Journal of International Research & Multidisciplinary Innovation

Data Privacy Challenges in the Age of Artificial Intelligence
Under India's Digital Personal Data Protection Act, 2023

Author: Gaurav Yadav

Affiliation: Mahatma Gandhi Kashi Vidyapeeth, Varanasi

Email: gauravyadav62@gmail.in

Submitted to: Journal of International Research & Multidisciplinary Innovation (JIRMI)

Date: 10 January 2026

ABSTRACT

India's Digital Personal Data Protection Act, 2023 (DPDP Act) arrived as the country's first comprehensive privacy statute, filling a regulatory void that ad hoc provisions under the Information Technology Act, 2000 had long struggled to cover. Yet the legislation was conceived primarily as a general data-protection instrument, not as a direct response to the governance challenges posed by artificial intelligence. This paper argues that the Act's architecture contains four structural gaps when assessed against AI-specific privacy risks: the legal status of machine-generated inferences about individuals remains undefined; there is no statutory right against consequential automated decisions; obligations on Data Fiduciaries regarding algorithmic transparency are thin; and the broad exemptions granted to State instrumentalities sit uncomfortably with the proportionality standard mandated by the Supreme Court's right-to-privacy ruling. Drawing on the EU General Data Protection

Journal of International Research & Multidisciplinary Innovation

Regulation and the EU AI Act as comparative reference points, the paper identifies concrete amendments and subordinate legislation that could make the DPDP Act meaningfully fit for the AI era without hampering India's growth ambitions.

Keywords: *DPDP Act 2023, Artificial Intelligence, Data Privacy, Algorithmic Accountability, GDPR, Automated Decision-Making, Data Fiduciary, India, Data Protection Board.*

I. INTRODUCTION

There is something inherently uncomfortable about asking a general statute to do specialised work. India's Digital Personal Data Protection Act, 2023 is a carefully constructed privacy framework—but it was not written with machine learning pipelines or recommender systems in mind. The result is a law that governs the fuel (personal data) powering some of the most consequential technology of our time, without saying anything meaningful about the engine itself.

This matters because AI-driven data processing is not merely a faster or larger-scale version of what came before. It is qualitatively different. An AI system trained on location traces, purchase histories, and browsing patterns can reliably infer an individual's health conditions, political beliefs, or financial vulnerabilities—attributes the person never disclosed, and attributes the platform never claimed to collect. The consent framework built into the DPDP Act, however carefully designed, cannot address harms that arise downstream of the original data collection and that were not foreseeable at the moment consent was given.

Journal of International Research & Multidisciplinary Innovation

This paper proceeds in five parts. Part II briefly contextualises the DPDP Act within its constitutional and legislative lineage. Part III identifies the specific AI-driven privacy risks the Act must confront. Part IV examines how the Act's current provisions measure up against those risks.

Part V draws on EU law as a comparative reference. Part VI sets out targeted recommendations.

II. LEGISLATIVE AND CONSTITUTIONAL BACKGROUND

The DPDP Act came into force on 11 August 2023, ending a legislative saga that began with the Justice B.N. Srikrishna Committee's 2018 report and passed through two withdrawn bills and a Joint Parliamentary Committee review. The final statute is noticeably leaner than its predecessors: the right to data portability, mandatory data localisation for sensitive categories, and separate regulation of non-personal data were either dropped or deferred to future rules.

The Act's constitutional anchor is the Supreme Court's nine-judge ruling in *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2017) 10 SCC 1, which held privacy—including the right to control one's personal data—to be a fundamental right under Article 21. That ruling imposes a proportionality requirement on any state interference with privacy, which bears directly on the sweeping exemptions the DPDP Act grants to government agencies under Section 17. Whether those exemptions survive proportionality scrutiny is a question that will eventually reach the courts.

On the institutional side, the Act establishes a Data Protection Board of India (DPBI) with adjudicatory functions, designates Data Fiduciaries as the primary duty-bearers, and carves out a sub-category of Significant Data Fiduciaries (SDFs) for enhanced obligations including Data

Journal of International Research & Multidisciplinary Innovation

Protection Impact Assessments (DPIAs) and third-party audits. These are solid foundations. The problem lies in what the Act does not say about AI.

III. HOW AI CREATES PRIVACY CHALLENGES CONVENTIONAL LAW CANNOT ABSORB

A. The Inference Problem

Modern AI models do not merely process data that individuals hand over—they generate new data about those individuals. Feed a model enough seemingly innocuous inputs and it will produce inferences about health, sexuality, political belief, or financial stress that are often more accurate than self-reported information. These inferences are personal data in every meaningful sense: they are about an identifiable person and they can affect that person's life. But they were never collected, never consented to, and—under a strict reading of the DPDP Act's definition of "personal data"—may not even fall clearly within the statute's scope.

This is the aggregation problem dressed in algorithmic clothes. The harm is not located in any single data point but emerges from the combination of many, filtered through a statistical model. Consent given for one purpose (say, a supermarket loyalty programme) is entirely disconnected from the inference produced at the other end of the pipeline (say, a pregnancy prediction used to time advertising).

B. Automated Decisions and the Absence of Human Judgment

Increasingly, decisions that used to require a human being—whether to approve a loan, shortlist a job applicant, grant bail, or allocate a hospital bed—are made or substantially pre-determined by algorithmic systems. The privacy concern here is not just about data use: it is about what happens

Journal of International Research & Multidisciplinary Innovation

to a person when data about them is fed into an opaque model that produces a consequential output with no meaningful opportunity for the person to understand or contest it. The DPDP Act creates no right of explanation and no right to request human review. An individual rejected by an AIpowered credit-scoring system has, under current Indian law, no statutory hook to demand to know why.

C. The Opacity of Machine Learning Systems

Privacy law is built on a disclosure model: you tell people what you collect, why you collect it, and how you use it. AI undermines all three limbs of that model. What is collected may be far less important than what is inferred from it. The purpose declared at collection may bear little resemblance to the actual downstream use. And "how" a deep neural network processes data is, in the strict technical sense, often unknowable even to the model's operators. Transparency notices written before a model is trained cannot honestly describe outputs that are themselves emergent properties of the training process.

IV. THE DPDP ACT THROUGH AN AI LENS

A. Consent and Its Limits

Section 6 of the Act requires consent to be free, specific, informed, and unambiguous. These are the right words. The difficulty is that AI-driven processing frequently involves purposes that were not specified at collection and that could not have been anticipated by either the data principal or the data fiduciary. The Act does not resolve whether using personal data to train a model constitutes a new purpose—one that requires fresh consent—or whether it falls within the scope of the original consent. This interpretive gap is not trivial: it determines whether millions of Indians can

Journal of International Research & Multidisciplinary Innovation

meaningfully exercise any control over how their data feeds into the AI systems that increasingly shape their lives.

B. Rights of Data Principals

The Act grants four rights: access to information (Section 11), correction and erasure (Section 12), grievance redressal (Section 13), and nomination (Section 14). The right to erasure is particularly fraught in the AI context. Once personal data has been used to train a model, deleting the original data does not remove the data's influence from the model's learned parameters. The Act does not engage with this technical reality, meaning that the right to erasure—in an AI training context—may be unenforceable in practice, even where it is legally triggered.

More consequential is what the Act omits. There is no right to explanation for automated decisions, no right to object to profiling, and no right to data portability. These are not peripheral niceties—they are the practical instruments through which individuals can exercise meaningful control in an environment where data about them is continuously processed by systems they have no visibility into.

C. Significant Data Fiduciaries and DPIAs

The SDF framework under Section 10 is the Act's most promising AI-governance tool. By empowering the Central Government to designate large data processors as SDFs and requiring them to conduct DPIAs and algorithmic audits, the Act creates at least the scaffolding for substantive accountability. The problem is that this scaffolding is incomplete. The Act does not specify what a DPIA must contain, what methodological standards auditors must apply, or how audit results should be disclosed. Without that detail—which must come through rules under

Journal of International Research & Multidisciplinary Innovation

Section 40—the DPIA obligation risks becoming a procedural checkbox rather than a genuine riskassessment exercise.

D. The Section 17 Problem

Section 17 permits the Central Government to exempt State instrumentalities from any or all provisions of the Act, subject to national security and public order considerations. This is significant because the Indian government is simultaneously one of the country's largest data collectors—through Aadhaar, the National Population Register, and sectoral databases—and an active developer of AI systems for governance. An Act that can be disapplied to government AI deployments by executive notification, without specifying proportionality safeguards or parliamentary oversight, is difficult to reconcile with the proportionality standard the Supreme Court set in Puttaswamy.

V. WHAT THE EU FRAMEWORK OFFERS

The EU's GDPR, in force since 2018, is the most-cited comparative reference for Indian data protection law. Its Article 22 confers a right not to be subject to decisions based solely on automated processing that produce significant effects—the exact right the DPDP Act lacks. Article 35 mandates DPIAs for high-risk processing with considerably more methodological specificity than the DPDP Act's SDF provisions. And the GDPR's accountability architecture—mandatory records of processing activities, DPO appointments, and robust supervisory authority powers—has generated an enforcement practice that Indian policymakers can study rather than replicate wholesale.

Journal of International Research & Multidisciplinary Innovation

The EU AI Act, adopted in June 2024, adds a second regulatory layer explicitly targeting AI systems. It classifies AI applications by risk, subjects high-risk systems (including those used for credit scoring, employment, education, and law enforcement) to mandatory conformity assessments and post-market monitoring, and requires operators to maintain technical documentation adequate for regulatory review. The significance for India is not that it should copy this structure, but that the EU's experience illustrates why general data protection law alone—however well drafted—is insufficient to govern AI applications. A sector-agnostic AI governance instrument, or at minimum an AI-specific chapter within the DPDP framework, will eventually be necessary.

VI. RECOMMENDATIONS

Four targeted interventions would substantially improve the DPDP Act's effectiveness as an AI governance instrument.

First, the definition of personal data in Section 2 should be amended to expressly include data derived or inferred by automated systems about an identifiable individual, regardless of whether the source data itself constituted personal data. This closes the inference gap without restructuring the Act's architecture.

Second, a new provision—modelled on but simpler than GDPR Article 22—should give data principals the right to be informed when a significant decision affecting them has been made predominantly by automated means, and to request human review of that decision. "Significant" should be defined to cover decisions with material consequences for legal rights, financial standing, employment, or access to government services, keeping the compliance burden proportionate.

Journal of International Research & Multidisciplinary Innovation

Third, the Central Government should use its rule-making power under Section 40 to prescribe substantive DPIA standards for SDFs deploying AI systems. These standards should require an assessment of discriminatory risk, a description of the model's functioning adequate for regulatory review, and a post-deployment monitoring plan. Audit reports should be disclosed—at least in summary—to the DPBI.

Fourth, Section 17 exemptions should require, at minimum, time-limiting, notification to the DPBI, and a designated independent oversight mechanism, even if not the DPBI itself. This does not curtail legitimate State functions; it subjects them to a floor of accountability consistent with the constitutional proportionality requirement.

VII. CONCLUSION

The DPDP Act is a credible starting point, not a finished answer. Its core architecture—consentbased processing, a tiered fiduciary framework, a dedicated adjudicatory board—provides a workable foundation. What it does not provide is a clear account of how that foundation applies to systems that generate personal data from other data, make consequential decisions without human involvement, and operate through mechanisms that are, by design, difficult to inspect or explain.

These are not obscure edge cases. They describe the normal operation of AI systems already deployed at scale across Indian banking, healthcare, government service delivery, and e-commerce. The legislative gap is real and its consequences are accumulating. The recommendations in this paper are modest in legislative ambition—none of them require the Act to be redrafted—but they

Journal of International Research & Multidisciplinary Innovation

would meaningfully shift the balance back toward the data principals the Act is meant to protect. Whether Parliament and the executive choose to act on them is, ultimately, a question about what kind of digital society India intends to build.

REFERENCES

A. Primary Sources

1. Digital Personal Data Protection Act, 2023 (No. 22 of 2023), Ministry of Electronics and Information Technology, Government of India.
2. Information Technology Act, 2000 (No. 21 of 2000), Government of India.
3. Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1 (Supreme Court of India).
4. Regulation (EU) 2016/679 (General Data Protection Regulation), OJ L 119/1, 4 May 2016.
5. Regulation (EU) 2024/1689 (EU AI Act), OJ L 2024/1689, 12 July 2024.
6. IT (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules, 2011.

B. Secondary Sources

7. Justice B.N. Srikrishna Committee, 'A Free and Fair Digital Economy: Protecting Privacy, Empowering Indians' (Ministry of Electronics and Information Technology, 2018).

Journal of International Research & Multidisciplinary Innovation

8. Joint Parliamentary Committee, Report on the Personal Data Protection Bill, 2019 (Lok Sabha Secretariat, December 2021).

9. NITI Aayog, 'Responsible AI for All: Adopting the Framework' (2022).

10. Meha Agarwal and Anubhav Pandey, 'The DPDP Act 2023: An Assessment of India's Data Protection Architecture' (2024) 36(1) National Law School of India Review 45.

11. Sandra Wachter, Brent Mittelstadt and Luciano Floridi, 'Why a Right to Explanation of Automated Decision-Making Does Not Exist in the GDPR' (2017) 7(2) International Data Privacy Law 76.

12. Mike Ananny and Kate Crawford, 'Seeing without Knowing: Limitations of the Transparency Ideal and its Application to Algorithmic Accountability' (2018) 20(3) New Media & Society 973.

13. Prateek Waghre, 'India's Draft Digital Data Protection Bill: An Analysis' (Internet Freedom Foundation, Working Paper No. 1/2022).

14. Paul Schwartz and Karl-Nikolaus Peifer, 'Transatlantic Data Privacy Law' (2017) 106(1) Georgetown Law Journal 115.

15. Shoshana Zuboff, The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power (Profile Books 2019).

16. AI Now Institute, 'AI Now Report 2023' (AI Now Institute, 2023) <https://ainowinstitute.org>.

Journal of International Research & Multidisciplinary Innovation

17. World Economic Forum, 'Data-Driven AI: Responsible Data Governance for the AI Economy' (WEF White Paper, 2023).

18. Brenda Leong and Patrick Hall, 'Artificial Intelligence and Data Privacy: Four Problems and Four Solutions' (Future of Privacy Forum, 2020).

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.