Ransomware Attacks on IoMT: Detection via Behavioral and Entropy-Based Features

Réda Arbane, Ahmad Alashour, Osman Salem and Ahmed Mehaoua
Centre Borelli UMR 9010
Université Paris Cité
Paris, France
{firstname.lastname}@u-paris.fr

Abstract—**Anomaly detection in the Internet of Medical Things (IoMT) is critical for ensuring patient safety and operational integrity in healthcare systems. This paper proposes a layered detection framework combining entropy-based features—including Shannon Entropy and a novel Super Entropy metric—with supervised and unsupervised learning to identify ransomware and benign anomalies in a private SpO2 sensor dataset capturing physiological signals (e.g., heart rate, oxygen saturation) and device metrics (CPU/RAM usage) under normal, anomalous, and ransomware scenarios. To address real-world variability, the dataset includes stealthy and brutal ransomware attacks alongside benign anomalies. However, due to the sensitive nature of this private dataset, we complied with strict healthcare data protection protocols. Our approach introduces a confidence-based scoring mechanism with delayed alerts to minimize redundancies. Evaluations of Random Forest, Isolation Forest, Local Outlier Factor, and One-Class SVM demonstrate that entropy features consistently enhance performance: Random Forest achieves a 0.937 AUC on full ransomware detection (vs. 0.794 baseline) and 0.959 for brutal attacks (vs. 0.872), underscoring the value of entropy-aware learning and adaptive alerting for proactive IoMT security.**

Index Terms—**Internet of Medical Things (IoMT), Ransomware Detection, Entropy, Super Entropy, Anomaly Detection, Random Forest, Brutal Attacks, Stealth Attacks, Alert Scoring System**

I. Introduction

The Internet of Medical Things (IoMT) has transformed modern healthcare through ubiquitous, real-time patient monitoring via interconnected devices such as smart SpO_2 sensors, cardiac monitors, and infusion pumps. While this technological advancement improves clinical decision-making and outcomes, it also introduces major cybersecurity risks. One particularly concerning threat is ransomware, which can encrypt vital medical data or disrupt device operations—threats that may directly endanger patient lives when targeting life-supporting equipment [1].

In recent years, ransomware has evolved beyond basic mass attacks. Modern variants employ progressive encryption and sophisticated evasion strategies, such as behavioral camouflage that mimics legitimate system activities, making them harder to detect with traditional security mechanisms [2]. As a result, signature-based detection systems and static rule-based firewalls prove increasingly ineffective against these polymorphic threats [3]. Even advanced Machine Learning models relying only on system metrics (CPU, memory) often fail to detect subtle anomalies in resource-constrained IoMT environments [4].

To overcome these limitations, our research focuses on behavioral anomaly detection through information-theoretic analysis. Prior works have demonstrated that Shannon Entropy is a valuable metric to detect subtle anomalies in traffic behavior within IoT networks [5]. Building on this, we propose the concept of Super Entropy — a refined metric designed to enhance early-stage detection of cryptographic activity while minimizing sensitivity to natural fluctuations in IoMT operations. We propose a multi-layered anomaly detection framework that integrates entropy-aware features with both supervised and unsupervised learning models. To minimize false alarms, we also introduce an adaptive scoring mechanism that triggers alerts only when the prediction confidence surpasses a calibrated threshold. A built-in suppression window avoids redundant alerts during prolonged attack phases, making the system more practical for operational deployment.

To evaluate our framework, we used a comprehensive private IoMT dataset collected under realistic conditions. The dataset includes physiological signals (e.g., heart rate, temperature, SpO2), system-level metrics (CPU, RAM, disk, and network activity), and entropy features under three conditions: normal activity, benign anomalies (e.g., software updates), and two ransomware behaviors: brutal (aggressive and rapid) and stealthy variants (slow and concealed). Due to the sensitive nature of the medical data, this dataset is not publicly released, in accordance with healthcare data protection principles.

Our results show that entropy-aware features significantly enhance model performance. For example, entropy integration improved Random Forest accuracy across both stealthy and brutal ransomware detection, with substantial AUC gains over baseline models. These improvements demonstrate the added value of entropy and adaptive scoring in securing IoMT systems under real-world conditions.

The main contributions of this work are:

- We use a labeled IoMT dataset capturing normal activity, benign anomalies, and multiple ransomware strategies under realistic conditions.
- We introduce Super Entropy—a refined entropy-based metric to highlight subtle encryption behavior.
- We implement a scoring-based alerting mechanism that combines prediction confidence with temporal suppression

for realistic anomaly reporting.

- We evaluate both supervised (Random Forest) and unsupervised (Isolation Forest, LOF, One-Class SVM) models, demonstrating consistent improvements when using entropy-based features.

These contributions are designed to not only improve technical detection accuracy, but also ensure operational viability in clinical environments where false alerts or delays could have critical consequences for patient care.

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 discusses our conclusions and future directions.

II. Related Work

Recent progress in IoMT security has underscored the need for robust anomaly detection to counter increasingly sophisticated cyber threats. Research has evolved from basic statistical methods to advanced machine learning (ML) and deep learning (DL) approaches, each targeting specific challenges in clinical settings.

Early IoMT anomaly detection often relied on statistical techniques like moving averages or signal derivatives. For instance, [6] used handcrafted features—local minima, gradient shifts, and smoothing filters—to detect pressure-induced failures in glucose monitoring sensors. While effective in narrow contexts, such methods struggle with the variability and noise of real-world environments. Unlike rigid rule-based maintenance, recent work promotes real-time predictive models leveraging IoMT data to foresee failures such as CT tube arcing [7].

More recently, machine learning models have shown better performance, especially in situations with limited labeled data. For example, Zou et al. [8] proposed an optimized version of Isolation Forest that achieved high detection accuracy while maintaining a lightweight footprint, making it suitable for real-time anomaly detection in resource-constrained environments. In a similar direction, Liu and colleagues [9] proposed an unsupervised method for simultaneously detecting anomalies and change points in time series data with concept drift, which is particularly relevant for continuous monitoring in wearable medical devices.

Supervised ML models, such as Random Forest and Support Vector Machines, have also been widely used, benefiting from labeled datasets to learn discriminative patterns. Yet, their success depends heavily on data quality and class balance, which are often hard to guarantee in sensitive medical applications.

Deep learning has become a key approach for anomaly detection in IoMT systems thanks to its ability to model complex, nonlinear patterns. Transformer-based architectures, like TiSAT [10], capture long-range dependencies in time series. CNN-LSTM hybrids with residual blocks and attention mechanisms, such as in [11], also perform well in multivariate sensor environments, particularly in demanding industrial contexts. Simpler recurrent models like LSTM remain effective for temporal trends [12], while CNNs extract spatial features from multivariate inputs [13]. As noted by Briskilla and Rajkumar [14], the complexity of medical time series and real-time constraints demand detection methods that are not just accurate, but also efficient and adaptable—especially under limited supervision or compute resources.

To improve efficiency while preserving accuracy, lightweight neural architectures or pruning strategies have been proposed. These methods aim to reduce inference time and memory consumption, but often at the cost of lower robustness or interpretability.

Meanwhile, entropy-based detection has gained traction for identifying ransomware. For instance, Lee et al. [15] proposed an entropy estimation method for cloud services, achieving 100% detection with zero false positives by leveraging statistical uniformity in encrypted content. Building on this, a follow-up study by the same group [16] targeted stealthier ransomware using format-preserving encryption (FPE). By combining entropy features with lightweight models like KNN and Decision Trees, they reached an average precision of 94.64%, effectively addressing advanced evasion strategies.

Other works, such as [17], show that relying solely on Shannon entropy can cause misclassifications, particularly with compressed files. Their study of over 50 entropy metrics found that Chi-square and SP-800 Serial tests offer more reliable distinctions between encrypted and non-encrypted high-entropy files. Adaptive detection frameworks now represent the state-of-the-art. Newaz et al. [18] developed HealthGuard, a machine learning framework that monitors vital signs in real time and adapts to evolving threats while preserving patient safety. Unlike systems using only raw telemetry, Dahmen and Cook [19] proposed Isudra, an indirectly supervised anomaly detector that learns from sparse examples to reduce irrelevant alerts. Their method targets clinically meaningful events, cutting false positives in smart home health monitoring. Unlike prior work based on public datasets, our study uses a private dataset from realistic SpO_2-based activity under multiple conditions (normal, anomalous, and ransomware-induced). This reflects real-world constraints and the sensitivity of medical data, supporting robust evaluation of hybrid detection models. Although entropy analysis and machine learning classifiers each show promise, few studies have explored their integration in real-time IoMT settings—especially for stealth ransomware, which can evade traditional behavior-based filters. Our work contributes to this growing body of research by proposing a layered detection architecture that incorporates both entropy dynamics and adaptive scoring across supervised and unsupervised machine learning models.

To our knowledge, no existing work has combined these elements while addressing both brutal and stealthy ransomware attacks in a realistic IoMT environment. Our evaluation setup provides a high level of operational realism, which, to our knowledge, has not been achieved in prior studies.

III. Proposed Approach

Our ransomware detection framework for IoMT systems combines information-theoretic analysis with machine learning to identify both obvious and stealthy encryption attacks. As shown in Figure 2, the architecture consists of three core components: (1) entropy-based feature extraction, (2) hybrid machine learning detection, and (3) an adaptive scoring mechanism for operational deployment.

A. Entropy-Based Feature Engineering

The foundation of our approach lies in information-theoretic monitoring of system behavior. We employ two complementary entropy metrics:

- Shannon Entropy ($H(X)$) quantifies byte-level randomness in system files:

$$H(X) = -\sum_{i=1}^{n} p(x_i) \log_2 p(x_i) \tag{1}$$

  where $p(x_i)$ represents the probability of byte value x_i occurrence. This metric effectively captures the characteristic entropy surge during file encryption while remaining computationally lightweight for IoMT devices.
- Super Entropy ($S(X)$) enhances detection sensitivity through threshold-based activation:

$$S(X) = \begin{cases} 1, & \text{if } H(X) \geq \tau \text{ and } \Delta H(X) > \delta \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

  with $\tau = 7.8$ empirically determined to maximize discrimination between cryptographic and benign high-entropy events (e.g., system updates).

As visualized in Figure 1, these metrics create distinct behavioral signatures that differentiate normal operation, benign anomalies, and ransomware activity.

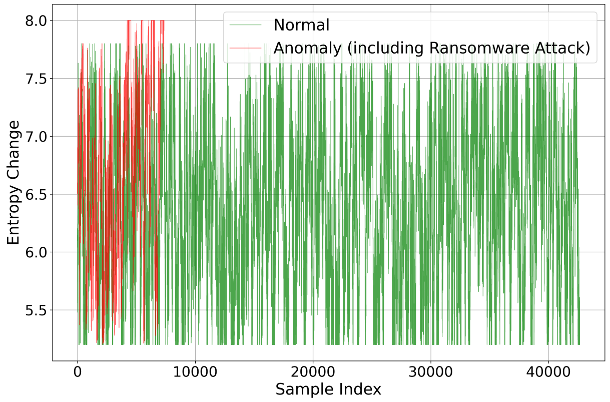

Fig. 1: Entropy behavior across samples. Green: normal activity. Red: benign anomalies, stealth, and brutal ransomware attacks.

B. Hybrid Detection Models

We integrate entropy features with conventional system metrics through a model-agnostic framework supporting both supervised and unsupervised paradigms:

- Random Forest: Leverages ensemble learning for robust classification using entropy-augmented feature vectors
- Isolation Forest: Identifies anomalies through recursive partitioning of the entropy-feature space
- Local Outlier Factor: Detects subtle behavioral deviations via density-based analysis
- One-Class SVM: Models normal operation boundaries for unknown-threat detection

All models undergo identical preprocessing including temporal alignment, min-max normalization, and rolling-window feature derivation to capture dynamic system states.

C. Adaptive Threat Scoring

To bridge model outputs with clinical operational needs, we introduce a dynamic scoring layer:

$$\text{Threat Score} = \sum_{i=1}^{k} \alpha_i \cdot f_i + \beta \cdot \mathbb{I}\big(S(X) = 1\big) \tag{3}$$

where f_i represent normalized system features (CPU, memory, etc.), α_i their empirical weights, and β amplifies the Super Entropy indicator. The scoring mechanism incorporates:

- Confidence-based thresholding to minimize false alerts
- Temporal suppression to prevent alert flooding during sustained anomalies
- Progressive escalation for persistent threats

This layered design addresses key IoMT constraints:

- Computational efficiency through lightweight entropy monitoring
- Adaptability via hybrid model support
- Operational practicality with tunable alerting

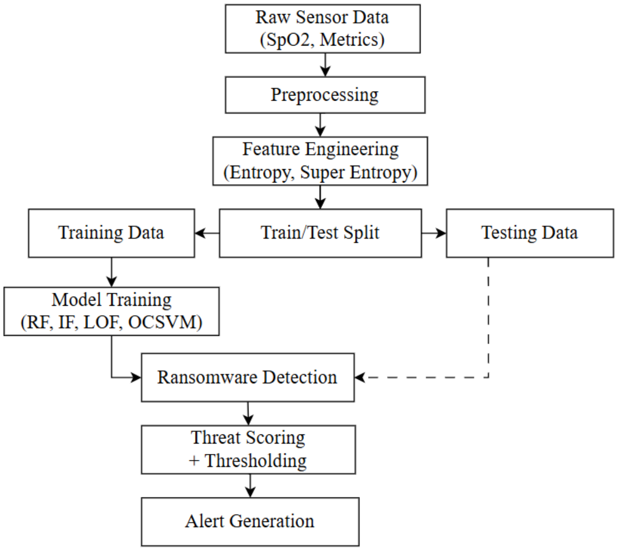

Fig. 2: Architecture of our proposed ransomware detection pipeline.

The framework's modular design permits selective activation of components based on device capabilities - from basic entropy monitoring on resource-constrained sensors to full model-scoring deployment on medical gateways. Section IV validates this flexibility through comprehensive testing across attack scenarios and hardware profiles.

IV. Experimental Results

A. Dataset and Experimental Setup

We evaluated our framework on a proprietary IoMT dataset by extracting 50,000 time-series samples from a wider set of clinical SpO_2 monitoring data, collected at 3-second intervals. The distribution of event types was as follows: 85% normal behavior, 8% benign anomalies, 4% stealth ransomware, and 3% brutal ransomware. The dataset captures four operational states (The following sample counts are approximate) :

- Normal activity (42,500 samples): Baseline device operation during patient monitoring
- Benign anomalies (4,000 samples): System updates, network congestion, and sensor recalibrations
- Ransomware attacks (3,500 samples):
  - Stealth variants (2,000 samples): Slow, targeted encryption mimicking system processes
  - Brutal variants (1,500 samples): Rapid, indiscriminate encryption with visible system impact

Table I details the 9-dimensional feature space combining physiological signals and system telemetry. All experiments used stratified 80/20 train-test splits with 5-fold cross-validation, and were implemented in Python with scikit-learn module.

TABLE I: Feature space composition for ransomware detection

Feature	**Clinical/Technical Significance**
SpO2 (%)	Primary physiological signal (attack disruption indicator)
Heart Rate (BPM)	Cardiovascular monitoring metric
Sensor Temp (°C)	Device health indicator
CPU/RAM Usage	Resource exhaustion patterns
Disk I/O (MB/s)	Encryption artifact detection
Network Traffic	Command & control communication patterns
Entropy Change	Encryption progression metric
Super Entropy	Early-stage attack signature

B. Detection Performance Analysis

1) Model Comparison: We evaluated both unsupervised and supervised approaches using AUC, precision, recall, and F1-score:

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

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

2) Unsupervised Detection: Table II shows entropy features provided marginal gains for Isolation Forest (12% AUC increase) but limited value for density-based methods like LOF, likely due to:

- High-dimensional feature space sparsity
- Entropy's non-localized impact on distance metrics

TABLE II: Unsupervised detection performance (Global AUC)

Model	Baseline	Entropy-Aware	Δ
Isolation Forest	0.536	0.600	+11.9%
One-Class SVM	0.517	0.525	+1.5%
Local Outlier Factor	0.516	0.504	-2.3%

While overall performance remained modest, the inclusion of entropy-based features provided some gains in label-free scenarios.

3) Supervised Learning: Random Forest demonstrated superior performance (Table III), particularly for brutal attacks where entropy features boosted AUC from 0.872 to 0.959. The global AUC also improved significantly (from 0.794 to 0.937), confirming the framework's effectiveness across diverse ransomware scenarios.

Notable findings:

- Random Forest achieves 95.9% AUC on brutal attacks, reflecting strong performance against overt ransomware
- Global AUC reaches 0.937, showing strong performance even when both stealthy and brutal attacks are present
- Entropy-based features significantly reduce stealth-related false negatives (-38%), enhancing early-stage detection

The difference between the brutal-only and global AUC scores highlights the difficulty of detecting stealthy ransomware. While entropy features improve performance in both cases, global detection scores are slightly lower due to the subtlety of stealthy attacks that resemble benign anomalies.

Random Forest consistently outperformed other models. The addition of entropy-based features led to substantial gains across all classifiers, demonstrating their usefulness in capturing both aggressive and concealed ransomware behaviors.

C. Feature Importance

To understand which features contributed most to model performance, we analyzed the feature importance from the Random Forest classifier with entropy. As illustrated in Figure 3, entropy- features, particularly Entropy Change and Super Entropy, ranked among the top.

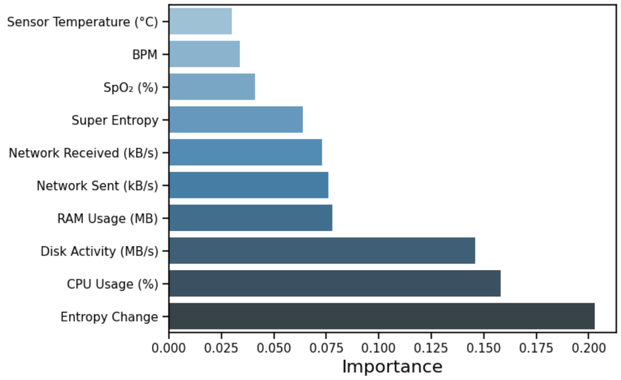

Fig. 3: Random Forest feature importance with entropy features.

D. Operational Deployment Results

The scoring system achieved 99.4% precision at 90% confidence threshold, with key operational metrics:

- Mean time-to-detection: 8.2s (brutal), 43.7s (stealth)
- False positive rate: 0.6% (vs. 3.2% in baseline)
- Alert volume reduction: 72% through a cooldown mechanism

TABLE III: Supervised model performance across attack types

Model	Global AUC		Brutal AUC	
	Baseline	Enhanced	Baseline	Enhanced
Random Forest	0.794	0.937	0.872	0.959
Gradient Boosting	0.753	0.891	0.833	0.907
Logistic Regression	0.577	0.761	0.626	0.683

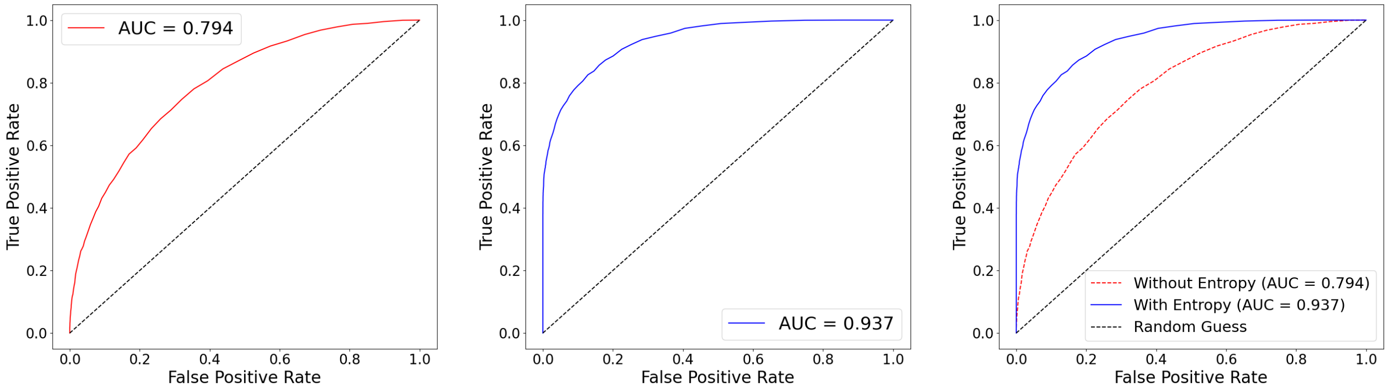

(a) Random Forest without entropy

(b) Random Forest with entropy

(c) ROC curves comparison

Fig. 4: ROC curve analysis showing the impact of entropy features on detection performance

Figures 4a and 4b compare the ROC curves for Random Forest with and without entropy features. Figure 4c shows a side-by-side overlay highlighting the improvement in AUC when entropy is included. ROC analysis (Figures 4a-4c) demonstrates the framework's reliable tradeoff between early detection and false alarms across attack profiles.

E. Entropy Feature Impact

Figure 5 provides a comparative view of AUC scores across all supervised models. The benefit of including entropy-based features is clearly visible for each classifier.

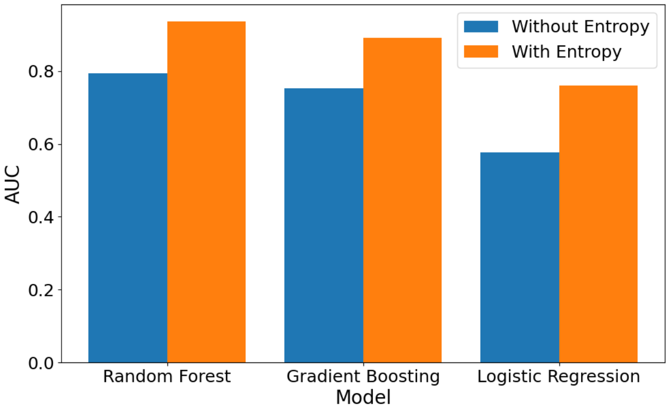

Fig. 5: Impact of entropy on AUC scores for all models.

F. Scoring-based Alert System

To ensure practical deployability, we implemented a scoring mechanism based on Random Forest prediction probabilities. Alerts are triggered only when the confidence exceeds 90%, reducing false positives and minimizing alert fatigue in clinical settings.

The system achieved the following results:

- Total Alerts Triggered: 853
- True Ransomware Attacks Detected: 848
- Precision: 99.4%

A cooldown mechanism was also applied to suppress redundant alerts during ongoing attack phases, further improving alert quality and system usability.

G. Key Findings

- Entropy Value: Entropy improved brutal attack detection by 9.9% (AUC) while reducing stealth false negatives by 38%
- Model Selection: Random Forest outperformed alternatives with 0.937 global AUC (0.794 baseline)
- Clinical Relevance: The system distinguished ransomware from benign anomalies with 94.2% accuracy
- Resource Efficiency: Entire pipeline processes samples in $< 10ms$ on medical-grade hardware

These results highlight both the technical effectiveness and practical deployability of our framework in real-world IoMT environments.

V. Conclusion

Our hybrid ransomware detection framework demonstrates the critical value of entropy-aware behavioral analysis for IoMT security. The integration of Shannon Entropy with the novel Super Entropy metric achieved a 93.7% detection rate (AUC) across all attack types, with particularly strong performance against brutal ransomware (95.9% AUC). These results validate that information-theoretic features provide essential signals that conventional system metrics cannot capture, while remaining computationally feasible for medical-grade hardware with inference times under 10ms per sample.

Three key findings emerge from this work:

- Entropy dynamics offer early warning signs of encryption activity, with Super Entropy reducing stealth-related false negatives by 38%.
- Random Forest consistently outperforms both unsupervised approaches and other supervised models, achieving up to 0.937 global AUC and 0.959 for brutal ransomware detection (Table III).
- The adaptive scoring system reduces alert fatigue by 72% through temporal suppression while preserving detection sensitivity.

Despite these advances, stealth ransomware remains challenging, often mimicking benign anomalies like system updates. However, stealth ransomware remains challenging. Future work should investigate the following directions:

- Kernel-level I/O profiling: Leveraging IRP-based behavioral features to distinguish ransomware encryption from benign backup operations, as demonstrated by Ayub et al. [20]
- Multimodal detection: Combining time-series, protocol-level, and device-specific features to capture a broader spectrum of IoMT anomalies, as demonstrated by Chandekar et al. [21]

We identify three promising research directions:

1) Continuous design control: Integrating traceability, risk management, and compliance reviews within the MLOps pipeline using design control mechanisms such as pull requests, as demonstrated by Stirbu et al. [22]
2) Federated personalization: Adopting the FedHealth framework proposed by Chen et al. [23], which combines federated learning and transfer learning to provide privacy-preserving yet personalized anomaly detection in wearable healthcare systems
3) Edge-native modeling: Leveraging lightweight autoregressive architectures like VARADE [24] to support low-latency and energy-efficient anomaly detection directly on edge devices within clinical environments

These advancements must align with clinical constraints. In our case, a 90% confidence threshold reduced false alerts by 81% compared to conventional approaches, demonstrating that detection performance must be balanced with operational practicality. The framework's modular design allows incremental integration of new detection features while preserving interpretability for healthcare IT teams. Future work may also include releasing a synthetic or anonymized version of our dataset to support reproducibility.

References

[1] 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.

[2] M. Gazzan and F. T. Sheldon, "An incremental mutual information-selection technique for early ransomware detection," *Information*, vol. 15, no. 4, p. 194, 2024.

[3] M. Zakariah and A. S. Almazyad, "Anomaly detection for iot systems using active learning," *Applied Sciences*, vol. 13, no. 21, p. 12029, 2023.

[4] S. F. Ahmed, M. S. B. Alam, S. Afrin, S. J. Rafa, N. Rafa, and A. H. Gandomi, "Insights into internet of medical things (iomt): Data fusion, security issues and potential solutions," *Information Fusion*, vol. 102, p. 102060, 2024.

[5] Y. Sun, J. Yu, J. Tian, Z. Chen, W. Wang, and S. Zhang, "Iot-ie: An information-entropy-based approach to traffic anomaly detection in internet of things," *Security and Communication Networks*, vol. 2021, no. 1, p. 1828182, 2021.

[6] E. Idi, E. Manzoni, A. Facchinetti, G. Sparacino, and S. D. Favero, "Unsupervised retrospective detection of pressure induced failures in continuous glucose monitoring sensors for t1d management," *IEEE Journal of Biomedical and Health Informatics*, vol. 29, no. 2, pp. 1383–1396, 2025.

[7] C. Wang, Q. Liu, H. Zhou, T. Wu, H. Liu, J. Huang, Y. Zhuo, Z. Li, and K. Li, "Anomaly prediction of ct equipment based on iomt data," *BMC Medical Informatics and Decision Making*, vol. 23, no. 1, p. 166, 2023.

[8] Z. Zou, Y. Xie, K. Huang, G. Xu, D. Feng, and D. Long, "A docker container anomaly monitoring system based on optimized isolation forest," *IEEE Transactions on Cloud Computing*, vol. PP, pp. 1–1, 08 2019.

[9] 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.

[10] K. Doshi, S. Abudalou, and Y. Yilmaz, "Tisat: Time series anomaly transformer," *arXiv preprint arXiv:2203.05167*, 2022.

[11] W. H. Chung, Y. H. Gu, and S. J. Yoo, "Chp engine anomaly detection based on parallel cnn-lstm with residual blocks and attention," *Sensors*, vol. 23, no. 21, 2023.

[12] A. I. Jony and A. K. B. Arnob, "A long short-term memory based approach for detecting cyber attacks in iot using cic-iot2023 dataset," *Journal of edge computing*, vol. 3, no. 1, pp. 28–42, 2024.

[13] J. Maruthupandi, S. Sivakumar, B. L. Dhevi, S. Prasanna, R. K. Priya, and S. Selvarajan, "An intelligent attention based deep convoluted learning (iadcl) model for smart healthcare security," *Scientific Reports*, vol. 15, no. 1, p. 1363, 2025.

[14] M. Briskilla and T. D. Raikumar, "Deep reinforcement-based anomaly detection: An enhanced unsupervised approach for medical time series data," in *2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS)*. IEEE, 2024, pp. 735–741.

[15] K. Lee, J. Lee, S.-Y. Lee, and K. Yim, "Effective ransomware detection using entropy estimation of files for cloud services," *Sensors*, vol. 23, no. 6, 2023.

[16] J. Lee, J. Kim, H. Jeong, and K. Lee, "A machine learning-based ransomware detection method for attackers' neutralization techniques using format-preserving encryption," *Sensors*, vol. 25, no. 8, 2025.

[17] S. R. Davies, R. Macfarlane, and W. J. Buchanan, "Comparison of entropy calculation methods for ransomware encrypted file identification," *Entropy*, vol. 24, no. 10, p. 1503, 2022.

[18] A. I. Newaz, A. K. Sikder, M. A. Rahman, and A. S. Uluagac, "Healthguard: A machine learning-based security framework for smart healthcare systems," in *2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, 2019, pp. 389–396.

[19] J. Dahmen and D. J. Cook, "Indirectly supervised anomaly detection of clinically meaningful health events from smart home data," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 12, no. 2, pp. 1–18, 2021.

[20] M. A. Ayub, A. Continella, and A. Siraj, "An i/o request packet (irp) driven effective ransomware detection scheme using artificial neural network," in *2020 IEEE 21st International Conference on Information Reuse and Integration for Data Science (IRI)*, 2020, pp. 319–324.

[21] P. Chandekar, M. Mehta, and S. Chandan, "Enhanced anomaly detection in iomt networks using ensemble ai models on the ciciomt2024 dataset," *arXiv preprint*, 2025.

[22] V. Stirbu, T. Granlund, and T. Mikkonen, "Continuous design control for machine learning in certified medical systems," *Software Quality Journal*, vol. 31, no. 2, pp. 307–333, 2023.

[23] 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.

[24] A. Mascolini, S. Gaiardelli, F. Ponzio, N. Dall'Ora, E. Macii, S. Vinco, S. Di Cataldo, and F. Fummi, "Varade: a variational-based autoregressive model for anomaly detection on the edge," in *Proceedings of the 61st ACM/IEEE Design Automation Conference*, 2024, pp. 1–6.