© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715742

The Digital Dissent: AI, Algorithmic Bias, And Constitutional Guarantees

RAJAT SRIVASTAVA[1], MANEESH A. SRIVASTAVA[2]
[1]*LL.M. 1st Year (1st Semester), Business and Corporate Law, Faculty of Juridical Sciences, Rama University, Kanpur*
[2]*Assistant Professor, Faculty of Juridical Sciences, Rama University, Kanpur*

Abstract- ***"Technology can be a great enabler but it can also be a great threat to privacy unless we ensure that our constitutional guarantees are not rendered illusory in the digital age"***
– JUSTICE D.Y. CHANDRACHUD

The escalating integration of Artificial Intelligence (AI) and advanced technology into governance and society poses profound, immediate threats to established constitutional rights. This research paper offers a critical analysis of the constitutional challenges arising from AI's deployment across three critical areas: Equality, Privacy, and Due Process. Regarding Equality and Non-Discrimination, the paper examines how opaque AI algorithms, often trained on biased datasets, perpetuate systemic prejudice in high-stakes decisions like judicial sentencing and resource allocation. This lack of algorithmic transparency fundamentally undermines the constitutional guarantee of equal protection, necessitating the recognition of algorithmic bias as a direct form of constitutional harm. The proliferation of AI-driven mass surveillance technologies, such as facial recognition and predictive policing, directly assaults the Right to Digital Privacy. Current constitutional safeguards, designed for an earlier technological era, prove inadequate against ubiquitous, constant data collection and analysis. The research advocates for the evolution of the right to privacy into a robust 'Right to Informational Self-Determination,' placing strict constitutional limits on state and corporate data exploitation. Finally, the paper addresses Due Process concerns stemming from autonomous, unreviewable decision-making. Ensuring a fair hearing and the Right to an Effective Remedy becomes impossible when fundamental rights are affected by AI systems whose internal reasoning is incomprehensible. Ultimately, reconciling AI's transformative power with enduring constitutional values requires proactive judicial innovation and legislative reform to establish clear, enforceable standards for technologically informed constitutional governance. This reconciliation is essential for preserving the rule of law in the digital age.

Keywords- ***Technology, Artificial Intelligence (AI), Constitutional Rights, Equality (or NonDiscrimination), Privacy (or Digital Privacy), Due Process, Equal Protection, Algorithmic Transparency***

I. INTRODUCTION: THE CONSTITUTIONAL CROSSROADS

1.1. The Promise and Peril of AI: A New Constitutional Moment

We are witnessing what can be termed a "constitutional moment" – a period where the fundamental relationship between the citizen and the state is being radically redefined by technology. The Indian Constitution, a majestic parchment conceived in the mid-20th century, now faces its most formidable challenge from the rise of 21st-century Artificial Intelligence. The government's ambitious 'Digital India' initiative underscores a wholehearted embrace of AI for governance, or 'AI for Good'. This includes using AI for optimizing agricultural supply chains, diagnosing diseases in remote areas, personalizing education, and streamlining the delivery of public services. This "great enabler" potential of technology, as acknowledged by Justice Chandrachud, promises efficiency, objectivity, and scalability. Proponents argue that AI can eliminate human bias, reduce bureaucratic red tape, and usher in an era of data-driven, rational governance.

However, embedded within this promise is a profound peril. The same systems that can efficiently allocate resources can also systematically exclude the most vulnerable. The tools that can predict crime can also become instruments of social control and discrimination. The algorithms that promise objectivity can, in fact, codify and amplify the deepest biases of our society. This paper argues that the

unregulated and opaque deployment of AI poses an existential threat to the foundational principles of the Indian Constitution: Liberty, Equality, and Fraternity. The "digital dissent" is not a Luddite rejection of progress, but a critical, constitutionalist assertion that technological advancement must be subservient to democratic values and fundamental rights. It is a call to ensure that the digital future is not a dystopia of automated inequality and pervasive surveillance, but a society where technology amplifies, rather than diminishes, human dignity and freedom.

1.2. Defining the Terrain: Artificial Intelligence and Machine Learning

To legally regulate a phenomenon, one must first understand it. At its core, AI refers to the ability of machines to perform tasks that typically require human intelligence, such as visual perception, speech recognition, and decision-making. The most prevalent and legally significant form of AI today is Machine Learning (ML). Unlike traditional software that follows explicit, pre-programmed rules, ML systems learn patterns and correlations from vast amounts of data. They are trained on "training datasets," and their performance is heavily dependent on the quality, quantity, and representativeness of this data. This data-centric nature is the primary source of both its power and its peril. When we speak of algorithmic decisionmaking in governance, we are largely referring to the outputs of such ML models. A critical sub-type is deep learning, which uses complex neural networks to make sense of data. While incredibly powerful, these networks are often "black boxes," meaning their internal decisionmaking processes are opaque and difficult for even their creators to interpret. This inherent opacity lies at the heart of the constitutional due process crisis.

1.3. The Indian Context: Digital India and the Regulatory Vacuum

India is at a critical juncture. On one hand, it is rapidly deploying AI systems across various domains. The National Crime Records Bureau (NCRB) is implementing a centralized Facial Recognition System (FRS). The Income Tax Department uses AI for risk profiling and scrutiny assessment. States like Punjab and Telangana have experimented with predictive policing. The Aadhaar ecosystem, the world's largest biometric ID system, provides the foundational data infrastructure for many of these initiatives. However, this rapid deployment is occurring in a significant regulatory vacuum. The Digital Personal Data Protection Act, 2023, while a step forward, is primarily focused on data processing and lacks specific, robust provisions to tackle algorithmic bias, ensure explainability, or regulate mass surveillance. This legislative lag creates a situation where the state is accumulating unprecedented power through technology, without the corresponding constitutional checks and balances to prevent its abuse.

1.4. Research Questions, Methodology, and Structure of the Paper

This paper is structured around three pivotal research questions:

1. How does the phenomenon of algorithmic bias in AI systems not merely replicate but structurally reinforce discrimination, thereby violating the substantive and procedural guarantees of equality under Article 14?
2. In what ways do AI-driven mass surveillance technologies create a digital panopticon that infringes upon the right to privacy (Puttaswamy) and other associated freedoms, and why must the judiciary evolve the right to privacy into a more robust 'Right to Informational SelfDetermination'?
3. How does the inherent opacity ("black box" nature) of complex AI systems create an insurmountable barrier to fulfilling the core tenets of due process and natural justice under Articles 14 and 21, and what doctrinal shifts are required to address this?

The methodology is primarily doctrinal, employing a critical analysis of Indian constitutional law, landmark judgments, and comparative jurisprudence from the European Union and the United States. It also incorporates insights from computer science ethics, sociology, and political philosophy to build a comprehensive legal critique. The paper is structured to first diagnose the constitutional harm in three key areas and then propose a comprehensive framework for redressal.

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715742

II. ALGORITHMIC BIAS AND THE EROSION OF SUBSTANTIVE EQUALITY (ARTICLE 14)

2.1. The Constitutional Architecture of Equality and Non-Arbitrariness

Article 14 of the Indian Constitution guarantees every person "the equal protection of the laws within the territory of India." The Supreme Court has transformed this provision from a mere formal guarantee into a dynamic and powerful tool for ensuring substantive equality.

“Equality is a dynamic concept with many aspects and dimensions and it cannot be imprisoned within traditional and doctrinaire limits[1], And that “arbitrariness is the antithesis of equality.”[2]
BY- JUSTICE BHAHWATI

This principle was cemented in, establishing that state action must not be arbitrary and must be fair. The doctrine of non-arbitrariness pervades Article 14, meaning that any state action that is capricious, irrational, or without adequate determining principle is constitutionally invalid. This expansive interpretation is crucial because it moves the focus from mere classification to the quality and rationality of the state action itself. Furthermore, the Supreme Court has consistently held that Article 14 embraces the concept of non-discrimination, which is explicitly enumerated in Article 15. Any state action that discriminates on the grounds of race, caste, sex, place of birth, or other analogous grounds would be a clear violation of this constitutional mandate.

2.2. Deconstructing the 'Black Box': The Technical Genesis of Algorithmic Bias

Algorithmic bias is not a minor technical glitch; it is a fundamental structural flaw that arises from the very process of how ML systems are built. The sources are multifarious and often interlinked:

Historical Bias in Training Data: This is the most pernicious source. An AI model is a reflection of the data it is fed. If trained on historical data of police arrests, it will learn the patterns of policing, not necessarily of crime. If the historical data reflects decades of overpolicing in predominantly minority neighborhoods or against certain caste groups, the algorithm will interpret this as those neighborhoods or groups being "high risk," thus perpetuating and automating the cycle of discrimination. The algorithm, in this sense, is a perfect mirror of our past injustices, giving them a veneer of technological objectivity.

Representation Bias: This occurs when the training data is not representative of the population the AI will serve. A seminal example is commercial facial recognition technology. If a system is trained predominantly on light-skinned male faces, it will perform poorly in identifying women and people with darker skin, leading to higher false-positive rates for these groups. This is not just an accuracy issue; it is an equality issue, as it subjects certain populations to a higher risk of wrongful identification and consequent state action.

Proxy Discrimination: This is a particularly insidious form of bias. Even if sensitive attributes like caste, religion, or gender are explicitly removed from the data, algorithms are adept at finding correlated proxies. Postal code can be a proxy for race and socio-economic status; university name can be a proxy for caste (given historical disparities in access); and online purchase history can be a proxy for religious affiliation. The algorithm can thus discriminate without ever being explicitly told the protected category, making it difficult to detect and prove.

Aggregation Bias: This occurs when a one-size-fits-all model is applied to diverse populations, failing to account for regional, cultural, or socio-economic differences. A creditscoring algorithm designed for urban, salaried populations may unfairly penalize rural applicants whose financial behavior (e.g., seasonal income, informal credit) follows different patterns. This imposes a metropolitan standard on diverse contexts, leading to systematic exclusion.

2.3. Panorama of Prejudice: Case Studies in Algorithmic Discrimination

2.3.1. Predictive Policing and the Feedback Loop of Injustice: In India, states like Punjab, Delhi, and Telangana have experimented with predictive policing systems like CMAPS and PAS. These systems use

historical crime data to generate "heat maps" of likely future crime. The fundamental flaw is the feedback loop, a concept termed "runaway feedback loops" by researcher Danielle Keats Citron. Increased police presence in a "high-risk" area leads to more arrests (as more police find more crime, even if crime rates are static), which feeds back into the system as "confirmed" crime data, justifying even more policing. This creates a selffulfilling prophecy that systematically targets and criminalizes certain communities. This violates Article 14 by creating an irrational and arbitrary classification based on biased data. It is the digital equivalent of the concept of "profiling," which the judiciary has often viewed with suspicion for its potential to violate equality.

2.3.2. Judicial Risk Assessment Tools: Automating Injustice from the Bench: While India has not formally adopted tools like the COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) used in the United States, the allure of "data-driven judiciary" is growing. COMPAS, which assesses a defendant's likelihood of reoffending, was found by ProPublica to be racially biased. The constitutional harm here is twofold. First, it violates substantive equality by disproportionately labeling Black defendants as "high risk." Second, and more fundamentally, it violates the principle of individualized justice, a cornerstone of criminal jurisprudence. Sentencing a person based on the statistical behavior of a group they belong to is the very definition of arbitrariness prohibited by Article 14. It reduces the individual to a data point, negating their unique circumstances and the court's duty to deliver personalized justice, as emphasized in cases like Bachan Singh v. State of Punjab[3].

2.3.3. The Aadhaar Ecosystem: Exclusion by Design in Welfare Distribution: The Aadhaar-based biometric authentication for welfare schemes like the Public Distribution System (PDS) and the Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA) is a form of automated decision-making. While not a complex AI, its logic is rigidly algorithmic. The "algorithm" is simple: no biometric match = no benefits. This has led to widespread "exclusion errors," where genuine beneficiaries—often the elderly with faded fingerprints, manual laborers with worn-out fingerprints, or those in humid conditions that affect scanner accuracy—are denied essential rations or wages. This is a clear violation of the right to equality, as it arbitrarily excludes the most vulnerable from the social safety net, a benefit they are legally entitled to under social welfare legislation. The system's design fails the test of proportionality and non-arbitrariness, as it prioritizes the prevention of fraud (a legitimate aim) over the fundamental right to life and food of millions, without providing adequate, accessible, and reliable alternatives for authentication failures.

2.3.4. Private Sector Bias: Hiring Algorithms and Credit Scoring: The constitutional mandate of Article 14 applies directly to the State. However, under Article 15(2), the State is empowered to make laws prohibiting discrimination by private individuals. The pervasive use of AI in the private sector creates a new frontier for discrimination. AI-powered hiring tools that filter resumes based on patterns from past successful hires can learn to discriminate against women (e.g., penalizing gaps in employment that may be due to childcare) or candidates from less prestigious colleges, which can have a caste and class dimension. Similarly, AIdriven credit scoring can create a "digital redlining" effect, denying loans to people from specific neighborhoods or with certain consumption patterns. This necessitates legislative action under Article 15(2) to enforce constitutional values in the private sphere and prevent the emergence of a new, digitally encoded caste and class system.

2.4. The Legal-Philosophical Underpinnings of Discrimination: From Intent to Effect

A critical legal shift is required to address algorithmic bias. Traditional anti-discrimination law often requires proving discriminatory intent (mens rea). However, algorithmic bias is often a product of negligence, oversight, or complex technical processes, not malicious intent. Therefore, the legal framework must evolve to focus on disparate impact. This doctrine, wellestablished in U.S. employment law, holds that a seemingly neutral policy or practice that has a disproportionately adverse effect on a protected group is discriminatory, regardless of intent. Indian courts have acknowledged this concept in environmental law (absolute liability) and should now import it into the realm of algorithmic governance. The question should not be whether the programmer intended to

discriminate, but whether the algorithm's effect is discriminatory.

2.5. Algorithmic Bias as a Distinct Constitutional Harm: Forging a New Jurisprudence

The judiciary must evolve its understanding of discrimination to formally recognize algorithmic bias as a distinct constitutional harm. Courts should develop a "Strict ScrutinyLite" standard for state algorithms. When a petitioner demonstrates a prima facie case of algorithmic bias (through statistical evidence of disparate impact), the burden should shift to the State to prove:

1. Proportionality: The use of the algorithm is necessary and the least rights-restrictive means to achieve a legitimate state aim.
2. Fairness: The algorithm has been rigorously audited for bias, both pre- and postdeployment, using standardized metrics, and the results are publicly available in an anonymized form.
3. Transparency: A meaningful, understandable explanation for any adverse decision can be provided to the affected individual.
4. Accountability: A clear line of accountability is established for the outcomes of the algorithmic system, with designated officials responsible for its functioning and redressal.

2.6. Comparative Perspectives: The EU AI Act and U.S. Algorithmic Accountability Acts Learning from global developments is crucial. The European Union's pioneering AI Act adopts a risk-based approach. It classifies AI systems by risk, and "high-risk" AI systems, including those used in critical infrastructure, employment, essential private and public services, law enforcement, and administration of justice, are subject to strict obligations. These include risk assessment, high-quality data governance, technical documentation, human oversight, and accuracy and robustness requirements. In the United States, proposed laws like the Algorithmic Accountability Act seek to mandate impact assessments for automated systems used by large corporations. India can draw from these frameworks to create a uniquely Indian model that is sensitive to its specific social cleavages (like caste) and constitutional ethos, perhaps by explicitly listing "caste" and "religion" as protected categories in any algorithmic accountability legislation.

III. THE SURVEILLANCE PANOPTICON AND THE RIGHT TO DIGITAL PRIVACY (ARTICLE 21)

3.1. The Jurisprudential March: From Kharak Singh to Puttaswamy

The right to privacy in India has traversed a long and arduous path. In Kharak Singh v. State of Uttar Pradesh (1962)[4], the Supreme Court, while invalidating nocturnal surveillance, narrowly interpreted privacy and failed to recognize it as a standalone fundamental right.

"right to be let alone," was a voice of foresight.[5]
BY- Justice Subba Rao,

This vision was finally realized in the historic Justice K.S. Puttaswamy (Retd.) v. Union of India (2017)[6] judgment.

The nine-judge bench unanimously declared that privacy is an intrinsic part of the right to life and personal liberty under Article 21. The Court eloquently described it as the "ultimate expression of the sanctity of the individual" and a necessary precondition for the exercise of other freedoms. It noted that privacy includes at its core: (a) bodily autonomy; (b) informational privacy; and (c) the right to make intimate personal choices.

Crucially, the Court laid down a three-pronged test for any state action infringing privacy: it must be (i) prescribed by law; (ii) in pursuit of a legitimate state interest; and (iii) proportionate. This test is the primary legal shield against state surveillance.

3.2. The Anatomy of the Digital Panopticon: FRT, Social Media Scraping, and IOT

AI has enabled a form of surveillance that is qualitatively different from anything before. It is mass, indiscriminate, persistent, and predictive, creating a modern-day digital panopticon where citizens never know if they are being watched, and thus must assume they always are.

Facial Recognition Technology (FRT): Deployed in airports, railway stations, and public spaces in cities like Delhi and Hyderabad, FRT enables persistent tracking. It creates a permanent record of an individual's movements, associations, and activities. The danger is compounded by its documented inaccuracy; studies like the one by the National Institute of Standards and Technology (NIST) confirm that FRT has higher false positive rates for women and people of color, leading to a higher risk of wrongful identification and detention for these groups, adding an equality dimension to the privacy violation.

Social Media Monitoring and Scraping: Tools like the Delhi Police's "DTN" or the NCRB's proposed "Cyber Crime Prevention against Women and Children" portal involve "scraping" social media platforms to analyze posts, networks, and sentiments to identify "potential threats" or "radicalization." This subjects freedom of speech and association to a chilling, pre-emptive scrutiny, turning online expression into a source of perpetual risk assessment.

Internet of Things (IOT) and Data Fusion: Smart city infrastructure, combined with data from mobile phones, creates a comprehensive digital footprint of every citizen. When fused and analyzed by AI, this data can reveal intimate details about a person's health (from fitness trackers), political views (from browsing history), sexual orientation (from app usage), and religious beliefs (from location data near places of worship). This creates a "digital twin" of the citizen that the state can monitor and manipulate.

3.3. The Chilling Effect: How Mass Surveillance Erodes Democracy and Fundamental Freedoms (Article 19)

The most profound harm of mass surveillance is the "chilling effect" it has on fundamental freedoms guaranteed under Article 19. When citizens know they are being constantly watched, they are less likely to express dissenting opinions, associate with controversial groups, or participate in protests for fear of being flagged, harassed, or worse.

This self-censorship strikes at the very heart of a vibrant democracy. As the Supreme Court noted in S. Rangarajan v. P. Jagjivan Ram (1989)[7], freedom of speech is essential for a "mature society." Mass surveillance fosters immaturity, conformity, and fear, not the dissent, debate, and pluralism necessary for democratic evolution. The right to assemble peacefully is rendered meaningless if every participant knows their face is being logged and their social networks mapped by the state. The constitutional rights under Articles 19(1)(a), (b), (c), and (d) are thus indirectly but effectively nullified by the architecture of surveillance.

3.4. The Inadequacy of Puttaswamy's Tripartite Test in the AI Era

While the Puttaswamy test is a powerful tool, its application to AI surveillance is fraught with challenges that often render it ineffective.
Legality: The "law" authorizing mass surveillance is often vague and outdated. The broad powers under the Indian Telegraph Act, 1885, or the procedural rules under the IT Act, 2000, are not "law" in the Puttaswamy sense, which requires precision, foreseeability, and safeguards against abuse. They provide executive discretion without legislative guidance, failing the test of legality.

Legitimate Aim: While "national security" and "public order" are legitimate aims, they are often used as a blanket, non-justiciable justification for indiscriminate surveillance of the entire population. The state must be required to demonstrate a specific and imminent threat that justifies such a wide-ranging intrusion, which it consistently fails to do.

Proportionality: The current deployment of FRT and social media scraping fails the proportionality test spectacularly. The marginal, speculative gains in security are vastly outweighed by the colossal invasion of privacy and the chilling of fundamental freedoms. There is no demonstration of necessity, and less intrusive alternatives (e.g., targeted surveillance based on specific, credible intelligence) are not adequately considered or implemented. The indiscriminate nature of the surveillance makes it inherently disproportionate.

3.5. The Path Forward: Elevating Informational Self-Determination to a Fundamental Right

To meet this challenge, the right to privacy must be concretized into the Right to Informational Self-Determination. This German constitutional concept, which the Puttaswamy judgment itself referenced approvingly, holds that an individual has the right to decide for themselves when and to what extent their personal data is disclosed and used. It is based on the idea that control over one's data is essential for individual autonomy in the digital age.

This would entail the following constitutional imperatives:
Data Minimization: The state should only collect data that is strictly necessary for a specific, lawful purpose. The collection of data "just in case" or for vague future purposes must be prohibited.

Purpose Limitation: Data collected for one purpose (e.g., issuing a driver's license) cannot be repurposed for another (e.g., FRT surveillance or data analytics) without fresh, specific, and informed consent.

Strong Rights for Data Subjects: This includes the right to access, correct, and, most importantly, the right to erasure (right to be forgotten), allowing individuals to control their digital footprint over time and have outdated or irrelevant data deleted.

The judiciary must explicitly recognize this right as a constitutional imperative, forcing the state to redesign its data-intensive AI systems around the principle of individual autonomy, not state control. Any law that enables mass surveillance would have to be tested against this robust standard.

3.6. Beyond the State: Corporate Data Exploitation and the Need for Horizontal Application

The threat to privacy is not only vertical (state-citizen) but also horizontal (corporation-citizen). Tech companies like Google and Meta engage in pervasive data collection for targeted advertising, creating detailed profiles of users. This corporate surveillance is often the foundation for state surveillance, as seen in data brokerage markets. While the Digital Personal Data Protection Act, 2023, attempts to regulate this, a constitutional perspective is needed. The Supreme Court should consider the horizontal application of fundamental rights, as done in other jurisdictions, to ensure that private entities are also bound by the principles of informational self-determination, especially when they perform public functions or hold data that can be exploited to infringe upon citizens' rights.

IV. AUTOMATED ADJUDICATION AND THE PROCEDURAL JUSTICE DEFICIT (ARTICLES 14 & 21)

4.1. The Golden Triangle: Due Process, Natural Justice, and the Rule of Law

Articles 14 and 21 together form the bedrock of procedural justice in India, creating a "golden triangle" of fundamental rights. The Maneka Gandhi case fused the two, holding that any procedure depriving a person of life or liberty must be "right, just and fair," and not "arbitrary, fanciful or oppressive." The principles of natural justice—audi alteram partem (hear the other side) and nemo judex in causa sua (no one shall be a judge in his own cause)—are integral to this guarantee. The right to a fair hearing implies the right to know the evidence against one, to have a reasonable opportunity to respond, and to be judged by an impartial arbiter. These principles are not mere technicalities; they are the essence of a government of laws and not of men.

4.2. The 'Black Box' Problem and the Evisceration of Audi Alteram Partem

The core of the due process crisis is the "black box" nature of complex AI models like deep neural networks. When an AI system denies a visa, rejects a loan application, flags a citizen for enhanced tax scrutiny, or recommends denying bail, the decision-making process is opaque and non-intuitive.

This eviscerates the right to a fair hearing. If an individual cannot understand the reasons for an adverse decision, how can they possibly contest it?
A notice that states, "Your application was rejected by an algorithm based on our internal risk model," is a constitutional nullity. It violates the very essence of audi alteram partem. The individual is rendered a passive subject of an unaccountable process, unable to

challenge the facts, logic, or relevance of the factors that led to the decision. For instance, if a predictive algorithm denies bail based on a defendant's zip code and social media friends, the defendant has a right to challenge the accuracy of that data and the rationality of using it as a proxy for flight risk. Opacity makes this impossible.

4.3. The Right to an Effective Remedy: When the Judge is an Incomprehensible Machine

Closely tied is the right to an effective remedy before a neutral forum, a right guaranteed under Article 32 and inherent in the rule of law. Judicial review is a cornerstone of our constitutional scheme. However, how is a judge to review the legality or rationality of an AI's decision if its reasoning is a mystery? The judge cannot perform their constitutional function of being a check on executive arbitrariness if the basis of the impugned decision is locked inside a black box. Is the judge to simply trust the machine? This would be an abdication of the judicial role. This disempowers the judiciary and nullifies the right to an effective remedy, creating an "algorithmic black hole" where state decisions are immune from meaningful scrutiny.

4.4. The Doctrine of Proportionality as a Tool for Scrutinizing State AI

The well-established doctrine of proportionality, used by Indian courts in cases like KS Puttaswamy and Modern Dental College v. State of MP[8], provides a robust framework for evaluating state AI. The use of an opaque AI system in administrative decision-making affecting rights would likely fail at multiple steps of the proportionality analysis:

1. Legitimate Goal: While the goal (efficient tax collection, crime prevention) may be legitimate...
2. Suitability: ...an opaque system is not a suitable means to a fair end, as fairness, by definition, requires explainability and the ability to be contested.
3. Necessity: ...a less rights-restrictive alternative (e.g., a transparent system, a system with detailed rules, or a system with a human final decision-maker) is almost always available. The state's argument of "efficiency" cannot trump fundamental due process.
4. Proportionality Stricto Sensu (Balancing): ...the severe curtailment of the right to a fair hearing and an effective remedy is grossly disproportionate to the marginal gains in administrative efficiency. The harm to the constitutional right far outweighs the benefit to the state.

4.5. Re-engineering Due Process: Explainable AI (XAI) and the 'Humanin-the-Loop' Mandate

To preserve due process, the law must mandate two key requirements:

1. The Right to a Meaningful Explanation: This is a positive obligation on the state. It is not a right to the source code, which is often a trade secret, but to a reason that is comprehensible, specific, and actionable. For instance, "Your loan was denied due to a combination of factors, the most significant being: your high debt-to-income ratio (70%) and two recent late payments on your credit card in the last six months." Techniques from the field of Explainable AI (XAI), such as LIME (Local Interpretable Model-agnostic Explanations) or SHAP (Shapley Additive Explanations), can generate such post-hoc explanations. This right must be read into Articles 14 and 21 by the judiciary.
2. The Human-in-the-Loop Mandate: For high-stakes decisions that fundamentally alter a person's life or liberty—such as bail, welfare benefits, blacklisting, or deportation—there must be a mandatory review by a human official. This human must have the authority, understanding, and responsibility to override the algorithmic recommendation. The human must be more than a rubber stamp; they must be a meaningful part of the decision-making chain, capable of applying empathy, context, and discretion that algorithms lack. The final decision, and the accountability for it, must rest with a human being.

4.6. The Burden of Proof in Algorithmic Decision-Making: Shifting the Onus to the State

Given the information asymmetry between the citizen and the state regarding the AI system, the traditional burden of proof must be recalibrated. Once a petitioner demonstrates that a decision affecting their rights was made by an AI system, the burden should shift to the state to prove that the process was fair, the explanation provided was adequate, and the decision was not

arbitrary or discriminatory. This would force the state to maintain detailed documentation and audit trails for its algorithmic systems, making judicial review a feasible exercise.

V. SYNTHESIS AND THE PATH FORWARD: A CONSTITUTIONAL FRAMEWORK FOR THE AI AGE

5.1. Recapitulation: The Tripartite Constitutional Crisis

This paper has demonstrated that the unregulated proliferation of AI creates a tripartite constitutional crisis. It erodes substantive equality (Article 14) by automating and scaling discrimination through biased algorithms. It assaults individual privacy (Article 21) and associated freedoms (Article 19) by enabling a architecture of mass surveillance that chills dissent. Finally, it nullifies procedural due process (Articles 14 & 21) by making administrative and quasi-judicial decision-making opaque and unreviewable. Together, these challenges threaten to hollow out the core of the Indian Constitution and undermine the social contract between the citizen and the state.

5.2. A Blueprint for Reform: Judicial, Legislative, and Administrative Actions

A passive approach will lead to the gradual erosion of our rights. A proactive, multilayered strategy is essential to build a constitutional moat against digital tyranny.

5.2.1. Judicial Doctrinal Innovation:

The Judiciary as a Catalyst for Change-

The Supreme Court, as the ultimate guardian of the Constitution, must act as a catalyst. It should issue a landmark judgment, a Puttaswamy for the AI age, that:

Explicitly recognizes algorithmic bias with disparate impact as a violation of Article 14.

Elevates the Right to Informational Self-Determination to a fundamental right under Article 21, incorporating principles of data minimization and purpose limitation.

Reads into Articles 14 and 21 a procedural right to a meaningful explanation for any significant automated decision.

Mandates a proportionate and human-in-the-loop approach for high-stakes AI systems.

Applies a strict standard of scrutiny to all state uses of high-risk AI systems, shifting the burden of proof to the state to justify their use.

5.2.2. Comprehensive Legislative Framework: A Proposed 'Digital India Rights Act'

India needs a comprehensive, rights-centric law that goes beyond the current data protection framework. This 'Digital India Rights Act' should:

Prohibit Algorithmic Discrimination: Make it illegal for the state and regulated private entities to use AI systems that have an unjustified disparate impact on protected groups, with "caste" and "religion" explicitly listed.

Mandate Fundamental Rights Impact Assessments (FRIAs): Require rigorous, publicly disclosed audits before deploying any high-risk AI system in governance. These assessments must evaluate potential impacts on equality, privacy, and due process.

Establish a Right to Explanation and a Human-in-the-Loop for critical decisions affecting rights.

Create a Strong Data Governance Regime based on data minimization, purpose limitation, and rights to access, correction, and erasure.

Impose a Moratorium on Mass Surveillance Technologies like live FRT in public spaces until a constitutional and legislative framework is in place.

5.2.3. Institutional Architecture: An AI Regulation Authority and Auditing Framework

An independent AI Regulation Authority (AIRA) should be established, staffed with legal experts, data scientists, sociologists, and ethicists. Its mandate would be to:

Set technical standards for fairness, accountability, and transparency.

Conduct and certify independent algorithmic audits.

Investigate public complaints of algorithmic harm.

Maintain a public register of high-risk AI systems used by the government.

The Bureau of Indian Standards (BIS) should develop Indian Standards for Algorithmic Auditing (ISAA) to provide a standardized, replicable methodology for testing AI systems for bias, accuracy, and robustness.

5.2.4. Ethical Governance and Capacity Building

The government must adopt and publish binding ethical guidelines for AI development and procurement, based on constitutional morality.
Legal and technical education must be integrated. Law schools need courses on law and technology, and engineering schools need mandatory modules on AI ethics and human rights.
A robust public discourse on the constitutional implications of AI is necessary to create democratic accountability and ensure that citizens are aware of their rights in the digital age.

5.3. Conclusion: Reasserting Constitutional Sovereignty Over Digital Systems

The challenge of AI is not a technological problem to be solved by engineers alone. It is a constitutional problem that demands a legal and democratic solution. The "digital dissent" is a call to arms for lawyers, judges, legislators, and citizens to ensure that the future of our society is shaped by the enduring values of our Constitution, not by the transient logic of algorithms. The choices we make today will determine whether our digital future is one of empowered citizenship or automated subjugation. As we stand at this crossroads, we must choose a path of constitutional resilience. We must build a future where AI serves as a tool for the empowerment of the individual, the strengthening of democracy, and the realization of the constitutional vision of a just, egalitarian, and free society. The sovereignty of the Constitution, and the inalienable rights it guarantees, must remain inviolable in the digital age. The time for a constitutional conversation on AI is not tomorrow; it is today.

FOOTNOTES

[1] *E.P. Royappa v. State of Tamil Nadu* (1974) SC 555

[2] *Maneka Gandhi v. Union of India* (1978) SC 597

[3] Bachan Singh v. State of Punjab (1980) 2 SCC 684

[4] Kharak Singh v. State of Uttar Pradesh (1962) SCR (1) 332

[5] K.S. Puttaswamy (Retd.) v. Union of India (2017) 10 SCC 1

[6] S. Rangarajan v. P. Jagjivan Ram (1989) 2 SCC 574

[7] Modern Dental College v. State of MP (2016) 4 SCC 346

REFERENCES

[1] Cases:

- Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1.
- Maneka Gandhi v. Union of India, (1978) 1 SCC 248.
- E.P. Royappa v. State of Tamil Nadu, (1974) 4 SCC 3.
- Kharak Singh v. State of Uttar Pradesh, AIR 1963 SC 1295.
- S. Rangarajan v. P. Jagjivan Ram, (1989) 2 SCC 574.
- Modern Dental College v. State of MP, (2016) 7 SCC 353.
- Bachan Singh v. State of Punjab, (1980) 2 SCC 684.

[2] Legislation and Reports:

- The Constitution of India.
- The Information Technology Act, 2000.
- The Digital Personal Data Protection Act, 2023.
- European Union Artificial Intelligence Act, 2021.
- Report of the Committee of Experts on a Data Protection Framework for India (Chair:
- Justice B.N. Srikrishna), 2018.
- "The Age of Surveillance Capitalism" - Shoshana Zuboff, PublicAffairs (2019).
- "Weapons of Math Destruction" - Cathy O'Neil, Crown Publishing Group (2016).

[3] Scholarly Articles:

- Barocas, Solon, & Selbst, Andrew D. "Big Data's Disparate Impact." California Law Review, Vol. 104, p. 671 (2016).
- Citron, Danielle Keats. "Technological Due Process." Washington University Law Review, Vol. 85, p. 1249 (2008).

- Pasquale, Frank. "The Black Box Society: The Secret Algorithms That Control Money and Information." Harvard University Press (2015).
- Satpathy, Jyotsna. "Algorithmic State: The Future of Governance and the Challenge to Constitutional Rights." Indian Law Review, Vol. 6, Issue 2 (2022).
- Singh, M.P. "The Constitution and Technology: A Jurisprudential Perspective." Journal of the Indian Law Institute, Vol. 62, No. 1 (2020).
- Wachter, S., Mittelstadt, B., & Floridi, L. "Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation." International Data Privacy Law, Vol. 7, Issue 2 (2017).
- Shivnarayan, R. & Desai, D. "The Invisible Caste: Algorithmic Bias and Social Inequality in India." NUJS Law Review, Vol. 14, Issue 3 (2021).