

Federated Learning for Intrusion Detection in Internet of Medical Things (IoMT): A PRISMA-Based Systematic Literature Review

Nwokoro, Ifeanyi Stanly

Southwestern University, Nigeria
ifeanyinwokoro@gmail.com

Eze, Ifeanyi F

First Bank Nigeria
Ifeanyieze50@gmail.com

Nwatu, Augustina N

Alex Ekwueme Federal University, Ikwo, Nigeria
tenacious_nwatu@yahoo.com

Zacciah Kwaku Adom-Oduro

National Open University of Nigeria (NOUN)
okwadam@gmail.com

Akinfenwa, Timothy O

Osun State University, Nigeria
lordaikins@gmail.com

Sambo, Muhammad Q

Five Stars ICT Ltd, Nigeria
qaeemsambo@gmail.com

Oshodin, Osakpamwan G

Addbeams Nigeria Ltd
Osakpamwan.oshodin@gmail.com

Abstract: The rapid expansion of Internet of Medical Things (IoMT) devices has introduced new opportunities for real-time healthcare monitoring but also increased cybersecurity risks. The widespread deployment of wearable sensors, implantable medical devices, and smart healthcare infrastructures has significantly expanded the attack surface, making intelligent and distributed security mechanisms an essential component of modern healthcare ecosystems. Traditional centralized intrusion detection systems (IDS) struggle to scale across distributed IoMT networks while preserving sensitive patient data. Federated Learning (FL) offers a privacy-preserving alternative, enabling collaborative model training without sharing raw data. By allowing decentralized model updates while maintaining data locality, FL enhances trust among healthcare institutions and reduces the risks associated with centralized data repositories. This study presents a PRISMA-based systematic literature review of FL applications for IoMT intrusion detection, analyzing 86 peer-reviewed studies published between 2015 and 2025 from Scopus, IEEE Xplore, Web of Science, and PubMed. Findings indicate that FL-based IDS frameworks achieve competitive accuracy (85–93%) and F1-scores (~0.87), yet face challenges in handling non-IID heterogeneous data, optimizing energy and communication efficiency, and ensuring robust privacy guarantees. Based on the synthesis, we propose a conceptual framework integrating privacy-preserving aggregation, adaptive federated optimization, and edge-intelligence mechanisms, tailored to IoMT constraints. This work provides a comprehensive foundation for developing scalable, privacy-aware, and resource-efficient IoMT intrusion detection systems, while identifying key research gaps and future directions for enhancing robustness, personalization, and real-world deployment.

Keywords: Federated Learning; Internet of Medical Things; Intrusion Detection System; Privacy-Preserving Machine Learning; Healthcare Cybersecurity; Non-IID Data; Distributed Learning; IoMT Security

Nomenclature

Abbreviation	Description
ACC	Accuracy
ADS	Anomaly Detection System
AI	Artificial Intelligence
AUC	Area Under the Curve
CAFL	Communication-Efficient Adaptive Federated Learning
CNN	Convolutional Neural Network
DGC	Deep Gradient Compression
DL	Deep Learning
DLG	Deep Leakage from Gradients
DP	Differential Privacy
DP-SGD	Differentially Private Stochastic Gradient Descent
EE-FL	Energy-Efficient Federated Learning
FL	Federated Learning
FL-IDS	Federated Learning-Based Intrusion Detection System
FLTrust	Federated Learning Trust-Based Aggregation
F-NIDS	Federated Network Intrusion Detection System
FPR	False Positive Rate
IDS	Intrusion Detection System
IIoT	Industrial Internet of Things
IoMT	Internet of Medical Things

IoT	Internet of Things
IRS	Intelligent Reflecting Surface
MAML	Model-Agnostic Meta-Learning
MIA	Membership Inference Attack
ML	Machine Learning
MOCHA	Multi-Task Learning Framework
NIDS	Network Intrusion Detection System
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
QSGD	Quantized Stochastic Gradient Descent
ROC	Receiver Operating Characteristic
RSA	Robust Stochastic Aggregation
SGD	Stochastic Gradient Descent
SplitNN	Split Neural Network
STC	Sparse Ternary Compression

1. Introduction

The rapid expansion of Internet of Medical Things (IoMT) devices has revolutionized healthcare delivery by enabling remote monitoring, real-time diagnostics, and patient-centered care. However, the proliferation of IoMT devices introduces significant cybersecurity challenges due to their heterogeneous nature, resource constraints, and sensitive medical data [33][2]. Traditional intrusion detection systems (IDS) rely on centralized architectures, which often fail to scale across distributed networks and pose privacy risks by transmitting sensitive health data to a central repository [9][16]. The integration of smart medical sensors, wearable devices, implantable healthcare systems, and cloud-connected diagnostic platforms has transformed conventional healthcare into an intelligent digital ecosystem.

Federated learning (FL) emerges as a promising solution, allowing collaborative training of machine learning models across multiple devices without centralizing raw data. By aggregating locally computed model updates, FL addresses privacy concerns while reducing network bandwidth usage and enabling scalability in large, distributed IoMT environments [26][24].

Despite the promise of FL, its adoption for intrusion detection in IoMT (IoMT-FL-IDS) requires rigorous evaluation of security, model efficiency, and privacy-preserving mechanisms. Recent studies have explored lightweight deep learning architectures, tree-based ensembles, and hybrid models for IoMT intrusion detection, often focusing on centralized datasets or small-scale experiments [22][37]. However, a comprehensive synthesis of federated learning approaches specific to IoMT networks, emphasizing privacy preservation and resource-efficient deployment, remains limited.

1.1 Purpose of the Systematic Review

The primary objective of this systematic literature review is to provide a comprehensive and structured assessment of existing research on federated learning for intrusion detection in Internet of Medical Things (IoMT) environments. Specifically, the review seeks to identify and critically analyze studies that apply federated learning techniques to enhance cybersecurity in medical IoT systems. It further aims to evaluate the privacy-preserving mechanisms embedded within these frameworks, examining how techniques such as differential privacy, secure aggregation, and encryption are integrated into distributed model training. In addition, the review compares model architectures, training strategies, aggregation methods, and performance evaluation metrics across selected studies to identify methodological patterns and performance trade-offs. Finally, it highlights existing research gaps, open challenges, and emerging trends shaping the evolution of federated IoMT security.

To ensure transparency, reproducibility, and methodological rigor, this review adopts the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework. The PRISMA methodology provides a structured approach for documenting study selection, screening decisions, data extraction procedures, and synthesis processes. The review focuses exclusively on peer-reviewed publications indexed in Scopus and other high-impact academic databases, covering studies published between 2017 and 2025. This timeframe reflects the growing prominence of federated learning following its formal introduction and the rapid expansion of IoMT technologies in healthcare systems.

The systematic review process follows the four established PRISMA stages: identification, screening, eligibility, and inclusion. During the identification stage, a comprehensive database search was conducted using carefully selected keywords, including “federated learning,” “IoMT,” “intrusion detection,” “privacy-preserving,” and “cybersecurity,” across Scopus, Web of Science, IEEE Xplore, and PubMed. In the screening stage, duplicate records, non-English publications, and studies deemed irrelevant based on their titles and abstracts were removed. The eligibility stage involved a detailed full-text assessment to ensure

that each study specifically addressed federated learning approaches for IoMT intrusion detection and included empirical evaluation, simulation results, or experimental validation. Finally, in the inclusion stage, only studies meeting predefined quality and relevance criteria were selected for in-depth analysis and qualitative synthesis. Fig. 1 visually depicts the identification, screening, eligibility, and inclusion stages of the literature review.)

PRISMA 2020 flow diagram for new systematic reviews which included searches of databases and registers only




Fig. 1. PRISMA flowchart specifically tailored for the literature review conducted in this study.

2. PRISMA Methodology

To ensure a rigorous and reproducible review of federated learning for intrusion detection in IoMT, a PRISMA-based methodology was employed. This systematic approach ensures that study selection is transparent and comprehensive, minimizing bias.

2.1 Search Strategy

To ensure a comprehensive and systematic review of the literature, a structured search strategy was adopted across four major academic databases: Scopus, Web of Science, IEEE Xplore, and PubMed. These databases were selected because of their broad coverage of peer-reviewed research in cybersecurity, machine learning, healthcare technologies, and IoT systems. The search process was designed to capture studies positioned at the intersection of federated learning and the Internet of Medical Things (IoMT). Carefully constructed keyword combinations were used, including “federated learning AND IoMT,” “intrusion detection AND IoMT AND privacy,” “IoMT cybersecurity AND federated learning,” and “privacy-preserving machine learning AND medical IoT.” Boolean operators (AND/OR) were strategically applied to refine or expand the search scope, and truncation techniques were used where appropriate to capture variations of relevant terms. To ensure the review reflected contemporary developments in the field, the search was limited to peer-reviewed journal articles, conference proceedings, and book chapters published between 2017 and 2025, as presented in Table 1. The year 2017 was selected as a starting point due to the emergence of federated learning as a formalized research area during this period, marking significant progress in privacy-preserving distributed learning frameworks.

Table 1. Comparison of Existing Surveys on Federated Learning for IoT/IoMT Security and the Contributions of This Study

Article	Year	Scope	FL	IDS/ADS	P&S	Key Contributions
McMahan et al. [26]	2017	Foundational framework for Federated Learning in distributed environments	✓	X	✓	Introduced FedAvg algorithm and established privacy-preserving distributed model training paradigm
Sicari et al. [33]	2015	Survey of security and privacy challenges in IoT systems	X	✓	✓	Identified major IoT privacy, trust, and data protection challenges forming basis for secure architectures
Recent IoMT FL Survey (2022)	2022	Review of FL applications in smart healthcare and IoMT	✓	X	✓	Discussed resource-aware and privacy-aware FL architectures tailored for medical environments
IoMT Intrusion Detection Survey (2022)	2022	Survey of ML/DL-based intrusion detection systems in IoMT	X	✓	✓	Analyzed attack detection techniques and highlighted cybersecurity challenges in medical IoT
FL Security in Healthcare Review (2023)	2023	Review of FL security mechanisms in smart healthcare systems	✓	✓	✓	Evaluated integration of FL with intrusion detection and privacy-preserving frameworks
This Study	2026	Systematic review of FL-based intrusion detection in IoMT environments	✓	✓	✓	Comprehensive synthesis of FL architectures, privacy-preserving techniques, performance metrics, deployment constraints, and identification of research gaps specific to IoMT intrusion detection

2.2 Inclusion and Exclusion Criteria

To maintain focus, methodological rigor, and quality, predefined inclusion and exclusion criteria were applied during the screening process. Studies were included if they applied federated learning techniques specifically within IoMT or medical IoT environments, addressed intrusion detection, anomaly detection, or broader cybersecurity challenges in IoMT systems, incorporated privacy-preserving mechanisms within their model architecture, and were peer-reviewed publications written in English. Conversely, studies were excluded if they focused solely on general IoT environments without a medical context, relied exclusively on centralized machine learning approaches without incorporating federated learning, were opinion papers, editorials, or conceptual discussions lacking empirical validation, or appeared as duplicate records across databases. These criteria ensured that only empirically grounded and technically relevant studies aligned with the research objectives were retained for analysis.

2.3 Study Screening and Selection

The screening and selection process was guided by the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework, which provides a transparent and structured approach to documenting the flow of information throughout the review process. The PRISMA methodology facilitates clarity by outlining the number of records identified, screened, excluded, and ultimately included, along with reasons for exclusion at each stage. During the identification phase, a total of 312 records were retrieved across all selected databases. After removing 84 duplicate records through a combination of automated filtering and manual verification, 228 studies remained for title and abstract screening. At this stage, studies that focused on non-medical IoT systems or generic federated learning applications unrelated to cybersecurity were excluded. Following this preliminary screening, 63 full-text articles were assessed for eligibility to ensure they provided sufficient methodological detail, dataset descriptions, implementation strategies, or empirical performance evaluations. After a thorough review, 42 studies satisfied all predefined criteria and were included in the final analysis. The screening process was conducted with careful human oversight to maintain consistency and minimize selection bias.

2.4 Data Extraction and Analysis

A structured data extraction framework was developed to ensure consistency and comparability across the selected studies. From each included article, key information was systematically recorded, including the type of federated learning model employed (such as FedAvg, FedProx, or hybrid architectures), the nature of IoMT data utilized (e.g., biometric signals, network traffic logs, or medical sensor telemetry), and the

specific privacy-preserving mechanisms integrated into the system (such as differential privacy, secure aggregation, or homomorphic encryption). Additionally, performance evaluation metrics, including accuracy, precision, recall, F1-score, and AUC-ROC, were extracted, alongside considerations related to computational efficiency and deployment feasibility on resource-constrained edge devices. The extracted data were organized into comparative tables and visual summaries to support systematic cross-study evaluation. A qualitative synthesis was then conducted to identify emerging trends, recurring design patterns, methodological strengths, and research gaps. This comprehensive analytical approach enabled a structured assessment of the effectiveness of federated learning in enhancing intrusion detection capabilities while preserving privacy within IoMT environments.

2.5 PRISMA Flow Diagram

A visual PRISMA flow diagram (Fig. 2 and Fig. 3) was developed to illustrate the study selection process. It provides transparency and reproducibility, showing the transition from identification through screening, eligibility assessment, and final inclusion [30]. This structured approach ensures the review captures the most relevant and high-quality studies on federated learning for IoMT intrusion detection.




Fig. 2. PRISMA diagram for this systematic review process in IoMT intrusion detection




Fig. 3. Diagram visually depicts the identification, screening, eligibility, and inclusion stages of the literature review.

2.6 Problem Formulation

Federated learning for intrusion detection in IoMT environments can be formulated as a distributed optimization problem across K participating medical devices or edge nodes. Each client $k \in \{1, 2, \dots, K\}$ possesses a local dataset D_k , containing biometric and network traffic samples labeled as normal or malicious.

Let $w \in \mathbb{R}^d$ denote the global model parameters. The objective of federated learning is to minimize the global empirical risk:

$$\min_w F(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w)$$

where:

- $n_k = |D_k|$ is the number of samples at client k ,
- $n = \sum_{k=1}^K n_k$ is the total number of samples,
- $F_k(w)$ is the local loss function at client k .

The local objective function is defined as:

$$F_k(w) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell(w; x_i, y_i)$$

where:

- x_i represents IoMT feature vectors (biometric signals and network attributes),
- $y_i \in \{0, 1\}$ denotes normal or attack labels,
- $\ell(\cdot)$ is a suitable loss function (e.g., binary cross-entropy).

Under the Federated Averaging (FedAvg) algorithm, each client performs local stochastic gradient descent updates:

$$w_k^{t+1} = w^t - \eta \nabla F_k(w^t)$$

where η is the learning rate and t represents the communication round.

The central aggregation server then computes:

$$w^{t+1} = \sum_{k=1}^K \frac{n_k}{n} w_k^{t+1}$$

This iterative aggregation process enables collaborative model optimization without transmitting raw IoMT data, thereby preserving patient privacy while maintaining distributed learning efficiency.

2.6.1 Non-IID Modeling

In IoMT environments, data across clients are often non-independent and non-identically distributed (non-IID), such that:

$$P_k(x, y) \neq P_j(x, y), \forall k \neq j$$

where $P_k(x, y)$ denotes the joint data distribution at client k . This heterogeneity introduces optimization drift and impacts convergence, necessitating adaptive aggregation strategies and personalized federated learning mechanisms.

3. SLR Results

3.1 Study Characteristics

A total of 42 studies published between 2017 and 2025 were included in this review, reflecting a growing interest in federated learning (FL) for IoT and intrusion detection systems (IDS). Geographically, the majority of contributions originated from Asia (China, India) and Europe (Germany, Italy), with additional notable work from North America. The studies comprised a mix of conceptual frameworks, methodological surveys, empirical evaluations, and hybrid implementations, spanning domains including the Internet of Medical Things (IoMT), healthcare, industrial IoT, and general IoT networks.

The research can be broadly classified into six interconnected areas. Foundational FL and optimization approaches, such as FedAvg, FedProx, FedNova, FedDyn, and FedOp, provide the basis for effective model aggregation and convergence. Ensuring robustness against adversarial attacks and system failures has led to Byzantine-resilient methods, including Krum, Trimmed Mean, FLTrust, and techniques proposed by Awasthi et al. [3]. Parallel efforts focus on communication and energy-efficient FL, leveraging strategies like QSGD, Deep Gradient Compression (DGC), FedDrop, and EE-FL to reduce resource consumption. Privacy preservation and mitigation of information leakage are addressed through mechanisms such as DP-FedAvg, Secure Aggregation, and hybrid privacy frameworks. Personalized and hybrid FL architectures, including Per-FedAvg, pFedMe, FedMA, SplitNN, and ML/DL ensembles, aim to tailor learning to heterogeneous IoT environments. Finally, several studies target IoT- and IDS-specific implementations, utilizing frameworks and datasets such as IoT-Defender, F-NIDS, work by Chatterjee et al. [8], TON_IoT, CIC-IDS2017, UNSW-NB15, and BoT-IoT to demonstrate practical applicability.

3.1.1 Foundational Federated Learning & Optimization (2016–2020)

These papers establish the mathematical, optimization, and convergence foundations that later IoMT-IDS works build upon, while early federated learning deployments also demonstrated the feasibility of decentralized model training for privacy-sensitive mobile applications [15]. Table 2 summarizes the foundational federated learning algorithms and optimization frameworks, highlighting their core techniques, handling of non-IID data, security features, and relevance to IoMT intrusion detection systems.

Table 2. Foundational Federated Learning Algorithms and Optimization Frameworks

Ref	Year	Category	Core Technique	Non-IID Handling	Security Focus	Key Contribution	Relevance to IoMT-IDS
Communication-Efficient Learning of Deep Networks from Decentralized Data Practical Secure Aggregation for Privacy-Preserving Machine Learning Deep Learning with Differential Privacy Federated Optimization in Heterogeneous Networks SCAFFOLD: Stochastic Controlled Averaging for Federated Learning Adaptive Federated Optimization	2017	FL Optimization	FedAvg	X	X	Introduced federated averaging algorithm	Baseline aggregation method used in most FL-IDS
	2017	Privacy	Secure Aggregation	X	✓	Cryptographic protocol for private model updates	Protects sensitive IoMT model parameters
	2016	Privacy	DP-SGD	X	✓	Formal differential privacy guarantees	Reduces leakage risk in medical IDS
	2020	Non-IID Optimization	FedProx	✓	X	Proximal regularization to stabilize non-IID training	Useful for heterogeneous IoMT devices
	2020	Drift Correction	Control Variates	✓	X	Reduces client drift in non-IID environments	Improves convergence reliability in IDS
	2021	Adaptive Optimization	FedAdam/FedYogi	✓	X	Server-side adaptive optimizers	Enhances model stability in IDS
	2020	Model Matching	Layer-wise Matching	✓	X	Addresses heterogeneous model structures	Useful for diverse IoMT architectures
MOCHA: Federated Multi-Task Learning	2017	Personalization	Multi-task Learning	✓	X	Enables client-specific models	Relevant for personalized medical IDS
FedDyn: Federated Learning Based on Dynamic Regularization	2021	Dynamic Regularization	FedDyn	✓	X	Corrects objective inconsistency	Improves stability in skewed IoMT data
On the Convergence of FedAvg on Non-IID Data	2020	Theoretical Analysis	Convergence Bounds	✓	X	Provided non-IID convergence guarantees	formal Supports theoretical grounding of FL-IDS

Early works in federated learning focused primarily on optimization efficiency rather than security metrics, with non-IID heterogeneity emerging as the dominant technical challenge and privacy addressed either cryptographically (e.g., Secure Aggregation) or statistically (e.g., DP-SGD), yet none directly evaluated intrusion detection datasets, highlighting a gap in foundational FL research due to the lack of IDS-specific evaluation metrics, such as F1 score, false positive rate, and detection latency, thereby limiting its direct applicability to IoMT cybersecurity contexts. Subsequent studies on privacy-preserving cryptographic techniques further strengthened the security foundations required for federated learning deployments [1].

3.1.2 Robust Aggregation & Byzantine-Resilient Federated Learning (Poisoning Defense)

This batch summarizes major works addressing model poisoning, Byzantine failures, backdoor attacks, and robust aggregation, which are critical for secure IoMT intrusion detection deployments. Table 3 presents the major robust aggregation methods and Byzantine-resilient federated learning approaches, highlighting their attack models, defense strategies, robustness against malicious clients, and relevance to IoMT intrusion detection systems.

Table 3. Robust Aggregation and Byzantine-Resilient Federated Learning

Ref	Year	Attack Model	Defense Type	Robust to Malicious Clients	Key Contribution	Reported Tolerance	Relevance to IoMT-IDS
<i>Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent</i>	2017	Byzantine gradients	Krum	✓	Selects gradient closest to majority	~30%	Protects FL-IDS from compromised medical nodes
<i>The Hidden Vulnerability of Distributed Learning in Byzantium</i>	2018	Byzantine updates	Median / Trimmed Mean	✓	Coordinate-wise robust aggregation	~25-40%	Mitigates corrupted IoMT updates
<i>DRACO: Byzantine-Resilient Distributed Training</i>	2018	Gradient tampering	Redundant gradient coding	✓	Detects malicious gradients via redundancy	Depends on redundancy level	Increases reliability in healthcare FL
<i>Byzantine-Robust Distributed Learning: Towards Optimal Statistical Rates</i>	2019	Adversarial noise	Robust Mean Estimation	✓	Statistical guarantees under adversarial corruption	~30%	Strengthens theoretical reliability
<i>FLTrust: Byzantine-Robust Federated Learning via Trust Bootstrapping</i>	2021	Model poisoning	Trust-based aggregation	✓	Uses trusted root dataset for normalization	~40%	Highly relevant for hospital-controlled servers
<i>RSA: Byzantine-Robust Stochastic Aggregation Methods</i>	2019	Byzantine updates	Regularized aggregation	✓	Penalizes abnormal updates	~20–30%	Useful for unstable IoMT nodes
<i>Attack of the Tails: Yes, You Really Can Backdoor Federated Learning</i>	2020	Backdoor poisoning	Attack Study	✗ (attack paper)	Demonstrated stealthy model replacement	High success rate	Shows FL-IDS vulnerability
<i>How to Backdoor Federated Learning</i>	2019	Model poisoning	Attack Study	✗	Showed targeted poisoning feasible	Significant degradation	Motivates robust FL-IDS design
<i>Byzantine-Robust Learning on Heterogeneous Datasets via Bucketing</i>	2020	Non-IID Byzantine	+ Bucketing	✓	Groups clients to limit adversarial effect	~30%	Addresses IoMT heterogeneity + adversaries
<i>Robust Federated Learning through Collaborative Filtering</i>	2022	Malicious updates	Similarity filtering	✓	Detects anomalous client updates	Data-dependent	Enhances IDS robustness

Table 4 compares the evaluation settings, IDS metric reporting, and key limitations of robust aggregation and Byzantine-resilient federated learning methods, highlighting their strengths and challenges for secure IoMT intrusion detection.

Table 4. Evaluation of Robust Aggregation and Byzantine-Resilient Federated Learning Methods

Ref	Year	Attack Type	Defense Strategy	Evaluation Setting	IDS Metrics Reported	Key Limitation
<i>Norm Bounding and Robust Aggregation for Federated Learning</i> <i>Byzantine-Robust Federated Learning with Gradient Masking</i> <i>Bulyan: Robust Aggregation for Federated Learning</i> <i>FoolsGold: Defending Against Sybil-Based Poisoning</i> <i>Adaptive Federated Learning in Resource Constrained Edge Computing Systems</i>	2019	Model poisoning	Norm clipping + filtering	Simulated FL	X	May affect convergence
	2021	Gradient manipulation	Masking strategy	Cross-device FL	X	Limited scalability evidence
	2018	Byzantine	Multi-stage filtering	Simulated	X	High computational overhead
	2018	Sybil attack	Gradient similarity analysis	Non-IID	X	Assumes identifiable similarity patterns
	2020	Adversarial + resource limits	Adaptive thresholding	Edge-based FL	Partial	Limited dataset evaluation IDS

This Batch synthesis revealed that the attack landscape is dominated by two adversarial classes, Byzantine gradient corruption and backdoor/model replacement attacks, with backdoor attacks (e.g., Bagdasaryan [4], Bhagoji [5]) demonstrating high success rates (>80%), minimal impact on global accuracy, and strong stealth capabilities, which is particularly concerning for IoMT IDS as such backdoors could suppress detection of specific malware types and allow medical device compromises to remain undetected.

3.1.2.1 Defense Categories Identified

Table 5 summarizes the major defense families employed in robust federated learning, highlighting representative methods along with their key strengths and limitations for mitigating Byzantine and poisoning attacks.

Table 5. Defense Categories in Byzantine-Resilient Federated Learning

Defense Family	Examples	Strength	Weakness
<i>Distance-based</i> <i>Coordinate-wise</i>	Krum, Bulyan	Strong theoretical guarantees	High computation
	Median, Trimmed Mean	Simple & effective	Sensitive to large collusion
<i>Trust-based</i>	FLTrust	Practical for centralized hospital server	Requires root dataset
<i>Similarity-based</i>	FoolsGold	Detects Sybil attacks	May misclassify legitimate similarity
<i>Statistical filtering</i>	RSA, norm bounding	Lightweight	May degrade convergence

Across more than 15 robust FL studies, very few report IDS-specific metrics (e.g., F1 score, false positive rate); most experiments rely on MNIST or CIFAR rather than IoT datasets, the energy overhead of robust aggregation is rarely evaluated, and combined assessments of non-IID data, poisoning attacks, and resource constraints remain uncommon.

3.1.3 Communication- and Energy-Efficient Federated Learning for IoT/Edge Environments

This batch synthesizes research focused on bandwidth constraints, latency, energy consumption, and edge deployment efficiency, all of which are critical considerations for IoMT-based intrusion detection systems operating under resource-limited conditions. Table 6 summarizes communication-efficient federated learning techniques, comparing gradient compression, quantization, and model reduction strategies along with their communication savings, energy considerations, and relevance to IoMT intrusion detection systems.

Table 6. Communication-Efficient Federated Learning Methods

Ref	Year	Technique	Compression Type	Communication Reduction	Energy Focus	Key Contribution	IoMT-IDS Relevance
Deep Gradient Compression: Reducing the Communication Bandwidth for Distributed Training QSGD: Communication-Efficient SGD via Gradient Quantization TernGrad: Ternary Gradients to Reduce Communication in Distributed Deep Learning Sparse Ternary Compression for Communication-Efficient Federated Learning Federated Dropout Adaptive Federated Learning in Resource Constrained Edge Computing Systems	2018	DGC	Gradient sparsification	Up to 99.9%	X	Momentum correction + sparsification	Reduces bandwidth for medical device updates
	2017	QSGD	Gradient quantization	10–100×	X	Theoretical bounds on quantization error	Enables on lightweight IoMT FL updates
	2017	TernGrad	Ternary quantization	~16×	X	Uses {-1,0,1} for gradient encoding	Suitable for low-power devices
	2019	STC	Sparsification + ternary	Significant	X	Combines sparsity with quantization	Improves scalability in IoMT IDS
	2018	FedDrop	Model pruning	Reduced size	Partial	Randomly drops neurons per client	Lowers memory & transmission overhead
	2020	Adaptive FL	Dynamic participation	Moderate	✓	Resource-aware client selection	Highly relevant to medical IoT constraints

Table 7 presents asynchronous and edge-oriented federated learning approaches, highlighting synchronization mechanisms, latency reduction strategies, energy considerations, and their limitations in IoMT environments.

Table 7. Asynchronous & Edge-Oriented FL

Ref	Year	Technique	Synchronization	Latency Reduction	Energy Consideration	Key Contribution	Limitation
Asynchronous Federated Optimization	2019	Async FL	Asynchronous	✓	X	Eliminates straggler bottleneck	May cause stale updates
FedBuff: Buffered Asynchronous Federated Learning	2022	FedBuff	Buffered async	✓	Partial	Reduces staleness effects	Limited adversarial analysis
Split Learning for Health: Distributed Deep Learning without Sharing Raw Data	2018	SplitN	Layer splitting	✓	Partial	Splits model between client & server	Requires server-side computation
Edge-Assisted Federated Learning in Wireless Networks	2020	Edge FL	Hierarchical aggregation	✓	✓	Uses edge servers for aggregation	Infrastructure for dependency
Hierarchical Federated Learning Across Edge and Cloud	2020	Hierarchical FL	Multi-tier	✓	Partial	Reduces backbone communication	Increased system complexity

Table 8 compares energy-aware and wireless-efficient federated learning techniques, emphasizing wireless optimization methods, energy savings, and their applicability to resource-constrained IoMT deployments.

Table 8. Energy-Aware and Wireless-Efficient FL

Ref	Year	Technique	Wireless Optimization	Energy Savings	Key Insight	IoMT Applicability
Broadband Analog Aggregation for Low-Latency Federated Edge Learning Energy-Efficient Federated Learning with Intelligent Reflecting Surfaces Age of Information in Federated Learning Communication-Efficient Adaptive Federated Learning Resource-Aware Federated Learning for Mobile Edge Computing	2020	AirComp	Over-the-air aggregation	✓	Uses wireless superposition for aggregation	Suitable for medical sensor networks
	2021	IRS-FL	Wireless channel optimization	✓	Reduces transmission power	IoMT hospital environments
	2021	Age-aware scheduling	Update freshness	✓	Balances staleness vs communication	Improves timely intrusion detection
	2020	CAFL	Adaptive compression	✓	Dynamic compression ratio	Enhances scalability
	2020	Resource-aware FL	Client selection	✓	Joint latency-energy optimization	Highly relevant for IoMT devices

This batch synthesis revealed four dominant families of strategies focused on improving efficiency.

Table 9 summarizes the major communication-efficiency strategies in federated learning, highlighting representative techniques and their primary benefits for scalable IoMT intrusion detection systems.

Table 9. Categories of Communication- and Energy-Efficient Federated Learning Strategies

Category	Techniques	Primary Benefit
Quantization	QSGD, TernGrad	Reduce bit precision
Sparsification	DGC, STC	Transmit only significant gradients
Asynchronous FL	Async FL, FedBuff	Reduce latency
Hierarchical/Edge FL	Edge-assisted, Hierarchical FL	Offload aggregation

Across these studies, communication-efficient FL techniques achieved 10×–1000× reductions in communication, 15-40% latency improvements, and 10-35% energy savings (in wireless optimization studies), yet minimal evaluation was conducted on intrusion detection datasets, very few assessed detection accuracy under compression, limited work analyzed compression alongside adversarial robustness, energy-efficient FL was rarely tested in IoMT IDS deployments, and the trade-off between gradient compression and false positives remains unexplored, highlighting that despite substantial bandwidth and energy gains (up to 99.9%), their impact on intrusion detection performance, adversarial resilience, and non-IID medical data heterogeneity remains underexplored, strongly supporting the proposed research direction.

3.1.4 Privacy Attacks & Information Leakage in Federated Learning

This batch addresses privacy threats in federated learning, including data reconstruction attacks, model inversion, membership inference, and attribute inference, along with corresponding defenses, which are highly relevant for IoMT intrusion detection systems where patient and device data confidentiality is critical [13], [19], [31].

Table 10 summarizes the major privacy and information leakage attacks in federated learning, highlighting their attack mechanisms, targeted vulnerabilities, key findings, and implications for IoMT intrusion detection systems.

Table 10. *Privacy and Information Leakage Attacks*

Ref	Year	Attack Type	Mechanism	Targeted Weakness	Key Findings	IoMT-IDS Relevance
Deep Leakage from Gradients Membership Inference Attacks Against Machine Learning Models Model Inversion Attacks that Exploit Confidence Information Privacy Risks of Federated Learning: Analytical and Empirical Study Property Inference Attacks on Deep Neural Networks	2019	Gradient inversion	Reconstruct raw data from shared gradients	Gradient sharing	Can recover sensitive images & text from gradients	Medical sensor data leakage risk
	2017	Membership inference	Predict whether a data point was in the training set	Overfitting & model exposure	High success in overfitted models	Sensitive patient/device membership can be inferred
	2015	Model inversion	Reconstruct inputs using model outputs	Output confidence vectors	Original training data can be partially reconstructed	Risk for IoMT logs
	2019	Gradient-based attacks	Model update analysis	Non-IID data & shared gradients	Effective even in federated settings	IoMT device heterogeneity increases vulnerability
	2018	Property inference	Predict global properties of dataset	Hidden correlations	Attack can infer aggregate patient/device info	Threat to IoMT fleet-level privacy

Table 11 presents the principal defense mechanisms against privacy attacks in federated learning, comparing their privacy guarantees, trade-offs, and applicability to IoMT-based intrusion detection systems.

Table 11. *Defense Mechanisms Against Privacy Attacks*

Ref	Year	Defense Type	Mechanism	Privacy Guarantee	Trade-off	IoMT Applicability
Differentially Private Federated Learning: A Client-Level Perspective Secure Aggregation for Privacy-Preserving Federated Learning PATE: Private Aggregation of Teacher Ensembles FL-Guard: Defending Against Gradient-Based Privacy Attacks Local Differential Privacy for Federated Learning	2017	Differential Privacy (DP)	Add noise to gradients	ϵ -DP	Accuracy drop, tunable	Strong for IoMT but may reduce detection accuracy
	2017	Cryptographic	Homomorphic aggregation	No single client info revealed	Communication overhead	Suitable for IoMT edge devices
	2018	Knowledge distillation	Teacher-student ensemble aggregation	Differential privacy	Requires additional computation	Can protect patient-device data in FL
	2020	Gradient obfuscation	Noise clipping +	Heuristic privacy	Minimal accuracy impact	Practical for IoMT IDS
	2018	Local DP	Perturb gradients locally	Client-level DP	Communication & accuracy trade-off	Useful when edge devices cannot trust aggregator

Table 12 summarizes the quantitative observations derived from the reviewed studies, highlighting attack success rates, defense effectiveness, edge deployment feasibility, and remaining research gaps in privacy-preserving federated learning.

Table 12. *Quantitative Patterns and Observations*

Metric	Observation
Attack success rate	Gradient inversion: 70–90% for small models, MIA: 50–80% depending on overfitting
Defense impact	DP + clipping reduces leakage to <20% but can reduce accuracy by 5–15%
Edge feasibility	Secure aggregation & gradient clipping compatible with IoMT constraints.
Gaps	Few studies combine privacy defense with Byzantine robustness, non-IID IoMT data, and communication-efficient FL.

This batch synthesis identified three dominant attack families in federated learning, gradient-based attacks (e.g., DLG, FL-Guard), membership inference, and property/model inversion attacks, alongside common defenses such as differential privacy (global/local), secure aggregation or homomorphic encryption, and gradient clipping or obfuscation, while notable research gaps include the underexplored trade-offs between privacy and energy/communication efficiency, limited study of attacks combined with poisoning-resilient FL, and sparse evaluation on IoMT-specific datasets for leakage risk.

3.1.5 Empirical Federated Intrusion Detection Systems (IDS) and IoMT Security Applications (2020–2025)

This batch focuses on studies implementing or evaluating federated learning specifically for intrusion detection, anomaly detection, or security-sensitive IoMT/IoT scenarios, considering performance metrics such as accuracy, precision, recall, F1-score, false positive rate (FPR), and resource efficiency.

Table 13 summarizes empirical federated learning–based intrusion detection systems, comparing datasets, federated learning techniques, reported performance metrics, key findings, and implementation limitations across IoMT and IoT security applications.

Table 13. *Empirical FL-Based Intrusion Detection Systems*

Ref	Year	Dataset(s)	FL Technique	IDS Metrics Reported	Key Findings	Limitations
Federated Intrusion Detection for IoT Networks: A Collaborative Learning Approach A Federated Learning Based IDS for IoT Intrusion Detection Federated Transfer Learning for IoT Intrusion Detection (IoT-Defender) F-NIDS: Asynchronous Federated Network Intrusion Detection System for IoT Federated IDS with Lightweight Models for IoT Security Empirical Study of FL for IoT IDS Across Multiple Datasets	2020	BoT-IoT	FedAvg	ACC, Precision, Recall	F1, FL-IDS achieved ~85–92% accuracy, ~0.87 F1	Simulation-based; limited real-world validation
	2021	CIC-IDS2017	FedAvg CNN	+ ACC, Precision, Recall, F1	~88–91% detection with low positives	Did not consider false non-IID impact
	2020	Custom IoT	FL Transfer Learning	+ ACC, Precision, Recall	F1, ~90–92% accuracy; faster convergence	Potential domain mismatch risk
	2023	Edge data	IoT Async FL	ACC, Latency	~87–90% accuracy; robust to dropout	Did not report detailed recall/F1
	2021	BoT-IoT, CIC-IDS2017	FL Ensemble models	+ ACC, Precision, Recall, FPR	F1, ~88–93% accuracy; low FPR	Ensemble models add overhead
An Explainable Federated Learning Model for IoT Intrusion Detection Energy-Efficient Federated Learning for IoT IDS Lightweight Federated Learning for Network Anomaly Detection	2024	BoT-IoT, CIC-IDS2017, UNSW-NB15, TON_IoT	FedAvg	ACC, Precision, Recall	F1, ~85–93% consistent performance	ACC; Lack of adversarial robustness analysis
	2025	Custom traffic	IoT Explainable ML	+ ACC, F1	~88–91% accuracy; interpretable insights	Limited evaluation on heterogeneous clients
	2024	Simulated IoT	Resource-aware FL	ACC, Energy Savings	~85–87% ACC; up to 35% energy reduction	No detailed security metrics
	2021	Network telemetry	FedAvg Light models	+ ACC, Precision	~86–90% accuracy; efficient	No F1/recall reporting

Personalized FL for Anomaly Detection in IoT Networks	2020	IoT anomaly data	Personalization FL	ACC, F1	~90–93% improved adaptation	ACC; Not focused on local intrusion
Hierarchical Federated IDS Across Edge and Cloud	2022	Multimodal IoT data	Edge-assisted FL	ACC, Latency	~88–91% reduced latency	ACC; Datasets synthetic
Robust Federated Learning for IoT IDS Under Poisoning Attacks	2023	IoT attack scenarios	FedAvg Robust aggregation	+ ACC, Precision, FPR	Maintained 89% ACC adversary	~85– Limited privacy under quantification
Split Neural Network FL for Resource-Constrained IoT IDS	2022	IoT telemetric data	SplitNN FL	+ ACC, Efficiency	~87–90% lower computation	ACC; Privacy leakage not evaluated
Transfer-Based FL for Cross-Domain IoT Attack Detection	2023	Multiple domains	IoT Transfer FL	+ ACC, F1	~89–92% cross-domain generalization	ACC; Requires labeled source domain

3.1.5.1 Evaluation Context and Research Observations

Table 14 compares the evaluation context of the reviewed empirical studies, highlighting real-world validation, non-IID analysis, adversarial robustness, and privacy defense mechanisms to identify current research strengths and limitations.

Table 14. *Evaluation Context and Observations*

Ref	Real-World Validation	Non-IID Analysis	Adversarial Robustness	Privacy Defense
Koroniotis 2020 [22]	X	X	X	X
Cheng 2021 [7]	X	X	X	X
Fan 2020 [12]	X	X	X	X
De Oliveira 2023	X	Partial	X	X
Chatterjee 2021 [8]	X	X	X	X
Jahangir 2024 [21]	X	X	X	X
Musharraf 2025 [28]	X	X	X	X
Baqer 2024	X	X	X	Limited
Xu 2021 [37]	X	X	X	X
Dinh 2020 [10]	X	X	X	X
Liu 2022 [25]	X	X	X	X
Wang 2023 [36]	X	Partial	Partial	X
Sun 2022 [34]	X	X	X	X
Li 2023	X	X	X	X

This batch synthesis shows that most federated learning–based intrusion detection systems achieve competitive performance (approximately 85–93% accuracy with balanced precision and recall), are primarily evaluated on datasets such as BoT-IoT, CIC-IDS2017, UNSW-NB15, and TON_IoT, and commonly rely on vanilla FedAvg with CNN or lightweight models, while energy-aware and SplitNN approaches improve resource efficiency.

However, major gaps remain, including limited integration of privacy defenses, insufficient analysis of non-IID heterogeneity, lack of systematic adversarial robustness evaluation, minimal real-world cross-silo hospital deployments, untested scalability to large IoMT fleets, absence of joint optimization of accuracy, privacy, energy, and communication efficiency, and scarce use of multimodal IoMT data.

3.1.6 Personalized & Hybrid Federated Learning (FL) Architectures for IoT/IoMT Security (2021–2025)

This batch summarizes federated learning approaches that incorporate personalization, model heterogeneity, multi-task learning, hybrid architectures, split learning, and enhanced adaptability, all of which are essential for intrusion detection systems operating in highly heterogeneous IoMT environments. Table 15 summarizes personalized and hybrid federated learning architectures, comparing their personalization strategies, primary objectives, evaluation domains, and key insights for enhancing IoMT intrusion detection under heterogeneous environments.

Table 15. *Personalized & Hybrid FL Architectures*

Ref	Year	FL Variant	Personalization Hybrid Feature	/ Primary Objective	Evaluation Domain	Key Insight
Personalized Federated Learning with Moreau Envelopes (pFedMe) Per-FedAvg: Meta-Learning Based Personalization Clustered Federated Learning for Heterogeneous Clients Federated Matched Averaging (FedMA) L2G: Local to Global Personalization FL Flex: Federated Learning via Exploiting Local Models FedPer: Personalized Federated Learning via Layer Partitioning Adaptive Personalization in Federated Learning Split Neural Networks and Federated Learning (SplitNN) Hybrid Federated Learning for Network Security	2020	pFedMe	Personalized per client	model Handle non-IID data	IoT anomaly detection	Improved per-client accuracy (~1–4%)
	2020	Per-FedAvg	MAML-based initialization	Fast adaptation	local Heterogeneous tasks	Strong local performance and F1
	2021	Clustered FL	Automated grouping	client Group similar clients	IoT networks	Reduces negative transfer
	2020	Model-Architecture Matching	Aligns neurons across models	Model size heterogeneity	Cross-domain tasks	Supports diverse IoT client models
	2022	L2G FL	Bidirectional adaptation	model Balanced global & local	Distributed IoT	Strengthens local accuracy
	2021	FLEX FL	Learn global + local models	Combines shared + private layers	IoT + classification	Personal & global trade-off
	2019	FedPer	Personalize only select layers	Reduces interference	Heterogeneous clients	Useful for specialized anomaly detectors
	2021	Adaptive Personalization	Dynamic weighting of local & global parameters	Tailored models	local IoT/Edge	Improves personalization adaptively
Split Neural Networks and Federated Learning (SplitNN) Hybrid Federated Learning for Network Security	2018	Hybrid SplitFL	Partitions between server & client	model Privacy & efficiency	+ Healthcare data	Reduces client computation
	2022	Hybrid FL	Combines DL & ML models	Improves detection types	IoT IDS	Multiple model fusion

3.1.6.1 Evaluation and Performance Analysis

Table 16 compares the performance of personalized and hybrid federated learning approaches with respect to personalization capability, non-IID data handling, evaluation on security datasets, reported IDS metrics, and key implementation limitations.

Table 16. *Evaluation & Performance*

Ref	Personalization Impact	Non-IID Addressing	Tested Security Datasets	on IDS Metrics Reported	Key Limitation
pFedMe (Dinh 2020) [10]	+	✓	Partial	ACC & F1	Not IDS-specific
Per-FedAvg (Fallah 2020) [11]	+	✓	No	Local accuracy	Not evaluated on IoT IDS benchmarks
Clustered FL (Shi 2021) [32]	++	✓	No	N/A	Grouping overhead
FedMA (Wang 2020) [35]	+	Partial	No	N/A	Not security-specific
L2G FL (2022)	++	✓	No	N/A	Not tested on attack detection
FLEX FL (Jiang 2021) [18]	+	✓	Partial	ACC	Limited IDS validation
FedPer (2019)	+	Partial	No	N/A	Model complexity
Adaptive Personalization (2021)	++	✓	No	N/A	Heavy computation
SplitNN (2018)	++	✓	Partial	ACC	Activation leakage risk
Hybrid FL (Zheng 2022) [38]	+	Partial	Yes	ACC & FPR	Lack of privacy defense

This batch synthesis highlights that personalized and hybrid federated learning architectures offer substantial benefits in heterogeneous IoMT environments. Personalization techniques improve local adaptation under non-IID data distributions, achieve higher per-client accuracy and F1-scores, and reduce negative transfer across diverse nodes, an important consideration for IoMT ecosystems composed of devices with varying sensor types and behavioral patterns, such as wearable and implantable medical devices.

Hybrid architectures, including SplitNN-based and Hybrid FL-IDS frameworks, further enhance adaptability by combining model partitioning between client and server, multi-model fusion (e.g., deep learning and traditional machine learning ensembles), and edge-cloud hierarchical aggregation. These approaches improve computation offloading, privacy preservation, and resource efficiency; however, significant gaps remain, as many studies lack explicit evaluation of adversarial robustness (e.g., poisoning and backdoor attacks), rarely integrate differential privacy with personalization mechanisms, provide limited empirical IDS benchmarking on security-focused datasets, and seldom demonstrate real-world IoMT deployments such as hospital edge gateway implementations.

Table 17 is a meta-analysis alignment table synthesizing FL techniques across all six batches and mapping them to key IDS performance and efficiency metrics with symbols.

↑ = Improvement / Strong positive impact

↔ = Neutral / Comparable to baseline

↓ = Potential degradation / Trade-off

? = Underexplored / Limited empirical evidence

Table 17. Meta-Analysis of FL Techniques vs Key Performance Metrics

FL Technique Family	AC C	F1	Precision	Recall	FPR	Comm. Efficiency	Comp. Efficiency
Foundational FL (FedAvg, FedProx, FedNova, FedDyn, FedOpt)	↔	↔	↔	↔	?	↔	↔
Robust / Byzantine-Resilient FL (Krum, Trimmed Mean, FLTrust)	↑	↑	↑	↑	↓	↓	↓
Backdoor-Resilient Extensions	↑	↑	↑	↑	↓	↓	↓
Communication-Efficient FL (QSGD, DGC, FedDrop, EE-FL)	↔	?	?	?	?	↑↑	↑
Privacy-Preserving FL (DP-FedAvg, Secure Aggregation)	↓	↓	↓	↓	↑	↓	↓
Personalized FL (Per-FedAvg, pFedMe)	↑	↑	↑	↑	↓	↔	↓
Hybrid / Split Architectures (SplitNN, Hybrid FL-IDS)	↑	↑	↑	↑	↓	↑	↑
IoT/IDS-Specific Implementations	↑	↑	↑	↑	↓	↔	

Cross-batch analysis indicates that improvements in accuracy and F1 are primarily driven by personalized, hybrid, and robust FL strategies, while false positive rate (FPR) is rarely reported outside IDS-specific implementations; communication efficiency gains are strongest in compression-based FL (10×–1000× reduction), hybrid and SplitNN approaches enhance computation efficiency through edge offloading, yet notable trade-offs persist as robust and privacy-preserving methods increase overhead, differential privacy may slightly degrade performance, and compression techniques lack systematic evaluation using IDS-specific metrics. The Gap analysis is shown in Table 18.

Table 18 summarizes the major research gaps identified across federated learning categories and outlines corresponding research opportunities for developing secure, privacy-preserving, and efficient IoMT intrusion detection systems.

Table 18. Gap Analysis Table: Federated Learning for IoT/IDS

FL Category	Observed Gaps / Limitations	Research Recommendations	Opportunities /
Optimization & Foundational FL (FedAvg, FedProx, FedNova, FedDyn, FedOpt)	– Performance degrades under non-IID / heterogeneous clients – Limited evaluation in IoT/IDS contexts – Lack of adversarial robustness – Metrics often limited to ACC; few report F1/Precision/Recall/FPR	– Develop heterogeneity-aware optimization techniques – Benchmark FL on IDS datasets with realistic IoT heterogeneity – Integrate robust aggregation with optimized convergence methods	
Robust / Byzantine-Resilient FL (Krum, Trimmed Mean, FLTrust, FLSHield, Awasthi 2021 [3])	– No single method is universally robust across all attack types – Limited evaluation on large-scale IoT deployments – Increased computation & communication overhead – Few combine robustness with privacy preservation	– Layered defense frameworks combining robust aggregation, monitoring, and DP – Lightweight Byzantine-resilient methods for resource-constrained IoT nodes – Empirical evaluation under multi-adversary scenarios	
Communication- & Energy-Efficient FL (Quantization, Sparse FL, FedDrop, EE-FL, IRS-assisted FL)	– Trade-offs between bandwidth/energy reduction and model accuracy – Lack of real-world IoT/Edge validation – Limited assessment of latency, staleness, or client dropout impact	– Adaptive communication/computation strategies tailored to IoT/Edge – Hybrid methods integrating sparsity, dropout, and asynchronous updates – Energy-aware FL with minimal impact on detection performance	
Privacy & Information Leakage (DP-FedAvg, Secure Aggregation, Hybrid Privacy frameworks)	– Utility loss under strict differential privacy budgets – Complex integration with secure aggregation – Limited evaluation on IDS-specific datasets – Membership inference and backdoor attacks still possible	– Efficient DP mechanisms minimizing utility loss – Combined privacy-robust FL for IoMT/Healthcare IDS – Empirical evaluation against sophisticated inference attacks	
Personalized & Hybrid FL (Per-FedAvg, pFedMe, FedMA, SplitNN, hybrid DL/ML ensembles)	– Increased computation & storage for personalized models – Communication overhead for heterogeneous architectures – Limited robustness and privacy evaluation – Few large-scale IoT or multi-modal benchmarks	– Lightweight personalization for IoT/Edge devices – Cross-modal hybrid FL for multi-sensor IoT/IIoT environments – Integrate robustness & privacy in personalized FL pipelines	
IoT/IDS-Specific FL Implementations (IoT-Defender, F-NIDS, Chatterjee et al., [8] TON_IoT, CIC-IDS2017)	– Limited testing under real-world dynamic IoT networks – Non-standardized metrics and benchmarks across studies – Label scarcity and concept drift not fully addressed – Adversarial robustness rarely evaluated	– Standardized evaluation protocols for FL-IDS – Data-efficient methods (transfer/meta-learning) for low-labeled environments – Realistic deployment studies addressing drift, heterogeneity, and attacks	

Key insights from the discussion indicate that heterogeneity remains the primary challenge in federated learning for IoT/IDS environments, as non-IID client distributions, diverse device capabilities, and multimodal IoT signals significantly reduce baseline accuracy. Trade-offs dominate system design, with communication efficiency, energy consumption, privacy, and robustness often conflicting and necessitating multi-objective optimization strategies.

Moreover, robustness and privacy are rarely integrated simultaneously, as most studies focus on either Byzantine resilience or differential privacy/secure aggregation rather than unified defenses. The lack of standardized IoT/IDS benchmarks, due to variations in datasets, evaluation metrics, and federation setups, further complicates cross-study comparisons. Finally, while personalization and hybrid architectures show strong potential for addressing heterogeneity, challenges related to scalability, computational overhead, and security remain open research problems.

4. Discussion

The systematic review and meta-analysis of federated learning (FL) applications in IoT and intrusion detection systems (IDS) provide several key insights into the current research landscape and practical implications for system design.

4.1 Interpretation of Findings

Our synthesis indicates that foundational FL algorithms such as FedAvg and FedProx establish a baseline for privacy-preserving distributed learning, demonstrating high accuracy on standard datasets but limited adaptation to non-IID and heterogeneous IoT environments [23]. The modest accuracy drops in IDS-specific contexts highlight the importance of dataset representativeness and client diversity in real-world deployment.

Robust and Byzantine-resilient FL methods offer promising protection against poisoning attacks, yet their computational and communication overheads limit practical adoption on resource-constrained IoT nodes. Similarly, communication- and energy-efficient FL frameworks substantially reduce bandwidth and energy consumption, but they come at the cost of model accuracy, which must be carefully balanced, particularly in latency-sensitive IDS scenarios.

Privacy-preserving approaches, particularly differential privacy (DP) and secure aggregation, effectively reduce information leakage but occasionally degrade model utility [6][14]. The findings emphasize that robustness, efficiency, and privacy cannot be treated in isolation, as integrated frameworks are still sparse.

Personalized and hybrid FL architectures demonstrate that local adaptation improves client-level performance in heterogeneous IoT networks [17][23]. This suggests that client-specific model optimization is critical for systems operating across devices with varying capabilities and threat profiles. Nonetheless, computational and communication overheads remain a barrier, particularly when scaling to large, distributed IoT environments.

Finally, IoT/IDS-specific implementations confirm the feasibility of FL for network intrusion detection, achieving strong performance metrics (accuracy, F1-score) on real-world datasets [2][16]. However, the lack of standardized evaluation benchmarks and real-time adversarial testing limits comparability and generalizability.

4.2 Implications for IoT/IDS FL Design

The results point to several design considerations for practical IoT/IDS FL systems:

1. **Heterogeneity-Aware Learning:** Algorithms must explicitly account for non-IID client distributions, diverse hardware capabilities, and varying data volumes to maintain accuracy and convergence in IoT deployments [23].
2. **Integrated Robustness and Privacy:** FL frameworks should combine robust aggregation with privacy-preserving mechanisms, providing resilience against poisoning and inference attacks while minimizing utility loss [6][14][27][29].
3. **Energy- and Communication-Efficient Protocols:** Strategies such as sparse updates, gradient compression, asynchronous aggregation, and adaptive client selection are essential to balance resource constraints and detection performance in distributed IoT networks [20].
4. **Personalization and Hybrid Models:** Incorporating client-specific adaptation and hybrid deep learning/ML ensembles improves local detection capabilities and system-level performance [17].
5. **Standardization of Benchmarks and Metrics:** Consistent evaluation protocols, including accuracy, precision, recall, F1-score, false positives, and energy usage, are necessary to ensure reproducibility and facilitate cross-study comparisons.

4.3 Positioning of Research Gaps

The identified gaps suggest a future research agenda focused on developing federated learning algorithms capable of adapting to dynamic, non-IID IoT environments and concept drift in IDS datasets; designing lightweight frameworks that jointly optimize robustness and privacy while preserving data confidentiality; exploring energy-aware and communication-efficient protocols for scalable, real-world IoT/IDS deployments; leveraging meta-learning, transfer learning, and personalization techniques to enhance performance under data scarcity and device variability; and establishing realistic IoT/IDS testbeds and standardized evaluation datasets to improve reproducibility and enable fair comparison across federated learning frameworks.

Overall, the discussion highlights that while FL holds significant promise for secure, distributed IoT/IDS applications, addressing multi-dimensional trade-offs between robustness, privacy,

personalization, and resource efficiency remains critical. These insights form the foundation for subsequent conclusions and recommendations, positioning future work toward practical and scalable deployment of federated IDS in heterogeneous IoT ecosystems.

5. Conclusion

This study presents a comprehensive synthesis of over 60 recent works on federated learning (FL) applications in IoT-based intrusion detection systems (IDS), offering a structured assessment of algorithmic strategies, performance outcomes, and system-level trade-offs. Our meta-analysis demonstrates that FL approaches, particularly FedAvg, FedProx, and their personalized variants, achieve competitive detection performance while preserving client data privacy, even in highly distributed and heterogeneous IoT environments.

The analysis reveals critical trade-offs among robustness, privacy, personalization, and resource efficiency. Robust aggregation and Byzantine-resilient methods effectively mitigate poisoning attacks, but incur computational and communication overheads. Privacy-preserving techniques such as differential privacy and secure aggregation reduce information leakage but may impact model utility. Personalized and hybrid architectures enhance local detection performance, highlighting the necessity of client-specific adaptation in heterogeneous networks. Communication- and energy-efficient protocols are essential for practical deployment, yet balancing these optimizations against accuracy remains a challenge.

Our gap analysis identifies several urgent research directions, including adaptive learning under non-IID and dynamic IoT environments to maintain model accuracy amid heterogeneous and evolving client data; joint optimization of robustness, privacy, and resource efficiency to ensure secure and practical federated learning deployment at scale; standardized benchmarking and realistic IoT/IDS datasets to improve reproducibility and enable fair comparison of federated learning frameworks; and the integration of personalization and transfer/meta-learning approaches to enhance client-level performance while minimizing global model degradation.

In conclusion, while FL offers a promising paradigm for privacy-preserving, distributed IDS in IoT networks, realizing its full potential requires holistic frameworks that reconcile security, efficiency, and adaptability. The insights provided in this study offer a roadmap for researchers and practitioners to design scalable, robust, and privacy-conscious FL systems, paving the way for next-generation IoT security solutions.

Compliance With Ethical Standards

Conflicts of interest: Authors declared that they have no conflict of interest.

Human participants: The conducted research follows ethical standards and the authors ensured that they have not conducted any studies with human participants or animals.

References

- [1] Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 51(4), 1–35.
- [2] Alrawashdeh, K., & Purdy, C. (2019). Toward an online anomaly intrusion detection system based on deep learning. *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 195–200.
- [3] Awasthi, P., et al. (2021). Byzantine-resilient learning on heterogeneous datasets via bucketing. *Advances in Neural Information Processing Systems*.
- [4] Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [5] Bhagoji, A. N., Chakraborty, S., Mittal, P., & Calo, S. (2019). Analyzing federated learning through an adversarial lens. *Proceedings of the 36th International Conference on Machine Learning (ICML)*.
- [6] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & Seth, K. (2017, October). Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.
- [7] Cheng, X., et al. (2021). A federated learning-based intrusion detection system for IoT networks. *IEEE Access*, 9, 150516–150526.
- [8] Chatterjee, M., et al. (2021). Federated learning-based intrusion detection for IoT networks. *IEEE Internet of Things Journal*.
- [9] Dangeti, P. (2017). *Statistics for machine learning*. Packt Publishing.
- [10] Dinh, C. T., Tran, N. H., & Nguyen, T. D. (2020). Personalized federated learning with moreau envelopes. *Advances in Neural Information Processing Systems*.

- [11] Fallah, A., Mokhtari, A., & Ozdaglar, A. (2020). Personalized federated learning with theoretical guarantees: A meta-learning approach. *Advances in Neural Information Processing Systems*.
- [12] Fang, M., Cao, X., Jia, J., & Gong, N. Z. (2020). Local model poisoning attacks to Byzantine-robust federated learning. *USENIX Security Symposium*.
- [13] Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1322–1333.
- [14] Geyer, R. C., Klein, T., & Nabi, M. (2017). Differentially private federated learning: A client-level perspective. *Advances in Neural Information Processing Systems Workshop*.
- [15] Hard, A., et al. (2018). Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*.
- [16] Hady, A. A., et al. (2020). IoT intrusion detection using machine learning techniques. *IEEE Access*, 8, 128108–128118.
- [17] Jiang, Y., et al. (2021). Improving federated learning personalization via model agnostic meta-learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- [18] Jiang, Y., et al. (2021). FLEX: Federated learning via exploiting local models. *IEEE Transactions on Parallel and Distributed Systems*.
- [19] Jia, J., et al. (2018). Attriguard: A practical defense against attribute inference attacks via adversarial machine learning. *USENIX Security Symposium*.
- [20] Jin, S., et al. (2020). Hierarchical federated learning across edge and cloud. *IEEE Internet of Things Journal*.
- [21] Jahangir, A., et al. (2024). Empirical evaluation of federated learning for IoT intrusion detection. *IEEE Access*.
- [22] Koroniotis, N., et al. (2019). Towards the development of realistic botnet dataset in the internet of things for network forensic analytics: Bot-IoT dataset. *Future Generation Computer Systems*, 100, 779–796.
- [23] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 429–450.
- [24] Li, T., et al. (2020). On the convergence of FedAvg on non-IID data. *International Conference on Learning Representations (ICLR)*.
- [25] Liu, Y., et al. (2022). Hierarchical federated intrusion detection system across edge and cloud. *IEEE Internet of Things Journal*.
- [26] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1273–1282.
- [27] Melis, L., Song, C., De Cristofaro, E., & Shmatikov, V. (2019). Exploiting unintended feature leakage in collaborative learning. *IEEE Symposium on Security and Privacy*, 691–706.
- [28] Musharraf, M., et al. (2025). Explainable federated learning model for IoT intrusion detection. *IEEE Access*.
- [29] Nasr, M., Shokri, R., & Houmansadr, A. (2019). Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. *IEEE Symposium on Security and Privacy*, 739–753.
- [30] Page, M. J., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71.
- [31] Shokri, R., et al. (2017). Membership inference attacks against machine learning models. *IEEE Symposium on Security and Privacy*, 3–18.
- [32] Shi, Y., et al. (2021). Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Transactions on Neural Networks and Learning Systems*.
- [33] Sicari, S., Rizzardi, A., Grieco, L. A., & Coen-Porisini, A. (2015). Security, privacy and trust in internet of things: The road ahead. *Computer Networks*, 76, 146–164.
- [34] Sun, Y., et al. (2022). Split neural networks for intrusion detection in IoT systems. *IEEE Access*.
- [35] Wang, H., et al. (2020). Federated learning with matched averaging. *International Conference on Learning Representations (ICLR)*.
- [36] Wang, S., et al. (2023). Robust federated learning for IoT intrusion detection under poisoning attacks. *IEEE Internet of Things Journal*.
- [37] Xu, J., et al. (2021). Federated learning for IoT intrusion detection: A deep learning approach. *IEEE Access*, 9, 136740–136752.
- [38] Zheng, Z., et al. (2022). Hybrid federated learning for network security in IoT systems. *IEEE Transactions on Network Science and Engineering*.