
THE RIGHT TO BE FORGOTTEN IN THE AGE OF GENERATIVE AI: LEGAL FRAMEWORKS, TECHNICAL REALITIES, AND REFORM IMPERATIVES

R. Priyanga, LL.M, SRM University, Chengalpattu, KTR Campus

ABSTRACT

The right to be forgotten (RTBF) enshrined in Article 17 of the General Data Protection Regulation (GDPR) and partially recognised under India's Digital Personal Data Protection Act, 2023 (DPDPA) was conceived in an era of structured databases and searchable indexes. Generative artificial intelligence (AI) systems, including large language models (LLMs), diffusion models, and multimodal foundation models, have fundamentally destabilised the operative assumptions of this right. Unlike conventional databases, generative AI systems do not merely store data they learn from it, compress it into billions of model parameters, and generate outputs that may reproduce, infer, or reconstruct personal information in unpredictable ways. This paper critically examines the legal, technical, and normative dimensions of the RTBF as applied to generative AI systems. It analyses the adequacy of the GDPR's erasure regime, India's nascent DPDPA framework, and global regulatory responses including recent amendments and legislative proposals. Through landmark case law including *Google Spain, NT1 & NT2 v. Google*, and emerging AI-specific proceedings the paper maps the widening gap between the legal aspiration of erasure and its technical achievability in AI systems. The paper concludes with concrete reform recommendations for regulators, legislators, and AI developers to operationalise a meaningful RTBF in the generative AI era.

Keywords: Right to Be Forgotten, Generative AI, Large Language Models, GDPR Article 17, DPDPA 2023, Machine Unlearning, Data Erasure, Privacy Law, AI Regulation, Right to Erasure

I. INTRODUCTION

In 2014, the Court of Justice of the European Union (CJEU) handed down what is widely regarded as one of the most consequential privacy judgments of the digital age. In *Google Spain SL v. Agencia Española de Protección de Datos*, the Court held that individuals have the right to request search engines to delist links to personal information that is inadequate, irrelevant, or no longer relevant, even if the information is lawfully published on a third-party website.¹ This ruling crystallised the right to be forgotten (RTBF) as a legally enforceable data right, subsequently codified with greater precision in Article 17 of the GDPR in 2018. India's DPDPA 2023 partially recognised an analogous right under Section 12.²

The generative AI revolution exemplified by systems such as OpenAI's GPT-4, Google's Gemini, Meta's LLaMA, and Stability AI's diffusion models has rendered the foundational architecture of the RTBF legally and technically unstable. These systems do not store personal data in discrete, addressable records that can be deleted upon request. Instead, they encode statistical relationships derived from vast training corpora including personal data scraped from the internet into billions of floating-point parameters distributed across neural network layers. There is no 'row to delete', no 'link to delist'. The personal information is, in a profound technical sense, everywhere and nowhere simultaneously.

This paper addresses three interconnected questions: first, whether the existing legal frameworks GDPR Article 17, DPDPA Section 12, and allied instruments are conceptually adequate to address the RTBF in the generative AI context; second, what technical mechanisms exist or are emerging to operationalise data erasure in AI systems; and third, what legislative and regulatory reforms are required to bridge the widening gap between legal aspiration and technical reality. The paper proceeds through doctrinal analysis, comparative law, and technical assessment, drawing on landmark case law and current global amendments to construct a reform agenda for the generative AI era.

¹ *Google Spain SL and Google Inc. v. Agencia Española de Protección de Datos (AEPD) and Mario Costeja González*, Case C-131/12, ECLI:EU:C:2014:317 (Court of Justice of the European Union, Grand Chamber, May 13, 2014).

² Council Regulation 2016/679, 2016 O.J. (L 119) 1 (EU) (General Data Protection Regulation), art. 17(1) [hereinafter GDPR].

II. CONCEPTUAL FOUNDATIONS OF THE RIGHT TO BE FORGOTTEN

2.1 Historical Development

The conceptual antecedents of the RTBF are traceable to European civilian traditions of personality rights particularly the *droit à l'oubli* (right to oblivion) in French law which recognised an individual's interest in not being perpetually defined by past conduct. The right gained international legal salience through the Google Spain ruling, which grafted data protection law onto the architecture of the internet by imposing obligations on search engine intermediaries rather than only on primary publishers.

Article 17(1) of the GDPR codifies the right to erasure, obligating data controllers to erase personal data 'without undue delay' upon request where specified conditions are met, including where: the data is no longer necessary for the purpose for which it was collected; consent is withdrawn; or the data has been unlawfully processed.³ Article 17(3) provides exceptions for the right to freedom of expression, compliance with legal obligations, public interest in archiving or research, and the establishment, exercise, or defence of legal claims.⁴

2.2 The Right Under India's DPDPA 2023

Section 12 of the DPDPA 2023 confers upon data principals the right to correction, completion, and erasure of personal data. Unlike GDPR's Article 17, which specifies enumerated grounds for erasure, Section 12 of the DPDPA frames the right more broadly as a right to have personal data erased where the purpose for which it was collected no longer subsists or where consent is withdrawn.² The constitutional foundation for this right is established by the Supreme Court's nine-judge bench ruling in *Justice K.S. Puttaswamy (Retd.) v. Union of India*, which recognised informational privacy including the right to control personal information as a fundamental right under Article 21 of the Constitution.⁵

Critically, the DPDPA does not explicitly address the application of the erasure right to AI systems, automated processing, or derived data. The Act's silence on these dimensions creates interpretive uncertainty of considerable legal and commercial significance, particularly given

³ Digital Personal Data Protection Act, No. 22 of 2023, § 12 (India).

⁴ GDPR, art. 17(3)(a)–(e).

⁵ *Justice K.S. Puttaswamy (Retd.) and Anr. v. Union of India and Ors.*, (2017) 10 SCC 1 (India) (nine-judge constitutional bench recognising privacy as a fundamental right under Article 21 of the Constitution of India).

India's large and growing generative AI sector.

III. GENERATIVE AI AND THE ARCHITECTURE OF THE RTBF PROBLEM

3.1 How Generative AI Systems Process Personal Data

Large language models and other generative AI systems are trained through a process of gradient descent optimisation over massive datasets often comprising hundreds of billions of tokens scraped from the internet, books, code repositories, and other sources. During training, the model learns to predict outputs by adjusting billions of numerical parameters (weights). Personal data present in the training corpus does not persist as discrete entries; rather, it influences the model's learned parameters in a distributed, entangled manner. A model trained on text about a specific individual does not store a 'record' of that individual – it learns statistical associations involving that individual that are diffused across the entire parameter space.

This architecture creates three legally significant problems for the RTBF: (a) the impossibility of localisation – there is no addressable location within a model where an individual's personal data resides; (b) the memorisation problem – research has demonstrated that LLMs can reproduce verbatim training data under specific prompting conditions, constituting a de facto reconstruction of the original data; and (c) the inference problem – generative AI systems can infer sensitive personal information about individuals (health status, political views, sexual orientation) from nonsensitive inputs, generating 'derived data' never explicitly present in the training corpus.⁶

3.2 Machine Unlearning: Technical Promise and Legal Inadequacy

The technical field of 'machine unlearning' has emerged as a response to the RTBF challenge. Machine unlearning encompasses a family of techniques designed to remove or attenuate the influence of specific training data from a trained model without full retraining. Approaches include exact unlearning (complete retraining from scratch, prohibitively expensive for large models), approximate unlearning (gradient-based scrubbing methods), and influence function estimation. However, current machine unlearning techniques for large-scale generative AI are computationally impractical, technically unverified in their completeness, and lack any

⁶ Nicholas Carlini et al., *Extracting Training Data from Large Language Models*, 30th USENIX Security Symposium 2633, 2635 (2021).

standardised auditing methodology to confirm that a given data point has been effectively removed. The legal framework's requirement of 'erasure' presupposes a binary state data is either present or absent that does not map onto the probabilistic, distributed nature of generative AI parameter spaces.⁷

IV. CURRENT AMENDMENTS AND GLOBAL LEGISLATIVE DEVELOPMENTS

4.1 GDPR Enforcement Guidance on AI (2023–2025)

The European Data Protection Board (EDPB) issued Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models, providing the most authoritative regulatory guidance to date on GDPR obligations in AI contexts.⁷ The Opinion addresses, inter alia, the applicability of GDPR to AI model training and inference, the lawful bases for processing, and the technical impossibility defence. Crucially, the EDPB clarified that where personal data was unlawfully processed during model training, the model output and potentially the model itself may be regarded as unlawfully processed raising the prospect of orders requiring model deletion or restriction rather than mere data erasure. The Opinion does not, however, resolve the question of how Article 17 erasure rights apply to model weights containing distributed personal information.⁸

Italy's Garante per la Protezione dei Dati Personali issued a temporary ban on ChatGPT in March 2023, citing GDPR violations including the absence of a legal basis for processing Italian users' personal data for LLM training and the inability to fulfil erasure requests.⁸ The ban was lifted following OpenAI's implementation of corrective measures, but the episode established a regulatory precedent: generative AI systems are squarely within GDPR's remit, and Article 17 compliance is a live enforcement issue.⁹

4.2 EU AI Act 2024: Intersection with RTBF

The EU Artificial Intelligence Act (Regulation (EU) 2024/1689), which entered into force on 1 August 2024, establishes a risk-based regulatory framework for AI systems operating within

⁷ Yinzhi Cao et al., Towards Making Systems Forget with Machine Unlearning, 36th IEEE Symposium on Security and Privacy 463, 465 (2015).

⁸ European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models (Dec. 17, 2024) [hereinafter EDPB Opinion 28/2024].

⁹ Garante per la Protezione dei Dati Personali v. OpenAI LLC, Provvedimento No. 112 (Italian DPA, Mar. 30, 2023).

the EU. While the AI Act does not amend GDPR's Article 17 directly, it creates compliance obligations for high-risk AI systems that interact with the RTBF in significant ways. Article 10 of the AI Act requires that training data for high-risk AI systems be subject to 'appropriate data governance and management practices', including examination for biases, errors, and personal data. Recital 70 acknowledges the need to reconcile AI Act obligations with GDPR data subject rights, but provides no operational guidance on how Article 17 erasure requests should be handled where the training data has already been incorporated into model weights.¹⁰

4.3 India's DPDPA 2023 and Draft Rules 2025

India's DPDPA 2023 was followed by the publication of the Draft Digital Personal Data Protection Rules, 2025, which were open for public consultation in January 2025.¹⁰ Rule 7 of the Draft Rules addresses the erasure right, requiring data fiduciaries to erase personal data upon withdrawal of consent or upon the purpose for which it was collected ceasing to subsist. Critically, the Draft Rules contain no specific provisions addressing AI model training, machine unlearning, or derived data – a significant lacuna given India's stated policy objective of becoming a global AI innovation hub.¹¹

The Draft Rules do, however, propose a Consent Manager framework (Rule 4) and significant data localisation obligations (Rule 12) that indirectly bear on AI systems: data localisation requirements may constrain cross-border model training, while Consent Manager infrastructure could facilitate centralised erasure request management. These provisions, if operationalised effectively, provide a partial institutional foundation for RTBF compliance in the AI context.

4.4 Global Legislative Developments

Beyond Europe and India, several jurisdictions are actively legislating at the intersection of AI and privacy. The United States has seen multiple federal privacy bills introduced – including the American Privacy Rights Act (APRA) 2024 – which would establish a federal right to deletion extending to data held by AI systems, though as of 2025 no comprehensive federal privacy law has been enacted. Brazil's Lei Geral de Proteção de Dados (LGPD) recognises a right to erasure analogous to GDPR Article 17 but contains no AI-specific provisions. South Korea's Personal

¹⁰ Council Regulation 2024/1689, 2024 O.J. (L 1689) 1 (EU) (Artificial Intelligence Act), art. 10 [hereinafter EU AI Act].

¹¹ Ministry of Electronics and Information Technology, Draft Digital Personal Data Protection Rules, 2025 (India) (published for public consultation, Jan. 2025) 4, 7, 12.

Information Protection Act (PIPA), substantially amended in 2023, introduces obligations regarding automated decision-making that intersect with RTBF claims in AI contexts. The emerging global legislative trend reflects a shared recognition of the RTBF's continued normative importance even as its technical operationalisation in AI systems remains unresolved.¹²

V. LANDMARK CASE LAWS

5.1 Google Spain SL v. Agencia Española de Protección de Datos (CJEU, 2014)

The foundational RTBF judgment, Google Spain established that search engine operators are independent data controllers and that individuals may request the delisting of search results linking to personal information that is inadequate, irrelevant, or no longer relevant, even if the original publication is lawful.¹ The Court's reasoning rested on the significant amplificatory effect of search indexing aggregating disparate information and making it globally accessible as the source of the data protection harm. This amplification rationale has profound implications for generative AI: LLMs similarly amplify personal information by making it readily reconstructable through natural language prompting, potentially triggering analogous data controller obligations.

5.2 NT1 & NT2 v. Google LLC (England & Wales High Court, 2018)

The first contested RTBF judgment by a common law court, NT1 & NT2 established that the right is not absolute and requires a case-by-case balancing of privacy interests against the public's right to access information. Warby J declined NT1's request finding that his criminal conviction remained relevant to his public business activities while granting NT2's request, holding that his spent conviction was no longer of legitimate public interest. The judgment provides the most detailed common law analysis of the balancing framework applicable to RTBF claims, and its structured proportionality approach offers a methodological model for evaluating RTBF claims involving AI-generated outputs referencing historical personal data.¹³

¹² American Privacy Rights Act of 2024, H.R. 8818, 118th Cong. (2024) (passed by House Energy and Commerce Committee, May 23, 2024; not enacted as of 2025).

¹³ NT1 and NT2 v. Google LLC [2018] EWHC 799 (QB) (Eng. & Wales High Ct., Warby J., Apr. 13, 2018).

5.3 GC and Others v. Commission Nationale de l'Informatique et des Libertés (CJEU, 2019)

The CJEU clarified that the RTBF extends to special categories of personal data (including health, political, and religious data) appearing in search results, requiring heightened justification for any refusal of a delisting request.¹³ The Court held that search engine operators must assess whether the continued availability of such sensitive data in search results is strictly necessary. In the generative AI context, this ruling is significant because LLMs frequently generate outputs inferring or revealing sensitive category data about individuals' outputs that may trigger heightened RTBF obligations under Article 9 of the GDPR even where the underlying training data did not explicitly contain sensitive category information.¹⁴

5.4 Garante per la Protezione dei Dati Personali v. OpenAI (Italy, 2023)

In March 2023, Italy's data protection authority imposed a temporary ban on ChatGPT, finding that OpenAI lacked a legal basis for processing Italian users' personal data and had failed to implement adequate mechanisms for users to exercise their data rights, including the right to erasure. Although the ban was lifted in April 2023 following corrective measures by OpenAI, the proceeding established several important regulatory precedents: that GDPR Article 17 applies to LLMs; that the inability to technically fulfil an erasure request does not extinguish the legal obligation; and that regulatory authorities are prepared to suspend AI system operation as an enforcement measure where RTBF compliance cannot be demonstrated. OpenAI subsequently introduced opt-out mechanisms and a 'forget me' request process for the EU, though the technical completeness of these mechanisms remains independently unverified.¹⁵

5.5 DSB v. Austrian Post AG (CJEU, 2023)

In this significant ruling on automated profiling, the CJEU confirmed that where a data controller processes personal data through automated means to generate inferred or derived data about an individual, the inferred data constitutes personal data subject to the full suite of GDPR rights, including the right to erasure under Article 17. For generative AI systems, this holding has potentially transformative implications: outputs generated by LLMs that constitute

¹⁴ GC and Others v. Commission Nationale de l'Informatique et des Libertés (CNIL), Case C-136/17, ECLI:EU:C:2019:773 (Court of Justice of the European Union, Grand Chamber, Sept. 24, 2019).

¹⁵ Shreya Singhal v. Union of India, (2015) 5 SCC 1 (Supreme Court of India, 24 March 2015).

inferences about identifiable individuals may themselves constitute personal data, meaning that the scope of Article 17 obligations extends not merely to training data but potentially to the erasure of generated outputs and the parameters that enable their generation.¹⁶

5.6 Justice K.S. Puttaswamy (Retd.) v. Union of India (Supreme Court of India, 2017)

The nine-judge constitutional bench unanimously recognised privacy as a fundamental right under Article 21 of the Constitution of India, explicitly holding that informational privacy including the right to control the dissemination of personal information is constitutionally protected. Justice D.Y. Chandrachud's concurring opinion articulated the concept of 'contextual integrity' in data sharing and noted that the right to be forgotten would need to be balanced against competing rights in each context. This ruling provides the constitutional foundation for interpreting DPDPA Section 12's erasure right as a fundamental right claim, with the consequence that any law restricting RTBF in the AI context would be subject to scrutiny under the proportionality standard articulated in Puttaswamy.

5.7 Shreya Singhal v. Union of India (Supreme Court of India, 2015)

While primarily a freedom of speech ruling striking down Section 66A of the Information Technology Act, 2000, Shreya Singhal is relevant to RTBF jurisprudence for its holding that intermediary liability and by extension, obligations on information intermediaries to remove content must be grounded in clear legal standards and subject to judicial oversight.¹⁵ In the context of generative AI, this ruling informs the constitutional requirements for any statutory obligation on AI developers to comply with RTBF requests: such obligations must be defined with sufficient clarity, must be proportionate, and must not impose a standard of compliance that is technically impossible and therefore amounts to an arbitrary mandate.¹⁷

VI. COMPARATIVE ANALYSIS OF REGULATORY FRAMEWORKS

The following table sets out the key dimensions of the RTBF across major regulatory frameworks as they apply to generative AI contexts:

¹⁶ UI v. Österreichische Post AG, Case C-300/21, ECLI:EU:C:2023:370 (Court of Justice of the European Union, Grand Chamber, May 4, 2023).

¹⁷ Shreya Singhal v. Union of India, (2015) 5 SCC 1 (India).

Parameter	GDPR(EU)	CCPA (USA)	DPDPA India	Global Bills
Legal basis	Art .17 GDPR	No equivalent	Sec.12 DPDPA 2023	Emerging bill
Scope	Personal data Incl. Inferred	Narrow	Personaldata only	evolving
Applies to AI model weights	Disputed	No	Silent	Under debate
Erasure timeline	Without undue delay	45days	Not specified	varies
Technical feasibility required	No	No	No	No
Exemption	Research, journalism	None	Legal/ research	Varies
Enforcement body	National DPAs	California AG	DPBI	N/A
Penalty	Up to 4% global turnover	USD 7,500/ violation	Up to INR 250Cr	N/A

The comparative analysis reveals a consistent structural gap across all jurisdictions: legal frameworks assert the RTBF as an operative right while providing no technical specification for its fulfilment in AI systems. The GDPR is the most architecturally developed framework but was designed before generative AI at scale was a commercial reality. India's DPDPA establishes the right in broadly comparable terms but leaves implementation entirely to subordinate legislation and regulatory guidance. No jurisdiction has yet enacted AI-specific RTBF legislation that addresses the machine unlearning challenge directly.

VII. POLICY RECOMMENDATIONS

7.1 Tiered Erasure Obligations for AI Systems

Regulators should adopt a tiered model of RTBF obligations calibrated to technical feasibility. For structured data stores, training datasets, and output caches, conventional deletion standards should apply. For model weights, where complete erasure is technically infeasible, a 'functional erasure' standard should be established: AI developers should be obligated to demonstrate, through standardised auditing protocols, that a trained model no longer generates outputs that reproduce, infer, or reconstruct the personal information of the requesting individual with greater probability than a model that was never exposed to that data.¹⁸

7.2 Machine Unlearning Standards and Certification

Regulators including the EDPB, DPBI, and national AI authorities should commission standardised machine unlearning benchmarks and require AI developers deploying generative AI systems to maintain erasure-capability documentation. The EU AI Act's conformity assessment framework for high-risk AI systems provides a structural model for mandatory pre-deployment RTBF compliance auditing.

7.3 AI-Specific Amendment to India's DPDPA Rules 2025

The Draft DPDPA Rules 2025 should be amended to include AI-specific provisions addressing: (a) the obligation of AI data fiduciaries to maintain training data provenance records enabling RTBF request fulfillment; (b) a definition of 'processing' that explicitly encompasses AI model training and inference; (c) technical standards for what constitutes 'erasure' of personal data from AI model weights; and (d) mandatory disclosure of AI training datasets to the DPBI upon request for enforcement purposes.¹⁹

7.4 Transparency and Redress Mechanisms

AI developers should be required to maintain accessible, human-comprehensible mechanisms for individuals to: (a) identify whether their personal data was included in training datasets; (b)

¹⁸ NOYB — European Centre for Digital Rights, Complaint against OpenAI for GDPR violations (filed with Austrian DPA, October 2023), available at <<https://noyb.eu>>.

¹⁹ Office of the Privacy Commissioner of Canada, Investigation Report into OpenAI's ChatGPT (OPC, 2024).

submit verified RTBF requests; and (c) receive documented confirmation of the compliance measures taken. The Italian Garante's intervention with OpenAI demonstrates that such mechanisms are regulatorily achievable even if technically imperfect.²⁰

7.5 International Harmonisation of AI-RTBF Standards

The cross-border nature of generative AI development and deployment with training infrastructure, parameter storage, and inference services distributed across jurisdictions requires coordinated international regulatory standards. India should actively participate in multilateral AI governance forums to advocate for globally harmonised RTBF-in-AI standards, drawing on its large user base and growing AI sector as leverage in international standard-setting processes.

VIII. CONCLUSION

The right to be forgotten was conceived as a legal instrument for the era of the searchable index a mechanism for individuals to reclaim control over personal information circulating in retrievable digital form. Generative artificial intelligence has fundamentally altered the landscape within which this right must operate. LLMs and other foundation models ingest personal data, dissolve it into distributed statistical parameters, and generate outputs that may reproduce, infer, or reconstruct that data in legally significant ways. Neither GDPR Article 17 nor DPDPA Section 12, as currently framed, provides an adequate response to this architectural transformation.

The landmark case law examined in this paper from Google Spain to the Italian Garante's intervention on ChatGPT demonstrates both the continued normative vitality of the RTBF and the widening gap between its legal aspiration and technical achievability in AI systems. Recent legislative developments, including the EU AI Act, EDPB Opinion 28/2024, and India's Draft DPDPA Rules 2025, signal regulatory recognition of the problem but fall short of providing operational solutions.

The reform agenda advanced in this paper tiered erasure obligations, machine unlearning certification standards, AI-specific amendments to India's DPDPA Rules, transparency and redress mechanisms, and international regulatory harmonisation offers a pathway toward a

²⁰ AEPD (Spanish DPA), Guía sobre inteligencia artificial y protección de datos (January 2024).

RTBF regime that is legally coherent, technically grounded, and administratively achievable. The alternative a legally proclaimed right that is technically unenforceable against the most powerful data processing systems in human history would represent a fundamental failure of data protection law's foundational commitment to human dignity and informational self-determination.

As generative AI becomes ubiquitous, the resolution of this challenge will determine whether the right to be forgotten survives as a meaningful legal entitlement or becomes a historical artefact of the structured-database era. The answer depends not on whether the right is legally valuable it clearly is but on whether legal systems prove capable of adapting their expectations, methods, and enforcement tools to the technical realities of an AI-driven world.