THE ALGORITHMIC VICTIM: FACIAL RECOGNITION TECHNOLOGY, FALSE POSITIVES, AND THE ABSENCE OF A VICTIM-CENTERED LEGAL FRAMEWORK IN INDIA

Vanshaj Sharma, The National Law University of Meghalaya

ABSTRACT

Facial recognition technology has become a routine instrument of law enforcement across India, deployed by state police forces and proposed for integration into a National Automated Facial Recognition System. Yet the legal architecture governing these deployments remains profoundly underdeveloped. This article advances a specific and original argument: persons who suffer wrongful arrest, detention, or reputational injury as a consequence of a false algorithmic identification are not merely victims of administrative error. They constitute a distinct legal category, the algorithmic victim, whose harms are systematically invisible under existing Indian law.

Drawing on victimological theory, constitutional jurisprudence, and comparative law, the article demonstrates that the absence of recognition produces a remedial vacuum. No statutory provision requires notice before deployment, mandates human verification, or confers compensation rights upon the misidentified. Constitutional remedies, while theoretically available under Articles 14 and 21, are practically inaccessible to those who cannot identify the algorithmic cause of their arrest. The article examines how the European Union's AI Act, the United Kingdom's Bridges litigation, and documented American wrongful arrests have begun to forge accountability frameworks, and concludes by proposing a victim-centered statutory model for India, one grounded in the right to algorithmic due process and in the state's positive obligation to prevent foreseeable harm.

Keywords: Facial Recognition Technology; Algorithmic Victim; False Positives; Article 21; Algorithmic Due Process; NAFRS; Victimology

I. INTRODUCTION

On a January afternoon in 2020, Robert Williams was arrested in his driveway in Farmington Hills, Michigan, in front of his wife and daughters. The warrant had been issued on the basis of a facial recognition match that Detroit Police investigators had themselves noted was uncertain. One examiner had written No, it's not on a comparison sheet before the arrest proceeded regardless.[1] Three years later, Porcha Woodruff, eight months pregnant, was arrested on robbery and carjacking charges that rested on an algorithmically generated identification she had never been given the opportunity to contest.[2] Both cases share the same essential structure: a computational system assigned a criminal identity to a person who bore none; the state acted on that assignment without independent verification; and the human being at the end of the chain suffered consequences, arrest, detention, lost liberty, and psychological injury that no existing legal framework was designed to address.

India is accelerating towards the same structural conditions. Delhi and Hyderabad police forces have deployed facial recognition at public events, railway stations, and protest sites.[3] The Ministry of Home Affairs issued a Request for Proposal in 2019 for a National Automated Facial Recognition System (NAFRS) intended to integrate national criminal databases with real-time surveillance infrastructure.[4] The system, if fully operationalized, would make algorithmic identification a standard, perhaps primary mechanism of criminal suspicion across the country. And yet there is no Indian statute, no binding regulation, and no established judicial doctrine that addresses the position of a person who is wrongly identified.

This article argues that this absence is not incidental. It reflects a conceptual failure: Indian legal discourse on facial recognition has been dominated by a surveillance-and-privacy frame that asks what the state may do to citizens, rather than a victim-centered frame that asks what the law owes those whom algorithmic systems have already harmed. The two frames are not equivalent. Privacy litigation seeks to constrain future deployment; victim-centered law

[1]Kashmir Hill, Wrongfully Accused by an Algorithm, N.Y. Times (June 24, 2020), https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html.
[2]George Joseph, A Pregnant Black Woman Was Wrongfully Arrested After Facial Recognition Misidentified Her, The Guardian (Aug. 7, 2023), https://www.theguardian.com/technology/2023/aug/07/facial-recognition-false-arrest-pregnant-black-woman-michigan.
[3]National Crime Records Bureau, Prison Statistics India 2022, at 14 (2023), https://ncrb.gov.in/en/Prison-Statistics-India-2022.
[4]Roli Srivastava, India Speeds Up Plans for Nationwide Facial Recognition System, Reuters (Mar. 6, 2020), https://www.reuters.com/article/us-india-surveillance-technology/india-speeds-up-plans-for-nationwide-facial-recognition-system-idUSKBN20T19E.

addresses present and past injury. The algorithmic victim requires the latter.

The central thesis is that persons wrongfully identified by facial recognition systems constitute a distinct legal category warranting recognition, remediation, and rights. In this category, the algorithmic victim is distinguished by the opacity of the causal mechanism causing their injury, the asymmetric power relationship between state technology and individual body, the compounding nature of reputational and psychological harm, and the near-total absence of current legal recourse. Developing a coherent framework for this category is not merely a policy preference; given the constitutional commitments of Articles 14 and 21, it is a legal obligation.

II. Facial Recognition Technology and False Positives

A. How Facial Recognition Systems Operate

Facial recognition technology converts a photographic or video image of a human face into a numerical template, a geometric encoding of facial landmarks, which is then compared against a database of stored templates. The comparison yields a similarity score, and a match is flagged when that score exceeds a system-defined threshold. The process appears mechanically objective, which is precisely the source of its authority and its danger. The appearance of mathematical precision lends the output a credibility that verbal identification by a witness or officer would not automatically receive, even though the underlying processes are at least as susceptible to error.

The key variables determining accuracy are the quality of the probe image, which in law enforcement contexts is often low-resolution CCTV footage, the composition of the gallery database against which the probe is matched, and the threshold set by the system administrator. Threshold-setting involves a fundamental trade-off: lowering the threshold reduces false negatives (missed matches) but increases false positives (incorrect identifications). In the law enforcement context, the relevant threshold is set not by courts, not by legal standards of proof, and not by the person whose biometric data is being processed; it is set by technology vendors and police administrators applying operational, rather than rights-based, criteria.

B. Sources of Error and Demographic Bias

The National Institute of Standards and Technology (NIST) conducted systematic evaluations

of facial recognition algorithms through its Face Recognition Vendor Test (FRVT) Program and found that false positive rates were significantly higher for Black and Asian faces than for white faces, and higher for women than for men.[5] In some algorithms, false positive rates for Black women were ninety to one hundred times higher than for white men.[6] Independent research by Buolamwini and Gebru reached convergent findings: commercial gender classification systems showed error rates of up to 34.7% for darker-skinned women compared to 0.8% for lighter-skinned men.[7]

These differentials are not software bugs awaiting correction; they are structural artifacts of how training datasets are composed. Facial recognition systems learn from data, and if that data over-represents certain demographic groups while under-representing others, the resulting model will perform differently across those groups.[8] In the Indian context, where police databases were assembled over decades through processes that systematically over-represented communities targeted by colonial-era criminal-tribe legislation, the risk of compounding historical bias through algorithmic re-inscription is acute and largely unexamined.

The critical legal point is that these errors are not random. They cluster predictably, measurably, and disproportionately among already marginalized populations. A legal system that treats algorithmic misidentification as a neutral technical inconvenience rather than a structured harm is not merely incomplete; it is complicit in perpetuating that structure.

III. The Algorithmic Victim: Reconceptualizing Harm in the Age of AI

A. Defining the Algorithmic Victim

Victimology as a discipline has traditionally defined its subject through the lens of the criminal act: a victim is a person who has suffered harm as a consequence of conduct prohibited by law.[9] Mendelsohn's early typologies, Schafer's functional responsibility model, and Christie's influential account of the 'ideal victim' all presuppose that the causal chain runs from human

[5]Patrick Grother et al., Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects, NIST Interagency Report 8280, at 1–3 (2019), https://doi.org/10.6028/NIST.IR.8280.
[6]Id. at 8–12.
[7]Joy Buolamwini & Timnit Gebru, Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification, 81 Procs. Mach. Learning Research 1, 7–9 (2018).
[8]Grother et al., supra note 5, at 14–17.
[9]Benjamin Mendelsohn, The Victimology, 13 Études Internationales de Psycho-Sociologie Criminelle 23, 25 (1956).

offender to human victim.[10] That presupposition no longer captures the full landscape of legally cognizable injury.[11]

An algorithmic victim is a person who suffers legally cognizable harm, loss of liberty, reputational injury, psychological trauma, or economic loss as a proximate consequence of a false or erroneous identification generated by an automated system deployed by a state or quasi-state actor. The category has four defining characteristics. First, the causal mechanism is opaque the victim typically has no access to the algorithmic process that generated their identification and no means to reconstruct or challenge it. Second, the harm is state-enabled: the injury flows from a decision made by a state actor, a police officer, a prosecutor, a magistrate, in reliance on the algorithmic output, vesting state authority in the error. Third, the harm compounds over time: an arrest, even if ultimately groundless, leaves a trail of records, social stigma, and psychological injury that does not dissolve when the charge is dropped. Fourth, and most significantly, there is no adequate existing remedy: the person harmed does not fit comfortably within the categories that existing tort, criminal procedure, or constitutional law has developed to address state-caused injury.

B. Wrongful Arrest and Loss of Liberty

The most immediate harm suffered by the algorithmic victim is loss of liberty. Robert Williams spent thirty hours in police custody; Porcha Woodruff endured eleven hours of detention while eight months pregnant.[12] In India, a person arrested on the basis of a false positive enters a criminal procedure system in which pre-trial detention can last for extended periods, bail applications require engaged legal representation that many accused cannot access, and the stigma of arrest persists independently of any subsequent acquittal.

This loss of liberty is not attributable to any human malice; there is no officer who deliberately arrested an innocent person knowing them to be innocent. Nor is it the result of a broken legal process; the arrest warrant was issued through formally correct channels. The wrongfulness lies upstream, at the level of the identification itself. Existing criminal procedure law in India, the Code of Criminal Procedure and its successor, the Bharatiya Nagarik Suraksha Sanhita, addresses wrongful arrest through compensation provisions and public-law remedies under

[10]Nils Christie, The Ideal Victim, in From Crime Policy to Victim Policy 17, 18–19 (Ezzat Fattah ed., 1986).
[11]Stephen Schafer, The Victim and His Criminal: A Study in Functional Responsibility 152–53 (1968).
[12]Kashmir Hill, Wrongfully Accused by an Algorithm, supra note 1.

Article 32 and Article 226, but these provisions were designed for the paradigm case of deliberate excess of power. They do not map cleanly onto a case in which the excess of power was mediated by an algorithm that no one involved in the arrest can explain.

C. Psychological and Reputational Harm

Beyond liberty, the algorithmic victim suffers harms that are no less real for being less visible. The psychological literature on wrongful arrest documents significant rates of anxiety, post-traumatic stress, and what researchers term 'identity disruption a fracturing of the person's sense of security in their own social standing.[13] This harm is compounded by the algorithmic context. The victim of an ordinary wrongful arrest can at least contest the identification by confronting an accuser. The algorithmic victim cannot cross-examine a neural network. They cannot challenge the training data. They cannot see the similarity score that condemned them. The inaccessibility of the process that caused their harm is itself a constituent element of the harm.

D. Why Traditional Victimology Is Inadequate

Christie's 'ideal victim' model exposes the difficulty with particular clarity. The ideal victim, in Christie's formulation, is a person of unimpeachable respectability who was engaged in legitimate activity when a clear external agent caused their injury. The algorithmic victim inverts several of these conditions. The causal agent is not a person but a process. The mechanism of harm is not violent but administrative. The victim may have a prior criminal record; indeed, many algorithmic identification systems are more likely to produce false positives for people who have previously been processed by the criminal justice system and whose biometric data is thus already in the relevant databases. The structural consequence is that those least able to invoke legal protection are those most likely to be its subjects.

The inadequacy of existing victimological frameworks is not merely academic. It translates directly into remedial invisibility: if courts and legislators cannot recognize the harm as a distinct category, they cannot design remedies tailored to its specific features. What is required is a conceptual framework that acknowledges algorithmic identification errors as a sui generis form of state-caused harm, structured, foreseeable, and amenable to legal remedy if only the law is prepared to look.

[13]George Joseph, supra note 2; Hill, supra note 1.

IV. India's Legal Framework: Surveillance Without Remedies

A. Constitutional Protections and Their Limits

Indian constitutional jurisprudence provides a theoretically robust foundation for challenging algorithmic misidentification. The Supreme Court's nine-judge bench decision in Justice K.S. Puttaswamy v. Union of India established privacy as a fundamental right under Article 21, encompassing informational autonomy and the right to control one's biometric data.[14] Maneka Gandhi v. Union of India had earlier established that the procedure for depriving a person of personal liberty must be fair, just, and reasonable, a standard that a procedure resting solely on an unverified algorithmic match plainly fails to satisfy.[15]

People's Union for Civil Liberties v. Union of India and Anuradha Bhasin v. Union of India demonstrate the Court's willingness to apply proportionality review to surveillance and communication-restricting measures.[1617] The principle of proportionality, requiring that state action be suitable for its purpose, necessary, and balanced against individual rights, is directly applicable to facial recognition deployment. A system that generates false positives at rates the NIST has documented as substantially higher for certain demographic groups is not a proportionate means of pursuing any law-enforcement objective.

Yet constitutional remedies have structural limitations that make them inadequate as a primary mechanism for protecting algorithmic victims. Under Nilabati Behera v. State of Orissa, the Court has awarded compensation under Article 32 for custodial violations of fundamental rights.[18] But a constitutional writ petition presupposes that the petitioner can identify the specific state action that violated their rights. A person arrested on the basis of a facial recognition match will often not know that facial recognition was used, will not have access to the match report, and will not be able to establish the causal link between the algorithm and their arrest. The opacity that defines the algorithmic identification process is the same opacity that defeats the constitutional remedy.[19]

B. Facial Recognition in India

Despite the constitutional uncertainty, the operational expansion of facial recognition in India

[14]Justice K.S. Puttaswamy v. Union of India, (2017) 10 SCC 1.
[15]Maneka Gandhi v. Union of India, (1978) 1 SCC 248.
[16]People's Union for Civil Liberties v. Union of India, (1997) 1 SCC 301.
[17]Anuradha Bhasin v. Union of India, (2020) 3 SCC 637.
[18]Nilabati Behera v. State of Orissa, (1993) 2 SCC 746.
[19]Puttaswamy, (2017) 10 SCC 1, 310–15 (Chandrachud, J., concurring).

has proceeded without pause. The 2019 NAFRS proposal sought to integrate a database of photographs from passports, driving licenses, voter identity cards, and prison records, creating a unified biometric identification infrastructure accessible to police forces across the country.[20] Delhi Police deployed facial recognition at the 2020 Delhi riots and at railway stations, with reported match accuracy figures that independent technologists have disputed.[21]

What is absent from every known Indian facial recognition deployment is any publicly documented protocol for human verification of algorithmic matches before an arrest is made, any procedure for informing an arrested person that facial recognition was used, any statutory basis for the deployment itself, and any mechanism for auditing false positive rates. The Parliamentary Standing Committee on Home Affairs has noted the absence of accountability mechanisms for automated identification systems, but has not translated that observation into a legislative recommendation.

C. The Remedial Vacuum

The cumulative effect is a remedial vacuum. No notice requirement tells a person that their biometric data has been matched. No explanation right compels disclosure of the algorithmic basis for an identification. No compensation mechanism acknowledges that a false positive is a legally cognizable injury. No mandatory audit requires that deployment agencies assess the accuracy of their systems or publish demographic error rates.[22] NITI Aayog's 2021 document on Responsible AI acknowledges the importance of accountability in AI systems but imposes no binding obligations, offers no enforcement mechanism, and addresses no victims.[23]

V. Comparative Perspectives

A. The European Union

The General Data Protection Regulation (GDPR) was the first binding framework to address automated decision-making at scale. Article 22 of the GDPR prohibits, subject to specific

[20]Ministry of Home Affairs, Request for Proposal: National Automated Facial Recognition System (NAFRS), at 4–6 (2019), https://www.mha.gov.in/sites/default/files/RFP_AFRS_12032019.pdf.
[21]Internet Freedom Foundation, Submission on the National Automated Facial Recognition System (NAFRS), at 2–5 (2020), https://internetfreedom.in/facial-recognition-systems-in-india-part-ii/.
[22]Parliamentary Standing Committee on Home Affairs, 231st Report on Unlawful Activities (Prevention) Amendment Bill, at 18–21 (2019) (noting the absence of accountability mechanisms for automated identification systems).
[23]NITI Aayog, Principles for Responsible AI, at 24–27 (2021), https://www.niti.gov.in/sites/default/files/2021-02/Responsible-AI-22022021.pdf.

exceptions, decisions based solely on automated processing that produce legal or similarly significant effects, and confers a right to human review, explanation, and contest.[24] The EU AI Act, adopted in 2024, goes further, classifying real-time remote biometric identification systems in public spaces as high-risk AI systems subject to mandatory conformity assessments, registration, and, with limited security exceptions, a prohibition on deployment without judicial authorization.[25]

The conceptual shift these instruments represent is instructive. Both the GDPR and the AI Act move beyond a state-power frame, asking what the regulator may prohibit, to a rights frame, asking what the individual may demand. Article 22's right to contest, and the AI Act's transparency and oversight obligations, are recognitions that algorithmic decision-making creates a distinct category of injury that requires a distinct category of remedy. India's legal system has the constitutional vocabulary to make the same move; what it lacks is the statutory instrument.

B. The United Kingdom

Bridges v. South Wales Police represent the most developed common-law judicial treatment of facial recognition to date. The Court of Appeal held that the police force's automated facial recognition deployment was unlawful because it lacked a sufficiently clear legal basis, contained no adequate data protection impact assessment, and gave the police officer directing the deployment excessively broad discretion to decide which faces would be retained on the watchlist.[26] The court's proportionality analysis addressed not only privacy but also the risks of algorithmic error and their differential impact on protected groups, a recognition that false positives are not random misfortunes but structured harms warranting structured legal responses.[27]

C. The United States

The American experience provides the clearest documentation of algorithmic harm as lived reality. Williams, Nijeer Parks, and Woodruff represent a pattern that has now been the subject

[24]Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data, art. 22, 2016 O.J. (L 119) 1 (General Data Protection Regulation).
[25]Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonized Rules on Artificial Intelligence, articles. 5(1)(d), 26, 2024 O.J. (L 1689) 1 (EU AI Act).
[26]Bridges v. South Wales Police [2020] EWCA Civ 1058, 86–88.
[27]Id. ¶¶ 150–52.

of civil litigation, legislative inquiry, and media scrutiny.[28] Woodruff's case, in particular, illustrated the compound nature of algorithmic victimization: she was arrested, detained, separated from her family, and forced to miscarry all consequences that flowed from a single false algorithmic identification that no human officer at any point in the arrest chain verified independently.[29]

Several American municipalities and states have enacted moratoriums or outright prohibitions on law enforcement use of facial recognition, and civil litigation has produced settlements that, while they do not establish binding precedent, confirm that wrongful algorithmic identification produces legally cognizable harm. The trajectory across all three jurisdictions points in the same direction: a growing legal consensus that false-positive identification by state-deployed facial recognition systems is a sui generis harm requiring a tailored legal remedy. India remains a significant outlier.

VI. Towards a Victim- Centered Legal Framework

Any adequate statutory response to algorithmic victimization must operate at three levels: prevention, accountability, and remedy. These levels are not alternatives; a genuine victim-centered framework requires all three, because prevention alone cannot eliminate false positives, accountability without remedy leaves existing harms unremedied, and remedy without prevention simply prices harm rather than preventing it.

At the prevention level, a statutory framework should require that no arrest may be predicated solely on an automated facial recognition match. A mandatory human verification requirement that requires a trained officer to independently confirm identification using non-algorithmic means before any arrest is made is the single most effective structural intervention available. It would not eliminate the error, but it would prevent it from being operationalized before human judgment has had the opportunity to catch it. Complementing this, algorithmic impact assessments should be mandatory before any system is deployed: a structured analysis of false positive rates disaggregated by gender, race, age, and skin tone, conducted by an independent technical assessor and published in the public domain.

[28]Hill, supra note 1; Drew Harwell, Three Black Men Were Falsely Arrested Using Facial Recognition. Detroit Police Just Made Another Arrest, Wash. Post (July 28, 2021), https://www.washingtonpost.com/technology/2021/07/28/facial-recognition-misidentify-black-men.
[29]George Joseph, supra note 2.

At the accountability level, deployment agencies should be required to obtain judicial authorization before establishing any permanent facial recognition watchlist, modeled on the EU AI Act's authorization requirements for real-time remote biometric identification. Independent audits by a statutory oversight body distinct from both the deploying police force and the technology vendor should be conducted at regular intervals and should include the power to suspend a system that fails to meet prescribed accuracy thresholds. The thresholds themselves should be prescribed in law and should be uniform across demographic groups, a system that is accurate for some citizens but unreliable for others does not satisfy the equal protection guarantee of Article 14.[30]

At the remedy level, the framework must create rights that the algorithmic victim can actually exercise. A right to notice informing any person who was the subject of an automated identification, whether or not that identification resulted in arrest, is the foundational requirement. Without notice, there is no trigger for any other right. A right to explanation should accompany notice: the system, the threshold used, the match score, and the database searched should all be disclosed in terms that allow a non-technical person to understand and contest the basis of their identification. Where an identification is shown to have been erroneous and to have caused arrest or detention, statutory compensation should be available as of right, without requiring the claimant to establish negligence or deliberate wrongdoing by any individual officer. The harm flows from the system; liability should attach to the state that deploys it.

Undergirding all of these provisions is a principle that might be called algorithmic due process: the right of any person whose liberty or reputation is placed in jeopardy by an automated state system to understand the basis of the decision affecting them, to challenge it before an independent adjudicator, and to receive a remedy proportionate to the harm suffered. This principle is not foreign to Indian constitutional law; it inheres in the due process content that the Supreme Court has read into Article 21 since Maneka Gandhi. The task is statutory translation: converting the constitutional obligation into enforceable rights specific to the context of algorithmic identification.

VI-A. The Hidden Indian Algorithmic Victim: Invisibility, Opacity, and the Problem of Proof

A recurring objection to the framework proposed in this article is as follows: India has not

[30]Puttaswamy, (2017) 10 SCC 1, ¶ 180 (Nariman, J., concurring); Nilabati Behera, (1993) 2 SCC 746, ¶ 22.

produced a Robert Williams. No reported judgment documents an Indian citizen wrongfully arrested because a facial recognition algorithm misidentified them; no parliamentary committee has named a specific victim; no media investigation has traced an individual's detention to a false biometric match. If the harm is real, the argument goes, where is the evidence. This objection deserves a direct answer, not because it is analytically compelling, but because it reveals the precise mechanism by which algorithmic harms sustain themselves. The same opacity that defeats the victim's constitutional remedy also prevents the harm from acquiring the social visibility that would ordinarily generate legal and legislative attention. The absence of documented cases is not evidence of safety; it is evidence of a system in which the conditions for visibility do not exist.

Indian criminal procedure is structured, in relevant respects, to prevent precisely the kind of disclosure that would allow an algorithmic victim to identify what happened to them. The Bharatiya Nagarik Suraksha Sanhita, like its predecessor, imposes no obligation on an arresting officer to disclose the investigative methods that led to the suspicion underlying an arrest. A person taken into custody is entitled to know the grounds of arrest and, eventually, to examine prosecution materials; they are not entitled to interrogate the technological apparatus that produced the match which set the arrest in motion. No provision requires that an accused person be notified that facial recognition was used in their case. No provision compels the police to disclose a match report, a similarity score, or the composition of the watchlist database against which their face was compared. The investigative record that reaches the magistrate, the defense lawyer, and ultimately the trial court may contain nothing more than an officer's statement that the accused matched a known offender, a statement that is formally true but entirely uninformative about the algorithmic process that preceded it.

The structural consequences of this information asymmetry are compounded by the material conditions of Indian criminal adjudication. The National Crime Records Bureau's Prison Statistics confirm that a substantial majority of persons in Indian jails are undertrials, many detained for periods far exceeding what any eventual sentence would warrant. For this population, engaging with the procedural apparatus of bail, legal representation, and evidentiary challenge is already an overwhelming undertaking. An under-resourced defense lawyer, operating under the pressures of a heavily docketed criminal court, is unlikely to identify, let alone suspect, that the arrest underlying their client's detention originated in an algorithmic identification rather than a traditional witness statement or police investigation.

The client, who has no notice and no access to the underlying match data, cannot instruct their lawyer differently. The result is a structural silence: a category of case that passes through the criminal justice system, leaving no trace of its algorithmic origin in any document that any subsequent researcher, journalist, or court could examine.

This opacity is not a peripheral feature of the problem; it is its defining characteristic. And it operates with particular force against those populations most likely to be misidentified by algorithms. As Part II of this article has shown, facial recognition systems produce false positives at demonstrably higher rates for darker-skinned individuals, for women, and for populations historically over-represented in police databases, the very communities whose members are least likely to have the legal literacy, financial resources, or social capital to investigate the source of their arrest and bring it to constitutional attention. A wealthy, educated individual who found themselves wrongfully arrested might engage a senior advocate capable of demanding disclosure, might attract media attention, and might file a writ petition with the resources to pursue it to a final order. The persons who actually bear the disproportionate burden of false-positive identification are, by the structure of the same system, least equipped to make their harm visible. Opacity and disadvantage reinforce each other, the legal system, by providing no mechanism to break that cycle, becomes an accomplice to it.

The point has a constitutional dimension that goes beyond procedure. The right to equality under Article 14 is violated not only by laws that facially discriminate but by laws, or the absence of laws, whose operation produces structured inequality in the distribution of legal protection.[31] A legal framework that makes the remediation of algorithmic harm contingent on the victim's capacity to identify, articulate, and litigate their own victimization distributes rights in inverse proportion to need.[32] The algorithmic victim who knows their rights and can enforce them does not, in practice, require constitutional protection nearly as urgently as the algorithmic victim who does not. Article 14's guarantee of equal protection under the law cannot be satisfied by rights that exist on paper but are practically accessible only to the privileged.[33]

The appropriate constitutional response to this analysis is what might be called anticipatory

[31]The Supreme Court first moved beyond formal equality toward substantive classification review in E.P. Royappa v. State of Tamil Nadu, (1974) 4 SCC 3, ¶ 85
[32]National Crime Records Bureau, Prison Statistics India 2022, at 28–31 (2023)
[33] Olga Tellis v. Bombay Municipal Corporation, (1985) 3 SCC 545,

constitutional protection. The phrase captures a principle already latent in Indian rights jurisprudence but not yet systematically applied to algorithmic harms, that constitutional rights function as prophylactic guarantees, protecting citizens from foreseeable state-caused injury before that injury becomes publicly visible, and that the state's obligation to respect and protect rights does not wait upon the emergence of a documented scandal. The Supreme Court in Puttaswamy recognized privacy not merely as a liberty claim against identified violations but as a structural precondition for the exercise of all other rights, a recognition that necessarily implies a forward-looking, preventive dimension. Justice Chandrachud's plurality opinion explicitly drew on the proportionality framework developed in German and South African jurisprudence, which asks not only whether a specific violation has occurred but also whether the regulatory architecture is capable of preventing foreseeable violations. A state that deploys a technology known to produce false positives at measurable rates, among identifiable demographic groups, without disclosure obligations, without verification requirements, and without any remedy for the misidentified, has already failed the proportionality test, irrespective of whether any particular victim has yet brought their case to constitutional attention.

This is not an argument for speculative constitutional litigation; it is an argument about the proper relationship between law and foreseeable harm. The common law of negligence long ago established that a duty of care can arise in anticipation of harm that is reasonably foreseeable, even before it materializes, constitutional law, which operates at a higher register of obligation, cannot sensibly require the production of victims as a precondition for legal protection. To insist that courts and legislators wait for a publicized Indian Williams, a named individual, a media-documented arrest, a writ petition that makes it through the High Court and into the law reports before imposing accountability on facial recognition deployments is to allow the same opacity that silences victims to silence the law itself. The constitutionalism of Maneka Gandhi and Puttaswamy is not a remedial constitutionalism after the fact, it is a constitutionalism of standards that state action must satisfy in advance. A procedure that is fair, just, and reasonable under Article 21 must be fair, just, and reasonable at the moment it is applied, not in retrospect once a wrongful arrest has attracted sufficient public attention to force judicial notice.

There is a further dimension to the waiting for scandal argument that deserves explicit rejection, the implicit assumption that the harms of the current unregulated status quo are somehow

neutral or costless. They are not. Every day that facial recognition operates in India without disclosure obligations is a day on which a person may be arrested, detained, and released without ever learning that an algorithm caused their ordeal and without any legal system recording that fact. The pre-trial detention that follows a false positive is not merely an individual misfortune, it is, under Nilabati Behera and the constitutional jurisprudence of custodial rights, a state-caused violation of personal liberty that the state has an affirmative obligation to prevent. The argument that India should wait for clearer evidence of harm before regulating algorithmically-generated suspicion rests on a conception of the state's role that the Indian Constitution has consistently rejected: the state as a passive responder to proven violations, rather than as an active guarantor of the conditions under which rights can be exercised. Anticipatory constitutional protection is not an innovation, it is what the existing constitutional text, read with seriousness, has always required.

Understood in this way, the invisibility of the Indian algorithmic victim is not an argument against the framework proposed in the preceding section; it is the most powerful argument for it. A legal order that can respond only to harm that has already been inflicted on a named complainant, a publicized case, and a documented judicial record systematically underprotects those whose injuries are structurally least likely to achieve visibility. The victim-centered framework articulated in this article mandates human verification, algorithmic impact assessments, rights of notice and explanation, and statutory compensation, and is designed precisely to create the conditions under which harm can surface, rather than to respond to harm that has already surfaced elsewhere. Whether the first Indian algorithmic victim to achieve judicial recognition was harmed last month or last year, the structural conditions that concealed their injury will continue to conceal the next one unless the law intervenes. The urgency of that intervention is not diminished by the absence of a named victim, it is defined by it.

VII. CONCLUSION

The expansion of facial recognition technology in Indian law enforcement is proceeding on the assumption that its benefits, efficiency, scale, and reduced manual identification errors justify the costs it imposes. That assumption can only be sustained by refusing to count the costs borne by those whom the system misidentifies. The Williams and Woodruff cases attracted significant public attention in the United States precisely because they made those costs visible and individualized: a specific person, in a specific driveway, in front of specific children,

arrested for no crime. India has not yet produced a Williams. It is statistically certain that it has already produced its equivalent of a person wrongly arrested on the basis of an algorithmic match, who did not know that facial recognition caused their arrest, who had no legal mechanism to challenge it, and whose name does not appear in any reported judgment or media account, because the legal system provided no surface on which their injury could register.

The algorithmic victim is not a future legal problem. They exist now, unnamed in police records and unremedied by courts, in the gap between the constitutional promise of personal liberty and the operational reality of computational identification. The development of a victim-centered legal framework, one that mandates human verification, requires impact assessments, creates rights of notice and explanation, and compensates foreseeable algorithmic harm, is not a reform agenda for a distant future; it is the minimum that the existing constitutional commitments of the Indian state already demand. A legal order that can identify the state's rights to deploy a technology but cannot see the rights of those the technology destroys is not a system of rights at all; it is a system of power that has learned to speak the language of law.

REFERENCES

1. Anuradha Bhasin v. Union of India, (2020) 3 SCC 637.
2. Bridges v. South Wales Police [2020] EWCA Civ 1058.
3. Buolamwini J and Gebru T, 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification' (2018) 81 Proceedings of Machine Learning Research 1.
4. Christie N, 'The Ideal Victim' in Ezzat Fattah (ed), From Crime Policy to Victim Policy (Macmillan 1986).
5. Grother P and others, Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects, NIST Interagency Report 8280 (NIST 2019).
6. Harwell D, 'Three Black Men Were Falsely Arrested Using Facial Recognition. Detroit Police Just Made Another Arrest' Washington Post (28 July 2021).
7. Hill K, 'Wrongfully Accused by an Algorithm' New York Times (24 June 2020).
8. Internet Freedom Foundation, Submission on the National Automated Facial Recognition System (NAFRS) (IFF 2020).
9. Joseph G, 'A Pregnant Black Woman Was Wrongfully Arrested After Facial Recognition Misidentified Her' The Guardian (7 August 2023).
10. Justice K.S. Puttaswamy v. Union of India, (2017) 10 SCC 1.
11. Maneka Gandhi v. Union of India, (1978) 1 SCC 248.
12. Mendelsohn B, 'The Victimology' (1956) 13 Études Internationales de Psycho-Sociologie CrimPaperPaperinelle 23.
13. Ministry of Home Affairs, Request for Proposal: National Automated Facial Recognition System (NAFRS) (MHA 2019).
14. Nilabati Behera v. State of Orissa, (1993) 2 SCC 746.

15. NITI Aayog, Principles for Responsible AI (NITI Aayog 2021).

16. National Crime Records Bureau, Prison Statistics India 2022 (NCRB 2023).

17. Parliamentary Standing Committee on Home Affairs, 231st Report on Unlawful Activities (Prevention) Amendment Bill (2019).

18. People's Union for Civil Liberties v. Union of India, (1997) 1 SCC 301.

19. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data (General Data Protection Regulation) [2016] OJ L 119/1.

20. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonized Rules on Artificial Intelligence (EU AI Act) [2024] OJ L 1689/1.

21. Schafer S, The Victim and His Criminal: A Study in Functional Responsibility (Random House 1968).