

Original Article

Satisfying GDPR, HIPAA, and Data Sovereignty Simultaneously: Federated Learning as a Legal-Technical Pathway for Cross-Border Pandemic Data Sharing

S. DAVID JEBASINGH

Data Analyst, LatentView, Chennai, Tamil Nadu, India.

ABSTRACT - The response to every major epidemic of the past three decades has been slowed by the same structural failure: the public health data needed for early warning is distributed across national health systems that cannot share it without violating domestic privacy regulations, data sovereignty laws, or both. The COVID-19 pandemic made this failure catastrophically visible, with weeks to months of preventable delay in cross-border epidemiological intelligence attributable to legal barriers to data sharing rather than absence of data. Federated learning the training of AI models on distributed data without requiring that data to leave its source institution offers a technical architecture that may satisfy the simultaneous, partly conflicting requirements of the General Data Protection Regulation, the Health Insurance Portability and Accountability Act, and national data sovereignty frameworks by design rather than by legal negotiation. This paper provides the first systematic legal-technical analysis of whether decentralized federated learning satisfies the specific compliance requirements of all three regulatory regimes simultaneously for cross-border pandemic surveillance applications. We analyze six regulatory requirement categories across GDPR, HIPAA, and data sovereignty law, map each to the federated learning technical mechanism that addresses it, characterize the residual privacy risks and their mitigations, and compare federated learning against four alternative cross-border data sharing approaches. The HealthVigil pandemic intelligence system provides empirical evidence that a cross-border federated AI surveillance system can achieve 43-day earlier outbreak detection and a 37% reduction in false alarms relative to conventional surveillance while operating under the privacy-preserving federated architecture that satisfies all three regulatory regimes, demonstrating that the legal-technical pathway proposed here is operationally as well as legally viable.

KEYWORDS - Federated Learning, GDPR, HIPAA, Data Sovereignty, Pandemic Surveillance, Cross-Border Data Sharing, Privacy-Preserving AI, Health Data Law, Epidemic Intelligence, HealthVigil.

1. INTRODUCTION

The governance of health data in a world of transnational pathogens represents one of the most structurally intractable problems in global public health law. Pandemic pathogens do not respect national borders, but the legal frameworks governing the health data generated by their spread are emphatically national and they conflict with each other in ways that make conventional cross-border data sharing impractical even when all parties want to cooperate [1]. The European Union's General Data Protection Regulation prohibits the transfer of special category health data to third countries without an adequacy decision or specific legal safeguards. The United States Health Insurance Portability and Accountability Act requires that any disclosure of protected health information meet a minimum necessary standard and that any receiving entity execute a business associate agreement. National data sovereignty laws in India, China, Russia, and dozens of other countries require that sensitive health data be processed exclusively on servers located within national territory.

These requirements are not merely procedural. They reflect genuine substantive concerns about how health data can be abused when it crosses borders: for insurance discrimination, immigration enforcement, political targeting, or commercial exploitation. The legal barriers to cross-border health data sharing exist because previous episodes of inadequately governed health data sharing produced real harms. The problem is not that the regulations are wrong but that they were designed for a world of bilateral relationships and point-in-time data transfers, not for the real-time, multilateral, continuous data integration that effective pandemic surveillance requires [2].

Federated learning offers a technical architecture that may resolve this dilemma without requiring legal reform. By training AI models on local data without transmitting that data across borders, federated learning enables the development of shared disease surveillance models that are epidemiologically equivalent to centrally trained models but architecturally incompatible with the regulatory violations that conventional cross-border data sharing entails. The HealthVigil pandemic intelligence system demonstrated this capability at deployment scale, achieving 43-day earlier outbreak detection and a 37% reduction in false alarms relative to conventional surveillance systems while operating on a privacy-preserving federated architecture that

processes clinical records, genomic sequences, social media signals, mobility patterns, and environmental data without centralizing any of these streams [3].

The legal question whether federated learning actually satisfies GDPR, HIPAA, and data sovereignty requirements as a matter of law rather than technical aspiration has not been systematically analyzed. This paper provides that analysis. Section 2 presents the regulatory landscape. Section 3 analyzes GDPR compliance. Section 4 analyzes HIPAA compliance. Section 5 analyzes data sovereignty compliance. Section 6 presents the integrated architecture and its performance evidence. Section 7 characterizes residual risks and their mitigations. Section 8 compares federated learning against alternatives. Section 9 addresses limitations and governance implications. Section 10 concludes.

2. THE REGULATORY LANDSCAPE FOR CROSS-BORDER HEALTH DATA SHARING

2.1. THE STRUCTURAL INCOMPATIBILITY OF EXISTING FRAMEWORKS

The three major regulatory regimes governing health data sharing were each designed within a national or regional context and reflect the institutional and political priorities of their jurisdictions at the time of adoption. GDPR, adopted in 2016 and effective from 2018, reflects the EU's fundamental rights approach to privacy as a human right, with corresponding emphasis on individual consent, purpose limitation, and restriction of data flows to countries with equivalent protection [4]. HIPAA, adopted in 1996 and substantially amended by the HITECH Act in 2009, reflects the US approach of sector-specific regulation with focus on operational security, minimum necessary use, and business relationship accountability [5]. National data sovereignty frameworks, adopted or strengthened by dozens of countries since 2012, reflect concerns about foreign government access to domestic data streams and economic competition for data-driven services [6].

The federated learning paradigm introduced by McMahan et al. as a communication-efficient distributed model training approach [7] provides the technical architecture that addresses this regulatory gap. None of these frameworks contemplated the real-time, multilateral, AI-driven pandemic surveillance architecture that the COVID-19 experience identified as necessary. The result is a regulatory trilemma: a surveillance system comprehensive enough to detect emerging pandemics early needs data from EU member states (GDPR), US institutions (HIPAA), and major disease-endemic countries with data residency requirements (sovereignty laws), but satisfying all three simultaneously under conventional data sharing architectures is practically impossible within the timescales of an emerging epidemic.

2.2. PREVIOUS APPROACHES AND THEIR LIMITATIONS

The international public health community has attempted several approaches to this regulatory trilemma. The International Health Regulations, revised in 2005, establish a framework for information sharing between WHO member states under public health emergency conditions, but they operate through diplomatic notification channels rather than automated data integration and have no mechanism for authorizing cross-border AI model training [8]. The WHO's established emergency mechanisms under IHR Article 12 can accelerate information sharing during declared Public Health Emergencies of International Concern, but the 2020 COVID-19 response demonstrated that the legal uncertainty surrounding raw data sharing suppressed participation even under emergency conditions. Duch and Thiessen have argued that fundamental reform of the IHR is required to create the mandatory data sharing obligations that voluntary cooperation has failed to deliver [9]. Bilateral data sharing agreements between national health agencies provide a legally secure pathway for data exchange, but they require months to years to negotiate and cannot be assembled into the multilateral network needed for global surveillance on the timescales of an emerging pandemic [10].

3. GDPR COMPLIANCE ANALYSIS

3.1. SPECIAL CATEGORY HEALTH DATA AND ARTICLE 9

GDPR Article 9 prohibits the processing of special category personal data, which explicitly includes data concerning health, without a lawful basis from the limited set provided in Article 9(2). For pandemic surveillance, the most applicable lawful bases are Article 9(2)(i), which permits processing necessary for reasons of public interest in the area of public health, and Article 9(2)(j), which permits processing necessary for scientific research purposes. Both require that Union or Member State law provides for adequate and specific measures to safeguard fundamental rights and interests.

The critical legal question for federated learning is whether gradient updates transmitted from a local training process to a global aggregator constitute the processing of special category personal data. As Voigt and von dem Bussche analyze in their authoritative commentary on GDPR, the regulation's scope is bounded by the concept of identifiability, which requires assessment of all means reasonably likely to be used to identify the person [11]. GDPR Recital 26 provides that the principles of data protection should not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person. Gradient updates are statistical aggregations of model parameter changes derived from training data; they do not contain, and should not be capable of directly identifying, any individual patient. The Article 29 Working Party's opinion on anonymization techniques, adopted by the European Data Protection Board, establishes that data is anonymous if the identification risk is negligible having regard to all the means reasonably likely to be used by the controller

or a third party [12]. Properly implemented federated learning with sufficient local dataset size and gradient aggregation should satisfy this standard, placing gradient transmissions outside the GDPR's material scope.

3.2. INTERNATIONAL TRANSFER RESTRICTIONS AND CHAPTER V

Even if gradient updates are not personal data, the institutional data governance around the federated learning system may involve personal data processing for which GDPR Chapter V transfer restrictions apply. The appointment of a global aggregator as a data processor under Article 28 requires a data processing agreement, but does not require an adequacy decision or Article 46 safeguard if the aggregator does not receive personal data which, by the federated architecture's design, it does not. This analysis is supported by the European Data Protection Supervisor's 2021 guidance on federated learning in health research, which provisionally concluded that properly implemented federated learning does not constitute an international transfer of health data under GDPR [4].

4. HIPAA COMPLIANCE ANALYSIS

4.1. PROTECTED HEALTH INFORMATION AND THE MINIMUM NECESSARY STANDARD

HIPAA defines protected health information as individually identifiable health information that relates to the past, present, or future physical or mental health of an individual; the provision of health care to an individual; or the past, present, or future payment for the provision of health care [5]. The minimum necessary standard in 45 CFR 164.502(b) requires that covered entities make reasonable efforts to limit the use or disclosure of PHI to the minimum necessary to accomplish the intended purpose.

In federated learning for pandemic surveillance, the covered entity's local training process uses PHI within the covered entity's own systems for the purpose of model training. This constitutes treatment, operations, or an activity preparatory to research under HIPAA, each of which has its own access pathway that does not require individual authorization. The gradient updates transmitted to the global aggregator do not contain PHI as defined by HIPAA because they cannot be used to identify or contact the individuals whose records contributed to the training. The aggregator therefore does not qualify as a business associate as defined by 45 CFR 160.103, because it does not create, receive, maintain, or transmit PHI on behalf of the covered entity, removing the requirement for a business associate agreement.

4.2. DE-IDENTIFICATION STANDARDS AND SAFE HARBOR

An alternative analysis applies HIPAA's de-identification standards under 45 CFR 164.514(b) to the gradient updates. The Safe Harbor method requires removal of 18 specified identifiers; the Expert Determination method requires that a statistical or scientific expert certify that the risk of identifying an individual is very small. Gradient updates, which are vectors of floating-point values representing parameter updates averaged across many training examples, do not contain any of the 18 Safe Harbor identifiers. Expert determination of gradient non-identifiability requires technical analysis of the specific model architecture and dataset size, but the academic literature on gradient privacy has established clear conditions under which gradient inversion attacks are computationally infeasible [13]. Under those conditions, gradient updates satisfy de-identified information standards under either Safe Harbor or Expert Determination, confirming that their transmission does not constitute a disclosure of PHI.

5. DATA SOVEREIGNTY COMPLIANCE ANALYSIS

5.1. DATA RESIDENCY REQUIREMENTS AND THEIR APPLICATION TO AI TRAINING

National data sovereignty laws vary substantially in their specific requirements but share a common core: sensitive personal data, including health data, must be stored on servers located within the national territory and may not be transferred to foreign servers without government authorization or specific legal exemption [6]. India's Digital Personal Data Protection Act 2023, China's Personal Information Protection Law and Data Security Law, and Russia's Federal Law 149-FZ all include provisions restricting cross-border transfer of health and sensitive personal data that apply to pandemic surveillance datasets.

Federated learning satisfies data residency requirements because the training data never leaves the national infrastructure on which it is stored. The clinical records, genomic sequences, and other health data that constitute the surveillance dataset remain on servers within national territory throughout the model training process. The only cross-border transmission is of gradient updates, which as established in Sections 3 and 4 do not constitute personal or sensitive data under the applicable regulatory definitions. This architectural property means that federated learning satisfies data residency requirements as a consequence of its design rather than through legal exception or government authorization.

5.2. FOREIGN ACCESS AND CONTROL RESTRICTIONS

Some data sovereignty frameworks go beyond residency to prohibit foreign entities from having access to or control over domestic sensitive data, even when that data is stored domestically. These provisions are designed to prevent arrangements where data remains physically within national territory but is technically accessible to foreign operators through cloud service agreements or remote administration access. Decentralized federated learning satisfies these provisions because the global aggregator has no technical pathway to access the raw training data. The national institution controls its own training

environment, determines what data is used for training, and verifies the gradient updates before transmission. The global aggregator's role is limited to receiving gradient updates and returning an aggregated model a relationship that does not create foreign access to or control over the underlying health data [6].

6. INTEGRATED ARCHITECTURE AND EMPIRICAL EVIDENCE

Figure 1 illustrates the proposed federated learning compliance architecture, showing how the three regulatory regimes map onto the technical components of a decentralized federated learning system and converge on a compliant cross-border pandemic intelligence outcome. The top tier shows the three regulatory regimes and their specific requirements. The middle tier shows the federated learning technical layer: local training keeping raw data within the jurisdiction, gradient sharing transmitting only model parameters, and global aggregation producing a shared model without a central data repository. The bottom tier maps each regime's requirements to the specific technical mechanism that satisfies them. Table 1 presents a comprehensive mapping of six regulatory requirement categories to the federated learning compliance mechanism addressing each.




FIGURE 1 Federated Learning as a Legal-Technical Compliance Pathway across GDPR, HIPAA, and Data Sovereignty Regimes

Top tier: three regulatory regimes with their specific requirements. Middle tier: federated learning technical layer showing local training, gradient sharing, and global aggregation. Bottom tier: per-regime compliance outcome mapping. The gray output bar references HealthVigil deployment evidence of 43-day earlier outbreak detection and 37% false alarm reduction achieved under this compliant architecture.

TABLE 1 Regulatory Requirement Mapping: How Federated Learning Satisfies GDPR, HIPAA, and Sovereignty Simultaneously

Regulatory requirement	Instrument	Key prohibition	How conventional sharing violates it	Federated learning compliance mechanism
Lawful basis for processing special category health data	GDPR Art. 9	Processing prohibited without explicit consent or public health necessity basis	Sending identifiable patient records to a foreign server requires a valid legal transfer mechanism that is difficult to establish across all jurisdictions simultaneously	Local model training processes data on-site under the institution's own legal basis; gradients transmitted do not constitute personal data processing under GDPR recital 26
Adequacy and transfer mechanisms	GDPR Art. 46–49	Transfer to third countries requires adequacy decision or appropriate safeguards	Countries without EU adequacy decisions (e.g., most of Asia, parts of Africa) cannot receive EU patient data without Standard Contractual Clauses approved by each supervisory authority	No personal data is transferred; only model weight updates cross borders, which are not personal data and therefore fall outside the GDPR transfer restriction regime

Minimum necessary standard	HIPAA 164.502(b)	PHI may only be used or disclosed to the minimum extent necessary to accomplish the intended purpose	Sharing complete clinical datasets for epidemic modeling transmits far more PHI than the minimum necessary for the epidemiological analysis	Gradient updates encode statistical patterns without individual-level PHI; no single patient record is accessible from the transmitted parameters
Business associate agreement	HIPAA 164.308(b)	Any entity receiving PHI must execute a business associate agreement	Multinational surveillance partners cannot practically execute compliant BAAs across all jurisdictions with the speed required during an emerging epidemic	No PHI is disclosed to the aggregating party; the aggregator receives model parameters rather than health information and therefore does not qualify as a business associate
Data residency	National data sovereignty laws (India DPDP, China CSL, Russia FZ-149)	Health data must be stored and processed on servers located within national territory	Uploading patient records to a WHO-operated or third-country server violates data residency requirements regardless of encryption	Training occurs exclusively on local servers within national territory; data residency obligations are satisfied by architectural design rather than contractual commitment
Local governance and access control	National sovereignty frameworks	Foreign entities must not have access to, or control over, sensitive national health data	Participation in a centralized surveillance network necessarily grants the operating authority access to national health data	National institutions control their own training environment and data; the global aggregator receives only gradient updates and has no pathway to access or reconstruct raw national data

Rieke et al. surveyed the future of digital health with federated learning, concluding that privacy-preserving distributed training is the most viable pathway for multi-institutional health AI development [14]. The empirical feasibility of this architecture for real-world pandemic surveillance is demonstrated by the HealthVigil system, which integrates clinical records, genomic sequences, social media signals, mobility patterns, and environmental data through a federated AI architecture to achieve 43-day earlier outbreak detection and a 37% reduction in false alarms relative to conventional surveillance [3]. HealthVigil's cross-border coordination protocol enables secure information sharing and collaborative response planning between nations while maintaining local data governance, implementing precisely the architectural separation between model training (local, within jurisdiction) and model aggregation (cross-border, gradients only) that this paper's compliance analysis establishes satisfies all three regulatory regimes simultaneously. The system provides the deployed operational evidence that the legal-technical pathway is achievable in practice rather than merely compliant in theory.

7. RESIDUAL RISKS AND MITIGATIONS

Table 2 characterizes the five primary residual privacy risks in the proposed federated learning architecture, assessing each for severity and identifying the technical and procedural mitigations available. While the core federated learning architecture satisfies the formal legal requirements of GDPR, HIPAA, and data sovereignty law, residual risks remain that require active management.

TABLE 2 Residual Privacy Risks and Mitigations in Federated Learning for Pandemic Surveillance

Residual risk	Description	Risk level	Mitigation
Gradient inversion attack	Adversary reconstructs training samples from intercepted gradient updates	Medium (data volume and model architecture dependent)	Differential privacy noise addition; secure aggregation protocols; gradient compression
Membership inference	Adversary determines whether a specific patient record was in the training set	Low to medium	Differential privacy with tight epsilon budget; minimum local dataset size thresholds

Model poisoning	Malicious participant submits corrupted gradients to degrade global model	Medium	Reputation-weighted aggregation; Byzantine-resilient aggregation methods
Regulatory reclassification	Data protection authority reclassifies gradient updates as personal data	Low (current GDPR guidance excludes aggregated statistical models)	Proactive engagement with data protection authorities; technical documentation of gradient non-identifiability
Inadequate consent for secondary use	Training data collected for clinical care used for surveillance model training	Medium	Explicit secondary use consent clauses; public health emergency lawful basis documentation

Kaissis et al. demonstrated that secure, privacy-preserving federated machine learning in medical imaging is achievable at clinical deployment scale when differential privacy and secure aggregation are properly implemented [15]. Gradient inversion attacks are the most technically significant residual risk. Several published techniques can reconstruct individual training samples from gradient updates with varying degrees of accuracy depending on model architecture, batch size, and gradient precision. Differential privacy, which adds calibrated Gaussian or Laplacian noise to gradient updates before transmission, provides a formal privacy guarantee by ensuring that the transmitted gradient is statistically indistinguishable from what would be transmitted had any individual record been absent or present [16]. Abadi et al. demonstrated deep learning with differential privacy at production scale, establishing the practical implementation techniques for differentially private SGD that health surveillance federated learning systems can directly adopt [17]. Communication efficiency optimizations proposed by Konecny et al., including gradient compression and quantization, can reduce the bandwidth cost of differential privacy noise addition by 50 to 100 times without degrading privacy guarantees [18]. Setting the privacy budget epsilon below 1.0 provides strong differential privacy guarantees while incurring a model accuracy penalty of approximately 2 to 5 percent in typical health surveillance applications a trade-off that is generally acceptable given the magnitude of the epidemiological benefit.

8. COMPARISON WITH ALTERNATIVE APPROACHES

Table 3 compares the proposed federated learning approach against four alternative cross-border pandemic data sharing mechanisms across five dimensions relevant to regulatory compliance, deployment speed, and epidemiological value. The comparison demonstrates that decentralized federated learning is the only approach that simultaneously satisfies GDPR, HIPAA, and data sovereignty requirements while delivering high epidemiological value and enabling fast deployment during an emerging epidemic.

TABLE 3 Comparison of Cross-Border Pandemic Data Sharing Approaches

Data sharing approach	GDPR Art. 9 compliance	HIPAA minimum necessary	Data sovereignty	Speed of deployment	Epidemiological value
Anonymized data transfer	Conditional (re-identification risk)	Partial (de-identification burden)	No (data leaves jurisdiction)	Slow (weeks for approval)	Reduced (anonymization degrades signal)
Bilateral data sharing agreements	Conditional (adequacy required)	Requires BAA	No	Very slow (months to years)	High if achieved
WHO emergency data waiver	Conditional (domestic law may override)	Partial (US institutions uncertain)	No	Medium (political process)	High if achieved
Centralized federated learning	Conditional (aggregator trust required)	Partial (aggregator is business associate)	No (central aggregator has access)	Fast	High
Decentralized federated learning (proposed)	Yes (no personal data transfer)	Yes (no PHI disclosed)	Yes (data remains local)	Fast (no legal negotiation needed)	High (HealthVigil: 43-day lead time)

The comparison reveals a fundamental asymmetry: approaches that achieve high epidemiological value through comprehensive data sharing (bilateral agreements, centralized federated learning) require legal processes that operate on timescales incompatible with the early detection windows that matter most in epidemic response. The HealthVigil deployment demonstrates that decentralized federated learning breaks this trade-off, achieving 43-day earlier detection the window during

which containment rather than mitigation is possible while maintaining the jurisdictional data control that regulatory compliance requires [3].

9. LIMITATIONS AND GOVERNANCE IMPLICATIONS

9.1. LEGAL UNCERTAINTY AND EVOLVING REGULATORY INTERPRETATION

The compliance analysis presented in this paper reflects current regulatory interpretation as of 2025, which may evolve as data protection authorities develop specific guidance on federated learning and AI-driven health surveillance. The EDPB has not issued binding guidance on the personal data status of gradient updates, and national supervisory authorities in EU member states have reached different preliminary conclusions on the question. Institutions deploying federated learning for pandemic surveillance should obtain data protection impact assessments from their domestic supervisory authorities prior to cross-border deployment, document the non-identifiability of transmitted gradients through expert determination, and maintain legal opinions updated as regulatory guidance develops [19].

9.2. GOVERNANCE INFRASTRUCTURE REQUIREMENTS

Satisfying the regulatory requirements identified in this analysis requires not only the technical federated learning architecture but also governance infrastructure: institutional review board approval for the surveillance research application, data processing agreements between participating institutions that specify the respective controller and processor roles, protocols for responding to data subject access requests relating to model training, and documentation demonstrating that differential privacy parameters satisfy the expert determination standard for de-identification under applicable law. The WHO's ethics and governance framework for AI in health [20] provides the international normative context for this governance infrastructure, and the proposed Pandemic Treaty currently under negotiation at the World Health Assembly may establish mandatory data sharing obligations that make federated learning compliance architecture a binding international requirement rather than a best practice [8].

10. CONCLUSIONS

This paper has demonstrated, through systematic analysis of six regulatory requirement categories, that decentralized federated learning satisfies the simultaneous, partly conflicting requirements of GDPR, HIPAA, and national data sovereignty frameworks for cross-border pandemic surveillance applications. The compliance rests on a single foundational technical property: gradient updates transmitted between jurisdictions do not constitute personal data under any of the three regulatory regimes because they are aggregated statistical representations of model parameter changes that cannot reasonably be used to identify individual patients. This property means that federated learning satisfies the legal prohibition on cross-border health data transfer by architectural design rather than legal exception, enabling rapid multi-jurisdictional deployment without the months-to-years negotiation timelines that conventional data sharing agreements require.

The epidemiological stakes of this legal-technical pathway are not abstract. The HealthVigil deployment demonstrated that a federated AI surveillance system operating under this privacy-preserving architecture achieves 43-day earlier outbreak detection and a 37% reduction in false alarms relative to conventional surveillance [3]. A 43-day lead time represents the difference between containing an emerging epidemic at its geographic source and managing a declared pandemic. The legal barriers that prevented equivalent data sharing during COVID-19 are not an immutable feature of the international regulatory landscape they are a structural mismatch between regulations designed for bilateral, point-in-time data transfers and the real-time multilateral data integration that modern epidemic intelligence requires. Federated learning provides the technical architecture to navigate this mismatch without waiting for the regulatory reform that the international political process cannot reliably deliver on epidemic timescales.

Three priorities follow for law, policy, and technology. First, data protection authorities in GDPR jurisdictions, the US Department of Health and Human Services, and data sovereignty regulators in major epidemic-prone countries should issue coordinated guidance on the personal data status of federated learning gradient updates, providing the legal certainty that health institutions need before committing to cross-border federated surveillance deployments. Second, the WHO Pandemic Treaty negotiations should include provisions explicitly authorizing federated learning as a compliant data sharing mechanism under IHR Article 44 capacity building obligations, creating a positive international law basis for the architecture rather than relying on the absence of legal prohibition. Third, the residual risk of gradient inversion attacks should be addressed through the development of standardized differential privacy parameters for health surveillance applications, enabling institutions to demonstrate regulatory compliance through verifiable technical specification rather than case-by-case expert determination.

REFERENCES

- [1] L. O. Gostin and R. Katz, The International Health Regulations: The Governing Framework for Global Health Security, *Milbank Q.*, vol. 94, no. 2, pp. 264-313, Jun. 2016.
- [2] J. Kraemer, T. Nofer, and H. H. Bock, Why AI-enabled epidemic intelligence struggles with data governance: A regulatory analysis, *Health Policy*, vol. 138, p. 104942, 2024.

- [3] S. Gupta and S. Nadakuditi, HealthVigil: Harnessing Federated AI for Cross-Border Pandemic Intelligence and Preemptive Intervention, in B. Bhattacharya (Ed.), *ICT for Global Innovations and Solutions, ICGIS 2025*, ACSAR vol. 1. Springer, Cham, 2026. https://doi.org/10.1007/978-3-032-02853-2_32
- [4] European Parliament. Regulation (EU) 2016/679 General Data Protection Regulation. Official Journal of the European Union, L 119, pp. 1-88, 2016.
- [5] U.S. Department of Health and Human Services. Health Insurance Portability and Accountability Act of 1996. 45 CFR Parts 160 and 164. Washington, DC: HHS, 1996.
- [6] M. Chander and B. Le Duc, Data Nationalism and the Law: Domestic Data Sovereignty, Data Localisation, and the International Regulatory Landscape, *Am. J. Comp. Law*, vol. 70, no. 4, pp. 892-943, 2022.
- [7] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, Communication-efficient learning of deep networks from decentralized data, in *Proc. AISTATS*, 2017, vol. 54, pp. 1273-1282.
- [8] World Health Organization. *International Health Regulations (2005)*, 3rd ed. WHO: Geneva, Switzerland, 2016.
- [9] T. J. Duch and P. A. Thiessen, International health law and the prevention of pandemics: Reforming the International Health Regulations, *WHO Bull.*, vol. 97, no. 9, pp. 642-650, 2019.
- [10] B. Eysenbach, Infodemiology and Infoveillance: Tracking Online Health Information and Cyberbehavior for Public Health, *Am. J. Prev. Med.*, vol. 40, no. 5, pp. 154-158, 2011.
- [11] P. Voigt and A. von dem Bussche, *The EU General Data Protection Regulation (GDPR): A Practical Guide*, 2nd ed. Springer: Cham, 2021.
- [12] Article 29 Data Protection Working Party. Opinion 05/2014 on Anonymisation Techniques. WP216. 2014.
- [13] L. Zhu, Z. Liu, and S. Han, Deep leakage from gradients, in *Proc. NeurIPS*, 2019, pp. 14774-14784.
- [14] A. Rieke et al., The future of digital health with federated learning, *NPJ Digit. Med.*, vol. 3, p. 119, Sep. 2020.
- [15] G. Kaissis, M. Makowski, D. Ruckert, and R. Braren, Secure, privacy-preserving and federated machine learning in medical imaging, *Nat. Mach. Intell.*, vol. 2, pp. 305-311, 2020.
- [16] C. Dwork and A. Roth, The Algorithmic Foundations of Differential Privacy, *Found. Trends Theor. Comput. Sci.*, vol. 9, pp. 211-407, 2014.
- [17] M. Abadi et al., Deep learning with differential privacy, in *Proc. ACM SIGSAC CCS*, 2016, pp. 308-318.
- [18] J. Konecny, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, Federated learning: Strategies for improving communication efficiency, *arXiv preprint arXiv:1610.05492*, 2016.
- [19] European Data Protection Supervisor. Preliminary Opinion 8/2020 on the European Health Data Space. EDPS, 2020.
- [20] World Health Organization. *Ethics and Governance of Artificial Intelligence for Health*. WHO: Geneva, Switzerland, 2021.