

Artificial Intelligence and the Right to Privacy: Legal Challenges in the Protection of Personal Data

Dr.Nitesh Kumar Srivastava¹

Abstract:

The rapid diffusion of artificial intelligence (AI) systems across commerce, government, and daily life has intensified long-standing tensions between technological innovation and the right to privacy. AI systems depend on the large-scale collection, aggregation, and inferential processing of personal data, often in ways that strain the foundational principles of data protection law, including purpose limitation, data minimisation, consent, and transparency. This paper examines the legal challenges that artificial intelligence poses to the protection of personal data, drawing on comparative analysis of the European Union's General Data Protection Regulation and Artificial Intelligence Act, the sectoral and state-based framework of the United States, and emerging approaches in other jurisdictions. It identifies recurring points of friction, including the opacity of algorithmic decision-making, the difficulty of applying consent-based models to machine learning training, the phenomenon of inference of sensitive attributes from ostensibly non-sensitive data, the tension between data minimisation and data-hungry model development, and the practical impossibility of full erasure once personal data has been absorbed into a trained model. The paper further considers enforcement fragmentation between data protection authorities and AI-specific regulators, cross-border data transfer complications, and the growing use of biometric and profiling technologies. Through analysis of statutory text, regulatory guidance, and enforcement trends current as of mid-2026, the paper argues that existing privacy frameworks, though adaptable in principle, require targeted reform, technical standard-setting, and closer institutional coordination to remain effective against AI-driven risks to personal data. The paper concludes with policy recommendations aimed at reconciling innovation incentives with the enduring legal and normative value of privacy as a fundamental right.

Keywords: artificial intelligence, data protection, privacy law, GDPR, AI Act, algorithmic accountability, automated decision-making, data minimisation, cross-border data transfer, biometric data

1. Introduction:

Privacy has long been recognised as a fundamental legal interest, protected under international instruments such as Article 12 of the Universal Declaration of Human Rights and Article 17 of the International Covenant on Civil and Political Rights, and given concrete regulatory expression through domestic data protection statutes across the globe.

The emergence of artificial intelligence as a general-purpose technology has transformed the scale, speed, and sophistication with which personal data can be collected, combined, and analysed. Machine learning systems, and large-scale generative models in particular, are trained on datasets that frequently include personal information scraped from public and semi-public sources, generating novel legal questions that traditional data protection law was not designed to answer. Regulators worldwide have responded with a wave of new instruments, most notably the European Union's Artificial Intelligence Act, which reached full applicability for high-risk systems in August 2026 and layers a second, AI-specific compliance and penalty structure on top of the General Data Protection Regulation.²

This paper addresses the central research question: to what extent do existing legal frameworks for privacy and data protection adequately address the distinctive risks that artificial intelligence poses to personal data, and what reforms are needed to close the resulting gaps? The analysis proceeds in six further parts. Part 2 sets out the conceptual relationship between privacy and data protection as it applies to AI. Part 3 surveys the principal legal frameworks in the European Union, the United States, and other representative jurisdictions. Part 4 identifies the specific legal challenges

¹ Assistant Professor, Department of Law, Koneru Lakshmaiah Educational Foundation, University, Vijayawada, Andhra Pradesh, India.

² See EU AI Act, Regulation (EU) 2024/1689; full enforcement obligations for high-risk AI systems became applicable on 2 August 2026, introducing penalties of up to €35 million or 7% of global annual turnover, a materially higher ceiling than the GDPR's own maximum of €20 million or 4% of global turnover.

that AI poses to personal data protection. Part 5 discusses illustrative enforcement actions and case law. Part 6 offers recommendations, and Part 7 concludes.

2. Conceptual Framework: Privacy, Data Protection, and Artificial Intelligence:

Privacy and data protection, while closely related, are conceptually distinct. Privacy is generally understood as a right to be free from unwarranted intrusion into one's personal life, decisions, and identity, whereas data protection is a narrower, procedural regime governing how personal data is collected, processed, stored, and shared. The Court of Justice of the European Union treats the two as related but separate fundamental rights, with privacy protected under Article 7 of the EU Charter and data protection under Article 8.³ Artificial intelligence complicates this distinction because AI systems do not merely store or transmit personal data; they transform it, extracting statistical patterns from training data to generate predictions or content about individuals never directly observed by the system. This implicates two categories of data: personal data used to train a model, and personal data a deployed model infers about a data subject during use. Commentators argue that such inferred data deserves protection comparable to directly collected data, since profiling can reveal sensitive characteristics even where those characteristics were never disclosed.⁴

Data protection law traditionally rests on principles of lawfulness, fairness and transparency, purpose limitation, data minimisation, accuracy, storage limitation, integrity and confidentiality, and accountability. These principles, codified in Article 5 of the GDPR, presuppose a linear data lifecycle in which data is collected for a defined purpose, processed accordingly, and eventually deleted.⁵ AI development departs from that model: training data is repurposed, aggregated at scale, and embedded within model parameters in ways that are difficult to isolate, correct, or delete.

3. Legal Frameworks Governing Artificial Intelligence and Privacy:

3.1 The European Union: GDPR and the AI Act

The European Union combines the horizontally applicable GDPR with the sector-spanning AI Act. The GDPR applies to any processing of personal data by controllers or processors established in, or targeting individuals in, the EU, and grants data subjects rights of access, rectification, erasure, restriction, portability, and objection, together with protections against solely automated decisions producing legal or similarly significant effects.⁶ The AI Act, formally Regulation (EU) 2024/1689, adopts a risk-based structure that prohibits certain practices outright, imposes heightened obligations on 'high-risk' systems, and applies lighter transparency duties to general-purpose AI models.⁷ Prohibited practices, applicable since February 2025, include harmful subliminal manipulation, exploitation of vulnerable groups, public-authority social scoring, and, with narrow exceptions, real-time remote biometric identification in public spaces. Obligations for general-purpose AI providers became applicable in August 2025, and the bulk of high-risk system obligations reached full enforcement on 2 August 2026.⁸

The two instruments are enforced by different authorities — data protection authorities for the GDPR, and national market-surveillance authorities together with the European AI Office for the AI Act — creating a risk of divergent interpretation. The EDPB and EDPS have issued joint opinions attempting to reconcile the regimes, including opinions in January and February 2026 on the Commission's Digital Omnibus proposals, which would revisit the definition of personal data and the treatment of pseudonymised training data; the Authorities have opposed changes that would narrow that definition.⁹ The AI Act reinforces GDPR Article 9 special-category protections by prohibiting workplace and

³ Charter of Fundamental Rights of the European Union, 2012/C 326/02, Arts. 7–8; see also Cases C-293/12 and C-594/12, *Digital Rights Ireland* (CJEU, 8 April 2014).

⁴ See generally S. Wachter and B. Mittelstadt, 'A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI', *Columbia Business Law Review* (2019).

⁵ Regulation (EU) 2016/679 (General Data Protection Regulation), Art. 5.

⁶ GDPR, Arts. 3, 15–22.

⁷ Regulation (EU) 2024/1689 (Artificial Intelligence Act), Arts. 5–6 and Annex III.

⁸ TechGDPR, 'Data protection digest 1 Aug 2026: EU AI Act becomes applicable, with high-risk AI temporarily excluded' (Aug. 2026).

⁹ EDPB–EDPS, Joint Opinion on the Digital Omnibus on AI (20 January 2026); EDPB–EDPS, Joint Opinion on the Digital Omnibus Regulation (11 February 2026).

educational emotion-recognition and tightly restricting biometric categorisation, though AI systems can still infer sensitive attributes indirectly from neutral proxy data such as purchasing patterns.¹⁰

3.2 The United States: A Sectoral and State-Based Patchwork

The United States lacks a single comprehensive federal privacy or AI statute, relying instead on sector-specific federal laws, general consumer-protection enforcement, and an expanding patchwork of state legislation. The Federal Trade Commission relies principally on Section 5 of the FTC Act to police unfair or deceptive AI-related practices, including 'AI washing', undisclosed AI tools, and data practices tied to automated decision-making.¹¹ By early 2026, nineteen states had comprehensive consumer privacy laws in effect, alongside targeted AI statutes such as the Colorado AI Act, which requires developers and deployers of high-risk AI systems to exercise reasonable care against algorithmic discrimination and to conduct impact assessments; California, Texas, Illinois, and Utah have enacted their own AI transparency, biometric-privacy, or automated-decision-making rules.¹²

This state-level activity has generated friction with federal policy: a December 2025 executive order favoured a single national AI framework over what it characterised as fifty discordant state regimes, directing a DOJ task force to challenge state AI laws and the FTC to address preemption of state rules that compel alteration of truthful AI outputs.¹³ As of mid-2026, no statute has preempted state AI or privacy law, and existing state laws remain enforceable pending further legislative or judicial action. Biometric privacy litigation has been especially active under statutes like the Illinois Biometric Information Privacy Act, and Congress has enacted the TAKE IT DOWN Act, requiring platforms to remove non-consensual intimate imagery, including AI-generated deepfakes, within forty-eight hours of a report, with FTC enforcement of the notice-and-removal process running since May 2026.¹⁴

3.3 Other Jurisdictions

A growing number of jurisdictions have adopted dedicated AI governance instruments layered atop existing data protection statutes. China's Personal Information Protection Law and algorithm-specific regulations impose registration, transparency, and security-assessment obligations on recommendation and generative-AI systems. India's Digital Personal Data Protection Act establishes consent-based processing and data-fiduciary duties applicable to AI systems, though India has not yet enacted a standalone AI statute. Brazil, South Korea, Canada, and the United Kingdom have pursued variants of a risk-based approach, and bodies such as the OECD and Council of Europe promote non-binding principles intended to harmonise minimum standards, though this proliferation still raises the prospect of conflicting substantive obligations for transnational AI developers.

4. Key Legal Challenges in Protecting Personal Data from AI:

4.1 Consent, Purpose Limitation, and Automated Decision-Making

Consent-based processing is poorly suited to AI training pipelines, since large-scale models are frequently trained on data scraped from the open web, where obtaining specific, informed consent from every data subject is practically infeasible. Controllers often rely on legitimate interest instead, but regulators require a documented balancing test against data subjects' rights, and this basis cannot justify indiscriminate scraping for model training.¹⁵ Purpose limitation is similarly strained when a model trained for one purpose is later fine-tuned for another, since the embedded personal data was never processed with that secondary purpose in view. Separately, Article 22 of the GDPR restricts decisions based solely on automated processing that produce legal or similarly significant effects, subject to a meaningful-human-oversight requirement, but the line between genuine human review and a nominal rubber stamp is frequently contested,

¹⁰ AI Act, Art. 5(1)(f)–(h).

¹¹ 15 U.S.C. § 45; Morgan Lewis, 'AI Enforcement Accelerates as Federal Policy Stalls and States Step In' (2 April 2026).

¹² Kiteworks, 'GDPR Fines Hit €7.1 Billion: Data Privacy Enforcement Trends in 2026' (25 March 2026); Colorado AI Act (SB 24-205), reenacted by SB 26-189 (14 May 2026).

¹³ Executive Order 14365 (11 December 2025); FTC, Policy Statement Concerning the Suppression of Accuracy in Artificial Intelligence Systems, 91 Fed. Reg. (7 July 2026).

¹⁴ 740 ILCS 14/1 et seq.; TAKE IT DOWN Act, Pub. L. No. 119-12 (19 May 2025); Baker Botts, 'U.S. Artificial Intelligence Law Update' (January 2026).

¹⁵ GDPR, Art. 6(1)(f); see EDPB, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models (adopted 17 December 2024).

and the limited interpretability of many deep-learning systems makes it difficult to give data subjects meaningful information about the logic involved.¹⁶

4.2 Inference of Sensitive Data and the Data-Minimisation Problem

AI systems can infer sensitive attributes — health status, sexual orientation, political opinion — from ostensibly neutral proxy data, creating a structural gap in a special-category regime under Article 9 of the GDPR that is triggered by the processing of data revealing a protected characteristic rather than by inference. The Court of Justice of the European Union has confirmed, in the context of automated creditworthiness scoring, that a probabilistic value derived from personal data can itself constitute personal data subject to the Regulation's protections.¹⁷ This tension compounds the broader conflict between data minimisation and the data-hungry paradigm of machine learning, where model performance generally improves with larger training datasets. Privacy-enhancing techniques such as federated learning, differential privacy, and synthetic data offer partial mitigation, but regulators caution that anonymisation claims require rigorous technical validation, since models trained on identifiable data rarely achieve genuine anonymisation.¹⁸

4.3 Transparency, Trade Secrets, and the Right to Erasure

Meaningful transparency about AI processing is complicated by the tension between disclosure obligations and protection of trade secrets and proprietary model architecture, and this tension has made transparency a top enforcement priority under the European Data Protection Board's 2026 Coordinated Enforcement Action.¹⁹ The right to erasure raises a distinct problem once personal data has been used to train a model: unlike a database record, data absorbed into model parameters is diffused throughout the model and cannot easily be deleted in isolation. Full retraining is often the only reliable route to complete erasure but is computationally expensive; emerging 'machine unlearning' techniques attempt to approximate erasure without retraining, though no regulator has yet settled a compliance standard for erasure requests directed at a trained model rather than a database.

4.4 Cross-Border Transfers and Biometric Data

AI supply chains routinely span multiple jurisdictions — data collection in one country, training in another, deployment in a third — amplifying exposure under transfer restrictions such as those in the GDPR. The EU–U.S. Data Privacy Framework, which underpins many such transfers, remains under legal challenge; the EU General Court dismissed one challenge in September 2025, with an appeal pending before the CJEU as of mid-2026.²⁰ Biometric identification raises acute concerns because identifiers are immutable and uniquely personal: the AI Act prohibits real-time remote biometric identification in public spaces for law enforcement subject only to narrow exceptions, and several U.S. states have added neural and biological data to their statutory definitions of sensitive data.²¹

4.5 Institutional and Enforcement Fragmentation

Data protection authorities, AI market-surveillance authorities, and consumer-protection regulators frequently assert overlapping jurisdiction over the same AI deployment, with limited formal coordination mechanisms. Complaint volumes reaching European data protection authorities have grown substantially as a consequence of AI adoption, prompting calls for enhanced cooperation among supervisory authorities.²² This fragmentation is compounded internationally by the absence of a single global standard, leaving multinational AI developers to reconcile materially different rights-based, sectoral, and risk-based regulatory models.

¹⁶GDPR, Art. 22.

¹⁷Case C-634/21, OQ v. Land Hessen (CJEU, 7 December 2023) (SCHUFA).

¹⁸Secure Privacy, 'GDPR Compliance in 2026: The Complete Guide' (2026).

¹⁹European Data Protection Board, 2026 Coordinated Enforcement Framework Action on Transparency, as reported in Inside Privacy, 'What to Watch in 2026' (28 January 2026).

²⁰Case T-553/23 (EU General Court, 3 September 2025), on appeal as Case C-703/25 P (CJEU, filed 31 October 2025).

²¹AI Act, Art. 5(1)(h); Nixon Peabody, 'Data Privacy, Cybersecurity, AI developments shaping 2026' (9 February 2026).

²²TechGDPR, 'Data protection digest 1 Aug 2026', reporting EDPB discussion of rising complaint volumes attributable to AI.

5. Illustrative Enforcement Actions and Case Law:

The Italian Data Protection Authority's 2023 temporary restriction on a major generative AI chatbot, imposed over the lawful basis for training-data processing, the absence of age verification, and the accuracy of outputs about identifiable individuals, remains a foundational precedent for applying existing GDPR powers to generative AI.²³ In the United States, FTC actions against Rite Aid, over AI facial recognition deployed for loss prevention without safeguards against discriminatory misidentification, and against Rytr, over an AI writing tool alleged to facilitate deceptive reviews, illustrate reliance on general unfairness and deception authority absent a comprehensive federal AI statute.²⁴ The CJEU's SCHUFA decision, discussed above, confirmed that an automated probability score itself constitutes personal data and can trigger Article 22 protections even where formally only one input into a third party's decision, with direct implications for AI-driven credit-scoring and hiring systems.²⁵ Dark-pattern enforcement against consent interfaces used by AI-enabled advertising platforms has also intensified, reflecting recognition that formally available privacy controls have limited value if their design nudges users away from exercising them.

6. Towards Solutions: Policy and Legal Recommendations:

Several reforms would help close the gaps identified above. Lawmakers should clarify the lawful basis for training AI models on personal data, whether through a dedicated statutory basis or binding guidance on the legitimate-interest balancing test, reducing reliance on ad hoc interpretation. Regulators should endorse technical standards for machine unlearning and training-data provenance documentation, enabling meaningful compliance with erasure and rectification rights without disproportionate full retraining. Data protection and AI-specific regulators should build on the EDPB–EDPS joint-opinion model to establish binding coordination mechanisms that prevent duplicative investigations and conflicting determinations. The treatment of inferred data should be clarified across jurisdictions, extending accuracy, contestability, and special-category protections to attributes an AI system predicts, not only to attributes an individual discloses. Cross-border transfer regimes should be adapted to the distributed nature of AI supply chains through sector-specific certification or mutual-recognition arrangements. Finally, transparency obligations should be operationalised through standardised, machine-readable model and dataset documentation — such as model cards and datasheets — that regulators and data subjects can consult without requiring disclosure of trade-secret-protected architectures.

7. Conclusion:

Artificial intelligence does not render existing privacy and data protection law obsolete, but it exposes structural assumptions — a linear data lifecycle, a stable line between collected and inferred data, and an assumption that deletion is straightforward — that no longer hold reliably in an AI-driven data ecosystem. The EU's layered GDPR–AI Act architecture is the most developed response to date, while the United States' sectoral, state-driven approach shows both the flexibility and the fragmentation that result from relying on general consumer-protection enforcement absent comprehensive federal legislation. Across jurisdictions, the recurring challenges identified here — consent, automated decision-making, inference, data minimisation, erasure, cross-border transfer, biometric identification, and enforcement fragmentation — point to a common conclusion: protecting personal data in the age of AI requires not only vigorous application of existing legal principles but targeted statutory reform, technical standard-setting, and sustained institutional coordination, which together will determine whether privacy retains practical meaning as AI becomes further embedded in everyday data processing.

²³ Garante per la protezione dei dati personali, Provvedimento del 30 marzo 2023 (Italy).

²⁴ FTC v. Rite Aid Corp. (Dec. 2023 settlement); FTC v. Rytr LLC (2024 enforcement action); White & Case, 'AI Watch: Global regulatory tracker – United States' (2026).

²⁵ Case C-634/21, OQ v. Land Hessen (CJEU, 7 December 2023).