Machine Learning and Temporal Convolutional Networks for Enhanced Ransomware Detection in IoMT Environments

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

***Abstract*—Ransomware attacks targeting Internet of Medical Things (IoMT) devices are an escalating threat to modern healthcare environments, where operational continuity and patient safety are mission-critical. Existing anomaly detection approaches—ranging from signature-based techniques to classical machine learning algorithms—often fall short when confronted with gradual or concealed attack behaviors. In this paper, we propose a hybrid detection framework that couples traditional supervised learning with temporal sequence modeling. Initially, a machine learning model (Random Forest or XGBoost) is trained on telemetry data enriched with crafted features such as risk signal scores and alert amplitudes. The resulting per-sample risk probabilities are then treated as temporal signals and fed into a Temporal Convolutional Network (TCN) to capture contextual progression and subtle threat dynamics. Experiments conducted on a synthetic IoMT dataset—emulating both rapid-impact and low-profile ransomware—demonstrate that the proposed approach achieves 0.9024 accuracy and 0.8441 AUC with XGBoost+TCN, and 0.9124 accuracy with 0.8223 AUC using Random Forest + TCN. Both configurations outperform their respective baselines (XGBoost: 0.8218 accuracy, 0.7509 AUC; Random Forest: 0.8695 accuracy, 0.7574 AUC). These findings confirm that incorporating temporal modeling of risk scores significantly improves ransomware detection, while remaining suitable for deployment in medical-grade environments.**

***Index Terms*—Internet of Medical Things (IoMT), Ransomware Detection, Machine Learning Models, Temporal Convolutional Network (TCN), XGBoost, Random Forest, Hybrid Detection Framework, Risk Scoring, Time-Aware Analysis, Feature Engineering, Sequence Modeling, Medical Cybersecurity, Anomaly Detection, Embedded AI**

I. Introduction

The Internet of Medical Things (IoMT) has transformed modern healthcare by enabling real-time data acquisition from interconnected medical devices such as infusion pumps, patient monitors, and diagnostic systems. While this connectivity has enhanced clinical decision-making and operational efficiency, it has also introduced new cybersecurity vulnerabilities. Among them, ransomware attacks are particularly disruptive, with the potential to encrypt critical patient data or disable safety-critical systems—directly endangering both care delivery and patient safety.

Recent ransomware campaigns have shifted from brute-force infections to more targeted and concealed strategies. These evolving behaviors include gradual encryption routines, minimal system footprint, and evasive execution, making traditional detection methods less effective. Within IoMT environments, this challenge is amplified by several constraints: limited device computational power, variability in normal telemetry patterns, and the presence of transient anomalies unrelated to attacks. Classic signature-based tools and rule-based systems are often blind to such subtle or evolving threats. Even supervised machine learning (ML) models, although more adaptive, often operate on isolated telemetry snapshots and rely on instantaneous metrics such as CPU usage or memory load, thereby overlooking long-term behavioral transitions or multi-step escalation patterns.

In this paper, we advocate for a shift toward temporal-aware detection strategies. We propose a hybrid two-phase architecture that first applies a classical ML classifier (XGBoost or Random Forest) to assign a probability of maliciousness to each telemetry instance, based on engineered features including *risk signal score* and *alert amplitude*. These risk probabilities are then restructured as a time series and processed by a Temporal Convolutional Network (TCN), a deep learning model well-suited for extracting temporal dependencies and uncovering evolving attack signatures. This architecture enables the system to detect not only abrupt attacks but also progressive malicious behaviors that manifest across time.

To evaluate our framework, we created a synthetic but realistic IoMT telemetry dataset that simulates a BD Alaris™ infusion pump operating under three operational contexts: normal usage, benign anomalies, and ransomware infection. The benign anomalies represent non-malicious but irregular system states (e.g., temporary high disk usage or CPU load due to software updates or maintenance), which can resemble ransomware-like behavior and thus pose a detection challenge. Because access to real medical telemetry is highly restricted due to privacy and regulatory constraints, simulation-based evaluation is a common alternative. Following the simulation-based approach proposed in [1], we designed this dataset to reflect plausible operational and adversarial profiles of medical devices without relying on clinical environments. This simulation-based design enables rigorous testing while preserving patient safety and data confidentiality.

The main contributions of this paper are as follows:

- We introduce a hybrid detection framework that integrates supervised ML classification with temporal modeling via

a TCN.
- We demonstrate that feeding the model-generated per-sample risk scores into the TCN allows the capture of temporal threat dynamics beyond what static classifiers can achieve.
- We validate our approach on a telemetry dataset simulating realistic IoMT attack conditions, and show improved performance metrics over baseline classifiers.

The remainder of this paper is organized as follows: section II discusses related work on IoMT cybersecurity and time-series anomaly detection. section III details our detection framework. section IV presents the evaluation setup and metrics. section V concludes the paper and outlines future directions.

II. Related Work

The growing sophistication of cyber threats targeting the IoMT has transformed anomaly detection from a reactive security task into a critical research problem in proactive threat modeling. Initial approaches relied heavily on static heuristics, such as threshold-based filters and handcrafted signal analysis, to detect unusual device behavior. Statistical tools including moving averages, gradient shifts, and extrema detection were applied to biomedical sensors with moderate success [2]. However, these techniques struggle to detect polymorphic or stealth ransomware that mimics benign system patterns. Furthermore, rule-based logic, while interpretable, lacks the adaptability required to handle dynamic telemetry changes in realistic IoMT deployments [3].

To address these shortcomings, many researchers have turned to ML models capable of capturing complex decision boundaries in high-dimensional, multivariate data. Supervised methods like Random Forests and Support Vector Machines offer high classification accuracy when trained on well-labeled datasets, but they are often sensitive to class imbalance, real-time fluctuations, and the absence of temporal continuity. Unsupervised strategies, such as Isolation Forest, have proven more scalable and lightweight. For instance, El Khairi et al. [4] proposed a contextual system call analysis approach tailored for anomaly detection in containerized IoT systems. Similarly, Liu et al. [5] introduced a change-point detection method for non-stationary medical data, highlighting the importance of tracking behavioral drift in long-term monitoring. Nevertheless, most of these approaches operate on static snapshots, which limits their ability to detect attacks that evolve gradually or span across time.

To address these temporal limitations, researchers have investigated sequence-based learning techniques. LSTM networks and Transformer variants such as TiSAT [6] are effective at modeling long-term dependencies in sequential data, but are often constrained by high latency, energy demands, and memory consumption—factors incompatible with embedded IoMT environments. CNN-LSTM hybrids [7] provide good generalization performance on high-frequency data streams, yet they typically suffer from tuning overhead, lack of interpretability, and a heavier inference footprint. In contrast, TCNs offer a lightweight, parallelizable architecture with competitive results on time-series data. TCNs have been deployed successfully in multivariate anomaly detection for IoT [8], semi-supervised event classification [9], and DDoS mitigation in medical network traffic [10]. Wang et al. [11] extended their use to log-based detection, while Mulia et al. [12] coupled TCNs with attention for interpretable time-series modeling.

Entropy-based anomaly detection has also been explored. Lee et al. [13] proposed an entropy estimation method to detect malicious file encryption in cloud services, aiming to prevent the synchronization of ransomware-infected files. While Davies et al. [14] demonstrated that entropy-based methods can produce false positives on legitimate high-entropy formats such as DICOM or HL7, they also showed how combining multiple entropy tests can improve ransomware detection. These structural analyses, while orthogonal to telemetry modeling, offer complementary insights when paired with behavioral features. Lightweight learning strategies such as model pruning or distributed learning (e.g., FedHealth [15]) have also been proposed to support inference on resource-constrained devices, but often lack the expressiveness to capture escalating threat patterns.

Probabilistic profiling has been proposed as a complementary detection layer. The study by [16] explored the idea of modeling threat behavior through probabilistic code anomaly scoring. However, this typically occurs at a static code or execution trace level. In contrast, very few studies have considered the notion of using ML-generated risk scores—such as per-sample probabilities—as temporal inputs for downstream models. When probabilistic outputs are used in some traditional models, they are typically treated as static scalar features embedded in a larger input vector, thereby discarding their temporal evolution. Conversely, models like that of Khan et al. [17] leverage RNNs to maintain sequence awareness, but do not explicitly model the evolution of risk scores over time.

Hybrid modular designs have been proposed to improve robustness. For instance, Almotiri [18] proposed a dual-layer static/dynamic ML pipeline for SD-IoMT infrastructure. Alzakari et al. [19] employed attention-based RNNs to detect early-stage ransomware in traffic flows. Berguiga et al. [20] suggested a hybrid deep learning model for IoMT threat detection, though their system lacked inter-sample sequence modeling. While these architectures represent steps toward multi-phase detection, they still lack temporal reasoning capabilities over predictive uncertainty or risk signals produced by upstream classifiers.

To the best of our knowledge, no prior work has proposed a ransomware detection pipeline that explicitly treats ML-derived per-sample risk probabilities as a structured time series input for a Temporal Convolutional Network. Our proposed method fills this gap by converting predictive outputs from classical models into sequential inputs, enabling the TCN to learn behavioral transitions and risk accumulation over time. This layered architecture supports anticipatory threat detection while maintaining compatibility with real-world IoMT deployments that are constrained by latency, memory, and power requirements. In doing so, our approach provides a deployable and interpretable solution for medical device cybersecurity, and establishes a

foundation for real-time, context-aware ransomware mitigation in constrained medical environments.

III. Proposed Approach

Our goal is to construct a ransomware detection framework capable of capturing both abrupt and stealthy patterns in telemetry generated by IoMT devices. To this end, we adopt a modular, two-phase architecture that combines a supervised ML classifier with a TCN. Figure 1 illustrates the complete detection pipeline, which unfolds in sequential stages: (A) Dataset Preparation and Feature Engineering, (B) Risk Probability Estimation, (C) Time-Series Structuring, (D) TCN-Based Threat Modeling, and (E) Pipeline Feature Traceability.

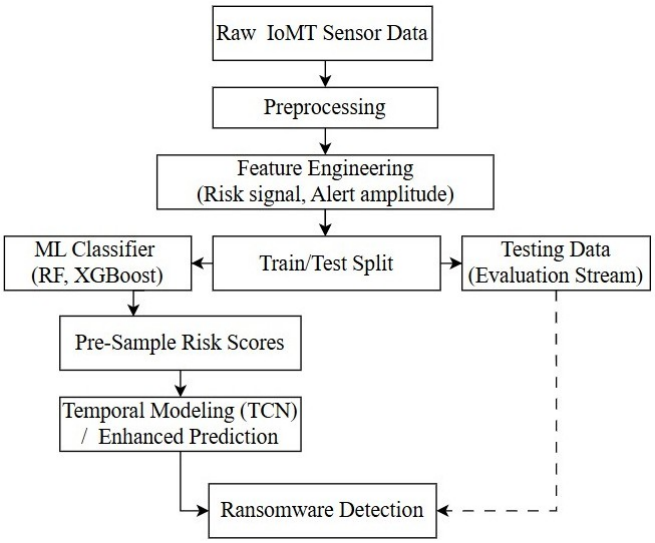

Fig. 1: Proposed Hybrid Detection Pipeline for IoMT Ransomware

A. Dataset Preparation and Feature Engineering

We define our dataset $\mathcal{D} = (x_i, y_i)_{i=1}^{N}$, where each $x_i \in \mathbb{R}^d$ is a telemetry snapshot composed of system metrics extracted from an IoMT device, and $y_i \in 0, 1$ is a binary label indicating whether the sample is benign or corresponds to ransomware behavior. The raw features include CPU usage, RAM consumption, and disk I/O levels—standard indicators of system workload and responsiveness.

To complement these low-level signals, we incorporate two composite features designed to highlight abnormal behavior patterns:

- Risk Signal Score: a continuous metric derived to reflect subtle behavioral deviations over time.
- Alert Amplitude: a feature sensitive to abrupt changes in system state, often indicative of ransomware-triggered anomalies.

All features are standardized using z-score normalization:

$$x'_{i,j} = \frac{x_{i,j} - \mu_j}{\sigma_j} \tag{1}$$

where μ_j and σ_j represent the empirical mean and standard deviation of feature j over the training set. This normalization ensures consistent feature scaling and facilitates effective model training downstream.

These enriched representations serve as input to the ML-based classifier and help improve the detection of both rapid and progressive attack behaviors.

B. Risk Probability Estimation via ML Classifier

We train a supervised ML classifier f_{ML}, such as XGBoost or Random Forest, on the engineered feature vectors $x_i \in \mathcal{D}$. The model outputs a probability score P_i indicating the likelihood of ransomware activity for each input:

$$P_i = f_{\text{ML}}(x_i), \quad P_i \in [0, 1] \tag{2}$$

These scores reflect soft confidence levels rather than hard classifications, and are subsequently used as temporal input signals in the next modeling phase. This intermediate representation enables the system to capture uncertainty and subtle patterns that may only become meaningful when analyzed across time.

To mitigate dataset imbalance, particularly the underrepresentation of ransomware events, we apply class-weight scaling during training:

$$\text{scale\_pos\_weight} = \frac{N_0}{N_1} \tag{3}$$

where N_0 and N_1 denote the number of benign and ransomware-labeled samples, respectively. This correction ensures that the model remains sensitive to rare but high-impact ransomware indicators during training.

C. Time-Series Structuring of Risk Scores

The sequence of ML-derived probabilities $P_1, P_2, \ldots, P_T$ is treated as a univariate time-series signal S, representing the evolving likelihood of ransomware activity over time. To prepare this data for sequential modeling, we segment S into overlapping sliding windows of fixed length w:

$$\mathbf{X}_t = [P_{t-w+1}, P_{t-w+2}, \ldots, P_t] \tag{4}$$

Each window $\mathbf{X}_t$ serves as an input to the temporal model, with the label y_t derived from the ground truth associated with the most recent element in the window. This formulation allows the model to learn not just from isolated risk scores, but from their progression and accumulation across time—key for detecting stealthy or slowly escalating threats.

D. Temporal Convolutional Modeling with TCN

The TCN processes each probability window $\mathbf{X}_t$ to produce a soft prediction $\hat{y}_t$ (where $\hat{y}_t \in [0, 1]$ means predicted ransomware probability) indicating the likelihood of ransomware presence:

$$\hat{y}_t = \text{TCN}(\mathbf{X}_t; \theta) \tag{5}$$

Here, θ denotes the learnable parameters of the network. The TCN architecture is composed of 1D dilated causal convolutions, which ensure that predictions at time t are only influenced

by past values, thus respecting the causal nature of time-series forecasting. Dilation allows the model to expand its receptive field without increasing depth, enabling it to learn long-term dependencies while maintaining computational efficiency.

Training is performed using the binary cross-entropy loss function:

$$\mathcal{L} = -\sum_{t=1}^{T'} [y_t \log(\hat{y}_t) + (1 - y_t) \log(1 - \hat{y}_t)] \tag{6}$$

We adopt the following hyperparameters in our implementation:

- `input_chunk_length = 12`: each input sequence includes 12 consecutive probability scores.
- `output_chunk_length = 4`: the model forecasts the next 4 risk values.
- `n_epochs = 8`: the model is trained over 8 complete passes through the dataset.
- `dropout = 0.35`: dropout is applied between layers to reduce overfitting.

This configuration strikes a balance between context depth and runtime efficiency, making the system suitable for resource-constrained IoMT devices.

E. Pipeline Feature Traceability

To promote transparency and reproducibility, we summarize in Table I the role of each key feature throughout the pipeline. This mapping clarifies the origin, transformation, and final usage of both raw and engineered signals within the detection architecture.

TABLE I: Traceability of Features Across the Detection Pipeline

Feature	Role in Pipeline
CPU, RAM, Disk I/O	Input features for ML classifier
Risk Signal Score	Augments ML classifier with semantic risk context
Alert Amplitude	Highlights anomaly spikes relevant to ransomware
P_i (ML probability)	Structured into sequences for TCN input
$\hat{y}_t$ (TCN output)	Final binary ransomware decision

IV. Experimental Results

We evaluate the effectiveness of our proposed detection framework on a custom-simulated IoMT telemetry dataset designed to reflect realistic operational and adversarial scenarios. Our objective is to quantify the added value of temporal modeling (via TCN) in enhancing detection performance, especially under stealth ransomware scenarios.

A. Dataset Description and Preprocessing

Our dataset simulates telemetry from a BD Alaris™ infusion pump, structured around three behavioral modes: *normal*, *benign anomaly*, and *ransomware*. Normal samples represent routine device behavior. Benign anomalies include non-malicious anomalies such as firmware updates or high CPU activity, while the ransomware category comprises both overt and stealth variants.

Figure 2 illustrates the distribution of event types.

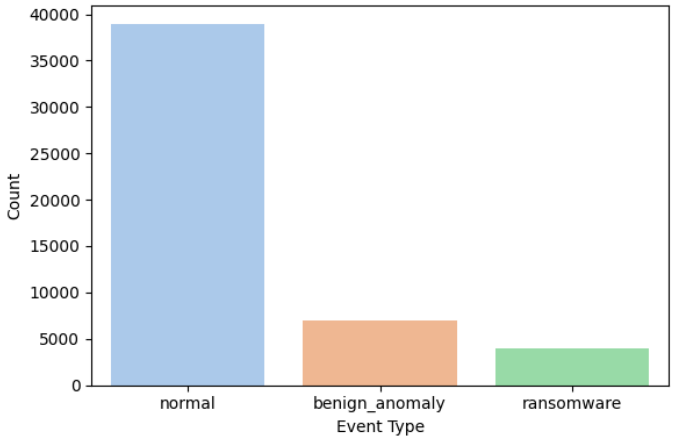

Fig. 2: Distribution of behavioral categories in the dataset.

The final dataset includes approximately 50,000 labeled instances and was split into 80% training and 20% test sets. All features were normalized using z-score scaling. To mitigate class imbalance, we used weighted loss functions during classifier training.

B. ROC Curve Analysis

To assess the impact of temporal modeling, we first analyze the ROC curves generated by the baseline and TCN-enhanced models. Figure 3a shows the ROC curve of the XGBoost classifier without temporal context. Figure 3b displays the ROC curve of the same model augmented with TCN-based sequence modeling. Finally, Figure 3c directly compares both curves to visualize the improvement in detection performance.

As shown, the TCN-enhanced model achieves a higher true positive rate across all thresholds, leading to a noticeable gain in AUC. This confirms that temporal modeling helps capture latent risk escalation and delayed activation patterns—especially valuable when detecting ransomware variants that evade static inspection.

C. Model Training and Evaluation Metrics

To highlight the effectiveness of our hybrid framework, we trained two widely-used classical classifiers—Random Forest and XGBoost—on the engineered telemetry feature vectors. These models act as the first stage of our pipeline and produce a per-sample probability score P_i, indicating the likelihood that a given telemetry snapshot is associated with ransomware activity.

These outputs serve two distinct purposes:

1) They are directly thresholded (typically at 0.5) to produce static binary predictions, forming the baseline evaluation.
2) They are restructured as a temporal sequence and passed into the TCN, enabling dynamic modeling of risk evolution.

The following evaluation metrics were used to quantify detection performance:

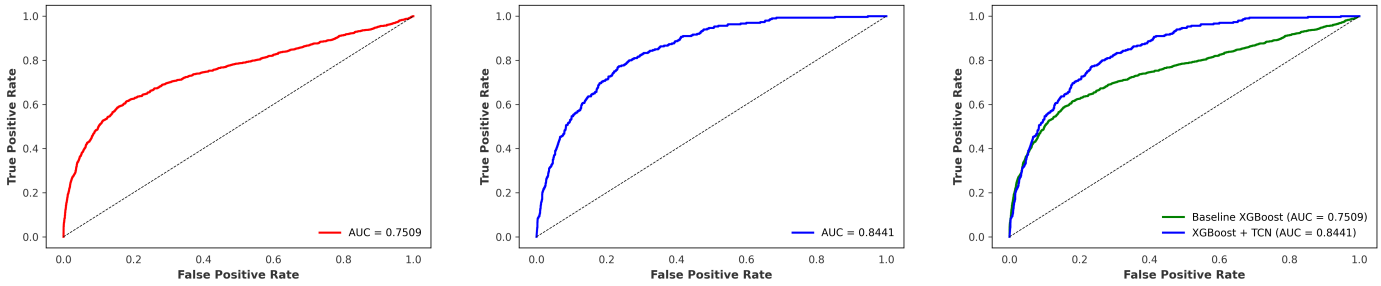

(a) XGBoost baseline (b) XGBoost + TCN (c) Comparative ROC overlay

Fig. 3: ROC analysis: baseline vs. TCN-enhanced detection.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \qquad \text{F1} = \frac{2TP}{2TP + FP + FN}$$

$$\text{AUC} = \int_0^1 \text{ROC}(t)dt$$

Here, TP (True Positives) denotes correctly identified ransomware samples, TN (True Negatives) refers to correctly identified benign samples, FP (False Positives) represents benign samples incorrectly flagged as ransomware, and FN (False Negatives) captures missed ransomware detections.

- Accuracy captures the overall correctness of the model across both classes.
- F1-Score balances precision and recall, making it a crucial metric in imbalanced settings where ransomware events are rare.
- AUC (Area Under the ROC Curve) measures the model's ability to rank positive samples above negative ones over varying decision thresholds. A higher AUC indicates better separability between the two classes.

Table II reports the results across four model configurations: the two static ML baselines, and their TCN-enhanced counterparts. Notably, both XGBoost and Random Forest achieve significant improvements when augmented with TCN, particularly in terms of F1-score, which nearly doubles, highlighting improved recall and reduced false negatives.

TABLE II: Performance Comparison Across Baselines and TCN-Enhanced Models

Method	Accuracy	F1-Score	AUC
XGBoost Baseline	0.8218	0.4371	0.7509
Random Forest Baseline	0.8695	0.2939	0.7574
XGBoost + TCN	0.9024	0.6402	0.8441
Random Forest + TCN	0.9124	0.6525	0.8223

D. Classifier-Wise AUC Comparison

To highlight the benefit of temporal modeling across classifiers, we present a comparative AUC barplot in Figure 4. Both Random Forest and XGBoost show considerable gains when paired with TCNs.

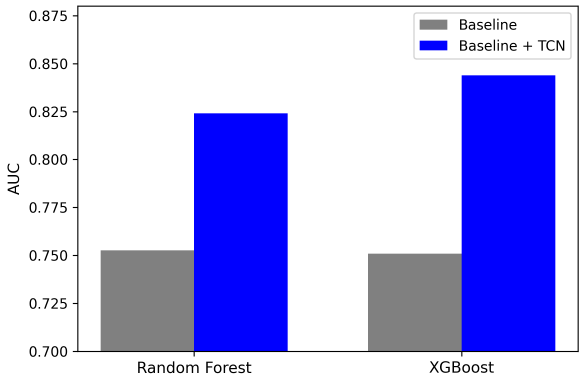

Fig. 4: Comparative AUC scores: baseline vs. TCN-enhanced classifiers.

This comparison confirms that risk probability trajectories, when modeled temporally, provide complementary signals that static classifiers fail to leverage.

E. Feature Importance and Interpretation

To better understand the behavior of the ML classifiers, we analyzed feature importance from the Random Forest model. As seen in Figure 5, composite features like Alert Amplitude and Risk Signal Score were ranked highest—confirming their relevance in identifying ransomware behavior.

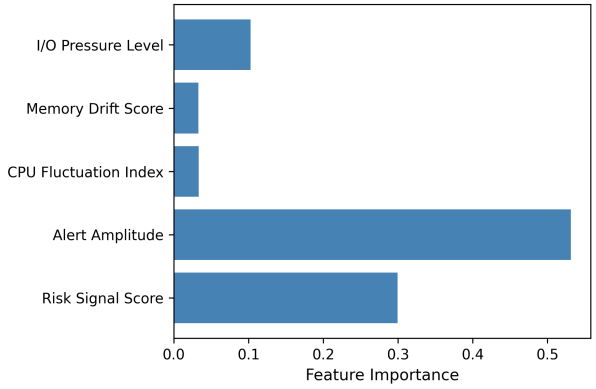

Fig. 5: Random Forest feature importance for ransomware classification.

F. Key Observations

- Temporal modeling improves AUC by over 9% for XGBoost, and nearly 7% for Random Forest.
  This demonstrates that even strong tree-based models benefit from sequential context, especially when subtle behavioral transitions occur.
- The TCN-enhanced model exhibits better separation between benign and malicious sequences across all thresholds.
  This is evident in the ROC curves, where the TCN consistently shifts the decision boundary toward higher true positive rates.
- Comparative curves indicate higher recall and reduced false negatives in stealth ransomware cases.
  This is critical in clinical environments, where undetected threats can escalate silently and jeopardize patient safety.
- Performance gains come with minimal computational overhead, preserving real-time inference viability. The architecture remains deployable on constrained medical devices.

V. Conclusion

The increasing frequency and sophistication of ransomware attacks on IoMT devices call for detection strategies that go beyond static inspection and per-sample anomaly classification. In this paper, we presented a hybrid detection pipeline that combines classical supervised learning with temporal sequence modeling via a TCN.

Our approach begins by training an ML classifier (e.g., XGBoost or Random Forest) to generate per-sample risk probabilities from telemetry signals. These probabilities are then structured as a univariate time series and passed to the TCN, which captures temporal dynamics indicative of progressive or stealthy threats. This two-stage architecture is modular, interpretable, and lightweight—making it suitable for real-time use on constrained medical hardware.

Experiments on a custom synthetic dataset simulating BD Alaris™ infusion pump telemetry demonstrated substantial performance improvements. The XGBoost + TCN configuration achieved an F1-score of 0.6402 and an AUC of 0.8441, outperforming static baselines by a wide margin. These results confirm that temporal reasoning over risk trajectories enhances detection sensitivity, particularly in subtle attack scenarios.

Unlike deep end-to-end pipelines, our method maintains modularity by decoupling feature-based scoring from temporal inference. This not only improves interpretability but also allows for flexible adaptation to heterogeneous IoMT platforms. In addition, the probabilistic interface between the ML and TCN layers enables plug-and-play substitution of components, opening the door to rapid experimentation or domain-specific tuning.

Future work includes extending the pipeline to multivariate temporal modeling—leveraging raw system metrics in parallel with risk scores—and exploring alternative temporal architectures such as Transformers. Validating the framework on real-world hospital telemetry remains an important next step toward clinical-grade deployment. This also implies addressing challenges such as data labeling scarcity and strict compliance with healthcare data protection standards (e.g., HIPAA or GDPR).

In summary, this work introduces a scalable, interpretable, and effective solution for ransomware detection in IoMT ecosystems, demonstrating that modeling risk progression over time can significantly improve threat visibility without compromising deployability.

References

[1] L. Aversano, M. L. Bernardi, M. Cimitile, D. Montano, R. Pecori, L. Veltri *et al.*, "Anomaly detection of medical iot traffic using machine learning." in *DATA*, 2023, pp. 173–182.

[2] Y. Zhu, S. Zhao, Y. Zhang, C. Zhang, and J. Wu, "A review of statistical-based fault detection and diagnosis with probabilistic models," *Symmetry*, vol. 16, no. 4, p. 455, 2024.

[3] F. Hussain, R. Hussain, S. A. Hassan, and E. Hossain, "Machine learning in iot security: Current solutions and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 3, pp. 1686–1721, 2020.

[4] A. El Khairi, M. Caselli, C. Knierim, A. Peter, and A. Continella, "Contextualizing system calls in containers for anomaly-based intrusion detection," in *Proceedings of the 2022 on Cloud Computing Security Workshop*, 2022, pp. 9–21.

[5] Y. Liu *et al.*, "Anomaly and change point detection for time series with concept drift," *World Wide Web*, vol. 26, pp. 3229–3252, 2023.

[6] K. Doshi *et al.*, "Tisat: Time series anomaly transformer," *arXiv preprint arXiv:2203.05167*, 2022.

[7] H. C. Altunay and Z. Albayrak, "A hybrid cnn+lstm-based intrusion detection system for industrial iot networks," *Engineering Science and Technology, an International Journal*, vol. 38, p. 101322, 2023.

[8] H. Yu and L. Zhang, "Dtaad: Dual tcn-attention networks for anomaly detection in multivariate time series," *Knowledge-Based Systems*, vol. 275, pp. 110–122, 2023.

[9] Y. Xu and Y. Cheng, "Semi-supervised variational tcn for multi-anomaly detection in iot," *arXiv preprint arXiv:2104.01813*, 2021.

[10] Z. Wang, Z. Guan, X. Liu, C. Li, X. Sun, and J. Li, "Sdn anomalous traffic detection based on temporal convolutional network," *Applied Sciences*, vol. 15, no. 8, p. 4317, 2025.

[11] Z. Wang, J. Tian, H. Fang, L. Chen, and J. Qin, "Lightlog: A lightweight temporal convolutional network for log anomaly detection on the edge," *Computer Networks*, vol. 203, p. 108616, 2022.

[12] M. A. Mulia, M. B. Bahy, M. Z. F. N. Siswantoro, N. R. D. Riyanto, N. R. Sudianjaya, and A. A. A. Shiddiqi, "Kbjnet: Kinematic bi-joint temporal convolutional network attention for anomaly detection in multivariate time series data," *Data Science Journal*, vol. 23, no. 1, pp. 1–22, 2024.

[13] K. Lee *et al.*, "Effective ransomware detection using entropy estimation of files," *Sensors*, vol. 23, no. 6, p. 3023, 2023.

[14] 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.

[15] M. Chen *et al.*, "Fedhealth: A federated transfer learning framework for wearable healthcare," *IEEE Network*, vol. 34, no. 4, pp. 16–23, 2020.

[16] R. Leisner, S. Kensington, L. Abernathy, and D. MacAllister, "Ransomware detection through probabilistic code anomaly profiling," 2024.

[17] N. W. Khan, M. S. Alshehri, M. A. Khan, S. Almakdi, N. Moradpoor, A. Alazeb, S. Ullah, N. Naz, and J. Ahmad, "A hybrid deep learning-based intrusion detection system for iot networks," *Mathematical Biosciences and Engineering*, vol. 20, no. 8, pp. 13 491–13 520, 2023.

[18] J. Almotiri, "Ai-driven iomt security framework for advanced malware and ransomware detection," *Journal of Cloud Computing*, vol. 14, no. 1, pp. 1–15, 2025.

[19] S. A. Alzakari, M. Aljebreen, N. Ahmad, A. A. Alhashmi, S. Alahmari, O. Alrusaini, A. M. Al-Sharafi, and W. S. Almukadi, "An intelligent ransomware based cyberthreat detection model using multi-head attention-based recurrent neural networks with optimization algorithm in iot environment," *Scientific Reports*, vol. 15, no. 1, p. 8259, 2025.

[20] A. Berguiga, A. Harchay, and A. Massaoudi, "Hids-iomt: A deep learning-based intelligent intrusion detection system for the internet of medical things," *IEEE Access*, vol. PP, pp. 1–1, 2025.