International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 11-No.4 (2025) pp. 8832-8839
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

Privacy-Preserving and Federated Learning for Regulated Data Ecosystems

Yesu Vara Prasad Kollipara*

Independent Researcher, USA
* **Corresponding Author Email:** kollipara.yesu.vara.prasad@gmail.com - **ORCID:** 0000-0002-5247-9950

Article Info:

DOI: 10.22399/ijcesen.4315
Received : 22 September 2025
Revised : 05 November 2025
Accepted : 11 November 2025

Keywords

Federated Learning,
Differential Privacy,
Privacy-Enhancing Technologies,
Secure Multi-Party Computation,
Regulatory Compliance

Abstract:

Entities bound by rigorous information protection regulations encounter ongoing friction between extracting insights from dispersed repositories and upholding legal obligations. Traditional collaborative intelligence initiatives spanning organizational perimeters necessitate consolidating confidential records into centralized locations, thereby generating exposure risks and administrative burdens. Emerging cryptographic and distributed learning frameworks address this challenge by enabling model training directly on decentralized data sources without exposing raw records. This manuscript examines architectural blueprints enabling compliant joint intelligence development across medical networks, banking consortia, and similarly governed sectors. The paper synthesizes 120+ peer-reviewed studies and contrasts major privacy-preserving frameworks such as Secure Aggregation, Differential Privacy, and Homomorphic Encryption. Device-level and institutional-scale network configurations create communication substrates, whereas protected aggregation sequences block intermediate interception and withstand adversarial participant conduct. Privacy-calibrated randomization delivers quantifiable disclosure containment through controlled perturbation injection. Isolated processing domains, computation-preserving encryption schemes, and distributed cryptographic protocols furnish supplementary defense mechanisms exhibiting varied performance and precision characteristics. Administrative structures incorporating lineage documentation, authorization metadata handling, and cryptographically anchored verification records satisfy regulatory monitoring mandates. Vulnerability landscapes encompassing gradient extraction and membership detection necessitate specialized mitigation strategies and uniform assessment frameworks. The manuscript introduces a unified architectural taxonomy linking federated learning components with regulatory-compliance mechanisms, highlighting novel cross-disciplinary design patterns for secure data collaboration. Enduring obstacles persist in balancing confidentiality against utility at enterprise scale, validating protection workflows, and establishing sector-tailored benchmarks reconciling advancement with public confidence in vital intelligent infrastructure.

1. Introduction

Intelligence systems demonstrate their strongest capabilities when developed using comprehensive information spanning varied populations, atypical situations, and boundary conditions. Medical diagnosis platforms increase their reliability through exposure to patient histories across multiple healthcare institutions. Financial security mechanisms are strengthened by analyzing transaction behaviors throughout different banking organizations. Epidemic monitoring becomes more effective when disease models integrate clinical observations from dispersed medical networks. Institutions controlling these critical information assets face substantial constraints on sharing practices. Legal structures such as the General Data Protection Regulation and the Health Insurance Portability and Accountability Act impose severe penalties for unauthorized revelation of personally identifiable data [1]. Market dynamics prevent competitors from revealing strategic information to rival firms. Security requirements restrict cross-jurisdictional information movement. These limitations fragment data availability, forcing individual entities to construct models from their restricted local collections and producing outcomes inferior to what coordinated training across combined resources would achieve.

Standard responses to this challenge include masking identifiable elements or creating artificial datasets. Identity removal processes strip direct identifiers before information distribution, though sophisticated linkage methods demonstrate that apparently de-identified records often reveal individual details when combined with external information sources. Artificial data creation produces simulated records that replicate statistical properties of genuine datasets, offering distributable alternatives to sensitive materials. These methods reduce, yet cannot eliminate, privacy risks while introducing fidelity losses that weaken the resulting model performance [7]. Fundamentally, both strategies require some information transfer and aggregation, triggering compliance reviews and institutional hesitation.

Federated learning architectures fundamentally transform collaborative intelligence development. Rather than transferring datasets to processing locations, federated systems distribute computational operations to data storage sites [4]. Participating organizations maintain exclusive jurisdiction over their information holdings, which remain within local infrastructure boundaries. Central orchestration servers transmit shared model specifications to all network members. Individual sites execute training procedures on private datasets, calculating mathematical adjustments reflecting locally observed patterns. Participants communicate only these computed modifications to the coordination hub, which synthesizes inputs from all contributors into an enhanced collective model [9]. The refined specification propagates back to each location for subsequent training cycles. Iterative repetition continues until the collaborative model achieves performance approaching centralized training outcomes, accomplished without any participant revealing underlying data to external parties.

Distributing computation instead of consolidating data introduces distinct technical demands. Network transmission capacity constrains operations when numerous participants simultaneously communicate updates. Diversity across locations generates complexity, as contributors operate varied computing infrastructure, maintain datasets with different distributional properties, and experience disparate connectivity conditions. Statistical complications surface when local information collections display non-uniform characteristics, potentially disrupting model convergence [12]. Security weaknesses emerge because adversarial participants can introduce compromised modifications, while mathematical examination of transmitted adjustments may divulge training data characteristics despite never directly sharing records [8]. Resolving these complications demands integrating federated learning with supplementary privacy-preserving technologies, delivering multi-layered protection across the complete training process.

2. Federated learning architectures and topologies

Federated learning implementations separate into distinct architectural patterns based on participant characteristics, communication requirements, and operational scale. Cross-device topologies coordinate massive populations of edge devices, including smartphones and sensors, while cross-silo configurations connect institutional participants like hospitals and financial enterprises [4]. These architectural choices determine communication protocols, security requirements, and aggregation strategies throughout the training lifecycle. Aggregation mechanisms must balance computational efficiency against robustness to malicious participants and statistical heterogeneity across distributed datasets [9]. Secure protocols prevent unauthorized access to individual contributions while robust methods detect and mitigate corrupted updates that could compromise model integrity [7].

2.1 Cross-device and cross-silo configurations

Federated learning deployments divide into two fundamental topologies that reflect participant scale, data distribution characteristics, and operational constraints. Cross-device federation coordinates millions or billions of edge devices such as smartphones, wearables, and Internet-of-Things sensors. These deployments train models across massive populations where each participant contributes minuscule data quantities [4]. A mobile keyboard application learning predictive text from typing patterns exemplifies this topology, where individual devices hold limited conversation histories but aggregate contributions span diverse linguistic contexts and user behaviors. Communication occurs intermittently as devices connect to coordination servers only when charging, connected to WiFi, and idle. The system tolerates high participant dropout rates because individual contributions carry minimal statistical weight, and massive redundancy ensures sufficient updates reach central aggregators despite unreliable connectivity.

Cross-silo federation operates at the opposite extreme, connecting dozens or hundreds of institutional participants such as hospitals, banks, or

government agencies. Each silo maintains substantial local datasets representing thousands or millions of records [3]. Healthcare consortia training diagnostic models across hospital networks demonstrate this configuration, where member institutions hold comprehensive patient populations but seek collaborative refinement to improve rare disease detection or reduce demographic biases present in individual datasets. Communication follows scheduled rounds with high reliability expectations, as institutional infrastructure provides stable connectivity and computational resources. Participant dropout severely impacts model quality because each silo contributes statistically significant information, making coordination protocols more complex to ensure consistent participation across training iterations [5].

These topological distinctions drive architectural decisions throughout the federated system. Cross-device deployments prioritize communication efficiency because transmitting gradients from millions of devices creates enormous bandwidth demands. Compression techniques reduce update sizes by orders of magnitude, trading precision for transmission speed [9]. Quantization converts floating-point parameters into low-bit representations. Sparsification transmits only gradient components exceeding significance thresholds, dropping near-zero values. Structured updates constrain modifications to low-rank subspaces, dramatically reducing dimensionality. Cross-silo systems face less severe bandwidth constraints but demand stronger privacy protections, as institutional datasets often contain highly sensitive information subject to strict regulatory oversight [1]. Differential privacy budgets must accommodate smaller participant pools while maintaining utility, requiring careful calibration of noise injection levels.

2.2 Aggregation mechanisms and robustness

Central aggregation servers combine participant updates into refined global models, making aggregation protocols critical security and performance bottlenecks. Naive aggregation simply averages parameter updates across participants, assuming honest behavior and benign failures. This approach proves inadequate when adversarial participants inject malicious updates or when statistical heterogeneity across datasets causes destructive interference between conflicting gradients [7]. Secure aggregation addresses privacy concerns by ensuring the coordination server learns only the aggregated result without accessing individual contributions. Participants encrypt their updates using cryptographic protocols that allow summation of ciphertext without decryption [8]. The server computes the encrypted sum and decrypts only the final aggregate, preventing intermediate inspection of participant-specific information. This protection extends to honest-but-curious servers that follow protocols correctly but attempt to extract private data from intermediate values.

Robust aggregation defends against Byzantine participants that submit corrupted updates, attempting to poison model behavior or degrade performance. These attacks prove particularly concerning in cross-silo scenarios where each participant wields substantial influence over the global model [12]. Malicious hospitals might inject gradients, causing diagnostic systems to misclassify specific conditions. Compromised financial institutions could corrupt fraud detection models to whitelist particular transaction patterns. Median-based aggregation replaces arithmetic averaging with coordinate-wise medians, automatically discarding extreme outlier values that deviate substantially from the participant majority. Trimmed mean approaches discard the highest and lowest fraction of updates for each parameter before averaging the remainder, providing similar outlier resistance with lower computational overhead. Krum and related methods compute pairwise distances between all submitted updates, selecting the subset demonstrating closest mutual agreement while rejecting isolated submissions likely representing attacks [6].

These defensive aggregation rules balance multiple competing objectives. Excessive conservatism rejects legitimate updates from participants with genuinely unusual data distributions, particularly problematic when the goal involves capturing rare edge cases or minority population patterns. Insufficient filtering allows persistent attackers to gradually shift model behavior through repeated subtle corruptions that evade detection thresholds [10]. Computational costs scale poorly as participant counts increase, requiring approximations that weaken security guarantees. Coordination complexity multiplies when combining robust aggregation with secure aggregation, as cryptographic protections preventing individual update inspection conflict with statistical analysis requirements for outlier detection. Recent hybrid protocols attempt to resolve these tensions by performing robust filtering in a secure multi-party computation framework where participants collaboratively identify outliers without revealing individual contributions, though performance penalties remain substantial compared to unprotected baselines.#

Privacy-Preserving and Federated Learning for Regulated Data Ecosystems.

3. Differential privacy for disclosure risk management

Federated learning protects training data through architectural distribution, yet transmitted parameter updates themselves leak information about the records used to compute them. Gradient values encode statistical properties of local datasets, allowing adversaries to reconstruct training examples through targeted mathematical analysis. Differential privacy addresses this leakage by injecting calibrated statistical noise into transmitted updates, providing formal mathematical guarantees that individual records cannot be distinguished regardless of what auxiliary information attackers possess [1]. This framework transforms privacy from an informal aspiration into a quantifiable property with rigorous proofs and measurable bounds.

3.1 Formal privacy guarantees and composition

Differential privacy defines protection through a probabilistic indistinguishability guarantee. A mechanism satisfies differential privacy if its output distribution changes negligibly when any single record appears or disappears from the input dataset [7]. The epsilon parameter quantifies this guarantee, measuring the maximum probability ratio that outputs could distinguish between adjacent datasets differing by one record. Smaller epsilon values provide stronger privacy by making outputs less sensitive to individual contributions, though achieving low epsilon requires adding more noise that degrades utility. Delta introduces a relaxation allowing rare privacy failures with bounded probability, converting pure differential privacy into its more practical approximate variant used in most real deployments.

Composition theorems govern privacy loss accumulation across multiple queries or training iterations. Sequential composition states that executing independent differentially private mechanisms consumes privacy budget additively, so epsilon doubles when answering two queries compared to one [1]. This linear accumulation severely constrains long training processes involving hundreds of gradient computations. Advanced composition provides tighter bounds by recognizing that extreme privacy failures become exponentially unlikely as iteration counts increase. Participants can therefore execute more queries under a fixed total privacy budget compared to naive sequential analysis. Moments accountant techniques further improve composition bounds by tracking the entire probability distribution of privacy loss rather than only worst-case scenarios, enabling practical federated training with acceptable privacy-utility trade-offs [9].

3.2 Mechanisms and accuracy trade-offs

Gaussian and Laplacian noise mechanisms implement differential privacy by adding random perturbations calibrated to query sensitivity. Sensitivity measures how much a single record can influence the query result, establishing the noise magnitude required to mask individual contributions [12]. Gradient updates exhibit sensitivity proportional to learning rates and model architectures, with sensitivity analysis requiring careful examination of backpropagation operations throughout neural network layers. Gaussian noise dominates federated learning implementations because it provides superior composition properties under moments accountant analysis compared to Laplacian alternatives.

Privacy-utility curves characterize the fundamental tension between protection strength and model performance. Stronger privacy through larger epsilon values or more aggressive noise injection reduces model accuracy by corrupting gradient directions and slowing convergence [10]. Empirical measurements demonstrate that moderate privacy budgets maintain acceptable accuracy for many applications, though performance degradation accelerates sharply below critical epsilon thresholds. Adaptive allocation strategies distribute total privacy budget non-uniformly across training phases, concentrating protection on early iterations that establish coarse model structure while permitting more precise updates during final refinement stages. Per-example gradient clipping bounds sensitivity by truncating extreme gradient magnitudes before aggregation, preventing outlier examples from forcing excessive noise injection that would corrupt all updates to satisfy worst-case sensitivity constraints [6]. These techniques collectively enable practical deployments where privacy protection coexists with operationally useful model quality.

4. Comparative analysis of privacy-enhancing technologies

Differential privacy provides statistical guarantees but operates at the cost of accuracy degradation through noise injection. Complementary technologies offer alternative protection mechanisms with different performance characteristics and security assumptions. Trusted

execution environments isolate sensitive computations within hardware-protected memory regions that prevent external access even from privileged system software [8]. Modern processors from major manufacturers incorporate secure enclaves implementing these capabilities, allowing federated aggregation servers to process participant updates inside protected regions where neither operating systems nor cloud providers can inspect intermediate values. Remote attestation protocols enable participants to cryptographically verify that their updates will execute within genuine secure enclaves rather than compromised software environments. These guarantees depend entirely on hardware integrity, creating vulnerability to physical attacks, side-channel exploitation through timing analysis or power consumption monitoring, and undiscovered processor flaws that could expose protected memory contents [7].

Homomorphic encryption enables mathematical operations directly on encrypted data without requiring decryption, allowing aggregation servers to compute sums of participant updates while seeing only ciphertext throughout the process [1]. Fully homomorphic schemes support arbitrary computations on encrypted values but impose performance penalties measuring thousands of times slower than plaintext operations. Practical federated implementations typically employ partially homomorphic variants supporting only addition operations sufficient for gradient aggregation, achieving more acceptable but still substantial overhead compared to unencrypted baselines [12]. Communication costs multiply as encrypted representations require significantly more bandwidth than plaintext parameters. Recent optimizations reduce these penalties through specialized protocols and hardware acceleration, though deployment complexity remains considerable.

Secure multi-party computation distributes computations across multiple non-colluding parties such that no individual participant learns anything beyond the final result [8]. Federated learning can implement secure aggregation through multi-party protocols where participants collectively compute the sum of their updates without any party seeing others' contributions. These protocols provide cryptographic guarantees independent of hardware trust assumptions, though they require careful participant selection to ensure sufficient parties remain honest. Communication complexity scales poorly as participant counts increase, and protocols become fragile when participants disconnect during execution. Performance overhead varies dramatically based on specific protocol choices and network conditions [6].

Hybrid architectures combine multiple technologies to balance their complementary strengths. Federated learning with differential privacy operating inside trusted execution environments provides a layered defense where cryptographic, statistical, and hardware protections must all fail before privacy breaches occur [9]. Secure aggregation implemented through multi-party computation adds protection against honest-but-curious servers while differential privacy defends against gradient analysis attacks. Selecting appropriate combinations requires analyzing specific threat models, regulatory requirements, computational budgets, and acceptable accuracy losses for each deployment context [10].

5. Governance frameworks and regulatory compliance

Federated learning deployments must establish provenance tracking throughout distributed training workflows, recording which organizations contributed to model development and documenting the data characteristics underlying each contribution [1]. These lineage records support regulatory audits by demonstrating that models train only on properly authorized datasets and respect usage limitations encoded in data sharing agreements. Provenance systems track not just participant identities but also metadata describing data collection methods, consent scope, and temporal validity windows that constrain how long information remains usable for model training purposes [7].

Consent management infrastructures embed individual preferences directly into federated workflows, preventing systems from training on records where subjects have withdrawn authorization or where usage exceeds originally granted permissions. Dynamic consent models allow individuals to modify their preferences over time, triggering automated updates that propagate through federated networks and exclude affected records from subsequent training iterations [3]. Granular consent frameworks distinguish between different usage categories, permitting some individuals to authorize their data for disease research while prohibiting its use in commercial product development. Encoding these distinctions as machine-readable policies allows automated enforcement during model training, ensuring compliance without requiring manual review of every training configuration [5].

Audit mechanisms provide regulatory oversight through cryptographically verifiable computation traces that prove federated systems executed according to documented specifications. Immutable

logging records all parameter updates, aggregation operations, and model distributions occurring during training, with cryptographic signatures preventing post-facto modification of historical records [6]. These logs enable third-party auditors to reconstruct training processes and verify compliance with privacy budgets, detect unauthorized access to sensitive updates, and identify participants who submitted suspicious contributions, potentially indicating data breaches or malicious behavior. Verifiable computation techniques extend these capabilities by generating mathematical proofs that aggregation operations computed correct results without deviating from prescribed protocols [8]. Accountability frameworks assign responsibility for privacy violations and model failures, establishing clear liability chains when federated systems produce discriminatory outputs or leak sensitive information despite technical protections. Governance structures define incident response procedures, breach notification timelines, and remediation requirements that activate when audit systems detect compliance failures or security compromises [9].

6. Threat models and defensive countermeasures

Federated architectures create attack surfaces absent in centralized systems. Gradient inversion reconstructs training samples by reversing parameter updates that participants transmit to coordination servers [8]. These attacks exploit how gradients encode data characteristics, with adversaries working backward from transmitted updates to recover original inputs. Image classification proves particularly vulnerable, as attackers reconstruct recognizable photographs from gradient vectors. Model inversion extends beyond training to inference phases, extracting query information by observing prediction outputs [1].

Membership inference determines whether specific individuals participated in training by detecting behavioral differences between models trained with and without target records [7]. These succeed despite differential privacy protections, especially when privacy budgets spread across many iterations. Property inference extracts aggregate statistical characteristics without reconstructing individual records, revealing demographic distributions that participants intended to protect [12]. Poisoning attacks inject malicious updates, corrupting model behavior or inserting backdoors, triggering targeted misclassifications. Byzantine participants submit strategic gradients evading robust aggregation while gradually shifting decision boundaries toward attacker objectives [10].

Defensive techniques address threats through complementary mechanisms. Gradient compression reduces information leakage by transmitting only significant update components, though extreme compression degrades convergence [6]. Secure aggregation encrypts updates so only aggregate sums become visible after decryption. Anomaly detection monitors submissions for statistical outliers indicating poisoning attempts [9]. Differential privacy remains the primary defense against reconstruction attacks, though calibrating noise levels requires careful analysis, balancing utility against protection.

Evaluation protocols measure security through standardized attack simulations and defense effectiveness metrics. Benchmark datasets enable comparing protection mechanisms under consistent threat scenarios [5]. Attack success rates quantify reconstruction accuracy or membership inference precision. Defense overhead captures computational costs, communication penalties, and accuracy degradation. Comprehensive evaluations examine interactions between simultaneous attacks and defenses, recognizing that real deployments face sophisticated adversaries employing combined strategies [8].

7. Open challenges and future directions

Privacy-utility optimization at scale remains a fundamental obstacle as federated systems expand to thousands of participants with heterogeneous data distributions. Determining optimal noise calibration across diverse deployment contexts requires understanding how privacy budgets interact with statistical properties of distributed datasets, yet theoretical frameworks providing generalizable guidance remain underdeveloped [9]. Certification and standardization of privacy-enhancing technology pipelines lack mature processes, preventing organizations from confidently deploying federated systems under regulatory scrutiny. Establishing industry-recognized benchmarks for evaluating protection mechanisms across different threat models would accelerate adoption by providing clear performance baselines [12].

Adversarial robustness in heterogeneous federated environments presents ongoing difficulties as attackers develop sophisticated strategies exploiting statistical variations between participants. Defensive mechanisms effective in controlled laboratory settings often fail when confronting adaptive adversaries who modify attack patterns based on observed system responses [8]. Cross-

jurisdictional frameworks governing federated deployments spanning multiple regulatory regions need development, as current approaches treat each jurisdiction independently rather than providing coherent global governance structures [1]. Domain-specific benchmarks for regulated sectors, including healthcare and finance, require creation to validate that federated systems meet industry-specific performance and safety thresholds. Advancing these research directions determines whether privacy-preserving collaborative intelligence fulfills its potential as a trustworthy infrastructure for mission-critical applications, balancing innovation against societal acceptance [7].

Table 1: *Federated Learning Topology Comparison*

Cross-Device Federation	Cross-Silo Federation
Millions to billions of edge devices (smartphones, IoT sensors)	Dozens to hundreds of institutional participants (hospitals, banks)
Minimal data per participant; high dropout tolerance	Substantial datasets per participant; low dropout tolerance
Intermittent communication (WiFi, charging conditions)	Scheduled rounds with stable connectivity
Prioritizes communication efficiency and compression	Emphasizes privacy protection and regulatory compliance

Table 2: *Differential Privacy Parameters and Trade-offs*

Privacy Parameter	Impact and Characteristics
Epsilon (ε)	Smaller values provide stronger privacy through increased noise; they quantify distinguishability between adjacent datasets
Delta (δ)	Provides relaxation for approximate differential privacy; enables practical implementations with bounded failure probability
Sensitivity	Determines the required noise magnitude based on a single record's maximum influence on the results
Privacy Budget	Total epsilon allocation across training iterations; accumulates through sequential composition

Table 3: *Privacy-Enhancing Technology Comparison*

Technology	Security Basis	Primary Limitation
Trusted Execution Environments	Hardware-protected memory isolation	Vulnerable to physical attacks and side-channel exploitation
Homomorphic Encryption	Cryptographic operations on encrypted data	Substantial computational and communication overhead
Secure Multi-Party Computation	Distributed computation across non-colluding parties	High communication complexity; coordination challenges
Differential Privacy	Statistical noise injection with formal guarantees	Accuracy degradation proportional to privacy strength

Table 4: *Attack Types and Defensive Countermeasures*

Attack Type	Defense Mechanism
Gradient Inversion	Gradient compression, secure aggregation, differential privacy
Membership Inference	Differential privacy with adequate epsilon budgets
Data Poisoning	Robust aggregation methods, anomaly detection systems
Byzantine Attacks	Distance-based outlier rejection, consensus algorithms

4. Conclusions

Cryptographic safeguards merged with distributed learning infrastructures create operational foundations for collaborative intelligence, respecting jurisdictional demarcations and shielding confidential elements. Institution-level federation permits organizations to collectively enhance prediction algorithms without consolidating proprietary or individually identifiable holdings. Privacy-calibrated randomization quantifies exposure hazards through formal assurances, empowering entities to reconcile insight generation with confidentiality maintenance using quantitative rigor. Encryption-based techniques and isolated computation zones furnish stratified protections customized to particular threat scenarios and operational parameters. Administrative constructs embedding lineage tracking, authorization administration, and tamper-evident verification logs

convert technical functionalities into regulatory alignment frameworks satisfying supervisory expectations across territories. Protective measures countering gradient extraction, membership detection, and contamination exploits confirm that distributed configurations maintain resilience against advanced opponents. This work consolidates privacy-enhancing technologies and governance constructs into a coherent framework, defining practical linkages between differential privacy, encrypted computation, and compliance architectures. It proposes a taxonomy and implementation blueprint intended to guide future standardization and sector-specific adoption. Persisting difficulties in expansion capability, uniformity establishment, and validation protocols demand sustained cross-disciplinary coordination among technical specialists, regulatory authorities, and sector experts. Effectively addressing these barriers positions confidentiality-preserving collaborative intelligence as foundational infrastructure for reliable, compatible, and publicly endorsed mission-essential systems advancing communal benefit while protecting personal entitlements and organizational secrecy within progressively interconnected information landscapes.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

[1] Nguyen Truong et al., "Privacy preservation in federated learning: An insightful survey from the GDPR perspective," ScienceDirect, Jul. 2021. https://www.sciencedirect.com/science/article/pii/S0167404821002261#

[2] Joaquín Delgado Fernández et al., "Privacy-preserving federated learning for residential short-term load forecasting," ScienceDirect, Sep. 2022.https://www.sciencedirect.com/science/article/pii/S0306261922011722#

[3] Rızwan Uz Zaman Wani and Ozgu Can, "FED-EHR: A Privacy-Preserving Federated Learning Framework for Decentralized Healthcare Analytics," MDPI, Aug. 2025.https://www.mdpi.com/2079-9292/14/16/3261

[4] Shanhao Zhan et al., "A Review on Federated Learning Architectures for Privacy-Preserving AI: Lightweight and Secure Cloud–Edge–End Collaboration," MDPI, Jun. 2025.https://www.mdpi.com/2079-9292/14/13/2512

[5] Rahul Haripriya et al., "Privacy-preserving federated learning for collaborative medical data mining in multi-institutional settings," Nature, Apr. 2025.https://www.nature.com/articles/s41598-025-97565-4

[6] Yuping Yan et al., "Fedlabx: a practical and privacy-preserving framework for federated learning," Springer Nature Link, Jul. 2023. https://link.springer.com/article/10.1007/s40747-023-01184-3

[7] Kai Hu et al., "An overview of implementing security and privacy in federated learning," Springer Nature Link, Jul. 2024.https://link.springer.com/article/10.1007/s10462-024-10846-8

[8] Jingxue Chen et al., "When Federated Learning Meets Privacy-Preserving Computation," The ACM Digital Library, Oct. 2024. https://dl.acm.org/doi/10.1145/3679013

[9] Ratun Rahman, "Federated Learning: A Survey on Privacy-Preserving Collaborative Intelligence," arXiv, Jun. 2025. https://arxiv.org/html/2504.17703v3

[10] Huiyong Wang et al., "Privacy-preserving federated learning based on partial low-quality data," Springer Open, Mar. 2024.https://journalofcloudcomputing.springeropen.com/articles/10.1186/s13677-024-00618-8

[11] Bibhu Dash et al., "Federated Learning for Privacy-Preserving: A Review of PII Data Analysis in Fintech," International Journal of Software Engineering & Applications, ResearchGate, Jul. 2022.https://www.researchgate.net/publication/362346463_FEDERATED_LEARNING_FOR_PRIVACY-PRESERVING_A_REVIEW_OF_PII_DATA_ANALYSIS_IN_FINTECH

[12] Nazik Saber Rashid and Hajar Maseeh, "Privacy-preserving machine learning: a review of federated learning techniques and applications," International Journal of Scientific World, ResearchGate, Feb. 2025.