diagnostic algorithms for understaffed clinics, or chatbots that place legal information within reach of the poor, and this article does not seek to take away from that promise, which is real enough as far as it goes. For hundreds of millions of people across the Global South, however, the first sustained encounter with algorithmic decision-making tends to look rather different, and considerably more consequential: a fingerprint that fails to authenticate at the ration shop, a name that is quietly struck off a beneficiary list because a reduplication routine has flagged it, or a pension that stops arriving because a bank-seeding algorithm has linked the wrong account. Before it is anything else, then, artificial intelligence in development practice is, in the main, the automation of the welfare interface.

India remains the case most commonly cited in this connection. Over roughly the past decade and a half, the Indian state has assembled what is now marketed internationally as digital public infrastructure, or DPI, comprising the Aadhaar biometric identity system, which by now covers more

than 1.3 billion residents, the Unified Payments Interface, which processes billions of transactions each month, and a wider array of platforms serving direct benefit transfer, health records and credit. During its G20 presidency in 2023, India actively promoted this 'India Stack' as a model capable of being replicated elsewhere, and a number of countries, ranging from Ethiopia to the Philippines, have since taken up that offer **(Government of India 2023)**. Increasingly layered onto this infrastructure are machine-learning systems of various kinds, including fraud-detection models that flag beneficiaries as 'suspicious', facial recognition deployed alongside conventional fingerprint checks, predictive analytics used to determine subsidy eligibility, and algorithms that cross-match databases at a scale that would be impossible for any human clerk to replicate.

This article is concerned with a question that lies at the heart of this special issue, namely who is to be held accountable when such systems fail, and what development practice might realistically do about this in the near term rather than at some indeterminate point in the future. The question is not merely academic, since failure in this domain is by now well documented: a substantial body of empirical work records authentication failures and instances of quantity fraud in Aadhaar-linked ration delivery **(Drèze et al. 2017; Muralidharan, Niehaus, and Sukhtankar 2020)**, while journalists and researchers alike have traced a number of starvation deaths to ration denials that followed Aadhaar-related exclusion **(Khera 2019)**. These outcomes are not properly described as glitches capable of being resolved through a future software update; they are, rather, the predictable consequence of building systems on the underlying assumption that the algorithm is correct and that the burden lies on the claimant to prove otherwise.

The argument is anchored in a long-standing insight of the information-and-communication-technologies-for-development literature: technology amplifies existing institutional capacity and intent rather than substituting for it **(Toyama 2015; Heeks 2018)**. Where administrative systems are designed around distrust of claimants, algorithmic tools amplify that distrust with unprecedented efficiency; where grievance and audit institutions are weak, digitization renders their weakness invisible behind dashboards of transactional success. The question for development practice is therefore not whether AI 'works', but which institutional dispositions it will scale.

The article's contribution is threefold. First, it reframes the AI-and-development debate around welfare automation rather than frontier AI, which is where algorithmic power actually touches the poor. Second, it offers a legal-institutional analysis of India's regulatory response in particular the “Digital Personal Data Protection Act 2023” and shows why a data-protection statute, however necessary, is structurally incapable of governing automated decision-making in welfare delivery. Third, it turns this analysis into a five-part safeguard framework. Development practitioners, implementing agencies and governments can adopt each safeguard under the law as it already stands, rather than wait for a comprehensive AI Act that may be years away. The argument proceeds from the conviction that in the Global South the governance question is not whether to use AI in development, but on whose terms and with what remedies.

Literature Review: AI, Digital Public Infrastructure, and Development Practice

Discussions of AI governance in the Global South often import the framing of the transatlantic debate; foundation models, existential risk, generative AI and then asks how poorer countries should respond. This framing obscures the actually existing deployment landscape. In Indian development administration, the systems that matter are mostly unglamorous forms of automation and machine learning embedded in programme delivery.

A definitional clarification is necessary at the outset; because a predictable objection to this framing is that much of what follows concerns 'mere' automation rather than artificial intelligence proper. The

objection misunderstands both the technology and the stakes. The systems described here sit on a continuum: deterministic rule-based cross-matching at one end, probabilistic biometric matching and statistical fraud-scoring in the middle, and machine-learning models for targeting and language processing at the other. Regulatory instruments increasingly treat this continuum as a single object; the EU AI Act's definition of an 'AI system' turns on inference and autonomy of output, not on any particular technique, and the OECD definition it adapts is similarly capacious. More importantly, from the standpoint of the affected claimant the distinction is irrelevant: the phenomenology of denial an unexplained, unappealable 'no' from a machine; is identical whether the engine is a threshold rule or a neural network. This article therefore uses 'algorithmic systems' and 'AI' to cover the continuum, while noting where a specific argument depends on machine learning properly so called.

Four categories dominate. The first is biometric authentication and its statistical machinery: fingerprint and iris matching, and increasingly face authentication, used to verify identity at the point of service in the “Public Distribution System”, the “Mahatma Gandhi National Rural Employment Guarantee Scheme”, and pension disbursement. Such matching is inherently probabilistic in nature, and the error rate the system is prepared to tolerate, particularly with respect to false rejections, is in truth a policy choice, one that is encoded in a threshold setting that beneficiaries themselves are never permitted to see. The second category comprises reduplication and fraud analytics, that is to say algorithms that cross-match beneficiary databases in order to detect what are termed 'ghost' or 'duplicate' entries; governments routinely cite the resulting large-scale deletion of ration cards as evidence of savings, although researchers have contested the extent to which these deletions in fact removed genuine households rather than eligible ones **(Drèze and Khera 2017)**. The third category is eligibility scoring and targeting, involving the use of proxy-means testing, satellite data, and predictive models to determine who is sufficiently poor to qualify for support, a practice with a fairly long pedigree in development economics that machine learning has, in recent years, considerably industrialized. The fourth, growing fastest, is generative and conversational AI at the citizen interface: multilingual chatbots for scheme information, AI-assisted grievance triage, and document processing in local languages, supported by public language-technology missions such as Bhashini **(Ministry of Electronics and Information Technology 2023)**.

Nor is this landscape uniquely Indian, which is precisely why the Indian case merits close attention. Kenya's “Huduma Namba digital identity programme” was halted by the High Court in 2021 for proceeding without an adequate data protection impact assessment; South Africa's social grant disbursement has repeatedly raised concerns about private-vendor control over beneficiary data; and Togo's Novissi programme during COVID-19 used machine learning on satellite imagery and phone metadata to target emergency cash transfers; an innovation whose accuracy gains its own architects paired with warnings about exclusion of the unconnected poor **(Aiken et al. 2022)**. Digital identification systems now underpin welfare delivery across dozens of developing countries **(Gelb and Metz 2018)**. What distinguishes India is scale, the deliberate export of its stack, and the density of empirical research documenting outcomes – making it the best available window into dynamics that other DPI-adopting states will encounter with fewer resources to detect them.

Two features of this landscape deserve emphasis for development practice. First, most of these systems are procured or built by the state but designed and operated by private vendors and quasi-governmental technology bodies, blurring accountability from the outset. Second, these systems operate on populations with low digital literacy and unreliable connectivity, and, more importantly, on bodies that biometric machines often read badly. Manual laborers have fingerprints worn smooth by years of physical work. Elderly people have iris patterns that fade with age. Persons with disabilities are excluded by the design of the system itself **(Rao and Nair 2019)**. This error does not fall evenly across the population. It falls hardest on those least able to bear it. Alston (2019), writing his report to the UN General Assembly on digital welfare states, put the danger plainly: the risk of what he called

'stumbling zombie-like into a digital welfare dystopia', where technological sophistication and the steady erosion of claimants' rights exist side by side.

Methodology

This article rests on doctrinal legal analysis of Indian statutes, subordinate legislation and constitutional case law, read alongside a synthesis of published empirical studies, official evaluations and civil-society evidence on digital welfare delivery. The statutes chosen bear directly on algorithmic gatekeeping of entitlements: the Aadhaar Act 2016, the DPDP Act 2023, the National Food Security Act 2013, MGNREGA 2005, and the Puttaswamy and PUCL line of Supreme Court judgments. Among empirical studies, preference went to those built on independent household surveys rather than administrative data alone, because the two are known to diverge, sometimes sharply **(Muralidharan, Niehaus, and Sukhtankar 2020)**. The approach is kept deliberately non-technical, so that readers of this journal, who come from many disciplines, do not need a computer science background to follow the argument. No fieldwork was undertaken, no human participants were involved, and no personal data was collected; informed consent procedures therefore did not arise. The discussion that follows is organised around three plain questions: what does the evidence actually show, why does existing regulation fail to answer it, and what can practice do about it right now.

Results and Discussion

The Indian experience: exclusion by design, exclusion by default

(a) The promise and the evidence

The official case for algorithmic welfare rests on leakage reduction and convenience: biometric authentication ensures that grain reaches the person named on the card, and direct benefit transfer eliminates intermediaries. There is genuine evidence of gains in some programmes and states, particularly where reforms were accompanied by choice and robust exception handling **(Muralidharan, Niehaus, and Sukhtankar 2016)**. The point of this article is not that digitization has failed; it is that its costs and benefits are distributed by design choices that currently escape legal accountability.

The most careful empirical work illustrates the distribution problem. Studying the imposition of Aadhaar-based biometric authentication in Jharkhand's PDS, Drèze et al. (2017) documented significant transaction failures and the exclusion of vulnerable households, with little evidence of corresponding leakage gains in that setting. Muralidharan, Niehaus, and Sukhtankar (2020) found in a large-scale experimental evaluation that reforms reduced some forms of leakage while also generating non-trivial exclusion of genuine beneficiaries and, tellingly, that the state's administrative data systematically understated exclusion relative to independent household surveys. In parallel, civil society tracking associated a series of hunger deaths with ration denial following failures of Aadhaar seeding or authentication **(Khera 2019)**. Whatever one's overall assessment of the reforms, one conclusion is difficult to resist: the systems were engineered to make inclusion errors visible (as 'fraud detected') and exclusion errors invisible (as silence in the administrative data).

Exclusion is also structured by gender and social position, a dimension the efficiency literature largely ignores. Women in South Asia are considerably less likely than men to own a mobile phone or use mobile internet **(GSMA 2023)**. Yet OTP-based authentication, app-based grievance portals and phone-seeded bank transfers all assume connectivity as the starting point. Consider three examples. A married woman may still hold documents that carry her natal-village address. A widow's entitlement may have been recorded under her husband's name rather than her own. A woman worker's fingerprints may have worn thin after years of domestic or farm labour. Each faces a higher risk of authentication failure

than the average claimant. Where welfare systems are digitized without regard to these gradients, the resulting algorithm does not merely mirror pre-existing social exclusion; it recasts that exclusion in the more neutral-sounding language of technical error.

(b)The constitutional frame

Indian constitutional law has responded to this shift, but only in part. The doctrinal foundation predates the technology. In the continuing mandamus proceedings of "People's Union for Civil Liberties v. Union of India" (2001), the Supreme Court turned food schemes into enforceable legal entitlements under the right to life in Article 21. "The National Food Security Act 2013" later codified this position. Entitlements, on this view, are rights, not charity from the state. A system that makes a right depend on biometric success has changed the content of that right, and no legislature ever debated the change. The statutory hinge is Section 7 of the "Aadhaar (Targeted Delivery of Financial and Other Subsidies, Benefits and Services) Act 2016", which lets governments demand Aadhaar authentication as a condition for subsidy delivery. This is the clause that turns an identity platform into a gate.

In "Justice K.S. Puttaswamy v. Union of India" (2017), a nine-judge bench of the Supreme Court unanimously recognized privacy as a fundamental right under Article 21 and set out a proportionality standard for state intrusions. The Aadhaar judgment followed a year later, "Puttaswamy (Aadhaar) v. Union of India" (2018). The majority upheld Aadhaar for welfare delivery but struck down its extension to private services. Importantly, the Court held that no one could be denied benefits simply because authentication failed; alternative proof of identity had to be accepted. Justice Chandrachud dissented, and went further still, treating the entire architecture as constitutionally unsound.

The gap between this doctrine and what actually happens at the counter is the real lesson for development practice. The Court's direction that authentication failure must not mean exclusion is, on paper, exactly the human fallback that scholars of algorithmic accountability recommend. Field studies after 2018 kept finding denial at the point of service all the same. The direction was never turned into an enforceable operational standard. No fallback protocol was mandated with monitored compliance. No penalty attached to wrongful denial. No independent body audited exclusion rates **(Khera 2019)**. A constitutional right unaccompanied by administrative machinery is, in practical terms, little more than a press release, and this pattern recurs across much of the Global South, wherever a strong apex court sits above a comparatively weak street-level bureaucracy, suggesting that AI governance strategies premised solely on rights litigation are likely, over time, to under-deliver.

(c)Algorithmic opacity meets administrative law

Indian administrative law requires, broadly, that state action affecting a person's rights be reasoned, non-arbitrary, and consistent with the principles of natural justice, in particular audi alteram partem, or the right to be heard before an adverse decision is taken. The automatic deletion of a ration card sits uneasily with each of these requirements, given that the household affected typically receives no notice, no stated reasons, and no opportunity to be heard before the deletion takes effect, and that the actual 'decision-maker' in such cases is, in truth, a matching threshold embedded within a vendor's software. Challenges to this are rare all the same, not because the law offers no remedy, but because the people affected rarely even know that an algorithm made the call, let alone have the money to litigate or the paperwork to prove, to the system's satisfaction, that they exist at all. The doctrinal tools exist; the practical capacity to invoke them does not. Development practitioners are often the only actors positioned to bridge that gap, which is why the safeguard framework proposed below assigns them a central role.

India's regulatory architecture: governing data, not decisions

(a)The DPDP Act 2023 and its silences

India's flagship response to the digital era is the "Digital Personal Data Protection Act 2023", which received presidential assent in August 2023 after a six-year legislative journey that began with the

Justice Srikrishna Committee. The Act establishes consent-based processing, purpose limitation, obligations on 'data fiduciaries', and a Data Protection Board to adjudicate breaches. As a data-protection instrument it is a milestone; as an instrument of algorithmic accountability in welfare delivery it is close to inert, for three structural reasons.

First, the Act contains sweeping exemptions for the state. Section 7(b) treats the provision of subsidies, benefits, services, certificates, licences, and permits by the state as a 'legitimate use' for which personal data may be processed without consent – precisely the welfare context this article concerns – and Section 17 empowers the central government to exempt state instrumentalities from most of the Act's provisions. The population most exposed to algorithmic decision-making is thus the population with the weakest statutory protection. Second, unlike the EU General Data Protection Regulation, whose Article 22 confers a qualified right not to be subject to solely automated decisions with legal or similarly significant effects, the DPDP Act contains no provision on automated decision-making at all: no right to human review, no right to explanation, no algorithmic impact assessment. Third, the enforcement machinery is centralized and appointed by government. The Act provides no real compensation to the person actually harmed. Remedies take the form of penalties paid to the state, not to the beneficiary. A person wrongly struck off a scheme is left, under this Act, with almost nothing by way of practical recourse.

(b)Soft law, strategy documents, and the advisory turn

Outside the DPDP Act, Indian AI governance is mostly a collection of strategy papers and advisories, not binding law. This includes NITI Aayog's National Strategy for Artificial Intelligence of 2018 and its subsequent 'Responsible AI' approach papers of 2021, both of which articulate principles of safety, inclusivity and accountability without imposing any binding obligation; the India AI Mission of 2024, which is, in substance, primarily a capacity-building and compute programme; and a series of advisories issued by the Ministry of Electronics and Information Technology under the Information Technology Act 2000 and the Intermediary Guidelines 2021, including the contested advisory of March 2024 on generative AI models, which was subsequently revised following pushback from industry. This is a familiar trajectory for observers of Indian regulation. Principles come first. Infrastructure investment follows. Binding obligations get deferred indefinitely, always in the name of innovation. The contrast with the European Union's AI Act 2024; which classifies AI used to determine access to essential public services and benefits as high-risk, triggering mandatory risk management, human oversight, and transparency duties; could not be sharper. Whatever the EU Act's faults as a template for the Global South, it at least recognises welfare algorithms as a regulatory object. Indian law, as yet, does not.

(c)Why 'wait for the AI Act' is not a strategy

It is tempting to say the answer is a comprehensive AI law, and India may well pass one eventually; the Report of the Subcommittee on AI Governance under the IndiaAI Mission (2025) already points that way. But development practice cannot simply wait for Parliament's timetable. Consider the Dutch childcare benefits scandal. An algorithmic fraud-scoring system there wrongly branded tens of thousands of families as fraudsters; families with dual nationality bore the brunt of it, and the scandal eventually brought down a government in 2021. All of this happened in a country with the GDPR in full force **(Amnesty International 2021)**. The statute on the books did not stop it. What eventually did were institutional practice, audit, and the plain political embarrassment of the harm becoming public. The lesson for the Global South is not comforting: legal frameworks are necessary, but nowhere near sufficient. The real work of accountability happens at the operational layer, in procurement terms, programme design, grievance systems and monitoring. That is where this article turns next.

(d)A practice-oriented framework: five safeguards that do not require new legislation

The following framework distils the Indian experience into five safeguards addressed to the actors who actually build and run algorithmic welfare systems: government programme managers, implementing

agencies, multilateral funders, technology vendors, and the NGOs that monitor delivery. Each safeguard is implementable under existing Indian law indeed, several merely operationalise what Puttaswamy (Aadhaar) already requires; and each is transferable, with adaptation, to other Global South contexts adopting DPI-style architectures.

The framework does not write on a blank slate. The World Bank's Principles on Identification for Sustainable Development already commit endorsing institutions to inclusion, non-discrimination, and user redress **(World Bank 2021)**, and 'human oversight' appears in virtually every AI-ethics charter. The contribution here is specificity and enforceability: converting principle into design requirement, procurement clause, and monitored metric; the register at which welfare delivery actually operates, and at which most existing frameworks fall silent.

1. The human fallback guarantee

No essential entitlement should ever be capable of being denied by a machine acting alone. Every algorithmic checkpoint, whether it is authentication, reduplication or eligibility scoring, needs a manual alternative that is properly defined, staffed and monitored, and, importantly, the failure of the machine ought as a matter of design to default to provisional inclusion rather than exclusion, which would invert the present logic under which the claimant bears the risk of every system error. In operational terms, this requires an exception register maintained at every point of service, a mandated turnaround time for manual verification, and an explicit instruction that grain, wages or pensions should continue to flow while such verification is under way. Funders and practitioners would do well to treat the absence of a coasted fallback protocol in any AI-enabled delivery design as a warning sign comparable to the absence of a safeguarding policy, rather than as a detail to be addressed at some later stage.

The usual objection is that a manual fallback reopens the leakage the technology was meant to close, and that staffing exception-handling is too costly at India's scale. This objection deserves an answer, not a dismissal. First, the trade-off is usually overstated: independent evaluations found that the exclusion caused by authentication mandates was substantial, while leakage gains in the same settings were modest, sometimes absent altogether **(Drèze et al. 2017)**. The 'cost' of a fallback, in other words, is mostly the cost of not wrongly denying people what they are owed, and that is not a cost the state has any right to save. Second, a fallback does not require unaudited discretion. Exception transactions can be logged, capped per outlet, flagged for later verification, and checked for collusion, which turns the fallback channel into a source of monitored data rather than into a loophole. Third, the comparison is asymmetric in kind: leakage is a fiscal loss recoverable through audit and prosecution; the denial of subsistence to an entitled household is, at the margin documented in India, irreversible. A legal system that treats the second as an acceptable price of preventing the first has inverted the constitutional ordering of harms.

2. Exclusion audits with independent publication

What is measured is what is managed, and welfare algorithms are currently measured almost exclusively on inclusion errors (fraud) rather than exclusion errors (wrongful denial). Programmes deploying automated verification should be required by procurement condition or funder covenant where statute is silent to publish disaggregated exclusion metrics: authentication failure rates by age, gender, disability, occupation, and region; deletion volumes with reinstatement rates; and the results of periodic independent field audits that survey deleted or denied persons rather than relying on administrative data, which the evidence shows understates exclusion **(Muralidharan, Niehaus, and Sukhtankar 2020)**. India already has a homegrown template for this in social audit institutions. Section 17 of the Mahatma Gandhi National Rural Employment Guarantee Act 2005 already requires social audits of how the programme is implemented. That same architecture can be stretched, by rule or scheme guideline, to cover the algorithmic layer too, and no new legislation is needed to do it.

3. An explainability floor

Every adverse decision reached by an algorithm ought to be accompanied by a reason on which the claimant concerned can meaningfully act, and this need not extend to the internal workings of the underlying model, which would in any event be of little practical use to the claimant. What matters, rather, is the specific ground of denial, namely which database failed to match, which threshold was crossed, or which document was found wanting, communicated to the claimant in their own language, in writing, and at the point at which the denial is made. This is properly understood not as a technical demand but as an administrative-law demand, being simply an extension of the existing requirement that reasons be given for adverse orders to the algorithmic context, and it remains feasible even where the underlying model is otherwise opaque, since what the claimant requires is an account of the decision pipeline, that is, which data source produced the mismatch, which rule or threshold was triggered, and which document would cure the defect, rather than a technical account of the model's internal weights. If a model genuinely cannot support even this much explanation for a decision touching an essential entitlement, the conclusion should be that the model is unfit for that use. The duty to explain is not the thing that ought to be relaxed. In the absence of such an explanation, the right of appeal, whatever its formal status, remains largely decorative.

4. Grievance parity

The grievance system must, at a minimum, be as accessible as the system that gives rise to the grievance in the first place, so that where enrolment takes place at the doorstep through a digital device, the corresponding appeal cannot reasonably be expected to take the form of a handwritten petition submitted to a district office. In concrete terms, this requires grievance channels available in local languages across assisted physical, telephonic and digital modes, fixed deadlines for acknowledgement and resolution with a clear mechanism for escalation, compensation in cases where subsistence entitlements have been wrongly denied, drawing on the precedent already established by service-guarantee legislation in several Indian states, and public reporting of how many grievances came in and how they were resolved. Development organisations can add value immediately by funding and staffing 'algorithmic help desks' – paralegal-style intermediaries who translate between claimants and systems for which early Indian civil society experiments around Aadhaar exclusion offer working models.

5. Procurement transparency and vendor accountability

Most welfare algorithms enter the state through procurement, which makes contract the most underused regulatory instrument in the Global South. Contracts for AI-enabled delivery systems should mandatorily include: disclosure of the system's function, training data provenance, and known error rates; pre-deployment and periodic bias testing on the actual beneficiary population; audit rights for the state and designated independent bodies; liability and indemnity provisions that prevent vendors from contracting out of harm; and escrow of documentation sufficient for a successor or auditor to understand the system. None of this requires legislation; it requires model contract clauses, which multilateral institutions and national procurement bodies could standardize tomorrow. For countries importing DPI stacks – including those importing India's insisting on these clauses at adoption is far easier than retrofitting them later.

Implications for the Global South: exporting infrastructure, exporting gaps

India's DPI diplomacy gives this analysis significance beyond one country. When identity, payment, and benefit-delivery architectures travel to countries with weaker courts, thinner civil society, and no equivalent of India's research ecosystem documenting exclusion the accountability gaps travel with them, and the compensating pressures may not. The G20 consensus on DPI (2023) celebrated interoperability and scale; it said almost nothing about fallback, audit, or remedy. Development practice, and the multilateral institutions financing DPI adoption, therefore occupies a position of unusual leverage: they can make the five safeguards above a condition of engagement, effectively

building a floor of algorithmic accountability into the transfer of infrastructure itself. This would be a distinctively Global South contribution to AI governance grounded not in the risk taxonomies of Brussels but in the delivery realities of Jharkhand and exportable on the same channels as the code.

There is also a research agenda here to which this special issue can contribute. We know far too little, comparatively, about exclusion dynamics in DPI adopting countries; about how claimants actually experience and contest algorithmic denial; about the political economy of vendor–state relationships in welfare automation; and about which grievance designs actually restore entitlements at scale. Legal scholarship, including this author's, must in turn move beyond doctrinal critique toward empirical engagement with administrative practice – the orders, audits, and help-desk logs where algorithmic governance actually happens.

Conclusion

Three limitations need to be said plainly, not buried in a footnote. First, this analysis rests on published evidence and legal material rather than new fieldwork of its own, so it inherits the geographic bias of the underlying studies, which lean heavily on states like Jharkhand and Andhra Pradesh; exclusion elsewhere may look quite different. Second, the doctrinal method employed here is better suited to capturing what the law says, and how courts have interpreted it, than to capturing how that law actually operates at the street level, a task that only sustained empirical work is capable of undertaking. Third, the five safeguards proposed in this article have been arrived at analytically rather than through piloting, so that their real costs and likely failure modes remain, for the present, open questions. Each of these limitations points toward a corresponding direction for future research, whether in the form of comparative studies of exclusion across DPI-adopting countries, ethnographies of how claimants themselves experience algorithmic grievance, or design experiments testing audited fallback channels in practice.

None of this, however, is to suggest that the practical implications of the analysis must await the completion of that research agenda. For programme managers, the immediate requirement is a coasted fallback protocol together with an exception register at every algorithmic checkpoint. For funders and multilateral institutions, it is a matter of making exclusion audits and the model procurement clauses set out above a condition of DPI financing. For development organisations, it is a matter of staffing the intermediary layer, in the form of algorithmic help desks, capable of converting paper rights into entitlements actually restored. And for governments engaged in adopting or exporting DPI, it is, ultimately, a matter of recognizing that inclusion is not something to be read off a dashboard but a constraint that must be designed for from the outset.

Artificial intelligence in Global South development practice is not, at bottom, primarily a story about frontier models; it is, rather, a story about automated gatekeeping standing between poor citizens and the entitlements that sustain them. India's experience brings together, in a single setting, world-leading infrastructure, genuine efficiency gains, documented exclusion, a data-protection statute that is silent on automated decisions, and constitutional directions that remain unenforced at the level of the ration shop. Much of the developing world will recognize this picture. The argument of this article is that the right response is neither blind faith in technology nor a blanket refusal to use it, but the patient, disciplined building of five safeguards, human fallback, exclusion audits, explainability, grievance parity and procurement transparency, all of which can be put in place now, under the law as it stands, by the governments, funders and practitioners who actually build these systems. The measure of AI in development practice should never be the leakage it prevents. It should be the last claimant it wrongly turns away. Until our systems are designed, audited and put right around that one person, the algorithm's 'no' will remain the Global South's most consequential AI output.

References

- Aiken, E., S. Bellue, D. Karlan, C. Udry, and J. E. Blumenstock. 2022. "Machine Learning and Phone Data Can Improve Targeting of Humanitarian Aid." Nature 603: 864–870. https://doi.org/10.1038/s41586-022-04484-9
- Alston, P. 2019. Digital Welfare States and Human Rights: Report of the Special Rapporteur on Extreme Poverty and Human Rights. UN Doc A/74/493. New York: United Nations. https://undocs.org/A/74/493
- Amnesty International. 2021. Xenophobic Machines: Discrimination through Unregulated Use of Algorithms in the Dutch Childcare Benefits Scandal. London: Amnesty International. https://www.amnesty.org/en/documents/eur35/4686/2021/en/
- Burman, A. 2023. "Understanding India's New Data Protection Law." Washington, DC: Carnegie Endowment for International Peace. https://carnegieendowment.org/research/2023/10/understanding-indias-new-data-protection-law
- Drèze, J., N. Khalid, R. Khera, and A. Somanchi. 2017. "Aadhaar and Food Security in Jharkhand: Pain without Gain?" Economic and Political Weekly 52 (50): 50–59. https://www.jstor.org/stable/45132599
- Drèze, Jean, and Reetika Khera. 2017. "Recent Social Security Initiatives in India." *World Development* 98: 555–72. https://doi.org/10.1016/j.worlddev.2017.05.035.
- European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2016/679/oj
- Irene, M., Pascal, P. D., & Mishra, D. K. (2026). Invisible Advertising to Children in Gaming Ecosystems: Legal Regulation of Embedded AI Marketing. *Canadian Journal of Marketing Research*, *16*(2), 486-493.
- Dubey, B., & Das, L. (2025). Legal Framework for Mediation in Family Disputes in India: Analysis of Statutes, Case Laws, and Mediation Mechanisms. *Issue 3 Int'l JL Mgmt. & Human.*, *8*, 4476.
- Batar, S. A HISTORICAL OVERVIEW OF THE INDIAN JUDICIAL SYSTEM. *JUDICIAL INDEPENDENCE and ACCOUNTABILITY*, *45*.
- Raj, M., Khan, H., Kathuria, S., Chanti, Y., & Sahu, M. (2024, March). The use of artificial intelligence in anti-money laundering (AML). In *2024 3rd International conference on sentiment analysis and deep learning (ICSADL)* (pp. 272-277). IEEE.
- Solanki, S. S., Gupta, Y., Chaudhary, V., & Agarwal, A. (2016). The indirect bonding: a literature review. *Int J Med Dent Homeopat Nurs*, *1*, 1-8.
- Batar, S. LACK OF JUDICIAL ACCOUNTABILITY IN INDIA. *JUDICIAL INDEPENDENCE & ACCOUNTABILITY*, *6*.
- Raman, A., Khan, H., Pandey, S., Lande, J., Patet, N., & Sahu, M. (2023, December). Imperative Role of AI in Cyber Fraud Detection. In *2023 International Conference on Power Energy, Environment & Intelligent Control (PEEIC)* (pp. 203-207). IEEE.
- Bibhu, N. K., & Dubey, B. IMPLEMENTATION OF CONSUMER PROTECTION ACT 2019: A STUDY OF THE ROLE OF CONSUMERS. *BIHAR JOURNAL OF*, 439.
- Singh, A. K., Irene, M. M., Jha, P., & Dubey, M. B. (2025). Justice for the Voiceless: Evaluating the Efficacy of Criminal Justice Mechanisms in Combating Child Sexual Abuse.
- Batar, S. ISSUES AND CHALLENGES OF JUDICIAL ACCOUNTABILITY IN INDIA. *JUDICIAL INDEPENDENCE and ACCOUNTABILITY*, *7*.

- Gupta, Y. C. (2022). Impact of Technology on Modern Policing. *Law & Order Administration*, *77*, 422-427.
- Bhattacharya, S., & Bibhu, N. K. (2026). THE ARBITRATION ODYSSEY: NAVIGATING GLOBALIZATION'S INFLUENCE ON CROSS-BORDER DISPUTE RESOLUTION. *Evaluating the Efficacy of Alternative Dispute Resolution: A Modern Legal Perspective*, 36.
- Batar, S. CURRENT MECHANISM OF JUDICIAL ACCOUNTABILITY IN INDIA. *JUDICIAL INDEPENDENCE & ACCOUNTABILITY*, *39*.
- Moparty, H., & Irene, M. (2020). LIFE BETWIXT ON-LINE AND OFF-LINE–A PHARMACEUTICAL PARANOIA. *World Journal of Pharmaceutical Sciences*, *9*(10.20959).
- Gupta, N., Adawadkar, M., Asim, I., Verma, S., & Dubey, B. (2025). Predictive insights: Leveraging artificial intelligence for strategic business decision-making. *Advances in Consumer Research*, *2*(5), 98.
- Bibhu, N. K., & Gaur, S. (Eds.). (2026). *Evaluating the Efficacy of Alternative Dispute Resolution: A Modern Legal Perspective*. SB Knowledge Press.
- Wright, D. S. (1998). Federalism, intergovernmental relations, and intergovernmental management. *Handbook of public administration*, 381-447.
- Batar, S. (2021). Gender inequality in india: an overview. *ACADEMICIA: An International Multidisciplinary Research Journal*, *11*(12), 458-464.
- European Union. 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. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
- Gelb, A., and A. Diofasi Metz. 2018. Identification Revolution: Can Digital ID Be Harnessed for Development? Washington, DC: Center for Global Development. https://www.jstor.org/stable/10.7864/j.ctt21c4t40
- Government of India. 2023. G20 New Delhi Leaders' Declaration. New Delhi: Ministry of External Affairs. https://www.mea.gov.in/images/CPV/G20-New-Delhi-Leaders-Declaration.pdf
- GSMA. 2023. The Mobile Gender Gap Report 2023. London: GSM Association. https://www.gsma.com/gender-gap-2023/
- Heeks, R. 2018. Information and Communication Technology for Development (ICT4D). Abingdon: Routledge. https://www.routledge.com/Information-and-Communication-Technology-for-Development-ICT4D/Heeks/p/book/9781138101814
- Justice K.S. Puttaswamy (Retd.) v. Union of India. 2017. (2017) 10 SCC 1 (Supreme Court of India). https://indiankanoon.org/doc/91938676/
- Justice K.S. Puttaswamy (Retd.) v. Union of India (Aadhaar). 2018. (2019) 1 SCC 1 (Supreme Court of India). https://indiankanoon.org/doc/127517806/
- Khera, R. 2019. "Aadhaar Failures: A Tragedy of Errors." Economic and Political Weekly 54 (14) (Engage). https://www.epw.in/engage/article/aadhaar-failures-food-services-welfare
- Ministry of Electronics and Information Technology. 2023. Bhashini: National Language Translation Mission – Overview. New Delhi: Government of India. https://bhashini.gov.in/
- Muralidharan, K., P. Niehaus, and S. Sukhtankar. 2016. "Building State Capacity: Evidence from Biometric Smartcards in India." American Economic Review 106 (10): 2895–2929. https://doi.org/10.1257/aer.20141346
- Muralidharan, K., P. Niehaus, and S. Sukhtankar. 2020. "Identity Verification Standards in Welfare Programs: Experimental Evidence from India." NBER Working Paper 26744. Cambridge, MA: National Bureau of Economic Research. https://doi.org/10.3386/w26744
- NITI Aayog. 2018. National Strategy for Artificial Intelligence #AIForAll. New Delhi: Government of India. https://www.niti.gov.in/sites/default/files/2021-02/Responsible-AI-22022021.pdf

- NITI Aayog. 2021. Responsible AI: Approach Document for India, Part 1 – Principles for Responsible AI. New Delhi: Government of India. https://www.niti.gov.in/sites/default/files/2021-02/Responsible-AI-22022021.pdf
- People's Union for Civil Liberties v. Union of India. 2001. Writ Petition (Civil) No. 196 of 2001 (Supreme Court of India, interim orders). https://indiankanoon.org/doc/411836/
- Rao, U., and V. Nair. 2019. "Aadhaar: Governing with Biometrics." South Asia: Journal of South Asian Studies 42 (3): 469–481. https://doi.org/10.1080/00856401.2019.1595343
- The Aadhaar (Targeted Delivery of Financial and Other Subsidies, Benefits and Services) Act, 2016 (No. 18 of 2016). New Delhi: Government of India. https://www.indiacode.nic.in/handle/123456789/2160/
- The Digital Personal Data Protection Act, 2023 (No. 22 of 2023). New Delhi: Government of India. https://www.indiacode.nic.in/handle/123456789/22037?view_type=browse
- The Information Technology Act, 2000 (No. 21 of 2000). New Delhi: Government of India. https://www.indiacode.nic.in/bitstream/123456789/13116/1/it_act_2000_updated.pdf
- The Mahatma Gandhi National Rural Employment Guarantee Act, 2005 (No. 42 of 2005). New Delhi: Government of India. https://www.indiacode.nic.in/bitstream/123456789/6930/1/the_mahatma_gandhi_national_rural_employment_guarantee_act%2C_2005.pdf
- The National Food Security Act, 2013 (No. 20 of 2013). New Delhi: Government of India. https://nfsa.gov.in/portal/nfsa-act
- Toyama, K. 2015. Geek Heresy: Rescuing Social Change from the Cult of Technology. New York: PublicAffairs.
- World Bank. 2021. Principles on Identification for Sustainable Development: Toward the Digital Age. Washington, DC: World Bank. https://documents.worldbank.org/en/publication/documents-reports/documentdetail/213581486378184357

Declarations

Conflict of interest: The authors declare that they have no conflict of interest.

Funding: This review received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Ethical approval: Not applicable. This article is a review of published literature and did not involve human participants or animal subjects.

Data availability: No new data were generated or analyzed in support of this review. All sources relied upon are cited in the References.

Use of AI tools: No generative artificial intelligence tool was used in the conception, drafting or analysis of this article. The authors take full responsibility for the content.

Author contributions: Author contributed to the conception of the review, the survey of the literature, and the drafting and revision of the manuscript. Author read and approved the final version.