

Algorithmic Discrimination and Constitutional Equality: Reconstructing Human Rights Protection in the AI Era

Vikram Meghwal¹, Niharika Kumari²

^{1,2}Assistant Professor, Jaipur National University

Abstract:

Automated decision systems now sit inside credit approval, welfare targeting, recruitment and policing, yet Indian equality jurisprudence was built for human officials and identifiable acts of the State. This article asks whether Articles 14, 15 and 21 of the Constitution, read with the Digital Personal Data Protection Act, 2023, can carry the weight of algorithmic discrimination, and what a reconstructed framework would require. It pairs doctrinal analysis with a structured survey of 150 advocates, academics, research scholars and government or technology professionals. Awareness of algorithmic bias was high (82 per cent) but differed significantly across occupational groups. Respondents rejected the proposition that existing data-protection law is adequate (65 per cent disagreed) and strongly backed mandatory impact assessment (87 per cent) and a statutory right to explanation (75 per cent). The article proposes a four-part framework: a disparate-impact standard anchored in the non-arbitrariness limb of Article 14, horizontal duties on private deployers, procedural transparency rights, and an independent audit mechanism.

Keywords: algorithmic discrimination; Article 14; non-arbitrariness; automated decision-making; human rights; artificial intelligence; India.

I. Introduction

In 2019, researchers examining a risk-prediction tool used to allocate extra medical care across American hospitals found that Black patients given the same risk score as white patients were considerably sicker. The algorithm had used past healthcare spending as a stand-in for medical need, and because less money had historically been spent on Black patients, the tool read lower spending as lower need.¹ Race appeared nowhere in the model. The discrimination arrived through a proxy, and it stayed hidden from the clinicians who trusted the score. That pattern, neutral inputs producing unequal outputs, is the central difficulty automated decision-making poses for equality law, and it is not confined to any one country.

India is adopting such systems quickly and mostly without a dedicated statutory framework. Biometric authentication gates access to welfare entitlements, facial recognition tools are being piloted by police forces, lenders score applicants from mobile-phone behaviour, and employers filter résumés with

¹Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 Science 447 (2019).

automated screening software.² The Constitution promises equality before the law and equal protection of the laws in Article 14, forbids discrimination on specified grounds in Article 15, and guarantees life and personal liberty in Article 21. Yet each of these guarantees was elaborated by courts reviewing the reasoned decisions of human officials, whose motives and reasons can be questioned in a hearing. An opaque scoring model offers no such handle.

This article makes two contributions. First, it maps how algorithmic harm fits, and fails to fit, existing constitutional doctrine. Second, it reports an original survey of 150 respondents designed to measure how legal and technical professionals perceive that gap and what reforms they support. Three hypotheses guided the empirical work. H1: awareness of algorithmic bias differs across occupational groups. H2: respondents regard the Digital Personal Data Protection Act, 2023 as inadequate against algorithmic discrimination. H3: respondents favour ex ante procedural safeguards over reliance on judicial review alone. Part II sets out the doctrinal framework, Part III explains the method, Part IV presents findings, Part V proposes a reconstructed framework, and Parts VI and VII address limitations and conclude.

II. Equality Doctrine and the Algorithmic Problem

A. How Algorithms Discriminate

Scholars usually identify three routes by which automated systems reproduce inequality. The first is biased training data: a model learns from historical decisions and therefore inherits whatever prejudice shaped them. The second is proxy variables, where facially neutral attributes such as postcode, device type or spending history correlate with caste, religion, gender or class.³ The third is unequal accuracy: a commercial facial-analysis system tested in 2018 misclassified darker-skinned women at error rates of up to roughly 35 per cent, against under one per cent for lighter-skinned men.⁴ Feedback loops compound all three, since a model that directs more policing to one neighbourhood generates more recorded crime there, which then confirms the model.

Courts have begun to confront these problems unevenly. In *State v. Loomis*, the Wisconsin Supreme Court permitted a sentencing court to consider a proprietary risk score but required written warnings about its limits, including that the tool's methodology was a trade secret and that it compared the defendant to group data.⁵ Investigative reporting on the same family of tools alleged that Black defendants were nearly twice as likely as white defendants to be wrongly flagged as high risk.⁶ Whatever one makes of the methodological dispute that followed, the episode shows that discrimination claims against algorithms turn on evidence that the affected person is rarely able to obtain.

The concern is sharper in welfare administration, where an automated failure can mean a denied ration or pension. Authentication and eligibility systems that match biometrics or database entries will reject a share of legitimate claimants, and the people most likely to be rejected, manual labourers with worn fingerprints or residents of areas with poor connectivity, are not randomly distributed across the

²Vidushi Marda, *Artificial Intelligence Policy in India: A Framework for Engaging the Limits of Data-Driven Decision-Making*, 376 Phil. Trans. Royal Soc'y A 20180087 (2018).

³Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 Calif. L. Rev. 671 (2016).

⁴Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 Proc. Mach. Learning Rsch. 77 (2018).

⁵*State v. Loomis*, 881 N.W.2d 749 (Wis. 2016).

⁶Julia Angwin et al., *Machine Bias*, ProPublica (May 23, 2016).

population. The Court's Aadhaar decision upheld the scheme subject to safeguards while recognising that exclusion from benefits was a serious risk, but it did not frame the problem as one of discriminatory impact.⁷ Opacity compounds the difficulty. As Pasquale argues, decisions made inside systems the affected person cannot inspect leave little room for the contestation on which accountability ordinarily depends.⁸

B. The Indian Constitutional Toolkit

Indian equality doctrine has moved well beyond formal sameness. In *E.P. Royappa v. State of Tamil Nadu*, Bhagwati J. described equality as a dynamic concept and arbitrariness as its opposite, a formulation later folded into the test for reasonableness under Article 14.⁹ *Maneka Gandhi v. Union of India* tied Articles 14, 19 and 21 into a single guarantee of fair, just and reasonable procedure.¹⁰ More recently, *Shayara Bano v. Union of India* applied manifest arbitrariness to strike down legislation, and *Navtej Singh Johar v. Union of India* read Articles 14, 15 and 21 together to protect dignity and autonomy against majoritarian morality.¹¹ *K.S. Puttaswamy v. Union of India* recognised privacy as a fundamental right and, in the Aadhaar decision that followed, the Court accepted that data-driven governance must satisfy proportionality.¹²

Indian courts have also recognised that equal treatment can mask unequal effect. In *Anuj Garg v. Hotel Association of India*, the Supreme Court struck down a ban on women working in bars, reasoning that protective stereotypes cannot justify exclusion and that the law's real-world operation matters, not only its stated purpose.¹³ In *Nitisha v. Union of India*, the Court went further and acknowledged indirect discrimination, accepting that a facially neutral criterion may disadvantage a protected group and so breach Articles 14 and 15 without any discriminatory intent.¹⁴ That reasoning is the closest available doctrinal bridge to algorithmic proxy discrimination, because the injury is located in the pattern of outcomes and not in the decision-maker's state of mind.

Comparative practice has moved in a similar direction by legislating around the problem. The European Union's Artificial Intelligence Act classifies systems used in employment and in assessing creditworthiness or access to essential services as high-risk, attaching duties of data governance, documentation, human oversight and accuracy.¹⁵ The American experience has been more fragmented, relying on existing anti-discrimination statutes whose application to opaque models remains contested. Neither approach can simply be transplanted. Indian regulators and courts work under different institutional constraints, and a framework must be fitted to the constitutional text and to the realities of enforcement capacity.

These authorities supply useful principles, but three structural limits remain. The first is the requirement of State action: Articles 14 and 15(1) bind the State and its instrumentalities, while much algorithmic

⁷ *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2019) 1 SCC 1 (India).

⁸ Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (2015).

⁹ *E.P. Royappa v. State of Tamil Nadu*, (1974) 4 SCC 3, ¶ 85 (India).

¹⁰ *Maneka Gandhi v. Union of India*, (1978) 1 SCC 248 (India).

¹¹ *Shayara Bano v. Union of India*, (2017) 9 SCC 1 (India); *Navtej Singh Johar v. Union of India*, (2018) 10 SCC 1 (India).

¹² *K.S. Puttaswamy v. Union of India*, (2017) 10 SCC 1 (India); *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2019) 1 SCC 1 (India).

¹³ *Anuj Garg v. Hotel Ass'n of India*, (2008) 3 SCC 1 (India).

¹⁴ *Nitisha v. Union of India*, (2021) 15 SCC 125 (India).

¹⁵ Regulation (EU) 2024/1689, art. 6 & annex III, 2024 O.J. (L 2024/1689) 1.

decision-making occurs in private lending, hiring and platform work. Article 15(2) reaches private conduct only as regards access to public places and facilities, and Article 21 has been extended horizontally in a patchy way. The second limit is evidentiary. The non-arbitrariness test asks for reasons, but a model can output a score without any reasons a court could assess. The third limit is that classic discrimination analysis looks for a prohibited ground, whereas proxy discrimination operates through variables that are not themselves prohibited. Table 1 summarises how each constitutional guarantee meets algorithmic harm.

Table 1: Constitutional guarantees and their reach over algorithmic harm

Provision	Core protection	Algorithmic application	Principal gap
Art. 14	Equality before law; non-arbitrariness	Scoring or ranking by State bodies without reasons	Binds State only; no disparate-impact test
Art. 15(1)-(2)	No discrimination on listed grounds	Direct use of caste, sex or religion in model features	Proxy variables escape the listed grounds
Art. 16	Equality in public employment	Automated shortlisting in public recruitment	No audit access for candidates
Art. 21	Life, liberty, dignity, privacy, fair procedure	Welfare denial through authentication failure	Procedural content undefined for automated decisions
DPDP Act, 2023	Consent, purpose limitation, grievance redress	Data processing underlying profiling	No right to explanation; broad State exemptions

Source: Author's analysis of the Constitution of India and the Digital Personal Data Protection Act, 2023, No. 22 of 2023, India Code.

The statutory layer does not close these gaps. The Digital Personal Data Protection Act, 2023 regulates consent and processing but contains no provision comparable to Article 22 of the European Union's General Data Protection Regulation, which restricts decisions based solely on automated processing.¹⁶ The Act also permits wide exemptions for State instrumentalities. Its operative obligations are being phased in, with core data-fiduciary duties scheduled to commence in May 2027 under the 2025 Rules.¹⁷ Anti-discrimination law proper remains sector-specific and was drafted with human conduct in mind.

III. Research Design and Methodology

The empirical component used a cross-sectional, purposive survey. The target population comprised persons whose work brings them into regular contact with automated decision-making or with the law governing it. The sample (N = 150) included 50 advocates, 40 academics, 30 research scholars and 30

¹⁶Digital Personal Data Protection Act, 2023, No. 22 of 2023, India Code (India); Regulation (EU) 2016/679, art. 22, 2016 O.J. (L 119) 1.

¹⁷Digital Personal Data Protection Rules, 2025, G.S.R. 846(E) (Nov. 13, 2025) (India).

government or technology professionals. The instrument had three sections: a screening item on awareness of algorithmic bias (yes or no), five statements measured on a five-point Likert scale (1 = strongly disagree, 5 = strongly agree), and a block rating the perceived risk of discriminatory harm in six sectors on a five-point scale.

The five statements were drafted to test the hypotheses directly. S1 asked whether Article 14 should bind private deployers of algorithmic systems; S2 whether impact assessment should be mandatory for high-risk systems; S3 whether individuals should have a statutory right to explanation; S4 whether the 2023 Act adequately addresses algorithmic discrimination; and S5 whether judicial review alone is sufficient. For analysis, scores of four or five were treated as agreement and one or two as disagreement. Group differences in awareness were tested with a chi-square test of independence, and differences in sector ratings with a Kruskal-Wallis test, since Likert data are ordinal. A point-biserial correlation examined whether awareness predicted support for impact assessment. Significance was set at $p < 0.05$.

Note on data: The respondent-level dataset analysed in this draft is a simulated dataset, parameterised to resemble plausible response patterns and generated to demonstrate the analytical design. It is not field data. The figures, tables and test statistics must be replaced with results from actual administration of the instrument before this article is submitted or cited as empirical evidence.

IV. Findings and Analysis

A. Awareness of Algorithmic Bias

Overall, 82 per cent of respondents (123 of 150) reported awareness of algorithmic bias as a legal or policy problem. Awareness was highest among research scholars (93 per cent) and academics (90 per cent), followed by advocates (80 per cent), and lowest among government and technology professionals (63 per cent). The difference across groups was statistically significant ($\chi^2 = 11.56$, $df = 3$, $p = 0.009$), supporting H1. Figure 1 shows the distribution. The lower figure among government and technology professionals is notable because this group builds and procures the systems at issue; it suggests that technical familiarity with a tool does not translate into familiarity with its discriminatory potential.




Figure 1: Awareness of algorithmic bias by occupational group (N = 150)

B. Attitudes Toward Reform

Table 2 reports descriptive statistics for the five statements. Support for procedural safeguards was strong. Some 87 per cent of respondents agreed that impact assessment should be mandatory for high-risk systems (mean 4.32, SD 0.72), 75 per cent supported a statutory right to explanation (mean 4.04), and 73 per cent agreed that Article 14 should reach private employers (mean 3.98). Only 9 per cent considered the 2023 Act adequate, while 65 per cent disagreed (mean 2.19), supporting H2.

Table 2: Descriptive statistics for attitudinal statements (N = 150)

Statement	Mean	SD	Agree (%)	Neutral (%)	Disagree (%)
S1 Art. 14 should bind private employers	3.98	0.80	72.7	24.7	2.7
S2 Mandatory impact assessment for high-risk systems	4.32	0.72	86.7	12.7	0.7
S3 Statutory right to explanation	4.04	0.83	74.7	22.7	2.7
S4 DPDP Act, 2023 adequately addresses discrimination	2.19	0.94	8.7	26.7	64.7
S5 Judicial review alone is sufficient	2.83	1.07	26.0	36.7	37.3

Source: Author's survey (simulated dataset; see Note on data).

Agree = scores 4-5; Disagree = scores 1-2.




Figure 2: Distribution of responses to the five attitudinal statements (N = 150)

The responses to S5 deserve separate comment. Opinion on judicial review was divided: 37 per cent disagreed that it suffices, 37 per cent were neutral and 26 per cent agreed. The sizeable neutral group is consistent with a view that courts remain important but cannot work without upstream disclosure. This reading supports H3 only in part, because the evidence shows strong preference for ex ante safeguards (S2 and S3) rather than a clear rejection of litigation. Awareness had no meaningful association with

support for impact assessment ($r = 0.06$, $p = 0.44$), which suggests the preference for safeguards is broad-based and does not depend on prior familiarity with the problem.

A further question is whether these attitudes vary by occupation, since awareness did. They largely did not. Table 3 reports mean scores by group, and Kruskal-Wallis tests found no significant group differences on any of the five statements (all $p > 0.27$). Government and technology professionals, whose awareness was lowest, supported mandatory impact assessment (mean 4.30) about as strongly as academics (4.35) and advocates (4.26). The pattern suggests that the demand for safeguards does not depend on prior exposure to the literature on algorithmic bias. It may instead reflect a more general expectation, shared across the surveyed professions, that consequential decisions about persons should be explainable and reviewable.

Table 3: Mean agreement with attitudinal statements by group (1-5 scale)

Group	n	S1	S2	S3	S4	S5
Advocates	50	3.92	4.26	4.14	2.30	2.70
Academics	40	3.98	4.35	4.10	2.08	2.82
Research scholars	30	4.07	4.40	3.83	2.23	3.00
Govt./technology professionals	30	4.00	4.30	4.00	2.10	2.87

Source: Author's survey (simulated dataset; see Note on data). Kruskal-Wallis p -values for S1-S5 range from 0.28 to 0.93.

C. Perceived Risk Across Sectors

Respondents rated six sectors for the risk of discriminatory harm. Policing and surveillance ranked highest (mean 4.27), followed by credit scoring (3.96), welfare targeting (3.93), recruitment (3.72), healthcare triage (3.56) and education admissions (3.37). The differences across sectors were significant (Kruskal-Wallis $H = 105.83$, $p < 0.001$). Figure 3 presents means with 95 per cent confidence intervals. The ordering tracks the stakes to liberty and livelihood rather than the volume of automated decisions: policing and welfare, where the State is the deployer, rank above education, where outcomes are more contestable.




Figure 3: Perceived risk of discriminatory harm by sector (mean, 95% CI)

Taken together, the findings point in one direction. Professionals who engage with the subject see existing safeguards as thin, regard private deployers as appropriately within the reach of equality norms, and prefer rules that operate before harm occurs. That preference is also what the doctrinal analysis suggests: a person denied a loan or flagged by a policing model cannot meaningfully plead arbitrariness without access to the logic of the decision.

V. Reconstructing Constitutional Protection

The evidence and doctrine support a framework with four components, each building on principles Indian courts have already accepted.

A. A Disparate-Impact Standard Within Article 14

Courts should recognise that a system producing systematically unequal outcomes for a protected group raises a presumption of arbitrariness, even where no prohibited ground was used as an input. Indian law has not formally adopted disparate impact, but the manifest-arbitrariness doctrine in *Shayara Bano* already permits scrutiny of effects, not only intent.¹⁸ Once a claimant shows statistical disparity, the burden would shift to the deployer to demonstrate a legitimate purpose and the absence of a less discriminatory alternative. This mirrors the burden-shifting scholars have proposed in other jurisdictions and answers the evidentiary asymmetry between individual and system.¹⁹

B. Horizontal Duties on Private Deployers

Because most consequential automated decisions in credit, employment and platform work are private, constitutional norms need a statutory bridge. The Court has long accepted that Article 21 rights may be enforced against non-State actors discharging public functions, and the horizontal reach of dignity was a theme in *Navtej Singh Johar*.²⁰ Parliament can legislate duties of non-discrimination on deployers of high-risk systems, enforceable before a regulator and civil courts. The 73 per cent support for S1 suggests that the profession is prepared for that extension.

C. Procedural Transparency Rights

A statutory right to meaningful information about the logic of an automated decision, together with human review on request, would supply the reasons that Article 14 and Article 21 presuppose. This right need not require disclosure of source code. Disclosure of the principal features, the data categories used and the results of bias testing would let affected persons and courts assess arbitrariness. The *Maneka Gandhi* guarantee of fair procedure is the natural constitutional anchor.²¹

The scope of any explanation right deserves care. Scholars have debated whether the GDPR actually confers a right to an explanation of individual decisions, and have concluded that its text supports only a more limited right to information about the existence and general logic of automated processing.²² An

¹⁸*Shayara Bano*, (2017) 9 SCC 1.

¹⁹Sandra Wachter, Brent Mittelstadt & Chris Russell, *Why Fairness Cannot Be Automated: Bridging the Gap Between EU Non-Discrimination Law and AI*, 41 Comput. L. & Sec. Rev. 105567 (2021).

²⁰*Navtej Singh Johar*, (2018) 10 SCC 1.

²¹*Maneka Gandhi*, (1978) 1 SCC 248.

²²Sandra Wachter, Brent Mittelstadt & Luciano Floridi, *Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation*, 7 Int'l Data Priv. L. 76 (2017).

Indian provision should be drafted more explicitly. It should state what must be disclosed (the main factors, the data categories and the outcome of any bias testing), to whom (the affected person, and a tribunal or auditor on request), and when (before final adverse action where practicable). It should also guarantee human review by a person with authority to alter the outcome, because a review that cannot change the result gives only the appearance of procedural fairness.

D. Independent Audit and Impact Assessment

Mandatory algorithmic impact assessment, which 87 per cent of respondents endorsed, should be paired with an independent audit mechanism and public registers for State-deployed systems. The European Union's approach, which requires risk management, data governance and human oversight for high-risk systems and fundamental-rights impact assessments for certain deployers, offers a reference point, although India would need to calibrate it to regulatory capacity.²³ Impact assessment before deployment is cheaper to remedy than discrimination discovered after a welfare programme has excluded thousands of beneficiaries.

Table 4: Proposed framework and its doctrinal anchors

Component	Content	Doctrinal anchor	Survey support (%)
Disparate-impact standard	Presumption of arbitrariness on statistical disparity; burden shifts to deployer	Art. 14; Royappa; Shayara Bano	n/a (doctrinal)
Horizontal duties	Statutory non-discrimination duties on private deployers of high-risk systems	Arts. 14, 15, 21; Johar	72.7
Transparency rights	Meaningful information on logic; human review on request	Art. 21; Maneka Gandhi	74.7
Audit and assessment	Mandatory impact assessment; independent audit; public registers	Art. 14; proportionality (Puttaswamy)	86.7

Source: Author's synthesis; survey percentages from Table 2 (S1, S3, S2).

E. Anticipated Objections

Four objections deserve a reply. The first is that mandatory assessment and audit will burden small firms and slow innovation. The answer lies in proportionality: duties can be tiered by the risk of the application, so that a recommendation engine for entertainment faces lighter obligations than a model deciding access to credit, welfare or liberty. The second objection is that models are trade secrets. Transparency here means disclosure of features, data categories and test results to an auditor or tribunal under confidentiality, which courts already manage in commercial disputes, and *Loomis* shows that

²³Regulation (EU) 2024/1689, arts. 9-10, 14, 27, 2024 O.J. (L 2024/1689) 1.

secrecy is a weak ground for denying a defendant information about the tool used against them.²⁴ The third objection is that a burden-shifting presumption sits uneasily with the ordinary rule that the claimant must prove a breach. Indian law, however, already shifts evidential burdens where one party holds the relevant information, and the presumption would arise only after the claimant has shown a statistical disparity.

The fourth objection concerns capacity. Regulators and courts may lack the technical expertise to evaluate models. That concern is real, and it argues for sequencing and not for inaction. A reasonable path would begin with State-deployed systems, where Articles 14 and 21 apply directly and public registers are administratively simple, then extend duties to private deployers in the highest-risk sectors identified by respondents: policing, credit and welfare targeting. An expert body attached to the proposed audit mechanism could also advise tribunals, so that judges do not have to evaluate statistical evidence unaided. None of these steps requires abandoning the constitutional foundation. They give existing principles an institutional form capable of operating at the scale of automated decision-making.

VI. Limitations

Several limitations qualify these conclusions. The sample was purposive, not probability-based, so the results describe the surveyed professionals and cannot be generalised to the Indian population. Self-reported awareness is an imperfect proxy for knowledge. Attitudinal statements measure support in principle and not willingness to bear the cost of compliance. Finally, as noted above, the dataset reported in this draft is simulated and the statistics are placeholders for results from actual fieldwork. The doctrinal analysis does not depend on the survey and stands independently.

VII. Conclusion

Algorithmic systems test the equality guarantee at the point where it is least equipped to respond: the absence of a reasoning human decision-maker and the absence of a prohibited ground. Indian constitutional law already holds the resources to respond, namely non-arbitrariness, fair procedure and dignity, but those resources need translation into an effects-based standard, duties that reach private deployers, procedural rights that surface reasons, and institutional audit. The survey evidence, once validated with field data, indicates a professional constituency ready for that shift. Equality before the law cannot depend on whether a decision was made by a person or a model, and the legal framework should say so explicitly.

²⁴*State v. Loomis*, 881 N.W.2d 749 (Wis. 2016).