

Blockchain-Enabled Intrusion Detection for Medical IoT Devices: A Privacy-Preserving Security Framework

Hashim Ali*

Department of Computer Science, Abdul Wali Khan University Mardan, Pakistan.

Corresponding Author Email: Hashimali@awkum.edu.pk

Asif Rahim

School of Computer and Information Security, Guilin University of Electronic Technology, Guilin 541004, China

Aftab Ahmed

Department of Computer Science, Abdul Wali Khan University Mardan, Pakistan

Umer Tanveer

Department of Computer Science, Abdul Wali Khan University Mardan, Pakistan

Kiran Falak Sher

Department of Computer Science, Abdul Wali Khan University Mardan, Pakistan

Wali Samad Khan

Department of Biotechnology, Abdul Wali Khan University Mardan, Pakistan

Shaista Bibi

Roots Milinium School System

ABSTRACT

In modern health care infrastructure the internet of medical things (IoMT) is become a backbone due to devices have resource limitation of risks which consider the security design for the future. We have introduce the farmwork which instruct the detection with blockchain technology with privacy that preserve the leaning to secure the deployments of cross site IoMT. Our aim to target in current system to target three critical weakness: the in centralized cloud monitor the single point of failure that's inherited, the low power devices have bottlenecks communication and the de to challenges of maintain privacy of patients due to threat detection in collaboration. To validate our system we have use the comprehensive CIC IoMT 2024 benchmark dataset, using various protocols on 40 different medical devices to

capture real world traffic. We achieved a 99% accuracy rate in detecting attacks by applying an ensemble classifier to a dimensionality-reduced feature set. In addition we have knowledge distillation for cutting the communication inter hospital 87% without guarantees of formal privacy . To cover the evaluation of five classes include spoofing, DDoS and access that unauthorized along to demonstration of thins frame work is viable for the deployment of the real world on edge gateway and resource constrain wearable , this enable the institution for collaborate securely without data which is centralized sensitive.

Keywords: Internet of Medical Things, Intrusion Detection, Blockchain Security, Federated Learning, Privacy-Preserving Analytics, Healthcare Cybersecurity

1. Introduction

Patient care is revolutionized through rapid growth of medical devices that is insulin pumps, sensors of glucose and ECG monitors. There is concern on cybersecurity significance. In current view, in IoMT ecosystem more then 127 billion devices, that produce daily patient sensitive data that estimated 2.5 quintillion bytes. Contrasting traditional IT infrastructure, to use easy for clinically the medical devices are primarily designed, also measure the expense of strict security. To disrupt device operations that are occasionally in consequences of life that endangered and that make atterective targets for patient privacy compromise.

In the current security strategies that are cloud based infrastructure for centralized detection of threats. There are three key vulnerabilities create through architural paradigm.

1. Failure of single point: At the central node risks compromise the exposing the all connected hospital devices
2. Privacy Conflicts: there are a lot of collection patient data in centralized server, it often confect with strict regulatory frameworks such as HIPAA.
3. Overhead Communication: raw data distributed from thousands of devices which are continues transmitted which incur the costs.

Due to responsibility and leadership we have need to ensure the benefits for working in organization that are recognized and rewarded. In other way offers in blockchain technology valid the alternate way for decentralized records that's are validate and immutable transaction. Though, present the blockchain executed for hardly suitable

for edge devices by limited power budgets (which frequently range from 1-10 watts) due to scalability restrictions and high computational upstairs. Furthermore, several recent attempts at integrating machine learning with blockchains have also relied on either naive rule-based anomaly detection methods or globally trained models. These methods in general are incapable of grasping the complex spatial and temporal dependencies presented during coordinated cyber-attacks on medical networks.

In light of these limitations, this work proposes a practical framework that will incorporate Federated Learning with blockchain-based audit trails and privacy-preserving techniques to facilitate real-time anomaly detection. Contrasting the theoretical setting that is common in most related works, the proposed system has been designed for-and evaluated on-realistic IoMT deployments within hospital settings. Evaluation will be done using the publicly available benchmark dataset of CIC IoMT 2024, containing 4.8 million instances of real traffic from 40 medical devices under realistic attack scenarios, including DDoS, spoofing, and unauthorized access.

The specific contributions of this study are as follows:

- **Efficient Feature Selection:** A methodology that reduces communication requirements by 87% while maintaining a 99% attack detection accuracy.
- **Real-World Validation:** Experimental results derived from actual medical device network traffic rather than simulations.
- **Edge-Compatible Security:** A blockchain-based access control framework specifically optimized for resource-constrained edge devices.
- **Guaranteed Privacy:** Quantified privacy assurances using differential privacy techniques, making the system suitable for secure multi-hospital collaboration.

2. Related Work and Gap Analysis

Existing IoMT Security Approaches

Recent research on intrusion detection for medical device networks spans a wide range of strategies. Using the CIC IoMT 2024 dataset, Lipsa et al. used a Random Forest classifier to obtain an accuracy of 99%, following dimensionality reduction to 16 features.. Their feature analysis identified IAT, packet flags, and protocol-specific

features as the main indicators of malicious activity and thus set a baseline for comparison. The TON_IoT dataset contributed telemetry data from seven IoT/IIoT devices over nine attack types. Federated learning experiments on TON_IoT showed that Random Forest outperformed deep learning methods, with an accuracy of 99.86% on highly imbalanced data, and thus suggested that ensemble methods handle the class imbalance inherent in medical device networks (typically 1–5% attack instances) better than neural methods.

Limitations with Current Approaches

Despite notable advances, extant work reveals several critical gaps:

1. Centralization Risk: Most of the IoMT security research deploys their models on cloud servers, repeating the same type of centralized attack surface vulnerability they presumably want to mitigate. Multi-institutional threat sharing remains limited.
2. Privacy–Accuracy Tradeoff: Previous works in distributed learning setup do not provide formal privacy analysis, hence hospitals cannot quantify what adversaries can infer about patient populations from the shared models.
3. Practicality of Deployment: Published methodologies often do not consider resource constraints. Most edge gateways, such as the NVIDIA Jetson Xavier with 8 GB RAM and 15 W, cannot run most of the deep learning inferences within the 145 ms latency for real-time alerts.
4. Validation Realism: Most studies use simulated data or small proprietary datasets. Multi-protocol real-world scenarios such as Wi-Fi, MQTT, and Bluetooth provided within the benchmark CIC IoMT 2024 are much closer to a hospital network in production.

3. Methodology

Framework Architecture

Our architecture is composed of three hierarchical layers:

Edge Layer: Individual medical devices and local gateways conduct preliminary anomaly detection using lightweight models that are trained during the federated rounds. In this case, whenever there are any anomalies detected, encrypted logs are generated and recorded onto the blockchain.

Coordination Layer: A federation server orchestrates multi-hospital learning without aggregating raw data. Each hospital trains the local models and transmits only model updates (gradient vectors) encrypted by homomorphic encryption to

avoid observing intermediate computations from the coordination server.

Blockchain Layer: Validator nodes maintain the immutable audit trail of security events, access control decisions, and model updates. Role-based access control policies are implemented through smart contracts, with verification provided by cryptographic proofs.

Feature Engineering and Dimensionality Reduction

The initial CICIoMT 2024 dataset consisted of 45 features over a total of 4.8 million instances, including 17 attack and benign classes. In order to reduce the communication overhead in a federated environment, Shapely-value-based feature selection was used—a gametheoretic approach that calculates the marginal contribution of each feature toward model predictions.

Dimensionality reduction procedure:

1. Train a baseline Random Forest using the full 45-feature dataset.
2. For each feature i , calculate its Shapley value, $\phi_i = \sum$ (the contribution of feature i to any coalition S upon addition).
3. Retaining features for which $\phi_i > 0.05$, the original 45 features are reduced to 16 features.

This reduction eliminates 64% of features while preserving 99% of detection accuracy; a particularly good result for bandwidth-constrained federated settings where the cost of transmission supersedes computation cost. Some of the top features identified include inter-arrival time, TCP/UDP packet flags, protocol type, flow duration, and connection count. Interestingly enough, IAT came up as the single most predictive feature in over 12 classes of attacks, which suggests that it is temporal patterns, rather than individual packets, that mark attacks from normal device operation.

Detection Models

We evaluated four ensemble and tree-based classifiers standard in intrusion detection literature:

Model	Accuracy	Precision (avg)	Recall (avg)	F1-Score (avg)
Random Forest	99%	0.98	0.94	0.96
XGBoost	98%	0.97	0.93	0.95

Model	Accuracy	Precision (avg)	Recall (avg)	F1-Score (avg)
Decision Tree	97%	0.95	0.91	0.93
SVM	98%	0.96	0.92	0.94

Random Forest was selected as the primary model due to superior handling of imbalanced data (inherent in medical networks where benign traffic dominates). Its bootstrapping and probability-based decision-making naturally weight minority attack classes more appropriately than single decision trees.

[Figure 1: Key Results – Model Performance on CIC IoMT 2024 Dataset]

Privacy-Preserving Federated Learning

These enable hospitals to collaborate without sharing raw patient data using three mechanisms:

- 1) Secure Model Updates: Each hospital trains local Random Forest models on their own device traffic. Instead of sending large gradient vectors > 50 MB in size, institutions share only hard predictions (class labels) on unlabeled proxy samples. Knowledge distillation this way reduces communication by about 87% while maintaining an accuracy within 1.8% of that obtained via centralized training.
2. Differential Privacy: Local gradients are perturbed with calibrated noise before transmission: $g_i = g_i + \text{Laplace}(0, \Delta g/\epsilon)$, with $\epsilon = 1.2$ to confer formal privacy guarantees. Subsampling leads to privacy amplification over 50 federated rounds with 10% device participation, and the total privacy budget then becomes $\epsilon_{\text{final}} \leq 1.2$, showing that the privacy loss remains statistically imperceptible.
3. Homomorphic Encryption: Model updates arrive encrypted at the coordination server, where averaging is done without decryption, so that the server only learns the aggregate direction, not the predictions from individual hospitals.

Blockchain Integration

Each IoMT device produces a transaction record on a permissioned blockchain-say, Hyperledger Fabric-so as to establish an immutable log of device behavior, including access attempts and anomalies detected. Access control is implemented through smart contracts: a device obtains an access token only when authenticated-such authentication must be accompanied by cryptographic proof. As a specific example, suppose Device A queries patient data from Device B. In this case, the smart

contract verifies that (1) A's certificate is valid, (2) the role of the requesting user allows for such access, and (3) no recent anomalies were detected that may affect B's integrity. This setup distributes the decision of access control across the validator nodes so that no healthcare administrator can grant unauthorized access unilaterally.

4. Experimental Validation

Dataset and Testbed

We evaluated on the CIC IoMT 2024 benchmark—a comprehensive dataset created by the Canadian Institute for Cybersecurity simulating realistic hospital environments. The dataset spans:

- **40 medical devices** (25 real hardware, 15 virtualized): ECG monitors, insulin pumps, glucose sensors, Holter monitors, pulse oximeters, respirometers connected via Wi-Fi, MQTT, and Bluetooth protocols
- **5 attack classes**: DDoS (distributed attacks across multiple devices), DoS (single-source attacks), reconnaissance (network scanning), MQTT-based misuse (protocol-specific attacks), and spoofing (identity impersonation)
- **4.8 million instances** with realistic traffic imbalance: ~99% benign, ~1% attacks

Training/test splits: 377,718 training instances (70%) and 98,432 test instances (30%), provided separately to avoid data leakage.

Results: Attack Detection Performance

After dimensionality reduction, Random Forest achieved 99% accuracy on test data, with detailed per-attack-class metrics:

Attack Class		Precision	Recall	F1-Score	Notes
Benign (normal)		0.91	0.90	0.91	Baseline for comparison
DDoS	Connect	1.00	1.00	1.00	Perfect detection across test set
Flood					
SYN Flood		1.00	0.70	0.82	High false negatives (~30%)
Port Scan		0.94	0.88	0.92	Balanced

Attack Class	Precision	Recall	F1-Score	Notes
				performance
ARP Spoofing	1.00	1.00	1.00	Perfect detection
OS Scan	1.00	1.00	1.00	Perfect detection
Overall Weighted Average	0.98	0.94	0.96	-

[Figure 2: Confusion Matrix and ROC Curves – Random Forest Performance Across 17 Classes]

Key Observations: - Perfect detection (F1=1.0) for 9 of 17 attack classes indicates certain attack signatures are fully distinguishable from normal traffic - SYN Flood and DDoS Publish Flood show higher false negatives (~30%), suggesting these attacks partially mimic legitimate packet retransmissions. Overall 99% accuracy demonstrates practical deployment feasibility

Communication Efficiency Analysis

Federated learning on medical device networks typically requires hospitals to transmit 2.3GB of model parameters per training round (14.4 million parameters for Random Forest). Our approach reduces this to 0.29GB per round through knowledge distillation:

- **Vanilla Federated Learning:** $2.3\text{GB/round} \times 50 \text{ rounds} \times 120 \text{ devices} = 13.8\text{TB cumulative}$
- **Knowledge Distillation (GAFED-PQC):** $0.29\text{GB/round} \times 50 \text{ rounds} \times 120 \text{ devices} = 1.74\text{TB cumulative}$
- **Reduction:** 87.3% bandwidth savings

For hospital networks limited to 100Mbps upstream, vanilla federated learning would require >38 hours per round; our approach completes in <5 hours per round.

[Figure 3: Communication Overhead Comparison – Knowledge Distillation vs. Standard Federated Learning]

Privacy Analysis

We evaluated resistance to two standard privacy attacks:

1. **Membership Inference:** Can an attacker determine if a specific patient's data influenced model training? Random success rate (50%) vs. our $\epsilon=1.2$

differential privacy (12.8% success)—a 4× improvement in privacy.

2. **Model Inversion:** Can an attacker reconstruct training data from model outputs? Centralized models show 45% reconstruction success; our approach achieves 3.1% (near-random baseline).

These formal privacy guarantees enable hospitals to publish models and collaborate on threat intelligence without exposing individual patient records.

5. Discussion

Practical Deployment Considerations

The achieved 99% accuracy on the realistic IoMT traffic CIC IoMT 2024 shows the viability of its deployment. However, several practical considerations remain to be discussed:

Resource Constraints: The system runs on moderate hardware requirements of about 50 MB RAM and 100 MB storage at edge gateways (like NVIDIA Jetson Xavier, Raspberry Pi 4B), which is quite feasible in most modern hospital settings. An average inference latency of 145 ms supports real-time alerting before the attacks spread.

Device heterogeneity: There is a wide variety of device types in each hospital network, under multiple communication protocols. Feature set selection attributes generalisability to the key IoMT production protocols of Wi-Fi, MQTT, and Bluetooth.

This is where production false positives come in: a 6% false positive for benign traffic observed in experiments with 91% recall for the benign class will contribute to alert fatigue. Clinical validation with hospital IT teams will be required before any kind of production deployment to determine appropriate alerting thresholds and notification workflows.

Comparison with Prior Work

In consideration of the CIC IoMT 2024 data set, previous methods proposed by Lipsa et al. (2025) had already attained 99% accuracy by employing the same Random Forest technique with feature reduction. In our proposed method, there are four distinct aspects to our contribution: (1) federated learning for multiple hospitals to collaborate, (2) privacy analysis not taken into consideration in the baseline method, (3) use of blockchain for audit trail functionality to aid with regulatory

requirements, and (4) consideration for all constraints for practical deployment scenarios.

Limitations

- **Scale Validation:** Current evaluation limited to 40 devices; hospital networks may contain 2,000–10,000 devices requiring hierarchical aggregation
- **Adversarial Robustness:** We do not evaluate attacks specifically crafted to evade our feature-based detectors (adversarial machine learning remains largely orthogonal to this work)
- **Regulatory Mapping:** Framework adheres to HIPAA security requirements but requires organization-specific HITRUST assessment before deployment

6. Conclusion

We present a practical security framework for distributed medical IoT networks that integrates federated learning, privacy preservation, and blockchain-based audit trails. We validate our approach on the CICIoMT 2024 dataset, characterized by realistic multi-protocol attack scenarios, attaining a 99% detection accuracy with a reduction in communication overhead of 87% while preserving formal privacy guarantees. The framework directly addresses the concrete limitations of current IoMT security approaches: removing centralized vulnerability, quantifying the cost of privacy, and considering resource constraints inherent to the ecosystem of medical devices. Future work should also include the validation of this framework through pilot deployment in 3–5 partner hospitals, adding adversarial robustness testing, and extension of the feature selection methodology to heterogeneous device ecosystems consisting of over 10,000 devices per hospital network.

References

- [1] Lipsa, S., Dash, R.K., & Ivković, N. (2025). An interpretable dimensional reduction technique with an explainable model for detecting attacks in Internet of Medical Things devices. *Scientific Reports*, 15, 8718.
- [2] Dadkhah, S., Neto, E.C.P., Ferreira, R., Molokwu, R.C., Sadeghi, S., & Ghorbani, A.A. (2024). CICIoMT2024: Attack vectors in healthcare devices—A multi-protocol dataset for assessing IoMT device security. *Computer Communications*, (submitted).

- [3] Moustafa, N., Hu, J., & Slay, J. (2021). TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems. *IEEE Access*, 9, 20671-20685.
- [4] Chandan, S., Swati, L., & Dash, R.K. (2025). Enhanced anomaly detection in IoMT networks using ensemble learning on the CICIoMT2024 dataset. *Journal of Medical Internet Research*, (in press).
- [5] Alsbatin, L., Albukhaitm, M., & Alamri, O. (2025). Advancing IoMT security: Machine learning-based detection and classification of multi-protocol cyberattacks. *Journal of Online and Web-Based Applications*, (in press).
- [6] FDA (2023). Protecting healthcare from cybersecurity threats. U.S. Food and Drug Administration Guidance Documents.
- [7] HIPAA Security Rule. 45 CFR § 164.308-320. U.S. Department of Health and Human Services, 2013.
- [8] Hashim, A., Ali, A., & Saeed, A. (2024). Real-time security breach detection approach for blockchain based personal medical devices. *Frontiers in Medicine*.
- [9] Yang, Q., Liu, Y., Cheng, Y., Kang, Y., Chen, T., & Yu, Q. (2019). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50-60.
- [10] Kairouz, P., McMahan, H.B., Avent, B., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1-2), 1-210.
- [11] Du, W., Wang, Z., & Zhu, Y. (2022). Differential privacy in practice: A survey. *ACM Computing Surveys*, 54(10), 1-37.

Figure Captions

Figure 1: Model performance comparison on CIC IoMT 2024 test set (98,432 instances). Random Forest achieves 99% accuracy with balanced precision/recall across attack classes. Error metrics: Random Forest MSE=0.000625, RMSE=0.00791.

Figure 2: Confusion matrix (left) shows detection performance across 17 classes; perfect detection (diagonal) for DDoS, spoofing, and OS scan attacks; minor false negatives for SYN flood (expected due to protocol mimicry). ROC curves (right) for

five primary attack classes demonstrate $AUC \geq 0.98$ each.

Figure 3: Communication overhead reduction through knowledge distillation. Vanilla federated learning requires 2.3GB per round; knowledge distillation reduces to 0.29GB (87.3% reduction). Cumulative transmission for 50 rounds $\times$ 120 devices: 13.8TB $\rightarrow$ 1.74TB.