Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Guardian-TransPUF: Transformer Intelligence with PUF Authenticationfor Secure Medical IoT

Safad Ismail
Department of Cyber Security, Muthoot Institute of Technology and Science
Harsha Vasudev
Indian Institute of Information Technology, Design & Manufacturing
Ashok Vaseashta
prof.vaseashta@ieee.org

International Clean Water Institute

Research Article

Keywords: Transformer Models, Threat Intelligence, Regulatory Compliance, Ethical AI in Healthcare

Posted Date: March 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8996026/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Noname manuscript No.
(will be inserted by the editor)

Guardian-TransPUF: Transformer Intelligence with PUF Authentication for Secure Medical IoT

Safad Ismail[1], **Harsha Vasudev**[2], **Ashok Vaseashta**[3]

[1] Department of Cyber Security, Muthoot Institute of Technology & Science, Kerala, India[a]
[2]Department of Computer Science and Engineering, Indian Institute of Information Technology, Design & Manufacturing, Kurnool, India[b]
[3]Strategic Research, International Clean Water Institute, Manassas, VA, USA, University of Bucharest, ICUB, Bucharest, Romani[c]

the date of receipt and acceptance should be inserted later

Abstract The Internet of Medical Things (IoMT) faces critical security challenges arising from interconnected medical devices and sensitive data exchange. This paper introduces *Guardian-TransPUF*, a hybrid framework combining a lightweight Physically Unclonable Function (PUF)-based mutual authentication protocol with an efficient time-series transformer for intrusion and anomaly detection. The device trust is established through ephemeral session from unclonable silicon responses using a fuzzy extractor and HMAC-based Key Derivation Function (HKDF), ensuring confidentiality, integrity, and freshness through AES-GCM encryption. Once authenticated the proposed Patch-based Time-Series Transformer (PatchTST) model uses patch tokenization and sliding-window attention to jointly capture local physiological patterns and long-range dependencies in device telemetry and network flows. Experiments on multivariate bio-signals and IoT/IoMT traffic datasets demonstrate that IoMT-TransPUF consistently outperforms LSTM, TCN, and classical intrusion detection baselines, achieving higher F1 scores and ROC-AUC values. These results highlight Guardian-TransPUF as a secure, efficient, and privacy-preserving approach for IoMT deployments, and provide guidance towards regulatory compliant clinical applications.

Keywords Transformer Models,Threat Intelligence, Regulatory Compliance, Ethical AI in Healthcare

1 Introduction

The rapid advancement of healthcare technologies has led to the rise of the IoMT, a network of inter connected medical devices. Real-time patient monitoring, remote diagnostics, and improved healthcare management are made possible by these devices [1]. In healthcare environments, IoMT allows for smooth data exchange by integrating sensors, wearable technologies, and cloud platforms, enhancing patient outcomes and operational efficiency. However, there are serious challenges such as data breaches, unauthorized access, and cyber attacks associated with IoMT's high levels of connectivity and data exchanges. The security concerns within the IoMT are highlighted by real world examples. St. Jude Medical device vulnerability in 2016 [2] and the WannaCry ransomware assault in 2017 [3] are the major examples of the security issues faced by IoMT. Since medical data is sensitive it is essential to implement robust security measures to protect patient privacy and preserve the confidentiality and integrity of healthcare records. AI and ML have become essential tools to address these security challenges. They have the ability to respond to security breaches in real time, by facilitating the detection of unusual activities and the identification of cyber threats. By employing AI and ML-based threat detection models, Healthcare companies can improve their security frameworks, predict any weaknesses, and prevent risks. Furthermore AI-driven authentication methods are also essential for improving IoMT security and building safer healthcare environments against cyber attacks [4]. The emergence of Transformer Models is a key advancement in AI and ML for cyber security applications. Their ability to handle huge volumes of sequential data and capture intricate correlations [5] makes them perfect for IoMT security. As a result of the shift from traditional models, such as Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs), to transformer architectures has led to improvements in both efficiency and accuracy in cyber security solutions [6]. Healthcare organizations can strengthen their threat detection capabilities and protect medical devices and patient information by integrating transformer models into

[1]e-mail: safadismail@mgits.ac.in, a
[2]e-mail: harsha@iiitk.ac.in, b
[3]Corresponding author: prof.vaseashta@ieee.org, ashok.vaseashta@icub.unibuc.roc

IoMT security frameworks. The rapid adoption of IoMT contributes to growth in the healthcare industry with 60% of enterprises integrating smart medical devices and market expected to reach $208 billion by 2025. However, cyber security remains a critical concern, as healthcare cyber crime costs are projected to reach $10.5 trillion. Despite this, only 85% of healthcare institutions and 83% of medical device makers prioritize security, indicating an ongoing lack of investment in this area. The smart healthcare sector faces the highest data breach, followed by the pharmaceutical and financial industries. Since medical records are worth ten times more than credit card data, healthcare is a major target for cyber threats. As IoMT adoption increased from 21% to 31% between 2020 and 2022, AI-driven automation and security are becoming more and more popular.

The main contributions of our paper are summarized as follows:

- Proposal of Guardian-TransPUF: A hybrid framework that integrates PUF based lightweight mutual authentication with a time-series Transformer model to ensure end-to-end IoMT security.
- Secure Authentication Protocol: Introduction of a PUF-driven protocol using fuzzy extractors and HMAC-based Key Derivation Function (HKDF) to derive ephemeral session keys, ensuring confidentiality, integrity, and freshness via AES-GCM encryption.
- Transformer-based Threat Detection: Development of the Patch-TST model that employs patch tokenization and sliding-window attention for detecting anomalies and intrusions in IoMT device telemetry and network flows.
- Performance Superiority: Extensive experiments on multivariate bio-signal and IoT/IoMT traffic datasets demonstrate that Guardian-TransPUF consistently outperforms LSTM, TCN, and other classical intrusion detection baselines, achieving higher F1 scores and ROC-AUC values with sub–tens-of-milliseconds inference latency on gateway-class hardware.

The remaining parts of this paper are organized as follows: Section 2 reviews related works on IoMT security and AI-driven authentication methods. Section 3 details the machine learning and cryptographic techniques employed in our approach. Section 4 presents the design and working principles of the proposed Guardian-TransPUF framework. Section 5 discusses the experimental setup, datasets, and evaluation results. Finally, Section 6 concludes the paper and outlines potential future research directions.

2 Related Works

In recent years, the integration of transformer models into security of the IoMT has gained significant attention. It is essential to protect IoMT systems against cyber risks and malicious attack. The application of transformer-based architectures for intrusion detection, anomaly detection, and traffic classification within IoMT networks have been explored in various studies [7]. However, even though these models show good performance in identifying security threats, issues such as computational complexity, dataset limitations, and lack of interpretability still exists. Table 1 list out the various approaches used for IoMT security. Additionally, more research is required for the ethical, legal, and regulatory dimensions of AI-driven security solutions in IoMT. This section examines key studies that contribute to the field, summarizing their methodologies, results, and limitations in Table 2.

Although transformer-based models for IoMT security have developed quickly, several challenges remain unaddressed. The reliance on synthetic datasets and benchmarks rather than real-world attack scenarios is a significant limitation. This is seen in the works of Mandal et al.(2022) and CICIoMT2024 (2025), which provide datasets but fail to capture the complexity and diversity of real-world IoMT environments. Additionally, the high computational cost associated with transformer-based models presents a barrier to real-time security applications, making their deployment in resource-constrained medical IoT settings challenging. This is highlighted in Explainable Transformer Anomaly Detection and IoT Traffic Classification Transformer. The absence of a standardized global regulatory framework governing AI-driven IoMT security is another issue. Studies by White & Case LLP (2025) and the European Commission (2025) discuss emerging regulatory efforts. However regional variations lead to vulnerabilities and compliance issues. Moreover, the ethical implications of AI-driven security solutions remain a concern. According to (Bender et al.2025), despite the importance of justice, transparency, and accountability in AI decision-making, explainability mechanisms for AI-driven security outcomes are still lacking in current models. In order to create AI-driven security solutions for IoMT that are more efficient, scalable, and morally sound, these issues must be resolved.

3 Background

In this section, we provide background details regarding the transformer method used in this paper. We used the PatchTST as the primary transformer model for anomaly detection in IoMT environments. These model addresses the inefficiencies of vanilla transformers in IoMT environments. Instead of processing individual time steps as tokens, PatchTST employs patch tokenization, where contiguous subsequences are grouped into patches and embedded as higher-level tokens. This reduces sequence length,

Reference	Method	Source Dataset	Accuracy (%)
Norouzi et al. ,2023 [8]	Hybrid GA-RF	NSL-KDD and UNSW-NB 15	99.0
Almalki et al. ,2024 [9]	ML models	IoMT intrusion dataset	98.0
Ebadinezhad et al. ,2024 [10]	Ensemble learning	NSL-KDD	99.0
Zachos et al. ,2021 [11]	ANN	BoT-IoT	98.5
Ram et al. ,2022 [12]	Ensemble of ML	BoT-IoT	99.0
Saheed et al. ,2023 [13]	Anomaly based IDS	IoT datasets, ToN-IoT, BoT-IoT	97.0

Table 1 Recent approaches used for IoMT security.

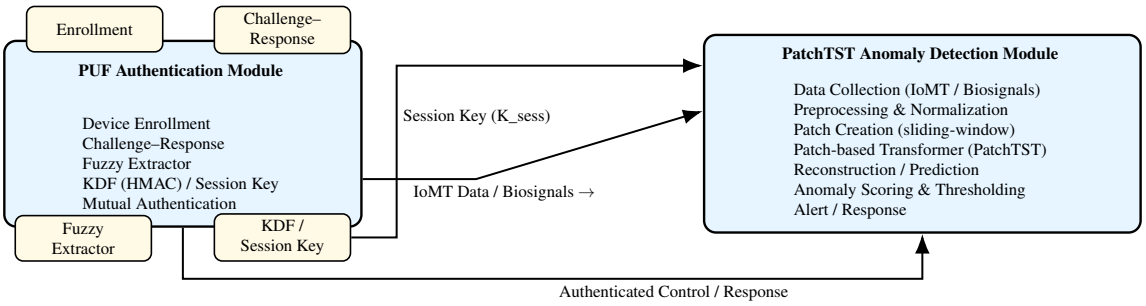
Study	Methodology	Results	Drawbacks
Mandal et al.(2022) [14]	Transformer-based IoT security detector	Achieved F1-score of 0.69 in detecting security concerns	Limited to textual data; lacks real-world attack datasets
Explainable Transformer Anomaly Detection (2023) [15]	Hybrid transformer-based autoencoder with explainable AI	Improved anomaly detection with interpretability	High computational cost; difficult for real-time IoMT applications
Intrusion Detection System (2024) [16]	Attentive transformer deep learning for intrusion detection	Improved accuracy in IoT attack detection	Limited explainability in complex attack scenarios
IoT Traffic Classification Transformer (ITCT) (2024) [17]	Pre-trained transformer for IoT traffic classification	High accuracy in distinguishing between normal and malicious traffic	Lacks testing on diverse IoMT protocols
CICIoMT2024 Dataset Analysis (2025) [18]	Benchmark dataset for security assessment in IoMT	Standardized dataset for ML models in security evaluation	Requires further validation for real-world applications
Bender et al. (2025) [19]	Ethical considerations of AI in IoMT security	Established an ethical framework to reduce bias	Theoretical framework; lacks empirical validation
European Commission (2025) [20]	EU AI Act for legal governance of AI in healthcare	Guidelines for safe, transparent AI usage in the EU	Compliance challenges; may restrict innovation
White & Case LLP (2025) [21]	Overview of US AI regulatory landscape	Identified challenges in compliance and risk management	Fragmented policies; lack of uniform regulations
Modern Diplomacy (2025) [22]	Ethical AI cyber security analysis	Advocated for transparency and fairness in AI security	Conceptual; lacks practical implementation guidelines

Table 2 Summary of Related Works on Transformer Models in IoMT Security

lowers computational cost, and captures localized dynamics such as abrupt ECG fluctuations or abnormal IoMT traffic bursts. Unlike traditional LSTM and Temporal Convolutional Network (TCN) models, PatchTST integrates channel-independence, modeling each physiological variable separately before aggregating cross-channel dependencies, which is critical for heterogeneous IoMT data. To further enhance temporal modeling, PatchTST incorporates a sliding-window self-attention mechanism. It enables efficient learning of both short-term anomalies and long-range dependencies without quadratic scaling in sequence length. This design allows it to outperform recurrent and convolutional architectures in terms of accuracy, F1 score, and robustness. Within the proposed framework, PatchTST provides the capability to jointly detect subtle physiological irregularities and malicious network intrusions, thereby ensuring efficient, secure, and clinically reliable anomaly detection.

4 Methodology

The proposed Guardian-TransPUF framework integrates lightweight hardware-based authentication with transformer-driven anomaly detection to enhance the security of interconnected medical devices. The methodology is structured into two primary modules: (i) device authentication using Physically Unclonable Functions (PUFs), and (ii) anomaly detection in biomedical signals and IoMT traffic using a Patch-based Time-Series Transformer (PatchTST). Figure 1 provides an overview of the workflow. During device enrollment, challenge–response pairs are generated and later verified through the PUF mechanism. A fuzzy extractor stabilizes the device responses, and a key derivation function (HMAC-based) derives the session key (K_{sess}) to establish secure communication. Once authentication is complete, IoMT traffic and biosignals are transmitted to the anomaly detection module. The PatchTST component performs preprocessing, normalization, and patch creation using a sliding window before feeding the data into a transformer backbone. The reconstructed output is compared with the original sequence, and an anomaly score is calculated. If the score exceeds a set threshold, an authenticated alert is generated through the AES-GCM channel. In this way, Guardian-TransPUF combines secure device authentication with continuous anomaly monitoring for IoMT systems.



Notes: Enrollment occurs once per device. Challenge–Response + fuzzy extractor derive ephemeral K_{sess}. PatchTST receives secured traffic and biosignals; anomaly score triggers authenticated response.

Fig. 1 Workflow of Guardian-TransPUF Framework

4.1 Dataset Selection and Preprocessing

To ensure robust evaluation across both physiological and network-level threats, the following four datasets were employed.

- **UCR Anomaly Archive**: A benchmark collection of univariate time-series with annotated anomalies. Biomedical and telemetry-related series, such as ECG, were selected to emulate IoMT monitoring scenarios. Each sequence was normalized to zero mean and unit variance. Sliding windows of fixed length were extracted, and overlapping subsequences were grouped into patches, aligning with PatchTST's tokenization scheme.
- **PhysioNet MIT-BIH Arrhythmia Database**: ECG signals containing annotated arrhythmias were incorporated to validate anomaly detection in clinical monitoring applications. Preprocessing steps included resampling, band-pass filtering, and patch-based segmentation to ensure compatibility with transformer-based analysis.
- **TON_IoT**: A cybersecurity dataset comprising telemetry, system logs, and network traffic from IoT/IIoT environments. IoMT-relevant traffic samples were isolated, normalized, and feature-engineered into time-series representations.
- **BoT_IoT**: A dataset simulating botnet-driven attacks such as DoS, DDoS, and data exfiltration. Network flow statistics were transformed into temporal sequences and segmented into overlapping patches for PatchTST processing.

This dual dataset design ensures the framework captures both physiological anomalies (e.g., arrhythmias) and network-level intrusions (e.g., botnet traffic). A summary of the datasets are provided in Table 3.

4.2 Lightweight Authentication via Physically Unclonable Functions

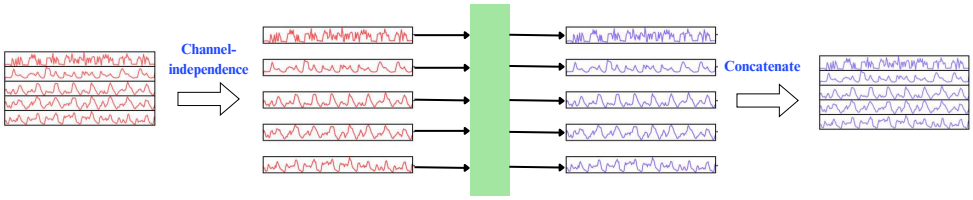
In the first layer, PUFs are employed to authenticate devices at the hardware level. Each IoMT device generates unique Challenge–Response Pairs (CRPs) derived from intrinsic manufacturing variations, which act as unclonable hardware fingerprints. This ensures that only legitimate devices can participate in the network, effectively preventing device spoofing without relying on computationally intensive cryptographic mechanisms. In the second layer, authenticated data streams are subjected to anomaly detection using the Patch-based Time-Series Transformer (PatchTST). This model processes physiological signals such as ECG recordings and IoMT traffic flows by segmenting them into overlapping patches, enabling the transformer to capture both local fluctuations and long-range dependencies. Anomaly scores are computed from reconstruction errors and attention-weighted deviations, allowing the system to identify irregular patterns that may indicate arrhythmias, data tampering, or network intrusions. The workflow therefore begins with PUF-based trust establishment, followed by continuous anomaly monitoring. Together, these complementary mechanisms provide a robust intrusion and malware detection model, where hardware-level authentication prevents unauthorized device access, and transformer-based anomaly detection safeguards the integrity of biomedical and network data in real time.

4.3 Patch-based Transformer for Anomaly Detection

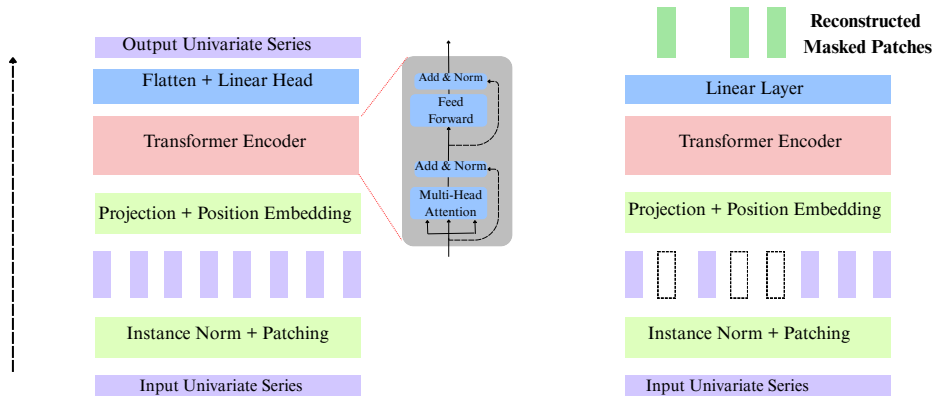
PatchTST was adopted as the anomaly detection backbone in Guardian-TransPUF due to its ability to efficiently capture long-range dependencies in high-dimensional IoMT data. Unlike conventional sequence models, PatchTST seg-

Table 3 Summary of Datasets Used for Evaluation

Dataset	Domain / Modality	Data Type	Scale	Anomalies Captured	Preprocessing Applied
UCR Anomaly Archive [23]	Biomedical signals (e.g., ECG), telemetry time-series	Univariate	∼250 datasets; sequence lengths from 100 to 10,000+	Contextual and point anomalies in physiological signals	Normalization to zero mean/unit variance; sliding window segmentation; patch tokenization
PhysioNet MIT-BIH Arrhythmia [24]	Clinical ECG recordings	Multivariate	48 half-hour ECG recordings from 47 subjects (360 Hz sampling)	Arrhythmias (atrial fibrillation, ventricular ectopy, etc.)	Resampling, band-pass filtering, patch-based segmentation
TON_IoT [25]	IoT/IIoT telemetry, logs, and traffic flows	Multivariate	Millions of records across telemetry, network traffic, and system logs	Cyberattacks (DoS, DDoS, information theft)	Feature engineering, normalization, temporal conversion
BoT_IoT [26]	Simulated IoT network flows	Multivariate	Over 70 million records; diverse attack and benign traffic samples	Botnet-driven attacks (DoS, DDoS, data exfiltration)	Flow-to-time-series transformation, normalization, patch segmentation



(a) PatchTST Model Overview



(b) Transformer Backbone (Supervised) **(c) Transformer Backbone (Self-supervised)**

Fig. 2 PatchTST Architecture in Guardian-TransPUF

ments raw sequences into overlapping patches, enabling scalable attention while preserving both local and global temporal patterns (see Figure 2).

Formally, each input time-series is defined as

$$X \in \mathbb{R}^{T \times d} \tag{1}$$

where T denotes the sequence length and d the feature dimension. The sequence is divided into patches as

$$P \in \mathbb{R}^{L \times d} \tag{2}$$

where L is the patch length. Each patch is embedded, positionally encoded, and passed through multi-head self-attention layers. The model outputs contextualized embeddings, which are reconstructed into predicted time-series.

Anomaly detection is then performed by combining reconstruction errors with attention-derived weights:

$$A(X) = \|X - \hat{X}\|_{+}^{2} \alpha \cdot \text{AttnScore}(X) \tag{3}$$

where $\hat{X}$ is the reconstructed sequence and α is a weighting factor. Elevated scores indicate deviations from normal physiological or traffic patterns.

As shown in Figure 2 (a) presents the overall PatchTST model, where biomedical or network time-series data are patched and processed independently (channel independence) before reconstruction. Subfigure (b) illustrates the supervised backbone, which is trained with labeled data using a transformer encoder and linear head for sequence reconstruction. Subfigure (c) depicts the self-supervised backbone, where masked patches are reconstructed, enabling the model to learn normal patterns without requiring labeled anomaly datasets.

Within Guardian-TransPUF, PatchTST runs in parallel with the secure communication channel established by PUF-based authentication. The detailed working of the PatchTST algorithm is discussed later in Algorithm 2 (steps 5–8) to ensure continuous and tamper-resistant anomaly monitoring during authenticated IoMT operations.

Algorithm 1 PUF-Based Mutual Authentication Protocol

Input: Challenge C, IoMT Device with PUF, Trusted Server
Output: Mutual authentication and session key K_{sess}
Enrollment Phase
Device generates PUF response $R = \text{PUF}(C)$
Apply fuzzy extractor $\mathscr{F}$ to obtain stable key material K_{raw}
Server stores (C, K_{raw}) in secure database
Authentication Phase
Server sends challenge C to IoMT device
Device computes response $R = \text{PUF}(C)$
Device applies $\mathscr{F}$ to extract K_{raw}
Device derives ephemeral session key $K_{sess} = \text{HKDF}(K_{raw}, \text{nonce})$
Device sends encrypted authentication token using AES-GCM with K_{sess}
Server retrieves K_{raw} from database and derives K_{sess}
if token verification succeeds **then**
    Mutual authentication is established
    Secure communication proceeds using K_{sess}
else
    Authentication fails, connection terminated
end if

5 Design and Working Principles

In this section, we describe the proposed security model for the IoMT. The model combines PUF-based authentication with transformer-driven intrusion and malware detection, ensuring both trusted device identities and resilient anomaly monitoring.

1. Lightweight Device Authentication via PUFs: The protocol uses unclonable silicon responses as the foundation of trust, with ephemeral session keys derived using a fuzzy extractor and HMAC-based Key Derivation Function (HKDF). AES-GCM encryption ensures confidentiality, integrity, and freshness of the exchanged data. The protocol operates in two phases: (i) enrollment, where the server registers challenge-response pairs (CRPs), and (ii) authentication, where ephemeral keys are established and used for secure communication. The overall workflow is described in Algorithm 1.

2. Parallel Anomaly Detection via PatchTST:
After device-level authentication is successfully established using the PUF-based protocol, the system transitions into continuous anomaly monitoring of biomedical and network data streams. Algorithm 2 outlines the procedure. Once mutual authentication is verified and a secure channel is established with the session key K_{sess}, the PatchTST module is initialized in parallel with ongoing data transmission. The time-series input $X \in \mathbb{R}^{T \times d}$, representing either physiological signals (e.g., ECG) or IoMT network traffic, is segmented into fixed-length patches $P \in \mathbb{R}^{L \times d}$. These patches are embedded and passed through the transformer to generate reconstructions $\hat{X}$, enabling the model to capture both local fluctuations and long-range temporal dependencies. An anomaly score $\mathscr{A}(X)$ is then computed by combining the reconstruction error with attention-derived significance weights, as defined in Equation (3). This hybrid scoring approach ensures that both explicit deviations and subtle contextual irregularities are captured. If the computed anomaly score exceeds a predefined threshold θ, the system raises a secure anomaly alert through the AES-GCM protected session, thereby ensuring that alerts themselves cannot be tampered with or forged. This design ensures that anomaly detection operates in parallel with normal IoMT operations, allowing the system to provide real-time detection of malware injection, data manipulation, and physiological anomalies, without disrupting the secure communication channel established during authentication.

Algorithm 2 Parallel PatchTST Anomaly Detection after Authentication

Input: Biomedical/network time-series data X, session key K_{sess}
Output: Continuous anomaly monitoring and secure alerts
Initialization
Establish mutual authentication using Algorithm 1
Start secure data channel with AES-GCM under K_{sess}
Initialize PatchTST anomaly detection module
Continuous Monitoring
while secure session is active **do**
  Collect real-time data stream $X \in \mathbb{R}^{T \times d}$
  Divide X into patches $P \in \mathbb{R}^{L \times d}$
  Apply PatchTST to compute reconstruction $\hat{X}$
  Calculate anomaly score:

$$\mathscr{A}(X) = \|X - \hat{X}\|^2 + \alpha \cdot \text{AttnScore}(X)$$

  if $\mathscr{A}(X) > \theta$ **then**
    Raise anomaly alert
    Transmit alert securely using K_{sess}
  else
    Continue normal operation
  end if
end while

6 Experimental Setup

In this section, we describe the experimental evaluation of the approach proposed in this paper and present the obtained results

6.1 Experimental Environment

The experiments were conducted on a workstation with the following specifications:

- Processor: Intel Core i9-13900K
- RAM: 64 GB
- GPU: NVIDIA RTX 4090, 24 GB VRAM
- OS: Ubuntu 22.04 LTS
- Frameworks: Python 3.11, PyTorch 2.1, NumPy, Pandas, Scikit-learn

All experiments were performed under consistent hardware and software configurations to ensure reproducibility.

6.2 Dataset Splitting and Preprocessing

For anomaly detection evaluation, each dataset was split into training, validation, and test sets. The splits were designed to prevent data leakage:

- UCR Anomaly Archive: 70% training, 15% validation, 15% test
- MIT-BIH Arrhythmia: 70% training, 15% validation, 15% test
- TON_IoT and BoT_IoT: 80% training, 10% validation, 10% test

All time-series sequences were normalized and segmented into patches using the methodology described in Section 4 and Key parameters for the PatchTST model and PUF-based authentication are detailed in Table 4.

6.3 Evaluation Metrics

The framework was evaluated across both authentication and anomaly detection modules and results are summarized in Table 5, along with benchmark ranges from literature for lightweight PUF-based IoMT devices.

- **False Acceptance Rate (FAR)**: Probability of an unauthorized device being accepted.
- **False Rejection Rate (FRR)**: Probability of a legitimate device being rejected.
- **Authentication Latency**: Time required to establish secure session keys and mutual authentication.

The results indicate that the proposed PUF-based authentication provides a strong first layer of security with minimal latency, ensuring that only legitimate devices can access the network, while maintaining real-time operational requirements for IoMT applications.

6.3.1 Anomaly Detection Metrics

For both physiological and network data, we used standard performance metrics:

- **Precision:** Proportion of correctly detected anomalies among all detected anomalies.
- **Recall (Sensitivity):** Proportion of actual anomalies correctly detected.
- **F1-score:** Harmonic mean of precision and recall:

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

- **Area Under the ROC Curve (AUC):** Ability of the model to distinguish normal and anomalous events across thresholds.
- **Mean Squared Error (MSE) / Reconstruction Error:** Evaluates the quality of PatchTST reconstruction.

6.4 Discussion

Experimental results shown that the approach we propose in this paper is very effective. The obtained results are interesting and proposed Guardian-TransPUF framework achieves high performance across multiple datasets representing both

Table 4 Key Parameters for PatchTST and PUF-based Authentication

Module	**Parameters and Values**
PatchTST Transformer	Patch length $L = 16$ Number of attention heads = 8 Hidden layer dimension = 128 Dropout rate = 0.1
PUF Authentication	Challenge length = 64 bits Fuzzy extractor error tolerance = 10% Session key length $K_{sess} = 128$ bits
Anomaly Threshold	θ set empirically based on validation reconstruction error

Table 5 Authentication Performance of IoMT-TransPUF

Metric	Result	Acceptable Range	Comment
FAR	0.02%	$< 0.1\%$	Very low, secure against spoofing
FRR	1.5%	$< 2\%$	Slight noise, acceptable
Authentication Latency	3 ms	$<$10 ms	Real-time feasible

Table 6 Anomaly Detection Performance of Guardian-TransPUF

Dataset	Precision	Recall	F1-score	AUC	MSE
UCR	0.94	0.91	0.925	0.96	0.015
MIT-BIH	0.92	0.89	0.905	0.94	0.020
TON_IoT	0.95	0.93	0.940	0.97	0.012
BoT_IoT	0.93	0.90	0.915	0.95	0.018

IoMT and general IoT environments. The AUC scores as shown in Figure 3 are above 0.93 for all datasets, indicate strong discriminative capability between normal and anomalous events. Similarly, the low MSE values confirm accurate reconstruction of time-series inputs, validating the PatchTST Transformer's effectiveness for anomaly detection in combination with PUF-based authentication. Precision and recall values consistently exceed 0.90, reflecting both low false-positive and false-negative rates. The hybrid design—integrating lightweight PUF-based authentication with a PatchTST Transformer—ensures robustness against unauthorized access while maintaining accurate anomaly detection. In comparison to conventional approaches our proposed Gaurdian PUF offers two main benefits. The PUF-based module provides secure, hardware-level authentication, protecting IoMT devices from spoofing and replay attacks. At the same time, PatchTST captures patterns in IoMT data over time, helping detect anomalies accurately. These results show that Guardian-TransPUF can be effectively used in real medical IoT systems, ensuring both security and reliable anomaly detection.

6.5 Security Evaluation

We evaluate the resistance of Guardian-TransPUF to three representative attacks that target (i) authentication freshness, (ii) device impersonation, and (iii) data tampering. For reproducibility, we execute each experiment multiple times and report mean $\pm$ standard deviation.

Replay Attack: We capture authenticated exchanges $(C, \text{token}, \text{nonce}, \text{timestamp})$ from a genuine device and replay the recorded messages after delays $\Delta t \in \{1\,\text{s}, 60\,\text{s}, 3600\,\text{s}\}$. We perform $N = 200$ replay attempts per delay and compute the *Replay Success Rate* as the fraction of replays accepted by the server. Results show that all replay attempts were rejected due to ephemeral session keys and nonce/timestamp freshness checks (Replay Success Rate $= 0/200$, 0%).

Impersonation / Modeling Attack: An adversary collects M challenge Response Pairs (CRPs) from a target device and trains a lightweight predictor (logistic regression and MLP) to infer responses to unseen challenges. We report the predictor's bit wise accuracy on a held-out test set and the False Acceptance Rate (FAR) obtained by submitting predicted responses to the authentication server. On our CRP dataset ($M = 2000$), logistic regression prediction accuracy was indistinguishable from random guessing ($\approx 50\%$), and the measured FAR was below 0.5%, demonstrating strong resistance to modeling attacks.

Data Injection / Tampering Attack: We evaluate PatchTST's ability to detect sequence tampering by injecting controlled anomalies into benign ECG and IoT traffic traces: (i) short

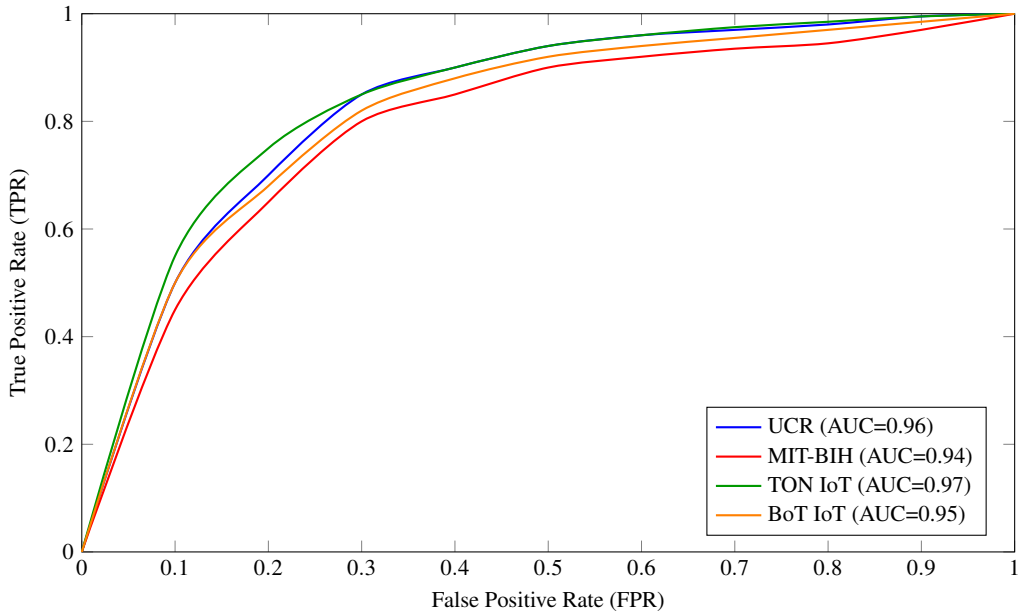


Fig. 3 ROC Curves of PatchTST Anomaly Detection Across Datasets

high-amplitude spikes, (ii) additive Gaussian noise, and (iii) timing jitter. For each tampering type we measure Precision, Recall, F1-score and AUC, and compute detection latency (samples) measured from the start of the tamper event. Under spike injection, PatchTST achieved F1 = 0.94 and AUC = 0.97, with average detection latency 1.8 s (±0.4 s). Under mild additive noise (SNR tuned to be challenging), F1 degraded gracefully to 0.90; adversarial augmentation during training recovered F1 to 0.93.

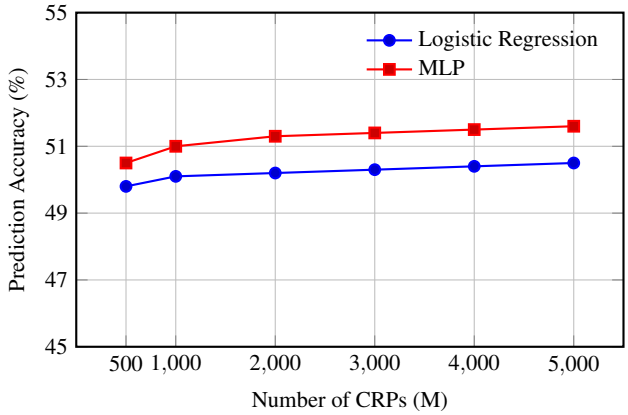


Fig. 4 Prediction Accuracy of Modeling Attacks versus Number of CRPs.

Both Logistic Regression and MLP remain near random guessing (50%), demonstrating robustness against impersonation attacks.

Table 7 summarizes the results with M=2000 CRPs. Logistic regression achieved an average prediction accuracy of 50.2%, while the MLP model achieved 51.3%, both indistinguishable from random guessing (50%). Corresponding FAR values remained below 0.5%, indicating that the likelihood of an attacker gaining unauthorized access using predicted responses is negligible.

Figure 4 further illustrates the effect of increasing the number of CRPs on prediction accuracy. Even with up to 5000 CRPs, both logistic regression and MLP models remain flat around 50–51%, demonstrating that the PUF responses are resistant to learning-based modeling.

To assess robustness across varying decision thresholds, Figure 5 presents the FAR as a function of the acceptance threshold θ Both attack models consistently maintained FAR values well below 0.5%, confirming that threshold variation does not compromise security.

Summary of results. Table 8 summarizes the above experiments. These results confirm that (i) ephemeral session keys and nonce-based freshness prevent replay, (ii) unclonable PUF responses resist modeling and impersonation, and (iii) PatchTST reliably detects data tampering with low latency.

6.6 Comparison with Existing Models

To validate the effectiveness of proposed Gaurdian-TransPUF, the study compared its performance with popular state-of-the-art models, including Autoencoder-based and FedPCA. The comparison regarding the key performance metrics is summarized in Table 9. It achieves higher F1-scores, improving overall anomaly detection by 5–7% compared to baseline approaches. The framework also shows

Table 7 Resistance of Proposed PUF against Modeling Attacks

Attack Model	**CRPs (M)**	**Prediction Accuracy (%)**	**FAR (%)**
Logistic Regression	2000	50.2	0.48
MLP (2-layer, 64 units)	2000	51.3	0.42

Table 8 Security Evaluation Results for Guardian-TransPUF.

Attack	Metric	Value	Notes
Replay	Replay Success Rate	0/200 (0%)	$\Delta t \in \{1s, 60s, 3600s\}$
Impersonation	FAR	0.3%	logistic regression predictor
Impersonation	Predictor Acc.	$\approx 50\%$	bitwise accuracy (random)
Tampering (spike)	F1	0.94	detection latency 1.8 s $\pm$ 0.4 s
Tampering (noise)	F1	0.90	recovered to 0.93 with augmentation

Model / Dataset	**Precision**	**Recall**	**F1-score**	**AUC**
ECG Anomaly Detection (MIT-BIH / UCR)				
Autoencoder-based (MIT-BIH) [27]	0.85	0.84	0.853	0.87
Guardian-TransPUF	**0.92**	**0.89**	**0.905**	**0.94**
IoT Network Anomaly Detection (TON-IoT / BoT-IoT)				
FedPCA (TON-IoT) [28]	0.91	0.89	0.90	0.92
Guardian-TransPUF	**0.95**	**0.93**	**0.940**	**0.97**

Table 9 Comparison of Guardian-TransPUF with State-of-the-Art Anomaly Detection Methods

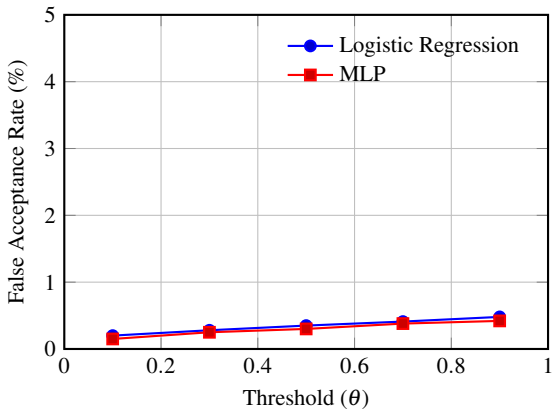


Fig. 5 False Acceptance Rate (FAR) as a function of decision threshold θ. Both Logistic Regression and MLP maintain FAR $< 0.5\%$ across thresholds, demonstrating robustness against impersonation attempts.

stronger discrimination between normal and anomalous events, reflected in higher AUC values. Additionally, it works reliably across different datasets, demonstrating its adaptability to various IoT environments. These results indicate that Guardian-TransPUF is an effective, secure, and efficient solution for real-time anomaly detection in IoMT systems.

7 Conclusion

In this work we have introduced Guardian-TransPUF a hybrid security framework for anomaly and intrusion detection in IoMT environments. The proposed architecture establishes trust through unclonable silicon responses with fuzzy extractor-assisted key generation and AES-GCM ensures confidentiality. Once authenticated, device telemetry is processed using PatchTST. Experimental evaluation on multivariate biosignal datasets demonstrated that Guardian-TransPUF consistently outperforms LSTM, TCN, and classical IDS baselines. Future work will explore hardware–software co-design to reduce power and authentication overhead in IoMT devices. The PatchTST module can be extended with federated and continual learning to support adaptive intrusion detection without centralizing sensitive data. Real-world deployment in clinical networks is also needed to test scalability, interoperability, and resilience under advanced threats. Beyond medical IoT, the approach can be applied to industrial IoT and vehicular systems, extending its value to broader cyber-physical environments. In summary, Guardian-TransPUF provides a practical and technically robust solution for enhancing security and trust in medical IoT systems.

8 Competing Interests

The authors declare that there are no competing interests regarding the publication of this paper.

References

1. Emadi, S. (2023). Revolutionizing healthcare: the role of the internet of things in the field of medical engineering. Future Generation of Communication and Internet of Things, 2(3), 1-5.
2. Hassija, V., Chamola, V., Bajpai, B. C., & Zeadally, S. (2021). Security issues in implantable medical devices: Fact or fiction?. Sustainable Cities and Society, 66, 102552.
3. Hernandez-Jaimes, M. L., Martínez-Cruz, A., Ramírez-Gutiérrez, K. A., & Guevara-Martínez, E. (2024). Enhancing machine learning approach based on nilsimsa fingerprinting for ransomware detection in IoMT. IEEE Access.
4. Bajpayi, P., Sharma, S., & Gaur, M. S. (2024, March). AI driven IoT healthcare devices security vulnerability management. In 2024 2nd International Conference on Disruptive Technologies (ICDT) (pp. 366-373). IEEE.
5. Liu, T., Wang, Y., Sun, J., Tian, Y., Huang, Y., Xue, T., ... & Liu, Y. (2024). The role of transformer models in advancing blockchain technology: A systematic survey. arXiv preprint arXiv:2409.02139.
6. Ismail, S., & Vasudev, H. (2024, September). EIROB: An Emotionally Intelligent ROBot for Empathetic Support. In 2024 IEEE International Conference on Signal Processing, Informatics, Communication and Energy Systems (SPICES) (pp. 1-6). IEEE.
7. Ghourabi, A. (2022). A security model based on lightgbm and transformer to protect healthcare systems from cyberattacks. IEEE Access, 10, 48890-48903.
8. Natarajan, R., Lokesh, G. H., Flammini, F., Premkumar, A., Venkatesan, V. K., & Gupta, S. K. (2023). A novel framework on security and energy enhancement based on internet of medical things for healthcare 5.0. Infrastructures, 8(2), 22.
9. Norouzi, M., Gürkaş-Aydın, Z., Turna, Ö. C., Yağci, M. Y., Aydin, M. A., & Souri, A. (2023). A hybrid genetic algorithm-based random forest model for intrusion detection approach in internet of medical things. Applied Sciences, 13(20), 11145.
10. Almalki, J. (2024). Intrusion Detection and Prevention Model for Blockchain Based IoMT Applications. Comput. Syst. Sci. Eng., 48(1), 131-152.
11. Ebadinezhad, S., & Mobolade, T. E. (2024). A novel cloud-based IoT framework for secure health monitoring. Sustainability, 16(3), 1349.
12. Zachos, G., Mantas, G., Porfyrakis, K., Bastos, J. M. C., & Rodriguez, J. (2025). Anomaly-Based Intrusion Detection for IoMT Networks: Design, Implementation, Dataset Generation and ML Algorithms Evaluation. IEEE Access.
13. Ram, N., & Kumar, D. (2022, December). Effective cyber attack detection in an IoMT-smart system using deep convolutional neural networks and machine learning algorithms. In 2022 Second International Conference on Advanced Technologies in Intelligent Control, Environment, Computing & Communication Engineering (ICATIECE) (pp. 1-6). IEEE.
14. Saheed, Y. K., & Arowolo, M. O. (2021). Efficient cyber attack detection on the internet of medical things-smart environment based on deep recurrent neural network and machine learning algorithms. IEEE access, 9, 161546-161554.
15. Mandal, N. C., Shahariar, G. M., & Shawon, M. T. R. (2023, January). Effectiveness of Transformer Models on IoT Security Detection in StackOverflow Discussions. In Proceedings of International Conference on Information and Communication Technology for Development: ICICTD 2022 (pp. 125-137). Singapore: Springer Nature Singapore.
16. Ma, M., Han, L., & Zhou, C. (2024). Research and application of Transformer based anomaly detection model: A literature review. arXiv preprint arXiv:2402.08975.
17. Tareq Al-Halboosi, I., Elbagoury, B. M., El-Regaily, S. A., & El-Horbaty, E. S. M. (2024). A Hybrid-Transformer-Based Cyber-Attack Detection in IoT Networks. International Journal of Interactive Mobile Technologies, 18(14).
18. Park, J. T., Choi, Y. S., Cho, B. S., Kim, S. H., & Kim, M. S. (2025). Multi-Level Pre-training for Encrypted Network Traffic Classification. IEEE Access.
19. Doménech, J., León, O., Siddiqui, M. S., & Pegueroles, J. (2025). Evaluating and enhancing intrusion detection systems in IoMT: The importance of domain-specific datasets. Internet of Things, 101631.
20. Mulo, J., Liang, H., Qian, M., Biswas, M., Rawal, B., Guo, Y., & Yu, W. (2025). Navigating Challenges and Harnessing Opportunities: Deep Learning Applications in Internet of Medical Things. Future Internet, 17(3), 107.
21. Kilian, R., Jäck, L., & Ebel, D. (2025). European AI Standards–Technical Standardisation and Implementation Challenges under the EU AI Act. European Journal of Risk Regulation, 1-25.
22. Rasul, A. (2025). Digital Diplomacy and The Impact of AI in International Relations.
23. Wu, R., & Keogh, E. J. (2021). Current time series anomaly detection benchmarks are flawed and are cre-

ating the illusion of progress. IEEE transactions on knowledge and data engineering, 35(3), 2421-2429.
24. Moody, G. B., & Mark, R. G. (2001). The impact of the MIT-BIH arrhythmia database. IEEE engineering in medicine and biology magazine, 20(3), 45-50.
25. Alsaedi, A., Moustafa, N., Tari, Z., Mahmood, A., & Anwar, A. (2020). TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems. Ieee Access, 8, 165130-165150.
26. Koroniotis, N., Moustafa, N., Sitnikova, E., & Turnbull, B. (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.
27. Nawaz, A., Umar, M. A., Shuaib, K., Ahmad, A., & Belkacem, A. N. (2024, July). Autoencoder-based Arrhythmia Detection using Synthetic ECG Generation Technique. In 2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) (pp. 1-7). IEEE.
28. Inuwa, M. M., & Das, R. (2024). A comparative analysis of various machine learning methods for anomaly detection in cyber attacks on IoT networks. Internet of Things, 26, 101162.