



The Indian Journal for Research in Law and Management

Open Access Law Journal – Copyright © 2026

Editor-in-Chief – Dr. Muktai Deb Chavan; Publisher – Alden Vas; ISSN: 2583-9896

This is an Open Access article distributed under the terms of the Creative Commons Attribution-Non-Commercial-Share Alike 4.0 International (CC-BY-NC-SA 4.0) License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium provided the original work is properly cited.

ALGORITHMIC DISCRIMINATION AND SOCIAL INEQUALITY: EMERGING CHALLENGES FOR INDIAN LAW

~ *Darshan.S*

Abstract

The rapid expansion of Artificial Intelligence (AI), automated decision-making systems and algorithmic technologies is transforming the manner in which individuals interact with institutions, markets and the State. While algorithms are often perceived as objective and technologically neutral, their design, training data and deployment may reproduce or amplify existing patterns of social inequality. In a socially diverse country such as India, where inequalities based on caste, gender, class, language, region, disability and economic status continue to influence access to opportunities and resources, algorithmic decision-making presents emerging sociological and legal challenges. This research paper examines algorithmic discrimination as a contemporary form of structural inequality and analyses how automated systems may perpetuate historical and institutional biases while making discriminatory outcomes less visible and more difficult to challenge.

The paper adopts a socio-legal and interdisciplinary approach to examine the relationship between algorithmic systems and constitutional principles of equality, non-discrimination, dignity and privacy. It analyses the relevance of Articles 14, 15, 16 and 21 of the Constitution of India, judicial developments concerning privacy and informational autonomy, and the emerging framework governing digital personal data and artificial intelligence in India. Particular attention is given to algorithmic applications in employment, financial services, welfare delivery, education, public administration and other areas where automated decisions may significantly affect individuals' social and economic opportunities. The study also examines the problems of

algorithmic opacity, lack of explainability, unequal access to technology, accountability gaps and difficulties in establishing discriminatory intent or causation.

From a sociological perspective, the paper argues that technological systems operate within existing social structures and therefore cannot be examined independently of the inequalities embedded in the societies in which they are developed and deployed. The study concludes by exploring possible legal and institutional responses, including algorithmic impact assessments, transparency requirements, human oversight, independent auditing, effective grievance mechanisms and stronger anti-discrimination safeguards. It seeks to contribute to the emerging Indian discourse on responsible AI by examining how law can respond to technological systems without losing sight of the underlying social structures that produce inequality.

Keywords: Algorithmic Discrimination, Artificial Intelligence, Social Inequality, Caste, Gender, Algorithmic Bias, Equality, Indian Constitution, Data Protection, Socio-Legal Studies.

1. INTRODUCTION

Technology has never developed independently of society. Every technological system is designed within a particular social, economic and institutional environment and consequently reflects, to some extent, the assumptions, priorities and limitations of that environment. The increasing use of Artificial Intelligence and algorithmic decision-making has made this relationship more significant. Decisions that were previously made directly by human beings are increasingly assisted or influenced by computer-based systems capable of processing large quantities of information and identifying patterns that may not be immediately visible to human decision-makers.

Algorithms are now associated with a wide range of activities. Recruitment systems may rank applicants, financial institutions may use automated systems to assess creditworthiness, educational institutions may use digital tools to evaluate students, online platforms may determine what information users see, and public authorities may employ digital systems for the administration of welfare and public services. Such applications may provide efficiency and

consistency, but they also raise an important question: what happens when the information on which an automated decision is based reflects an unequal society?

The assumption that an algorithm is neutral merely because it is mathematical or computer-based is problematic. An algorithm does not ordinarily create its own social understanding. It operates through data, categories, variables, objectives and rules determined by people and institutions. Historical data may contain the consequences of previous discrimination. A system trained on such data may learn patterns that correspond with existing inequalities. Even where a prohibited characteristic such as caste or gender is not expressly included, other variables may operate as substitutes or proxies for socially significant characteristics.

This issue is particularly important in India. Indian society has historically been characterised by multiple and overlapping forms of inequality. Caste, gender, economic status, language, geographical location, disability and access to education have influenced people's opportunities in different ways. The constitutional framework attempts to address these inequalities through equality, affirmative action, fundamental rights and other social-justice mechanisms. The introduction of automated decision-making therefore cannot be considered merely as a technological development. It raises questions about how constitutional equality should operate when decisions affecting individuals are increasingly mediated by technological systems.

The central concern of this paper is not that technology is inherently discriminatory. Rather, the concern is that technological systems may reproduce existing inequalities when they are designed or deployed without adequate consideration of the social environment in which they operate.

Algorithmic discrimination can therefore be understood as a problem of both technology and social structure. It concerns not only the functioning of algorithms but also the data selected, the categories used, the institutions employing the systems and the individuals affected by their decisions.

The legal challenge is equally complex. Traditional discrimination law has often been concerned with identifiable decision-makers and identifiable acts of discrimination. Automated systems complicate this model. A person affected by an algorithmic decision may not know why a particular decision was made, what data was considered, which variables influenced the result or who should be held responsible for the outcome.

The problem becomes even more difficult when discriminatory consequences occur without an express discriminatory intention. A criterion may appear neutral on its face but produce disproportionately adverse consequences for a particular social group. This distinction between discriminatory intention and discriminatory effect is increasingly important within Indian constitutional jurisprudence.

The Supreme Court's decision in *Lt. Col. Nitisha v. Union of India* provides an important constitutional foundation in this regard. The Court examined indirect discrimination and recognised that apparently neutral criteria can produce discriminatory consequences because they operate within existing social and institutional structures. This reasoning is particularly relevant to algorithmic systems, where apparently neutral variables may generate unequal outcomes.

The challenge for Indian law is therefore to develop mechanisms capable of examining not merely whether an algorithm was intentionally designed to discriminate, but also whether its operation produces unjustifiable patterns of disadvantage.

2. RESEARCH PROBLEM

The emergence of algorithmic decision-making creates a gap between traditional understandings of discrimination and the realities of technologically mediated decision-making.

Traditional discrimination is generally easier to identify where an individual or institution openly differentiates between persons on the basis of a prohibited ground. Algorithmic discrimination may occur through a much more complex chain involving data collection, data classification, model development, training, testing, deployment and subsequent decision-making.

An individual denied employment by a human decision-maker may ask why the application was rejected and challenge the decision. A person rejected by an automated recruitment system may not know whether the decision resulted from educational qualifications, employment history, geographical information, linguistic characteristics, online behaviour or some other variable.

The resulting research problem is therefore:

How should Indian law respond when apparently neutral algorithmic systems reproduce or intensify existing forms of social inequality without necessarily involving direct or intentional discrimination?

This question requires an examination of both legal doctrine and social structure.

3. OBJECTIVES OF THE STUDY

The study has the following objectives:

1. To examine the concept of algorithmic discrimination from a sociological and legal perspective.
2. To analyse the relationship between algorithmic decision-making and existing forms of social inequality in India.
3. To examine the relevance of Articles 14, 15, 16 and 21 of the Constitution of India to algorithmic decision-making.
4. To analyse the significance of indirect and structural discrimination for algorithmic systems.
5. To examine challenges relating to transparency, explainability, accountability and access to remedies.
6. To study the relevance and limitations of India's emerging data-protection and AI governance framework.
7. To suggest legal and institutional mechanisms for reducing discriminatory outcomes arising from automated decision-making.

4. RESEARCH QUESTIONS

The study is guided by the following questions:

1. What constitutes algorithmic discrimination?
2. How can algorithmic systems reproduce existing social inequalities?
3. Can a facially neutral algorithmic criterion result in indirect discrimination under Indian constitutional law?
4. How are Articles 14, 15, 16 and 21 relevant to algorithmic decision-making?
5. What challenges arise in establishing responsibility for discriminatory algorithmic outcomes?
6. Does data protection law by itself provide an adequate response to algorithmic discrimination?
7. What institutional safeguards are required to ensure that automated decision-making remains consistent with constitutional equality and human dignity?

5. HYPOTHESIS

The study proceeds on the hypothesis that algorithmic systems are not socially neutral instruments because their development and deployment occur within existing institutional and social structures. Where historical inequalities are reflected in datasets, institutional practices or decision-making objectives, automated systems may reproduce or amplify those inequalities even in the absence of an intention to discriminate.

It is further hypothesised that India's existing constitutional principles provide a foundation for addressing algorithmic discrimination, but specific mechanisms concerning algorithmic transparency, impact assessment, auditing, explanation and effective remedies require further development.

6. RESEARCH METHODOLOGY

The research adopts a doctrinal and socio-legal methodology.

The doctrinal component examines constitutional provisions, legislation, judicial decisions and policy documents relevant to equality, privacy, data protection and technological governance.

The socio-legal component examines algorithmic discrimination as a phenomenon situated within broader structures of social inequality. The paper draws upon concepts of structural discrimination, substantive equality, institutional inequality and digital exclusion.

The research is primarily qualitative and analytical. It does not attempt to measure algorithmic discrimination through an original statistical dataset. Instead, it examines the legal and sociological implications of algorithmic decision-making and identifies areas in which further empirical research may be required.

7. UNDERSTANDING ALGORITHMIC DISCRIMINATION

An algorithm may broadly be understood as a defined sequence of computational instructions used to process information and produce an output. In simple applications, the process may be relatively transparent. Contemporary AI systems, however, can involve highly complex statistical models capable of identifying patterns from large datasets.

Algorithmic discrimination occurs when the use of an algorithmic system results in differential or unequal treatment that disadvantages individuals or groups on grounds connected to protected or socially significant characteristics. The concept is wider than intentional discrimination. There are several ways in which discriminatory outcomes can arise.

7.1 Historical data bias

A system may be trained using historical information. If the historical decisions reflected social inequality, the resulting model may reproduce those patterns.

For example, if a particular group historically received fewer employment opportunities, a recruitment model trained on historical recruitment data may incorrectly interpret the under-representation of that group as evidence that members of the group are less suitable candidates.

The problem does not necessarily originate in the mathematical model. It may originate in the social history represented by the data.

7.2 Selection bias

The dataset itself may not represent the population adequately.

A facial-recognition system developed using an insufficiently diverse dataset may perform differently across demographic groups. Similarly, a language-processing system trained primarily on dominant languages may perform less effectively when confronted with regional or minority languages.

This creates an important sociological problem: the absence of adequate representation in data can itself become a source of disadvantage.

7.3 Proxy discrimination

An algorithm may not directly use caste, gender or another protected characteristic but may use variables that correlate with those characteristics.

Geographical location, educational institution, occupation, family background, language or patterns of consumption may sometimes function as indirect indicators of socioeconomic or demographic characteristics.

Thus, removing an explicitly discriminatory variable does not automatically eliminate discriminatory effects.

7.4 Feedback effects

Algorithms may also create feedback loops.

Suppose an automated system repeatedly identifies a particular neighbourhood as presenting a higher risk. Increased surveillance in that neighbourhood may generate more recorded incidents. Those additional records may then be fed into the system, reinforcing its original assessment.

In this manner, a prediction can influence the social reality from which future predictions are generated.

7.5 Automation of institutional practices

An algorithm may also reproduce the assumptions of the institution using it.

If an organisation has historically valued certain educational institutions, employment histories or social indicators, a system designed to imitate past institutional decisions may transform those preferences into apparently objective computational rules. The resulting discrimination becomes difficult to identify because the original institutional judgment has been translated into technological form.

8. ALGORITHMIC DISCRIMINATION AS A FORM OF STRUCTURAL INEQUALITY

Sociological analysis is essential because discrimination does not always arise from individual prejudice. Social inequality may be embedded within institutions, social practices and historical patterns. Individuals may experience disadvantage even where no particular person consciously intends to discriminate against them. Structural inequality therefore provides a useful framework for understanding algorithmic discrimination.

An algorithm operates within a structure. The structure determines what information is collected, which categories are considered relevant, what objective is pursued and what outcome is regarded as successful.

Consider recruitment.

If a company historically recruited individuals from a limited number of institutions, an automated system trained on past recruitment decisions may identify the characteristics of those applicants as indicators of future success. The algorithm may consequently prefer candidates with similar profiles.

The system may therefore reproduce an existing social pattern while appearing neutral.

The same concern can arise in financial services. Creditworthiness may be assessed using variables associated with income, employment history, financial transactions and geographical information. Persons from economically vulnerable communities may have less conventional financial histories and therefore receive poorer automated assessments. The resulting inequality may then restrict access to credit, housing, education or business opportunities.

The sociological importance of this process lies in the possibility that disadvantage becomes cumulative.

A person who has historically experienced educational disadvantage may receive fewer employment opportunities. Limited employment may result in weaker financial records. Weaker financial records may influence access to credit. Limited access to credit may restrict economic mobility.

If automated systems rely on these historical outcomes, the algorithm may reproduce the original inequality at each subsequent stage.

Technology can consequently transform temporary disadvantage into a continuing cycle.

9. CASTE AND ALGORITHMIC DECISION-MAKING IN INDIA

Caste presents a particularly important issue for Indian algorithmic governance.

The Indian Constitution expressly recognises caste as a prohibited ground of discrimination in Articles 15 and 16 while simultaneously permitting affirmative measures for historically disadvantaged groups.

Digital systems may encounter caste in several different ways.

First, caste may appear directly in datasets. Second, caste may be reflected indirectly through geographical, educational or socioeconomic variables. Third, datasets may exclude caste

information altogether, making it difficult to determine whether a system produces unequal outcomes across caste groups.

The third problem requires particular attention.

The absence of caste data does not necessarily mean the absence of caste discrimination. If social inequalities are already embedded in the variables used by an algorithm, ignoring caste may make the resulting inequality more difficult to identify.

This creates a paradox. Collecting sensitive information can create privacy concerns, while refusing to collect such information may make discrimination impossible to measure.

A meaningful legal framework must therefore distinguish between the collection of information for legitimate equality monitoring and the inappropriate use or disclosure of sensitive personal information.

The issue becomes especially significant where algorithmic systems are used in welfare administration, education, employment or public services.

An automated system that determines eligibility for a welfare programme should not merely be evaluated on whether it processes applications efficiently. It must also be examined for whether its operation disproportionately excludes communities that the programme was intended to benefit.

10. GENDER AND ALGORITHMIC DISCRIMINATION

Gender discrimination illustrates how neutral criteria can have unequal consequences.

Recruitment algorithms provide a useful example. If historical recruitment patterns reflect male dominance in particular occupations, a system trained on previous recruitment decisions may treat male-associated employment histories as stronger indicators of suitability.

Even without using the variable "gender", the system may rely on characteristics strongly correlated with gender. The problem is not merely technological. It reflects the historical distribution of opportunities between men and women. Indian constitutional jurisprudence has increasingly recognised that equality cannot always be understood as identical treatment.

In *Lt. Col. Nitisha v. Union of India*, the Supreme Court considered the doctrine of indirect discrimination and examined how facially neutral criteria could disproportionately disadvantage women. The Court emphasised that discrimination may be located in the effect of a criterion within an existing social structure rather than solely in the intention behind the criterion.

This reasoning has considerable relevance to algorithmic systems.

An algorithm may treat every applicant according to exactly the same computational rule and yet produce unequal consequences. Formal uniformity therefore cannot by itself establish substantive equality.

11. CLASS, POVERTY AND DIGITAL INEQUALITY

Economic inequality introduces another dimension.

Algorithmic systems frequently rely upon digital records. Persons with greater economic resources tend to generate more extensive and stable digital records through formal employment, banking, education and access to digital services.

Economically vulnerable persons may have fewer such records.

The result is a potential digital divide between people who are highly represented in datasets and people who are poorly represented.

This can create a form of "data disadvantage".

For example, a person employed in the informal economy may not possess the same conventional employment documentation as a formally employed person. A person without extensive banking activity may have a different financial profile from someone with a long history of formal financial transactions.

If an automated system interprets the absence of data as evidence of risk or low reliability, individuals already disadvantaged by economic circumstances may face an additional technological disadvantage.

The concern is therefore not merely whether the algorithm is accurate in statistical terms. The deeper question is whether the categories used to define accuracy are themselves socially neutral.

12. LANGUAGE, REGION AND DIGITAL EXCLUSION

India's linguistic diversity presents another challenge.

Many digital systems perform better in languages for which substantial training and evaluation data are available. Individuals communicating in less-resourced languages may therefore experience reduced access to automated services.

This can affect education, public administration, banking, legal services and welfare delivery. Language disadvantage may intersect with regional and socioeconomic inequality. If access to a public service increasingly depends upon an individual's ability to interact with a particular digital system, language competence can become a condition for exercising rights or obtaining benefits.

This raises questions of substantive equality and accessibility.

Digital inclusion must therefore mean more than providing internet connectivity. It must include meaningful access to digital services in forms that can be understood and used by diverse populations.

13. DISABILITY AND AUTOMATED SYSTEMS

Persons with disabilities may also experience specific forms of algorithmic exclusion.

Digital systems may assume particular patterns of communication, mobility, vision, hearing or cognitive processing. When such assumptions are incorporated into automated decision-making, persons with disabilities may be disadvantaged even where disability is not expressly considered.

The problem demonstrates the importance of reasonable accommodation.

Equality cannot always be achieved by applying identical technological procedures to everyone.

In some situations, equal participation requires adaptation of the system to different needs.

This principle is consistent with the broader development of substantive equality in Indian constitutional law.

14. CONSTITUTIONAL FRAMEWORK

14.1 Article 14

Article 14 guarantees equality before the law and equal protection of the laws.

Algorithmic decision-making raises questions of both classification and arbitrariness. Where an algorithm determines access to employment, public services or benefits, the criteria employed by the system must remain consistent with constitutional requirements. The traditional constitutional examination of classification may need to be supplemented by attention to the consequences of automated decision-making. An algorithmically generated classification may appear objective because it is based on numerical calculations. But numerical precision does not automatically establish constitutional validity. The relevant question remains whether the distinction has a legitimate basis and whether the means employed are constitutionally permissible.

Article 14 has also developed beyond a narrow formal conception of classification and has been associated with the constitutional prohibition against arbitrariness. This provides an important basis for examining automated decisions that cannot be adequately explained or justified.

14.2 Article 15

Article 15 prohibits discrimination by the State on specified grounds including religion, race, caste, sex and place of birth, while also permitting specified forms of affirmative action. The provision is significant because it demonstrates that the Constitution recognises both equality and the need to address historical disadvantage. Algorithmic governance should therefore not be evaluated only through the question of whether everyone was subjected to the same procedure. A uniform procedure can produce unequal outcomes where individuals begin from unequal social positions.

14.3 Article 16

Article 16 guarantees equality of opportunity in matters relating to public employment and prohibits discrimination on specified grounds. The increasing use of automated recruitment and assessment systems in public institutions creates a potential area for constitutional scrutiny. Where an automated system influences recruitment, promotion or selection, questions may arise

concerning the criteria used, the representativeness of the data and the possibility of indirect discrimination.

14.4 Article 21

Article 21 has developed into a broad constitutional protection of life and personal liberty. The Supreme Court's privacy jurisprudence has recognised privacy as an aspect of liberty, dignity and autonomy. The significance for algorithmic governance is substantial because automated systems may involve extensive collection, processing and inference from personal information. A system may not merely store information supplied by an individual. It may generate new conclusions about that person. This raises the issue of inferential privacy.

For example, a system may use apparently ordinary information to infer financial vulnerability, preferences, behavioural patterns or other characteristics.

The individual may never have directly disclosed such an inference.

Consequently, privacy in the algorithmic age cannot be reduced to secrecy. It also concerns control, autonomy, dignity and the ability of individuals to understand and challenge decisions based on information about them.

15. RIGHT TO PRIVACY AND ALGORITHMIC PROFILING

The Supreme Court's decision in Justice K.S. Puttaswamy (Retd.) v. Union of India established privacy as a constitutionally protected right grounded in liberty, dignity and autonomy. This development is important for algorithmic decision-making because algorithms frequently operate through personal data.

Data protection and equality are closely connected but are not identical.

A system may process personal data lawfully while still producing discriminatory outcomes. Conversely, an individual may suffer a privacy violation even where there is no discriminatory outcome. Therefore, a comprehensive legal approach must address both dimensions. The problem of profiling illustrates this distinction. Profiling involves analysing personal information in order to evaluate or predict characteristics, preferences or behaviour. The concern is not only the information that an individual knowingly provides but also the conclusions generated from that information.

The greater the capacity of an institution to generate predictions about individuals, the greater the need for accountability concerning how those predictions are made and used.

16. INDIRECT DISCRIMINATION AND ALGORITHMIC DECISIONS

The doctrine of indirect discrimination is particularly relevant to the subject.

Direct discrimination is generally easier to identify because the discriminatory distinction is visible.

Indirect discrimination may arise where a criterion is apparently neutral but places a particular group at a disadvantage. The Supreme Court's reasoning in *Lt. Col. Nitisha v. Union of India* is important because it recognises that discrimination must sometimes be examined through the consequences of a rule within a broader social context. This approach is highly relevant to algorithms. Suppose an employment algorithm gives greater weight to uninterrupted employment history. The criterion appears neutral. However, if women have historically experienced career interruptions because of unequal care responsibilities, the criterion may disadvantage women disproportionately.

Similarly, if a financial model gives significant weight to formal credit history, economically marginalised persons may be disadvantaged because they have historically had less access to formal financial institutions. The algorithm itself may contain no explicit discriminatory instruction. The inequality may arise from the interaction between a neutral criterion and an unequal social structure. This is precisely why algorithmic discrimination requires a substantive equality approach.

17. THE PROBLEM OF ALGORITHMIC OPACITY

One of the principal challenges is opacity. An individual may receive a decision without knowing how the decision was reached. This creates at least three problems.

First, the individual may be unable to understand the basis of the decision.

Second, the individual may be unable to identify discrimination.

Third, a court or regulatory authority may face difficulties in examining the decision.

The problem is sometimes described through the idea of the "black box".

However, opacity can arise for different reasons. Some systems are technically complex. Others may be protected by commercial confidentiality. In some cases, the organisation using the system may itself not understand all of the relationships identified by the model.

From a legal perspective, complexity cannot by itself eliminate accountability.

Where an automated decision has serious consequences for an individual's rights or opportunities, some degree of explanation and institutional responsibility is necessary.

18. EXPLAINABILITY AND THE RIGHT TO KNOW

The concept of explainability concerns the ability to provide an understandable account of how an automated system contributed to a decision. Explainability does not necessarily require disclosure of every line of computer code.

A meaningful explanation should enable an affected individual to understand the important factors that influenced the outcome and to challenge an erroneous or discriminatory decision.

This is particularly important in high-impact decisions.

For example, a person denied employment should ordinarily be able to know the relevant basis of the decision. A person denied access to a financial service should have an avenue to contest inaccurate information. A person excluded from a public welfare programme should have a mechanism through which the decision can be reviewed.

Without explanation, the right to challenge a decision becomes largely theoretical.

19. ACCOUNTABILITY: WHO IS RESPONSIBLE?

A central difficulty is determining responsibility.

An algorithmic system may involve several actors:

- * the organisation commissioning the system;
- * the developer;
- * the data provider;
- * the system operator;
- * the institution making the final decision; and
- * third-party service providers.

If discrimination occurs, responsibility may be distributed across this chain.

The affected individual should not be required to reconstruct the entire technological architecture merely to obtain a remedy.

Legal responsibility should therefore focus upon institutions exercising meaningful control over the decision and the consequences arising from its use.

The existence of an automated system should not permit an institution to escape responsibility by claiming that "the computer made the decision".

Automation changes the method of decision-making; it should not eliminate accountability.

20. DATA PROTECTION AND ITS LIMITATIONS

India's Digital Personal Data Protection Act, 2023 establishes a framework concerning the processing of digital personal data. The legislation recognises the individual's interest in protecting personal data while also recognising the need for lawful processing.

The framework is important but should not be treated as a complete answer to algorithmic discrimination.

Data protection primarily addresses questions relating to personal data and its processing.

Algorithmic discrimination involves a broader set of questions:

- * Was the data representative?
- * Was the model discriminatory?
- * Were proxy variables used?
- * Did the system produce disproportionate effects?
- * Was human oversight available?
- * Could the affected person challenge the outcome?
- * Was the decision based upon an appropriate criterion?

A system can therefore comply with data-processing requirements and nevertheless produce socially unequal outcomes.

Data protection and anti-discrimination law should consequently be understood as complementary rather than interchangeable fields.

21. THE EMERGING INDIAN AI GOVERNANCE FRAMEWORK

India's policy discussion concerning responsible AI increasingly emphasises transparency, explainability, accountability, fairness, human oversight and rights protection. The IndiaAI responsible-AI framework identifies transparency, explainability and accountability among the principles relevant to responsible deployment. The Digital Personal Data Protection Rules, 2025 also represent an important stage in the development of India's data-protection framework. However, AI governance remains a developing field.

The principal challenge is converting broad ethical principles into enforceable institutional practices.

A statement that an AI system should be fair is insufficient unless institutions can determine:

1. what fairness means in the particular context;
2. how discriminatory effects will be measured;
3. who will conduct the assessment;
4. what happens when discrimination is identified; and
5. what remedy is available to the affected person.

The legal framework must therefore move from general principles towards operational safeguards.

22. ALGORITHMIC IMPACT ASSESSMENTS

One possible mechanism is an Algorithmic Impact Assessment.

Before deploying a high-impact algorithmic system, an institution could be required to examine:

- * the purpose of the system;
- * categories of data used;
- * groups potentially affected;
- * foreseeable discriminatory effects;
- * privacy risks;
- * accessibility concerns;
- * human oversight mechanisms;
- * procedures for correcting errors; and
- * available grievance mechanisms.

The assessment should not be treated as a one-time document.

Algorithms may change as new data is introduced or the system is modified. Periodic assessment would therefore be necessary. An impact assessment could be particularly important for systems used in employment, public welfare, education, financial services, policing and other areas involving significant consequences for individuals.

23. INDEPENDENT ALGORITHMIC AUDITING

Independent auditing provides another possible safeguard. An algorithm should not be evaluated solely by the organisation that developed or purchased it. Independent assessment can examine whether the system produces materially different outcomes for relevant groups. Auditing should consider both statistical performance and social context.

For example, an algorithm may demonstrate high overall accuracy while performing significantly less accurately for a smaller demographic group. Overall accuracy may therefore conceal unequal performance.

Auditing must consequently be sensitive to distributional effects. The law may also need to protect auditors from commercial or institutional pressure where significant public interests are involved.

24. HUMAN OVERSIGHT

Human oversight remains important in high-impact decisions. The purpose of human oversight should not be merely to place a person somewhere in the decision-making chain.

A meaningful human review requires:

1. access to relevant information;
2. authority to question the automated outcome;
3. sufficient understanding of the system's limitations;
4. power to correct the outcome; and
5. responsibility for the final decision.

Otherwise, human oversight may become a procedural formality.

The presence of a human signature on an automated decision does not necessarily make that decision genuinely human.

25. GRIEVANCE REDRESSAL

A person affected by an automated decision should have an accessible means of challenging it.

A grievance mechanism should provide:

- * notice of the decision;
- * an understandable explanation;
- * an opportunity to submit contrary information;
- * review by a competent person;
- * correction of inaccurate data;
- * reconsideration where appropriate; and
- * access to an external legal remedy.

The importance of grievance mechanisms is particularly high for economically vulnerable individuals. A sophisticated legal remedy that requires expensive litigation may not provide meaningful protection to a person who lacks financial or technological resources.

Access to justice must therefore remain central to algorithmic governance.

26. ALGORITHMIC DISCRIMINATION AND ACCESS TO JUSTICE

Traditional legal systems assume that the affected person can identify the decision-maker and demonstrate how the decision affected them.

Algorithmic systems can complicate each stage.

The individual may not know:

- * whether an algorithm was used;
- * what information was considered;
- * whether the system made the final decision;
- * whether the result was generated by a third-party service;
- * whether other similarly situated individuals were treated differently.

These information asymmetries may make litigation difficult.

One possible response is to reconsider evidentiary burdens in appropriate cases. Where an individual establishes a credible pattern of differential impact and the relevant information concerning the operation of the system lies exclusively with the institution, legal mechanisms could provide appropriate disclosure or evidentiary assistance. Such an approach would

recognise that equality cannot be effectively enforced where the evidence necessary to demonstrate discrimination is inaccessible to the person who has suffered it.

27. TECHNOLOGY AND SUBSTANTIVE EQUALITY

The constitutional principle of substantive equality offers a useful framework for addressing algorithmic discrimination.

Formal equality asks whether the same rule is applied to everyone.

Substantive equality asks whether the rule operates fairly in a society characterised by unequal starting conditions. This distinction is critical. An algorithm can treat every person according to exactly the same mathematical formula and still produce unequal consequences.

For example, a requirement for uninterrupted employment may appear identical for all applicants but may disproportionately disadvantage persons who have experienced structural barriers to continuous employment.

The question should therefore not simply be:

"Was the same algorithm applied to everyone?"

It should also be:

"What social consequences does the algorithm produce, and are those consequences constitutionally justifiable?" This is where sociology and constitutional law intersect.

28. THE ROLE OF SOCIOLOGY IN AI GOVERNANCE

Sociology contributes something that purely technical approaches may overlook. A technically accurate algorithm may still produce socially problematic outcomes. The reason is that accuracy itself is not socially neutral.

Suppose a predictive system accurately reproduces a historical pattern of inequality. From a narrow statistical perspective, the model may be accurate. From the perspective of social justice, however, reproducing the pattern may perpetuate the very inequality that law seeks to eliminate.

Sociological inquiry therefore requires attention to:

- * power;
- * institutions;
- * social hierarchy;

- * historical disadvantage;
- * access to resources;
- * group identity;
- * social exclusion; and
- * consequences of technological change.

The question is not simply whether an algorithm works.

The question is *for whom it works, under what conditions, and with what consequences*.

29. THE DANGER OF TECHNOLOGICAL DETERMINISM

There is a tendency to treat technology as an independent force capable of determining social outcomes.

Such an approach can obscure human and institutional responsibility.

Algorithms are created by people, adopted by organisations and deployed for particular purposes. If an automated system produces discrimination, the appropriate response is therefore not simply to blame "the technology".

Attention must also be directed towards:

- * the institution that adopted the system;
- * the purpose for which it was adopted;
- * the data selected;
- * the criteria chosen;
- * the safeguards implemented; and
- * the remedies made available.

Technology should be treated as part of the social system rather than as an external force acting upon society.

30. RECOMMENDATIONS

30.1 Development of an algorithmic accountability framework

India would benefit from a coherent framework dealing specifically with high-impact automated decision-making. Such a framework should identify systems that require enhanced safeguards because of their potential impact upon fundamental rights.

30.2 Mandatory impact assessments for high-risk systems

Before deployment, institutions should assess foreseeable risks relating to discrimination, privacy, exclusion and accessibility.

30.3 Periodic independent audits

High-impact systems should be periodically examined for discriminatory outcomes, data quality and performance across relevant population groups.

30.4 Explanation of significant decisions

Individuals should receive meaningful and comprehensible explanations where automated systems materially influence decisions affecting their rights or opportunities.

30.5 Human review

High-impact decisions should remain subject to meaningful human review, particularly where the outcome may seriously affect livelihood, education, welfare, liberty or access to essential services.

30.6 Protection against proxy discrimination

Regulation should recognise that discriminatory outcomes can arise through variables that indirectly correspond with protected characteristics.

30.7 Greater representation in technology development

Development teams and oversight institutions should include individuals from diverse social, linguistic, geographical and professional backgrounds.

Diversity cannot by itself guarantee fairness, but exclusion of diverse perspectives can make certain forms of harm more difficult to identify.

30.8 Accessible grievance mechanisms

Persons affected by automated decisions should have inexpensive and accessible mechanisms for seeking review and correction.

30.9 Greater transparency in public-sector systems

Where automated systems are used by public authorities, information concerning their purpose, legal basis, significant decision criteria and accountability mechanisms should be disclosed subject to legitimate security and privacy considerations.

30.10 Digital literacy

Individuals should be made aware that automated systems may influence decisions affecting employment, finance, education, welfare and other areas of life. Digital literacy should therefore include not only the ability to use technology but also the ability to understand and question technological decision-making.

31. LIMITATIONS OF THE STUDY

The present study is primarily doctrinal and conceptual. It does not provide an empirical measurement of algorithmic discrimination across Indian institutions.

India also lacks comprehensive publicly accessible datasets that would permit systematic comparison of algorithmic outcomes across caste, gender, disability, language and socioeconomic groups. This limitation itself demonstrates the difficulty of studying algorithmic inequality. Future empirical research should examine actual algorithmic systems used in India and investigate their outcomes across different social groups. Interdisciplinary collaboration between lawyers, sociologists, computer scientists, statisticians and policymakers would be particularly valuable.

32. CONCLUSION

Algorithmic decision-making represents an important transformation in the relationship between individuals, institutions and the State. Its significance extends beyond technological innovation because automated systems increasingly influence access to opportunities, resources and public services.

The central concern is not that algorithms are inherently discriminatory. The deeper concern is that algorithms operate within societies that are themselves unequal.

Historical discrimination can enter datasets. Institutional preferences can become computational criteria. Social inequalities can appear as statistical patterns. Neutral variables can function as proxies for social characteristics. Decisions that were once visibly made by human beings can become embedded within technological systems and consequently more difficult to examine.

Indian constitutional law already provides important principles through equality, non-discrimination, dignity, privacy and substantive equality. The jurisprudence concerning indirect discrimination demonstrates that constitutional equality cannot be limited to the examination of

explicit discriminatory intention. The consequences of apparently neutral rules must also be examined within their social context.

This approach is particularly important for algorithmic decision-making.

An algorithm should not receive constitutional legitimacy merely because it applies the same mathematical formula to everyone. Equal treatment requires attention to the social consequences of that formula.

The Digital Personal Data Protection framework provides an important component of India's response to the data-driven economy, but data protection alone cannot address every form of algorithmic discrimination. Privacy, equality, transparency, accountability and access to remedies must operate together.

The future of AI governance in India therefore requires a shift from technological neutrality towards technological accountability. Institutions using automated systems should be able to demonstrate not only that their systems are efficient, but also that they are capable of operating consistently with constitutional values.

The sociological dimension is equally important. If algorithmic systems are allowed to reproduce existing inequalities without scrutiny, technological development may strengthen rather than reduce structural disadvantage. Conversely, if technological systems are designed with awareness of social diversity and subjected to meaningful legal oversight, they can become instruments that improve access and participation.

The central question for Indian law is therefore not whether society should accept or reject algorithmic technology. The more fundamental question is how technological systems can be integrated into a constitutional democracy in which equality, dignity and human autonomy remain meaningful.

Algorithmic governance should ultimately remain accountable to the people affected by it. Technology may assist decision-making, but responsibility for the social consequences of those decisions cannot be transferred entirely to the machine.

REFERENCES

A. Constitutional and Legal Materials

1. Constitution of India, 1950.

2. Digital Personal Data Protection Act, 2023.
3. Information Technology Act, 2000.
4. Digital Personal Data Protection Rules, 2025.
5. Rights of Persons with Disabilities Act, 2016.
6. Transgender Persons (Protection of Rights) Act, 2019.

B. Judicial Decisions

1. Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1.
2. K.S. Puttaswamy (Retd.) v. Union of India, (2018) 3 SCC 797.
3. Anuj Garg v. Hotel Association of India, (2008) 3 SCC 1.
4. Navtej Singh Johar v. Union of India, (2018) 10 SCC 1.
5. Lt. Col. Nitisha v. Union of India, 2021 SCC OnLine SC 261.
6. National Legal Services Authority v. Union of India, (2014) 5 SCC 438.
7. Jeeja Ghosh v. Union of India, (2016) 7 SCC 761.
8. Vikash Kumar v. Union Public Service Commission, 2021 SCC OnLine SC 84.