AICYBERLEX NEXUS

Journal of Artificial Intelligence, Cyber & Contemporary Law

Volume I · **Issue I** · **Jul–Sep 2026** · Article 1
ISSN: Applied for (National ISSN Centre, NIScPR, India) · e-ISSN: TBA
Peer-reviewed · Open Access · https://aicyberlexnexus.com

ARTICLE 1

Who Answers When the Algorithm Decides? Accountability and Liability for AI Decisions in India

Sundara

ABSTRACT

Indian banks, welfare departments, employers and police units now lean on automated systems to decide who gets a loan, a ration card, an interview or a risk score. When one of these systems gets it wrong, no Indian statute says clearly who must explain the decision or who must pay for the harm. This paper looks at that gap. It reads the Constitution, the Information Technology Act, 2000, the Consumer Protection Act, 2019 and the Digital Personal Data Protection Act, 2023 side by side, and then tests them against the American decision in Loomis, the Dutch SyRI judgment and the European Union's AI Act. The argument is modest. India does not need one grand AI code. It needs a duty to give reasons for automated decisions that touch legal rights, a liability rule that reaches the deployer first, and an audit trail that a court can actually open.

KEYWORDS: artificial intelligence; decision-making; Indian law

algorithmic

accountability;

liability;

automated

1. Introduction

A few years ago an automated decision was a curiosity. Today it is routine. A credit model screens a small-business loan before any officer sees the file. A welfare database drops a name because two records do not match. A hiring tool ranks a thousand résumés and shows a recruiter fifty. In each case a person loses something, and in most cases that person never learns why. The legal question is old, even if the technology is new: when a decision harms someone, who answers for it? Ordinary law has a ready reply for a human officer. An order can be challenged, reasons can be demanded, and a negligent act can be sued upon. The reply gets harder when the "officer" is a model trained by one company, fine-tuned by another, and bought by a third that has never seen its code. India has so far chosen not to pass a dedicated AI statute. The government's governance guidelines, issued by the Ministry of Electronics and Information Technology in late 2025, lean towards applying existing law and

adding rules only where a gap is shown.[1] That approach is defensible. But it only works if the existing law can carry the load. This paper asks whether it can. The paper has six parts. Part 2 explains why automated decisions strain the usual liability rules. Part 3 examines what Indian law offers today. Part 4 draws on three foreign examples. Part 5 sets out a workable model for India, and Part 6 concludes. The focus is on decisions made about individuals, not on wider questions of AI safety or market power. It also covers only systems that produce or heavily shape the outcome, not tools that merely assist a human who then decides independently, although the line between the two is often thinner than it looks.

1

2. Why automated decisions strain old liability rules

Three features of these systems make them awkward for courts. The first is opacity. Many models cannot give a reason in the way a clerk can. A tree of rules can be read back; a neural network produces a score from thousands of weighted signals. Even its builders may be unable to say why one applicant was refused and another approved. Companies add a second layer by claiming trade secrecy over the model and its training data. The second is the number of hands. A single automated decision usually involves a data supplier, a model developer, a platform that hosts the model, a deployer who uses it, and sometimes a human who clicks "approve". Tort law is built for a wrongdoer and a victim. It copes badly with five parties, each of whom can point to another. The third is drift. A system that was tested before launch may behave differently six months later, because the data it sees has changed or because it keeps learning. A defect is then not a fixed fact about a product but a moving target, and proof of causation becomes harder for the person harmed. Add to this a plain imbalance of information. The affected person has the outcome and nothing else. The deployer has the logs, the thresholds and the vendor contract. Any workable rule has to correct that imbalance, or the right to complain stays on paper.

3. What Indian law offers today

3.1 The Constitution The strongest tools are constitutional, and they apply mainly to the State. Article 14 forbids arbitrary action, and the Supreme Court has read arbitrariness broadly enough to cover decisions made without application of mind.[2] Article 21 requires that any procedure affecting life or liberty be fair, just and reasonable.[3] The Court has also insisted on reasons: in Kranti Associates v Masood Ahmed Khan it described the giving of reasons as a basic requirement of fairness in decision-making.[4] And Puttaswamy gave the proportionality test a central place for any measure that intrudes on privacy.[5] Put together, these cases suggest that a welfare benefit cut by an unexplained algorithm is open to challenge. The difficulty is practical. A writ lies against the State and its instrumentalities, not against a private lender or an employer. Most of the automated decisions that people feel in daily life are private ones. 3.2 The Information Technology Act, 2000 The IT Act was written for a different era. Section 79 gives intermediaries a safe harbour for content posted by others, but a company that deploys its own decision model is not an intermediary in that sense, and it should not be able to borrow the shelter. The Act's remedy for careless handling of sensitive data, section 43A, is being replaced by the data protection statute discussed below. Nothing in the IT Act speaks to the quality or fairness of an automated decision itself.

2

3.3 The Consumer Protection Act, 2019 The 2019 Act introduced product liability in Chapter VI, covering manufacturers, service providers and sellers for harm caused by a defective product or a deficient service.[6] This is the nearest thing in current law to a liability rule that could reach an AI vendor. Two problems stand in the way. It is unsettled whether software counts as a "product", and a consumer who is refused a loan by a model has not necessarily bought anything from the model's maker. A claim against the bank for deficiency in service is more plausible, but the consumer must still show the defect, and without access to the logs that is a hard task. 3.4 The Digital Personal Data Protection Act, 2023 The DPDP Act is the most relevant statute, though it is thin on automated decisions. Section 8(3) requires a data fiduciary to ensure completeness, accuracy and consistency of personal data that is likely to be used to make a decision affecting the data principal.[7] That is a useful hook, since many wrong automated decisions begin with wrong data. The Act also gives a right to information about processing and a right to grievance redress. What it does not contain is anything like the right in Article 22 of the European GDPR not to be subject to a solely automated decision. And the State enjoys wide exemptions, so the cases where explanation matters most are often the ones left outside. 3.5 Taking stock The result is a patchwork. Public bodies are bound by strong constitutional norms but few statutory duties. Private deployers face data accuracy rules and general consumer law but no duty to explain. Vendors face almost nothing. A person harmed by an automated decision therefore has to guess which door to knock on, and several of the doors are locked.

4. Lessons from abroad

4.1 Loomis In State v Loomis, the Wisconsin Supreme Court considered a sentencing court's use of COMPAS, a proprietary risk-scoring tool.[8] The defendant could not examine how the score was produced. The court allowed its use, but only with written warnings about its limits, and not as the deciding factor. The case is useful for what it shows about half-measures. A warning label does not give the affected person anything to argue with. 4.2 SyRI In 2020 the District Court of The Hague struck down the Dutch SyRI legislation, which allowed government agencies to link datasets and flag people as welfare-fraud risks.[9] The court held that the system was not sufficiently transparent or verifiable and so failed the proportionality balance under Article 8 of the European Convention on Human Rights. The court did not say algorithms are forbidden. It said that a system affecting people this way must be one that outsiders can scrutinise. 4.3 The EU AI Act

3

The EU's AI Act, Regulation (EU) 2024/1689, takes a risk-based route.[10] Systems used in areas such as education, employment, access to essential services and law enforcement fall in a high-risk class with duties of documentation, human oversight and logging. The Act also gives a person affected by a decision based on a high-risk system a right to a clear explanation of the system's role. This sits beside the Court of Justice's ruling in SCHUFA, which treated a credit score that strongly drives a lender's decision as itself an automated decision for GDPR purposes.[11] India need not copy this structure. But the three examples agree on one thing: accountability starts with being able to see inside the process.

5. A workable model for India

The proposal here is deliberately narrow. It builds on what exists, and it can be done through amendments and rules rather than one new code. First, a duty to give reasons. Any deployer, public or private, that uses an automated system to make or substantially shape a decision affecting a legal right, a benefit, credit, employment

or access to an essential service should have to give the affected person, on request, a plain statement of the main factors and the role the system played. The DPDP Rules or a sectoral regulator such as the Reserve Bank of India are natural places to put this. It would not require disclosure of source code, only of reasons a person can use to contest the outcome. Second, deployer-first liability. The person harmed should be able to claim against the entity that used the system, which is the one they dealt with and the one that profited from it. That entity can then recover from the vendor under contract. This mirrors how the 2019 Consumer Protection Act already treats a seller. It spares the victim the task of untangling a supply chain they cannot see. Third, a shifted burden of proof where logs are missing. If a deployer cannot produce the records of how a decision was made, a court should be free to draw an adverse inference. The incentive is simple: keep proper logs, or accept the risk of losing. Fourth, a human who can actually change the result. Many systems keep a human in the loop on paper, but the human only confirms the machine's output. Rules should ask whether the reviewer had the information, time and authority to disagree. A rubber stamp is not oversight. Fifth, proportion. Not every recommendation engine needs this machinery. The obligations should apply by reference to the stakes: money, liberty, livelihood, access to public services. A film suggestion is not in the same league as a rejected pension claim. None of this is costless, and deployers will say compliance burdens slow adoption. That worry deserves a fair hearing. But the burden falls mostly on keeping records and being able to explain, which a responsible deployer should be doing anyway.

6. Conclusion

4

Existing Indian law was not written with automated decisions in mind, and it shows. Constitutional doctrine can discipline the State but not private deployers. The IT Act is silent. Consumer law offers a path, but an uncertain one. The DPDP Act helps on data accuracy but gives no right to understand or challenge an automated outcome. The answer is not to wait for a comprehensive statute, nor to pretend the present patchwork is enough. A short list of targeted duties — reasons, deployer liability, adequate logs, meaningful human review — would close the worst gaps without choking innovation. The common thread in Loomis, SyRI and the AI Act is that a decision which cannot be examined cannot be fairly accepted. Indian law should say the same.

References [1] Ministry of Electronics and Information Technology, India AI Governance Guidelines (2025). [2] E.P. Royappa v State of Tamil Nadu (1974) 4 SCC 3; Shayara Bano v Union of India (2017) 9 SCC 1. [3] Maneka Gandhi v Union of India (1978) 1 SCC 248. [4] Kranti Associates Pvt Ltd v Masood Ahmed Khan (2010) 9 SCC 496. [5] Justice K.S. Puttaswamy (Retd) v Union of India (2017) 10 SCC 1. [6] Consumer Protection Act 2019, Chapter VI (ss 82-87). [7] Digital Personal Data Protection Act 2023, ss 8(3), 11-13, 17. [8] State v Loomis, 881 NW 2d 749 (Wis 2016). [9] NJCM c.s. v The Netherlands (SyRI), District Court of The Hague, 5 February 2020, ECLI:NL:RBDHA:2020:1878. [10] Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (AI Act). [11] Case C-634/21 SCHUFA Holding (Scoring) (CJEU, 7 December 2023).

5

Cite as: Sundara, “Who Answers When the Algorithm Decides? Accountability and Liability for AI Decisions in India”, *AICyberLex Nexus*, Vol. I, Issue I, Jul–Sep 2026, Article 1. Available at https://aicyberlexnexus.com
© 2026 AICyberLex Nexus. Open access under the terms stated on the journal website.