PRIVACY & DATA PROTECTION

De-Indexed but Not Deleted: The Right to Be Forgotten in the Age of *Generative AI*

AUTHORED BY

Anshul Verma
Partner

Atharv Khanna
Associate

The Legal Background

In *K.S. Puttaswamy (Retd.) v. Union of India*, Justice Kaul observed: "Humans forget, but the internet does not forget and does not let humans forget.....The footprints remain." That observation has acquired a sharper dimension in 2026, as AI systems have displaced conventional search as the primary medium through which information is accessed. This raises a question that privacy law had no occasion to confront in *Puttaswamy*: *whether such systems are even capable of forgetting, in any legally meaningful sense of that word*.

In 2014, the Court of Justice of the European Union in *Google Spain SL v. Agencia Española de Protección de Datos* (Case C-131/12) established that individuals may request search engines to delist links pertaining to outdated or irrelevant personal data which set forth the legal architecture for the Right to Be Forgotten (the "RTBF"). This was codified in **Article 17 of the General Data Protection Regulation (the "GDPR")**, conferring a statutory right to erasure subject to exceptions for freedom of expression, legal obligations, and public interest, requiring data controllers to balance erasure requests against competing rights. Indian jurisprudence developed along a parallel track, the Hon'ble Supreme Court in *Puttaswamy* affirmed informational privacy as a fundamental right under **Article 21**, and the **Digital Personal Data Protection Act, 2023 (the "DPDP Act")** operationalised this through **Section 12**, requiring data fiduciaries to erase personal data upon request unless retention is necessary for the original purpose or legal compliance. Notably, the DPDP Act does not use the phrase 'Right to Be Forgotten' which is materially distinct from the term 'right to erasure'. Statutory erasure under Section 12 is triggered by specific conditions: necessity having lapsed, consent withdrawn, or processing being unlawful. RTBF is a broader concept, allowing individuals to control their digital identity and disconnect from outdated or harmful information even

where no illegality exists, and it is this broader space that the Hon'ble Delhi High Court has now addressed.

What Laksh Vir Yadav Decided and What It Could Not

On 29th May 2026, the Hon'ble Delhi High Court in *Laksh Vir Yadav v. Union of India* explicitly recognised the **RTBF** as a constitutionally protected facet of informational privacy under Article 21, holding that it "*flows naturally and necessarily*" from the right to informational privacy. Following the EU approach, the Hon'ble Delhi High Court, after carving out exceptions, permitted relief for private individuals in respect of irrelevant or disproportionately harmful digital records by obligating search engines, digital intermediaries, and online legal databases to de-index and de-link specific judicial records and related online material from name-based searches.

While the precise scope and enforceability of RTBF against digital intermediaries, with respect to judicial records cannot be treated settled, since the Hon'ble Supreme Court's decision in *Ikanoon Software Development Pvt. Ltd. v. Karthick Theodore and Ors.* (SLP(C) No. 15311/2024) remains pending. The *Laksh Vir Yadav* judgment is a landmark in Indian privacy jurisprudence, and the framework it provides is a significant step forward. Similar to *Google Spain* and the GDPR, it rests, however, on an assumption that personal information resides in identifiable records that can be located, delisted, and confirmed as removed. Google's Gemini-powered search and Microsoft's Copilot now respond to natural language queries with AI-generated summaries that may draw on both live web retrieval and knowledge encoded during model training. De-indexing a URL therefore addresses only one of the channels through which personal information can surface, and not necessarily the most consequential one.

The search engines directed to comply with the Hon'ble Delhi High Court's directions are no longer simple URL-retrieval systems. They are AI-powered platforms whose generative features encode a person's legal history in distributed form, across billions of numerical parameters that collectively constitute the model's learned behaviour. Once personal data has been absorbed into its training, the familiar mechanics of erasure cease to function. The technical reality is that AI systems do not reliably forget unless they are fundamentally retrained.

The consequence is that de-indexing blocks information at the surface but cannot erase what the model has already learned. The directions from the *Laksh Vir Yadav* judgement are enforceable and meaningful for traditional search index and where AI assistants fetch live web results to generate

responses. However, where the data is used for training AI, the position is fundamentally different as the AI's knowledge of a person's past is not stored in a place, rather it is embedded in the model's behaviour. Therefore, in respect of model training data for AI Systems, the constitutional promise of forgetting remains technically undeliverable under the current legal framework.

The Compliance Gap Yet to be Closed

This technical limitation has not gone unnoticed by regulators, though legislative responses have struggled to keep pace. In Europe, the European Data Protection Board (the "EDPB") established a ChatGPT Taskforce and multiple national data protection authorities, including those in Italy, France, Spain, and Poland, opened investigations into generative AI platforms on grounds of transparency and legal basis for processing. Italy's Garante temporarily banned ChatGPT in 2023 and subsequently imposed a EUR 15 million fine on OpenAI, though the fine was later overturned on appeal. The broader regulatory questions it raised, however, remain live. The EDPB's guidance has focused primarily on transparency and data minimisation rather than providing concrete rules on erasure from AI systems, leaving a significant compliance gap at precisely the point where the technical problem is most acute. Regulators have gestured toward emerging concepts such as machine unlearning, but these remain experimental and no enforceable standard exists.

India's DPDP Act operates on the same assumptions as the GDPR's erasure framework. Section 12 requires data fiduciaries to erase data upon request, but it too is silent on what erasure means in the context of model weights, whether output suppression suffices, or how a data fiduciary would demonstrate that a model has genuinely forgotten a data point. The *Laksh Vir Yadav* framework is similarly limited in scope to search and delisting obligations and does not address data embedded in AI training. The result is that Indian law, like European law before it, leaves organisations in a compliance grey area on the most consequential aspect of the problem.

The problem, however, does not stop at individual privacy. For corporates especially regulated entities, the inability to guarantee true deletion of data absorbed into AI systems creates a distinct and immediate governance risk. Generative AI is now embedded in business workflows, often without full organisational awareness of the data flows involved, and the consequences of that invisibility are beginning to surface in courts. The UK Upper Tribunal in *Munir v. Secretary of State for the Home Department* [UKUT 81 (IAC)] held that uploading confidential documents to an open source AI tool such as ChatGPT amounts to placing that information in the public domain and constitutes a waiver of legal privilege. A United States federal court reached a comparable conclusion in *United*

States v. Heppner (820 F. Supp. 3d 292 S.D.N.Y., February 2026), holding that neither the prompts entered into nor the outputs generated through a public AI platform attracts attorney-client privilege or work product protection. Taken together, these decisions establish a principle with direct relevance to Indian practitioners and regulated entities i.e once confidential material enters a public AI system, it does not merely risk exposure, it loses its protected character entirely, and no subsequent deletion request can restore what has already been waived. The inability to guarantee true deletion of sensitive data fed into AI systems can trigger concurrent obligations under the DPDP Act, RBI outsourcing and KYC directions, CERT-In cybersecurity directions, and, in the case of listed entities, the SEBI insider trading and disclosure framework, particularly where AI systems have processed unpublished price sensitive information or confidential customer records. Legacy data loss prevention systems were not designed to detect a user pasting proprietary information into an AI chatbot, and the absence of an audit trail compounds the governance problem. What began as a question of individual privacy rights has therefore acquired a distinctly corporate dimension, where boards and compliance officers must now ask not only whether their data retention policies are adequate, but whether their AI usage policies prevent the creation of indelible, uncontrollable footprints that no deletion request can subsequently reach.

A Milestone, Not a Destination

Laksh Vir Yadav marks a genuine advancement in Indian privacy law and provides meaningful, enforceable relief for individuals whose legal history continues to surface through conventional search. For corporates and AI developers, however, the judgment is less a resolution than a signal of what is coming. The RTBF framework breaks down where information absorbed into AI training models has no delete function and no verified mechanism for removal. No Indian regulatory guidance currently addresses what erasure means in this context or how compliance would be demonstrated. Organisations that have fed personal data into AI systems face the possibility that data they are legally obliged to erase is data they are technically unable to erase. That is not a future risk. It is a present one.

The practical response requires treating AI usage policies as data governance instruments, building contractual protections with vendors that specifically prohibit the use of input data for model training, and anticipating that the DPDP Rules and future Data Protection Board guidance will extend erasure obligations to model training data. Developers who build with deletion capability from inception will be better positioned than those who attempt to retrofit it. The footprints Justice Kaul described are no longer merely left on the internet. They are encoded into the models that have

replaced it. The law has taken a significant step toward addressing the first problem. The second remains substantially unresolved.

DISCLAIMER

This article is intended for general information only and does not constitute legal advice. For specific queries, please write to us at contactus@skvlawoffices.com.