Ransomware Detection in ECG-Based IoMT Devices Using Secure Heartbeat Features

Anonymous Author(s)
Submission for double-blind review
IEEE LCN 2025

***Abstract*—Securing Internet of Medical Things (IoMT) devices against ransomware requires not only anomaly detection but also contextual awareness of system integrity. In this paper, we propose a hybrid detection approach that augments classical machine learning with behavioral security indicators derived from a secure heartbeat protocol. These indicators—heartbeat delay, hash consistency, and token validity—are extracted from a real-time monitoring system embedded in a private ECG-based IoMT dataset under normal, anomalous, and ransomware conditions. While standard metrics such as CPU and RAM usage often fail to capture stealthy threats, these additional features enable detection of anomalies that manipulate execution timing, alter critical files, or originate from unauthorized processes. We evaluate Random Forest models trained with and without these features: the base model achieves an AUC of 0.739, which rises to 0.894 with heartbeat delay, 0.918 with hash consistency, and 0.933 with all three features combined. This layered security-aware framework significantly improves the detection of both stealth and brutal ransomware attacks while preserving system efficiency, offering a proactive and lightweight solution for IoMT environments.**

***Index Terms*—Internet of Medical Things (IoMT), Ransomware Detection, Secure Heartbeat, Behavioral Features, Hash Consistency, Token Validation, Random Forest, Stealth Attacks, Brutal Attacks**

I. Introduction

The Internet of Medical Things (IoMT) has revolutionized healthcare by enabling real-time monitoring through smart and interconnected devices such as ECG sensors, pulse oximeters, and insulin pumps. These technologies improve healthcare outcomes but introduce critical security challenges. Among them, ransomware remains one of the most damaging threats, as shown in early studies that highlighted their ability to encrypt data and disrupt system operations [1].

Early work in adversarial evasion demonstrated that malware could bypass machine learning models through subtle, functionality-preserving transformations [2]. These techniques enable stealth ransomware to closely mimic benign behavior, complicating its detection. Meanwhile, traditional rule-based and signature-driven intrusion systems often fail to identify such attacks in IoMT environments. Their reliance on static patterns limits responsiveness to emerging or obfuscated threats—highlighting the need for adaptive, behavior-aware approaches tailored to real-world medical systems [3].

To address these limitations, we propose a hybrid detection framework that enriches machine learning with lightweight behavioral indicators derived from a *secure heartbeat* protocol. The extracted features—*heartbeat delay*, *hash consistency*, and *token validity*—capture timing anomalies, file tampering, and signature mismatches often seen in ransomware. This aligns with recent efforts toward secure and efficient anomaly detection in constrained IoT environments [4].

We evaluate this approach using a private ECG-based IoMT dataset collected under realistic conditions. It includes system activity, physiological signals, and behavioral security metrics under four operating modes: normal operation, benign anomalies, and two ransomware strategies (brutal and stealthy). By comparing models trained with and without the proposed features, we demonstrate that behavioral indicators offer substantial gains in detection accuracy and robustness against stealth techniques.

Our experiments show that a Random Forest model trained solely on traditional system metrics reaches an AUC of 0.739. Augmenting the model with heartbeat delay raises the AUC to 0.894; with delay and hash consistency, to 0.918; and with all three features, to 0.933. These results confirm behavioral signals enhance protection in medical devices.

The main contributions of this work are:

- We design a secure heartbeat monitoring protocol tailored for IoMT devices, generating real-time behavioral features related to timing, file integrity, and authentication.
- We construct a behavioral feature set including heartbeat delay, hash consistency, and token validity, integrated into a lightweight anomaly detection framework.
- We evaluate both standard and augmented Random Forest classifiers, and explore additional supervised and unsupervised models to assess generalizability.
- We provide an extensive analysis of detection performance under both brutal and stealth ransomware conditions using a private ECG-based dataset.

This work contributes to closing the gap between technical detection performance and real-world applicability. By embedding behavioral awareness into IoMT systems, we not only improve early detection but also reduce the risk of false alerts, helping maintain clinical safety and operational continuity.

The remainder of this paper is structured as follows: Section II reviews prior research. Section III details our proposed architecture. Section IV presents our experimental evaluation. Finally, Section V concludes and discusses future work.

II. Related Work

Recent developments in IoMT security emphasize the need for robust and context-aware anomaly detection, particularly to address the increasing sophistication of ransomware attacks.

Research efforts have transitioned from traditional threshold-based approaches to machine learning (ML) models capable of handling variability and noise in real-time clinical data streams.

Early approaches in anomaly detection relied heavily on handcrafted signal features and statistical thresholds. Salem et al. [5] employed statistical behavior modeling to detect failures in biomedical sensors, using local extrema and gradient shifts. While effective in constrained scenarios, such approaches fail to generalize to stealthy threats mimicking legitimate behavior. Rule-based schemes likewise suffer from static assumptions and poor adaptability in dynamic clinical environments [6].

More recent frameworks incorporate machine learning to better handle time-varying and multivariate data. Iacovazzi and Raza [7] proposed an ensemble model combining Random Forest and Isolation Forest, trained on system call graphs extracted from containerized environments, achieving high detection accuracy with low false positive rates. Liu et al. [8] addressed the challenge of concept drift in time series by proposing an unsupervised approach capable of detecting both anomalies and change points simultaneously. Their method leverages fluctuation-based features and rate-of-change transformations to adapt to dynamic environments. These works show the need for accurate, adaptive models under non-stationary conditions.

Behavioral learning has emerged as a promising alternative to metric-centric models. KitNET, developed by Mirsky et al. [9], utilizes an ensemble of autoencoders to learn normal network behavior and detect anomalies in an unsupervised, online manner. While effective for identifying traffic deviations, such models miss system-level integrity or authentication failures—critical when facing ransomware acting below the application layer.

Several IoMT-specific solutions focus on efficient real-time monitoring. Almotiri [10] proposed an AI-driven security framework that combines deep learning and machine learning techniques to detect advanced malware and ransomware in Software Defined Networks (SDN) deployed within IoMT environments. However, lacking behavioral context such as delays or integrity issues may limit detection of stealthy threats operating within normal resource bounds.

The notion of integrity verification has been explored in cyber-physical systems. Kook et al. [11] proposed a lightweight authentication protocol based on one-way hash functions for smart grids, designed to detect tampering during communication. This aligns with our use of *hash consistency* to detect silent file changes, especially in stealth ransomware where traditional metrics stay unchanged.

Authentication-based security is also gaining traction. Xiao et al. [12] proposed a hardware fingerprint-based authentication framework (MCU-Token) for IoT devices, generating request-specific tokens to prevent replay attacks and unauthorized access. This finding supports our use of *token validity* to detect unauthorized system activity and replay attacks in heartbeat transmissions.

Timing-based anomaly detection, including response time and heartbeat delay, has been applied in domains like edge computing and biomedical monitoring. Zhen et al. [13] demonstrated how analyzing heartbeat intervals can effectively detect anomalies in ECG signals on edge devices. This approach inspires our use of *heartbeat delay* to capture latency footprints from ransomware disrupting scheduling flows.

In the medical domain, deep learning has also been explored. Transformer-based architectures, such as the one proposed by Zia et al. [14], have demonstrated strong capabilities for detecting complex anomalies in IoT time-series data. Similarly, Gueriani et al. [15] developed a CNN-LSTM hybrid model that achieved high accuracy on recent IoT intrusion datasets. However, such methods need significant resources, limiting applicability in constrained healthcare settings.

Hybrid frameworks that combine machine learning with embedded security signals have shown promise in balancing detection accuracy and efficiency. For instance, Sharma et al. [16] proposed a blockchain-integrated fog computing framework combining ensemble learning for real-time intrusion detection in IoT networks. Though built for industry, the approach shows growing interest in merging trust, auditability, and learning—principles we extend to medical systems.

Entropy-based methods have gained traction for detecting encrypted payloads. Quince et al. [17] proposed a decentralized entropy-driven approach, achieving 97.3% accuracy across several ransomware families. Shadow et al. [18] introduced a dynamic entropy framework that monitors real-time fluctuations in system processes, sustaining over 90% accuracy over time. However, these features may misclassify compressed or encoded benign files, as shown by McIntosh et al. [19], who demonstrated how attackers can manipulate entropy values using Base64 encoding or partial encryption to evade detection, thus reducing reliability in heterogeneous IoMT environments.

In contrast to these methods, our approach combines the benefits of behavioral anomaly detection with lightweight, security-focused features embedded directly in the system layer. Rather than relying solely on resource metrics or entropy, we integrate *heartbeat delay*, *hash consistency*, and *token validity*—each designed to capture a specific class of ransomware behavior: aggressive slowdown, silent modification, and unauthorized access.

To the best of our knowledge, no prior work has unified these behavioral integrity checks within a machine learning pipeline tailored to the dual challenge of stealth and brutal ransomware in realistic IoMT settings. Furthermore, few studies report evaluations on private datasets collected from actual medical workflows, which we consider critical for producing clinically relevant results.

Our contribution thus extends the state-of-the-art by introducing a modular and proactive framework that embeds behavioral awareness into detection systems without compromising on performance or deployability in constrained medical environments.

III. Proposed Approach

Our proposed framework enhances ransomware detection in Internet of Medical Things (IoMT) systems by combining traditional machine learning classifiers with lightweight behavioral security features derived from a secure heartbeat protocol. As

illustrated in Figure 2, the architecture follows five sequential stages: (1) data preprocessing, (2) behavioral feature extraction, (3) dataset enrichment, (4) model training, and (5) real-time alert generation.

A. Secure Heartbeat Feature Extraction

Our approach introduces three behavioral security indicators based on signed heartbeat messages periodically transmitted by the IoMT device. These indicators—*Heartbeat Delay*, *Hash Consistency*, and *Token Validity*—aim to detect system integrity violations or communication spoofing that may accompany ransomware activity.

- Heartbeat Delay (Δt) reflects the time interval between two consecutive heartbeat messages:

$$\Delta t = t_i - t_{i-1} \tag{1}$$

  If Δt exceeds a predefined threshold (e.g., 6 seconds), it may suggest abnormal system latency induced by ransomware overload.
- Hash Consistency (h_c) evaluates the integrity of a critical system file by comparing its current hash to a reference:

$$h_c = \begin{cases} 1, & \text{if } H_{\text{received}} = H_{\text{expected}} \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

  where $H(\cdot)$ represents the SHA-256 hash of the file. A mismatch suggests file encryption or unauthorized tampering.
- Token Validity (τ_v) verifies message authenticity using an HMAC-SHA256 signature with a shared secret key k:

$$\tau_v = \begin{cases} 1, & \text{if HMAC}(k, m) = \text{Token}_{\text{received}} \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

  where m is the message payload, typically a concatenation of timestamp and hash. Invalid tokens indicate possible spoofed messages or malicious code injection.

Example: Suppose the monitored file `/etc/init.conf` originally has a reference hash $H_{\text{expected}} =$ `e3b0...c442`. If a heartbeat reports a current hash H_{received} that differs due to encryption or tampering, then `hash_consistency` is set to 0. Similarly, if the message timestamp deviates from the expected 3-second interval, `heartbeat_delay` increases. If the HMAC signature fails verification using the shared key, `token_validity` is also set to 0.

This simple logic allows the system to flag anomalies in timing, integrity, or authenticity without requiring deep packet inspection or heavy computation.

These features are computed in real-time and appended to each observation in the monitoring dataset. When integrated into a learning model, they substantially enhance detection performance, especially for stealthy or system-level ransomware behaviors.

Figure 1 illustrates which attack classes are primarily captured by each feature.

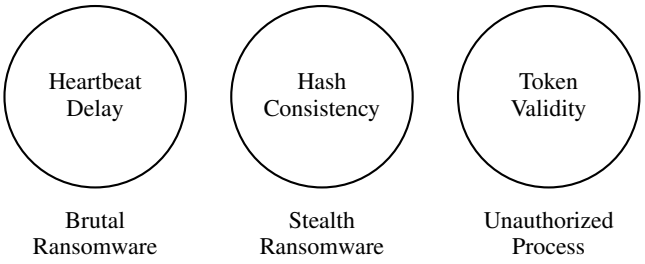

Fig. 1: Each behavioral feature contributes to detecting a specific class of ransomware attack.

B. Dataset Enrichment and Preprocessing

Once extracted, the three heartbeat-based behavioral features are merged with classical IoMT metrics such as CPU usage, RAM consumption, disk I/O, and SpO_2. This combination creates a unified feature space that captures both system performance and security integrity.

To prepare the dataset for temporal modeling, we apply the following preprocessing steps:

- *Min-max normalization*, which rescales all features to a common range, improving model convergence;
- *Timestamp alignment*, ensuring that data from all sensors and heartbeat messages are synchronized;
- *Rolling window statistics* (e.g., mean and standard deviation), computed over short time intervals to capture dynamic fluctuations and short-term trends.

These operations transform the raw input into structured, time-aware sequences suitable for ransomware detection. They also help highlight subtle changes that may indicate stealthy or progressive attacks.

C. Model Training: Supervised and Unsupervised

We evaluated a range of supervised and unsupervised models commonly used for anomaly-rich and imbalanced datasets:

- *Random Forest (RF)*: supervised ensemble classifier for structured data.
- *Gradient Boosting (GB)*: supervised boosting method capturing nonlinear feature interactions.
- *Logistic Regression (LR)*: supervised linear model with probabilistic outputs.
- *Isolation Forest (IF)*: unsupervised anomaly detector using recursive partitioning.
- *Local Outlier Factor (LOF)*: unsupervised density-based local anomaly detector.
- *One-Class SVM (OCSVM)*: unsupervised model defining a boundary around normal data.

All models were trained on a stratified 80/20 split of a private ECG-based IoMT dataset. Heartbeat-based features were included in all models to assess their impact on detection performance.

D. Real-Time Alerting Mechanism

Rather than using continuous scoring, our system triggers an alert as soon as a sample is flagged with high confidence.

This enables immediate actions such as isolation, logging, or backup—minimizing risk before spread.

This simple threshold-based logic ensures fast response with minimal complexity, ideal for clinical IoMT settings.

While inference and alerting are performed in real time, model training remains offline. This allows for robust learning using historical data while maintaining low-latency detection during deployment. Periodic retraining can be scheduled to reflect changes in device behavior or the emergence of novel threats.

Figure 2 shows the full pipeline from telemetry to alert generation.

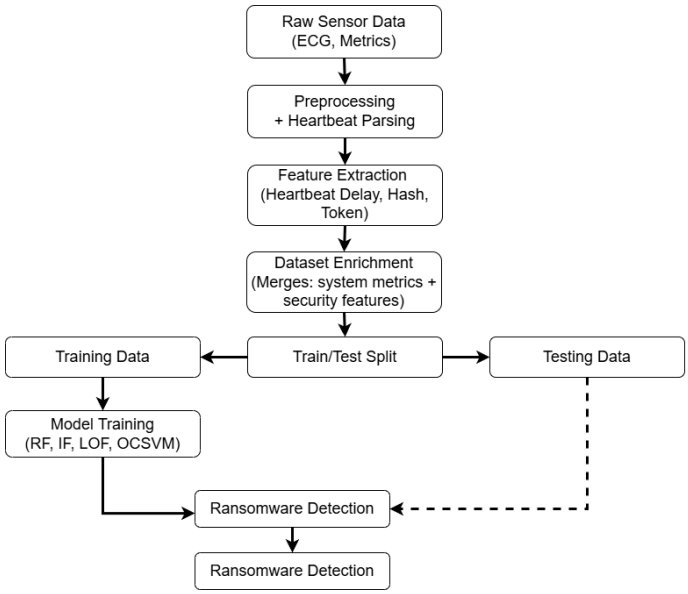

Fig. 2: Proposed detection pipeline combining classical system metrics with behavioral heartbeat integrity features.

This modular framework is suitable for both edge-level and gateway-level deployments. It improves ransomware detection precision and recall by extending traditional monitoring with behavioral integrity verification. Section IV provides an in-depth evaluation on our private dataset under realistic conditions.

IV. Experimental Results

A. Dataset and Experimental Setup

We conducted our evaluation using a proprietary IoMT dataset originally derived from ECG monitoring devices. The dataset contains 60,000 time-series samples captured at 3-second intervals, combining physiological signals (e.g., BPM, HRV) with system telemetry (e.g., CPU, RAM, Disk I/O, Network). To improve ransomware detection, three behavioral integrity indicators were appended to each sample: `heartbeat_delay`, `hash_consistency`, and `token_validity`, derived from a secure heartbeat protocol. This fusion creates a hybrid telemetry space representing both performance and security behavior.

The dataset includes four distinct types of events:

- Normal behavior (48,274 samples) — regular patient monitoring without abnormal activity.
- Benign anomalies (4,791 samples) — events such as software updates, sensor recalibrations, and network fluctuations. These resemble ransomware behavior in disk or CPU metrics but are non-malicious.
- Stealth ransomware (3,314 samples) — slow, disguised file encryption with minimal impact on system metrics.
- Brutal ransomware (3,621 samples) — rapid and aggressive encryption, visibly impacting resource usage and timing.

To clarify the nature of benign anomalies, Table I provides examples of common non-malicious disruptions present in the dataset.

TABLE I: Examples of Benign Anomalies

Anomaly Type	Description
Sensor Recalibration	Temporary data spike during manual adjustment
Software Update	Increased CPU/disk from update processes
Network Dropout	Short-term disconnection and reconnection
Backup Event	Legitimate high disk activity from snapshot

The class distribution is illustrated in Figure 3, showing a strong imbalance, with 80.5% normal samples and only 6% stealth ransomware. All models were trained using 80/20 stratified splits and evaluated with 5-fold cross-validation in scikit-learn.

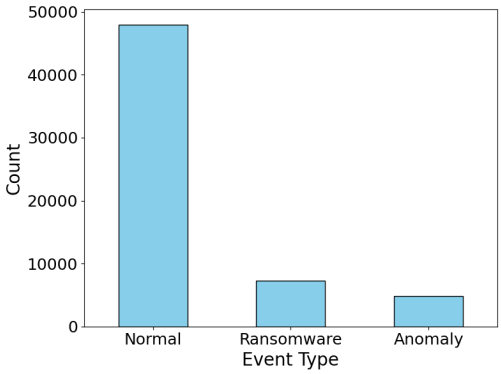

Fig. 3: Distribution of event types across the dataset.

B. Feature Overview

The feature space blends physiological, system-level, and behavioral dimensions. Table II presents a subset of key indicators used during training and evaluation.

TABLE II: Overview of Feature Space (selected examples)

Feature	Description
BPM	Heart rate monitoring signal
HRV (ms)	Heart rate variability
CPU Usage (%)	System load
RAM Usage (MB)	Memory pressure
Disk Activity (MB/s)	Ransomware activity indicator
Network Sent/Received (kB/s)	Communication patterns
heartbeat_delay	Irregularity in heartbeat timestamp intervals
hash_consistency	Integrity of critical system files
token_validity	Message authenticity check

To better understand feature behavior, we analyzed `Disk Activity (MB/s)` across the three main event types (Normal, Anomaly, Ransomware). As shown in Figure 4, ransomware events tend to have higher disk usage than normal activity, but significant overlap with anomalies makes this feature unreliable on its own.

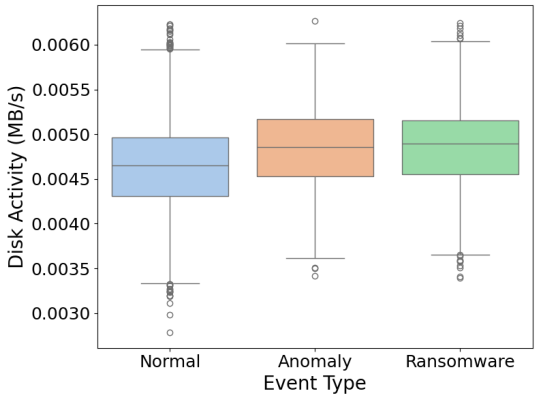

Fig. 4: Boxplot of disk activity across event types.

This observation highlights the need for behavioral features: since benign anomalies (e.g., software patches or backups) may spike disk usage without indicating ransomware, relying solely on classical metrics risks high false positives. Behavioral indicators such as `hash_consistency` and `token_validity` capture integrity violations and message authenticity failures that help disambiguate these cases.

C. Threat Model and Assumptions

Our detection framework assumes a threat landscape involving both stealthy and overt ransomware behaviors. The attacker is assumed to have gained sufficient access to the IoMT device to execute malicious code, encrypt or tamper with system files, or manipulate device scheduling. Two attack categories are considered:

- Brutal ransomware: rapidly encrypts files or overloads system resources, producing detectable anomalies such as CPU spikes, disk I/O bursts, or heartbeat delay.
- Stealth ransomware: operates at low intensity over extended periods, mimicking benign anomalies and making detection more difficult without integrity checks or authentication mechanisms.

The model assumes that the secure heartbeat mechanism—responsible for transmitting timestamped, signed telemetry with file hashes—is trustworthy. Specifically:

- The attacker cannot access the HMAC signing key used to validate heartbeat tokens.
- The attacker cannot spoof or forge valid heartbeat messages without detection.
- The baseline hash of monitored files is known and securely stored.

The system operates in a passive detection mode. It does not attempt to neutralize or block malicious activity, but rather detects and flags suspicious behavior for downstream mitigation. Real-time response mechanisms (e.g., isolation or backup) are implemented as a separate containment layer once an alert is triggered.

D. Detection Performance Overview

We evaluated the detection capability of the Random Forest model in four configurations: without any heartbeat features, with only `heartbeat_delay`, with both `heartbeat_delay` and `hash_consistency`, and with all three behavioral features (`heartbeat_delay`, `hash_consistency`, `token_validity`).

Figure 5 shows the ROC curve of the baseline model trained without any behavioral security indicators. Figure 6 illustrates the complete progression in detection performance as each heartbeat feature is successively added.

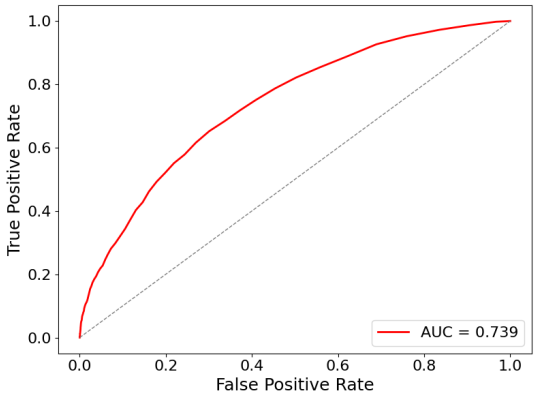

Fig. 5: ROC curve of the baseline Random Forest model without behavioral features (AUC = 0.739).

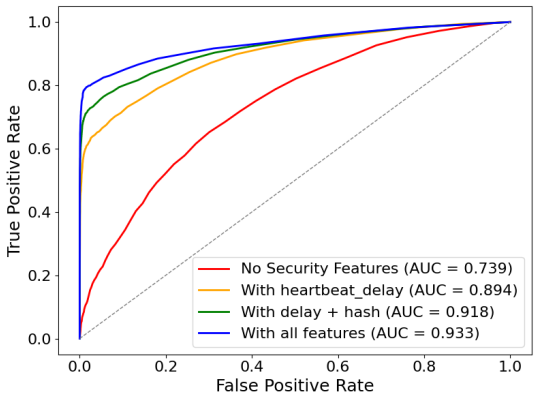

Fig. 6: Progressive impact of each heartbeat-based feature on Random Forest ROC curve.

As seen in Figure 5, the baseline model performs poorly in detecting ransomware, with an AUC of 0.739 and a very low recall (5.1%). However, as shown in Figure 6, detection performance improves substantially with each added security feature. The final model—including all heartbeat features—achieves an AUC of 0.933, confirming the effectiveness of this layered behavioral strategy.

The ROC (Receiver Operating Characteristic) curve plots the True Positive Rate (Recall) against the False Positive Rate across varying decision thresholds. A perfect model reaches the top-left corner (TPR = 1, FPR = 0), while random guessing yields a diagonal line with AUC = 0.5. As visible in Figure 6, each added behavioral feature shifts the curve closer to ideal separation. This progressive improvement confirms their relevance in distinguishing ransomware from benign activity under various thresholds.

E. Impact of Behavioral Features

We quantitatively evaluated the impact of each heartbeat-based feature by progressively adding them to the feature space and observing the evolution of detection performance. The model used for all evaluations was Random Forest, trained on an 80/20 stratified split with 5-fold cross-validation.

The performance was assessed using four main metrics: Precision, Recall, F1-score, and Area Under the ROC Curve (AUC), defined as:

$$\text{Precision} = \frac{TP}{TP + FP} \qquad \text{Recall} = \frac{TP}{TP + FN}$$

$$\text{F1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \qquad \text{AUC} = \int_0^1 \text{ROC}(t)dt$$

To interpret these metrics, we recall the definitions of core classification outcomes: true positives (TP) are correctly identified ransomware samples; false negatives (FN) are ransomware samples missed by the model; false positives (FP) are benign samples wrongly flagged as ransomware. Precision reflects how many predicted ransomware samples are truly malicious, while recall measures the proportion of actual ransomware detected. The F1-score balances both. Accuracy is the overall percentage of correct predictions, but may be misleading in imbalanced datasets where normal events dominate.

Table III summarizes how performance improves as we integrate each feature:

The baseline model, without any heartbeat features, performs poorly in detecting ransomware. Despite a decent overall accuracy, its recall for ransomware is only 5.1%, indicating that most attacks go undetected. The AUC of 0.739 confirms its limited ability to discriminate ransomware from benign activity.

Adding `heartbeat_delay`, derived from temporal irregularities in secured heartbeats, yields a major breakthrough. It raises the recall to 52.6% and boosts AUC to 0.894, capturing early-stage behavioral drift.

The addition of `hash_consistency`, which validates the stability of critical system files, further improves detection robustness. It increases recall to 65.5% and the F1-score to 0.774, indicating better balance between sensitivity and specificity.

Finally, including `token_validity`, which ensures the authenticity of transmitted data, completes the feature set. The model achieves a recall of 74.6%, with an AUC of 0.933 and an F1-score of 0.839 — demonstrating strong and reliable performance across all ransomware profiles.

Each behavioral feature targets a distinct ransomware trait: `heartbeat_delay` detects abnormal timing behaviors (especially in brutal attacks), `hash_consistency` captures subtle integrity drifts in stealth scenarios, and `token_validity` validates the authenticity of messages, often tampered with by sophisticated threats. Together, they form a complementary stack for robust detection.

F. Model Comparison

To assess the effectiveness of the proposed behavioral features, we evaluated three supervised classifiers under two configurations: a baseline using only traditional system and sensor metrics, and a full feature set including the secure heartbeat indicators.

Table IV summarizes the Area Under the Curve (AUC) scores for each model under both settings.

TABLE IV: Comparison of Supervised Models (AUC)

Model	Baseline (AUC)	Full (AUC)
Random Forest (RF)	0.739	0.933
Gradient Boosting (GB)	0.722	0.901
Logistic Regression (LR)	0.608	0.770

All classifiers showed a substantial improvement in AUC when enriched with heartbeat-based features, confirming the added value of behavioral indicators. Among them, the Random Forest consistently achieved the highest detection performance in both configurations.

Given its superior accuracy, robustness, and ability to capture non-linear feature interactions, Random Forest was selected as the primary model for the remaining experiments.

G. Impact Analysis

Figure 7 shows the feature importance scores from the Random Forest model trained with all features.

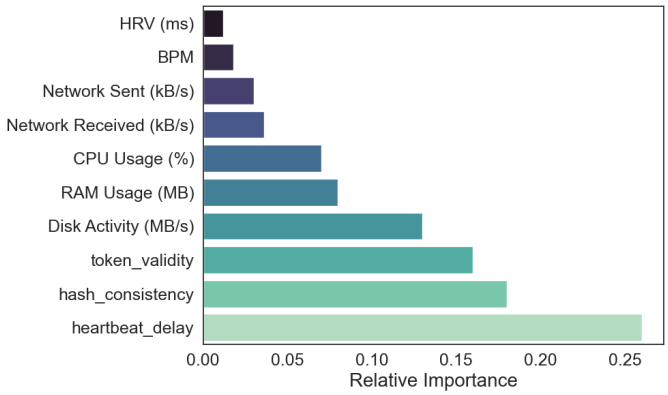

Fig. 7: Feature Importance - Random Forest (with heartbeat features)

The top three features (`heartbeat_delay`, `hash_consistency`, and `token_validity`) highlight the impact of behavioral security indicators. These features, extracted from a secure heartbeat mechanism, proved more

TABLE III: Random Forest Performance as Heartbeat Features Are Added

Features	AUC	Precision (Ransom)	Recall (Ransom)	F1-score
Without heartbeat features	0.739	0.655	0.051	0.095
+ `heartbeat_delay`	0.894	0.932	0.526	0.673
+ `hash_consistency`	0.918	0.946	0.655	0.774
+ `token_validity` (full)	0.933	0.958	0.746	0.839

discriminative than traditional system metrics. Lower-ranked features such as `BPM` and `HRV (ms)` confirm that physiological signals were less relevant for detecting ransomware activity.

H. Stealth vs. Brutal Detection

To better understand the behavioral features' contribution across attack types, we analyzed the AUC separately for brutal and stealth ransomware variants using Random Forest. Results are shown in Table V.

TABLE V: Detection AUC by Ransomware Type (Random Forest)

Feature Configuration	Brutal AUC	Stealth AUC
Without heartbeat features	0.834	0.646
+ `heartbeat_delay`	0.924	0.681
+ `hash_consistency`	0.935	0.742
+ `token_validity` (full)	0.957	0.823

Observations:

- `heartbeat_delay` significantly improved brutal ransomware detection by capturing anomalies in system activity timing, raising AUC from 0.834 to 0.924.
- `hash_consistency` provided early signs of stealth activity through subtle integrity changes in monitored files, leading to a notable increase in stealth AUC (from 0.681 to 0.742).
- `token_validity` further strengthened stealth detection by modeling message authenticity failures common in advanced attacks, raising stealth AUC to 0.823.

Overall, the full behavioral stack enhanced both detection fronts, with brutal attacks reaching 0.957 AUC and stealth variants showing the strongest relative gain—improving by +27.7% from baseline—while maintaining a low false negative rate.

Compared to existing intrusion detection approaches such as Kitsune [9], which uses the KitNET algorithm, and CNN-LSTM models [15], our method offers a lightweight and interpretable alternative tailored to real-time IoMT environments. Kitsune performs online anomaly detection at the network level using an ensemble of autoencoders and runs efficiently even on devices like Raspberry Pi. However, it lacks visibility into file integrity and system-level timing behaviors. CNN-LSTM approaches, while achieving high accuracy (98.57% F1-score on IoT datasets), rely on deep supervised architectures that are computationally intensive and less suited for constrained medical devices. In contrast, our approach integrates low-cost behavioral indicators such as `heartbeat_delay`, `hash_consistency`, and `token_validity`, enabling robust ransomware detection—including stealth attacks—without incurring high computational overhead.

I. Real-Time Response Strategy

To complement detection, we implemented a reactive response mechanism suited to clinical IoMT environments. When the model predicts a ransomware probability above a predefined threshold (e.g., 90%), it triggers a lightweight containment protocol to ensure high confidence while preserving system stability.

The response includes:

- Isolating the affected device from the network to prevent lateral spread;
- Notifying clinical staff via dashboard or secure messaging;
- Initiating a local backup or snapshot to preserve patient data.

These actions are managed by a lightweight agent at the gateway level, avoiding firmware changes and ensuring compatibility with existing infrastructure. Inference time remains under 10 ms, enabling rapid decisions even on constrained hardware.

To prevent alert fatigue during extended attacks, a cooldown policy suppresses repeated alerts for the same device. While the system does not actively block attacks, it enables fast containment with minimal disruption to clinical workflows.

J. Deployment Considerations

The proposed framework is designed for deployment at the gateway level, where telemetry streams are collected and processed before reaching centralized infrastructure. This architecture eliminates the need for firmware modifications on medical devices, while still enabling real-time anomaly detection.

In high-risk environments, heartbeat validation can optionally be embedded at the device level to enhance trust guarantees. All models are inference-ready and compatible with edge hardware supporting Python or other lightweight runtime environments.

K. Limitations and Future Directions

While the proposed framework achieves strong detection accuracy and real-time alerting, several limitations remain. First, it treats all ransomware types as a single class, which limits its forensic and attribution capabilities.

Second, the model is trained offline and may degrade over time in the presence of concept drift or emerging threats, unless adaptive learning is introduced.

Third, false positives may arise during legitimate updates or maintenance operations, particularly when `hash_consistency` or `token_validity` deviate from expected patterns. Context-aware whitelisting or dynamic baselines could mitigate this issue.

Future work will explore real-time model adaptation, per-device learning, and finer-grained classification of ransomware variants.

L. Key Findings

- *AUC Improvement:* Random Forest AUC improved from 0.739 (baseline) to 0.933 with full behavioral features.
- *Stealth Detection:* False negatives were reduced by over 70%, particularly for stealth attacks, due to the addition of heartbeat-based features.
- *Detection Quality:* The final model achieved strong ransomware classification performance, as confirmed by F1-score and balanced precision-recall metrics.
- *Real-Time Readiness:* The detection system enables actionable, low-latency responses suitable for clinical IoMT environments.

V. Conclusion

This paper introduced a lightweight and security-aware detection framework for IoMT ransomware, integrating traditional system metrics with behavioral indicators extracted from a secure heartbeat protocol. These features—heartbeat delay, hash consistency, and token validity—enabled the model to capture timing anomalies, file integrity breaches, and authentication failures associated with both stealthy and aggressive ransomware behaviors.

Our experimental results demonstrated that the Random Forest classifier, when enriched with heartbeat-based features, improved its AUC from 0.739 to 0.933 and significantly reduced false negatives, especially in stealth scenarios. Feature importance analysis further confirmed that behavioral features outperformed standard metrics in ransomware classification. The framework also supports low-latency alerts and response strategies, essential for medical environments.

This design aligns well with the operational constraints of modern healthcare environments, where detection speed, interpretability, and ease of deployment are often more critical than theoretical model complexity.

While the experiments focused on ECG-based telemetry, the proposed approach is transferable to other critical IoT infrastructures—such as industrial control systems or smart energy platforms—where behavioral drift and integrity violations are equally relevant. The framework could also serve as a modular layer within a larger intrusion detection system (IDS), complementing signature-based and network-level defenses.

For future work, deeper integration of temporal learning models such as LSTMs or hybrid CNN-LSTM architectures [20] could improve early-stage detection, as long as latency and computational constraints remain satisfied. Additionally, adopting federated learning or edge-native inference [21] would enhance scalability and data privacy, particularly in multi-device deployments. Finally, we plan to explore more advanced real-time response strategies triggered by the model, including adaptive device isolation, dynamic threat escalation, and rollback of encrypted data.

References

[1] A. Kharraz, W. Robertson, D. Balzarotti, L. Bilge, and E. Kirda, "Cutting the gordian knot: A look under the hood of ransomware attacks," in *Detection of Intrusions and Malware, and Vulnerability Assessment: 12th International Conference, DIMVA 2015, Milan, Italy, July 9-10, 2015, Proceedings 12*. Springer, 2015, pp. 3–24.

[2] H. S. Anderson, A. Kharkar, B. Filar, and P. Roth, "Evading machine learning malware detection," *black Hat*, vol. 2017, pp. 1–6, 2017.

[3] A. Naghib, F. S. Gharehchopogh, and A. Zamanifar, "A comprehensive and systematic literature review on intrusion detection systems in the internet of medical things: current status, challenges, and opportunities," *Artificial Intelligence Review*, vol. 58, no. 4, pp. 1–88, 2025.

[4] V. Prakash, O. Odedina, A. Kumar, L. Garg, and S. Bawa, "A secure framework for the internet of things anomalies using machine learning," *Discover Internet of Things*, vol. 4, no. 1, p. 33, 2024.

[5] O. Salem, A. Guerassimov, A. Mehaoua, A. Marcus, and B. Furht, "Anomaly detection scheme for medical wireless sensor networks," *Handbook of medical and healthcare technologies*, pp. 207–222, 2013.

[6] B. R. Kikissagbe and M. Adda, "Machine learning-based intrusion detection methods in iot systems: A comprehensive review," *Electronics*, vol. 13, no. 18, p. 3601, 2024.

[7] A. Iacovazzi and S. Raza, "Ensemble of random and isolation forests for graph-based intrusion detection in containers," in *2022 IEEE International Conference on Cyber Security and Resilience (CSR)*. IEEE, 2022, pp. 30–37.

[8] J. Liu, D. Yang, K. Zhang, H. Gao, and J. Li, "Anomaly and change point detection for time series with concept drift," *World Wide Web*, vol. 26, no. 5, pp. 3229–3252, 2023.

[9] Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai, "Kitsune: An ensemble of autoencoders for online network intrusion detection," 2018.

[10] S. H. Almotiri, "Ai driven iomt security framework for advanced malware and ransomware detection in sdn," *Journal of Cloud Computing*, vol. 14, no. 1, p. 19, 2025.

[11] S. Kook, K. Kim, J. Ryu, Y. Lee, and D. Won, "Lightweight hash-based authentication protocol for smart grids," *Sensors*, vol. 24, no. 10, p. 3085, 2024.

[12] Y. Xiao, Y. He, X. Zhang, Q. Wang, R. Xie, K. Sun, K. Xu, and Q. Li, "From hardware fingerprint to access token: Enhancing the authentication on iot devices," *arXiv preprint arXiv:2403.15271*, 2024.

[13] P. Zhen, Y. Han, A. Dong, and J. Yu, "Careedge: a lightweight edge intelligence framework for ecg-based heartbeat detection," *Procedia Computer Science*, vol. 187, pp. 329–334, 2021.

[14] S. Zia, N. Bibi, S. Alhazmi, N. Muhammad, and A. Alhazmi, "Enhanced anomaly detection in iot through transformer-based adversarial perturbations model," *Electronics*, vol. 14, no. 6, 2025.

[15] A. Gueriani, H. Kheddar, and A. C. Mazari, "Enhancing iot security with cnn and lstm-based intrusion detection systems," in *2024 6th International Conference on Pattern Analysis and Intelligent Systems (PAIS)*. IEEE, 2024, pp. 1–7.

[16] P. Sharma, A. Kalia, and H. Saini, "A hybrid ensemble framework for intrusion detection in iot using blockchain-integrated fog networks," *Journal of Information Systems Engineering and Management*, vol. 10, no. 2s, 2025.

[17] B. Quince, L. Gareth, S. Larkspur, T. Wobblethorn, and T. Quibble, "Decentralized entropy-based ransomware detection using autonomous feature resonance," *arXiv preprint arXiv:2502.09833*, 2025.

[18] K. Shadow, W. Fairbanks, N. Bexley, O. Radcliffe, and X. Langford, "An adaptive ransomware detection method using dynamic entropy analysis," *TechRxiv preprint*, 2024.

[19] T. McIntosh, J. Jang-Jaccard, P. Watters, and T. Susnjak, "The inadequacy of entropy-based ransomware detection," in *Neural Information Processing: 26th International Conference, ICONIP 2019, Sydney, NSW, Australia, December 12–15, 2019, Proceedings, Part V 26*. Springer, 2019, pp. 181–189.

[20] M. Jouhari and M. Guizani, "Lightweight cnn-bilstm based intrusion detection systems for resource-constrained iot devices," in *2024 International Wireless Communications and Mobile Computing (IWCMC)*. IEEE, 2024, pp. 1558–1563.

[21] Y. Chen, X. Qin, J. Wang, C. Yu, and W. Gao, "Fedhealth: A federated transfer learning framework for wearable healthcare," *IEEE Intelligent Systems*, vol. 35, no. 4, pp. 83–93, 2020.