INTERNATIONAL JOURNAL FOR CORPORATE AND COMPETITION LAW [IJCCL]

2026 || ISSN. 3049-3560 (Online) || Volume 2 Issue 1

ISSN: 3049-3560 (O) International Journal for Corporate and Competition Law Vol. 2 Issue 1 (Dec-Feb)

WHEN THE ALGORITHM DECIDES YOUR FUTURE: INDIA'S DPDP ACT AND THE AI GOVERNANCE GAP

Satyam[1]

INTRODUCTION

My neighbour Rajan teaches history at a municipal school in Nagpur. He is the sort of man who keeps every receipt, pays his bills three days early, and once returned ₹200 to a shopkeeper who had miscounted change in Rajan's favour. Last year, he applied for a small home loan through one of those sleek fintech apps that promise decisions in minutes.

The decision came in forty seconds. Declined.

No reason was given. No officer had looked at his file. An algorithm trained on patterns Rajan would never see, weighing signals he never consented to share had decided he was a credit risk. His phone model, perhaps. His location. The apps he had installed. We do not know. The app would not say.

Rajan asked the company to explain. He got a template email. He asked how to challenge the outcome. There was no process for that. Under India's law as it stands today, he had no clear legal right to either.

I keep thinking about Rajan because what happened to him is not a glitch. It is the system working as designed. Across India, algorithmic systems are making consequential decisions about people's lives who gets a loan, who clears a hiring filter, whose insurance premium quietly doubles and most of the people on the receiving end have no idea it is happening, let alone any way to push back.

The Digital Personal Data Protection Act, 2023 ("DPDP Act") arrived with genuine fanfare, and it deserves some credit. But it was built for a different problem. What the Act handles well is stopping companies from storing your data without permission. What it does not handle and what

[1] The author is a student of law at Kingston Law College, West Bengal.

Rajan's situation exposes is the question of what happens when a company uses your data lawfully, hands it to an algorithm, and the algorithm quietly ruins your afternoon.

This article takes the position that the DPDP Act is not worthless for AI governance. With some imaginative reading, parts of it can be made to do useful work. But imagination has limits, and beyond those limits there are gaps that no amount of purposive construction will fill. India needs to be honest about where those gaps are because pretending they do not exist will not make Rajan's problem go away.

TWO DIFFERENT PROBLEMS

I want to slow down here, because I think there is a conflation that runs through most Indian commentary on this subject, and it makes the debate harder than it needs to be.

Data protection and AI governance are related, but they are not the same thing. Data protection is fundamentally about information who collected it, whether they had permission, how long they kept it, whether you can see it or ask for it to be deleted. These are important questions. The DPDP Act answers most of them reasonably well.

AI governance is about something harder to pin down. It is about systems that take information and do things with it - infer, predict, score, rank, decide. The scary part is not that these systems have your data. The scary part is that they have your data, they have everyone else's data, and from that combined picture they draw conclusions about you that you cannot see, cannot verify, and in most cases cannot challenge. Whether those conclusions are accurate, whether they are fair, whether a human being ever paused to ask if the model was working correctly - these questions are essentially untouched by data protection law, which stops at the point the data is collected and does not follow it into the algorithm.

The EU AI Act (2024) was built around this insight. It classifies AI systems by the risk they pose to real people — the higher the stakes, the stricter the obligations and it gives individuals specific rights when automated systems make significant decisions about them: a right to explanation, a

right to human review, and a right to contest the outcome.[2] The NIST AI Risk Management Framework takes a similar view from the governance side, treating AI risk as something that needs to be mapped, measured, and managed on an ongoing basis, not ticked off once at deployment.[3]

India has neither of these. What India has is the DPDP Act, which was designed by people thinking about data protection and not, it seems, about what happens inside the black box.

WHERE THE DPDP ACT CAN ACTUALLY HELP

That said, there is more in the Act than people give it credit for, and I want to make that case properly before tearing it apart.

Section 8 is the most underappreciated provision in the Act. Section 8(4) requires every Data Fiduciary to put in place "appropriate technical and organisational measures" to ensure compliance and prevent data breaches. That is open-textured language, and open-textured language is an invitation for regulators and courts to fill it in. If India's Data Protection Board takes its mandate seriously, there is nothing in Section 8(4) that prevents it from deciding that operating an AI system without testing it for discriminatory outputs, or without maintaining logs of automated decisions, falls below the standard of "appropriate technical measures." That reading is not far-fetched — it is what the UK's Information Commissioner's Office has effectively done with equivalent language under the UK GDPR.[4]

Section 8(3) is narrower but more pointed. It says that where personal data is going to be used to make a decision affecting an individual, the Data Fiduciary must ensure that data is complete, accurate, and consistent. A credit-scoring algorithm fed on outdated income data, or a hiring filter trained on a biased sample, is in breach of this provision at least in principle. Getting a regulator to enforce it that way is a different story, but the textual hook is there.

[2]Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), art. 9, 13, 14, 2024 O.J. (L 1689) (EU) [hereinafter EU AI Act].
[3]Nat'l Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1 (Jan. 2023).
[4]Info. Comm'r's Office & Alan Turing Inst., *Explaining Decisions Made with AI*, pt. 1, at 6–14 (2020).

Section 10 is the Act's closest thing to risk-tiering, and it matters. The Central Government can designate entities as "Significant Data Fiduciaries" — SDFs based on how much data they process, how sensitive it is, and how much risk it poses to individual rights.[5] Once designated, an SDF has to conduct Data Protection Impact Assessments, submit to independent audits, and appoint a Data Protection Officer. These obligations are not specifically designed for AI, but they create exactly the kind of compliance infrastructure that AI governance requires. Facial recognition vendors, credit bureau APIs, behavioural analytics firms; any of them could be designated as SDFs. The question is whether the Central Government will use this power proactively or leave it dormant.

Sections 5 and 6 do more work in an AI context than most people notice. The Act requires that individuals receive clear, specific notice of why their data is being processed, and give consent that is genuinely free and informed. What this means in practice is that a company cannot get your consent to "improving our services" and then use your data to train a credit-scoring model — those are different purposes, and using data for one after obtaining consent for the other is a straightforward breach.[6] Section 6(4)'s right to withdraw consent is worth something too. It is not a right to contest an automated decision, but it is a right to say: stop using my data for this. In a world where most people have no idea their transaction histories are feeding AI training pipelines, that is at least a start.

WHERE THE ACT RUNS OUT OF ROAD

And yet. Rajan is still waiting for his explanation.

The gap I find hardest to argue around is the complete absence of any right to contest an automated decision. The GDPR — Europe's data protection law has Article 22, which gives individuals the right not to be subjected to a decision made solely by automated means when that decision significantly affects them.[7] They can ask for human review. They can ask for an explanation. They

[5]Digital Personal Data Protection Act, No. 22 of 2023, § 10(1)–(2) (India) [hereinafter DPDP Act].
[6]DPDP Act, supra note 4, §§ 5–6; cf. 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), art. 5(1)(b), 2016 O.J. (L 119) 1 (EU) [hereinafter GDPR] (for comparative context on how purpose-specific consent operates in practice).
[7]GDPR, supra note 5, art. 22.

can push back. The DPDP Act has nothing like this. Section 12 gives you the right to correct inaccurate data — your name, your address, your income figure. What it does not give you is any right to challenge what the algorithm concluded from that data. If an AI reads Rajan's pin code and decides it signals financial instability and the pin code is accurate; Section 12 is useless to him. The data is right. The inference is wrong. The Act simply does not go there.

The SDF mechanism has a problem that I do not think enough people are pointing out. It is triggered by data volume and data sensitivity not by the potential for harm. So a company running a high-stakes hiring algorithm that processes modest amounts of personal data might never qualify for SDF designation, while a supermarket loyalty programme processing millions of routine purchases might be captured. The Act is filtering for the wrong thing. Rajan's fintech lender might process relatively little data per user; a few hundred data points, perhaps and sail right past the SDF threshold, never subject to a DPIA, never audited, never required to explain itself.

Algorithmic transparency gets no serious treatment anywhere in the Act. Section 5 requires companies to disclose the purpose of processing. Rajan's lender can write "credit assessment" in its notice and be fully compliant. It does not have to tell him that the model treats battery-saving mode as a proxy for financial stress, or that people from his pin code are systematically scored lower, or that no human being reviewed his application at any point. The EU AI Act requires deployers of high-risk AI to provide information sufficient for individuals to understand and challenge outcomes.[8] The DPDP Act requires nothing of the kind.

Two other gaps I want to flag briefly, because they tend to get lost. First: the Act has no framework for AI training data. The Section 8(3) accuracy obligation kicks in at the point of an individual decision not at the point of training, which is where the biases that produce discriminatory decisions are actually baked in. By the time an individual is affected, the damage is long done. Second: when the AI model that assessed Rajan's loan was built by a US foundation model company, fine-tuned by a startup in Bengaluru, and deployed through a Mumbai fintech app, who exactly is the Data Fiduciary? The Act's binary of Fiduciary and Processor was not designed for this kind of layered supply chain, and it shows.

[8]EU AI Act, supra note 1, art. 13.

WHAT CAN ACTUALLY BE DONE

I want to resist the temptation to end with a neat list of three recommendations in matching sentences, because that would be too easy and probably not very useful. So let me try to be more honest about what I think is actually achievable.

The most tractable intervention, and the one that requires no new legislation, is for the Data Protection Board to develop a spine. Section 8(4)'s "appropriate technical measures" language is broad enough to support a reading that encompasses AI-specific obligations - bias auditing, output monitoring, automated decision logging. The ICO did exactly this in 2020 when it published its guidance on explaining AI decisions, treating algorithmic transparency as an implicit component of accountability under existing law.[9] India's DPBI could do the same. The question is whether it will, or whether it will spend its early years handling individual data breach complaints and leaving the harder structural questions for someone else.

The Central Government's rule-making power under Section 40 is the second lever. It is broad. It could be used to prescribe, for AI systems operated by SDFs, minimum requirements around human review in high-stakes contexts, disclosure of automated decision logic, and pre-deployment impact assessment. It will not happen without pressure from civil society, from affected individuals, from academics willing to make the case publicly. But the legal authority is there.

On the longer question of a dedicated AI governance framework: MEITY has been running consultations since at least 2023, and IndiaAI has produced discussion papers, but nothing binding has emerged.[10] Whatever framework does emerge needs to be legally enforceable — not a voluntary code of practice that companies can sign, display on their website, and ignore. It needs risk-based classification of AI systems, not just data-based classification of the companies that run

[9]Info. Comm'r's Office & Alan Turing Inst., supra note 3, at 6–14.
[10]Ministry of Elecs. & Info. Tech., *IndiaAI Mission Consultations and Discussion Papers* (2023–2024), https://indiaai.gov.in.

them. And it needs to give people like Rajan something they currently have nothing of: a legal right to know why an algorithm made a decision about their life, and a meaningful way to challenge it if the answer is not good enough.

CONCLUSION

The DPDP Act 2023 is a real achievement. Saying that is not flattery for years, India had no meaningful data protection law, and having one matters. The Data Protection Board, once it is fully operational, has the potential to become a serious institution. These things are worth defending.

But I keep coming back to Rajan. He did everything right. He paid on time, he kept records, he applied through proper channels. And a machine he had never heard of, running logic he could never access, made a decision about his future in less time than it takes to read this sentence. He has no remedy. He has no explanation. He has a template email.

That is not good enough. And until India builds a legal framework that takes algorithmic accountability as seriously as it takes data protection, it will keep not being good enough for Rajan, and for everyone who looks like him, lives where he lives, earns what he earns, and applies for things through apps that decide in forty seconds.

The DPDP Act is where we started. It is not where we should stop.