ALGORITHMIC MEMORY AND THE RIGHT TO FORGET: EMERGING PRIVACY CHALLENGES IN INDIA

Md Kaif, Student, B.A. LL.B. (Hons.) School of Law, Galgotias University.

ABSTRACT

Artificial intelligence systems learn from personal data and keep what they learn after the data itself is deleted. I call this algorithmic memory: the behavioural profiles and inferences a trained model can still produce about a person once her records are gone. This article asks whether Indian law gives people any way to deal with it. Reading the Supreme Court's decision in Puttaswamy, the Srikrishna Committee's 2018 report, the Digital Personal Data Protection Act, 2023 (DPDP Act) and the European Union's GDPR and AI Act side by side, I argue that a right to delete data means little unless it also reaches the inferences drawn from that data. The article then looks at state surveillance and political micro-targeting, and closes with five reform proposals.

Key Words: - *Algorithmic memory, Right to forget, DPDP Act 2023, Artificial intelligence, Inferential privacy, Behavioural profiling, GDPR, Surveillance*

I. Introduction

Jeremy Bentham designed the Panopticon so that prisoners, never sure whether anyone was watching, would police themselves. Michel Foucault later turned the design into a theory of modern power in *Discipline and Punish*: the gaze does not have to be constant, because the possibility of it is enough to change behaviour.

Today the watching is done by software. Every search, every paused scroll and every abandoned purchase is logged, cross-referenced and fed to machine learning models that build a predictive picture of a person: what she wants, how she will vote, what she fears. Algorithmic memory is what remains when that picture outlives the data it was built from. Delete the records and the model can still produce the profile.

Indian constitutional law has an answer to part of this. In *Justice K.S. Puttaswamy (Retd.) v Union of India*, a nine-judge bench held that privacy, including informational privacy, is protected under Article 21. The statute meant to give that holding effect, the DPDP Act, deals mostly with data as something collected, stored and deleted. It says very little about data a system infers.

The argument runs in a straight line. Part II explains how algorithmic memory works. Part III places the right to forget in Indian law. Parts IV and V cover state surveillance and political micro-targeting. Part VI marks the gaps in the DPDP Act, Part VII borrows from the EU, and Part VIII proposes changes. My claim is simple. A right to delete data, without a right to remove what was inferred from it, protects the form of privacy and gives up the substance.

II. How algorithmic memory works

A traditional database is an archive. A record goes in, and in principle it can be found again or erased. A large neural network works differently. It is trained on personal data, and training changes the model's internal parameters. Once that has happened, deleting the training set does not undo the change. The model keeps the pattern without keeping the record.

Computer scientists have shown this is more than a theory. Carlini and his co-authors extracted verbatim training data from GPT-2, including personal details, using nothing more than queries

to the model. Earlier work by Cao and Yang had already asked how a system could be made to forget what it learned. The legal consequence is easy to state. Suppose a person asks a company to erase her data under section 12 of the DPDP Act and the company complies. The model it trained on that data may still infer her traits, rebuild her profile and use it commercially.

Inferential identities

Models also produce attributes the person never disclosed. In a well-known study, Youyou, Kosinski and Stillwell found that a computer judging personality from about 300 Facebook Likes did better than the person's spouse. Insurers use driving data to estimate health risk. Employers use browsing history to guess who is about to resign. Campaigns use similar methods to find voters who can be moved.

Luciano Floridi argues in *The Ethics of Artificial Intelligence* that information does more than describe people, because it takes part in making them. Read that way, algorithmic memory builds a second self that nobody consented to, which runs alongside the living person and can drift a long way from what she actually thinks or wants.

III. Privacy and the right to forget in India

Puttaswamy did not create a right to forget, but it supplies the ground for one. Justice Chandrachud's plurality described privacy as having a spatial, a decisional and an informational side, and rooted all three in Articles 14, 19 and 21. Informational privacy, on that account, covers a person's control over how information about her is shared and used. The judgment spoke of data a person had given or that others had collected. It left open whether the same protection covers information that was never disclosed and was generated by a machine.

The Srikrishna Committee

The Committee of Experts chaired by Justice B.N. Srikrishna reported to the Ministry of Electronics and Information Technology in 2018. Its report recommended a right to be forgotten, framed as a right to restrict continuing disclosure of personal data rather than a right to wipe it. The Committee accepted that the right could not be absolute and had to be weighed against free speech, access to information and journalism. The report said little about inferred data or trained

models, which is understandable for 2018 and now looks like a gap.

Section 12 of the DPDP Act

Section 12 gives a data principal the right to correction, completion, updating and erasure of her personal data. Three features limit it in practice. First, the right attaches to data processed on the person's consent, and the Act says nothing about data a system derives. Second, a fiduciary may keep the data where retention is necessary for the specified purpose or to comply with any law, and the second ground is wide enough to swallow much of the rule. Third, nothing in the Act requires a fiduciary to audit or retire a model trained on data that has since been erased.

Courts have touched the idea only in reputational cases. In *Jorawer Singh Mundy v Union of India*, the Delhi High Court gave interim relief to a man acquitted in a criminal case, directing Indian Kanoon to block the judgment from search engines and relying on privacy under Article 21. It was a prima facie order about one document. It does not say what happens to a profile a model has built.

IV. Surveillance and predictive governance

Facial recognition is the clearest public-sector case of algorithmic memory. The National Crime Records Bureau's 2019 request for proposals for a National Automated Facial Recognition System describes a central system linking face matching to CCTV feeds, criminal records and other databases. Such a system does not compare a face with a stored photograph. It compares it with a mathematical embedding, a numerical summary of facial geometry that persists. A person who is wrongly matched, or who happened to be at a protest or a religious gathering, acquires an association in that system. Deleting the original image does not remove it. The Georgetown Law study on unregulated police face recognition in the United States documents the same risk abroad.

The Pegasus controversy of 2021 raised a related question about the limits of state surveillance. Journalists, activists and opposition politicians were alleged to have been targeted with spyware sold by the NSO Group. In *Manohar Lal Sharma v Union of India*, the Supreme Court appointed a technical committee, supervised by a retired judge, to investigate, after finding that the Union's response had been limited. The committee's findings were partly inconclusive, and the Court's

order treated unauthorised surveillance as a threat to rights under Articles 19 and 21. Pegasus is not an example of algorithmic memory as such, but it shows how little oversight exists over state surveillance tools.

Predictive policing adds a further problem. These tools use past crime records to rank people and places by their supposed risk of future offending. Several state police forces have piloted them. If the historical data reflects over-policing of particular communities, the model repeats that bias, and people who have committed no offence end up with a record of suspicion.

V. Micro-targeting and democratic risk

Political micro-targeting is the same machinery pointed at voters. The Facebook and Cambridge Analytica affair showed how platform data could be used to build psychological profiles and aim political messages at them. The UK Information Commissioner's Office later described a disregard for voters' privacy across the political campaigning system and fined Facebook £500,000.

Shoshana Zuboff, in *The Age of Surveillance Capitalism*, argues that platforms use behavioural predictions to steer what people do, and not just to record it. Curated feeds then narrow what each person sees, which strengthens existing beliefs and deepens polarisation. Jürgen Habermas argued in *The Structural Transformation of the Public Sphere* that democratic debate needs a shared space where citizens meet competing arguments. A feed built separately for each user removes that space.

India has more than 700 million internet users by industry estimates, and parties of every kind hire analytics firms to find swing constituencies and tailor messages to demographic and psychological profiles. Indian law has not yet asked whether AI-driven micro-targeting is electoral manipulation under section 123 of the Representation of the People Act, 1951, or whether the Model Code of Conduct reaches it. *Manohar Lal Sharma* did not decide the point.

VI. Gaps in the DPDP Act

The DPDP Act is India's first comprehensive data protection statute and a real improvement on the patchwork of the Information Technology Act, 2000 and its rules. Measured against

algorithmic memory, I count four gaps.

A. The definition of personal data

Section 2(t) defines personal data as any data about an individual who is identifiable by or in relation to that data. Read generously, that covers inferences. But the Act never says so, and the rights in section 12 are written around data a person gave or a fiduciary collected. A fiduciary that wants to treat its model's outputs as outside the Act has an arguable case, and the Act gives a regulator no clear text to answer it.

B. Consent and legitimate uses

Section 7 allows processing without fresh consent for listed legitimate uses, among them State functions, compliance with law and medical emergencies, and it also covers data a person has voluntarily provided. Large platforms can be expected to argue that much of their inferential processing falls under these heads. The Act does not draw the boundary, so the argument is open to them.

C. No rules specific to AI

The Act contains nothing on AI systems. A person has no right to be told when an automated system made a decision about her, no right to ask for human review, and no right to have a model checked for the accuracy of what it infers. The EU AI Act, by contrast, sets transparency, documentation and human oversight duties for high-risk systems.

D. The government exemption

Section 17 lets the Central Government exempt any instrumentality of the State from the Act in the interest of sovereignty, security or public order. Combined with other security laws, this could leave State-run surveillance systems, NAFRS among them, with almost no data protection scrutiny.

VII. What the EU offers

The GDPR has applied since 2018 and remains the most developed data rights regime. Article 17

sets out the right to erasure. In *Google Spain v AEPD* the Court of Justice held that a search engine must delist results about a person when the information is inadequate, irrelevant or no longer relevant, even if the original publication was lawful.

Article 22 gives people the right not to be subject to decisions based solely on automated processing, including profiling, where those decisions have legal or similarly significant effects. Read with the transparency provisions in Articles 13 to 15, it gives people some access to the logic behind such decisions. The EU AI Act, adopted in 2024, treats AI used in employment, education, law enforcement and access to essential services as high-risk, and attaches conformity assessment, transparency and human oversight requirements to it. California's privacy law goes at the problem from another side by counting inferences drawn about a consumer as personal information.

The GDPR's experience also teaches a limit. Erasure at the level of source data does not reach the model. The right to forget needs a technical counterpart, machine unlearning, meaning methods for removing the influence of specific data from a trained model. No regulator has yet settled how far erasure duties extend to trained models, and India has a chance to address the question in its own law from the start.

VIII. Proposed reforms

A. Treat inferences as personal data

The DPDP Act should be amended to say that inferred and derived data, including behavioural profiles and psychographic assessments, count as personal data. They would then carry the rights of access, correction and erasure.

B. Require machine unlearning

A new rule under the Act should require fiduciaries that deploy machine learning to keep documented unlearning procedures. After an erasure request, the fiduciary should have to show, within a fixed period, that the model has been retrained or modified to remove the person's influence, or explain why that is not technically possible.

C. Register high-risk systems

Following the EU model, India should require registration of AI systems used in law enforcement, employment screening, credit assessment and political communication. Registration should disclose training data sources, the kinds of inferences produced and accuracy rates across demographic groups.

D. Limit micro-targeting in elections

Parliament, or the Election Commission under its existing powers, should bar parties and candidates from using AI-driven psychological profiling to target voters. The Model Code of Conduct should cover digital advertising and require disclosure of targeting criteria and data sources.Section 17 should be cut back. State surveillance should need judicial authorisation and periodic review, consistent with the proportionality test in *Puttaswamy*. NAFRS should be put on hold until a dedicated surveillance oversight law is in place.

IX. Conclusion

Both common law and civil law systems assume that a right over information is a right over a particular, findable thing. Machine learning breaks that assumption. Information is absorbed into a model and comes back out in forms that deleting a file cannot reach.

Puttaswamy gives India the constitutional footing to protect people against profiles they never agreed to. The Srikrishna Committee wanted a protective, citizen-centred regime and did not foresee this problem. What is missing is legislative attention. The DPDP Act, as enacted, is written for a world where data is stored, shared and deleted in traceable ways. The most consequential facts about a person are now guessed by a model and never directly observed.

Lawrence Lessig wrote that code is law: the design of a system sets rules as effectively as a statute does. Indian lawmakers need to deal with that design directly. They should require machine unlearning, transparency about automated decisions and protection for inferred data, or the promise of *Puttaswamy* will not reach the systems that now shape people's lives.

References

1. The Constitution of India, arts. 14, 21, 22(1) and 39A.

2. The Constitution (Forty-Second Amendment) Act, 1976 (Act 91 of 1976).

3. The Legal Services Authorities Act, 1987 (Act 39 of 1987).

4. The Legal Services Authorities Act, 1987 (Act 39 of 1987), ss. 2(c), 12 and 13.

5. The National Legal Services Authority (Free and Competent Legal Services) Regulations, 2010.

6. M.H. Hoskot v. State of Maharashtra, (1978) 3 SCC 544.

7. Hussainara Khatoon (IV) v. Home Secretary, State of Bihar, (1980) 1 SCC 98.

8. Khatri (II) v. State of Bihar, (1981) 1 SCC 627.

9. Centre for Legal Research v. State of Kerala, (1986) 2 SCC 706.

10. Suk Das v. Union Territory of Arunachal Pradesh, (1986) 2 SCC 401.

11. State of Maharashtra v. Manubhai Pragaji Vashi, (1995) 5 SCC 730.

12. National Legal Services Authority, Legal Aid, available at: https://nalsa.gov.in/legal-aid/ (last visited on October 31, 2026).

13. National Legal Services Authority, Frequently Asked Questions, available at: https://nalsa.gov.in/faqs/ (last visited on October 31, 2026).

14. Government of India, The Constitution of India, Ministry of Law and Justice, available at: https://legislative.gov.in/ (last visited on October 31, 2026).

15. Press Information Bureau, Government of India, “Affordable and Accessible Justice to Poor People”, Ministry of Law and Justice, May 15, 2015, available at: https://www.pib.gov.in/newsite/PrintRelease.aspx?lang=2&reg=48&relid=121772 (last visited on October 31, 2026).