

THE GOLDBLOCKS STANDARD

Machine Unlearning and the Right to be Forgotten Under Emerging Legal Frameworks

-Pratiksha Ashok*

Abstract

This paper critically examines the legal frameworks that impose requirements for machine unlearning, the process by which AI systems forget previously learned information upon request. It focuses on the General Data Protection Regulation (GDPR) and the Data Act, highlighting how these laws have set early benchmarks that technical standards are only now beginning to meet. While these legal standards were, at the time of adoption, a well-balanced “Goldilocks” solution, neither too rigid nor too lenient, they face increasing strain under the weight of rapidly evolving AI technologies. The analysis explores whether these laws are truly future proof by identifying key legal and technical gaps that hinder the effective implementation of machine unlearning. The paper argues that while the legal frameworks provide essential momentum for developing unlearning capabilities, they fall short in anticipating future challenges. Ultimately, the paper calls for continuous regulatory adaptation and cross-disciplinary collaboration to ensure that laws governing AI remain both relevant and effective.

Key Words: Machine Unlearning, Goldilocks Standard, EU Law, GDPR, Data Act, AI Act.

*Pratiksha Ashok is a Post-Doctoral Researcher at Tilburg Institute for Law, Technology, and Society (TILT) and Tilburg Law and Economics Center (TILEC), Tilburg University. p.ashok@tilburguniversity.edu

1. Introduction: The evolution of machine unlearning technology

The emergence of machine unlearning techniques designed to remove the influence of specific data points from trained artificial intelligence (AI) models has introduced a critical new dimension to discussions on data governance.¹ These techniques are rapidly evolving in response to legal, technical, and ethical concerns over data persistence in AI systems.²

Unlike traditional deletion methods that remove raw data from storage, machine unlearning aims to reverse or neutralise the impact that such data has had on the model's parameters, outputs, and behaviour.³ In an era where AI systems are increasingly deployed in socially consequential domains healthcare, education, criminal justice the inability to remove personal data from learned representations raises significant concerns about autonomy, fairness, and accountability.⁴

European data protection law, most notably the General Data Protection Regulation (GDPR), has long anticipated the need for data subjects to control their digital traces through mechanisms like the right to erasure (Article 17).⁵ However, the GDPR was not drafted with machine learning or unlearning in mind. It assumes a model of data processing that remains close to the data subject, where deletion is straightforward, and influence is traceable. As a result, current law lacks clarity on how rights such as erasure or data minimisation apply once personal data has been absorbed into complex AI systems. The regulatory framework is therefore out of sync with technical reality, and with the legitimate expectations of individuals whose data fuels these systems.

Recent legal developments, including the 2024 EU Data Act and the Artificial Intelligence Act (AI Act), reflect the Union's ambition to modernise digital governance.⁶ Yet, even these newer instruments have not directly engaged with the problem of unlearning. The Data Act focuses on data access, portability, and interoperability, while the AI Act is primarily concerned with risk management and transparency. Neither provides explicit rights or obligations around reversing the influence of data within AI models. As such, there is growing uncertainty over

¹ Jiaao Chen and Diyi Yang, 'Unlearn What You Want to Forget: Efficient Unlearning for LLMs' (arXiv, 31 October 2023) <<http://arxiv.org/abs/2310.20150>> accessed 6 May 2025; Yaxuan Wang and others, 'LLM Unlearning via Loss Adjustment with Only Forget Data' (arXiv, 14 October 2024) <<http://arxiv.org/abs/2410.11143>> accessed 6 May 2025; Meg Ambrose, 'Speaking of Forgetting: Analysis of Possible Non-EU Responses to the Right to Be Forgotten and Speech Exception' 33.

² Chunxiao Li and others, 'An Overview of Machine Unlearning' [2024] High-Confidence Computing 100254; Weijia Shi and others, 'MUSE: Machine Unlearning Six-Way Evaluation for Language Models' (arXiv, 14 July 2024) <<http://arxiv.org/abs/2407.06460>> accessed 6 May 2025.

³ Li and others (n 2).

⁴ Chen and Yang (n 1); Ruiqi Zhang and others, 'Negative Preference Optimization: From Catastrophic Collapse to Effective Unlearning' (arXiv, 10 October 2024) <<http://arxiv.org/abs/2404.05868>> accessed 6 May 2025.

⁵ Regulation 2016/679/EU on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) 2016 (OJ L 119, 452016) Article 17.

⁶ Regulation 2024/1689 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (OJ L, 2024/1689, 1272024); Regulation 2023/2854/EU on harmonised rules on fair access to and use of data and amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act) (OJ L, 2023/2854, 22122023).

where unlearning fits within the EU’s regulatory ecosystem: is it an implied obligation under existing rights, a technical safeguard to be required under AI risk regulation, or a novel right that requires new legislative articulation?

This paper argues that unlearning is a crucial and increasingly feasible component of digital rights enforcement in AI contexts. However, its legal basis remains diffuse, split across overlapping and sometimes inconsistent regulatory frameworks. To clarify this ambiguity and build a coherent regulatory approach, this paper addresses the following research question: Should the EU legal framework explicitly recognise machine unlearning as a data subject right, and if so, under which instrument—the GDPR, the Data Act, or the AI Act?

Through an analysis of existing EU legislation and the technical limitations of unlearning, the paper evaluates which regulatory instrument is best suited to enforce and oversee this emerging capability. It ultimately argues that while the GDPR offers the most natural legal home for unlearning obligations, it cannot function in isolation. Instead, a coordinated regulatory strategy that combines the rights-based orientation of the GDPR with the technical oversight of the AI Act and the infrastructural reforms of the Data Act is required to ensure that unlearning becomes not merely a technical ideal, but a legally actionable right in the EU’s digital future.

2. Legal Challenges in a Shifting Technological Landscape.

Unlike many areas of AI governance, the right to erasure under the General Data Protection Regulation (GDPR) was developed before robust technical solutions for unlearning were widely available.⁷ These legal frameworks envision a "goldilocks standard" of data deletion: a world where individuals can demand that all traces of their data, including its embedded effects on algorithmic models, be completely and irreversibly erased. This is also observed in the new Data Act.

However, current technical realities reveal a significant gap between this legal aspiration and what can be practically achieved. Machine learning models often entangle personal data within complex parameters, making complete erasure technically infeasible without retraining models from scratch—a process that is computationally expensive, economically inefficient, and environmentally unsustainable. In this context, a rigid insistence on perfect deletion risks rendering compliance practically impossible and disproportionately burdensome.⁸

Machine unlearning challenges the assumption embedded in many legal texts that once data is used to train a model, it becomes practically inseparable from the model's knowledge. As unlearning technologies become more viable, regulators and lawmakers must grapple with

⁷ GDPR, OJ L 119, 4.5.2016.

⁸ Shi and others (n 2); Aengus Lynch and others, ‘Eight Methods to Evaluate Robust Unlearning in LLMs’ (arXiv, 26 February 2024) <<http://arxiv.org/abs/2402.16835>> accessed 6 May 2025; Jinghan Jia and others, ‘SOUL: Unlocking the Power of Second-Order Optimization for LLM Unlearning’ (arXiv, 24 June 2024) <<http://arxiv.org/abs/2404.18239>> accessed 6 May 2025; Pratyush Maini and others, ‘TOFU: A Task of Fictitious Unlearning for LLMs’ (arXiv, 11 January 2024) <<http://arxiv.org/abs/2401.06121>> accessed 6 May 2025; Jinghan Jia and others, ‘WAGLE: Strategic Weight Attribution for Effective and Modular Unlearning in Large Language Models’ (arXiv, 12 April 2025) <<http://arxiv.org/abs/2410.17509>> accessed 6 May 2025; Ronen Eldan and Mark Russinovich, ‘Who’s Harry Potter? Approximate Unlearning in LLMs’ (arXiv, 4 October 2023) <<http://arxiv.org/abs/2310.02238>> accessed 6 May 2025.

new questions: Should individuals have the right to compel the unlearning of their data from AI systems? What constitutes adequate erasure in a machine learning context? How do we ensure verifiability and accountability in the unlearning process? These are especially pressing in areas like algorithmic decision-making, where residual influence from deleted data could perpetuate discrimination or privacy harms.⁹

In this evolving context, the law must recognise the technical feasibility of machine unlearning and develop standards and enforcement mechanisms that align with it. Doing so could bridge the gap between data protection principles and AI development, ensuring that individual rights are preserved even in the age of increasingly complex models. This may require a rethinking of regulatory language and oversight tools, emphasising technical transparency, auditability, and the dynamic lifecycle of AI systems.

3. Current Legal Frameworks and Their Adequacy

3.1. GDPR, the Right to be Forgotten, and Machine Unlearning.

The General Data Protection Regulation (GDPR) enshrines the “right to erasure” under Article 17, often referred to as the “right to be forgotten.”¹⁰ This provision allows individuals to request the deletion of their personal data when it is no longer necessary for the purposes for which it was collected, or when the individual withdraws consent. At a surface level, machine unlearning appears to be a powerful technical tool that could affect this right in the context of AI. By enabling the removal of a data subject’s information not just from storage but also from the trained model itself, unlearning could potentially extend the GDPR’s erasure principle to the model’s internal logic and outputs. However, while promising in theory, this application is not yet explicitly addressed in the GDPR’s legal language or European Data Protection Board (EDPB) guidance.

3.1.1. Legal Gaps in the Application of the Right to Erasure to AI Models.

A key legal ambiguity lies in how the right to erasure applies when data is no longer stored in an identifiable or accessible form. Instead, it has been used to shape the learned parameters of an AI model. In traditional data processing systems, deletion involves removing rows from a database or clearing a user profile. In contrast, in machine learning systems, personal data can influence a model's weights and structure in complex and often non-traceable ways. The GDPR does not currently offer a framework for interpreting what it means to “erase” data when it has been absorbed into a model’s decision-making architecture. For example, it is unclear whether model retraining, approximate unlearning, or simply ceasing the model's use in specific contexts would satisfy compliance.¹¹ However, the GDPR embeds this right within a broader framework of balancing interests and technical feasibility. Article 17(1) is subject to numerous exceptions, including freedom of expression, compliance with legal obligations,

⁹ Alessandro Acquisti, Curtis Taylor and Liad Wagman, ‘The Economics of Privacy’ 64.

¹⁰ Ambrose (n 1); Paulan Korenhof and others, ‘Timing the Right to Be Forgotten’ 30; ‘Everything You Need to Know about the “Right to Be Forgotten”’ (*GDPR.eu*, 5 November 2018) <<https://gdpr.eu/right-to-be-forgotten/>> accessed 25 February 2022; ‘The Right to Erasure or Right to Be Forgotten under the GDPR Explained and Visualized’ (*i-SCOOP*, 25 February 2022) <<https://www.i-scoop.eu/gdpr/right-erasure-right-forgotten-gdpr/>> accessed 25 February 2022.

¹¹ Shi and others (n 2); Lynch and others (n 8).

and the exercise or defence of legal claims. Crucially, Article 17(2) and (3) recognise that complete erasure may not always be immediately or entirely possible, especially when personal data has been made public or embedded in complex systems.¹² Furthermore, the GDPR does not demand absolute technical perfection. Recital 66 clarifies that the obligation to erase includes taking "reasonable steps," considering available technology and the cost of implementation, to inform other controllers of the erasure request.¹³ This establishes a framework where the right to erasure must be fulfilled to the extent that is reasonable, proportionate, and feasible, setting important interpretive boundaries that are particularly relevant for evaluating machine unlearning. The accountability principle under Article 5(2) further obliges controllers to document and justify their erasure processes, emphasising procedural robustness rather than unattainable outcomes.

Furthermore, there is a legal vacuum regarding the technical limitations and trade-offs of unlearning. Machine unlearning is still a developing science, and depending on the architecture, it may not always be feasible to surgically remove one data point without degrading the performance of the model. This raises the question of proportionality and feasibility under Article 17(1)(c) and 17(3), which provide exceptions to erasure rights when processing is necessary for public interest or when deletion is technically impossible or requires disproportionate effort. However, these clauses were drafted without the specific context of machine learning models in mind, leaving room for conflicting interpretations.

In addition, Article 25 GDPR—requiring "data protection by design and by default"—stipulates that controllers must integrate data protection measures considering the "state of the art," the costs of implementation, and the nature, scope, context, and purposes of processing. This further confirms that data deletion obligations must be understood dynamically, evolving with technical possibilities but always measured against economic, practical, and societal realities. These doctrinal elements collectively suggest that the GDPR anticipates the need for compromise solutions like the Goldilocks standard when perfect erasure is not yet technologically feasible.

3.1.2. Are Existing Laws Sufficient?

The broader question is whether the GDPR, in its current form, is flexible enough to accommodate the emerging practice of machine unlearning, or whether legislative updates or new regulatory guidance will be required. On one hand, the GDPR is designed to be technology-neutral and principles-based, which offers some adaptability. Recitals 26 and 78, for instance, emphasise data minimisation, accountability, and privacy by design—principles that could be interpreted to support the deployment of unlearning where feasible. On the other hand, the absence of specific provisions or interpretative guidance on machine learning systems and their interaction with rights like erasure leaves a gap in enforcement. Supervisory authorities may lack the tools or technical frameworks to verify whether unlearning has taken place or whether residual data influence persists.

In sum, while the GDPR contains foundational principles that are highly relevant to the challenges of AI and data retention, it was not crafted with machine unlearning in mind. To

¹² GDPR, OJ L 119, 4.5.2016 Article 17(2) and (3).

¹³ *ibid.*

fully operationalise the right to be forgotten in the age of AI, EU regulators may need to issue targeted guidance, expand accountability mechanisms, and possibly amend legal texts to explicitly include model-level erasure as a component of data protection rights. Without such efforts, there is a risk that machine unlearning remains a theoretical capability. At the same time, the legal system continues to operate on assumptions that no longer reflect the technical realities of AI.

3.2. Machine Unlearning and the EU Data Act: Future-Proofing Data Governance?

The Data Act, adopted in 2024, extends data governance obligations beyond personal data to shared industrial, IoT, and user-generated data contexts. While it primarily concerns data access, interoperability, and portability, it also introduces requirements around the secure deletion of data, especially at the end of contractual relationships or upon user request. Article 4(10) of the final compromise text emphasises that data holders must ensure that "personal data are erased or anonymised," aligning with GDPR principles but acknowledging that technical and economic considerations play a role. Importantly, Recital 50 of the Data Act explicitly frames these deletion obligations in light of "the state of the art" and "proportionality in costs and burden," mirroring GDPR's flexibility.

The Data Act thus reinforces the idea that perfect deletion is not an absolute mandate; instead, controllers and data holders must make reasonable, demonstrable efforts, balancing technical feasibility, security, and economic proportionality. This is highly relevant for machine unlearning contexts, where the complete retraining of models to erase a single user's influence could be economically ruinous or environmentally unsustainable. In this legal environment, approaches like the Goldilocks standard—where deletion is sufficiently effective without demanding absolute erasure—are not only legally permissible but actively encouraged by the spirit of EU data governance law.¹⁴

The EU Data Act plays a central role in reshaping the European data economy by establishing more explicit rules for access, use, and sharing of both personal and non-personal data. It aims to enhance data availability while preserving fundamental rights, particularly in a context where data is increasingly integrated into AI and machine learning systems. While the GDPR provides individuals with the right to erasure of personal data, the Data Act introduces additional obligations and mechanisms related to the sharing and re-use of data, often in non-personal or industrial contexts. However, the boundaries between personal and non-personal data are increasingly blurred in machine learning applications, raising complex questions around data control, erasure, and technical feasibility.

Machine unlearning—understood as the technical process of removing the influence of specific data points from trained models—sits at the intersection of this evolving legal landscape.¹⁵ While unlearning is not explicitly regulated by the Data Act, the Act's provisions on access, control, and data sharing create new legal and operational implications for how

¹⁴ 'Data Governance Act- BEUC Position Paper' (The European Consumer Organisation- BEUC 2021) BEUC-X-2021-026; Ayush K Tarun and others, 'Fast Yet Effective Machine Unlearning' (2024) 35 IEEE Transactions on Neural Networks and Learning Systems 13046.

¹⁵ Li and others (n 2); Tarun and others (n 14); Zheyuan Liu and others, 'Machine Unlearning in Generative AI: A Survey' (arXiv, 30 July 2024) <<http://arxiv.org/abs/2407.20516>> accessed 6 May 2025.

data used in AI systems is governed, especially when the same data is accessed by multiple actors or reused across different contexts. These implications become particularly salient when machine learning models are built collaboratively or trained on shared data assets, as reversing the influence of a specific dataset may be legally required but technically difficult or economically costly.

3.2.1. Data Sharing and Control

The Data Act introduces robust rules for ensuring fair access to data generated by connected devices, industrial platforms, and services. It mandates that users—whether individuals or businesses—should have access to data they help generate and, in certain cases, be able to share it with third parties. This enhanced access regime is central to the EU’s vision of a competitive, interoperable digital single market. However, from the standpoint of machine unlearning, the decentralisation of data use and the proliferation of data recipients introduce significant challenges.

Once data has been shared with third parties, especially if used to train machine learning models, the ability to effectuate a data subject's right to erasure or request for unlearning becomes complicated. The Data Act does not provide detailed mechanisms to ensure traceability or reversibility of data influence across complex AI supply chains. This raises critical questions: Can a data subject reasonably expect their data to be "unlearned" from a model that has been trained by a third-party recipient of shared data? What obligations, if any, do data recipients have to accommodate such requests? These issues become even more opaque when the data in question is de-identified, anonymised, or aggregated before use. While such processing may appear to circumvent GDPR obligations, it does not necessarily eliminate the relevance of unlearning when these data points influence automated decision-making.

Furthermore, the Act’s emphasis on data portability and access for users must be reconciled with privacy rights, particularly the right to erasure.¹⁶ This creates a potential legal tension: promoting broader access and re-use of data may unintentionally limit the enforceability of rights that depend on the ability to control or retract data after dissemination. Addressing this tension will require a more nuanced legal framework that recognises the lifecycle of data use in AI models and the varying degrees of reversibility at each stage.

3.2.2. Unlearning and Data Access

Another core ambition of the Data Act is to make data more accessible and reusable by businesses, public authorities, and research institutions. It provides a structured framework for mandated data sharing under specific circumstances, such as emergencies or when public interest is involved. However, this push toward open access must be balanced against the rights of individuals and entities whose data is used in these processes, especially in AI systems where data contributes to model development in complex, non-transparent ways.

Machine unlearning introduces a critical friction point in this context. Unlearning is not simply a matter of deleting stored data—it involves identifying and eliminating the influence of that

¹⁶ Nicholas Carlini and others, ‘The Privacy Onion Effect: Memorization Is Relative’ (arXiv, 22 June 2022) <<http://arxiv.org/abs/2206.10469>> accessed 6 May 2025.

data on the structure and behaviour of a trained model. When data has been shared across organisational or even national boundaries and used in models trained by entities far removed from the original data collector, implementing unlearning becomes legally and technically burdensome. This is further complicated by the Data Act's recognition of the need to protect trade secrets and intellectual property. If data contributes to the competitive advantage of a business, either directly or by shaping a proprietary AI model, legal requests for unlearning could conflict with the protected interests of that business. The Act permits data holders to refuse access requests where trade secrets would be compromised, but it is less clear how this applies in reverse: Can a right-holder demand data unlearning if it affects protected trade interests?

Cross-border data flows and multi-stakeholder AI ecosystems only exacerbate this problem. For instance, if data is used in federated AI architectures or in models trained across EU and non-EU jurisdictions, what legal responsibilities apply to each actor in ensuring data erasure or unlearning? The technical feasibility of ensuring deletion or unlearning across distributed systems may vary, but the legal framework must still ensure consistent enforcement of data rights.

In light of these complexities, the Data Act's objectives to maximise data access and utility must be reconciled with an emerging need for legally enforceable and technically feasible unlearning protocols. This will likely require future regulatory clarifications and possibly harmonisation with existing data protection instruments such as the GDPR, as well as emerging AI-specific regulation under the AI Act. Without such alignment, machine unlearning risks being trapped between ambitious legal standards and insufficient technological readiness, especially in data ecosystems shaped by the European legal digital environment's expansive data sharing provisions.

3.2.3. Data Portability

The Data Act emphasises data portability as a cornerstone of the EU's effort to empower users and foster competition in the data economy.¹⁷ Much like Article 20 of the GDPR, which allows individuals to obtain and transfer their personal data to another controller, the Data Act extends this principle to non-personal data and data generated by connected products and services. The goal is to reduce vendor lock-in and promote data-driven innovation across platforms and sectors. However, the right to data portability is inherently forward-looking—it facilitates data movement and access, not its removal or erasure.

This limitation is particularly salient in the context of machine unlearning. While the ability to move data is essential for user autonomy, portability does not resolve the issue of how data already ingested and processed by AI systems should be retrospectively removed or "unlearned." AI models do not merely store data; they abstract, transform, and embed it into learned representations. Once a model is trained, the original data's influence may be distributed across multiple layers, making simple deletion technically infeasible without retraining or redesign. The Data Act, in its current form, does not mandate that platforms provide mechanisms for reverse inference or influence tracking—both of which are necessary for unlearning. This creates a legal and technical gap, leaving the right to erasure under the GDPR partially unenforceable in AI contexts and unaddressed by the Data Act's portability provisions.

¹⁷ G Zanfir, 'The Right to Data Portability in the Context of the EU Data Protection Reform' (2012) 2 International Data Privacy Law 149.

3.2.4. Interoperability

The Data Act also promotes interoperability, particularly through the development of common standards and open protocols that allow systems to exchange and use data effectively. This ambition aligns with broader EU initiatives on digital sovereignty and competitiveness. However, when viewed through the lens of machine unlearning, interoperability raises unique technical and legal questions. Interoperability frameworks typically focus on enabling data sharing and integration, but they do not necessarily ensure the reversibility of data influence across systems.

For machine unlearning to be meaningful, AI systems must be compatible with unlearning mechanisms, such as influence functions, model patching, or retraining pipelines, that can identify and remove the contribution of specific data points. If such systems are built using divergent or proprietary architectures, the implementation of unlearning across platforms becomes inconsistent or even infeasible. While the Data Act encourages interoperability in principle, it remains silent on whether this interoperability must include compatibility with data minimisation, erasure, or unlearning techniques. This silence creates uncertainty: will future interoperability standards require AI developers to build in unlearning capacity, or will this responsibility be left to market actors under principles of industry self-regulation? Without legal clarity, there's a risk that unlearning remains a theoretical right rather than a practically enforceable obligation.

3.2.5. Responsibility for Data Handling.

One of the important contributions of the Data Act is the clarification of roles and responsibilities in the data economy, particularly in multi-party environments where data is shared across actors. The Act assigns duties to data holders, data recipients, and third-party service providers, especially in cases of non-compliance, misuse, or breach. However, machine unlearning introduces new ambiguities around liability. If a data subject invokes their right to erasure, but the data has already been used to train a model, who is responsible for implementing unlearning, or for the failure to do so?

The issue is further complicated by the nature of AI development, which often involves multiple stakeholders: original data collectors, cloud infrastructure providers, AI model developers, and downstream users of trained models. In such distributed environments, there is no clear consensus on who holds the technical or legal responsibility for ensuring that data is forgotten. The Data Act provides liability mechanisms for data misuse but does not directly address responsibility in cases of residual influence, where the data is technically deleted but continues to shape algorithmic outcomes. Clarifying liability in such cases is critical. Without clear accountability mechanisms, data subjects may find their rights unenforceable, while companies face legal uncertainty and risk over unanticipated obligations.

3.2.6. Data Protection vs. Commercial Use.

The Data Act attempts to balance individual rights with the EU's ambition to harness the value of data for economic growth. This balancing act becomes especially delicate when machine unlearning conflicts with commercial interests. If data used to train a model must later be

unlearned, the integrity and utility of the model may be compromised. This can directly impact business operations, especially in sectors that rely heavily on AI, such as finance, healthcare, and mobility.

The core legal challenge here is reconciling the right to erasure (under the GDPR) with the legitimate interest in data retention for model development (as promoted by the Data Act). While the GDPR provides a set of exceptions, such as overriding public interest or legal obligations, these do not map neatly onto the operational requirements of AI-driven enterprises. The Data Act could serve as a bridging instrument by specifying how and when commercial interests can justify retaining data in models, and when they must yield to data protection rights. At present, however, the law does not offer this guidance, leaving developers and regulators alike in a zone of interpretive uncertainty.

3.2.7. Incorporating Unlearning into Data Governance: Is the Data Act Prepared for Emerging Technologies?

The Data Act, as part of the EU's broader strategy to govern the digital economy, introduces forward-looking provisions on data access, usage, and control. However, it stops short of addressing the lifecycle of data within AI systems. In particular, the Act does not engage with how data influence should be managed post-training, nor does it incorporate machine unlearning into its conception of responsible data governance. As AI becomes more deeply embedded into infrastructure and decision-making, the absence of such mechanisms becomes increasingly problematic.

A future-ready data governance regime should recognise that deletion in AI contexts requires more than database-level removal—it demands model-level corrections. Incorporating unlearning into the Data Act could involve mandating transparent documentation of model training data, auditing mechanisms to track data influence, and minimum technical standards for reversible data processing. These changes would signal to industry actors that data governance extends beyond access and sharing to include downstream influence and compliance with evolving rights-based obligations.

4. Potential Reforms

To align with the challenges posed by machine unlearning, the Data Act could undergo targeted reforms. One potential area of reform is the explicit regulation of AI developers and service providers as distinct actors with obligations related to data deletion and influence mitigation. Provisions could require providers to disclose whether their systems support unlearning and, if not, to justify such limitations under a proportionality test. Additionally, the Act could incentivise the adoption of technical solutions that facilitate unlearning, such as modular architectures or influence-tracking frameworks, through public procurement preferences or regulatory sandboxes.

Another reform pathway is the clarification of data controller and processor responsibilities in the context of AI. The Data Act could articulate how these roles apply when models are trained on shared or pooled data and specify unlearning obligations across the AI value chain. This would help ensure that data protection is not diluted simply because data has been transformed into model parameters.

4.1. Towards a More Future-Proof Law

As part of the European Commission's broader digital strategy, the Data Act aims to enable innovation while reinforcing fundamental rights. However, its silence on machine unlearning reveals a critical gap in its ability to future-proof data governance. Much like the Data Governance Act, which addresses trust, security, and access in data sharing, the Data Act could include provisions for how data should be removed from models once its lawful basis expires. This would prevent the creation of irreversible data dependencies and ensure that AI innovation does not outpace accountability.

Including unlearning in the legal text would also facilitate the creation of technical standards and certification schemes, allowing organisations to demonstrate compliance and foster trust. Without these measures, unlearning remains largely aspirational—a noble concept without concrete institutional support.

4.2. The Need for Coordination Between Data Protection Laws and AI Regulation.

Addressing machine unlearning requires coordination across the EU's regulatory architecture. The GDPR, the AI Act, and the Data Act each govern different aspects of data and AI, but their interplay remains underdeveloped. For instance, while the GDPR enshrines the right to erasure, and the AI Act introduces transparency and risk mitigation requirements for high-risk AI systems, neither directly explains how unlearning must be implemented or verified. The Data Act has the opportunity to fill this gap by linking rights, responsibilities, and technical feasibility into a coherent governance framework.

A coordinated approach could involve shared definitions, joint compliance mechanisms, and regulatory guidance clarifying how rights to unlearning intersect with AI lifecycle management. Without this, regulatory fragmentation may lead to inconsistent enforcement and gaps in protection, especially for individuals whose data powers opaque and decentralised AI systems.¹⁸

As AI systems become increasingly central to decision-making processes, the legal system must evolve to ensure individual rights are enforceable even in technically complex environments. One key reform proposal would be to explicitly integrate machine unlearning obligations into existing data protection legislation, most notably the GDPR. Currently, Article 17 of the GDPR establishes the "right to be forgotten," which allows individuals to request the deletion of personal data under specific circumstances. However, the GDPR does not explicitly require data controllers to remove that data from trained AI models, nor does it provide guidance on how such erasure should be implemented technically.

To address this, future legislative updates could introduce a specific provision mandating that data controllers implement unlearning mechanisms where technically feasible, or provide detailed explanations if unlearning is not possible. This requirement could be supported by transparency obligations, such as mandating that controllers disclose whether and how personal data influences automated systems, and whether those systems allow for erasure after training. A new article or recital could be added to the GDPR acknowledging the complexity of machine learning systems while establishing that the right to be forgotten

¹⁸ Bill Marino, Meghdad Kurmanji and Nicholas D Lane, 'Bridge the Gaps between Machine Unlearning and AI Regulation' (arXiv, 18 February 2025) <<http://arxiv.org/abs/2502.12430>> accessed 6 May 2025.

includes data embedded in AI models, thereby bridging the gap between legal rights and technical practice.

4.3. AI Governance.

The EU AI Act represents another regulatory frontier where machine unlearning could be formalised. The Act already introduces differentiated risk-based obligations for AI systems, including requirements for transparency, human oversight, and risk mitigation for "high-risk" AI applications. Given that such systems are often trained on sensitive or consequential data, ranging from health records to employment data, there is a strong case for embedding unlearning requirements in the governance of high-risk AI.¹⁹

Unlearning could be treated as a mandated safeguard under the risk mitigation and data governance provisions of the AI Act. For instance, model developers could be required to demonstrate that they can delete or mitigate the influence of individual data points, particularly when such data is processed without an adequate legal basis or becomes outdated or erroneous. As the AI Act evolves, introducing performance metrics or certification schemes around unlearning could help operationalise this requirement. If adopted, this would also set a global precedent for AI accountability and harmonise AI regulation with existing data protection law.

5. Conclusion: How “Future-Proof” Is the Law in Relation to Machine Unlearning?

5.1. Assessment of Current Laws.

Despite the growing prominence of AI in data processing, current legal frameworks remain ill-equipped to handle the technical and ethical complexities of machine unlearning. The GDPR lays down a strong foundation for individual data rights, particularly through the “right to be forgotten” under Article 17. However, the GDPR was drafted before the widespread use of machine learning technologies and does not directly address how this right should be enforced in systems where data is embedded in trained AI models. In such cases, simply deleting personal data from storage may not remove its influence from models that have already been trained, especially in deep learning systems where individual data points are not explicitly stored but abstracted into complex weight matrices.

Similarly, the EU AI Act represents a forward-looking, risk-based regulatory approach to the governance of AI systems. Yet, it does not currently specify how the right to erasure or data minimisation principles should be applied post-training. High-risk AI systems, which have significant impacts on individuals’ rights and safety, may still retain the influence of data long after it has been deleted, creating legal ambiguities.²⁰ Additionally, the Data Act—though ambitious in its goal to promote data sharing and user control—remains largely silent on the lifecycle of data within machine learning models. Without explicit provisions on post-processing and unlearning, data subjects remain vulnerable to residual harm from data they no longer consent to being used.

¹⁹ Liu and others (n 15).

²⁰ Jakub Łucki and others, ‘An Adversarial Perspective on Machine Unlearning for AI Safety’ (arXiv, 10 April 2025) <<http://arxiv.org/abs/2409.18025>> accessed 6 May 2025.

The core issue is that the legal enforceability of the right to erasure becomes murky in AI contexts. Non-linear and opaque models, such as large language models or recommender systems, make it difficult to trace or quantify the influence of a particular data point. This technological opacity challenges the capacity of existing laws to provide meaningful redress, leading to a growing gap between the promise of legal rights and their realisation in AI-driven environments.

5.2. Call for Reform

To ensure that legal frameworks remain robust and responsive in the face of rapid AI advancement, legislators must proactively incorporate machine unlearning into the architecture of data and AI governance. A central reform would be the explicit legal recognition of unlearning obligations within key regulations. Amendments to the GDPR, EU AI Act, and Data Act could clarify that the right to be forgotten extends to the removal of personal data influence from trained models, and impose duties on data controllers to ensure that unlearning processes are part of their compliance practices.

In tandem with legal recognition, there is a need for the development of technical standards that define and support machine unlearning. Legislators and standardisation bodies should invest in research that identifies feasible unlearning methods across different model architectures, such as parameter pruning, model reweighting, and influence functions, and establish modular, auditable system designs that can facilitate unlearning requests. This would not only make the process more tractable but also ensure that it meets regulatory requirements for proportionality and accountability.

Transparency requirements must also evolve to address the data lifecycle in AI systems. Regulations could mandate that AI developers disclose the influence of specific data points on system outputs, the feasibility of unlearning in their systems, and the mechanisms they use to track data provenance. Such disclosures would empower data subjects and regulators alike to assess whether rights are being respected and provide a foundation for accountability in enforcement actions.

To address ethical risks, fairness safeguards should be embedded into the legal framework for unlearning. There is a possibility that frequent unlearning of certain data subsets, especially from marginalised groups, may degrade model performance in ways that inadvertently introduce bias. Regular bias assessments and fairness audits should be conducted both before and after unlearning operations, ensuring that privacy rights are not exercised at the cost of discriminatory outcomes.

Finally, the cross-border nature of data flows necessitates international coordination. Multilateral frameworks or treaties, potentially developed under the auspices of the OECD, UNESCO, or the United Nations, could provide consistent standards for unlearning. These instruments should set minimum thresholds for technical feasibility, transparency, and procedural fairness, while encouraging mutual recognition of compliance mechanisms. Harmonising global practices would reduce regulatory fragmentation and support the development of interoperable AI systems that respect data rights across jurisdictions.

Flexibility must underpin these reforms. Laws should be adaptive enough to accommodate emerging technologies and evolving AI capabilities. Regulatory sandboxes, where new methods of unlearning can be tested in a controlled environment, could help accelerate innovation while ensuring legal compliance. Additionally, dynamic governance tools—such as regularly updated technical guidelines or sector-specific codes of practice—would ensure that the legal framework remains current without requiring constant legislative overhaul.

In conclusion, current data and AI laws are not yet “future-proof” when it comes to machine unlearning. However, with targeted reforms that align technical feasibility with legal enforceability, the law can evolve to support unlearning as a cornerstone of ethical and accountable AI. By embedding flexibility, transparency, and fairness into legal mandates, regulators can ensure that machine unlearning becomes a tool for empowering individuals, not a loophole for evading responsibility.