Computer Fraud and Security
ISSN (online): 1873-7056

Federated Machine Learning for IoT/IoMT Security and Network Intelligence in HIPAA-Regulated Environments

Nagappan Nagappan Palaniappan

Fynbosys, USA

Abstract

The increase in the deployment of Internet-of-Things (IoT) and Internet-of-Medical-Things (IoMT) devices in healthcare systems has posed unprecedented cybersecurity challenges that compromise patient safety and data integrity, on the one hand, but introduce demanding regulatory compliance provisions under HIPAA systems, on the other hand. Conventional centralized machine learning methods of network security are not sufficient in healthcare settings where privacy laws do not allow aggregation of uncoded patient data, device logs, and network traffic patterns in centralized repositories. The new concept called federated learning represents a groundbreaking solution that promotes the idea of collaborative intelligence between the distributed nodes without the need to place raw data in a centralized place. The suggested architecture coordinates distributed device fingerprinting, real-time anomaly detection, and behavioral analytics among heterogeneous network nodes, which include medical devices, clinical equipment, and enterprise IoT endpoints, by local model training and encrypted parameter aggregation. Privacy-preserving solutions such as differential privacy, homomorphic encryption, and secure multi-party computation are designed such that even updates of a model cannot be deanonymised to reveal sensitive information whilst maintaining the same detection ability as in centralized methods. Experimental analyses reveal that federated intelligence can significantly improve the detection of advanced multi-stage attacks and low-frequency anomalies by combining patterns that can be seen across many institutions and achieve high accuracy in device fingerprinting and anomaly detection with a low false positive rate, which would be appropriate in a clinical setting. The framework has been able to strike the right balance between the necessity to ensure security and the need to provide privacy, which allows healthcare institutions to combine their efforts to protect against changing cyber threats without violating data sovereignty and regulatory requirements. Application in a wide range of healthcare settings confirms that federated principles hold strong performance in the context of inherent heterogeneity of devices in terms of population, distribution, and computational resources, and can offer a technically viable route to improved network intelligence in controlled healthcare environments.

Keywords: Federated Learning, Internet-of-Medical-Things Security, HIPAA Compliance, Privacy-Preserving Machine Learning, Healthcare Cybersecurity

1. Introduction

The development of Internet-of-Things (IoT) and Internet-of-Medical-Things (IoMT) devices in healthcare systems has radically altered the processes of providing patient care, improving operational efficiency, and the process of clinical decision-making. In modern healthcare environments, the implementation of hundreds and thousands of connected medical devices (such as systems that monitor patients and administer medication, diagnostic apparatus, health sensors, etc.) is a routine practice. The medical device security environment has been significantly advanced, and the healthcare sector has witnessed unprecedented cybersecurity threats that are directly affecting patient safety and data safeguard. The transformation to interconnected medical gadgets has developed complicated security settings where the conventional security methods are not effective in dealing with multifaceted threats to the healthcare infrastructure [1].

The peculiarities of healthcare IoT/IoMT settings introduce tough security challenges that cannot be compared to enterprise networks. Medical devices have legacy software, restricted computing resources, and long working life cycles, which do not allow them to have frequent security updates. The enterprise IoT security environment presents promising trends in terms of device vulnerability and network exposure, and healthcare organizations have a problem keeping full visibility of their growing device ecosystems. The non-homogeneous nature of these environments, with devices of different vendors with different protocols and security postures, presents a huge attack surface that can be exploited. The recent studies on device security connectivity show that healthcare organizations struggle a lot in detecting, categorizing,

and safeguarding medical devices on their network, with some of them working without any proper security provisions or tracking functionalities [2].

The standard centralized machine learning models of network security and anomaly detection do not have any chance of success in the healthcare setting of the Health Insurance Portability and Accountability Act (HIPAA). These regulatory frameworks provide rigid safeguards to Protected Health Information (PHI) and forbid the recombination of unfinished patient information, equipment records, or network traffic patterns in central repositories. This leads to the fact that the traditional security analytics approaches, which rely on merging data from various sources to train the models, become not only legally but also ethically unsustainable, posing the essential contradiction between security needs and privacy demands. Federated Learning (FL) can be considered a paradigm shift to this dilemma, as it allows the nodes of machine learning to learn collaboratively without requiring the raw data to be located at a central location. FL ensures data sovereignty by training local models on individual devices or network segments, and model parameters are only aggregated in an encrypted form, creating a globally intelligent security system.

Security Dimension	Healthcare IoT/IoMT Characteristics	Enterprise Security Implications
Device Lifecycle Management	Extended operational lifecycles with manufacturing security responsibility shifts	Asset inventory maintenance challenges and vulnerability management complexity
Protocol Security	Unencrypted communication protocols are prevalent across device populations	Network segmentation deficiencies and monitoring capability gaps
Authentication Controls	Default credentials remain unchanged throughout operational deployment	Inadequate security controls enabling lateral movement opportunities
Visibility and Classification	Difficulties distinguishing medical devices from general IT equipment	Risk management capability limitations and exposure assessment challenges

Table 1: Medical Device Security Challenges and Enterprise IoT Vulnerabilities [1,2]

2. Healthcare IoT/IoMT Security Landscape and Regulatory Constraints

The modern healthcare tech environment represents a complex structure of interconnected systems that involves various levels of clinical and operational infrastructure. IoMT devices that engage directly with patients are cardiac monitors, insulin pumps, ventilators, and wearable biosensors, which constantly capture and transmit physiological information. Hiding behind these frontline devices lies a sophisticated supporting infrastructure of Picture Archiving and Communication Systems (PACS), Electronic Health Record (EHR) systems, nurse call systems, and building automation controls. The multi-layered structure offers a wide range of possible points of entry to malicious actors, and at the same time makes tracing, security monitoring, and threat identification hard. The security in the medical equipment has been changing drastically, where ownership and responsibility in device security have changed throughout healthcare institutions. There have been changes in the manufacturing practices to accommodate security concerns in the earlier stages of the lifecycle of the devices, but there is still a significant gap in the implementation process as healthcare providers endeavor to actualize security requirements into the clinical processes. Medical device ecosystems have become increasingly complicated in recent times, with organizations implementing more advanced connected technology, making it difficult to keep up with security teams trying to keep a complete list of assets and vulnerability management programs [1].

A number of circumstances combine to make the healthcare IoT/IoMT settings highly susceptible to cyberattacks. A variety of medical devices were created and implemented in times when little attention was paid to cybersecurity concerns, and created systems that lacked any fundamental security features, like authentication methods, encryption, or secure boot-up procedures. The urgency of healthcare processes places a strong demand on maintenance windows, so that it is difficult to implement security patches or firmware updates without interfering with patient care. Also, medical equipment has a long user life, which implies that devices are not obsolete because their manufacturers will no longer support them or update them by the time they are made. The analysis of enterprise IoT security indicates worrying trends in the areas of device visibility, classification accuracy, and risk management potential of healthcare networks. Organizations also often struggle to tell the difference between medical devices and general IT equipment, and this

results in insufficient security measures and monitoring loopholes. As the enterprise security environment evidences, the network of devices is usually unencrypted, not segmented correctly, and uses default passwords, which stay unchanged during the entire period of the device's lifetime. Such system-wide weaknesses make attackers' chances to exploit individual devices and use them as entry points to wider network intrusion, which may become critical care delivery systems and patient safety mechanisms [2].

The threat environment on healthcare infrastructure has changed radically with regard to sophistication and frequency. The APTs have a new target, which is medical facilities, because they have realized that healthcare organizations usually have valuable patient data, and they have limited security budgets and are operationally mandated to pay ransom. The attack vectors are the exploitation of unpatched vulnerabilities on medical devices, lateral movement via poorly segmented networks, and the compromise of parts of the supply chain. The impacts of effective attacks are not limited to data breaches but may include possible patient damage by controlling device functionality or derailing important care provision systems. The statistics of data breaches in healthcare indicate the growing scale of security events involving patient data and clinical activities. The rate of breaches has grown significantly, and healthcare facilities have had incidents that expose patient records of millions of patients each year. The economic value of such breaches is not only restricted to direct costs of remediation, but also regulatory fines, litigation fees, and reputational losses that result in loss of patient confidence and organizational sustainability. Patterns in breaches indicate that access to medical devices that are connected to the network and poor network security controls are often the initial attack points, which allow adversaries to develop a long-lasting access point to healthcare networks and steal sensitive information [3].

The HIPAA and other similar regulatory measures provide a detailed set of requirements to safeguard the confidentiality, integrity, and availability of PHI with strong limitations on security implementation strategies. The HIPAA Security Rule requires both administrative, physical, and technical protection, but does not offer much specific information on how to secure IoT/IoMT devices. Limitations of PHI use and disclosure by the Privacy Rule virtually forbid most of the conventional security analytics strategies, which would consolidate patient-related information to analyze threats. The Business Associate Agreements (BAAs) also complicate the process of cooperative security work among healthcare organizations because data disclosure can be inhibited, even with legitimate reasons to do so, by the contract regulations. These legal restrictions have an inherent paradox in that effective security analytics are usually best served by large, heterogeneous datasets that allow powerful pattern recognition and anomaly detection, but privacy regulations do not allow the centralization of the very types of data that are best suited to such purposes. HIPAA protects device logs with timestamps and patient identifiers, network flow logs that may disclose patterns of treatment, and behavioral analytics that may inform inferences on health conditions. This, therefore, means that conventional Security Information and Event Management (SIEM)-powered mechanisms and centralized machine learning-based solutions should not be allowed to perform their full capacity in healthcare settings, and this necessitates new methods to address the issue of balancing security effectiveness and the rule of law.

Privacy Dimension	Breach Characteristics	IoMT Data Protection Requirements
Attack Vectors	Connected medical device vulnerabilities and network security control inadequacies	Stringent protection mechanisms throughout the collection, transmission, and analysis processes
Breach Frequency	Escalating incident magnitude exposes patient records annually	Defense against inference attacks, model inversion attempts, and unauthorized reconstruction
Financial Impact	Regulatory penalties, litigation expenses, and reputational damage	Raw medical data retention at local facilities with cryptographic parameter protection
Data Sensitivity	Persistent network access enabling extended exfiltration periods	Privacy-preserving techniques prevent sensitive information exposure

Table 2: Healthcare Data Breach Patterns and IoMT Privacy Requirements [3,4]

3. Federated Learning Architecture for Distributed Healthcare Security

FL represents a model shift in the relation of the models and the data. Instead of bringing the data to the cloud for training the model, FL brings the models to the data where they are locally trained at distributed clients and only the model parameters are aggregated to form global intelligence. The architecture has been able to offer successful usage in multiple healthcare IoT/IoMT security domains where data sensitivity, regulations and heterogeneity of networks make customary centralized protocols lack efficiency. Feasible data security and privacy in the IoMT domain is challenging since the medical data contains extremely sensitive information that implies patients, who should be equipped with strict security practices during the collection, transmission and analysis of such data. To solve these issues, federated learning architectures are proposed such that raw medical data is never shared in the local health organizations. Instead, federated model parameters are shared after being cryptographically encrypted. IoMT data is too sensitive to avoid different forms of privacy-defending mechanisms against several attack vectors like inference, model inversion attacks and data reconstructions which are unauthorized by the system users [4].

The federated learning architecture will consist of a number of components that work together to enable privacy aware security analytics of the healthcare network. The components of the architecture could include local learning agents deployed at the edge tier either at the network segment, medical device endpoint or facility level infrastructure. These agents can extract features from device behavior and network traffic patterns without disclosing PHI or any other personally identifying information. The features are derived from device behavior including communication patterns, protocol adherence, time patterns, and resource consumption patterns. These don't disclose PHI; however, they do have security and privacy implications about how a device behaves. Fingerprinting of each device, creating baselines, and detection of anomalous behavior is done using local models trained at each participating node. Supervised, semi-supervised, or unsupervised learning techniques are utilized, depending on the available labeled data and security objectives of the implemented solution.

The aggregation layer represents the key innovation of the federated model and enables the aggregation of local updates on distributed nodes into a global model. Nodes do not send the raw data or time-series data, instead they transmit the model's parameters that encode the patterns the model has learned from its dataset. The parameters are encrypted using a homomorphic encryption or secure multi-party computation protocol and the updates sent by the devices are encrypted in such a way that no information can be leaked from the encoded message that could enable reconstruction of sensitive data. Research into federated learning techniques using differential privacy showed that it is possible to achieve accuracy and strong mathematical privacy guarantees against leakage. The aggregation should also allow for non-identically distributed and dependent data, as each healthcare facility has a different population of devices, demographics of patients and threats.

Nevertheless, aggregation algorithms are improved by incorporating aspects of local data quality, strong aggregation techniques that can discard malicious updates, and other personalized methods that allow nodes to maintain more specialized parts of the models to fit their specific use cases [5].

Finally, the cycle of federated learning concludes with model validation and the update of global models to each participating node. As described above, model validation can be performed using a holdout test set at each participating institution, adversarial testing to test the model's resistance to evasion attacks, or fairness testing to test the model's distributive fairness across devices of varying types and network conditions. Once validated, this new global model is sent back to the edge nodes making local detection capabilities of edge nodes stronger, by integrating information from the whole federated network.

This allows for continuous improvement as nodes encounter new device types, attack vectors and evolutions in normal operation. Identity aware metadata also allows for device-contextual information to be integrated into the framework in a privacy preserving manner. In addition, network flow capabilities provide additional discriminative capability in the form of statistical analysis of traffic volumes, protocol distributions, timing patterns and the structure of the communication graph. The architecture is designed to support heterogeneity on several dimensions to reflect the real-world healthcare setting. This heterogeneity consists of device heterogeneity (varying kinds of medical instruments), data heterogeneity (non-independent and identically distributed (IID) local datasets), and computational heterogeneity (widely varying computational power available on the participating nodes).

Architectural Component	Privacy Protection Technique	Heterogeneity Accommodation
Model Parameter Transmission	Differential privacy with mathematical leakage guarantees	Adaptive weighting based on local data quality assessment
Aggregation Process	Homomorphic encryption prevents intermediate decryption	Robust methods for detecting and excluding malicious updates
Node Participation	Secure multi-party computation distributing trust	Personalization mechanisms for specialized model components
Feature Extraction	Statistical behavioral features excluding application-layer content	Packet inter-arrival distributions and protocol usage patterns

Table 3: Federated Learning Privacy Mechanisms and Aggregation Strategies [5,6]

4. Implementation framework and technical methodology

In health care in IoT/IoMT security, real-world deployments of federated learning would require accommodating for managing network architecture, feature extraction, the model, and operationalization. Requirements include security efficiency, computational efficiency, regulatory compliance, and clinical adoption and use, while remaining secure against cyber attack or system failure. Network instrumentation is the foundational element of the security model. Key network instrumentation points, such as passive network taps and span ports, should be factored into the design of healthcare IT infrastructure to monitor network traffic at data center egress and ingress points, device segment and inter-facility network links. Deep packet inspection capabilities extract protocol-level characteristics without decrypting or exposing protected health information (PHI).

Feature engineering is one important technical challenge in security analytics, since a representation with a lot of information is usually desired, but may be difficult to construct in a privacy-preserving manner. Features used in device fingerprinting include MAC address patterns, protocol implementations, firmware signatures, and physical layer features. These properties allow to characterize and classify individual network entities without the need to access content or any of their identifiers at the level of application protocols. In addition, the set of network flow features abstracts flows by employing statistical distributions over packet sizes, inter-arrival times, port usage and protocol sequences. IoT device identification has been implemented with high accuracy through analysis of both traffic pattern properties and a set of statistical and behavioral characteristics using machine learning on the network flow data. These methods can be used for device fingerprinting without deep packet inspection or access to encrypted packet payloads. Metric-based data, such as packet inter-arrival time distributions, payload size distributions, protocol usage patterns and communication graph topologies, may be used to classify devices and detect compromised or malfunctioning devices [6].

The model architecture is a hybrid model that uses different ML algorithms tuned to their specific security objectives. The Convolutional Neural Networks take the sequential network traffic as input and detect patterns in the sequence of packets, the exchanges of packets, and the interactions of the various protocols on the network. Long Short-Term Memory networks model temporal dependencies on the devices to detect anomalies occurring over long timescales. Random forests provide interpretable decision logic for device classification and rule-based alerts. Using autoencoders to learn a compressed representation of normal behavior allows for unsupervised detection of new attacks without the need for large-scale datasets of attack types and has the potential to address the issue of real-world asynchrony and heterogeneity of devices in federated learning. In a hierarchical aggregation topology, facility-level aggregators collect updates on local network segments before aggregating for the organization or multi-institution federation.

Deep learning techniques have proven to be a powerful technique for identifying network intruders, malware and anomalies in complicated IT environments. Several studies have compared different implementations of deep learning architectures and found that ensemble models of different types outperform all others. Models which use convolutional layers to capture spatial features, recurrent layers to capture temporal features, and attention mechanisms to focus on important traffic areas have proved powerful in detecting subtle indicators of an attack in the noise of the network. These models are especially relevant in healthcare, where clinical workflows confound attack behavior and behavior varies by

device. Particularly, deep learning methods can be successfully applied to high dimensional feature space of network security data and processed efficiently to meet the real-time demands [7].

The models' update frequency can also be dynamically adjusted within the algorithm based on learning speed, communication and computational costs. For example, an adaptive scheduling algorithm can change the update frequency of the models when a large model drifts, anomalies occur, and environmental changes determine that retraining can further improve performance. Transport-layer encryption and aggregation of model updates rely on secure network communication techniques. To ensure the integrity of the communication channel, mutual authentication over TLS is used. Homomorphic encryption, which enables model parameters to be computed on encrypted data, allows aggregation while avoiding decryption on intermediate nodes. Although this adds computational cost, it is acceptable for the typical update frequency and model sizes of IoT/IoMT security applications, and transparent to federated learning algorithms.

Model Architecture	Security Detection Capability	Healthcare Context Application
Convolutional Neural Networks	Spatial feature extraction from sequential traffic data	Pattern identification in packet sequences and protocol exchanges
Long Short-Term Memory Networks	Temporal dependency capture across extended timeframes	Anomaly detection obscured by legitimate clinical workflows
Ensemble Approaches	Superior detection through a multiple classifier combination	High true positive rates with controlled false alarm frequencies
Attention Mechanisms	Critical traffic segment focus in noisy environments	Multi-stage attack detection across organizational boundaries

Table 4: Deep Learning Architectures and Intrusion Detection Performance [7,8]

5. Experimental evaluation and performance analysis

Realistic testing environments were created to assess the proposed federated learning framework, considering the complex architecture, heterogeneity, and strict operational requirements of the healthcare IoT/IoMT ecosystem. The performance was evaluated concerning detection performance, privacy preservation, computational efficiency, and regulatory compliance. We tested the system in our simulation environment and conducted tests in real-world healthcare facilities. The testbed supported a wide range of medical devices, some of which are among the most commonly used medical devices in patient environments. Continuous cardiac monitors and pulse oximeters, as well as multi-parameter vital signs monitors from different manufacturers, and infusion devices, such as smart and patient-controlled analgesia pumps and automated dispensing machines, were included. Other devices included ultrasound machines, ventilators, and dialysis machines. The heterogeneity of these devices gives us the ability to evaluate the accuracy of device fingerprinting and behavioral anomaly detection in realistic scenarios.

Network traffic was produced through the replay of captured and anonymized clinical network traffic and through synthetic traffic that simulated medical device communications and injections of simulated attacks. Expected network traffic included device registration, continuous medical device data streams, periodic medical device configuration updates, and maintenance activities that are consistent with clinical practice. Simulated scenarios include reconnaissance scans, exploiting device-specific vulnerabilities, command and control traffic, exfiltration of data, and controlling devices. The application of federated learning to IDS has shown that distributed training can achieve detection performance at par with centralized training with important privacy advantages. The sensitivity and accuracy of federated based intrusion detection systems in different configurations and environments have been evaluated. The rate of true positives was shown to be high while the false positive rate was at moderate levels and within acceptable operational limits. A particular strength of this method is in multi-subnet or multi-organization attack detection where standalone systems lack sufficient context. In practice, evaluations show that federated models generalize well and are capable of detecting new attack variants not included in the training data set using behavior modeling and anomaly detection techniques [8].

During detection, the federated learning approach outperformed local models. FL-based device fingerprinting achieved the best known results on the known device type, distinguishing between identical devices from different manufacturers

and helping to identify discrepancies in firmware versions that could indicate potential security relevance. This unknown device detection demonstrated a high degree of accuracy in identifying and allowing security personnel to investigate novel devices, and it was found to be at least acceptably sensitive for use in a simulated attack with few false positives. False positives in a clinical environment are excessive, create alert overload for staff, and may even be detrimental to operations and patient safety. The federated model also provides additional value in enabling the detection of more advanced multi-stage attacks and low-frequency anomalies for which local evidence alone would not be sufficient to achieve detection.

The computation requirements were determined for all nodes participating in the system. The edge agents on the resource-constrained IoT gateways have a small memory and CPU footprint under normal monitoring conditions and slightly increased resource requirements during local model training epochs. Facility-level aggregators required as much CPU and memory as conventional IT management platforms. Additionally, communication overhead was relatively low in terms of bandwidth and communication latency. Update sizes may also be reasonable, depending on model design and compression. Blockchain-based federated learning approaches show that distributed ledger technologies (DLTs) can improve security and trust in federated model aggregation by providing permanent audit trails for updates, cryptographic assurances of authentication and integrity of participants, and decentralized consensus mechanisms to ensure malicious behavior in creating or modifying the global model. These federations are well suited to multi-organization healthcare use cases with untrusted participants and a need for governance transparency. In trials the blockchain overhead for federated learning update frequencies typical of IoT security use cases was modest. The consensus latencies were seconds rather than the multi-minute cryptocurrency-style confirmations [9].

Membership inference attacks, a method of evaluating the privacy guarantees of the shared model parameters by determining whether certain behaviors are used as training data, have achieved nearly random performance, confirming that privacy is well-preserved. Model inversion attacks to infer properties of inputs based on model parameters were unsuccessful. Differential privacy analysis measured implicit privacy guarantees of noise injection approaches. Organizations can use these analyzes to weigh privacy and utility trade-offs of different approaches through explicit privacy guarantees or in the context of privacy regulations or risk tolerances. Evaluations for these models provide strong evidence that when compared to a centralized machine learning baseline, federated learning achieves a better privacy-utility trade-off, as the federated models' detection accuracies were comparable to the baseline, with meaningful privacy and regulatory compliance advantages. Additionally, the federated model was shown to be more strong to distributional shifts and new attacks, as decentralized training between decentralized organizations allows for better generalization than machine-learning models trained at a single institution that may easily overfit the characteristics of that institution. Federated learning applications in healthcare fields show that cooperative learning with privacy guarantees is a promising path for knowledge transfer between several healthcare institutions, all in compliance with existing data privacy legislation. This federation technique allows researchers to mine decentralized patient groups of information, without the need to concentrate patient health-related data. This allows predictive models and clinical decision support systems to be developed using the power of big data while still respecting the privacy of individual patients and the data sovereignty of institutions [10].

Conclusion

The Internet-of-Things and Internet-of-Medical-Things are growing in importance, while cyber threats increase and privacy regulation keeps developing. These realities present novel challenges in the development of cybersecurity strategies for healthcare systems that do not conform to the customary model. The technical notion of federated machine learning shows how security and privacy measures can be reconciled to allow healthcare organizations to defend against systemic cyber threats while upholding ethical and legal obligations to safeguard patient information security. The proposed architecture shows how privacy-preserving collaborative learning considerably improves threat detection performance with an effectiveness similar to a centralized model, while maintaining a high degree of privacy using differential privacy methods, homomorphic encryption and secure aggregation. The framework is likely to be most effective for advanced attacks and low-prevalence anomalies for which the evidence for the attack spreads across greater than two institutions. The proposed framework may be able to detect attacks that are difficult to detect because they are crossing institutional boundaries using federated intelligence. Experimental evaluation has demonstrated that the framework can be instantiated to achieve high utility across the heterogeneous device populations, data distributions, computational capabilities found in real-world healthcare settings and deployed across sites with heterogeneous resource

availability and operational constraints. The privacy-utility trade-offs of the framework can be configured to support scale according to institutional risk tolerance, regulatory interpretation, and a variety of compliance requirements and security objectives. Implementation challenges include governance model (what model to use, how to engage in models and policies regarding the usage of data), legal model (what laws to apply to model parameters), and economic model incentives (incident costs to discourage malicious behavior and encourage using provably secure models). These challenges can be transplanted to finance, critical infrastructure, and government agencies where privacy-security trade-offs exist and the need to decentralize private data and computational power creates an advantage for collaborative learning. Digital transformation will allow health care to achieve the clinical and operational efficiencies and patient outcomes needed to meet the needs of health care today, but requires security systems that can protect the critical medical infrastructure without sacrificing privacy. Federated learning solutions will ensure that the medical infrastructure on which the future of health care depends will be an intrinsic characteristic of the way health care works, and a security solution that preserves privacy will be a necessary condition of patient safety and the trust of the community in connected health care systems.

References

[1] Cybellum, "The State of Medical Device Security in 2024: Challenges, Trends, and Ownership Shifts," 2024. [Online]. Available: https://cybellum.com/blog/the-state-of-medical-device-security-in-2024-challenges-trends-and-ownership-shifts/

[2] Forescout Research Labs, "The Enterprise of Things Security Report: State of IoT Security." [Online]. Available: https://www.forescout.com/the-enterprise-of-things-security-report-state-of-iot-security/

[3]Steve Alder, "Healthcare Data Breach Statistics," HIPAA Journal LLC, 2025. [Online]. Available: https://www.hipaajournal.com/healthcare-data-breach-statistics/

[4] Min Chen, et al., "A 5G Cognitive System for Healthcare," MDPI, 2017. [Online]. Available: https://www.mdpi.com/2504-2289/1/1/2

[5] Adrien Banse, et al., "Federated Learning with Differential Privacy," arXiv, 2024. [Online]. Available: https://arxiv.org/abs/2402.02230

[6] Arunan Sivanathan et al., "Classifying IoT Devices in Smart Environments Using Network Traffic Characteristics," IEEE, 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8440758

[7] Mohamed Amine Ferrag et al., "Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study," ScienceDirect, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S2214212619305046

[8] Babatunde Olanrewaju-George, Bernardi Pranggono, "Federated learning-based intrusion detection system for the internet of things using unsupervised and supervised deep learning models," ScienceDirect, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2772918424000341

[9] Umer Majeed; Choong Seon Hong, "FLchain: Federated Learning via MEC-enabled Blockchain Network," IEEE, 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8892848

[10] Daniel J. Beutel, et al., "Flower: A Friendly Federated Learning Research Framework," Arxiv, 2020. [Online]. Available: https://arxiv.org/abs/1812.06127