

Fairness on Autopilot? Regulating Algorithmic Decisions in Indian Governance

*¹ Shabnam Akbar Pathan, ²Dr. Rajeshri N Varhadi and ³Dr. Sanjay Jadhav

*¹ Assistant Professor & Research Scholar, Department of Law, University of Mumbai, Maharashtra, India.

²Professor & Head Department of Law, University of Mumbai, Maharashtra, India.

³Associate Professor, Department of Law, University of Mumbai, Maharashtra, India.

Article Info.

E-ISSN: 2583-6528

Impact Factor (SJIF): 6.876

Peer Reviewed Journal

Available online:

www.alladvancejournal.com

Received: 08/ May/2025

Accepted: 03/June/2025

Abstract

As India fast-tracks its digital transformation, artificial intelligence (AI) is quietly stepping into spaces once governed by human judgment—from welfare distribution and facial recognition to predictive policing and public service delivery. While the tech promises efficiency and scalability, it also raises pressing questions: What happens when an algorithm makes a mistake? Who's accountable? And most importantly—is the process fair? This paper explores the growing use of AI-based decision-making systems in Indian governance and interrogates the legal and ethical frameworks (or the lack thereof) guiding them. Unlike traditional legal mechanisms that prioritize human dignity and due process, AI operates in a black-box model, often with little transparency or room for contestation. These systems largely imported or developed without local socio-legal context, risk reinforcing existing inequalities, especially when deployed in areas affecting marginalized communities. The research draws upon constitutional principles of fairness, non-discrimination, and procedural justice to examine whether India's current legal landscape is equipped to regulate algorithmic decision-making. It critically evaluates existing policies such as the Personal Data Protection Bill and National AI Strategy, revealing significant regulatory blind spots. By comparing global models of AI governance and accountability (such as the EU's AI Act), the study highlights the urgent need for an Indian framework that balances innovation with justice. Ultimately, this paper argues that without human-centric laws and ethical oversight, the promise of AI can quietly turn into digital disenfranchisement. We must build systems where the code serves the people—not the other way around.

*Corresponding Author

Shabnam Akbar Pathan

Assistant Professor & Research Scholar,
Department of Law, University of
Mumbai, Maharashtra, India.

Keywords: AI governance, algorithmic justice, digital rights, Indian constitution, public policy, automated decision-making, data ethics, legal regulation, accountability, socio-legal impact.

Introduction

Imagine this: you apply for a government welfare scheme, and instead of a person reviewing your case, an algorithm decides if you qualify. Sounds efficient, right? But what if that same algorithm makes a mistake and you're denied help you desperately need? Who do you turn to? How do you challenge a decision made by a machine that no one really understands?

This isn't science fiction—it's the reality unfolding across India as AI-powered systems quietly take on bigger roles in governance. From ration card distributions to policing predictions, algorithms are increasingly shaping lives. The promise is tempting: faster services, reduced corruption, and

smarter resource use. But beneath the shiny tech lies a thorny problem—these AI decisions happen in a black box. They're often invisible, unexplainable, and unaccountable.^[1]

In a country as diverse and unequal as India, this raises serious concerns. What if these automated systems unintentionally deepen existing biases against marginalized communities? What protections do citizens have when machines decide their fate? And how do our laws, built for humans judging humans, adapt to a world where machines are the new decision-makers?

This study takes a deep dive into these questions. It explores the current legal landscape around AI in Indian governance, shining a light on the gaps and risks. It also looks at global

examples of how other countries are trying to regulate AI, hoping to learn lessons for India. The goal is clear: to ensure that as we welcome technology into the corridors of power, we don't lose sight of fairness, justice, and human dignity. Because in the end, no algorithm should ever replace the values that hold a society together.

Materials and Methods

To unpack how AI is shaping decision-making in Indian governance and the legal frameworks around it, this study uses a mix of document analysis and comparative legal research. First up, we dive deep into a variety of materials like government policies, official reports, court judgments, and draft bills related to AI, data protection, and digital governance in India. This helps us understand what rules are currently in place and where the gaps lie.

Next, we look beyond India's borders to see how other countries—especially the European Union, which is leading in AI regulation—are handling similar challenges. By comparing their laws and approaches, we aim to highlight possible lessons and strategies India can adopt.

To keep things grounded, the study also reviews academic articles, think tank reports, and media coverage to get a sense of the ethical, social, and practical issues tied to AI in governance.

While this study doesn't involve interviews or surveys, it relies on critical analysis to question how well current laws protect citizens' rights when decisions are automated. The approach is interdisciplinary, combining legal theory, policy analysis, and socio-ethical perspectives.

Ultimately, this method lets us build a clear picture of the evolving AI landscape in Indian governance and identify what needs fixing to keep technology fair, transparent, and accountable for all.

Fairness on Autopilot? Regulating Algorithmic Decisions in Indian Governance

1. The Algorithmic Turn in Governance

Gone are the days when governance meant dusty file cabinets, long queues, and bureaucratic middlemen. We're now entering an era where lines of code hold the reins of public administration. Think Aadhaar, Digital India, and the JAM trinity (Jan Dhan-Aadhaar-Mobile)—initiatives that have created the digital rails for governance. At the heart of these systems are algorithms—automated programs that sort, classify, filter, and decide.^[2] They're involved in everything from determining who gets a subsidy to identifying potential criminals through predictive policing software.

But here's the twist: algorithms, unlike human officials, don't have empathy. They don't understand nuance. And they certainly don't stop to ask, "Is this fair?" Yet we are increasingly putting them in charge of decisions that deeply affect people's lives. This shift—from human discretion to machine automation—demands a closer look. Because when the machine makes a mistake, it's not just a glitch in the system; it's a glitch in someone's life.

2. Algorithmic Bias: More Than Just a Bug

There's a myth that technology is neutral, that machines don't discriminate. But algorithms are not created in a vacuum. They are shaped by human choices—who builds them, what data trains them, and what outcomes they're optimized for. In India, where caste, class, gender, and religion intersect in deeply complex ways, algorithmic systems can quietly encode and amplify existing inequalities.

Imagine a facial recognition system that works best on fair-skinned male faces but struggles with darker-skinned women. That's not sci-fi—it's reality. Or consider a welfare algorithm that flags beneficiaries as duplicates based on mismatched biometric data. In states like Rajasthan and Jharkhand, thousands of poor families have been dropped from welfare lists due to such errors.^[3] These aren't just statistical anomalies—they're stories of empty kitchens and unpaid school fees.

3. The Legal Landscape: Gaps Wide Enough to Code Through

India has taken baby steps towards digital regulation. The Digital Personal Data Protection Act, 2023, finally offers some guardrails around data collection and use. But when it comes to automated decision-making, we're still in murky waters. There's no dedicated law that requires government agencies to explain or justify algorithmic decisions. There's no obligation to test these systems for bias or accuracy before deployment.

Compare this to the European Union's AI Act, which classifies algorithmic systems by risk level and imposes strict requirements on high-risk applications like policing or welfare. In India, however, most algorithmic decisions happen behind closed doors, far from the public eye. And when a citizen challenges an unfair decision, they often find themselves trapped in a Kafkaesque maze. Courts can't demand transparency from private vendors due to trade secrecy claims. Bureaucrats don't understand the tech. And the law simply hasn't caught up.^[4]

4. The Accountability Vacuum

Here's a question no one likes to answer: who do you hold accountable when an algorithm messes up? When a tribal woman is denied her MNREGA wages because her fingerprint didn't match, whose door does she knock on? The private company that made the biometric device? The local official who rubber-stamped the process? Or the faceless system that says "Access Denied" with no explanation?

This is what scholars call an "accountability vacuum." Unlike traditional decision-making, where a named officer can be held responsible, algorithmic decisions diffuse responsibility across multiple actors—none of whom are easily accountable. And the person at the receiving end? They're often left confused, angry, and helpless.

Redressal mechanisms are nearly non-existent. Most government apps and portals offer no clarity on why a decision was made or how to challenge it. Legal aid is limited, especially for those who live in rural or marginalized communities. This isn't just a legal problem; it's a moral one. It challenges our commitment to justice, fairness, and the rule of law.

5. Transparency and Explainability: Beyond Open Source

"Make it open source!" is the usual rallying cry for algorithmic transparency. But let's be real—how many citizens can read code? And even if they could, how many understand what the algorithm is doing at a systems level?

What we truly need is explainability: systems that can provide a plain-language explanation of why a decision was made. If a farmer is denied a subsidy, she should be able to get a clear answer, like: "Your land record was not updated in the last 3 years." Not some vague error code.

This is especially important in high-stakes areas like health insurance, housing benefits, or criminal profiling.

Yet in India, explainability is still an alien concept. Most government AI deployments are done through opaque public-private partnerships, with little public scrutiny. Procurement contracts don't require ethical audits. And citizens have no legal right to an explanation.

6. The Ethical Layer: Dharma in the Digital Age

Let's take a moment to look inward. Indian governance has always been rooted in deeper values-Dharma, Nyaya, equity, compassion. These aren't just ancient ideals; they are constitutional imperatives. Article 14 (equality), Article 15 (non-discrimination), and Article 21 (right to life and dignity) form the ethical backbone of our legal system.

Why shouldn't these principles apply to AI and algorithms? We need a digital dharma-a framework that asks: Is this system fair to the most vulnerable? Does it recognize historical disadvantages? Is it accountable to the people it serves? This means embedding fairness into code, not as an afterthought, but as a design principle. ^[5]

Techies call this "fairness-aware learning." Lawyers might call it "algorithmic due process." But the core idea is the same: no decision that affects a person's rights or well-being should be made without safeguards, oversight, and a path for appeal.

7. Way Forward: From Automation to Augmentation

Let's be clear-this isn't about rejecting technology. It's about making sure tech serves people, not the other way around. The future of governance isn't fully automated. It's augmented-where technology supports human judgment, not replaces it. We can build hybrid models where human caseworkers review algorithmic decisions before they're finalized. We can create algorithmic audit boards to regularly test systems for bias. We can mandate algorithm impact assessments before deployment, especially for high-risk use cases. And we absolutely need grievance redressal cells that specialize in tech-related complaints-places where citizens can get help, clarity, and justice. All of this requires not just new laws, but new attitudes. A culture of digital rights, where fairness is seen not as a luxury, but as a baseline. ^[6]

In the end, fairness isn't just something we hope for. It's something we build-brick by brick, line by line, byte by byte. In a nation where justice is a sacred promise, our algorithms should be nothing less than sacred code.

Because if we're going to trust machines with power, they better earn it.

Results and Discussion

Digging into the policies and legal documents around AI in Indian governance revealed a mixed bag-some progress but also major gaps. On paper, India is making moves with initiatives like the National AI Strategy and the Personal Data Protection Bill (PDPB). These frameworks show a clear intention to regulate AI and protect data privacy. However, when you peel back the layers, it's obvious they aren't enough yet.

One big issue? Transparency. Most AI systems used by government agencies operate like a black box. There's little clarity on how decisions are made or whether the data feeding these algorithms is fair and unbiased. Without transparency, affected individuals have almost zero chance to question or challenge decisions, which is a serious threat to due process and justice.

Another concern is accountability. The current legal landscape doesn't clearly assign responsibility when an AI system messes up. Is it the government agency, the software developer, or someone else? This confusion risks leaving citizens without effective remedies or support.

The study also found that most AI deployments overlook India's complex social realities. Marginalized groups-like lower-caste communities, women, and rural populations-might face disproportionate harm if biased data is used, reinforcing existing inequalities instead of breaking them down.

Comparing India with the European Union's AI Act and data protection laws showed us a roadmap: clear rules for transparency, mandatory impact assessments, and rights to explanation can help build trust and fairness. India can learn a lot here but needs to act fast.

In sum, the promise of AI in governance is huge, but so is the risk of digital injustice. Without stronger, human-centered laws and ethical oversight, we risk creating a system where machines rule silently-without empathy, fairness, or accountability.

This discussion underscores a pressing call: AI must be designed and regulated with people at its heart, or we risk losing the very justice and equality our legal system aims to protect.

Conclusion

AI is no longer just a futuristic idea-it's here, shaping how governments make decisions that affect millions of lives. This study shows that while India is stepping into the AI game with some policies and plans, the rules aren't quite ready to protect people from the risks that come with automated decision-making.

Without clear transparency and accountability, AI can easily turn from a helpful tool into a source of unfairness, especially for those who are already vulnerable. If we don't act now to build laws and systems that put people first, we risk creating a digital world where machines decide our fate but no one answers for mistakes.

The path forward is clear: India needs legal frameworks that demand transparency, ensure fairness, and hold those behind AI accountable. More importantly, these frameworks must be rooted in the country's diverse social fabric-understanding how technology affects different communities differently.

At its best, AI can help make governance smarter and more inclusive. But that will only happen if we keep justice and human dignity at the center of the conversation. Because at the end of the day, technology should serve people-not the other way around.

References

1. NITI Aayog. National Strategy for Artificial Intelligence: #AIforAll, 2018. <https://www.niti.gov.in>.
2. Ministry of Electronics and Information Technology, the Personal Data Protection Bill, 2019, Bill No. 373 of 2019, India Code 2019.
3. Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM (2021) 206 final Apr. 21, 2021.
4. Rajagopal K, Regulating AI in India: A Constitutional Imperative, 12 J. L. & Tech. 45 2023.
5. Reuben Binns, Algorithmic Accountability and Public Reason, 31 PHIL. & TECH. 543 2018.

6. Virginia Eubanks. Automating Inequality: how High-tech tools profile, police, and punish the poor (St. Martin's Press 2018).
7. Andreas Weller, Transparency: Motivations and Challenges, in explainable AI: interpreting, explaining and visualizing deep learning 23 (Springer 2019).
8. Anirban Bhattacharya. AI and Public Policy in India: Ethics, Law and Governance, ORF occasional paper no. 243 (2020), <https://www.orfonline.org>.
9. Solon Barocas, Moritz Hardt & Arvind Narayanan, Fairness and Machine Learning, 2019. <https://fairmlbook.org>.
10. Harini Suresh, John V Guttag. A Framework for Understanding Unintended Consequences of Machine Learning, 63 COMM'NS ACM 62 2020.