INDIAN JOURNAL OF LEGAL AFFAIRS AND RESEARCH

Peer-reviewed, open-access, refereed journal

IJLAR

+91 70421 48991
editor@ijlar.com
www.ijlar.com

DISCLAIMER

The views and opinions expressed in the articles published in the Indian Journal of Legal Affairs and Research are those of the respective authors and do not necessarily reflect the official policy or position of the IJLAR, its editorial board, or its affiliated institutions. The IJLAR assumes no responsibility for any errors or omissions in the content of the journal. The information provided in this journal is for general informational purposes only and should not be construed as legal advice. Readers are encouraged to seek professional legal counsel for specific legal issues. The IJLAR and its affiliates shall not be liable for any loss or damage arising from the use of the information contained in this journal.

Introduction

Welcome to the Indian Journal of Legal Affairs and Research (IJLAR), a distinguished platform dedicated to the dissemination of comprehensive legal scholarship and academic research. Our mission is to foster an environment where legal professionals, academics, and students can collaborate and contribute to the evolving discourse in the field of law. We strive to publish high-quality, peer-reviewed articles that provide insightful analysis, innovative perspectives, and practical solutions to contemporary legal challenges. The IJLAR is committed to advancing legal knowledge and practice by bridging the gap between theory and practice.

Preface

The Indian Journal of Legal Affairs and Research is a testament to our unwavering commitment to excellence in legal scholarship. This volume presents a curated selection of articles that reflect the diverse and dynamic nature of legal studies today. Our contributors, ranging from esteemed legal scholars to emerging academics, bring forward a rich tapestry of insights that address critical legal issues and offer novel contributions to the field. We are grateful to our editorial board, reviewers, and authors for their dedication and hard work, which have made this publication possible. It is our hope that this journal will serve as a valuable resource for researchers, practitioners, and policymakers, and will inspire further inquiry and debate within the legal community.

Description

The Indian Journal of Legal Affairs and Research is an academic journal that publishes peer-reviewed articles on a wide range of legal topics. Each issue is designed to provide a platform for legal scholars, practitioners, and students to share their research findings, theoretical explorations, and practical insights. Our journal covers various branches of law, including but not limited to constitutional law, international law, criminal law, commercial law, human rights, and environmental law. We are dedicated to ensuring that the articles published in our journal adhere to the highest standards of academic rigor and contribute meaningfully to the understanding and development of legal theories and practices.

ARTIFICIAL INTELLIGENCE AND DATA BREACHES: A CRITICAL LEGAL STUDY (A COMPARATIVE ANALYSIS WITH REFERENCE TO INDIA, U.S.A. AND E.U.)

AUTHORED BY - ASHISH KUMAR HIMDHAR & DR. SHRADHA PANDEY
Faculty of Law, Shri Rawatpura Sarkar University, Raipur, Chhattisgarh, India

Abstract

Artificial intelligence (AI) is altering the legal meaning of a data breach. AI systems now leak personal data through model memorisation, inference attacks, poisoned training pipelines and compromised prompt logs, while amplifying attackers' capacity to find and exploit vulnerabilities. This paper examines whether the breach-response regimes of India, the United States and the European Union are adequate to this shift. Using a doctrinal and comparative method, it analyses India's Digital Personal Data Protection Act, 2023 and Rules of 2025; the
American patchwork of sectoral statutes, state laws, Federal Trade Commission enforcement and standing doctrine; and the EU's General Data Protection Regulation, Artificial Intelligence Act (as amended in 2026) and NIS2 Directive. It argues that the three systems have converged on notification timelines but diverge on redress, and that none has resolved three AI-specific problems: when model-level exposure is a "breach," the causation and harm problem that defeats claimants, and accountability across multi-party AI supply chains. It closes with jurisdiction-specific and common reform proposals.

Keywords: artificial intelligence; data breach; DPDP Act; GDPR; AI Act; genomic data; standing; comparative data protection

I. Introduction

The law of data breaches was drafted with a particular picture in mind: a custodian holds records, a third party gets unauthorised access to them, and the law responds with security duties before the event and notification and compensation after it. That picture is still accurate for many

incidents, but it is becoming an incomplete account of how personal data is exposed in an age of machine learning. A large language model trained on scraped text can reproduce fragments of that text, including names, addresses and telephone numbers, when prompted in particular ways.[1] A classifier trained on patient records can reveal whether a given individual's record was in its training set.[2] A face-recognition model can be inverted to approximate the images on which it was trained.[3] None of these events involves a conventional intrusion into a database, yet each results in the disclosure of personal data to persons who were never meant to receive it.

At the same time, AI is a force multiplier on the offensive side. Law-enforcement analysts have warned that generative models lower the skill threshold for phishing, impersonation and the writing of malicious code.[4] And AI is increasingly the object of attack: training corpora, vector databases, fine-tuning datasets, model weights and conversation logs are all now repositories of personal data that require protection. In March 2023, a bug in an open-source library used by a leading chatbot service exposed some users' conversation titles and, for a small percentage of subscribers, payment-related information to other users.[5] The incident was modest in scale, but it showed that AI services inherit every classical breach vector while adding new ones.

The problem is most acute where the underlying data is irreplaceable. Basu and De Mukhopadhyay, writing on next-generation sequencing (NGS) of human genomes, note that modern AI, especially deep learning, holds promise for improving variant calling, variant prediction and the usability of electronic health record systems in NGS-based diagnostics, and they identify informed consent, privacy and outcome review as the three principal ethical domains raised by the technology.[6] Genomic data illustrates why breach law must be reconsidered. A password can be reset and a card can be cancelled, but a genome cannot be changed, and it discloses information about biological relatives who never consented to anything. The re-identification of supposedly anonymous research participants from surname inference and public genealogical databases shows that de-identification offers weaker protection than legislators assumed.[7] When AI is layered on such data, both the benefits and the consequences of failure are magnified.

This paper asks how three major legal systems respond. India, the United States and the European Union have been chosen because they represent three distinct regulatory philosophies. India has recently enacted a comprehensive, notice-and-consent statute that is still being phased in and that contains no AI-specific provisions. The United States relies on a fragmented combination of sectoral federal law, state breach statutes, consumer-protection enforcement and private litigation,

with an AI policy that has shifted considerably between administrations. The European Union has the most developed framework, pairing the General Data Protection Regulation (GDPR) with the Artificial Intelligence Act and cybersecurity instruments. The three also differ in institutional capacity and in the weight they give to individual redress, which makes the comparison productive.

The central argument is threefold. First, the jurisdictions have converged on the mechanics of breach notification, with short regulatory deadlines now the norm, but they have not converged on the question of who is compensated and how. Second, none of the regimes has a satisfactory answer to the definitional question whether, and when, the extractability of personal data from a trained model amounts to a personal data breach. Third, existing law allocates responsibility on the basis of a binary distinction between controller and processor (or fiduciary and processor) that fits poorly with the layered supply chains of foundation-model development, fine-tuning, integration and deployment. The paper concludes that an "AI-aware" breach regime requires targeted amendments rather than wholesale replacement.

The method is doctrinal and comparative: primary materials are analysed in their own terms and then compared functionally. The law is stated as at 1 October 2026; because several instruments are being phased in or amended, readers should verify their status before relying on them.

II. Conceptual Framework

A. What Is a "Data Breach"?

Statutory definitions of a breach differ in wording but share a core. The GDPR defines a personal data breach as a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data.[8] India's Digital Personal Data Protection Act, 2023 (DPDP Act) defines the term more broadly still, covering unauthorised processing and accidental disclosure, acquisition, sharing, use, alteration, destruction or loss of access to personal data that compromises its confidentiality, integrity or availability.[9] American state statutes, by contrast, typically define a breach as the unauthorised acquisition of computerised data that compromises the security, confidentiality or integrity of personal information, and many apply only to enumerated categories of data such as name combined with a social security number or financial account number.[10]

B. Four Points of Intersection Between AI and Breach

It is useful to distinguish four ways in which AI and data breaches interact.

AI as target. Training datasets, feature stores, vector databases, model weights and prompt logs concentrate personal data in new places. These assets are often built quickly, by teams whose security practices lag behind those applied to core transactional systems, and they are frequently hosted with third-party providers. The 2023 chatbot incident referred to above is an example.[11]

AI as attack amplifier. Generative models reduce the cost of social engineering, including voice cloning and personalised phishing at scale, and may speed up vulnerability discovery.[12] The legal consequence is that the standard of "reasonable" security is a moving target: as attackers adopt AI, what a prudent custodian must do rises correspondingly.

AI as model-level leakage. Research has demonstrated training-data extraction from language models,[13] membership inference against machine-learning classifiers,[14] model inversion that exploits confidence information,[15] and scalable extraction from deployed production models.[16] Data poisoning, in which an adversary corrupts training data so as to alter model behaviour, has been shown to be practical against web-scale datasets.[17] These are the cases in which the traditional definition of a breach is most strained.

C. Genomic and Biometric Data as a Stress Test

Genomic and biometric data are the hard cases for breach law, and it is no accident that the legislation of all three jurisdictions singles out at least some of them for special treatment, or that scholars use them to expose gaps. Basu and De Mukhopadhyay concentrate on genomic data, and their reference list draws together much of the literature on re-identification, data sharing and group privacy.[18] Three insights from that literature are relevant here.

First, genetic data is relational. An individual's genome discloses information about parents, siblings and children. Quinn and Quinn argue that big genetic data presents data protection challenges that individual-centred regimes strain to accommodate.[19] Floridi and Mantelero have likewise argued for recognition of group or collective privacy interests in the era of big data.[20] If a breach exposes one person's genome, the harm is distributed among persons who are not "data subjects" of the record in a formal sense, and who therefore have no standing under statutes that confer rights only on the data subject or data principal.

Second, anonymisation is fragile. Surname inference from Y-chromosome markers and

genealogical databases demonstrated that apparently anonymous genomic records could be traced to named individuals.[21] Where AI models are trained on such data, the claim that a model is "anonymous" must be treated with considerable scepticism, a point that returns in the discussion of the European Data Protection Board's opinion on AI models.

Third, context matters. Nissenbaum's theory of contextual integrity holds that privacy is violated when information flows contravene the norms of the context in which it was shared.[22] A patient who provides a sample for diagnosis has not consented to its use in training a commercial model, still less to its extraction by an attacker. Contextual integrity thus offers a principled explanation of why a breach of AI-held data is harmful even when no financial loss follows.

D. Theories of Breach Harm

A recurrent obstacle in breach litigation, discussed in Part IV, is the demand that claimants show concrete injury. Solove and Citron have argued that data breaches create real harms of risk and anxiety that the law should recognise: the increased risk of future misuse, and the costs of the defensive measures and psychological distress that follow.[23] This paper adopts that view as its normative premise, and extends it: where the exposed data is immutable (genomes, biometrics) or where the exposure occurs through a model that will continue to leak after the incident, the risk is not transient, and the law of remedies must be able to account for it.

III. India

A. Constitutional Foundations

Indian data protection law rests on the constitutional recognition of privacy as a fundamental right in *Justice K.S. Puttaswamy (Retd.) v. Union of India*, where a nine-judge bench held that privacy is protected as an intrinsic part of the right to life and personal liberty under Article 21.[24] The Court identified informational privacy as a facet of that right and indicated that any restriction must satisfy legality, legitimate aim and proportionality. The judgment also contemplated that the State has a positive obligation to put in place a data protection regime, a point that supplies a constitutional justification for security and breach-notification duties imposed on private entities. The Court's later decision upholding the Aadhaar scheme with limitations confirmed that proportionality governs large-scale state data systems.[25]

B. The Information Technology Act, 2000 and Its Rules

Until the DPDP Act is fully operative, the primary legislative source for breach liability is the Information Technology Act, 2000 (IT Act). Section 43A makes a body corporate that possesses, deals with or handles sensitive personal data or information liable to pay compensation if it is negligent in implementing and maintaining reasonable security practices and thereby causes wrongful loss or gain.[26] The 2011 Reasonable Security Practices Rules define sensitive personal data to include health information, biometric information and financial data, and treat compliance with a recognised standard such as IS/ISO/IEC 27001 as one route to demonstrating reasonable security.[27] Section 72A criminalises disclosure of personal information in breach of a lawful contract, and sections 66 and 66C and related provisions address unauthorised access and identity theft.

Two further instruments deserve mention. Section 70B establishes the Indian Computer Emergency Response Team (CERT-In), and directions issued under it in April 2022 require covered entities to report specified categories of cyber incidents to CERT-In within six hours of noticing them, and to retain logs for 180 days.[28] Those directions are not confined to personal data breaches and apply to incidents such as data breaches, data leaks and attacks on AI and machine-learning systems, which the directions list among reportable categories. This makes them, perhaps unexpectedly, the only binding Indian instrument that names AI systems expressly in a security-incident context.

C. The Digital Personal Data Protection Act, 2023 and the 2025 Rules

The DPDP Act reorients Indian law around the concepts of data fiduciary and data principal. It applies to digital personal data processed within India and to processing outside India in connection with offering goods or services to persons in India. Three aspects are central to breach law.

Security and notification duties. Section 8(5) requires a data fiduciary to protect personal data in its possession or under its control, including data processed by a data processor on its behalf, by taking reasonable security safeguards to prevent a personal data breach. Section 8(6) requires the fiduciary, in the event of a breach, to give intimation to the Data Protection Board of India and to each affected data principal in the prescribed form and manner.[29] The Digital Personal Data Protection Rules, 2025, notified in November 2025, give content to these duties. They specify

minimum safeguards, including encryption or similar measures, access control, monitoring and the retention of logs, and require breach intimation to affected persons without delay and to the Board without delay, followed by a detailed report within 72 hours unless the Board allows longer.[30] Rules under the Act commence in phases, with the substantive obligations on fiduciaries scheduled to apply about eighteen months after notification, that is, from May 2027, and the existing IT Act regime continues in the interim.[31]

Penalties. The Schedule to the Act provides for monetary penalties of up to INR 250 crore for failure to take reasonable security safeguards and up to INR 200 crore for failure to give notice of a breach.[32] These are among the highest statutory ceilings in the Indian regulatory landscape, but two features undercut their significance for victims. Penalties are imposed by the Board and paid to the State, and the Act confers on data principals no right to compensation. Section 43A, which did provide a compensatory remedy, is to be omitted on commencement of the relevant provision.[33] The result is that, once the DPDP regime is fully in force, a person whose data is exposed through a fiduciary's negligence may have a right to complain to the Board but will have to look to general civil law for damages.

Exemptions and gaps relevant to AI. The Act excludes from its scope personal data that the data principal herself has made publicly available, and personal data made publicly available under a legal obligation.[34] On a plain reading, this exempts a substantial portion of web-scraped text and images from the Act's obligations, including its security duty. Since scraped public data is the principal raw material for foundation models, the exemption means that a leak of such data from a training set might not be a "personal data breach" at all. Section 17(2)(b) further exempts processing for research, archiving or statistical purposes, subject to prescribed standards, which may be invoked by developers of research models. The Act does not recognise special categories of data: the "sensitive personal data" category in the 2011 Rules is not replicated, and genomic and biometric data therefore receive no enhanced statutory protection beyond what the general reasonable-safeguards duty implies.[35] The Act has no analogue to the right not to be subjected to solely automated decisions found in the GDPR, and no obligation to explain algorithmic outputs.

D. Assessment

The Indian regime is in transition. The 72-hour reporting requirement, the codified safeguards and the penalty ceilings align India with international norms. Four gaps stand out: no compensation

for victims; a public-data exemption that may remove much AI training data from the Act's protection; no special treatment of genetic and biometric data; and no provision on model-level leakage. The Board's technical capacity to investigate machine-learning systems also remains to be demonstrated.

IV. The United States

A. A Fragmented Architecture

The United States has no comprehensive federal data protection statute. Instead, obligations arise from sectoral federal statutes, state breach-notification and security laws, federal agency rules and the general prohibition of unfair or deceptive practices.

All fifty states have breach-notification statutes, beginning with California's law, which took effect in 2003.[36] These laws differ in their definitions, triggers (some require a risk-of-harm assessment) and deadlines, and a company that suffers a nationwide breach must in effect comply with the strictest. Federal sectoral rules include the Health Insurance Portability and Accountability Act (HIPAA) Breach Notification Rule, which requires covered entities to notify affected individuals without unreasonable delay and in no case later than sixty days after discovery;[37] and the Gramm-Leach-Bliley Act and the Federal Trade Commission's Safeguards Rule, which now requires notification to the Commission of certain events affecting at least 500 consumers.[38]

Two federal developments are specific to cyber incidents. The Securities and Exchange Commission adopted rules in 2023 requiring public companies to disclose material cybersecurity incidents on Form 8-K within four business days of determining materiality, and to describe their cybersecurity risk management and governance annually.[39] And the Cyber Incident Reporting for Critical Infrastructure Act of 2022 directs the Cybersecurity and Infrastructure Security Agency to require covered entities to report cyber incidents within 72 hours and ransom payments within 24 hours; the rulemaking was proposed in 2024, and the finality and scope of the rule should be checked as of the date of reliance.[40]

B. FTC Enforcement and the Reasonable Security Standard

In the absence of a general statute, the Federal Trade Commission (FTC) has been the principal federal enforcer. Section 5 of the FTC Act prohibits unfair or deceptive acts or practices,[41] and the Commission has used it to challenge inadequate data security as unfair and misleading security

representations as deceptive. In *FTC v. Wyndham Worldwide Corp.*, the Third Circuit upheld the Commission's authority to bring unfairness claims in respect of data security.[42] The Eleventh Circuit, in *LabMD, Inc. v. FTC*, accepted the Commission's authority for argument's sake but held that its cease-and-desist order, which required a general "reasonable" security programme, was insufficiently specific to be enforceable.[43] The two decisions together illustrate the structural tension in American law: a standard of "reasonableness" is flexible enough to adapt to AI, but too indeterminate to guide conduct or sustain remedies.

The Commission has begun to apply these tools to AI. In its settlement with Everalbum, a photo-storage company that had used customers' images to develop facial-recognition technology without adequate consent, the FTC required deletion not only of the data but of the models and algorithms derived from it, a remedy sometimes described as algorithmic disgorgement.[44] In 2023, the Commission banned a retailer from using facial-recognition technology for surveillance for five years after alleging a failure to take reasonable measures to prevent harm from its use.[45] These actions are significant because they treat the model itself as an artefact that can be tainted by unlawful data practices and that may be ordered destroyed. They stop short of treating the leakage of data from a model as a breach, but the logic of the remedy would support doing so.

C. Private Litigation and the Standing Barrier

Private enforcement is where American breach law is most distinctive, and most problematic. Because Article III of the Constitution limits federal courts to actual cases or controversies, plaintiffs must show an injury in fact that is concrete, particularised and actual or imminent. In *Clapper v. Amnesty International USA*, the Supreme Court held that a risk of future harm must be "certainly impending" or involve a substantial risk to confer standing.[46] *Spokeo, Inc. v. Robins* added that a bare statutory violation is insufficient without concrete harm,[47] and *TransUnion LLC v. Ramirez* held that, in a damages suit, the mere risk of future harm does not suffice unless the harm materialises or the plaintiff shows a separate concrete harm such as emotional distress from exposure to the risk.[48]

The courts of appeals have applied these cases inconsistently to data breaches. The Seventh and Ninth Circuits have found standing where stolen data has been or is likely to be misused,[49] while the Fourth and Eighth Circuits have been more demanding where there was no evidence of actual misuse.[50] After *TransUnion*, the Second Circuit set out a three-factor test turning on whether the

data was targeted, whether it had been misused, and whether it was sensitive,[51] the Third Circuit held that the risk of misuse of stolen data can be concrete where the plaintiff alleges present harms such as mitigation costs,[52] and the Eleventh Circuit has required allegations of a substantial risk and concrete consequences.[53]

Those who clear the standing hurdle then face problems of causation and damages. In the litigation arising from the 2017 Equifax breach, which exposed the personal information of approximately 147 million people, the Eleventh Circuit largely upheld a class settlement while addressing standing and cy pres issues,[54] but the case also showed how difficult it is to prove that a specific identity theft was caused by a specific breach, given the number of other breaches that have exposed overlapping data.

For AI-related breaches the barrier is steeper. A plaintiff whose personal data was memorised by a language model and emitted in an output to another user must show injury. If the output was seen by only one other person and never misused, a court applying *TransUnion* may find no concrete harm. If the plaintiff cannot even determine whether the model has memorised her data, the causation problem is acute. Solove and Citron's argument that the law should recognise risk and anxiety as cognisable harms is a response to this precisely, but it has not been adopted in federal doctrine.[55]

State law is more receptive. The California Consumer Privacy Act (CCPA) gives consumers a limited private right of action, with statutory damages, where nonencrypted personal information is exfiltrated or disclosed because of a business's failure to maintain reasonable security.[56] Statutory damages avoid the problem of proving loss, and state courts are not bound by Article III. The Illinois Biometric Information Privacy Act (BIPA) has been the most consequential state statute for AI: it regulates the collection, retention and disclosure of biometric identifiers and provides liquidated damages.[57] The Illinois Supreme Court has held that a plaintiff need not allege harm beyond a statutory violation[58] and that a claim accrues with each scan or transmission,[59] although the legislature has since limited recovery to one per person.[60] BIPA litigation has produced substantial settlements against facial-recognition developers, including one against Clearview AI in which class members received an interest in the company's equity.[61]

D. AI Governance: Volatile and Largely Voluntary

Federal AI policy has changed course. Executive Order 14110 of October 2023, which had directed

federal agencies to develop safety, security and privacy standards for AI, was rescinded in January 2025,[62] and was replaced by an order that emphasised removing barriers to American AI leadership.[63] The principal non-binding federal resource is the National Institute of Standards and Technology's AI Risk Management Framework, which identifies privacy and security as characteristics of trustworthy AI and offers a voluntary process for mapping, measuring and managing risk.[64] A voluntary framework can inform what is "reasonable" in FTC enforcement or negligence litigation, but it creates no obligation.

E. Assessment

The American system is strong in enforcement tools and litigation volume, and weak in coherence. Standing doctrine denies a remedy to many whose data is exposed, and reliance on enumerated data elements and on "reasonableness" leaves AI-specific leakage largely unaddressed. Its distinctive contributions are model disgorgement and, in BIPA, statutory damages that sidestep the proof-of-harm problem.

V. The European Union

A. The GDPR as the Baseline

The GDPR is the most developed general regime for personal data breach, and it applies to AI systems without any amendment because it regulates processing of personal data regardless of the technology used. Its relevant provisions can be grouped into principles, security duties, breach response and remedies.

Principles and design duties. Controllers must process personal data in a manner that ensures appropriate security, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage (integrity and confidentiality), and must be able to demonstrate compliance (accountability).[65] They must implement data protection by design and by default,[66] and conduct a data protection impact assessment where processing is likely to result in high risk, which is expected for large-scale processing of special categories or novel technologies.[67] Genetic and biometric data used for identification are special categories whose processing is prohibited unless a specific ground applies.[68] The Court of Justice has held that data from which sensitive characteristics can be indirectly inferred are themselves special category data,[69] a holding of considerable importance for AI systems that routinely derive sensitive

inferences from non-sensitive inputs.

Security. Article 32 requires controllers and processors to implement technical and organisational measures appropriate to the risk, taking into account the state of the art, cost, and the nature, scope, context and purposes of processing, and gives examples including pseudonymisation, encryption, resilience of systems and regular testing.[70] The "state of the art" language gives the provision a dynamic character that is well suited to AI-enabled threats.

Breach response. Controllers must notify the supervisory authority of a personal data breach without undue delay and, where feasible, within 72 hours of becoming aware of it, unless the breach is unlikely to result in a risk to individuals' rights and freedoms; where the breach is likely to result in a high risk, they must also communicate it to affected data subjects without undue delay.[71] Processors must notify controllers without undue delay.[72]

Remedies and sanctions. Article 82 confers a right to compensation for material or non-material damage resulting from an infringement, against controllers and processors, and Article 83 permits administrative fines of up to EUR 10 million or 2 per cent of worldwide annual turnover for infringements of the security and breach-notification provisions, and up to EUR 20 million or 4 per cent for infringements of the basic principles and special-category rules.[73]

B. The Court of Justice and the Meaning of Damage

The Court of Justice of the European Union (CJEU) has given Article 82 an interpretation that is markedly more claimant-friendly than American standing doctrine, though it has not made liability automatic. In *Österreichische Post*, the Court held that a mere infringement of the GDPR does not itself confer a right to compensation, which requires damage and a causal link, but that the right is not subject to a minimum threshold of seriousness.[74] In *VB v. Natsionalna agentsia za prihodite*, concerning a hacking incident at the Bulgarian tax authority that exposed the data of millions of persons, the Court held that unauthorised disclosure of or access to

personal data by third parties does not alone mean that the controller's measures were inappropriate; that the controller bears the burden of showing that the measures were appropriate; and that fear of potential misuse of one's data can itself constitute non-material damage provided it is well founded.[75]

This reasoning is highly relevant to AI. First, it treats anxiety about future misuse as a compensable harm, going considerably further than the American federal courts and in the direction urged by

Solove and Citron. Second, the allocation of the burden of proof to the controller addresses the information asymmetry that is acute in AI contexts, where only the developer can know how a model was trained and secured. Third, the refusal to infer inadequate security from the mere fact of a breach preserves a negligence-like standard and prevents the regime from becoming one of strict liability. A claimant who suffers a model-leakage incident would, on these authorities, have a considerably clearer path to recovery in the EU than in the United States or India.

C. Is an AI Model Personal Data? The EDPB Opinion

The most important unresolved question under the GDPR for AI breaches is whether a trained model contains personal data. If it does, then the model is itself subject to security duties, its unauthorised disclosure (for example, theft of weights) is a personal data breach, and individuals have rights in respect of it. If it does not, model-level leakage may be outside the regime until it materialises in an output.

Supervisory authorities have not been of one mind. A discussion paper by the Hamburg data protection authority in 2024 took the view that large language models do not, by their storage of tokens and statistical relationships, store personal data as such, so that the model's existence is not itself a processing of personal data, although outputs may be.[76] The European Data Protection Board (EDPB), in its Opinion 28/2024, took a more cautious position: AI models trained on personal data cannot in all cases be considered anonymous, and anonymity must be assessed case by case by asking whether the likelihood of extracting personal data directly or through queries is insignificant, taking account of means reasonably likely to be used.[77] The Opinion also considers the consequences where a model was developed using unlawfully processed data, and indicates that this may affect the lawfulness of its subsequent deployment depending on the circumstances.

The EDPB approach is preferable. The research on extraction and inference attacks cited in Part II shows that whether a model "contains" personal data is an empirical rather than a metaphysical question.[78] A test that asks whether extraction is reasonably likely, rather than whether data is stored in a human-readable form, is the only one consistent with the GDPR's technology-neutral definition of personal data. But the test also imposes a significant evidentiary burden on developers, who must be able to demonstrate resistance to extraction, and it leaves open what testing is sufficient.

D. Enforcement Against AI Developers

National authorities have begun to act. The Italian data protection authority temporarily restricted a major chatbot in 2023 and later fined its developer EUR 15 million, citing, among other things, failure to notify a personal data breach to the authority, absence of an adequate lawful basis for training, and transparency and age-verification shortcomings.[79] The proceeding is significant because the 2023 incident described in Part I was treated as a notifiable breach,

which confirms that the GDPR applies to AI services' operational incidents. The Court of Justice has also held that automated credit scoring can itself constitute an automated individual decision under Article 22,[80] and that data subjects may be entitled to meaningful information about the logic involved in automated decisions.[81]

E. The Artificial Intelligence Act

The Artificial Intelligence Act, Regulation (EU) 2024/1689, adds an AI-specific layer.[82] It adopts a risk-based classification, prohibits certain practices, imposes extensive requirements on high-risk systems, and creates obligations for providers of general-purpose AI models, including those with systemic risk.

Several provisions bear directly on breach. High-risk systems must be designed to achieve an appropriate level of accuracy, robustness and cybersecurity, including resilience against attempts by third parties to alter their use, outputs or performance by exploiting vulnerabilities, with technical solutions to address AI-specific vulnerabilities such as data poisoning, model poisoning, adversarial examples and confidentiality attacks.[83] Data governance requirements apply to training, validation and testing datasets,[84] and providers of general-purpose models with systemic risk must ensure adequate cybersecurity protection for the model and its physical infrastructure and must report serious incidents.[85] Providers of high-risk systems must report serious incidents to market surveillance authorities,[86] and penalties can reach EUR 35 million or 7 per cent of worldwide turnover for prohibited practices, and EUR 15 million or 3 per cent for most other violations.[87]

Two points are critical. The first is that the AI Act expressly recognises "confidentiality attacks" as a vulnerability to be addressed, which is the first explicit legislative acknowledgement among the three jurisdictions of the model-level leakage described in Part II. The second is timing. The Act's application dates have been reshaped by Regulation (EU) 2026/1744, the "AI Omnibus,"

which entered into force in July 2026 and postponed the date from which obligations for stand-alone high-risk systems apply to 2 December 2027, with a later date for systems embedded in regulated products, while leaving the prohibitions and the general-purpose model rules in place.[88] The breach-relevant obligations for high-risk systems are therefore not yet enforceable, and for the present the GDPR remains the operative instrument for AI-related breaches.

F. Cybersecurity Instruments and Reform Proposals

The NIS2 Directive imposes cyber risk-management and reporting duties on essential and important entities across many sectors, with an early warning within 24 hours of becoming aware of a significant incident, an incident notification within 72 hours, and a final report within one month.[89] The Cyber Resilience Act introduces cybersecurity requirements for products with digital elements, which may extend to AI-enabled software.[90] The overlap among the GDPR, NIS2, the AI Act and the Cyber Resilience Act is a recognised source of compliance burden.

G. Assessment

The European Union offers the most complete treatment, and its strengths are the technology-neutral definition of personal data, a dynamic "state of the art" security standard, a developing body of case law on damage, specialised supervisory authorities, and the AI Act's recognition of AI-specific attack vectors. Its weaknesses are the complexity created by overlapping instruments; delay in implementing the AI Act; uneven enforcement across Member States and the slowness of cross-border cooperation under the one-stop-shop mechanism; and the continuing uncertainty regarding model anonymity. It is also exposed to the political pressure for deregulation reflected in the Omnibus process.

VI. Comparative Analysis

A. Overview

The following table summarises the principal features of the three regimes on the matters examined above. It is a simplification, and its entries carry the qualifications set out in Parts III to V.

Feature	India	United States	European Union
Core instruments		Breach definition	

Regulator notice	DPDP Act 2023 and Rules 2025 (phased); IT Act 2000 s. 43A and 2011 Rules in the interim; CERT-In Directions	Sectoral federal statutes; 50 state breach laws; FTC Act s. 5; SEC rules; CIRCIA	GDPR; AI Act (as amended 2026); NIS2; Cyber Resilience Act
	Broad: unauthorised processing, accidental disclosure, loss of confidentiality, integrity or availability	Varies by state; tied to unauthorised acquisition of enumerated data	Breach of security leading to destruction, loss, alteration, disclosure or access
	Without delay, plus detailed report within 72 hours; CERT-In six hours	Fragmented: 60 days (HIPAA); four business days (SEC, material incidents); state deadlines	72 hours (GDPR); 24 hours/72 hours/one month (NIS2)

IJLAR

Feature	India	United States	European Union
Individual redress	None under DPDP Act; residual civil remedies	Class actions subject to standing; statutory damages in CCPA and BIPA	Art. 82 damages, including non-material damage; no minimum threshold
Special data	No special categories in DPDP Act	Sectoral (health, financial, biometric in some states)	Special categories incl. genetic and biometric
AI-specific rules	None binding; soft-law guidance; CERT-In names AI incidents	Voluntary NIST framework; state laws; FTC remedies	AI Act, including confidentiality attacks and incident reporting

IJLAR

B. Convergence on Notification

The clearest convergence is on timing. The EU's 72-hour rule has been adopted in substance in India's Rules, and the SEC and CIRCIA rules show a similar trend towards days rather than weeks. The convergence is shallower than it looks: the EU requires notification only where there is a risk to rights and freedoms, the Indian Rules appear to require notification of every personal data breach, and the United States has a tiered, sector-specific structure.

C. Divergence on Redress

The sharpest divergence is on the right of victims. The EU treats compensation as an individual right with a low threshold of damage, supported by burden-shifting. The United States provides a remedy through class actions and statutory damages but conditions access on a constitutional standing doctrine that increasingly excludes risk-based injuries. India's new law provides none. This divergence matters in an AI setting because the harms of model-level leakage are, by nature, probabilistic and diffuse. A system that compensates only demonstrated pecuniary loss will fail to compensate them. A regime that accepts non-material damage, as the CJEU has done, is better equipped.

D. Definitions and Scope Versus the AI Reality

None of the three systems defines a breach so as to capture model-level leakage expressly. The EU comes closest through the combination of the EDPB's opinion and Article 15 of the AI Act. In India, the public-data exemption and the omission of special categories point in the opposite direction. In the United States, enumerated-element definitions and the focus on unauthorised acquisition exclude most AI-specific harms. These differences have a practical consequence: a developer that trains a model on scraped data and releases it globally faces materially different exposure depending on whether the leaked individual is in Frankfurt, Chicago or Kolkata. This regulatory asymmetry may encourage forum shopping and a compliance race to the bottom, and it exposes people in jurisdictions with weaker law to risks that are tolerated less elsewhere.

E. Accountability Along the AI Supply Chain

All three regimes allocate responsibility principally to the data controller, fiduciary or covered entity, with subsidiary duties for processors. A modern AI system is typically produced by a chain of actors: a foundation-model developer, a fine-tuner who adds proprietary data, an integrator, a cloud host and a deploying enterprise. The classification of each actor as

controller, joint controller or processor is contested and context-dependent. The AI Act addresses this through the roles of provider and deployer, but its allocation of obligations concerns product safety more than data protection. The result is that a breach arising from model leakage may be the subject of reciprocal blame among actors, with the claimant unable to identify the responsible party. Basu and De Mukhopadhyay's emphasis on informed consent and outcome review in genomic AI, by analogy, highlights the difficulty of attributing responsibility in a chain where

the patient's consent is given at the start and the consequences materialise far downstream.[91]

VII. Critical Assessment and Recommendations

A. Common Reforms

The comparison suggests four reforms that are broadly applicable across jurisdictions.

First, define an "AI-mediated personal data breach." Statutes or authoritative guidance should provide that the unauthorised disclosure of personal data through the operation of an AI system, including by extraction from a model's parameters or outputs, is a personal data breach, regardless of whether there was an intrusion into a conventional database. To avoid over-inclusion, the trigger should be demonstrated or reasonably foreseeable extractability, consistent with the EDPB's approach of asking whether extraction is reasonably likely, rather than the mere fact that personal data was in the training set. Developers should be required to conduct and document adversarial testing for extraction and inference attacks proportionate to the sensitivity of the training data.

Second, adopt a harm standard that recognises risk. Whether through legislation or interpretation, remedies should extend to well-founded fear of misuse, mitigation costs and the loss of control over irreplaceable data such as genomes. The CJEU's reasoning in *VB* is a workable template. Statutory or liquidated damages, as in BIPA, can be used where proof of individual loss is impractical, with caps or per-person limits to avoid disproportionate exposure.

Third, allocate accountability across the supply chain. Legislators should provide that the developer of a model, the party that fine-tunes it and the deployer each bear obligations proportionate to their control, with joint and several liability to the data subject and rights of contribution among themselves. Contractual allocation should not be allowed to defeat the claimant. Transparency duties, for example a documented data lineage and a record of the datasets used, are a precondition to any such allocation.

Fourth, strengthen protection of genetic and biometric data and recognise group interests.

Given the permanence of such data and its relational character, the law should require enhanced safeguards (encryption, strict access control, purpose limitation and independent audit) for AI systems that process it, and should explore collective redress mechanisms that allow representative bodies to act for affected families and communities.

B. Jurisdiction-Specific Recommendations

India. The Rules should be supplemented, or the Act amended, in four respects. A right to compensation for data principals should be established, since a penalty regime that pays the State and leaves the victim unrecompensed does not align with the constitutional vindication of privacy. The exemption for publicly available personal data should be narrowed so that it does not dispense with security and breach duties for such data once aggregated into a dataset or model, an approach consistent with the contextual integrity analysis that information shared in one context does not become free for all uses.[92] Genetic, biometric and health data should be recognised as a category warranting enhanced safeguards, restoring and extending the protection of the 2011 Rules. And the Board should be given the technical capacity, and a statutory mandate, to audit AI systems, with the Significant Data Fiduciary regime used as the vehicle for AI-specific obligations such as extraction testing.

United States. Congress should enact a federal baseline for breach notification and security that pre-empts only weaker state standards and provides a private remedy for risk-based harm, addressing the standing problem by legislative definition of injury. The FTC's use of model disgorgement should be placed on firmer statutory footing.

VIII. Conclusion

The advent of AI has not made the existing law of data breach obsolete, but it has exposed its dependence on a model of the world in which data sits in identifiable places and is stolen by identifiable intruders. In AI systems, data is dispersed across parameters, embeddings and logs; it can be extracted without intrusion; and the parties who are responsible for its security are numerous and interdependent. The three systems studied here respond to that challenge at different speeds and with different tools. The European Union has the most comprehensive and the most victim-oriented framework, although it is complex and its AI Act timetable has slipped. The United States has powerful enforcement and litigation tools that are constrained by standing doctrine and by the absence of a coherent federal framework. India has recently built a modern statutory foundation that will matter greatly for the largest population of data

principals in the world, but it has deferred, so far, the questions of compensation, AI-specific safeguards and special categories of data.

The reforms proposed in this paper (a definition of AI-mediated breach tied to reasonably likely extractability, a risk-sensitive harm standard, supply-chain accountability and enhanced protection for irreplaceable data) are incremental and capable of being absorbed into existing frameworks. They also respond to the central insight offered by the genomic-data literature on which this paper has drawn: that when personal data is irreplaceable, relational and increasingly processed by systems that few people can inspect, the protection of individuals requires a law that attends to risk before harm materialises, and that treats consent, privacy and outcome review as continuing obligations rather than one-time formalities.[93]

[1] Nicholas Carlini et al., *Extracting Training Data from Large Language Models*, 30 USENIX Sec. Symp. 2633 (2021).
[2] Reza Shokri et al., *Membership Inference Attacks Against Machine Learning Models*, 2017 IEEE Symp. on Sec. & Priv. 3.
[3] Matt Fredrikson, Somesh Jha & Thomas Ristenpart, *Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures*, 22 Proc. ACM SIGSAC Conf. on Comput. & Commc'ns Sec. 1322 (2015).
[4] Europol Innovation Lab, *ChatGPT: The Impact of Large Language Models on Law Enforcement* (2023).
[5] *March 20 ChatGPT Outage: Here's What Happened*, OpenAI (Mar. 24, 2023), https://openai.com/index/march-20-chatgpt-outage/.
[6] Abhisikta Basu & Keya De Mukhopadhyay, *Next-Generation Sequencing of Human Genomes: Ethico-Legal Challenges with Regard to Data Security*, *in* Proceedings of Sixth International Ethical Hacking Conference 149 (Mohuya Chakraborty et al. eds., 2026), https://doi.org/10.1007/978-981-96-8632-2_9.
[7] Melissa Gymrek et al., *Identifying Personal Genomes by Surname Inference*, 339 Science 321 (2013); *see also* Basu & De Mukhopadhyay, *supra* note 6, at 149 (listing the literature on genomic re-identification).
[8] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, art. 4(12), 2016 O.J. (L 119) 1 [hereinafter GDPR].
[9] Digital Personal Data Protection Act, 2023, No. 22 of 2023, § 2(u), India Code (2023) [hereinafter DPDP Act].
[10] Cal. Civ. Code § 1798.82(g) (West 2024).
[11] *March 20 ChatGPT Outage: Here's What Happened*, OpenAI (Mar. 24, 2023), https://openai.com/index/march-20-chatgpt-outage/.
[12] Europol Innovation Lab, *ChatGPT: The Impact of Large Language Models on Law Enforcement* (2023).
[13] Nicholas Carlini et al., *Extracting Training Data from Large Language Models*, 30 USENIX Sec. Symp. 2633 (2021).
[14] Reza Shokri et al., *Membership Inference Attacks Against Machine Learning Models*, 2017 IEEE Symp. on Sec. & Priv. 3.
[15] Matt Fredrikson, Somesh Jha & Thomas Ristenpart, *Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures*, 22 Proc. ACM SIGSAC Conf. on Comput. & Commc'ns Sec. 1322 (2015).
[16] Milad Nasr et al., *Scalable Extraction of Training Data from (Production) Language Models*, arXiv:2311.17035 (Nov. 28, 2023).
[17] Nicholas Carlini et al., *Poisoning Web-Scale Training Datasets Is Practical*, arXiv:2302.10149 (Feb. 20, 2023).
[18] Basu & De Mukhopadhyay, *supra* note 6.
[19] Paul Quinn & Liam Quinn, *Big Genetic Data and Its Big Data Protection Challenges*, 34 Comput. L. & Sec. Rev. 1000 (2018).
[20] Luciano Floridi, *Group Privacy: A Defence and an Interpretation*, *in* Group Privacy: New Challenges of Data

Technologies 83 (Linnet Taylor, Luciano Floridi & Bart van der Sloot eds., 2017); Alessandro Mantelero, *From Group Privacy to Collective Privacy: Towards a New Dimension of Privacy and Data Protection in the Big Data Era*, *in id.* at 139.
[21] Gymrek et al., *supra* note 7; *see also* Mahsa Shabani & Pascal Borry, *Rules for Processing Genetic Data for Research Purposes in View of the New EU General Data Protection Regulation*, 26 Eur. J. Hum. Genetics 149 (2018).
[22] Helen Nissenbaum, Privacy in Context: Technology, Policy, and the Integrity of Social Life (2010).
[23] Daniel J. Solove & Danielle Keats Citron, *Risk and Anxiety: A Theory of Data-Breach Harms*, 96 Tex. L. Rev. 737 (2018).
[24] Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 S.C.C. 1 (India).
[25] Justice K.S. Puttaswamy (Retd.) v. Union of India, (2019) 1 S.C.C. 1 (India).
[26] Information Technology Act, 2000, No. 21 of 2000, § 43A, India Code (2000).
[27] Information Technology (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules, 2011, rr. 3, 8, G.S.R. 313(E) (Apr. 11, 2011) (India).
[28] Directions Under Sub-Section (6) of Section 70B of the Information Technology Act, 2000 Relating to Information Security Practices, Procedure, Prevention, Response and Reporting of Cyber Incidents for Safe & Trusted Internet, No. 20(3)/2022-CERT-In (Apr. 28, 2022) (India).
[29] DPDP Act, *supra* note 9, § 8(5)–(6).
[30] Digital Personal Data Protection Rules, 2025, rr. 6–7 (India) (notified Nov. 13, 2025).
[31] *Id.* r. 1 (providing for phased commencement); *see* Press Release, Ministry of Electronics & Information Technology, Digital Personal Data Protection Rules, 2025 Notified (Nov. 14, 2025) [exact citation to be confirmed against the Gazette notification].
[32] DPDP Act, *supra* note 9, § 33 & sched.
[33] *Id.* § 44(2).
[34] *Id.* § 3(c)(ii).
[35] *Compare supra* note 22, r. 3 (defining sensitive personal data or information), *with* DPDP Act, *supra* note 9 (containing no corresponding category).
[36] Cal. Civ. Code § 1798.82 (West 2024) (enacted 2002, effective July 1, 2003).
[37] 45 C.F.R. §§ 164.400–.414 (2024); *id.* § 164.404(b).
[38] Gramm-Leach-Bliley Act, 15 U.S.C. §§ 6801–6809; 16 C.F.R. § 314.4(j) (2024).
[39] Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure, 88 Fed. Reg. 51,896 (Aug. 4, 2023) (codified at 17 C.F.R. pts. 229, 232, 239, 240, 249). [40] Cyber Incident Reporting for Critical Infrastructure Act of 2022, Pub. L. No. 117-103, div. Y, 136 Stat. 1038 (codified at 6 U.S.C. §§ 681–681g); Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) Reporting Requirements, 89 Fed. Reg. 23,644 (proposed Apr. 4, 2024).
[41] 15 U.S.C. § 45(a).
[42] FTC v. Wyndham Worldwide Corp., 799 F.3d 236 (3d Cir. 2015).
[43] LabMD, Inc. v. FTC, 894 F.3d 1221 (11th Cir. 2018).
[44] *In re* Everalbum, Inc., FTC File No. 192-3172 (May 7, 2021) (decision and order).
[45] Press Release, Fed. Trade Comm'n, Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology Without Reasonable Safeguards (Dec. 19, 2023).
[46] Clapper v. Amnesty Int'l USA, 568 U.S. 398, 409, 414 n.5 (2013).
[47] Spokeo, Inc. v. Robins, 578 U.S. 330 (2016).
[48] TransUnion LLC v. Ramirez, 594 U.S. 413 (2021).
[49] Remijas v. Neiman Marcus Grp., LLC, 794 F.3d 688 (7th Cir. 2015); *In re* Zappos.com, Inc., 888 F.3d 1020 (9th Cir. 2018).
[50] Beck v. McDonald, 848 F.3d 262 (4th Cir. 2017); *In re* SuperValu, Inc., 870 F.3d 763 (8th Cir. 2017).
[51] McMorris v. Carlos Lopez & Assocs., LLC, 995 F.3d 295 (2d Cir. 2021).
[52] Clemens v. ExecuPharm Inc., 48 F.4th 146 (3d Cir. 2022).
[53] Green-Cooper v. Brinker Int'l, Inc., 73 F.4th 883 (11th Cir. 2023).
[54] *In re* Equifax Inc. Customer Data Sec. Breach Litig., 999 F.3d 1247 (11th Cir. 2021).
[55] Solove & Citron, *supra* note 18, at 737.
[56] Cal. Civ. Code § 1798.150 (West 2024).
[57] Biometric Information Privacy Act, 740 Ill. Comp. Stat. 14/1–99 (2024).
[58] Rosenbach v. Six Flags Ent. Corp., 2019 IL 123186, 129 N.E.3d 1197.
[59] Cothron v. White Castle Sys., Inc., 2023 IL 128004, 216 N.E.3d 918.
[60] 740 Ill. Comp. Stat. 14/20 (as amended by Pub. Act 103-0769 (2024)).
[61] *In re* Clearview AI, Inc. Consumer Priv. Litig., No. 1:21-cv-00135 (N.D. Ill.) [final approval and citation to be confirmed].

[62] Exec. Order No. 14,110, 88 Fed. Reg. 75,191 (Oct. 30, 2023), *revoked by* Exec. Order No. 14,148, 90 Fed. Reg. 8237 (Jan. 20, 2025).
[63] Exec. Order No. 14,179, 90 Fed. Reg. 8741 (Jan. 23, 2025).
[64] Nat'l Inst. of Standards & Tech., NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (2023).
[65] GDPR, *supra* note 8, art. 5(1)(f), (2).
[66] *Id.* art. 25.
[67] *Id.* art. 35.
[68] *Id.* arts. 4(13)–(14), 9.
[69] Case C-184/20, OT v. Vyriausioji tarnybinės etikos komisija, ECLI:EU:C:2022:601 (Aug. 1, 2022).
[70] GDPR, *supra* note 8, art. 32.
[71] *Id.* arts. 33–34.
[72] *Id.* art. 33(2); *see also* Eur. Data Prot. Bd., Guidelines 9/2022 on Personal Data Breach Notification Under GDPR (version 2.0, Mar. 28, 2023).
[73] GDPR, *supra* note 8, arts. 82, 83(4)–(5).
[74] Case C-300/21, UI v. Österreichische Post AG, ECLI:EU:C:2023:370 (May 4, 2023).
[75] Case C-340/21, VB v. Natsionalna agentsia za prihodite, ECLI:EU:C:2023:986 (Dec. 14, 2023).
[76] Der Hamburgische Beauftragte für Datenschutz und Informationsfreiheit, Discussion Paper: Large Language Models and Personal Data (July 15, 2024).
[77] Eur. Data Prot. Bd., Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models (Dec. 17, 2024).
[78] Carlini et al., *supra* note 1; Shokri et al., *supra* note 2; Fredrikson et al., *supra* note 3.
[79] Garante per la protezione dei dati personali [Italian Data Protection Authority], Press Release (Dec. 20, 2024) (fine against OpenAI OpCo, LLC) [exact provision number to be confirmed].
[80] Case C-634/21, SCHUFA Holding (Scoring), ECLI:EU:C:2023:957 (Dec. 7, 2023).
[81] Case C-203/22, CK v. Dun & Bradstreet Austria GmbH, ECLI:EU:C:2025:117 (Feb. 27, 2025).
[82] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence, 2024 O.J. (L) 2024/1689 [hereinafter AI Act].
[83] *Id.* art. 15(4)–(5).
[84] *Id.* art. 10.
[85] *Id.* art. 55(1).
[86] *Id.* art. 73.
[87] *Id.* art. 99.
[88] Regulation (EU) 2026/1744 of the European Parliament and of the Council Amending Regulation (EU) 2024/1689 (Digital Omnibus on AI), 2026 O.J. (L) (July 24, 2026) (entered into force July 27, 2026) [Official Journal series number to be confirmed].
[89] Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on Measures for a High Common Level of Cybersecurity Across the Union, art. 23, 2022 O.J. (L 333) 80.
[90] Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on Horizontal Cybersecurity Requirements for Products with Digital Elements, 2024 O.J. (L) 2024/2847.
[91] Basu & De Mukhopadhyay, *supra* note 6.
[92] Nissenbaum, *supra* note 17.
[93] Basu & De Mukhopadhyay, *supra* note 6.