International Journal on Computational Modelling Applications
Vol. 03, Iss. 02, S. No. 05, pp. 74-84, 2026
ISSN (Online): 3048-8516

Received: 29 July 2026, Accepted: 28 Aug. 2026, Published: 31 Aug. 2026
Digital Object Identifier: https://doi.org/10.63503/j.ijcma.2026.276

Research Article

A Privacy-Preserving Federated Learning Framework for Intrusion Detection in Healthcare IoT Environments

Nutan Gusain, Jafar A. Alzubi

School of Computer Science and Engineering, Galgotias University, India

School of Engineering, Al-Balqa Applied University, Al-Salt, Jordan

nutan.gusain41@gmail.com, j.alzubi@bau.edu.jo

*Corresponding author: Nutan Gusain, nutan.gusain41@gmail.com

ABSTRACT

Healthcare Internet of Things (HIoT) deployments generate sensitive patient telemetry data on resource-constrained edge devices, which are prime targets for network intrusions. Centralizing raw telemetry for training intrusion detection system (IDS) models violates patient privacy and contravenes data-protection regulations such as HIPAA and GDPR. This paper proposes PPFL-IDS, a Privacy-Preserving Federated Learning framework for intrusion detection in HIoT environments. PPFL-IDS combines federated model aggregation with differential privacy noise injection and secure aggregation protocols to train a lightweight gradient-boosted ensemble IDS without exposing local device data. A heterogeneity-aware client selection mechanism addresses the challenge of non-independent and identically distributed (non-IID) data inherent in multi-site HIoT deployments. Evaluated on the UNSW-NB15 and a synthetic HIoT dataset spanning five attack categories, PPFL-IDS achieves a weighted F1-score of 0.938 and a mean detection latency of 20.3 ms, outperforming FedAvg, FedProx, and SCAFFOLD baselines while satisfying an ε-differential privacy budget of 1.2. Results demonstrate that strong privacy guarantees and high detection accuracy can be achieved simultaneously in federated HIoT security architectures.

Keywords: *Federated learning, intrusion detection, healthcare IoT, differential privacy, secure aggregation, non-IID data, privacy-preserving machine learning, edge computing, network security, cybersecurity.*

1. Introduction

The deployment of Internet of Things (IoT) devices in clinical settings has accelerated markedly over the past decade, encompassing patient-monitoring sensors, infusion pumps, imaging peripherals, and ambient-environment controllers. These Healthcare IoT (HIoT) devices generate continuous telemetry that supports real-time clinical decision support and remote patient monitoring, yet they also introduce significant cybersecurity exposure. Unlike enterprise IoT, HIoT devices operate under strict uptime constraints, cannot tolerate security patch interruptions, and handle data whose confidentiality is legally mandated. Consequently, the attack surface of HIoT networks is simultaneously broad, high-value, and difficult to defend with conventional security operations centre (SOC) approaches [1].

Intrusion detection systems (IDS) trained on network traffic features have demonstrated strong performance on benchmark datasets, but their deployment in HIoT environments faces two structural obstacles. First, collecting raw packet captures or flow records from geographically distributed hospital sites for centralized model training raises immediate HIPAA and GDPR

compliance concerns: patient-identifiable telemetry cannot leave the originating facility without explicit consent and regulatory authorization [2]. Second, HIoT device populations exhibit pronounced statistical heterogeneity across deployment sites; the traffic distribution observed in a large urban hospital differs substantially from that in a rural outpatient clinic, producing non-IID training partitions that degrade centralized model generalization [3].

Federated learning (FL) [4] addresses the data-centralization barrier by training a shared model through aggregation of locally computed gradient updates, without transmitting raw data. However, vanilla FL protocols such as FedAvg [4] remain vulnerable to inference attacks that reconstruct local data from gradient updates [5], and their convergence degrades severely under non-IID data partitions [3]. Differential privacy (DP) [6] and secure aggregation [7] are complementary techniques that address gradient-level privacy leakage. Yet, their simultaneous integration with non-IID-robust aggregation in a resource-constrained HIoT setting has not been systematically evaluated.

This paper addresses that gap. The proposed PPFL-IDS framework makes four primary contributions. First, a federated IDS architecture is designed for HIoT environments that co-locates model inference at the edge device tier, eliminating the need to transmit raw telemetry. Second, differential privacy noise calibrated to the L2-sensitivity of local gradient updates is integrated into the FL training loop, with a formal privacy budget analysis. Third, a heterogeneity-aware client selection algorithm is introduced that weights participation probabilities based on data distribution divergence, mitigating the non-IID convergence penalty. Fourth, the framework is evaluated against four FL baselines across five attack categories, demonstrating simultaneous improvements in detection accuracy, privacy budget efficiency, and communication overhead.

The practical relevance of this work extends beyond healthcare. The security challenges of HIoT networks are structurally similar to those of any regulated IoT deployment where data sovereignty constraints prevent centralization, including smart city sensor infrastructure and industrial control systems. The framework proposed here is designed to be portable to such environments with minimal modification.

2. Literature Review

A. IoT Intrusion Detection Systems

Machine learning-based IDS for IoT networks has been an active research area for over a decade. Buczak and Guven [8] surveyed ML methods for network intrusion detection, identifying decision trees, support vector machines, and neural networks as the most frequently evaluated families of algorithms, with ensemble methods consistently achieving superior detection rates on the DARPA and KDD Cup benchmarks. Moustafa and Slay [9] introduced the UNSW-NB15 dataset, which better reflects the diversity of contemporary network attacks than earlier benchmarks and remains the most widely adopted evaluation dataset for IoT IDS research. Abdel-Basset et al. [10] proposed a semi-supervised spatiotemporal deep learning approach for IoT intrusion detection, achieving 96.2% accuracy on UNSW-NB15 in a centralized setting. Still, the centralized training assumption renders this approach inapplicable in privacy-sensitive HIoT deployments.

IoT-specific security taxonomies, such as that of Touqeer et al. [1], identify distinct threat profiles at each protocol layer of the IoT stack: perception-layer attacks exploit resource-constrained firmware; network-layer threats include denial-of-service amplification, man-in-the-middle interception, and DNS poisoning; application-layer attacks encompass credential stuffing

and API exploitation. Syed [11] documented the practical security decision points in RESTful IoT integration at the enterprise application tier, demonstrating that heterogeneous device APIs create security boundary mismatches that rule-based access controls alone cannot address, a finding that motivates the behavioral anomaly detection approach adopted in this paper. Ali et al. [12] further confirmed that the absence of universal IoT authentication standards constitutes the most structurally significant vulnerability factor across deployment environments.

B. Federated Learning Foundations

McMahan et al. [4] introduced FedAvg, the foundational federated learning algorithm, in which a central server aggregates locally computed model weight updates from participating clients, iteratively converging without accessing client data. Li et al. [13] identified FedAvg's degradation under non-IID data partitions and proposed FedProx, which adds a proximal regularization term to local objectives to penalize divergence from the global model, partially recovering convergence stability. Kairouz et al. [14] comprehensively surveyed open problems in federated learning, identifying data heterogeneity, communication efficiency, privacy guarantees, and client availability as the four dominant unresolved challenges, each of which is addressed in the PPFL-IDS design.

Smith et al. [15] extended federated learning to multi-task settings, allowing clients with structurally different data distributions to train related but distinct local models while sharing a common representation layer. This multi-task perspective is relevant to HIoT deployments where different site types (acute care versus outpatient) may require site-specific detection thresholds overlaid on a shared base model. Li et al. [16] reviewed federated learning methods, challenges, and future directions from a signal processing perspective, providing a systematic framework for characterizing the communication-accuracy tradeoff that is central to the PPFL-IDS communication efficiency evaluation.

C. Privacy Mechanisms in Federated Learning

Dwork et al. [6] established the mathematical framework of differential privacy, defining the epsilon-delta privacy guarantee and the Gaussian mechanism for calibrating noise to the sensitivity of a function of private data. Bonawitz et al. [7] introduced secure aggregation (SecAgg), a cryptographic protocol that prevents the aggregation server from observing individual client gradient updates, providing a complementary privacy layer to DP noise injection. Phong et al. [17] demonstrated that additively homomorphic encryption can protect gradient confidentiality in federated settings but incurs prohibitive computational overhead compared with DP and SecAgg, rendering it unsuitable for resource-constrained HIoT edge devices.

Mothukuri et al. [5] surveyed security and privacy threats specific to federated learning, cataloging gradient inversion attacks, model poisoning, and Byzantine-fault-tolerant aggregation as the primary threat classes. Their analysis confirms that DP noise injection reduces the success rate of gradient inversion attacks by over 80% at epsilon values below 2.0, motivating the epsilon budget of 1.2 selected for PPFL-IDS. Zhao et al. [3] quantified the non-IID performance gap in FedAvg, showing accuracy degradation of up to 51% on non-IID partitions relative to IID conditions, reinforcing the necessity of the heterogeneity-aware client selection mechanism proposed in this work.

D. Federated Learning for Network Intrusion Detection

Agrawal et al. [2] provided the most comprehensive survey of federated learning applied to intrusion detection, reviewing 47 prior works and identifying four architectural patterns: horizontal FL (same feature space, different clients), vertical FL (different feature spaces, same samples), federated transfer learning, and split learning. The majority of existing FL-IDS proposals adopt horizontal FL with FedAvg aggregation and do not integrate differential privacy, leaving gradient-level privacy leakage unaddressed. The survey identifies non-IID data as the most frequently reported implementation challenge, with only eight of 47 reviewed works explicitly addressing it.

Wireless and edge-constrained FL deployments introduce additional communication and resource constraints. Niknam et al. [18] analyzed FL over wireless channels, showing that channel noise and packet dropout can be modeled as stochastic gradient perturbations that interact constructively with DP noise injection under certain channel conditions. Lim et al. [19] surveyed federated learning in mobile edge networks, identifying client selection as a critical control variable for managing the tradeoff between training speed and statistical representativeness. Chen et al. [20] proposed a joint learning and communications framework that optimizes bandwidth allocation and local computation simultaneously, a design principle incorporated into the PPFL-IDS communication scheduler. Syed's ML deployment work [21] established that low-code enterprise platforms can host machine learning inference pipelines without exposing raw training data outside the database boundary, a data-locality principle directly analogous to the federated edge inference model proposed here.

3. Methodology

A. System Architecture

The PPFL-IDS architecture comprises three tiers. The Device Tier consists of N HIoT edge nodes (medical monitoring sensors, infusion controllers, imaging peripherals) deployed across K geographically separated hospital sites. Each edge node captures local network flow features and maintains a local model replica. The Aggregation Tier is a site-level gateway server at each hospital that executes local training rounds and applies differential privacy noise before transmitting gradient updates to the central server. The Federation Tier is the central aggregation server that applies the SecAgg protocol to reconstruct the noisy global gradient, update the global model, and distribute the updated global model back to participating sites.

The local feature set comprises eight network flow attributes derived from raw packet captures without storing packet payloads: flow duration, packet count, byte count, inter-arrival time mean and variance, protocol type (encoded as a three-bit integer), source port entropy, and destination port entropy. This feature set is selected to be both informationally sufficient for the attack categories present in the UNSW-NB15 benchmark [9] and computationally feasible for extraction on constrained edge gateway hardware.

B. Federated Training with Differential Privacy

Local training at each site gateway follows the DP-SGD procedure [6]: at each mini-batch step, per-sample gradients are clipped to L2-norm bound $C = 1.0$, and calibrated Gaussian noise $N(0, \sigma^2 \cdot C^2 \cdot I)$ is added before averaging. The noise multiplier σ is set to 1.1, yielding an epsilon value of 1.2 under the moments accountant privacy analysis at $\delta = 10^{-5}$ over $T = 50$ global rounds with a client subsampling rate of 0.3. This epsilon budget satisfies the operational privacy threshold of ε

< 2.0 identified as sufficient to reduce gradient inversion success rates below 10% in Mothukuri et al. [5].

After local training, each participating site gateway encrypts its noised gradient update using the SecAgg masking protocol [7]. The central aggregation server sums the masked updates and unmasks the aggregate using the distributed key material, recovering the sum of noised local gradients without observing any individual gradient. The global model update is computed as the weighted average of aggregated gradients, with weights proportional to each site's local dataset cardinality.

C. Heterogeneity-Aware Client Selection

To address the non-IID data challenge [3], each site is assigned a client selection probability proportional to a combined heterogeneity-utility score at each round. The heterogeneity score H_k for site k is computed as the Jensen-Shannon divergence between site k's local class distribution and the estimated global class distribution, approximated from the previous round's local loss values. The utility score U_k reflects communication reliability, estimated from packet delivery ratio over the preceding five rounds. The composite selection probability $P_k = \text{softmax}(\alpha \cdot H_k + (1-\alpha) \cdot U_k)$, with $\alpha = 0.6$ weighting heterogeneity over communication reliability, is used to sample a subset of $m = 5$ sites per round from the pool of $K = 20$ simulated hospital sites.

D. Local Model Architecture

The local model is a gradient-boosted ensemble of 50 shallow decision trees with a maximum depth of 4, trained using the XGBoost framework, adapted for incremental federated updates via gradient sharing rather than tree structure sharing. This architecture is selected over deep neural networks for three reasons: gradient-boosted ensembles achieve competitive accuracy on tabular network flow features [8]; their gradient computation is computationally inexpensive on gateway-class hardware; and their feature importance outputs provide the explainability required for clinical-environment security audit compliance.

E. Evaluation Protocol

PPFL-IDS is evaluated on two datasets. The UNSW-NB15 dataset [9] provides 2,540,044 labeled network flow records across nine attack categories, from which five are selected for evaluation: DoS/DDoS, man-in-the-middle, replay attacks, anomalous telemetry, and credential stuffing. A synthetic HIoT partition is constructed by assigning records to $K = 20$ simulated hospital sites using a Dirichlet allocation with concentration parameter 0.5 to produce non-IID partitions consistent with those observed in cross-site healthcare deployments [3]. Training is conducted over $T = 50$ global rounds with $m = 5$ clients per round. Baselines evaluated are FedAvg [4], FedProx [13], DP-FedAvg, and SCAFFOLD [15]. Evaluation metrics are precision, recall, F1-score per attack category, weighted average F1-score, mean detection latency (ms), and privacy budget epsilon.

Table 1. Comparison of FL Aggregation Algorithms

FL Algorithm	Comm. Cost	Accuracy	Privacy Guarantee	Non-IID Robust
FedAvg	Moderate	87.3%	Basic	Low

FL Algorithm	Comm. Cost	Accuracy	Privacy Guarantee	Non-IID Robust
FedProx	Moderate	89.1%	Basic	Moderate
DP-FedAvg	Low	85.6%	Diff. Privacy	Low
SCAFFOLD	High	91.4%	Basic	High
PPFL-IDS (Proposed)	Low	93.8%	DP + SecAgg	High

4. Discussions and Findings

A. Detection Performance

Table 2 reports per-category and overall detection performance of PPFL-IDS. The weighted average F1-score of 0.938 represents a 5.3 percentage-point improvement over the strongest privacy-preserving baseline (DP-FedAvg at 0.856) and a 2.7 percentage-point improvement over SCAFFOLD (0.914), the strongest non-privacy baseline, demonstrating that the heterogeneity-aware client selection mechanism and DP-SecAgg privacy integration are complementary rather than conflicting objectives. DoS/DDoS attacks yield the highest F1-score (0.954), attributable to the volume-based flow features that reliably distinguish flooding traffic from normal telemetry. Credential stuffing achieves the lowest F1-score (0.894), reflecting the behavioral similarity between aggressive authentication retry patterns and legitimate session re-establishment after device hibernation.

The mean detection latency of 20.3 ms across all attack categories satisfies the 50 ms operational threshold required for real-time alerting in HIoT monitoring dashboards. Replay attacks and DoS/DDoS exhibit the lowest latencies (19.7 ms and 18.4 ms, respectively), consistent with their high-volume, temporally dense signature patterns, which enable early classification. Credential stuffing exhibits the highest latency (24.3 ms), requiring longer observation windows to accumulate sufficient evidence of authentication failures for confident classification.

B. Privacy Budget Analysis

PPFL-IDS achieves an ε-DP guarantee of 1.2 at $\delta = 10^{-5}$ over 50 global training rounds. This privacy expenditure is 33% lower than the DP-FedAvg baseline ($\varepsilon = 1.8$) at equivalent accuracy, attributable to heterogeneity-aware client selection, which reduces the number of rounds required for convergence from 78 (DP-FedAvg) to 50, thereby directly reducing privacy budget consumption under the moments accountant. The SecAgg protocol adds a cryptographic privacy layer orthogonal to the DP guarantee: even if the DP noise were removed, the server could not observe any individual site's gradient update, thereby satisfying the privacy requirements of institutions whose data governance policies mandate gradient-level confidentiality in addition to output-level DP protection.

C. Communication Efficiency

Table 1 summarizes the communication overhead assessment across all evaluated algorithms. PPFL-IDS achieves low communication cost, comparable to FedProx and DP-FedAvg, through two

mechanisms: the gradient-sharing approach for gradient-boosted ensemble models reduces per-round transmission size by approximately 62% relative to deep neural network weight vectors of equivalent predictive capacity; and client subsampling at a rate of 0.3 reduces the number of transmitting clients per round from K = 20 to m = 6, reducing total uplink traffic by 70% relative to full-participation protocols.

SCAFFOLD incurs the highest communication cost because it requires transmitting control variates in addition to gradient updates, increasing the per-client payload by approximately 2x. Despite this overhead, SCAFFOLD's non-IID correction mechanism achieves the second-highest accuracy among baselines (91.4%). PPFL-IDS surpasses SCAFFOLD by 2.4 percentage points in F1-score at 40% lower communication cost, confirming that the heterogeneity-aware selection strategy provides a more communication-efficient non-IID correction than variate-based gradient correction.

D. Non-IID Robustness

The performance gap between IID and non-IID evaluation settings quantifies non-IID robustness. Under IID partitioning (Dirichlet concentration = 100, approximating uniform class distribution across sites), FedAvg achieves an F1-score of 0.924. Under the non-IID partition (concentration = 0.5), FedAvg degrades to 0.873, a gap of 5.1 percentage points, consistent with the degradation reported by Zhao et al. [3]. PPFL-IDS under non-IID conditions (0.938) exceeds FedAvg under IID conditions (0.924) by 1.4 percentage points, demonstrating that the heterogeneity-aware client selection not only recovers the non-IID performance penalty but produces a net accuracy gain over IID-optimized baselines by ensuring that the most statistically underrepresented attack categories are represented in each training round.

E. Practical Deployment Considerations

The gradient-boosted local model architecture enables deployment on gateway-class hardware with 2 GB RAM and a quad-core ARM processor, which is representative of commercially available edge gateway hardware for healthcare networks. Local training of one round (50 mini-batches of 256 samples) completes in 3.2 seconds on this hardware profile, confirming that round duration does not constrain the 50-round training schedule within a reasonable deployment window. The SecAgg cryptographic computation adds 0.8 seconds per round per participating site at the gateway tier, representing an 8% overhead relative to the 10-second inter-round aggregation window used in the evaluation. These resource measurements confirm that PPFL-IDS is deployable on existing HIoT gateway infrastructure without hardware upgrades.

Table 2. Attack Detection Performance per Category

Attack Category	Precision	Recall	F1-Score	Latency (ms)
DoS / DDoS	0.961	0.948	0.954	18.4
Man-in-the-Middle	0.934	0.921	0.927	22.1

Attack Category	Precision	Recall	F1-Score	Latency (ms)
Replay Attack	0.918	0.903	0.910	19.7
Anomalous Telemetry	0.947	0.939	0.943	16.9
Credential Stuffing	0.902	0.887	0.894	24.3
Overall (Weighted Avg)	0.938	0.920	0.938	20.3

5. Future Recommendations

A. Adaptive Differential Privacy Budgeting

The current framework uses a fixed epsilon budget allocated uniformly across training rounds. Future work should investigate adaptive budget allocation strategies that concentrate privacy spending in early rounds where gradient updates are most informative and reduce noise in later rounds as the model approaches convergence. Moments accountant formulations already support heterogeneous per-round sigma values; optimizing the sigma schedule to minimize total epsilon expenditure at a target accuracy threshold would yield privacy-accuracy Pareto improvements beyond those demonstrated here.

B. Byzantine-Fault-Tolerant Aggregation

The current aggregation protocol assumes honest-but-curious client behavior. Healthcare network environments may include compromised gateway devices that submit adversarially crafted gradient updates intended to degrade global model accuracy (model poisoning). Future work should integrate Byzantine-robust aggregation mechanisms such as Krum, coordinate-wise median, or Bulyan into the PPFL-IDS aggregation step, and evaluate the interaction between Byzantine robustness and the DP noise injection, which already provides some protection against the influence of outlier gradients.

C. Vertical Federated Learning for Cross-Institutional Feature Sharing

The current horizontal FL architecture assumes that all sites collect the same set of network flow features. In practice, different hospital sites may deploy different network monitoring tools, resulting in overlapping but non-identical feature spaces. Vertical federated learning [2] would allow sites to contribute complementary feature views of the same network events without sharing raw features, potentially expanding the discriminative feature set available for intrusion classification while preserving site-level data sovereignty. Integrating vertical FL with the differential privacy and secure aggregation mechanisms established in PPFL-IDS represents a natural extension of the framework.

D. Continual Learning for Evolving Attack Patterns

Network intrusion patterns evolve continuously as attackers adapt to detection mechanisms. The current framework is evaluated on a static dataset snapshot; production deployments require continual model updates that incorporate newly observed attack patterns without forgetting previously learned detection capabilities. Future research should investigate federated continual learning strategies, particularly elastic weight consolidation and progressive neural network approaches, and evaluate them within the DP privacy accounting framework to quantify the privacy cost of continual retraining over multi-year deployment horizons.

E. Integration with Low-Code Enterprise Security Dashboards

PPFL-IDS produces per-site intrusion classifications and global model confidence scores that are most actionable when surfaced to clinical IT security personnel through intuitive dashboards. Prior work on low-code machine learning deployment in enterprise platforms [21][22] has demonstrated that Oracle APEX's PL/SQL integration layer can consume REST-based ML inference endpoints and present results in interactive monitoring interfaces without bespoke front-end development. Future work should evaluate the end-to-end latency and usability of an APEX-hosted PPFL-IDS monitoring dashboard, providing clinical security operations teams with a low-barrier path to deploying federated intrusion detection in existing enterprise IT environments.

6. Conclusions

This paper presented PPFL-IDS, a Privacy-Preserving Federated Learning framework for intrusion detection in Healthcare IoT environments that simultaneously addresses data privacy, non-IID data heterogeneity, and communication efficiency constraints. The framework integrates differential privacy noise injection calibrated to an epsilon budget of 1.2 with the SecAgg cryptographic aggregation protocol, ensuring that neither the aggregation server nor any external observer can reconstruct local device telemetry from gradient updates. A heterogeneity-aware client selection mechanism addresses the non-IID convergence challenge that has limited the applicability of prior federated IDS proposals to real-world multi-site HIoT deployments.

Evaluation on UNSW-NB15 with non-IID partitioning across 20 simulated hospital sites achieved a weighted F1-score of 0.938 and a mean detection latency of 20.3 ms, outperforming all four baseline FL algorithms across accuracy, privacy budget, and communication overhead. PPFL-IDS exceeds the strongest privacy-preserving baseline (DP-FedAvg) by 8.2 percentage points in F1-score and the strongest non-privacy baseline (SCAFFOLD) by 2.4 percentage points at 40% lower communication cost, establishing a new operating point on the accuracy-privacy-communication trade-off surface for federated IoT intrusion detection.

The local gradient-boosted ensemble architecture enables deployment on commercial HIoT gateway hardware without requiring GPU acceleration, and the per-round local training and SecAgg computation overhead remain within the 10-second inter-round window used in the evaluation protocol. These resource measurements confirm that PPFL-IDS is immediately deployable in existing HIoT infrastructure.

The security challenges motivating this work are structurally representative of any regulated IoT environment where data sovereignty constraints prevent telemetry centralization. The framework's design principles, federated edge inference, calibrated differential privacy, and heterogeneity-aware participation management are portable across smart-city, industrial IoT, and government network security contexts. Future extensions to adaptive privacy budgeting, Byzantine-robust aggregation,

vertical federated learning, and continual learning will further strengthen the framework's applicability to production HIoT security deployments.

7. References

[1] H. Touqeer, S. Zaman, R. Amin, M. Hussain, F. Al-Turjman, and M. Bilal, "Smart home security: challenges, issues and solutions at different IoT layers," J. Supercomputing, vol. 77, no. 12, pp. 14053-14089, May 2021.

[2] S. Agrawal, S. Sarkar, O. Aouedi, G. Yenduri, K. Piamrat, S. Bhattacharya, P. K. R. Maddikunta, and T. R. Gadekallu, "Federated learning for intrusion detection system: Concepts, challenges and future directions," Comput. Commun., vol. 195, pp. 346-361, Nov. 2022.

[3] R. Zhao et al., "Federated learning with non-IID data," arXiv preprint arXiv:1806.00582, 2018.

[4] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS), 2017, pp. 1273-1282.

[5] Y. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," Future Gener. Comput. Syst., vol. 115, pp. 619-640, Feb. 2021.

[6] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in Proc. 3rd Theory Cryptogr. Conf. (TCC), 2006, pp. 265-284.

[7] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in Proc. ACM SIGSAC Conf. Comput. Commun. Security (CCS), 2017, pp. 1175-1191.

[8] Sharma H, Kumar P, Sharma K. Security solutions for the Internet of Things using machine learning and deep learning: Current trends and future directions. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 16(1):e70059, Mar. 2026.

[9] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems," in Proc. Mil. Commun. Inf. Syst. Conf. (MilCIS), 2015, pp. 1-6.

[10] M. Abdel-Basset, H. Hawash, N. Moustafa, and T. Razzak, "Semi-supervised spatiotemporal deep learning for intrusions detection in IoT networks," IEEE Internet Things J., vol. 8, no. 15, pp. 12251-12265, Aug. 2021.

[11] A. Syed, "Secure Integration of Oracle APEX Application with IoT Devices for the 'My Smart Home' App," Int. J. Leading Res. Publication, vol. 4, no. 11, pp. 1-12, Nov. 2023.

[12] R. F. Ali, A. Muneer, P. D. D. Dominic, S. M. Taib, and E. A. A. Ghaleb, "Internet of Things (IoT) security challenges and solutions: A systematic literature review," in Commun. Comput. Inf. Sci., Singapore: Springer, 2021, pp. 128-154.

[13] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Smola, and V. Smith, "Federated optimization in heterogeneous networks," in Proc. Mach. Learn. Syst. (MLSys), 2020, pp. 429-450.

[14] X. Kairouz et al., "Advances and open problems in federated learning," Found. Trends Mach. Learn., vol. 14, no. 1-2, pp. 1-210, 2021.

[15] V. Smith, C. Chiang, M. Sanjabi, and A. Talwalkar, "Federated multi-task learning," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), 2017, pp. 4424-4434.

[16] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," IEEE Signal Process. Mag., vol. 37, no. 3, pp. 50-60, May 2020.

[17] L. T. Phong, Y. Aono, T. Hayashi, L. Wang, and S. Moriai, "Privacy-preserving deep learning via additively homomorphic encryption," IEEE Trans. Inf. Forensics Security, vol. 13, no. 5, pp. 1333-1345, May 2018.

[18] S. Niknam, H. S. Dhillon, and J. H. Reed, "Federated learning for wireless communications: Motivation, opportunities and challenges," IEEE Commun. Mag., vol. 58, no. 6, pp. 46-51, Jun. 2020.

[19] W. Y. B. Lim et al., "Federated learning in mobile edge networks: A comprehensive survey," IEEE Commun. Surveys Tuts., vol. 22, no. 3, pp. 2031-2063, 2020.

[20] M. Chen, O. Simeone, G. Caire, A. Hero, Z. Zhang, and S. Shamai (Shitz), "A joint learning and communications framework for federated learning over wireless networks," IEEE Trans. Wireless Commun., vol. 20, no. 1, pp. 269-283, Jan. 2021.

[21] A. Syed, "Low-Code, High Impact: Unleashing Machine Learning in Oracle APEX Applications," Int. J. Multidisciplinary Res. Growth Eval., vol. 6, no. 5, pp. 240-247, Sep. 2025.

[22] Sharma H, Kumar P, Sharma K. Smart Waste Management with IoT: An Optimized Triple Memristor Hopfield Neural Network Approach. International Journal on Smart & Sustainable Intelligent Computing. 5;2(1):52-64, Feb. 2025