
ARTIFICIAL INTELLIGENCE, DATA PROTECTION AND PRIVACY: EMERGING CHALLENGES UNDER INDIAN CYBER LAW

Sanskriti Sinha & Rudra Jaiswal, B.A. LL.B. (Hons.), Amity Law School, Amity
University, Gwalior, Madhya Pradesh, India

ABSTRACT

India now regulates artificial intelligence with a statute that was not written for it. The Digital Personal Data Protection Act, 2023, whose substantive provisions become enforceable on 13 May 2027, rests on a transactional model in which an identified fiduciary collects identified data from an identified individual for a stated purpose. Machine learning inverts almost every element of that model. Corpora are assembled by scraping rather than solicitation, purposes are discovered after collection rather than declared before it, and the data most likely to injure a person is often not data collected at all but data the model has inferred. This paper examines that mismatch across the Indian framework, reading the DPDP Act with the Rules of 2025, the surviving provisions of the Information Technology Act, 2000, the intermediary guidelines as amended in February 2026, and the India AI Governance Guidelines. It identifies a set of specific difficulties, among them the exclusion of publicly available personal data under section 3(c)(ii), the impossibility of erasure once data has been absorbed into model weights, the absence of any statutory definition of anonymisation, the omission of a right to an explanation for automated decisions, and the withdrawal of the compensation remedy that section 43A of the Information Technology Act had provided. It argues that the courts, working through personality rights litigation, have outpaced the legislature, and that this is an unstable substitute for statutory protection because it is available in practice only to the famous. It closes with reforms that fit India's sector-led approach.

Keywords: Artificial Intelligence; Data Protection; Privacy; DPDP Act 2023; Information Technology Act 2000; Deepfakes; Algorithmic Accountability; Cyber Law.

I. INTRODUCTION

There is a particular kind of legislative misfortune in which a law arrives just late enough to be obsolete on the day it commences. The Digital Personal Data Protection Act, 2023 is not quite in that position, but it is uncomfortably close to it. The Act was conceived during a decade in which the dominant privacy anxiety concerned what companies did with information that people had handed over to them. By the time the Rules were notified in November 2025 and the compliance clock began to run, the more pressing anxiety had shifted to what machines were doing with information that nobody had handed over to anybody. The Government's own framing acknowledges the shift; the India AI Governance Guidelines describe artificial intelligence as the defining force of the present industrial moment and commit India to building the entire AI stack.¹ What the Guidelines do not do, because they are not a statute, is tell a person whose face has been used to train a facial recognition model what she may do about it.

The constitutional starting point is settled. A nine-judge Bench in *Justice K.S. Puttaswamy (Retd.) v. Union of India* held that privacy is intrinsic to Article 21 and located informational privacy squarely within it.² The statutory position is more complicated. The DPDP Act³ and the Digital Personal Data Protection Rules, 2025⁴ together supply a compliance architecture, but most of its substantive obligations do not bite until 13 May 2027.⁵ Meanwhile the Information Technology Act, 2000 continues to do the practical work of Indian cyber law, and the intermediary guidelines were amended in February 2026 to bring synthetically generated information within their scope.⁶ The result is a regime assembled in layers, each responding to a different technological moment, and none of them written with foundation models in mind.

This paper does not argue that India needs a comprehensive artificial intelligence statute. That argument has been made often and, in the authors' view, made too quickly. It argues something narrower: that a number of specific provisions in the existing framework produce results that

¹ Ministry of Electronics and Information Technology, Government of India, India AI Governance Guidelines (5 November 2025) 4.

² Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1.

³ The Digital Personal Data Protection Act, 2023, No. 22 of 2023, Acts of Parliament, 2023 (India) [hereinafter DPDP Act].

⁴ The Digital Personal Data Protection Rules, 2025, G.S.R. 846(E), Gazette of India, Extraordinary, Part II, Section 3(i) (13 November 2025) [hereinafter DPDP Rules].

⁵ Ministry of Electronics and Information Technology, Notification, G.S.R. 843(E) (13 November 2025). Sections 3 to 17, 28 to 34, 36 and 37 of the DPDP Act come into force eighteen months from the date of publication, that is, on 13 May 2027.

⁶ The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, G.S.R. 120(E) (10 February 2026), in force from 20 February 2026 [hereinafter 2026 Amendment Rules].

are either incoherent or unjust when applied to machine learning systems, and that these can be corrected without a new Act. Part II sets out why AI systems strain the ordinary logic of data protection law. Part III recovers the constitutional baseline. Part IV describes the statutory framework as it stands in the middle of 2026. Part V, which is the substance of the paper, works through the emerging challenges one by one. Part VI places the Indian position against the European, American and Chinese approaches. Part VII examines how the courts have responded, largely through personality rights litigation. Part VIII proposes reforms.

II. WHY ARTIFICIAL INTELLIGENCE SITS AWKWARDLY WITH DATA PROTECTION LAW

Data protection statutes across jurisdictions share a common structure. Somebody decides to collect information about somebody else, tells them why, obtains permission or relies on a defined lawful basis, uses the information for the declared purpose, keeps it secure, and deletes it when the purpose is exhausted. Every one of those steps assumes a bounded, deliberate act of collection. Machine learning does not proceed that way, and three features of it explain most of the difficulty.

A. The Appetite of the Training Corpus

Modern models are trained on datasets whose defining quality is indiscriminateness. A web crawl does not distinguish between a press release, a court judgment, a health forum post and a photograph uploaded to a wedding album that was never meant to leave the family. The material is gathered because it exists and is reachable, not because anyone assessed its relevance. The developer generally cannot say, at the point of collection, which individuals appear in the corpus, and often cannot say afterwards either. Notice and consent presuppose a knowable counterparty. Scraping is defined by its indifference to who the counterparty is.

A corpus is also not merely a storage problem. It shapes what the model does, permanently. Where the training material reflects historical patterns of exclusion, the model reproduces those patterns in its outputs, and it does so in a way that survives the deletion of any particular record.⁷ Bias in this setting is not a defect that can be patched; it is a property that has been learned.

⁷ See generally Solon Barocas and Andrew D. Selbst, 'Big Data's Disparate Impact' (2016) 104 California Law Review 671.

B. Memorisation and the Illusion of Aggregation

For a long time it was assumed that once individual records were absorbed into a statistical model they lost their individual character. The assumption has not held up. Large language models have been shown to reproduce verbatim sequences from their training data, including names, telephone numbers and addresses, when prompted appropriately.⁸ A model is therefore not always a summary of its training data. Sometimes it is a lossy but recoverable copy of it. Any legal framework that treats trained parameters as inherently non-personal rests on a technical premise that does not consistently hold.

C. Inference as a Method of Collection

The third feature is the one Indian law handles worst. AI systems do not merely process personal data. They manufacture it. From ordinary behavioural traces such as purchase records, browsing patterns or location history, models derive conclusions about sexual orientation, political affiliation, religious practice, pregnancy and mental health.⁹ None of it was disclosed by the person concerned, and none of it was collected in any ordinary sense of that word.

Section 2(t) of the DPDP Act defines personal data as data about an individual who is identifiable by or in relation to such data.¹⁰ On a plain reading, inferred data satisfies that definition, and the better view is that it falls within the Act. But the Act's machinery is built around the notion of data being *provided* by a Data Principal. The notice requirement under section 5 is framed in terms of the personal data proposed to be processed. The consent under section 6 is limited to such personal data as is necessary for the specified purpose. Neither provision contemplates that the fiduciary will end up holding categories of information about the person that neither party anticipated at the outset. The Act does not exclude inferred data. It simply does not see it.

D. Opacity

Then there is the problem of explanation. Opacity in machine learning arises from several

⁸ Nicholas Carlini et al., 'Extracting Training Data from Large Language Models' (2021) Proceedings of the 30th USENIX Security Symposium 2633.

⁹ Michal Kosinski, David Stillwell and Thore Graepel, 'Private Traits and Attributes are Predictable from Digital Records of Human Behaviour' (2013) 110(15) Proceedings of the National Academy of Sciences 5802.

¹⁰ DPDP Act, § 2(t), defines personal data as any data about an individual who is identifiable by or in relation to such data.

distinct sources, and it is worth separating them. Some is deliberate, arising from trade secrecy. Some is a function of the technical literacy required to read model documentation. But some is intrinsic: a system with billions of parameters produces outputs through interactions that no human reviewer traces step by step.¹¹ Disclosure obligations in data protection law were designed against the first two kinds of opacity. They assume that the processor knows the reason and is merely being made to say it. Where the processor does not know either, a disclosure duty produces documentation without understanding.

III. THE CONSTITUTIONAL BASELINE

Before the DPDP Act, Indian privacy law was constitutional law. It is worth recalling how unpromising the beginning was. In *M.P. Sharma* and *Kharak Singh* the Supreme Court declined to read a general right of privacy into the Constitution, though Subba Rao J.'s dissent in the latter has worn considerably better than the majority.¹² A line of smaller Benches then built the right up incrementally, through decisions on domiciliary visits, publication of private facts and telephone interception.¹³

Puttaswamy settled the question. Chandrachud J. held that informational privacy is a facet of the right to privacy and observed that the dangers to privacy in an age of information can originate not only from the State but from non-State actors as well.¹⁴ The Court expressly recognised that the ability to make choices about one's own information is part of the autonomy that Article 21 protects, and it directed the Union Government to bring a data protection regime into existence.¹⁵ That direction produced, after six years and three drafts, the DPDP Act.

The second contribution of the constitutional jurisprudence is the proportionality standard. In the *Aadhaar* judgment the Court adopted a four-part test: a legitimate State aim, a rational connection between the measure and the aim, the absence of a less restrictive alternative, and a balance between the extent of the interference and the benefit sought.¹⁶ *Anuradha Bhasin* applied the same discipline to internet restrictions and insisted that the least intrusive measure

¹¹ Jenna Burrell, 'How the Machine Thinks: Understanding Opacity in Machine Learning Algorithms' (2016) 3(1) *Big Data and Society* 1.

¹² *M.P. Sharma v. Satish Chandra*, AIR 1954 SC 300; *Kharak Singh v. State of Uttar Pradesh*, AIR 1963 SC 1295.

¹³ *Gobind v. State of Madhya Pradesh*, (1975) 2 SCC 148; *R. Rajagopal v. State of Tamil Nadu*, (1994) 6 SCC 632; *People's Union for Civil Liberties v. Union of India*, (1997) 1 SCC 301.

¹⁴ *Puttaswamy* (n 2), ¶ 297 (Chandrachud J.).

¹⁵ *Ibid*, ¶ 325.

¹⁶ *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2019) 1 SCC 1, ¶ 157 (Sikri J.).

be identified and justified.¹⁷

The significance of this for algorithmic governance is easy to overlook. Proportionality is a demanding standard when applied honestly to AI deployment by the State. A facial recognition system installed at scale in public spaces must be shown to be necessary, and shown to be the least restrictive means available. Very few of the systems presently deployed by Indian public authorities have been subjected to that inquiry, largely because no statutory instrument requires it and because judicial review is available only when somebody with standing and resources initiates it.

IV. THE STATUTORY FRAMEWORK

A. The Information Technology Act, 2000 and the SPDI Rules

For twenty-three years, section 43A of the IT Act was India's data protection statute.¹⁸ It required a body corporate handling sensitive personal data to maintain reasonable security practices, and made it liable to pay compensation for negligence resulting in wrongful loss or gain. The Information Technology (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules, 2011 gave the section content, defining sensitive personal data to include passwords, financial information, health data, sexual orientation and biometric information, and requiring a published privacy policy.¹⁹

This regime is being dismantled. Section 44(2) of the DPDP Act omits section 43A along with the rule-making power that supported the 2011 Rules.²⁰ The consequence deserves emphasis, because it is not widely appreciated. Section 43A was the only provision in Indian law under which an individual could claim compensation for a data breach. The DPDP Act contains no equivalent. Its Schedule provides for penalties payable to the Consolidated Fund of India, not damages payable to the person whose data was compromised. A regime that was already weak on individual remedies is being replaced by one with none at all.

What survives of the IT Act still matters. Sections 66C and 66D address identity theft and cheating by personation, both of which are directly engaged by AI-driven impersonation;

¹⁷ Anuradha Bhasin v. Union of India, (2020) 3 SCC 637, ¶ 71.

¹⁸ The Information Technology Act, 2000, No. 21 of 2000, Acts of Parliament, 2000 (India) [hereinafter IT Act].

¹⁹ The Information Technology (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules, 2011, G.S.R. 313(E) (11 April 2011), r. 3.

²⁰ DPDP Act, § 44(2), omits section 43A and clause (ob) of section 87(2) of the IT Act.

section 66E addresses the capture and transmission of images of a private area; and section 72A punishes disclosure of information in breach of a lawful contract.²¹ The Supreme Court has held that where the IT Act deals with a subject specially, its provisions prevail over the general criminal law.²²

B. The DPDP Act, 2023 and the Rules of 2025

The Act applies to digital personal data processed within India and, extraterritorially, to processing outside India connected with offering goods or services to Data Principals in India. Processing must rest either on consent or on one of the legitimate uses enumerated in section 7.²³ The legitimate uses include voluntary provision of data by the Data Principal for a purpose, provision of subsidies and services by the State, medical emergencies, employment purposes, and disaster response.²⁴

The obligations on a Data Fiduciary are conventional in shape: notice, accuracy, security safeguards, breach notification, and erasure once the purpose is served.²⁵ Rule 7 of the DPDP Rules requires intimation to the Data Protection Board without delay, with detailed particulars to follow within seventy-two hours.²⁶ Data Principals have rights of access, correction, erasure and grievance redressal, together with a right to nominate.²⁷

The genuinely novel element is the treatment of Significant Data Fiduciaries. Entities notified as such under section 10 must appoint a Data Protection Officer based in India, engage an independent data auditor, and conduct an annual Data Protection Impact Assessment and audit.²⁸ Rule 13(3) goes further and requires such fiduciaries to observe due diligence to verify that algorithmic software deployed for hosting, display, uploading, modification, publishing, transmission, storage, updating or sharing of personal data does not pose a risk to the rights of Data Principals.²⁹ This is the first binding provision in Indian law that speaks directly to algorithmic risk. Its limitations are discussed in Part V. Penalties run to two hundred and fifty

²¹ IT Act, §§ 66C, 66D, 66E, 72A.

²² Sharat Babu Digumarti v. Government (NCT of Delhi), (2017) 2 SCC 18, ¶ 30.

²³ DPDP Act, §§ 4, 6(1).

²⁴ DPDP Act, § 7.

²⁵ DPDP Act, § 8(5) read with DPDP Rules, r. 6 and Schedule I.

²⁶ DPDP Rules, r. 7, requiring intimation to the Board without delay and detailed particulars within seventy-two hours.

²⁷ DPDP Act, §§ 11 to 14; DPDP Rules, r. 14.

²⁸ DPDP Act, § 10(2); DPDP Rules, r. 13.

²⁹ DPDP Rules, r. 13(3).

crore rupees for a failure of security safeguards.³⁰

C. The Intermediary Guidelines and the 2026 Amendment

On 10 February 2026 the Ministry of Electronics and Information Technology notified amendments to the intermediary guidelines that came into force ten days later. The amendments insert a definition of synthetically generated information, covering audio, visual or audio-visual content created or altered by computational means so as to appear authentic or indistinguishable from a real person or a real event, with carve-outs for routine edits such as camera filters and accessibility enhancements.³¹

Three changes matter. First, the window for acting on a court order or a government notification under rule 3(1)(d) was cut from thirty-six hours to three.³² Second, complaints concerning non-consensual intimate imagery, including morphed and synthetically generated nudity, must be acted upon within two hours.³³ Third, synthetically generated information must be labelled prominently and, where technically feasible, must carry embedded provenance metadata that users are forbidden to strip or suppress; significant social media intermediaries additionally bear proactive verification duties.³⁴

The definition is technology-neutral, which is sensible drafting, but it is also broad. It catches synthetic media regardless of whether artificial intelligence was involved, and the assessment of whether content is realistic enough to fall within it is left to platforms operating under a three-hour clock. The predictable behavioural response is over-removal, which is precisely the outcome *Shreya Singhal* was concerned to prevent.

D. Sectoral Regulation and the AI Governance Guidelines

India has chosen not to enact an omnibus AI statute. The Reserve Bank's FREE-AI Committee reported in August 2025 with a framework for the financial sector built on seven principles and

³⁰ DPDP Act, § 33 read with the Schedule.

³¹ 2026 Amendment Rules (n 6), inserting r. 2(1)(wa) in the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021.

³² Ibid, amending r. 3(1)(d) of the 2021 Rules to substitute a period of three hours for the earlier period of thirty-six hours.

³³ Ibid, amending r. 3(2)(b) of the 2021 Rules.

³⁴ Ibid, amending r. 3(3) and r. 4 of the 2021 Rules. The draft amendment circulated in October 2025 had proposed that a label cover at least ten per cent of the surface area of visual content; the numerical threshold was dropped in the notified text in favour of a qualitative standard of prominence.

twenty-six recommendations.³⁵ In November 2025 the Ministry of Electronics and Information Technology released the India AI Governance Guidelines, which adapted those seven principles as national sutras and organised recommendations across six pillars grouped under enablement, regulation and oversight.³⁶ The Guidelines recommend an AI Governance Group, a Technology and Policy Expert Committee and an AI Safety Institute, and expressly favour voluntary commitments and the application of existing law over new legislation.³⁷ India's hosting of the AI Impact Summit in February 2026, which concluded with the New Delhi Declaration, consolidated this positioning internationally.³⁸

There is a good deal to be said for this. Sectoral regulators understand their sectors, and one horizontal statute risks imposing financial-sector controls on an agricultural advisory tool. The difficulty is that the model works only if the sectoral regulators actually regulate. The Reserve Bank has moved. Regulators in health, education, employment and policing largely have not, and it is in those sectors that algorithmic decisions bear most heavily on people with the fewest resources to contest them.

V. EMERGING CHALLENGES

A. Consent and the Training Corpus

Section 6(1) sets a high bar. Consent must be free, specific, informed, unconditional and unambiguous, signified by a clear affirmative action, and confined to the personal data necessary for the specified purpose.³⁹ Applied strictly to AI training, the provision is close to unworkable. A developer assembling a corpus from the open web has no relationship with the individuals whose data it contains and no means of describing the purpose with the specificity the section demands. The honest notice would say that the data may contribute to any capability the model eventually develops, which is not a specified purpose but the negation of one.

The deeper problem is that consent as a regulatory technique does not scale. Solove's objection remains unanswered: individuals cannot meaningfully evaluate the consequences of each

³⁵ Reserve Bank of India, Report of the Committee on Framework for Responsible and Ethical Enablement of Artificial Intelligence (13 August 2025).

³⁶ India AI Governance Guidelines (n 1), Part I.

³⁷ Ibid, Part II and Part IV.

³⁸ Ministry of External Affairs, Government of India, New Delhi Declaration on AI Impact (19 February 2026), adopted at the India AI Impact Summit held at Bharat Mandapam, New Delhi, from 16 to 21 February 2026.

³⁹ DPDP Act, § 6(1). Consent must be free, specific, informed, unconditional and unambiguous with a clear affirmative action, and must be limited to such personal data as is necessary for the specified purpose.

disclosure, and the aggregate burden of doing so across hundreds of services exceeds anyone's capacity.⁴⁰ Consent then becomes a formality that transfers risk to the person least able to assess it. Section 6 raises the formal standard without addressing the structural problem, and the practical consequence is that AI developers will either rely on section 7 legitimate uses, which were not designed for this, or on the section 3(c)(ii) exclusion, to which we now turn.

B. The Publicly Available Data Exclusion

Section 3(c)(ii) provides that the Act does not apply to personal data that the Data Principal has made publicly available, or that any person is under a legal obligation to make public.⁴¹ In an era of web-scale scraping, this is the single most consequential provision in the statute, and it appears to have received the least attention during its passage.

Consider what it means. A person who posts a photograph on a public social media profile, writes a blog about a medical condition, or appears in a news report has made that data publicly available. Under section 3(c)(ii), the Act simply does not apply. No notice is owed. No consent is required. No purpose limitation attaches. No right of erasure exists. The data may be scraped, aggregated with other sources, used to train a commercial model and used to generate inferences about the person, and nothing in the DPDP Act constrains any of it.

The provision reflects a conception of privacy as secrecy, under which information once disclosed is exhausted. Nissenbaum's account of contextual integrity is the better one: a disclosure made in one context carries norms about its onward flow, and a photograph shared with a social circle has not been released to a training pipeline.⁴² The European position illustrates the contrast. Public availability under the GDPR does not remove data from the Regulation's scope; it may bear on the balancing exercise under the legitimate interests basis, but the controller still owes duties. The Indian exclusion is absolute, and it is difficult to see it surviving a serious constitutional challenge under *Puttaswamy*, though no such challenge has yet been mounted.

⁴⁰ Daniel J. Solove, 'Privacy Self-Management and the Consent Dilemma' (2013) 126 Harvard Law Review 1880, 1888.

⁴¹ DPDP Act, § 3(c)(ii).

⁴² Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford University Press 2010) 140.

C. Purpose Limitation and the Reuse of Models

Purpose limitation requires that data collected for one purpose not be used for an incompatible one.⁴³ Foundation models are built on the opposite principle. Their commercial value lies in the fact that a model trained on general data can be adapted to tasks that were unknown when training began. A model trained on medical literature and clinical notes for diagnostic support can be fine-tuned for insurance underwriting. The data was the same. The purpose has changed entirely, and the second purpose may be one the Data Principal would have refused.

The Act does not address this because it regulates the flow of data and not the artefacts that data produces. Once personal data has been used in training, the resulting model is treated as outside the framework. That treatment is defensible only if the model contains no personal data, and the memorisation literature indicates that this cannot be assumed.

D. Erasure and the Trained Model

Section 8(7) obliges a Data Fiduciary to erase personal data on withdrawal of consent or when the purpose is no longer served.⁴⁴ The corresponding right under the GDPR has generated a substantial literature on its application to models.⁴⁵ The Indian provision has not, but the technical problem is identical and it is severe.

Deleting a record from a training database does not remove its influence from a model already trained on it. The information persists in the weights, distributed across parameters in a form that cannot be located and excised. Machine unlearning is an active research field precisely because the naive solution, retraining from scratch, is prohibitively expensive for large models. And where memorisation has occurred, deletion from the database leaves a model that can still emit the deleted data on the right prompt.⁴⁶

So what should erasure mean here? One answer, which the United States Federal Trade Commission has adopted in consent orders, is algorithmic disgorgement: where a model was trained on unlawfully obtained data, the model itself must be destroyed.⁴⁷ The remedy is blunt,

⁴³ Regulation (EU) 2016/679 (General Data Protection Regulation), art. 5(1)(b) [hereinafter GDPR].

⁴⁴ DPDP Act, § 8(7).

⁴⁵ GDPR, art. 17.

⁴⁶ Carlini et al. (n 8), 2634.

⁴⁷ Federal Trade Commission, In the Matter of Everalbum, Inc., Decision and Order, File No. 192 3172 (11 May 2021).

and in many cases it would be disproportionate. What it does get right is the target: it acts on the thing that carries the harm forward. Neither the DPDP Act nor the Rules contain any comparable power, and the Data Protection Board has no express authority to order that a model be retrained, restricted or destroyed.

E. Anonymisation without a Definition

The DPDP Act applies to personal data. It contains no definition of anonymisation, no standard by which data is to be treated as sufficiently de-identified to fall outside the Act, and no consequence for re-identification.⁴⁸ This is a significant gap, because the boundary of the statute depends on it.

The gap is not merely technical. Re-identification research has repeatedly shown that individuals can be re-identified from datasets stripped of direct identifiers, with a small number of quasi-identifiers frequently sufficient in large populations.⁴⁹ Where a Data Fiduciary asserts that a dataset has been anonymised, it is making a legal claim that removes the dataset from the Act's coverage, and at present the claim is unverifiable by any published standard and unpunishable if false. A fiduciary facing an erasure request has an obvious incentive to characterise its holdings as anonymised. Nothing in the current framework tests that characterisation.

F. Automated Decisions and the Missing Right to Explanation

The most conspicuous omission in the DPDP Act is automated decision-making. There is no right to be informed that a decision was made by an algorithm, no right to meaningful information about the logic involved, no right to human review, and no right to contest. The GDPR provides all four in some measure.⁵⁰ The Indian statute provides none.

The practical effect is that a loan application refused by a credit-scoring model, a job application filtered by a screening system, or an insurance claim declined by an automated assessment can be dealt with in India without the applicant learning that a model was involved, still less why it decided as it did. The right to correction under section 12 is of no assistance,

⁴⁸ DPDP Act, § 2(t); compare GDPR, recital 26, which addresses anonymous information expressly.

⁴⁹ Luc Rocher, Julien M. Hendrickx and Yves-Alexandre de Montjoye, 'Estimating the Success of Reidentifications in Incomplete Datasets Using Generative Models' (2019) 10 Nature Communications 3069.

⁵⁰ GDPR, arts. 13(2)(f), 14(2)(g), 15(1)(h) and 22.

because the input data may be perfectly accurate and the objection may be to the inference drawn from it.

Indian anti-discrimination law compounds the problem. Articles 14 and 15 bind the State, and the State action doctrine has not been extended to cover ordinary commercial decision-making by private firms.⁵¹ A person refused credit by a model that has learned to disfavour applicants from particular postal codes, which in Indian conditions may correlate closely with caste and religion, has no discrimination claim under the Constitution and no statutory claim under the DPDP Act. The Consumer Protection Act, 2019 offers an indirect route through the concept of unfair trade practice, but it was not designed for this and the evidentiary burden on a complainant who cannot see the model is formidable.⁵²

Rule 13(3) deserves credit here. It is the first provision in Indian law to name algorithmic risk at all, and it was drafted without any obligation to do so. Its weaknesses are structural rather than careless. It applies only to entities notified as Significant Data Fiduciaries, and no notification had been issued at the time of writing. It is framed as a due diligence obligation owed to the regulator rather than as a right vested in the individual. And it does not require that any part of the assessment be disclosed to the persons whose interests are being assessed.

G. Synthetic Media, Dignity and Non-Consensual Imagery

In under five years, deepfake technology has travelled from a research curiosity to a routine instrument of harassment, fraud and political manipulation. The harm it produces is not principally a data protection harm. It is a dignitary harm, and the DPDP Act is poorly placed to address it because the injury lies in the fabricated output rather than in the processing of the input.

Indian courts have filled the gap through personality rights litigation. In *Anil Kapoor v. Simply Life India* the Delhi High Court restrained the unauthorised use of the actor's name, voice, image and signature phrase, and dealt expressly with AI-generated content and morphed imagery.⁵³ Similar orders were granted in favour of Amitabh Bachchan and Jackie Shroff.⁵⁴

⁵¹ The Constitution of India, arts. 14 and 12. On the limits of the State action doctrine, see *Zee Telefilms Ltd. v. Union of India*, (2005) 4 SCC 649.

⁵² The Consumer Protection Act, 2019, No. 35 of 2019, § 2(47).

⁵³ *Anil Kapoor v. Simply Life India*, CS(COMM) 652/2023 (Delhi High Court, order dated 20 September 2023).

⁵⁴ *Amitabh Bachchan v. Rajat Nagi*, CS(COMM) 819/2022 (Delhi High Court, order dated 25 November 2022); *Jackie Shroff v. The Peppy Store*, CS(COMM) 389/2024 (Delhi High Court, order dated 15 May 2024).

The Bombay High Court extended the reasoning to voice cloning in *Arijit Singh v. Codible Ventures*, restraining the operation of platforms that permitted users to synthesise the singer's voice.⁵⁵ In 2025 the Delhi High Court granted wide-ranging injunctions to Aishwarya Rai Bachchan and Abhishek Bachchan, directing intermediaries to remove infringing material within seventy-two hours and directing the Ministry of Electronics and Information Technology to block identified addresses.⁵⁶

The courts responded quickly to a harm the legislature had not touched, and the orders they have passed are inventive. The foundation, though, is uncomfortable. Personality rights in India are uncodified and have developed principally as a commercial right, protecting the economic value of a well-known persona against misappropriation. The doctrinal engine is passing off. A schoolteacher whose face has been used to generate sexual imagery has no commercial persona to protect and no realistic prospect of obtaining a John Doe injunction from a High Court. The remedy exists in inverse proportion to the vulnerability of the victim.

The criminal law offers something, though it is patchwork. Sections 67 and 67A of the IT Act address obscene and sexually explicit electronic material; sections 66C and 66D address identity theft and personation; and the Bharatiya Nyaya Sanhita provides for obscenity and defamation.⁵⁷ None of these was drafted for synthetic imagery, and the definitional questions they raise, such as whether a fabricated image of a person is an image *of* that person at all, remain unresolved. The two-hour takedown obligation introduced in February 2026 is a meaningful improvement in speed, but it operates at the level of platform compliance and gives the victim no remedy against the person who made the content.

H. The State Exemption and Algorithmic Surveillance

Section 17(2)(a) permits the Central Government to exempt any instrumentality of the State from the Act by notification, in the interests of the sovereignty and integrity of India, the security of the State, friendly relations with foreign States, public order, or preventing incitement to a cognisable offence.⁵⁸ Section 17(4) additionally disapplies the storage limitation

⁵⁵ *Arijit Singh v. Codible Ventures LLP*, Interim Application (L) No. 23443 of 2024 (Bombay High Court, order dated 26 July 2024).

⁵⁶ *Aishwarya Rai Bachchan v. Aishwaryaworld.com*, CS(COMM) 956/2025 (Delhi High Court, order dated 9 September 2025); *Abhishek Bachchan v. The Bollywood Tee Shop* (Delhi High Court, 2025).

⁵⁷ IT Act, §§ 67 and 67A; The Bharatiya Nyaya Sanhita, 2023, No. 45 of 2023, §§ 79 and 356.

⁵⁸ DPDP Act, § 17(2)(a).

and erasure obligations to such exempted bodies.⁵⁹

It is hard to overstate how wide this is. The exemption operates by executive notification, without prior legislative or judicial scrutiny, and the enumerated grounds track Article 19(2) without importing the reasonableness inquiry that accompanies it. Once notified, an instrumentality is not merely relieved of consent obligations; it is released from purpose limitation, retention limits and erasure. It may retain personal data indefinitely for purposes never specified. Section 36 separately empowers the Government to require any Data Fiduciary to furnish such information as it may call for.⁶⁰

Set this against the proportionality standard the Supreme Court laid down in *Aadhaar* and applied in *Anuradha Bhasin*. A blanket executive exemption is the antithesis of a leastrestrictive-means analysis. When the exempted processing involves AI systems, facial recognition deployed across public spaces, predictive policing tools, or automated analysis of communications metadata, the individual has no notice, no access right, no correction right and no route to the Data Protection Board. Several such systems are already operational in Indian cities without any statutory basis specific to them. The DPDP Act does not authorise them; it simply declines to reach them.

I. Cross-Border Transfers and Model Training Abroad

Rule 15 adopts a negative list. Personal data may be transferred outside India except to countries or territories that the Central Government restricts by notification.⁶¹ This is a liberal and, in the authors' view, sensible position, considerably more permissive than the adequacy regime under the GDPR. A separate obligation applies to Significant Data Fiduciaries, who may be required to keep specified categories of personal data within India as determined by the Government on the recommendation of a committee.⁶²

For AI development the negative list is close to a practical necessity: training runs on distributed infrastructure, and few developers could function under a strict localisation rule. There is a blind spot in it all the same. It regulates the transfer of personal data and says nothing

⁵⁹ DPDP Act, § 17(4), which disappplies the storage limitation and erasure obligations in respect of instrumentalities of the State exempted under § 17(2)(a).

⁶⁰ DPDP Act, § 36.

⁶¹ DPDP Rules, r. 15.

⁶² DPDP Rules, r. 13(4).

about the transfer of models trained on Indian personal data. If memorisation means that a model can contain recoverable personal information, then exporting model weights may achieve indirectly what a restriction on data transfer was meant to prevent. No Indian instrument presently addresses this.

J. Section 44(3) and the Right to Information

Section 44(3) of the DPDP Act substitutes clause (j) of section 8(1) of the Right to Information Act, 2005, so that information relating to personal information is exempt from disclosure without qualification.⁶³ The previous clause exempted personal information having no relationship to any public activity or interest, or which would cause an unwarranted invasion of privacy, and permitted disclosure where the public interest justified it.⁶⁴

The public interest override is gone. That is not an AI question in itself, but it becomes one because information about government procurement and deployment of algorithmic systems, the identity of vendors, the composition of training datasets, and the results of internal accuracy testing frequently contains or is intertwined with personal information. The Right to Information Act has been the principal tool through which Indian civil society has obtained information about State surveillance systems. Narrowing it at the moment when algorithmic State capacity is expanding fastest is, at best, unfortunate timing.

K. The Capacity of the Data Protection Board

The Board is constituted by the Central Government, which appoints its Chairperson and Members and prescribes their terms of service.⁶⁵ It functions as a digital office and must complete inquiries within six months, extendable in three-month increments for reasons recorded in writing.⁶⁶ It may direct mediation and accept voluntary undertakings in lieu of proceedings.⁶⁷

Two concerns follow. The first is independence. A body whose composition and service conditions are determined by the executive is being asked to adjudicate complaints against

⁶³ DPDP Act, § 44(3), which substitutes clause (j) of § 8(1) of the Right to Information Act, 2005.

⁶⁴ The Right to Information Act, 2005, No. 22 of 2005, § 8(1)(j), as it stood before the substitution.

⁶⁵ DPDP Act, §§ 18 to 26; DPDP Rules, rr. 16 to 21.

⁶⁶ DPDP Rules, r. 19, which requires an inquiry to be completed within six months, extendable for reasons recorded in writing by periods not exceeding three months at a time.

⁶⁷ DPDP Act, §§ 31 and 32.

instrumentalities of that executive, subject to the exemption power in section 17(2)(a). The second is capacity. Assessing whether an algorithmic system poses a risk to the rights of Data Principals requires technical expertise that regulators worldwide have struggled to acquire. Nothing in the Act or the Rules requires that the Board include members with such expertise, and nothing provides for a standing technical secretariat.

VI. COMPARATIVE PERSPECTIVES

The European Union has taken the opposite structural approach, layering a dedicated AI statute on top of an existing data protection regulation. The AI Act classifies systems by risk, prohibits a defined set of practices including untargeted scraping of facial images to build recognition databases, and imposes conformity assessment, data governance, logging and human oversight obligations on high-risk systems.⁶⁸ Transparency duties apply to generative systems, including disclosure of AI-generated content and publication of summaries of training data.⁶⁹

European enforcement has also been considerably more active. The Italian data protection authority imposed a penalty of fifteen million euro on OpenAI in 2024 over the lawful basis for training and the absence of age verification.⁷⁰ Clearview AI has been fined by regulators in several Member States for scraping facial images without a lawful basis.⁷¹ Whatever one thinks of the merits of individual decisions, the European regulators have established that scraping is a regulated activity. Indian law, through section 3(c)(ii), has established close to the opposite.

China regulates generative AI through the Interim Measures of 2023, which require providers to ensure the legality of training data, to respect the rights of others in the data used, and to label generated content, all within a framework that also imposes content-alignment obligations that would be constitutionally impermissible in India.⁷² The United States has no federal statute, and regulation has proceeded through Federal Trade Commission enforcement against unfair and deceptive practices and through State legislation, of which the Colorado Artificial

⁶⁸ Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence [hereinafter EU AI Act], arts. 6 and 5.

⁶⁹ EU AI Act, arts. 50 and 53.

⁷⁰ Garante per la protezione dei dati personali (Italy), Provvedimento del 2 novembre 2024, doc. web n. 10085432, imposing a penalty of EUR 15 million on OpenAI.

⁷¹ Commission nationale de l'informatique et des libertes (France), Deliberation SAN-2022-019 of 17 October 2022.

⁷² Cyberspace Administration of China, Interim Measures for the Management of Generative Artificial Intelligence Services (effective 15 August 2023), arts. 7 and 12.

Intelligence Act is the most developed example in addressing algorithmic discrimination.⁷³

India sits oddly against all three. It has a strong constitutional privacy right, a data protection statute whose substantive provisions are not yet in force, binding rules on synthetic content that arrived before binding rules on data processing, and a declared policy preference against horizontal AI legislation. Content regulation has run ahead of data regulation. Whether that ordering reflects a considered judgment or the relative political salience of deepfakes as against training data is a question worth asking.

VII. THE JUDICIAL RESPONSE

The courts have reached AI-related harms well before Parliament did. They have done it by borrowing doctrines built for other purposes, and the seams are beginning to show.

The personality rights cases discussed above are the clearest illustration. They provide effective relief, including dynamic injunctions that extend to addresses identified after the order, and they demonstrate judicial willingness to engage with the technology rather than defer to it. But the doctrine is proprietary in character. It protects a persona because the persona has commercial value. Extending it to a person with no commercial persona would require reconceiving it as a dignitary right, which the courts have gestured towards but not yet done squarely.

Beyond personality rights, three lines of authority are likely to matter. *Shreya Singhal* established that intermediaries cannot be required to adjudicate the lawfulness of content on private complaint, and that obligations to act must be triggered by a court order or a government notification.⁷⁴ The 2026 Amendment Rules, which require proactive detection and labelling of synthetic content by significant social media intermediaries, sit awkwardly with that holding and will very likely be tested. *Ritesh Sinha* and *Selvi* address compelled extraction of biometric and mental information and supply the framework within which voice cloning and biometric inference will be assessed.⁷⁵ And the proportionality line will govern challenges to State AI deployment, though such challenges depend on the existence of a petitioner who knows the system exists, which brings us back to the narrowing of the Right to Information Act.

⁷³ Colorado Revised Statutes, § 6-1-1701 et seq. (Colorado Artificial Intelligence Act, 2024).

⁷⁴ *Shreya Singhal v. Union of India*, (2015) 5 SCC 1, ¶ 117.

⁷⁵ *Ritesh Sinha v. State of Uttar Pradesh*, (2019) 8 SCC 1; *Selvi v. State of Karnataka*, (2010) 7 SCC 263.

VIII. RECOMMENDATIONS

What follows is deliberately incremental. It assumes that India will not pass a horizontal AI statute any time soon, and that reform therefore has to work inside the architecture that already exists.

First, section 3(c)(ii) should be qualified. Public availability should not extinguish the Act's application. It should be a factor relevant to the lawful basis for processing, so that a fiduciary relying on publicly available data is required to demonstrate that its use is compatible with the context in which the data was made available. Data made public by a person about themselves should be distinguished from data made public by others about them.

Second, the Act should define anonymisation, prescribe a standard against which a claim of anonymisation is to be tested, and make deliberate re-identification of anonymised data an offence. Without this the boundary of the statute is set by the assertions of the entities it regulates.

Third, a provision on automated decision-making should be inserted. At a minimum it should confer a right to be informed that a significant decision was made by automated means, a right to meaningful information about the logic and the principal factors involved, and a right to obtain human review of decisions with legal or similarly significant effects. The obligation should be scoped by the consequence of the decision, not by the size of the entity making it.

Fourth, the Act should state expressly that inferred and derived personal data falls within section 2(t), and that the rights of access, correction and erasure extend to it.

Fifth, the Board should be given express power, where a model has been trained on personal data processed in contravention of the Act, to direct measures proportionate to the contravention, including restriction on deployment, retraining, or in appropriate cases deletion of the model. A penalty payable to the Consolidated Fund does not address a continuing harm embedded in an artefact that remains in use.

Sixth, the compensation remedy removed by the omission of section 43A should be restored in some form. The absence of any route by which an individual can recover for a data breach is a substantial regression, and it distinguishes the Indian regime unfavourably from every comparable framework.

Seventh, Rule 13(3) should be strengthened. The obligation of algorithmic due diligence should extend beyond Significant Data Fiduciaries to any entity deploying automated systems for decisions with legal or similarly significant effects, and a summary of each assessment, redacted for trade secrets, should be published.

Eighth, non-consensual synthetic intimate imagery should be made a distinct offence, with the offence defined by the absence of consent and the sexual character of the depiction rather than by obscenity, and with a statutory civil remedy available to the victim in addition to platform takedown. The dignitary interest that the High Courts have been protecting through commercial doctrines should be given a statutory home available to everyone.

Ninth, the exemption power under section 17(2)(a) should be subject to procedural discipline: written reasons, a defined period of operation, periodic review, and an annual report to Parliament identifying the instrumentalities exempted and the categories of processing covered. The power itself may be necessary. Its present form is not.

Tenth, the Data Protection Board should be provided with a technical secretariat and its appointment process placed on a footing that does not rest entirely on executive discretion. The institutions proposed in the India AI Governance Guidelines, particularly the AI Safety Institute, should be established on a statutory footing rather than as administrative arrangements, and their evaluations made available to the Board.⁷⁶

IX. CONCLUSION

India has assembled a data protection framework of respectable quality for a world in which data is collected from identified individuals for stated purposes. That world is receding. The systems that now pose the greatest risk to informational privacy in India acquire data by scraping rather than collection, generate personal information rather than receiving it, and retain it in a form from which it cannot readily be removed. The DPDP Act does not address these systems well, not because its drafters were careless but because they were solving a different problem.

The most serious defects identified in this paper are not defects of ambition. They are specific

⁷⁶ India AI Governance Guidelines (n 1), Part III, which contemplates a phased action plan and the establishment of an AI Governance Group, a Technology and Policy Expert Committee and an AI Safety Institute.

and remediable: an exclusion for publicly available data that hands the entire training corpus problem to the market, an erasure right that cannot reach trained models, an undefined concept of anonymisation that sets the statute's own boundary, an absent right to explanation, and a State exemption broad enough to swallow the constitutional protection the Act was directed to implement. Each can be repaired by amendment. None requires a new Act, and none requires India to abandon the sectoral approach it has chosen.

What the framework should not do is continue to rely on the courts to supply what the statute omits. The personality rights jurisprudence has been resourceful, and it has produced real protection for those able to reach the Delhi and Bombay High Courts. It is not a substitute for a statutory right, because a remedy that requires a commercially valuable persona and a senior advocate is not a remedy for most people. The eighteen-month period before the substantive provisions of the DPDP Act take effect on 13 May 2027 is an opportunity to make these corrections before the compliance architecture calcifies around them. It is unlikely to be repeated soon.