applied sciences

MDPI

Article

Handling Imbalanced IoMT Network Data for Intrusion Detection via PCA and One-Class SVM

Eren Gencturk, Beste Ustubioglu * and Guzin Ulutas

Department of Computer Engineering, Karadeniz Technical University, Trabzon 61080, Türkiye; ktu@erengencturk.com.tr (E.G.)
* Correspondence: bustubioglu@ktu.edu.tr

Abstract

The Internet of Medical Things (IoMT) has become integral to modern healthcare, yet its always-connected and resource-constrained nature enlarges the attack surface and complicates timely intrusion detection. This study presents a deployment-oriented, two-stage anomaly-detection pipeline. First, Principal Component Analysis (PCA) is employed to reduce the dimensionality of network traffic data, capturing the most significant variance. Subsequently, a One-Class Support Vector Machine (OC-SVM) is trained exclusively on these principal components of normal traffic. This approach prioritizes computational efficiency for resource-constrained IoMT devices while maintaining high model robustness. By modeling the principal components of normal behavior, our method achieves state-of-the-art performance across diverse attack families. We adopt a uniform protocol across four public IoMT corpora—BoT-IoT, CICIoMT2024, ECU-IoHT, and IoMT-TrafficData. The model's hyperparameters, including the optimal number of principal components determined by explained variance, are tuned via randomized search. Despite using no attack labels during training, the proposed PCA-enhanced detector achieves state-of-the-art performance across diverse attack families: on BoT-IoT we obtain 99.92% F1-score (99.84% accuracy), on CICIoMT2024 we obtain 99.88% F1-score (99.77% accuracy), on ECU-IoHT 99.25% F1-score (98.58% accuracy), and on IoMT-TrafficData 99.19% F1-score (98.66% accuracy). The compact model size, enabled by PCA, makes the approach highly amenable to edge or gateway deployment in clinical networks, while the normal-only training paradigm improves robustness to zero-day threats. The results demonstrate that modeling the principal components of routine network behavior is a highly effective and efficient strategy for reliable, low-latency threat detection in realistic IoMT settings.

Keywords: Internet of Medical Things (IoMT); one-class support vector machine; principal component analysis; anomaly detection; intrusion detection; class imbalance; cybersecurity; healthcare

Academic Editor: Paolo Renna

Received: 31 December 2025
Revised: 26 January 2026
Accepted: 29 January 2026
Published: 9 May 2026

1. Introduction

IoT (Internet of Things) refers to outfitting physical objects—such as devices, vehicles, household appliances, and other artifacts—with sensors, software, and connectivity so they can exchange data with one another over the internet [1]. Through this capability, the everyday devices we rely on are able to gather and analyze information and then act intelligently based on those insights [2].

The IoT landscape comprises several fundamental layers: devices and sensors, communication/connectivity technologies, data-processing backends, and user-facing interfaces [3]. Sensors capture environmental phenomena and translate them into digital signals,

while communication mechanisms convey these signals across networks. Within this ecosystem, protocols such as MQTT, CoAP, and LoRaWAN address different requirements [4]. For instance, MQTT (Message Queuing Telemetry Transport) is favored for its lightweight design and modest bandwidth demands, making it common in use cases that need rapid transmission, including industrial sensing or real-time monitoring. By contrast, the Constrained Application Protocol (CoAP) targets devices with limited energy and compute resources (e.g., wearables or smart-home nodes); it mirrors HTTP semantics but with improved efficiency for constrained contexts [5]. LoRaWAN, meanwhile, is distinguished by long-range (up to roughly 10 km) and low-power operation, which suits wide-area deployments like smart-city installations (environmental sensing, street lighting) and large-scale agricultural telemetry [6]. Complementary connectivity options—Wi-Fi, Bluetooth, Zigbee, and 5G—also play key roles in device-to-device and device-to-cloud interaction. Collected data is processed by various algorithms to yield actionable information, and users engage with these systems through dedicated interfaces [7].

Today, IoT technologies span numerous domains. In residential settings, smart thermostats, surveillance cameras, and lighting solutions enhance convenience and efficiency. In healthcare, IoT-enabled wearables provide continuous observation of an individual's health status, contributing to more efficient and personalized care [8]. Among the principal advantages of IoT are increased automation (reducing human error), remote supervision and control, optimized resource utilization, expedited workflows, and broader support for data-driven decision-making. It is estimated that approximately 18.8 billion IoT endpoints are currently online worldwide, with projections suggesting rapid growth to 30.9 billion by the end of 2025 [9]. Within healthcare specifically, IoT device adoption is accelerating, with forecasts indicating the installed base will approach 1 billion by 2025. Driven by expanding application areas and rising investment, the overall IoT market size is also expected to surpass 1 trillion USD by 2025 [7]. While such momentum underscores IoT's central role in digital transformation for individuals and organizations alike, significant challenges persist—chief among them cybersecurity and personal-data privacy. Continuous, internet-based data exchange enlarges the attack surface. Furthermore, inconsistent standards across manufacturers can hinder interoperability, creating integration and compatibility issues. These obstacles are widely recognized as core barriers to broader IoT adoption [10].

IoT devices are pivotal in healthcare delivery. A prime use case is patient monitoring: tracking diverse physiological signals is essential in multi-modal treatment plans, and continuous observation can improve therapeutic outcomes. IoMT systems can function as physiological monitors that are straightforward to deploy and able to furnish clinicians with critical, timely information; they also afford mobility and ease of operation in hospital settings [8]. Another key application is smart medication delivery, where dosing prescribed by clinicians is automated. Infusion pumps, for example, are increasingly prevalent in medical facilities, providing precise and reliable administration. Telemedicine, which has risen sharply in popularity, enables patients to be treated and followed remotely—improving overall care by offering comfort and continuity. In such scenarios, IoMT devices act as sensors to collect salient health data and can provide localized services that reduce the need for frequent in-person visits.

However, the broad deployment of these systems introduces a multifaceted cybersecurity risk landscape. Attacks targeting IoT/IoMT infrastructures range from Smurf assaults and ARP spoofing to MQTT-based DoS/DDoS campaigns and a variety of reconnaissance activities [11]. In a Smurf attack—an amplification-style DDoS—ICMP echo requests are broadcast so that numerous hosts respond to the victim's IP, saturating bandwidth and potentially disrupting critical healthcare services. ARP spoofing, in contrast, injects

forged ARP responses to insert an adversary between communicating parties, enabling interception or modification of sensitive information such as patient data.

MQTT-specific threats include Connect Flood and Publish Flood attacks, wherein adversaries inundate brokers with spurious connection attempts or published messages, exhausting system resources and degrading device availability [12]. MQTT-Malformed Data attacks—sending syntactically incorrect packets—can further destabilize or crash resource-constrained medical devices. Reconnaissance operations, while not immediately destructive, set the stage for later intrusions [13]. Typical methods include port scans to enumerate open services, OS fingerprinting to identify platforms in use, vulnerability scans to pinpoint weaknesses, and ping sweeps to map active hosts—each step sharpening the precision of subsequent targeted attacks.

Given this layered threat environment, it is essential to harden IoMT solutions at both hardware and software tiers. Strong cryptographic protections, continuous monitoring of network traffic, anomaly-detection-oriented security tooling, and regularly updated protocol defenses are all vital to preventing intrusions, safeguarding patient privacy, and maintaining service continuity.

In summary, IoMT devices occupy a critical position in modern healthcare yet are inherently more exposed to cyber risk due to their constrained processing and memory resources. Successful attacks may jeopardize confidentiality, integrity, and the uninterrupted delivery of care. Consequently, there is a growing emphasis on employing advanced approaches—such as machine learning-based techniques—to bolster IoMT security and resilience.

Departing from the supervised learning approaches prevalent in the literature, we study a lightweight one-class anomaly-detection pipeline. Our method first applies Principal Component Analysis (PCA) to compress the feature space and isolate the most stable patterns of network behavior; we then train a One-Class Support Vector Machine (OC-SVM) solely on principal components extracted from benign traffic. We evaluate generalization on four IoMT datasets—BoT-IoT, ECU-IoHT, CICIoMT2024, and IoMT-TrafficData—under a unified preprocessing and tuning protocol. This normal-only, dimension-reduced design minimizes dependence on scarce attack labels and yields a compact model amenable to edge deployment in clinical networks.

This work makes the following contributions:

- We adopt a normal-only training regime that reduces reliance on attack labels, addressing severe class imbalance and improving resilience to unseen (zero-day) attacks.
- We conduct a multi-dataset evaluation on four public IoMT corpora within a consistent preprocessing and model-selection protocol, enabling fair comparison across datasets.
- We propose a PCA-OC-SVM pipeline tailored to IoMT traffic that is lightweight and deployment-friendly for resource-constrained devices and gateways.
- We present an ablation quantifying the impact of PCA on accuracy, false-alarm behavior, and model compactness, clarifying when dimensionality reduction helps without harming detection quality.
- We provide per-attack breakdowns and error analyses to guide thresholding and false-alarm control in operational clinical networks.
- To the best of our knowledge, this is among the first unified, normal-only one-class baselines reported across multiple IoMT datasets, emphasizing practicality and efficiency.

The remainder of this paper is organized as follows: Section 2 reviews related work on IoMT security and anomaly detection. Section 3 presents our proposed two-stage method. Section 4 outlines the experimental setup, including datasets, metrics, and the evaluation protocol. Section 5 reports the experimental results, including ablation and per-dataset analyses, computational footprint, and state-of-the-art comparisons. Section 6 discusses the

main findings and deployment implications. Finally, Section 7 concludes the paper and outlines directions for future work.

2. Related Work

The security of Internet of Medical Things (IoMT) devices has become a pressing concern as device fleets and the sensitivity of clinical telemetry continue to grow. The literature broadly spans (i) IoMT-focused intrusion detection systems (IDS) that tailor learning architectures to multi-protocol medical traffic and edge constraints, (ii) strategies to cope with severe class imbalance—where rare and evolving attacks appear amid abundant benign activity—and (iii) anomaly/novelty detection paradigms that dispense with full supervision by modeling "normal" behavior only. Anomaly detection is a well-established paradigm for settings where labeled attacks are scarce or evolving [14]. In this study, we align with the third line: we adopt a one-class support vector machine (OC-SVM) trained exclusively on benign traffic, and we select its decision boundary via randomized cross-validated search over ν (training fraction treated as outliers) and γ (RBF kernel bandwidth), choosing the best configuration by CV F1-score on held-out benign-only splits. We then evaluate the resulting detector on four public corpora representative of medical environments (BoT-IoT, ECU-IoHT, CICIoMT2024, and IoMT-TrafficData), reporting Accuracy, Precision/Recall/F1-score for the attack class to diagnose asymmetric errors between normal and attack classes. This normal-only training protocol reduces dependence on scarce or noisy attack labels while keeping the detector compact and deployable near the data source in clinical networks.

A large body of work adopts fully supervised deep or ensemble models tailored to IoMT traffic. Mosaiyebzadeh et al. propose a federated deep neural network IDS that learns from distributed hospitals without sharing raw records, demonstrating strong accuracy on WUSTL–EHMS–2020 and ECU–IoHT under privacy constraints [15]. Katib and Ragab design H3SC–DLIDS, a blockchain–assisted framework that combines meta–heuristic feature selection (HHO, SCA) with an LSTM autoencoder to detect DDoS attacks on BoT–IoT [16]. Kumar and Kim introduce EFL–LSTM, an embedded federated LSTM that captures temporal patterns in IoHT telemetry while keeping sensitive data local [17]. Other studies explore CNN–based IDS for medical traffic [18], autoencoder–plus–attention hybrids [19], hybrid CNN–LSTM or GRU pipelines [20–22], stacked or tree–based ensembles [23], and blockchain–integrated or Swin–Transformer–based schemes for cloud–facing analytics on IoMT-TrafficData [24]. Collectively, these works report high accuracy in supervised settings where abundant labeled attack data are available.

Several contributions focus more explicitly on deployment constraints, class imbalance, and real–time operation. Saheed et al. propose IoT–Defender, a lightweight IDS that uses a modified genetic algorithm for feature selection and an LSTM classifier deployed on resource–constrained edge servers such as Raspberry Pi, achieving high detection rates with focal loss to mitigate imbalance on BoT–IoT [25]. Kumar et al. present a hybrid edge–to–cloud framework coupling embedded feature selection, a 1D–CLSTM classifier, and a PoAh consensus layer, targeting medico–legal traceability as well as robust detection on ECU–IoHT and WUSTL–EHMS–2020 [26]. Alalwany et al. embed a stacking ensemble into a Kappa architecture to sustain low–latency analytics on streaming ECU–IoHT telemetry [27], while other studies quantify CPU, memory, or latency overheads for deep or ensemble IDS pipelines on IoMT networks [23,25]. These methods highlight the importance of resource awareness but still rely on supervised attack labels and often target a single dataset or deployment scenario.

Closer to our perspective, several works move toward anomaly–based or representation–learning–enhanced detection. Hernandez–Jaimes et al. leverage Nilsimsa fingerprint-

ing to derive compact similarity hashes from payloads and apply shallow ensembles for ransomware detection in connected healthcare environments, thereby reducing manual feature engineering even though the final training remains supervised [28]. Kirubavathi et al. combine autoencoders with TabNet to build an interpretable IDS that is relatively robust under class imbalance [19]. The recently released CICIoMT2024 dataset by Dadkhah et al. [29] provides a realistic multi–protocol environment with detailed device profiling and heterogeneous attack families, which has become a common benchmark for evaluating both supervised and anomaly–based IoMT IDS. Beyond IoMT, PCA has been used as a pre–processing step for feature reduction before supervised classifiers, but such pipelines are typically trained on labeled attacks and evaluated on one or two generic datasets.

As summarized in Table 1, the majority of IoMT intrusion-detection studies still adopt fully supervised learning and thus require substantial labeled attack corpora, which can be costly and difficult to maintain in real clinical environments. While several works explicitly consider deployment constraints and, in some cases, discuss edge considerations, they typically remain label-dependent and are often validated within a single dataset or a narrow operational scenario. Moreover, representation-learning components (e.g., autoencoders) are frequently paired with supervised classifiers rather than establishing a truly benign-only baseline, leaving limited systematic evidence on how far a carefully tuned, lightweight one-class detector can go across heterogeneous IoMT corpora. Motivated by these gaps, we study a deployment-oriented PCA-OC-SVM pipeline trained exclusively on benign traffic and evaluate a single, unified protocol across four public datasets (BoT-IoT, ECU-IoHT, CICIoMT2024, and IoMT-TrafficData), aiming to balance robust detection under severe imbalance with computational practicality for edge/gateway deployment.

Table 1. Representative IoMT IDS studies and the gaps addressed by this work.

Ref.	Paradigm	Core Idea	Datasets	Key Gap
[15]	Supervised (FL)	Federated DNN IDS	ECU-IoHT WUSTL-EHMS-2020	Supervised/label-dependent FL setting; emphasis is privacy/distributed learning rather than establishing a benign-only baseline; unified multi-corpus comparability under a single normal-only protocol is not the primary focus.
[16]	Supervised	HHO/SCA + LSTM-AE blockchain-assisted	BoT-IoT	Multi-component and operationally heavy pipeline (meta-heuristics + deep model + blockchain); still label-dependent in practice and typically evaluated in a scenario-specific manner rather than as a unified benign-only baseline.
[17]	Supervised (FL)	Embedded FL LSTM	ECU-IoHT	Requires labeled attacks and continual label maintenance; evaluation is largely dataset/scenario specific; novelty/zero-day capability is not explicitly targeted via a benign-only training paradigm.
[25]	Supervised (edge-aware)	Lightweight FS + LSTM focal loss	BoT-IoT UNSW-NB15 N-BaIoT	Edge-aware deployment is considered, yet the detector remains supervised and label-dependent; robustness to unseen/novel attacks is not explicitly validated under a normal-only training protocol.
[26]	Supervised (edge-to-cloud)	1D-CLSTM + PoAh imbalance handling	ECU-IoHT WUSTL-EHMS-2020	System-level edge-to-cloud framework introduces additional complexity; supervised training still depends on attack labels; evidence for unified cross-dataset performance under one consistent protocol is limited.
[27]	Supervised (streaming)	Stacking ensemble + Kappa architecture	ECU-IoHT	Streaming/stacking approach is supervised and operationally heavier; generalization beyond a single corpus is a concern, and a compact benign-only baseline is not the main contribution.
[18]	Supervised	CNN-based IDS	CICIoMT2024	CNN-based supervised classification relies on labeled attacks and typically reflects closed-world evaluation; a label-free benign-only detector baseline is not provided for evolving/novel threats.
[19]	Hybrid	AE representations + TabNet classifier	CICIoMT2024	Although representation learning is used, the final decision stage is supervised (classifier-based), not a purely benign-only detector; computational footprint is not the main focus.
[28]	Supervised	Nilsimsa fingerprints + shallow ensembles	CICIoMT2024 ICE Dataset	Feature/fingerprint-driven and label-dependent detection; transferability across heterogeneous IoMT corpora and a unified benign-only baseline are not the primary focus.
[23]	Supervised (ensemble)	Tree/boosting IDS	CICIoMT2024	Supervised boosting depends on labeled attacks; adaptability to unseen/zero-day behaviors is limited by the training paradigm, and a unified benign-only baseline is not established.

3. Proposed Method

Our proposed anomaly detection pipeline is a deployment-oriented, two-stage process designed to identify cyberattacks in IoMT networks efficiently. The core idea is to model the behavior of only normal (benign) network traffic and to detect any deviations from this learned norm as potential attacks. This approach circumvents the need for scarce and often incomplete labeled attack data, thereby enhancing the system's robustness against novel and zero-day threats. The workflow consists of an initial dimensionality reduction stage using Principal Component Analysis (PCA), followed by a one-class classification stage using a One-Class Support Vector Machine (OC-SVM) to build the model of normal behavior. The workflow of the method is shown in Figure 1, and the full training and evaluation procedure is summarized in Algorithm 1.

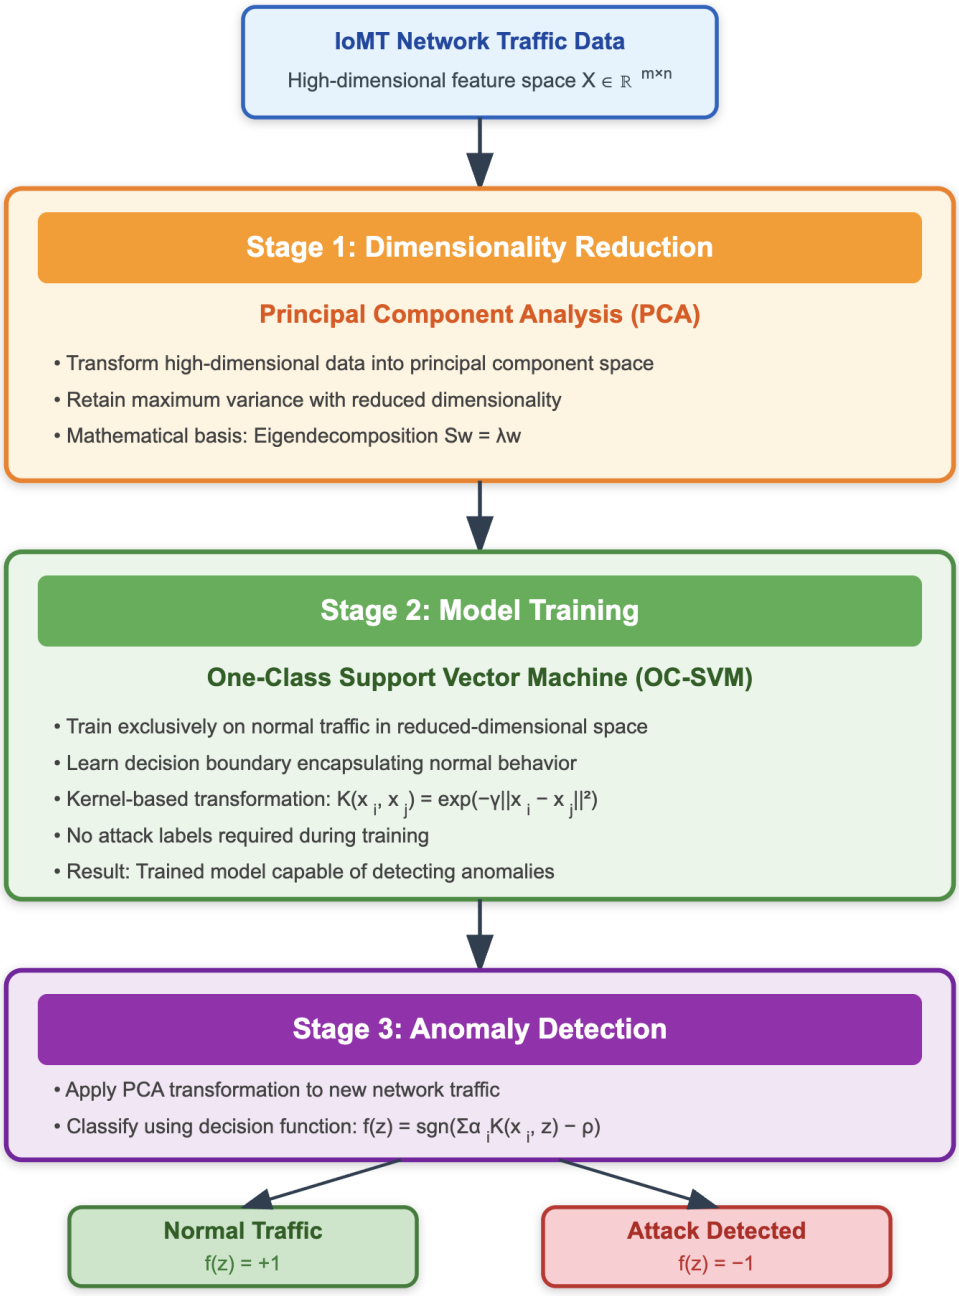

Figure 1. Flowchart of the proposed method.

Algorithm 1 Unified PCA + One-Class SVM protocol across multiple IoMT datasets

procedure PCA-OCSVM($\{\mathcal{D}_k\}_{k=1}^4$, τ, $\mathcal{R}$, $\mathcal{S}(\nu,\gamma)$, N_{iter}, s)
  Input: $\{\mathcal{D}_k\}_{k=1}^4$: four IoMT datasets
    $\tau = 0.10$: tuning fraction
    $\mathcal{R} = \{0.95, 0.90, \ldots, 0.60\}$: variance ratios to test
    $\mathcal{S}(\nu,\gamma)$: hyperparameter search space
    $N_{\text{iter}} = 100$: randomized search iterations
    $s = 42$: random seed
  Output: Dataset-wise performance metrics and models
  for all dataset $\mathcal{D} \in \{\mathcal{D}_k\}_{k=1}^4$ **do**
    Load $\mathcal{D}$; map labels to {`normal`, `attack`}; extract numeric features X
    $\mathcal{T} \leftarrow \text{Sample}(\mathcal{D}, \tau, s)$ ▷ 10% tuning subset
    Split normal samples in $\mathcal{T}$ into $\mathcal{T}_{tr}$ (80%) and $\mathcal{T}_{val}$ (20%)
    $\mathcal{V} \leftarrow \mathcal{T}_{val} \cup \{\text{all attack samples in } \mathcal{T}\}$ ▷ validation set
    Initialize: $F1_{\text{best}} \leftarrow -1$, $\rho^* \leftarrow 0$, $(\nu^*, \gamma^*) \leftarrow \emptyset$
    for all variance ratio $\rho \in \mathcal{R}$ **do**
      Fit $\text{Scaler}_{\text{tune}}$ on $\mathcal{T}_{tr}$ (normal-only)
      Transform: $\mathcal{T}_{tr}^{\text{scaled}} \leftarrow \text{Scaler}_{\text{tune}}(\mathcal{T}_{tr})$
        $\mathcal{V}^{\text{scaled}} \leftarrow \text{Scaler}_{\text{tune}}(\mathcal{V})$
      Determine n_{comp} to achieve variance $\geq \rho$ on $\mathcal{T}_{tr}^{\text{scaled}}$
      Fit $\text{PCA}_{\text{tune}}(n_{\text{comp}})$ on $\mathcal{T}_{tr}^{\text{scaled}}$ (normal-only)
      Transform: $\mathcal{T}_{tr}^{\text{PCA}} \leftarrow \text{PCA}_{\text{tune}}(\mathcal{T}_{tr}^{\text{scaled}})$
        $\mathcal{V}^{\text{PCA}} \leftarrow \text{PCA}_{\text{tune}}(\mathcal{V}^{\text{scaled}})$
      Create predefined split: train indices $= -1$, validation indices $= 0$
      $X_{\text{all}} \leftarrow [\mathcal{T}_{tr}^{\text{PCA}}; \mathcal{V}^{\text{PCA}}]$
      $(\nu', \gamma') \leftarrow \arg\max_{(\nu,\gamma)\in\mathcal{S}} \text{RandomizedSearchCV}(X_{\text{all}}, N_{\text{iter}}, \text{F1-scorer})$
      if $\text{F1-score}(\nu', \gamma', \mathcal{V}^{\text{PCA}}) > F1_{\text{best}}$ **then**
        $F1_{\text{best}} \leftarrow \text{F1-score}(\nu', \gamma', \mathcal{V}^{\text{PCA}})$
        $(\nu^*, \gamma^*) \leftarrow (\nu', \gamma')$
        $\rho^* \leftarrow \rho$, $n_{\text{comp}}^* \leftarrow n_{\text{comp}}$
      end if
    end for
    Split full $\mathcal{D}$: Train = 80% normal, Test = 20% normal $\cup$ all attacks
    Fit $\text{Scaler}_{\text{final}}$ on Train (normal-only)
    Transform: $\text{Train}^{\text{scaled}} \leftarrow \text{Scaler}_{\text{final}}(\text{Train})$
      $\text{Test}^{\text{scaled}} \leftarrow \text{Scaler}_{\text{final}}(\text{Test})$
    Fit $\text{PCA}_{\text{final}}(n_{\text{comp}}^*)$ on $\text{Train}^{\text{scaled}}$ (normal-only)
    Transform: $\text{Train}^{\text{final}} \leftarrow \text{PCA}_{\text{final}}(\text{Train}^{\text{scaled}})$
      $\text{Test}^{\text{final}} \leftarrow \text{PCA}_{\text{final}}(\text{Test}^{\text{scaled}})$
    Train OC-SVM(ν^*, γ^*, kernel = rbf) on $\text{Train}^{\text{final}}$
    Predict: $\hat{y}_{\text{test}} \leftarrow \text{OC-SVM.predict}(\text{Test}^{\text{final}})$
    Map: $\hat{y}_{\text{test}}[i] \leftarrow \begin{cases} \text{normal} & \text{if } \hat{y}_{\text{test}}[i] = 1 \\ \text{attack} & \text{if } \hat{y}_{\text{test}}[i] = -1 \end{cases}$
    Compute: Accuracy, Precision, Recall, F1, ROC-AUC, PR-AUC
    Generate: Confusion matrix, ROC/PR curves, scree plot, per-attack breakdown
    Measure: Model size (MB), inference time (ms/sample)
  end for
  return Aggregated results across all four datasets
end procedure

3.1. Dimensionality Reduction with PCA

Cybersecurity datasets often describe network traffic with a large number of correlated features, leading to high-dimensional data, where anomaly detection can become challenging due to sparsity and scalability issues [30]. Processing such data directly can be computationally expensive and may suffer from the "curse of dimensionality," where models struggle to

generalize [31]. To address this, the first stage of our pipeline employs Principal Component Analysis (PCA), a powerful unsupervised linear transformation technique.

The primary objective of PCA is to transform the original set of correlated features into a new, smaller set of uncorrelated variables known as principal components. These components are ordered such that the first few retain most of the variance present in the original data, effectively summarizing the information with minimal loss [32]. Given a dataset represented by a matrix X of size $m \times n$ (where m is the number of samples and n is the number of features), the first step is to center the data by subtracting the mean of each feature from all its values. Let this centered matrix be X_c.

The goal of PCA is to find a new set of orthogonal axes (principal components) onto which the projection of the data has the maximum variance. Let w be a unit vector representing the direction of the first principal component. Maximizing the variance of the projected data can be formulated as:

$$\text{maximize} \quad \frac{1}{m}\sum_{i=1}^{m}(X_{c,i}^{T}w)^2 \quad \text{subject to} \quad w^T w = 1 \tag{1}$$

This optimization is equivalent to maximizing $w^T S w$, where S is the covariance matrix of the data, $S = \frac{1}{m}X_c^T X_c$. The solution is found through eigenvalue decomposition of the covariance matrix. The vector w that maximizes the variance is the eigenvector corresponding to the largest eigenvalue λ, based on the fundamental equation of PCA:

$$Sw = \lambda w \tag{2}$$

By selecting the top k eigenvectors corresponding to the k largest eigenvalues, we construct a projection matrix to transform the original n-dimensional data into a new, more compact k-dimensional space, where $k < n$.

Let $X \in \mathbb{R}^{m \times n}$ denote the centered data matrix with m samples and n features. Computing the covariance matrix $S = \frac{1}{m}X_c^T X_c$ requires $\mathcal{O}(mn^2)$ operations. The eigenvalue decomposition of $S \in \mathbb{R}^{n \times n}$ scales as $\mathcal{O}(n^3)$. After selecting the top k eigenvectors ($k \ll n$), projecting the data into the reduced space costs $\mathcal{O}(mnk)$. In practice, using a small k substantially reduces the downstream learning and inference costs.

3.2. Anomaly Detection with One-Class SVM

After reducing the data's dimensionality, the second stage involves training a One-Class Support Vector Machine (OC-SVM) on these principal components. Unlike traditional supervised learning, which requires examples from all classes, our approach focuses exclusively on modeling the characteristics of normal network behavior. OC-SVM, first proposed by Scholkopf et al. [33], is an unsupervised algorithm designed for novelty and outlier detection. The one-class learning objective is closely related to Support Vector Data Description (SVDD), which learns a compact description of normal data and flags deviations as outliers [34]. It learns a decision boundary that encapsulates the normal data points, and any new data point falling outside this boundary is classified as an anomaly. This strategy is exceptionally well-suited for the IoMT security context, where attack data is often scarce and constantly evolving, making the modeling of normal behavior a more robust strategy [35].

The fundamental idea behind OC-SVM is to map the input data into a high-dimensional feature space and then find a hyperplane that separates the majority of these data points from the origin with a maximum possible margin, where the origin is treated as the sole member of the 'outlier' class [36]. This transformation, illustrated in Figure 2, is crucial as a complex, non-linear data distribution can become linearly separable in the

higher-dimensional space. This concept is formalized by solving the following primal optimization problem for a set of normal observations $\{x_1, x_2, \ldots, x_m\}$:

$$\min_{w,\xi,\rho} \frac{1}{2}\|w\|^2 + \frac{1}{\nu m}\sum_{i=1}^{m}\xi_i - \rho \tag{3}$$

subject to the constraints:

$$(w \cdot \phi(x_i)) \geq \rho - \xi_i, \quad \xi_i \geq 0 \quad \text{for } i = 1, \ldots, m \tag{4}$$

The hyperparameter $\nu \in (0, 1]$ controls the flexibility of the one-class decision boundary through two theoretical guarantees [33,35]:

- Upper bound on outliers: at most a fraction ν of training samples can have $\xi_i > 0$ (i.e., be treated as margin violations/outliers).
- Lower bound on support vectors: at least a fraction ν of training samples become support vectors ($\alpha_i > 0$), thus shaping the decision boundary.

Intuitively, ν sets the expected outlier rate within the benign training data and therefore trades off boundary tightness against robustness to benign variability. Smaller values (e.g., $\nu = 0.01$) enforce a tighter boundary that admits very few benign edge cases as normal, which can increase sensitivity to subtle deviations but may raise false alarms. Larger values (e.g., $\nu = 0.5$) yield a looser boundary that tolerates greater variability in normal traffic, potentially reducing false alarms but risking missed detections for weak anomalies.

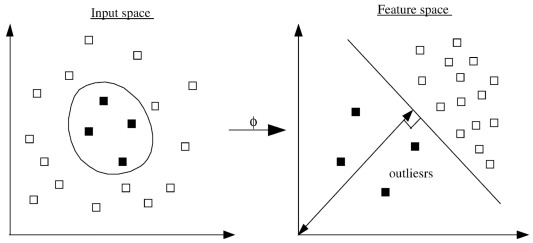

Figure 2. Illustration of the non-linear transformation from the input space to a high-dimensional feature space where data becomes linearly separable from the origin. Filled squares (blacksquare) denote the target class modeled by OC-SVM (normal data), whereas hollow squares (square) indicate samples lying outside the learned description (outliers/anomalies).

To solve this problem efficiently, OC-SVM employs the kernel trick, which avoids explicit computation of $\phi(x)$ by using a kernel function, such as the Radial Basis Function (RBF) kernel used in our study:

$$K(x_i, x_j) = \exp(-\gamma\|x_i - x_j\|^2) \tag{5}$$

This converts the primal problem into a solvable dual quadratic programming (QP) problem. The data points x_i with non-zero Lagrange multipliers are called support vectors, as they alone define the hyperplane. Once the model is trained, the decision function for a new test point z is given by:

$$f(z) = \text{sgn}\left(\sum_{i=1}^{m}\alpha_i K(x_i, z) - \rho\right) \tag{6}$$

A result of $+1$ indicates a 'normal' instance, while -1 indicates an 'attack'. A critical aspect of this approach is its sensitivity to outliers in the training data, which can become support vectors and distort the decision boundary, as illustrated in Figure 3. This sensitivity

underscores the importance of meticulous hyperparameter tuning to create a tight, yet robust, boundary around the true distribution of normal data.

Let N denote the number of training samples (benign-only) and d the feature dimension after PCA. Training an RBF-kernel OC-SVM requires computing pairwise kernel values, which scales as $\mathcal{O}(N^2 d)$ due to the $N \times N$ kernel matrix and $\mathcal{O}(d)$ cost per distance computation. The memory footprint is dominated by storing the kernel matrix, i.e., $\mathcal{O}(N^2)$. At inference time, scoring a test sample requires evaluating the kernel against the support vectors, yielding approximately $\mathcal{O}(N_{\text{sv}} d)$ per sample, where $N_{\text{sv}} \leq N$ is the number of support vectors.

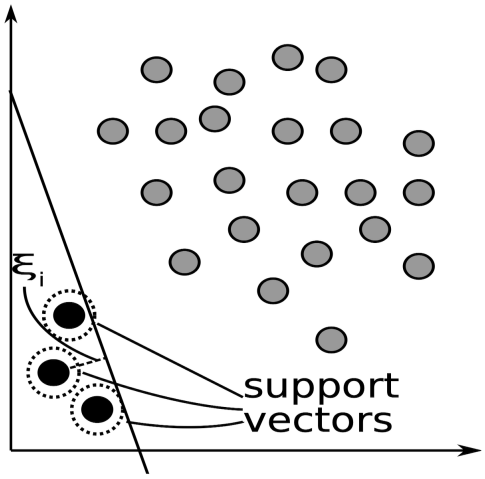

Figure 3. An illustration of how outliers (dark points) in the training data can become support vectors and shift the decision boundary.

4. Experimental Setup

4.1. Datasets

Datasets sit at the heart of machine learning and artificial intelligence: they are the foundational assets on which models are trained, tuned, and judged. Built from observations, measurements, and experience for a specific task, they provide the indispensable substrate for training, testing, and performance assessment [37]. With AI now pervasive in domains such as finance, healthcare, and cybersecurity, the salience of both dataset scale and dataset quality has grown markedly [38]. In practice, a model's effectiveness is largely governed by how rich, accurate, and representative its data is [39].

A dataset's primary role is to expose the patterns and relationships present in the real world so that learning algorithms can internalize them. This role extends far beyond amassing large quantities of samples; the data must encompass sufficient diversity across classes, conditions, and variations. In cybersecurity, for example, widely used corpora such as NSL-KDD and CIC-IDS2018 are curated to include heterogeneous attack families and network-traffic behaviors so that models encounter a realistic operating envelope [40,41].

Constructing such datasets is a multi-stage, nontrivial pipeline rather than a simple act of collection. One must first decide whether to source data from real deployments or simulations, and then apply feature engineering to distill informative signals from raw inputs [42]. Labeling—assigning correct ground truth at scale—can be costly and error-prone, particularly for large corpora. Subsequent cleaning removes missing, inconsistent, and noisy entries. In sensitive areas like IoT and healthcare, anonymization and security filtering are essential to protect personal information, and adhering to ethical and legal obligations can complicate both sharing and preparation [43].

Beyond their use in training, datasets now underpin fair algorithmic comparison, scientific reproducibility, community standards, and accuracy benchmarking [44]. Open, well-documented, and regularly maintained resources function as a shared lingua franca, enabling different methods to be evaluated under comparable conditions. The CICIoMT2024 dataset employed in this study exemplifies this trend: its multi-protocol traffic captured from IoMT devices supports both device profiling and attack detection, offering a distinct and timely testbed [29]. These developments highlight a qualitative evolution as well as a quantitative one—datasets are increasingly specialized to the needs of particular application areas.

In sum, datasets not only determine what a machine-learning system has learned; they also delimit what it can learn next. A comprehensive, balanced, and correctly labeled corpus that mirrors real-world distributions directly shapes model performance; conversely, even state-of-the-art learning algorithms can falter when fed inadequate data. This dependency is especially pronounced for one-class methods such as OC-SVM, which are trained only on "normal" data—coverage gaps or distribution shift in the normal class directly inflate false alarms or missed detections. Consequently, selecting and preparing datasets should be treated with the same rigor as algorithm selection, feature engineering, and training protocol design. Looking ahead, continued maturation of domain-specific datasets will improve both the fidelity and practicality of learning-based solutions. In this sense, the trajectory of machine learning is tightly coupled to the ongoing evolution of datasets and the innovations surrounding them.

4.1.1. BoT-IoT

The BoT-IoT dataset, developed by Koroniotis et al. (2019), was created to provide a realistic and large-scale benchmark for network forensic analytics and intrusion detection in IoT environments [45]. It was specifically designed to overcome the limitations of previous datasets, such as a lack of modern attack scenarios, incomplete traffic capture, and the absence of legitimate IoT traffic patterns. The dataset was generated within a sophisticated testbed at the Cyber Range Lab of UNSW Canberra, comprising a mix of virtual machines and simulated IoT devices. The environment utilized Node-RED to emulate realistic smart home devices communicating via the MQTT protocol. Attack traffic was launched from dedicated Kali Linux VMs, ensuring a diverse and contemporary set of threats.

As detailed in Table 2, the dataset comprises over 73 million records, making it suitable for evaluating the scalability of detection models. A defining characteristic is its severe class imbalance, with approximately 73.36 million attack instances overwhelming the 9543 normal traffic samples. This imbalance presents a significant challenge for training robust machine learning models, as classifiers may become biased towards the majority (attack) class, potentially leading to poor performance in identifying normal traffic. The dataset contains 46 features, including flow-based statistics and connection details, providing a rich feature space for developing sophisticated detection models.

The dataset's features provide rich ground for analysis. Volumetric features such as pkts (total packets) and bytes (total bytes) are fundamental for identifying high-volume attacks. Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) attacks, which constitute the largest portion of the dataset, are characterized by an anomalous and sustained surge in these values, often saturating network bandwidth or exhausting server resources.

Temporal flow features, including rate (packets per second), srate (source-to-destination packets per second), and drate (destination-to-source packets per second), offer insights into traffic behavior over time. These are particularly effective for detecting flooding attacks, where an attacker attempts to saturate a network or service. An

abnormally high srate, for instance, is a strong indicator of an outgoing attack from a compromised device.

Protocol-level features like proto (e.g., TCP, UDP, ICMP) and connection state information allow for more granular analysis. For example, Reconnaissance attacks, such as service and OS scans, often manifest as a series of connection attempts across various ports or unusual protocol flag combinations. Analyzing these features can help distinguish legitimate network discovery from malicious scanning activities. Similarly, Information Theft attacks like Keylogging and Data Theft are stealthier and may be identified through subtle deviations in packet size, timing, and protocol usage that differ from baseline normal behavior.

Table 2. Sample counts in the BoT-IoT dataset.

Class	Category	Type of Attack	Count
Normal	-	-	9543
Attack	Information Gathering	Service Scanning (Nmap, Hping3)	1,463,364
		OS Fingerprinting (Nmap, Xprobe2)	358,275
	DDoS	TCP (hping3)	19,547,603
		UDP (hping3)	18,965,106
		HTTP (golden-eye)	19,771
	DoS	TCP (hping3)	12,315,997
		UDP (hping3)	20,659,491
		HTTP (golden-eye)	29,706
	Information Theft	Keylogging (Metasploit)	1469
		Data Theft (Metasploit)	118

4.1.2. CICIoMT2024

The Internet of Things (IoT) is becoming increasingly integrated into daily life, particularly in the field of healthcare through the Internet of Medical Things (IoMT). IoMT devices enhance healthcare services by supporting functionalities such as uninterrupted operation. However, these devices also raise serious cybersecurity concerns due to vulnerabilities that may arise during health monitoring. This issue is further exacerbated by the complexity of IoMT network traffic and the large volume of data generated.

In this context, Machine Learning techniques are being used to detect, prevent, and mitigate the effects of cyberattacks. However, existing benchmark datasets have some fundamental shortcomings, such as the presence of only a small number of real devices, limited diversity, and the inability to create comprehensive device profiles.

To address these shortcomings and to improve and evaluate IoMT security solutions, the CICIoMT2024 dataset has been proposed as a realistic benchmark dataset [29]. As part of this study, an IoMT test environment consisting of 40 devices (25 real, 15 simulated) was established.

The CICIoMT2024 dataset serves as a foundation that complements existing datasets and supports Machine Learning-based cybersecurity solutions. It provides a valuable resource for researchers to classify cyberattacks targeting IoMT devices and to develop secure healthcare systems, thus aiming to contribute to the widespread adoption of more reliable IoMT devices. The CICIoMT2024 training and test datasets are provided separately. The training set consists of 7,160,831 samples and the test set consists of 1,614,182 samples. Class, category, and attack counts are presented in Table 3 for the training set and Table 4 for the test set.

The dataset contains traces of attacks conducted over three main communication protocols: Wi-Fi, Bluetooth Low Energy (BLE), and MQTT. Under each protocol, both normal

and various attack scenarios were modeled. The 18 different attacks applied are classified under five categories: DDoS, DoS, Recon (reconnaissance attacks), MQTT-based attacks, and Spoofing. These attacks include ARP spoofing, port scanning, ping sweep, MQTT publish/connect flood, malformed packet transmissions, and various TCP/UDP/ICMP/SYN attacks. Thus, the dataset is suitable for both multiclass and binary classification studies.

Table 3. Sample counts in the CICIoMT2024 training set.

Class	Category	Type of Attack	Count
Normal	-	-	192,732
Attack	Spoofing	ARP Spoofing	16,047
	DDoS	TCP_IP-DDoS-UDP	1,635,956
		TCP_IP-DDoS-ICMP	1,537,476
		TCP_IP-DDoS-TCP	804,465
		TCP_IP-DDoS-SYN	801,962
	DoS	TCP_IP-DoS-UDP	566,950
		TCP_IP-DoS-SYN	441,903
		TCP_IP-DoS-ICMP	416,292
		TCP_IP-DoS-TCP	380,384
	MQTT	MQTT-DDoS-Connect_Flood	173,036
		MQTT-DoS-Publish_Flood	44,376
		MQTT-DDoS-Publish_Flood	27,623
		MQTT-DoS-Connect_Flood	12,773
		MQTT-Malformed_Data	5130
	RECON	Recon-Port_Scan	83,981
		Recon-OS_Scan	16,832
		Recon-VulScan	2173
		Recon-Ping_Sweep	740

Table 4. Sample counts in the CICIoMT2024 test set.

Class	Category	Type of Attack	Count
Normal	-	-	37,607
Attack	Spoofing	ARP Spoofing	1744
	DDoS	TCP_IP-DDoS-UDP	362,070
		TCP_IP-DDoS-ICMP	349,699
		TCP_IP-DDoS-TCP	182,598
		TCP_IP-DDoS-SYN	172,397
	DoS	TCP_IP-DoS-UDP	137,553
		TCP_IP-DoS-SYN	98,595
		TCP_IP-DoS-ICMP	98,432
		TCP_IP-DoS-TCP	82,096
	MQTT	MQTT-DDoS-Connect_Flood	41,916
		MQTT-DoS-Publish_Flood	8505
		MQTT-DDoS-Publish_Flood	8416
		MQTT-DoS-Connect_Flood	3131
		MQTT-Malformed_Data	1747
	RECON	Recon-Port_Scan	22,622
		Recon-OS_Scan	3834
		Recon-VulScan	1034

The dataset shows the count or average of packets belonging to relevant protocols such as TCP, UDP, ICMP, ARP, DHCP, DNS, and HTTP. It contains information such as fin_flag_number, syn_flag_number, and rst_flag_number, which are useful for detecting

DoS and port scanning (recon) attacks. It also includes statistics such as Min, Max, AVG, and Std for packet lengths, which can be used for network load analysis.

The features extracted from the dataset cover both technical details at the network level and time-based behaviors. The features include packet header lengths, Time-To-Live (TTL), flag information (ACK, SYN, FIN, etc.), average values for protocols such as IGMP, DNS, and HTTP, as well as statistical information regarding packet sizes (minimum, maximum, average, and standard deviation). In total, 45 features were defined, and these features were provided in PCAP and CSV formats.

4.1.3. ECU-IoHT

To strengthen the cybersecurity of IoHT, the ECU-IoHT dataset was developed, built upon an IoHT environment where various attacks exploiting different vulnerabilities were carried out [46]. This dataset was designed to help the healthcare security community analyze attack behaviors and develop robust countermeasures. In this dataset, an IoHT test environment was set up, and vulnerabilities were investigated to develop a realistic dataset. The dataset features two classes (attack and normal) as well as five attack types: Smurf attack, Nmap Port Scan, ARP Spoofing, DoS attack, and normal.

As seen in Table 5, the dataset consists of a total of 111,206 samples, with each row representing a network packet or event and containing nine features: No, Time, Source, Destination, Protocol, Length, Info, Type, and Type of attack. The "No" column uniquely identifies each network event, while the "Time" column shows the timestamp of events in seconds. The "Source" and "Destination" columns indicate the IP or MAC addresses from which the packets are sent and received. The "Protocol" column includes network protocols such as ARP, TCP, and DNS, while the "Length" column expresses the packet size in bytes. The "Info" column provides detailed information about the packet content, and the "Type" and "Type of attack" columns label the events as either normal or attack, and specify the attack types.

The "Time" column shows that events occurred within the time interval from 0.0 s to 10,109 s. This data provides a valuable opportunity to examine the temporal distribution of attacks. For example, patterns such as the concentration of ARP Spoofing attacks in certain intervals or the increase of DoS attacks at specific times can be detected.

The "Length" column indicates the size of packets in bytes and is an important metric for network traffic analysis. ARP packets are generally 42 bytes, while TCP and TLSv1.1 packets have larger sizes such as 58 bytes, 321 bytes, or 395 bytes. These differences can be used to distinguish between normal and abnormal traffic patterns.

The "Info" column provides detailed information about packet contents. ARP queries, for example, "Who has 192.168.43.1? Tell 192.168.43.186", demonstrate device authentication. DNS queries such as "Standard query 0x0c44 PTR 1.43.168.192.in-addr.arpa" contain domain name resolution requests, while TCP connections are indicated with statements like "[SYN] Seq=0 Win=1024 Len=0 MSS=1460", describing the process of establishing a connection.

Table 5. Sample counts in the ECU-IoHT dataset.

Class	Type of Attack	Count
Attack	Smurf Attack	77,920
	Nmap Port Scan	6836
	ARP Spoofing	2359
	DoS Attack	639
Normal	-	23,453

4.1.4. IoMT-TrafficData

The IoMT-TrafficData dataset was developed to address a critical gap in the field of cybersecurity for the Internet of Medical Things (IoMT): the lack of comprehensive, realistic, and publicly available datasets for benchmarking Intrusion Detection Systems (IDS) [47]. Existing datasets often fall short in representing the unique protocols, device heterogeneity, and specific attack vectors prevalent in modern healthcare environments. This dataset aims to provide a robust and representative foundation for researchers to develop, train, and validate machine learning (ML) and deep learning (DL) models for detecting malicious activities within IoMT networks.

The foundation of the dataset is a meticulously constructed testbed that emulates a real-world sports clinic scenario. This environment includes a diverse range of interconnected devices, categorized into a General IoT Area and a Wireless Body Area Network (WBAN). The setup featured sensors like a PIR motion detector and a DHT11 (temperature, humidity, and heat index) sensor, communicating via MQTT and CoAP protocols with an Arduino Yún and a Raspberry Pi, respectively. The Raspberry Pi also functioned as an HTTP server. The WBAN component consisted of a smartwatch and a heart rate belt, which communicated with a smartphone using both Bluetooth 5.0 and Bluetooth 4.0 Low Energy (BLE). This entire ecosystem was connected through a dedicated router and switch to create a private network, ensuring that the captured traffic was free from external noise and exclusively represented the experimental scenario.

Data collection was bifurcated into two primary phases: capturing benign and malicious traffic. Benign (normal) network traffic was collected over a continuous period of approximately 120 h, resulting in the capture of roughly 3 million packets to establish a baseline of normal operational behavior. For the malicious traffic, a series of targeted cyberattacks were executed using well-known tools. The dataset includes eight distinct types of IP-based attacks: four variants of Denial of Service (DoS) attacks (Apache Killer, R-U-Dead-Yet, Slow Read, and Slow Loris) executed with 'slowhttptest'; a CAM Table Overflow attack using 'macof'; ARP Spoofing and Network Scanning/Sniffing attacks performed with 'bettercap' and 'nmap'; and an MQTT Malaria attack using the 'mqtt-malaria' tool. Furthermore, to account for threats on other communication channels, two Bluetooth-specific attacks—Reconnaissance and Injection—were conducted using 'bettercap' and 'nrf' tools.

A distinguishing feature of IoMT-TrafficData is its hierarchical structure, which offers data at two different levels of granularity. The entire dataset is first divided into two main subsets: IP-based traffic and Bluetooth traffic. The IP-based subset is further processed into two distinct formats:

Packet-based Data: This contains the raw, granular information from individual network packets, extracted from '.pcap' files using the 'tshark' utility. This provides a deep-dive view into the network communications.

Flow-based Data: This represents aggregated data from sequences of packets that share common characteristics (e.g., same source/destination IP and port). Network flows were generated from the same '.pcap' files using the Zeek Flowmeter tool. This format is often more efficient for ML models and can reveal patterns not obvious at the packet level.

This dual-format approach for IP traffic is a key contribution, as it enables a direct and fair comparison of the efficacy of intrusion detection models when trained on packet-level features versus flow-level statistics.

Following the data capture and initial processing, a rigorous data arrangement and cleaning process was undertaken. This included removing redundant headers, labeling each record as either benign or with its specific attack type for supervised learning, and converting all data into a unified CSV format. The final composition of the dataset, detailing the

capture duration and the number of records for each class across the different approaches after pre-processing, is summarized in Table 6. The authors' evaluation demonstrated that this dataset is highly effective for training ML models, with findings indicating that a flow-based approach can improve multiclass classification performance by nearly 5% compared to a packet-based one. In this study, we adopt the flow-based representation of IoMT-TrafficData for all experiments. Consequently, IoMT-TrafficData stands as a valuable and multifaceted resource for advancing research in IoMT security.

Table 6. Sample counts in the IoMT-TrafficData dataset.

Class	Capture Duration (s)	No. Packets Generated (Approx.)	No. Records (Packets)	No. Records (Flows) *
IP-Based				
Normal (0)	57,600	3,032,261	3,032,261	300,000
Apache Killer (1)	1800	1,077,195	376,990	2334
R-U-Dead-YET (2)	1800	780,895	375,942	71,990
Slow Read (3)	1800	379,890	375,000	6740
Slow Loris (4)	1800	377,750	375,055	63,496
ARP Spoofing (5)	7200	765,500	375,055	11,222
CAM Table Overflow (6)	600	2,398,320	379,419	62,500
MQTT Malaria (7)	900	6,573,400	379,837	4961
Scanning (8)	900	3,287,890	379,407	62,500
Total of Instances (IP-Based)	114,000	15,382,341	6,048,938	585,743
Bluetooth				
Normal (0)	72,000	879,033	208,000	-
Bluetooth Reconnaissance (1)	10,800	43,812	104,000	-
Bluetooth Injection (2)	10,800	51,785	88,000	-
Total of Instances (Bluetooth)	114,000	974,630	400,000	-

* This study uses the flow-based representation.

4.2. Metrics

Evaluating model performance is critical for understanding generalization capabilities on unseen test data. This study employs accuracy and F1-score as the primary evaluation metrics. These metrics are used to assess different aspects of model prediction performance and are especially important for understanding the true performance of the model when working with imbalanced datasets.

Accuracy is the ratio of correctly classified samples to the total number of samples. As shown in Equation (7), the sum of true positives and true negatives is divided by the total number of cases. This metric shows the overall correctness of the model. However, if there is an imbalance between classes (e.g., one class is much larger than the other), this metric can be misleading.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \tag{7}$$

Precision, as shown in Equation (8), is the ratio of truly positive samples among those predicted as positive. It is important in cases where the cost of false positives is high (e.g., spam detection).

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{8}$$

Recall is the ratio of correctly predicted positive samples among all actual positive samples, as shown in Equation (9).

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{9}$$

F1-score is the harmonic mean of precision and recall, as in Equation (10), and provides a balance between these two metrics. The F1-score is preferred when both false positives and false negatives are important, or when classes are imbalanced.

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{10}$$

4.3. Training and Evaluation Protocol

To reduce overfitting and prevent data leakage, we adopt a two stage training procedure for each dataset. In the first stage, we draw a small tuning subset (10% of the full feature matrix) and use it exclusively for hyperparameter selection. Within this subset, only benign flows are used to fit the scaler and PCA, while validation folds contain both benign and attack traffic. The OC-SVM is trained on principal components extracted from benign flows, and its hyperparameters (ν, γ) together with the target PCA variance level are chosen via randomized search with a predefined split that maximises the F1-score on the internal validation partitions. The scaler, PCA and OC-SVM models fitted in this tuning stage are discarded afterwards; only the selected hyperparameters and PCA dimensionality are retained.

In the second stage, we train the final PCA enhanced OC-SVM detector from scratch on a separate benign training set and evaluate it on a held out test set that contains both benign and attack flows. The final scaler and PCA are fitted only on the benign training partition, and the OC-SVM is likewise trained exclusively on the corresponding principal components of benign traffic. The held out test flows are never used to fit any model parameters and are only passed through the already fitted scaler and PCA before being scored by the OC-SVM. Attack labels are thus used solely for validation and final evaluation, which prevents label leakage from the attack class into the one class model. Combined with the regularising effect of PCA and the bounded search ranges for ν and γ, this protocol limits overfitting and supports the strong generalisation performance observed on all four IoMT corpora.

5. Results

In this section, we present a comprehensive analysis of the performance of our proposed PCA-enhanced OC-SVM model. We begin with an ablation study to quantify the impact and trade-offs of the PCA stage. Subsequently, we provide a detailed examination of the model's performance on each of the four datasets, contextualizing our findings with comparisons to state-of-the-art methods. Finally, we present a consolidated table for a complete overview of all performance metrics.

5.1. Ablation Study

To validate the inclusion of Principal Component Analysis in our pipeline, we conducted an ablation study. The primary motivation for using PCA is to reduce the high dimensionality of network traffic data, which in turn decreases model complexity, training time, and inference latency, making it suitable for resource-constrained IoMT devices. This study evaluates how this dimensionality reduction impacts detection performance across multiple metrics. We compared our full pipeline (PCA + OC-SVM) against a baseline using only OC-SVM on the original high-dimensional features.

In addition to PCA, we also evaluated an Autoencoder (AE) as a non-linear representation learning alternative before OC-SVM, aiming to learn a compact latent space that may better capture complex dependencies in IoMT traffic. As reported in Table 7, the AE + OC-SVM pipeline remains highly competitive, but it does not provide consistent gains over the simpler PCA-based reduction: it slightly underperforms PCA on BoT-IoT and CICIoMT2024, and it degrades more noticeably on ECU-IoHT, while offering only a

marginal improvement over the OC-SVM-only baseline on IoMT-TrafficData. This suggests that optimizing reconstruction error alone may not preserve the most OC-SVM-separable structure of "normal" behavior across datasets, and the learned latent space can introduce dataset-specific distortions that affect the one-class boundary. Moreover, compared to PCA, AE-based reduction adds additional training/inference overhead and architectural hyperparameters, which weakens the deployment-oriented motivation of our pipeline. Overall, these findings reinforce PCA as the more robust and lightweight dimensionality reduction choice for a general-purpose, edge-friendly one-class IDS in IoMT settings.

Table 7. Ablation study on the impact of feature-space transformation (PCA vs. AE) prior to OC-SVM. The F1-Score is the primary metric for comparison due to class imbalance in the datasets.

Dataset	Model Configuration	Acc (%)	F1 (%)
BoT-IoT	OC-SVM only	**99.99**	**100.00**
	AE + OC-SVM	99.74	99.87
	PCA + OC-SVM	99.84	99.92
CICIoMT2024	OC-SVM only	**99.77**	**99.88**
	AE + OC-SVM	99.76	99.87
	PCA + OC-SVM	**99.77**	**99.88**
ECU-IoHT	OC-SVM only	98.27	99.09
	AE + OC-SVM	97.26	98.36
	PCA + OC-SVM	**98.58**	**99.25**
IoMT-TrafficData	OC-SVM only	97.28	98.37
	AE + OC-SVM	97.34	98.41
	PCA + OC-SVM	**98.66**	**99.19**

The results presented in Table 7 reveal a nuanced but important relationship between dimensionality reduction and model performance. For the ECU-IoHT and IoMT-TrafficData datasets, applying PCA provides a clear and significant improvement in the overall F1-Score. This confirms our hypothesis that for certain datasets, projecting features onto their principal components effectively filters noise and reduces complexity, enabling the OC-SVM to construct a more effective decision boundary.

Conversely, for the BoT-IoT dataset, which is characterized by an extreme class imbalance, the baseline model without PCA achieved near-perfect scores, outperforming the PCA-enhanced model. This suggests that the original high-dimensional feature space of BoT-IoT contains subtle but highly discriminative information that is partially lost during the linear projection of PCA. In this specific case of extreme imbalance, a powerful non-linear classifier like the RBF-kernel OC-SVM can be more effective when applied directly. The CICIoMT2024 dataset represents a neutral case where performance is largely maintained.

Despite the observed performance trade-off on one of the four datasets, the primary goal of our pipeline is to balance high performance with practical deployability. The universal and significant gains in computational efficiency from dimensionality reduction, combined with the clear performance benefits on half of the datasets, strongly justify the inclusion of PCA as a vital component in a general-purpose IoMT security pipeline. It provides a critical balance between achieving state-of-the-art detection performance and meeting the practical requirements of resource-constrained environments.

As an additional baseline, we replaced OC-SVM with Isolation Forest (IF) while keeping the same PCA representation and preprocessing pipeline. This comparison isolates the effect of the second-stage anomaly detector choice. Across datasets, IF yields slightly lower detection quality than OC-SVM (Table 8); for example, on IoMT-TrafficData, PCA + IF achieves 99.13% F1 compared to 99.19% with PCA + OC-SVM. This gap is consistent with the fact that

OC-SVM (with an RBF kernel) can form a smoother and more flexible boundary around benign behavior in the PCA space, whereas IF is more sensitive to its contamination/threshold setting and may under-separate subtle attack deviations or increase false alarms depending on the dataset. Therefore, we retain OC-SVM as the preferred second-stage detector in our deployment-oriented pipeline.

Table 8. Detector baseline comparison under the same PCA representation (PCA + OC-SVM vs. PCA + Isolation Forest). The F1-Score for the attack class is reported due to class imbalance.

Dataset	Model Configuration	Acc (%)	F1 (%)
BoT-IoT	PCA + OC-SVM	**99.84**	**99.92**
	PCA + IF	98.02	99.00
CICIoMT2024	PCA + OC-SVM	**99.77**	**99.88**
	PCA + IF	99.70	99.85
ECU-IoHT	PCA + OC-SVM	**98.58**	**99.25**
	PCA + IF	98.56	99.24
IoMT-TrafficData	PCA + OC-SVM	**98.66**	**99.19**
	PCA + IF	98.56	99.13

5.1.1. Hyperparameter Tuning and Sensitivity Analysis

To ensure a transparent and reproducible hyperparameter selection, we tune the OC-SVM parameters ν and γ on a small tuning split. Following our normal-only training assumption, the preprocessing steps (median imputation, standardization, and PCA) are fitted using only benign flows. Hyperparameters are then selected by evaluating the attack-detection F1-score on a validation fold that contains both benign and attack traffic (attack is treated as the positive class). We search ν in the range $[0.01, 0.8]$ and $\gamma \in \{scale, auto, 0.01, 0.1, 1\}$ for the RBF kernel, using a randomized search on the tuning split to obtain the best configuration. In addition, we report a sensitivity analysis to visualize how performance varies over the (ν, γ) space.

As shown in Figure 4, performance remains high within a compact region of the (ν, γ) grid, indicating that the method is not overly sensitive to a single parameter setting. This visualization complements the randomized search by explicitly showing the stability of the chosen configuration.

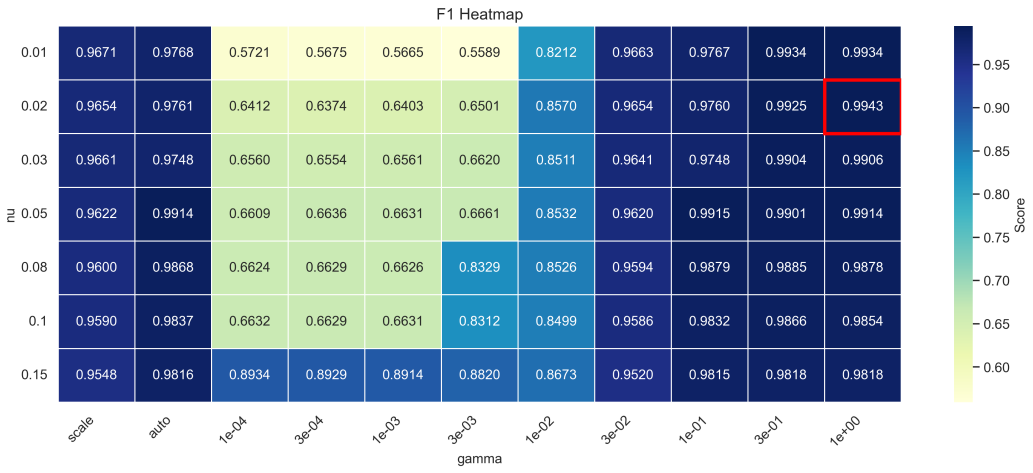

Figure 4. F1-score sensitivity of the RBF OC-SVM over ν and γ on the tuning split. The red box highlights the best-performing (ν, γ) pair (maximum F1).

We selected a broad yet practical range for ν and γ to cover both conservative and more sensitive anomaly settings without making the search computationally prohibitive.

The parameter ν controls the expected fraction of outliers and the tightness of the decision boundary; therefore, we scanned $\nu \in [0.01, 0.8]$ to capture scenarios from low false-alarm operation to more permissive detection. For the RBF kernel, γ determines the locality of the decision function; we included `scale` and `auto` as standard data-driven defaults, and added discrete values $\{0.01, 0.1, 1\}$ to represent coarse-to-fine kernel widths (from smoother boundaries to more localized, high-variance boundaries).

Although the heatmap is shown for a representative dataset, the same search space and tuning protocol were applied consistently across all datasets.

5.1.2. Statistical Analysis

To verify that the performance improvements introduced by the PCA stage are statistically robust and not artifacts of a particular random split, we conducted a paired significance test on the ECU-IoHT dataset. We repeated the full training and evaluation procedure for 10 independent runs using different random seeds (42–51). For each seed, we obtained paired results for two conditions: (i) the baseline OC-SVM pipeline operating on the original feature space, and (ii) the proposed PCA + OC-SVM pipeline. We then applied the Wilcoxon signed-rank test, a non-parametric paired-sample test suitable for repeated-measures comparisons without assuming normality (two-sided, $\alpha = 0.05$).

Table 9 presents the results of this statistical analysis. As can be seen from the ECU-IoHT results in Table 7, PCA improves both Accuracy (98.27% $\rightarrow$ 98.58%) and F1-score (99.09% $\rightarrow$ 99.25%). The Wilcoxon test confirms that these improvements are statistically significant for both metrics ($W = 0.0$, $p = 0.001953$), indicating that the PCA-enhanced pipeline consistently outperforms the baseline across all 10 random seeds. The effect size ($r = 1.0$) suggests strong and consistent improvement pattern.

Due to computational constraints, this formal statistical analysis was conducted only on ECU-IoHT. Extending the same repeated-run testing to the remaining three datasets is left for future work, though the consistent patterns observed in Table 7 suggest similar robustness across datasets.

Table 9. Wilcoxon signed-rank test results on ECU-IoHT over 10 paired runs (seeds 42–51), comparing the baseline OC-SVM pipeline against PCA + OC-SVM. Reported values are mean $\pm$ std across repeats.

Metric	Baseline (Mean $\pm$ Std)	PCA + OC-SVM (Mean $\pm$ Std)	W	p-Value	r
F1-score	99.09 ± 0.01	99.23 ± 0.01	0.0	0.001953	1.0
Accuracy	98.27 ± 0.02	98.54 ± 0.03	0.0	0.001953	1.0

5.1.3. Leave-One-Attack-Out (LOAO) Zero-Day Evaluation

While a normal-only training paradigm is often argued to be suitable for zero-day detection, a standard random split may still evaluate the model on attack instances drawn from the dataset's overall attack distribution. To provide a stricter zero-day stress test, we additionally conducted a leave-one-attack-out (LOAO) experiment on the ECU-IoHT dataset, which contains multiple distinct attack categories (Smurf, Nmap Port Scan, ARP Spoofing, and DoS) in addition to benign traffic. In LOAO, one attack type is completely withheld from the tuning/validation stage and is used only at test time. Concretely, for each held-out attack type b, we (i) fit the scaler and PCA using benign-only training flows, (ii) select (ν, γ) and the PCA target variance level on a validation set composed of benign flows plus seen attack types (all attacks except b), and (iii) evaluate on a zero-day test set containing benign test flows plus only the unseen attack type b. This protocol ensures that

the held-out attack type has no influence on preprocessing, hyperparameter selection, or model fitting, aligning with the leakage-safe training protocol.

Table 10 reports per-attack detection rates under the standard split (seen) versus LOAO (unseen), together with the overall LOAO test F1-score and the false positive rate (FPR) on benign traffic. The results indicate that the detector generalizes strongly to unseen Smurf and ARP Spoofing attacks, and maintains similar detection for Nmap Port Scan under the zero-day setting. DoS remains the most challenging category, where the zero-day F1-score is lower due to the combination of moderate detection rate and non-trivial false-alarm rate on benign traffic. Overall, LOAO supports the claim that a benign-only boundary can detect several attack types without prior exposure, while also highlighting that zero-day performance is attack-dependent and may degrade for attack patterns that overlap more with benign behavior.

Table 10. Leave-one-attack-out (LOAO) zero-day evaluation on ECU-IoHT. For each run, one attack type is fully withheld from tuning/validation and used only at test time (together with held-out benign traffic). "Seen TPR" denotes the per-attack detection rate under the standard evaluation split, whereas "Unseen TPR" reports detection on the held-out (zero-day) attack type. ρ is the target cumulative variance retained by PCA.

Holdout Attack	Seen TPR (%)	Unseen TPR (%)	Δ (pp)
Smurf Attack	100.00	100.00	+0.00
Nmap Port Scan	93.61	93.67	+0.06
ARP Spoofing	100.00	100.00	+0.00
DoS Attack	57.28	58.84	+1.56

5.2. Performance on BoT-IoT Dataset

On the BoT-IoT dataset, a benchmark known for its large scale and inclusion of diverse botnet attacks such as DDoS, DoS, and information theft, the proposed OC-SVM model demonstrates high performance. It achieves a 99.92% F1-score and 99.84% accuracy. This result is particularly noteworthy because the model was trained only on normal traffic, yet it could effectively identify a wide array of attack vectors. The performance surpasses or matches numerous state-of-the-art supervised models listed in the study's comparative analysis, underscoring the efficacy of the one-class anomaly detection approach, especially in a dataset with severe class imbalance. The model's success highlights its ability to establish a precise boundary for normal network behavior, enabling it to detect sophisticated attacks without prior exposure to them. The scree plot in Figure 5 (top) shows that explained variance saturates within a small number of principal components, indicating that BoT-IoT traffic admits a compact representation. The performance curves in Figure 5 (bottom) hug the ideal ROC and PR corners, evidencing strong separability at low false-positive rates and aligning with the reported 99.92% F1 and 99.84% accuracy.

The model's effectiveness is further detailed in the confusion matrix shown in Figure 6, which visually represents the classification results, highlighting the low number of misclassified instances.

The per-class performance breakdown in Figure 7 offers a more granular view of these results. The model demonstrates exceptional capability by achieving 100% detection rates for all high-volume DDoS and DoS attacks, as well as for targeted Theft attacks. Its performance is more nuanced for reconnaissance activities, detecting 'Recon-OS Fingerprint' attacks with an 80.9% rate, which highlights the stealthy nature of such scans. A critical observation from the right panel is the model's handling of normal traffic: 78.9% of normal instances are misclassified as attacks (False Positives). This indicates that in a scenario with severe class imbalance, the OC-SVM creates an extremely tight boundary around the sparse

normal data, prioritizing the detection of any deviation at the cost of a higher false alarm rate for benign traffic.

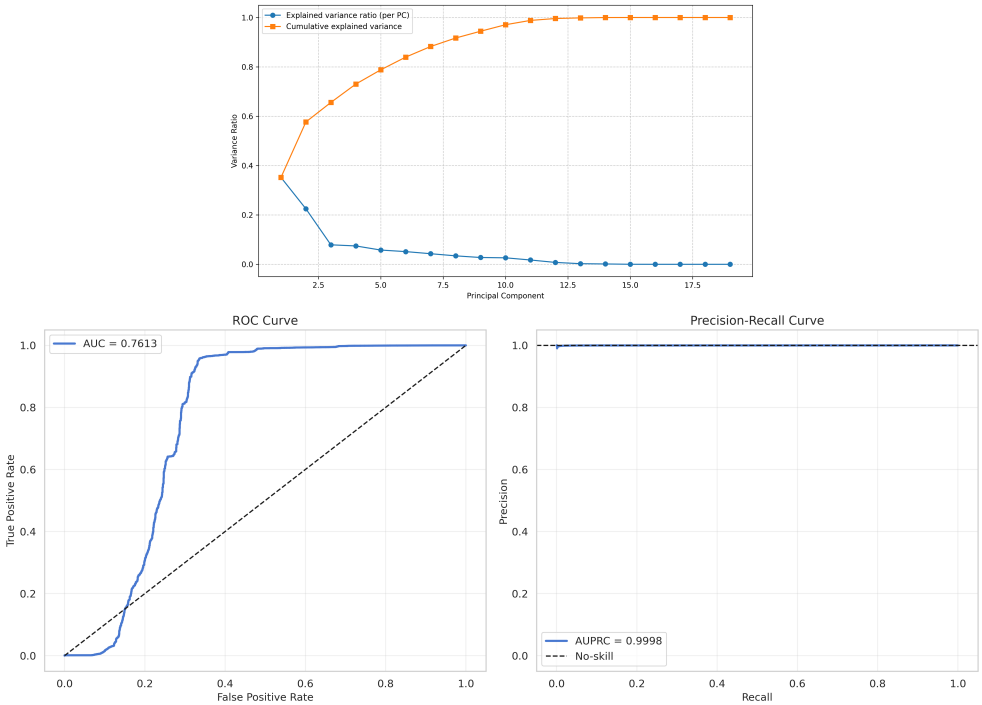

Figure 5. PCA scree plot (**top**) and model performance curves (**bottom**) for the BoT-IoT dataset.

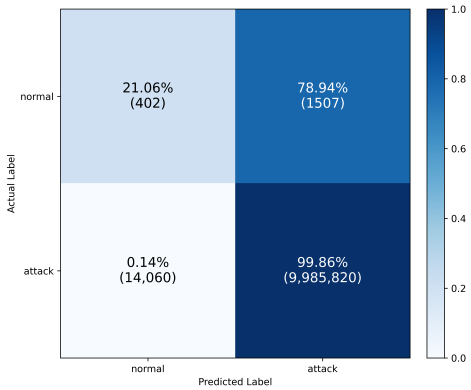

Figure 6. Confusion matrix for the BoT-IoT dataset.

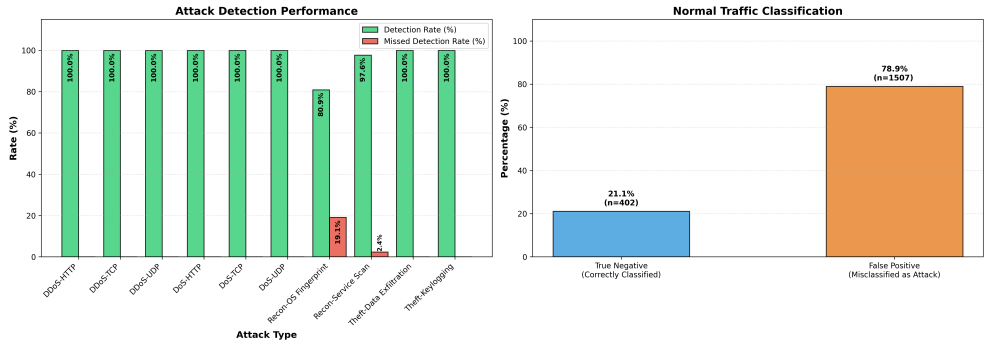

Figure 7. Per-class performance breakdown on the BoT-IoT dataset.

Despite the near-perfect detection of BoT–IoT attacks, this performance comes at the cost of a very high false positive rate on the extremely small normal class. In our

experiments, 78.9% of benign BoT–IoT flows are flagged as anomalous. This behavior reflects an explicit sensitivity–specificity trade-off under extreme class imbalance: the one-class objective encourages the OC–SVM to draw a tight boundary around sparse benign samples, prioritizing attack sensitivity at the expense of benign specificity. In practical deployments, this trade-off can be mitigated by calibrating the decision threshold toward a target FPR, augmenting the benign training set with additional in-situ normal traffic, or using the one-class detector as a high-recall pre-filter whose alarms are then triaged by a second, more precise stage. We therefore highlight the BoT–IoT false-positive behavior primarily as a limitation of the dataset's extreme imbalance rather than of the one-class paradigm itself.

5.3. Performance on CICIoMT2024 Dataset

The CICIoMT2024 dataset is the most comprehensive and modern benchmark, featuring 18 distinct attack types across diverse protocols (Wi-Fi, MQTT, BLE). In this highly complex environment, our model achieved an accuracy of 99.77% and an F1-score of 99.88%. This high performance indicates that the OC-SVM, when properly tuned, can create an extremely precise and tight decision boundary around complex, multi-protocol normal traffic. The model's ability to generalize from a single 'normal' class to effectively reject a wide array of sophisticated attacks—from DDoS floods to reconnaissance scans and malformed MQTT packets—is remarkable. These results are on par with, and in terms of F1-score, superior to, highly complex supervised deep learning models, demonstrating the power and efficiency of our approach. Figure 8 (top) demonstrates that multi-protocol normal traffic compresses efficiently, with variance concentrated in the first few components. Figure 8 (bottom) shows near-perfect ROC and PR behavior across a wide recall range, indicating that the OC-SVM generalizes well to diverse attack families despite training on normal data only.

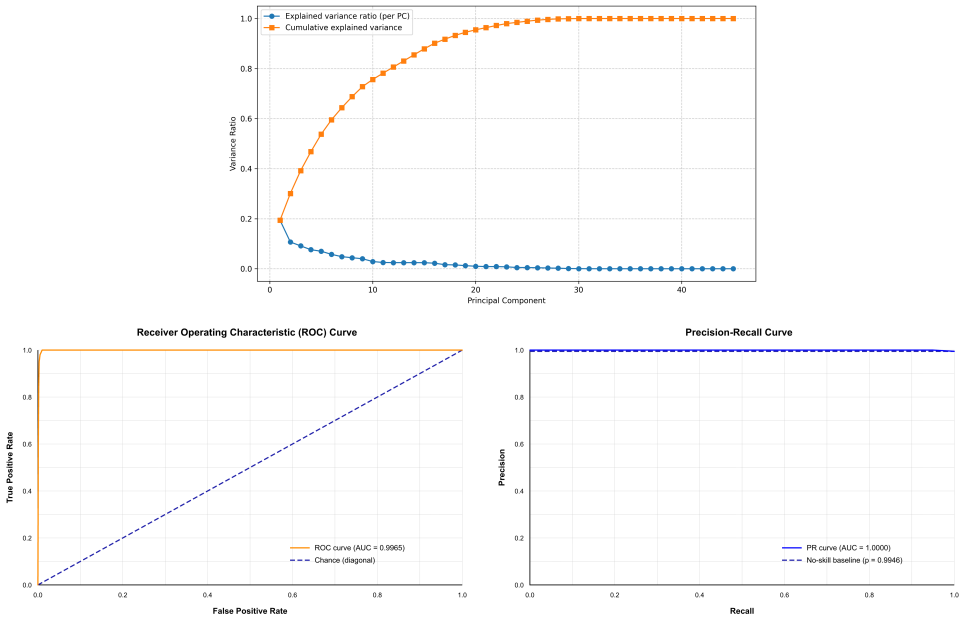

Figure 8. PCA scree plot (**top**) and model performance curves (**bottom**) for the CICIoMT2024 dataset.

The confusion matrix in Figure 9 provides a granular view of the classification performance, confirming the model's high accuracy in distinguishing between the vast number of attack instances and normal traffic.

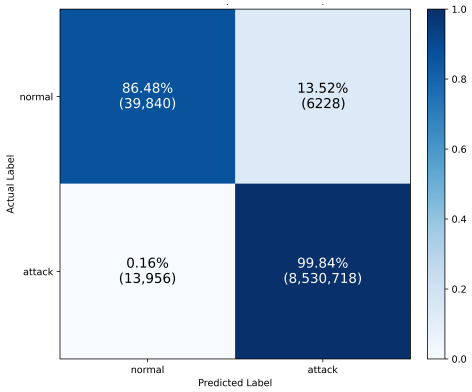

Figure 9. Confusion matrix for the CICIoMT2024 dataset.

The detailed performance breakdown in Figure 10 further clarifies this strong result. The model achieves perfect 100% detection rates across all eight TCP/IP-based DDoS and DoS attack variants, proving its efficacy against high-volume network floods. For more subtle, protocol-specific attacks, the results are varied: 'ARP Spoofing' is detected with a 64.9% rate and 'MQTT-Malformed' packets with 61.8%, suggesting these attack signatures are harder to distinguish from the normal behavior learned by the model. Crucially, the model demonstrates a well-balanced approach to normal traffic, correctly identifying 86.5% of instances (True Negatives) while maintaining a low False Positive rate of 13.5%. This balance is vital, showcasing the model's ability to form a robust yet flexible decision boundary in a highly complex, multi-protocol environment.

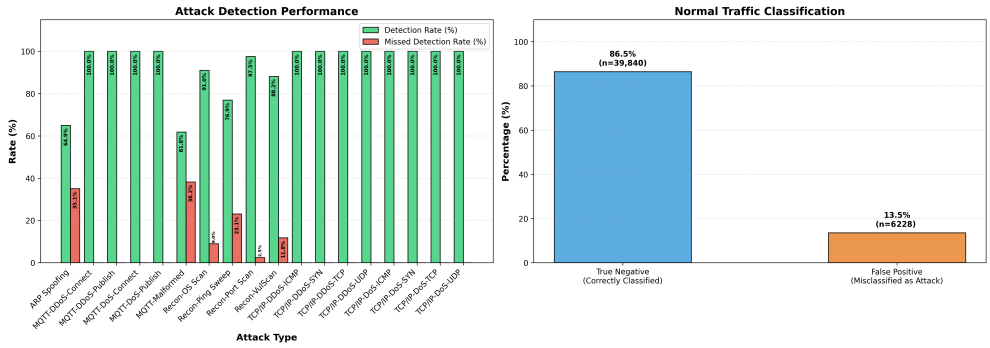

Figure 10. Per-class performance breakdown on the CICIoMT2024 dataset.

5.4. Performance on ECU-IoHT Dataset

For the ECU-IoHT dataset, which contains classic network attacks such as Smurf, ARP Spoofing, and DoS, the model demonstrated outstanding robustness by achieving 98.54% accuracy and a 99.25% F1-score. Of particular note is the recall of 99.19%, which signifies that the model correctly identified nearly all true attack instances, minimizing the critical risk of false negatives in a healthcare setting. This proves that our lightweight, unsupervised approach, which is far less demanding in terms of data labeling and computational resources, can match the detection capabilities of more complex, data-hungry architectures. As seen in Figure 11 (top), the scree plot exhibits a clear elbow, suggesting that a modest number of components is sufficient to capture the data structure. The curves in Figure 11 (bottom) maintain high true-positive rates at very low false-positive regions, with a PR curve close to unity—consistent with the high recall reported for classical attacks.

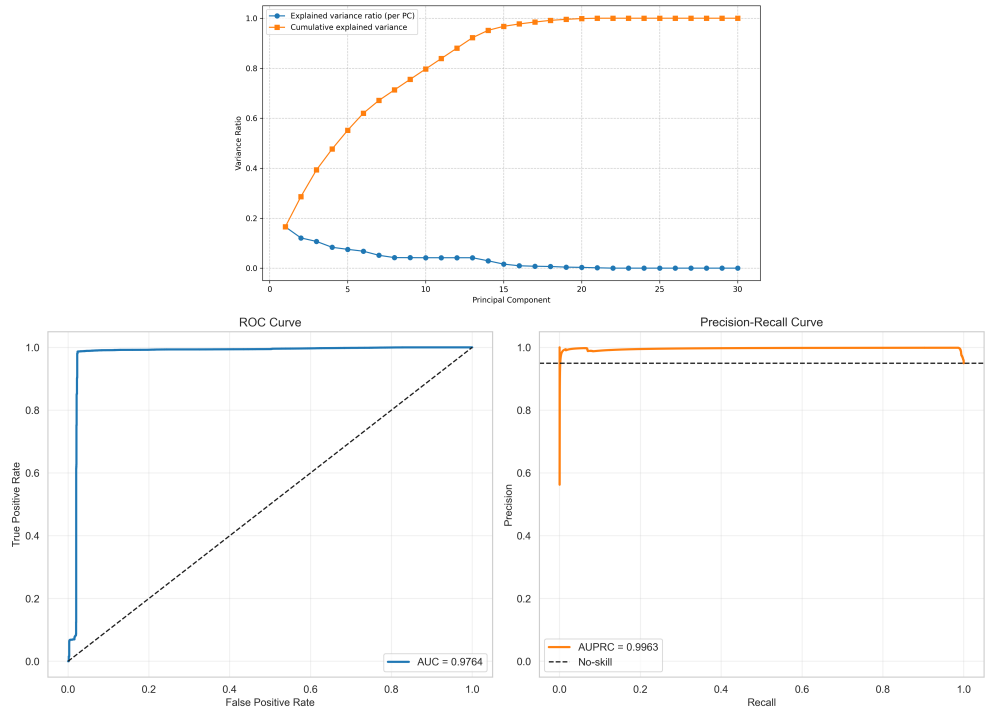

Figure 11. PCA scree plot (**top**) and model performance curves (**bottom**) for the ECU-IoHT dataset.

Figure 12 displays the confusion matrix, illustrating the model's strong performance with a very high number of correctly identified attacks (true positives) and normal instances (true negatives).

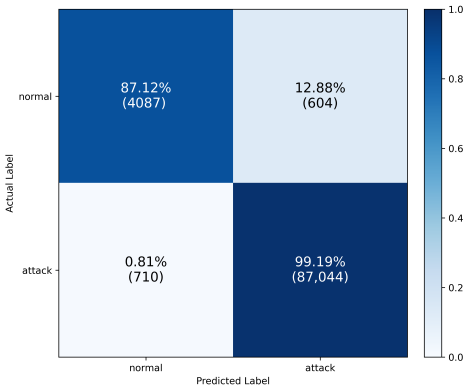

Figure 12. Confusion matrix for the ECU-IoHT dataset.

A detailed per-class analysis, presented in Figure 13, reveals specific strengths and weaknesses. The model provides 100% detection for both 'ARP Spoofing' and 'Smurf Attack', demonstrating its complete effectiveness against these classic threats. However, it faces a significant challenge with the 'DoS Attack' category, achieving only a 57.3% detection rate and consequently missing 42.7% of these attacks. This suggests that the traffic patterns of this particular DoS attack closely mimic the normal behavior profile learned by the OC-SVM. Regarding normal traffic, the model performs robustly, correctly classifying 87.1% of instances and maintaining a low false positive rate of 12.9%, which is essential for minimizing unnecessary alerts and operational disruptions in a real-world healthcare setting.

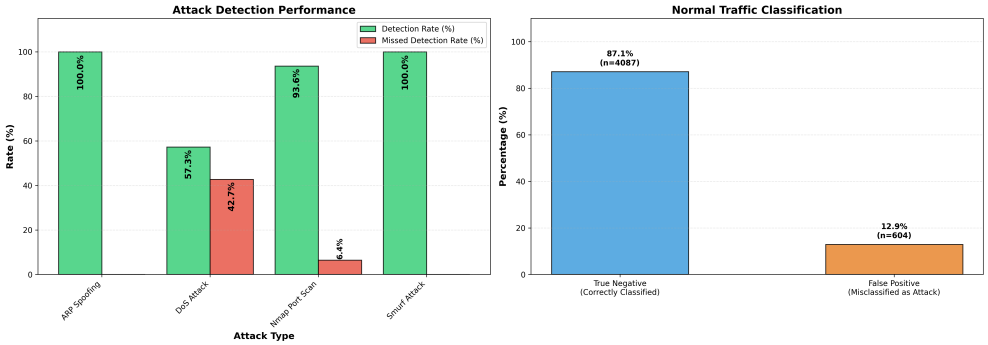

Figure 13. Per-class performance breakdown on the ECU-IoHT dataset.

5.5. Performance on IoMT-TrafficData Dataset

The IoMT-TrafficData dataset emulates a realistic scenario, providing a crucial test of the model's practical applicability. Here, our model achieved a strong accuracy of 98.66% and an F1-score of 99.19%. The model's success in this environment underscores its real-world viability, as it effectively learned the baseline of a specific clinical setup and detected a range of IP-based attacks. The high recall of 99.34% on this dataset further emphasizes the model's reliability in catching threats. This performance, achieved through a unified preprocessing and training pipeline, confirms that our methodology is not just a theoretical exercise but a robust and adaptable solution for securing real-world IoMT deployments. The scree plot in Figure 14 (top) indicates that flow-level features are highly compressible, enabling lightweight models. Correspondingly, Figure 14 (bottom) displays ROC and PR curves characteristic of strong class separation, underscoring the method's practical utility in realistic clinical network settings.

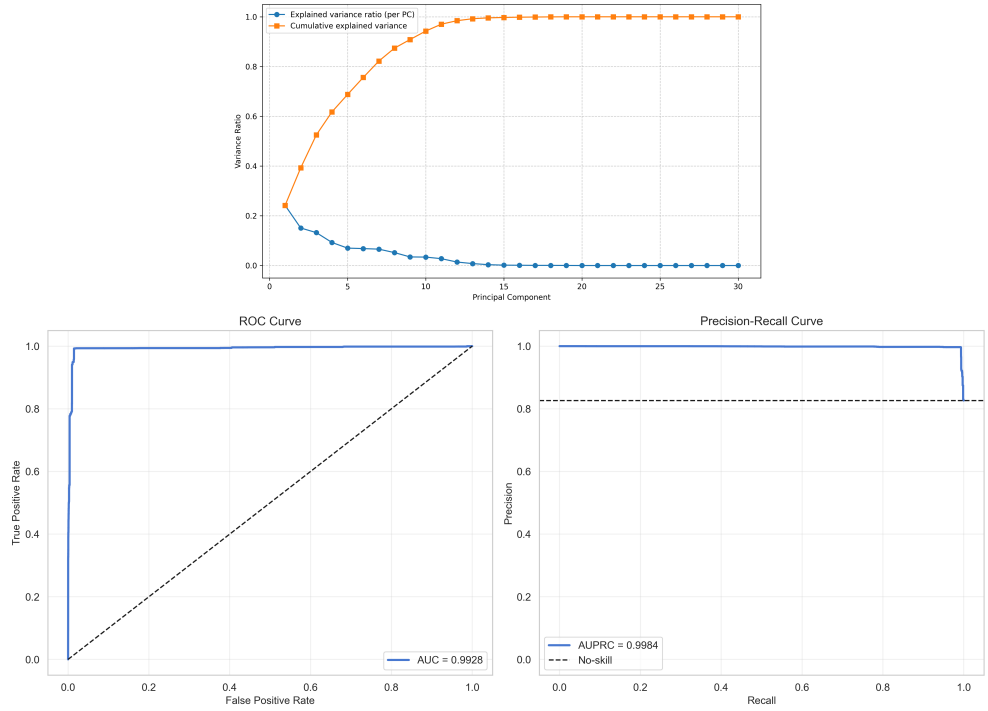

Figure 14. PCA scree plot (**top**) and model performance curves (**bottom**) for the IoMT-TrafficData dataset.

The corresponding confusion matrix for the IoMT-TrafficData dataset is presented in Figure 15, corroborating the high F1-score and accuracy metrics with a clear visualization of correct and incorrect classifications.

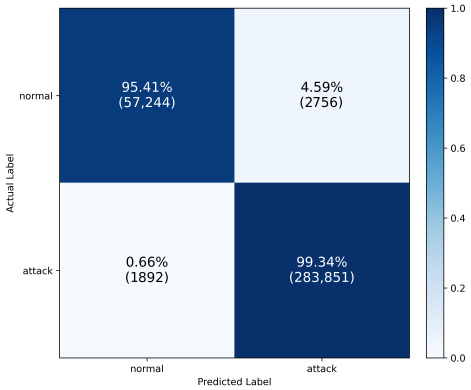

Figure 15. Confusion matrix for the IoMT-TrafficData dataset.

The per-class performance breakdown in Figure 16 underscores this real-world viability. The model's detection rates are exceptionally high and consistent across all eight attack types, with most exceeding 99% and perfect 100% detection for 'CAM Overflow' and 'Net Scan'. The minimal missed detection rates across the board highlight the model's reliability in a practical clinical environment. Most impressively, the model achieves its best performance on normal traffic classification with this dataset, attaining a 95.4% True Negative rate and an extremely low 4.6% False Positive rate. This outstanding balance between high sensitivity to attacks and high specificity for normal traffic strongly validates our methodology as a robust and adaptable solution for securing real-world IoMT deployments.

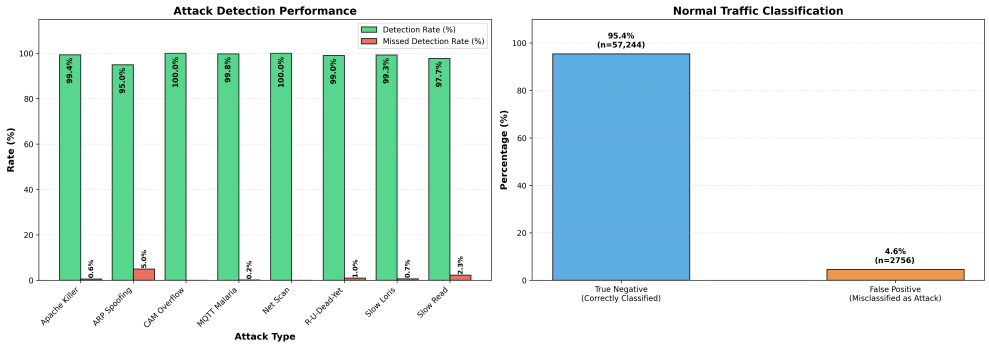

Figure 16. Per-class performance breakdown on the IoMT-TrafficData dataset.

5.6. Computational Efficiency and Deployment Footprint

Beyond detection quality, we also quantify the computational cost of the PCA + OC–SVM pipeline using the measurements reported in our computational-efficiency study. For each corpus we record the size of the serialized model (scaler, PCA and OC–SVM parameters), and the peak additional RAM usage during inference. Table 11 summarises these results.

Table 11. Computational footprint of the PCA + OC-SVM detector across IoMT datasets.

Dataset	Model Size (MB)	Avg Per-Sample (ms)	Total Batch Inference Time
BoT-IoT	0.161	0.1570	1570 s for 10,001,789 samples
CICIoMT2024	1.439	0.8418	7231 s for 8,590,742 samples
ECU-IoHT	0.130	0.0861	7 s for 92,445 samples
IoMT-TrafficData	1.091	0.3505	121 s for 345,743 samples

As shown in Table 11, training and tuning complete within practical time budgets for all four corpora, and inference remains fast enough to process IoMT flows in (near) real time on standard gateway hardware. The additional memory footprint and final model size

are modest, which makes the proposed PCA + OC–SVM pipeline suitable for deployment on resource-constrained edge devices and clinical gateways.

5.7. Comparison with State-of-the-Art Results

In this study, an effective system for detecting cyberattacks in IoMT environments has been developed using a One-Class Support Vector Machine (OC-SVM) model trained exclusively on normal traffic. As detailed in Table 12, experiments conducted on four distinct IoMT datasets—BoT-IoT, ECU-IoHT, CICIoMT2024, and IoMT-TrafficData—demonstrate that the proposed approach achieves high performance, often surpassing existing state-of-the-art methods that rely on more complex, supervised learning architectures.

Table 12. IoMT model performance comparison.

Reference	Dataset	Acc (%)	Prec (%)	Rec (%)	F1-Score (%)
Hussan et al. [48]	BoT-IoT	99.02	99.75	-	98.35
Katib & Ragab [16]		99.05	96.65	95.67	96.14
Saheed et al. [25]		99.41	-	99.78	-
Sagu et al. [22]		92.96	93.00	-	92.78
Our Model		**99.84**	**99.98**	**99.86**	**99.92**
Kavkas & Yildiz [20]	CICIoMT2024	99.00	99.00	99.00	99.00
Hernandez-Jaimes et al. [28]		96.25	98.46	95.57	96.99
Mohammadi et al. [18]		99.00	99.00	99.00	99.00
Shaikh et al. [21]		99.58	98.81	98.21	98.57
Berqia et al. [23]		99.00	89.00	97.00	92.00
Dadkhah et al. [29]		99.50	95.20	94.00	94.60
Kirubavathi et al. [19]		96.00	97.00	94.00	95.00
Our Model		**99.77**	**99.93**	**99.84**	**99.88**
Kumar et al. [26]	ECU-IoHT	98.55	90.08	92.03	87.77
Alalwany et al. [27]		98.00	98.00	99.00	98.00
Kumar & Kim [17]		97.16	93.96	79.33	82.37
Mosaiyebzadeh et al. [15]		98.47	93.88	99.18	-
Our Model		**98.58**	**99.31**	**99.19**	**99.25**
Ebinazer [24]	IoMT-TrafficData	95.00	96.32	99.00	94.00
Areia et al. [47]		87.90	87.90	90.09	88.89
Kouekam [49]		94.89	98.08	94.89	96.19
Our Model		**98.66**	**99.04**	**99.34**	**99.19**

A comparison with these studies, presented in Table 12, clearly shows that our OC-SVM model achieves results competitive with or superior to current state-of-the-art methods across all tested datasets. The overall success of our model stems from its core design principles:

Effectiveness of the One-Class Anomaly Detection Paradigm: Unlike supervised models that require comprehensive labels for both normal and attack classes, our approach models only the baseline of normal behavior. This makes it inherently capable of detecting novel or zero-day attacks not seen during training, addressing a fundamental limitation of supervised methods. This is critical in the ever-evolving landscape of IoMT threats.

Inherent Robustness to Class Imbalance: The most significant challenge in network intrusion detection is the severe imbalance between the vast amount of normal traffic and the rare occurrence of attacks. Our model elegantly bypasses this issue by not requiring attack samples for training; its performance is therefore not skewed by the minority attack class, a common pitfall for supervised classifiers.

Methodological Rigor in Hyperparameter Optimization: The high performance is not accidental but a direct result of our rigorous methodology. The use of RandomizedSearchCV with a PredefinedSplit, optimized specifically for the F1-score, ensures that the crucial ν and

γ hyperparameters are tuned to find a perfect balance between precision and recall, rather than just optimizing for raw accuracy which can be misleading in imbalanced scenarios.

Practicality for Edge Deployment: The OC-SVM approach yields a compact and computationally efficient model (Table 11), which makes the framework amenable to deployment on resource-constrained IoMT edge gateways. While these results indicate strong potential for real-time detection, the end-to-end throughput, on-device latency, and energy consumption on representative edge hardware are left to future work.

6. Discussion

The results across four IoMT corpora indicate that the PCA stage is most beneficial when the original feature space contains substantial redundancy and noise, because projecting traffic into a compact set of principal components simplifies the geometry that the one-class boundary must capture. This behavior is consistent with the ablation findings (Table 7), where PCA improves detection on some datasets while largely preserving performance on others.

A notable exception arises under extreme class imbalance and sparse benign coverage, where the one-class objective may form an overly tight description of normal behavior. In such regimes, the detector can prioritize sensitivity to deviations at the expense of benign specificity, yielding elevated false alarms (as observed in the BoT-IoT per-class analysis). In practice, this sensitivity-specificity trade-off can be managed by calibrating the decision threshold toward a target false-positive rate, expanding the benign training set with additional in-situ normal traffic, or using the one-class detector as a high-recall pre-filter followed by a more precise triage stage.

In practice, this sensitivity-specificity trade-off can be managed by calibrating the decision threshold toward a target false-positive rate, expanding the benign training set with additional in-situ normal traffic, or using the one-class detector as a high-recall pre-filter followed by a more precise triage stage. From a deployment perspective, the comparison of alternative representations and detectors further supports the proposed design. While non-linear representation learning (e.g., AE before OC-SVM) can be competitive, it introduces additional training/inference overhead and extra hyperparameters, weakening the lightweight motivation of the pipeline (Table 7). Likewise, when keeping the same PCA representation, OC-SVM remains a robust second-stage detector compared to alternatives such as Isolation Forest (Table 8), reinforcing PCA + OC-SVM as a practical accuracy-efficiency balance for resource-constrained IoMT environments.

7. Conclusions

This study introduced and systematically evaluated a two-stage anomaly detection pipeline combining Principal Component Analysis (PCA) and a One-Class Support Vector Machine (OC-SVM) for cyberattack detection in IoMT networks. By training the model exclusively on the principal components of normal traffic data, our approach successfully addresses key challenges in IoMT security. Extensive experiments on four diverse datasets (BoT-IoT, CICIoMT2024, ECU-IoHT, and IoMT-TrafficData) validated the model's high efficacy, achieving superior F1-scores of 99.92%, 99.88%, 99.25%, and 99.19%, respectively.

The success of this paradigm demonstrates that it is possible to achieve state-of-the-art performance without relying on scarce, incomplete, or evolving attack labels, thereby addressing a critical challenge in real-world cybersecurity applications. The integration of PCA proved to be a critical component, primarily by providing a substantial reduction in data dimensionality. This enhancement in computational efficiency is paramount for deployment on resource-constrained IoMT devices. Our ablation study demonstrated that this efficiency gain is often accompanied by a significant improvement in detection perfor-

mance, particularly on complex, high-dimensional datasets. While a minor performance trade-off was observed on one dataset, the overall results confirm that PCA effectively balances high performance with practical deployability. This methodology, combined with our rigorous, F1-score-driven hyperparameter tuning, results in a model that generalizes effectively across different network protocols and attack types, making it a robust and versatile solution.

The high F1-scores across all experiments underscore the model's balanced precision and recall, highlighting its potential as a practical, lightweight, and scalable security solution. Its computational efficiency makes it particularly suitable for deployment on resource-constrained edge devices within clinical networks, ensuring timely detection where it matters most. This work establishes that modeling normal behavior, when executed with a methodologically sound approach, is a powerful and sufficient strategy for securing modern IoMT infrastructures against both known and unknown threats.

Although the proposed PCA-OC-SVM framework achieves strong detection performance and improved efficiency across multiple public IoMT corpora, some limitations remain. First, our evaluation is conducted on offline datasets and general-purpose computing hardware; therefore, we have not yet validated end-to-end throughput, latency, and energy consumption on resource-constrained edge gateways in real clinical deployments. Second, the experiments rely on publicly available traces, which may not fully capture the diversity of devices, protocols, and operational drift observed in real hospital networks. Future work will address these limitations by (i) testing the pipeline on representative edge platforms and reporting real-time performance metrics, (ii) expanding the study with more natural, heterogeneous, and continuously evolving IoMT datasets, and (iii) integrating explainable AI components to improve transparency and practitioner trust. In addition, we plan to explore privacy-preserving federated learning to enable collaborative model development across decentralized institutions without sharing raw traffic, ultimately advancing the framework toward a more robust, ethical, and widely deployable IoMT security solution.

Author Contributions: E.G.: Writing—original draft preparation, software, formal analysis. B.U.: Writing—review and editing, methodology, formal analysis, conceptualization. G.U.: Writing—review and editing, methodology, supervision, formal analysis, conceptualization. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by Office of Scientific Research Projects of Karadeniz Technical University. Project number: FBA-2026-17392.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All datasets used in this study are publicly available benchmark datasets.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kumar, S.; Tiwari, P.; Zymbler, M. Internet of Things is a revolutionary approach for future technology enhancement: A review. *J. Big Data* **2019**, *6*, 111. [CrossRef]
2. Gubbi, J.; Buyya, R.; Marusic, S.; Palaniswami, M. Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Gener. Comput. Syst.* **2013**, *29*, 1645–1660. [CrossRef]
3. Miorandi, D.; Sicari, S.; De Pellegrini, F.; Chlamtac, I. Internet of things: Vision, applications and research challenges. *Ad Hoc Netw.* **2012**, *10*, 1497–1516. [CrossRef]
4. Al-Fuqaha, A.; Guizani, M.; Mohammadi, M.; Aledhari, M.; Ayyash, M. Internet of things: A survey on enabling technologies, protocols, and applications. *IEEE Commun. Surv. Tutor.* **2015**, *17*, 2347–2376. [CrossRef]
5. Atzori, L.; Iera, A.; Morabito, G. The Internet of Things: A survey. *Comput. Netw.* **2010**, *54*, 2787–2805. [CrossRef]

6. Zanella, A.; Bui, N.; Castellani, A.; Vangelista, L.; Zorzi, M. Internet of Things for smart cities. *IEEE Internet Things J.* **2014**, *1*, 22–32. [CrossRef]
7. Whitmore, A.; Agarwal, A.; Da Xu, L. The Internet of Things—A survey of topics and trends. *Inf. Syst. Front.* **2015**, *17*, 261–274. [CrossRef]
8. Islam, S.R.; Kwak, D.; Kabir, M.H.; Hossain, M.; Kwak, K.S. The internet of things for health care: A comprehensive survey. *IEEE Access* **2015**, *3*, 678–708. [CrossRef]
9. Statista Research Department. Internet of Things (IoT) Statistics & Facts. 2003. Available online: https://www.statista.com/outlook/tmo/internet-of-things/worldwide (accessed on 15 November 2025).
10. Tawalbeh, L.; Muheidat, F.; Tawalbeh, M.; Quwaider, M. IoT privacy and security: Challenges and solutions. *Appl. Sci.* **2020**, *10*, 4102. [CrossRef]
11. Andy, S.; Rahardjo, B.; Hanindhito, B. Attack scenarios and security analysis of MQTT communication protocol in IoT system. In Proceedings of the 2017 4th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI), Yogyakarta, Indonesia, 19–21 September 2017; pp. 1–6. [CrossRef]
12. Vaccari, I.; Aiello, M.; Cambiaso, E. SlowITe, a novel denial of service attack affecting MQTT. *Sensors* **2020**, *20*, 2932. [CrossRef]
13. Iqbal, W.; Abbas, H.; Daneshmand, M.; Rauf, B.; Bangash, Y.A. An in-depth analysis of IoT security requirements, challenges, and their countermeasures via software-defined security. *IEEE Internet Things J.* **2020**, *7*, 10250–10276. [CrossRef]
14. Chandola, V.; Banerjee, A.; Kumar, V. Anomaly Detection: A Survey. *ACM Comput. Surv.* **2009**, *41*, 15. [CrossRef]
15. Mosaiyebzadeh, F.; Pouriyeh, S.; Parizi, R.M.; Han, M.; Batista, D.M. Intrusion detection system for IoHT devices using federated learning. In Proceedings of the IEEE INFOCOM Workshops: AidTSP 2023 International Workshop on AI-Driven Trustworthy, Secure, and Privacy-Preserving Computing, Hoboken, NJ, USA, 20–20 May 2023; pp. 1–6. [CrossRef]
16. Katib, I.; Ragab, M. Blockchain-Assisted Hybrid Harris Hawks Optimization Based Deep DDoS Attack Detection in the IoT Environment. *Mathematics* **2023**, *11*, 1887. [CrossRef]
17. Kumar, M.; Kim, S. Securing the Internet of Health Things: Embedded federated learning-driven long short-term memory for cyberattack detection. *Electronics* **2024**, *13*, 3461. [CrossRef]
18. Mohammadi, A.; Aminian, M.; Ghahramani, H.; Asghari, S.A. Securing healthcare with deep learning: A CNN-based model for medical IoT threat detection. In Proceedings of the 2024 19th Iranian Conference on Intelligent Systems (ICIS), Sirjan, Iran, 23–24 October 2024; pp. 168–174. [CrossRef]
19. Kirubavathi, G.; Sivakumar, S.; Kumar, S.; Manickam, S. Optimized hybrid approach for anomaly detection of DDoS and network attacks in IoMT systems using autoencoders and TabNet. In Proceedings of the 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC), Tumkuru, India, 4–5 December 2024; pp. 1–6. [CrossRef]
20. Kavkas, N.C.; Yildiz, K. Enhancing IoMT security with deep learning based approach for medical IoT threat detection. In Proceedings of the 13th International Symposium on Digital Forensics and Security (ISDFS), Boston, MA, USA, 24–25 April 2025; pp. 1–6. [CrossRef]
21. Shaikh, J.A.; Wang, C.; Us Sima, M.W.; Arshad, M.; Owais, M.; Hassan, D.S.M.; Alkanhel, R.; Muthanna, M.S.A. A deep reinforcement learning-based robust intrusion detection system for securing IoMT healthcare networks. *Front. Med.* **2025**, *12*, 1524286. [CrossRef]
22. Sagu, A.; Gill, N.S.; Gulia, P.; Alduaiji, N.; Shukla, P.K.; Shah, M.A. Advances to IoT security using a GRU-CNN deep learning model trained on SUCMO algorithm. *Sci. Rep.* **2025**, *15*, 16485. [CrossRef]
23. Berqia, A.; Bouijij, H.; Chahbi, M.; Ismaili, O. Predicting cyberattacks on connected healthcare devices using LightGBM-based approach. In Proceedings of the 13th International Symposium on Digital Forensics and Security (ISDFS), Boston, MA, USA, 24–25 April 2025; pp. 1–7. [CrossRef]
24. Ebinazer, S.E. Hybrid encryption with greylag goose optimizer-based swin transformer for secured data deduplication and anomaly detection for cloud-based IoMT applications. *SSRN* **2025**. [CrossRef]
25. Saheed, Y.K.; Abdulganiyu, O.H.; Tchakoucht, T.A. Modified genetic algorithm and fine-tuned long short-term memory network for intrusion detection in the internet of things networks with edge capabilities. *Appl. Soft Comput. J.* **2024**, *155*, 111434. [CrossRef]
26. Kumar, M.; Singh, S.K.; Kim, S. Hybrid deep learning-based cyberthreat detection and IoMT data authentication model in smart healthcare. *Future Gener. Comput. Syst.* **2025**, *166*, 107711. [CrossRef]
27. Alalwany, E.; Alsharif, B.; Alotaibi, Y.; Alfahaid, A.; Mahgoub, I.; Ilyas, M. Stacking ensemble deep learning for real-time intrusion detection in IoMT environments. *Sensors* **2025**, *25*, 624. [CrossRef]
28. Hernandez-Jaimes, M.L.; Martínez-Cruz, A.; Ramírez-Gutiérrez, K.A.; Guevara-Martínez, E. Enhancing machine learning approach based on Nilsimsa fingerprinting for ransomware detection in IoMT. *IEEE Access* **2024**, *12*, 153886–153900. [CrossRef]
29. Dadkhah, R.; Neto, E.C.P.; Ferreira, R.; Molokwu, R.; Sadeghi, S.; Ghorbani, A.A. CICIoMT2024: A benchmark dataset for multi-protocol security assessment in IoMT. *Internet Things* **2024**, *28*, 101351. [CrossRef]
30. Erfani, S.M.; Rajasegarar, S.; Karunasekera, S.; Leckie, C. High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning. *Pattern Recognit.* **2016**, *58*, 121–134. [CrossRef]

31. Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer: New York, NY, USA, 2006.
32. Jolliffe, I.T. *Principal Component Analysis*, 2nd ed.; Springer Series in Statistics; Springer: New York, NY, USA, 2002. [CrossRef]
33. Schölkopf, B.; Platt, J.C.; Shawe-Taylor, J.; Smola, A.J.; Williamson, R.C. Estimating the support of a high-dimensional distribution. *Neural Comput.* **2001**, *13*, 1443–1471. [CrossRef] [PubMed]
34. Tax, D.M.J.; Duin, R.P.W. Support Vector Data Description. *Mach. Learn.* **2004**, *54*, 45–66. [CrossRef]
35. Shin, H.J.; Eom, D.H.; Kim, S.S. One-class support vector machines—An application in machine fault detection and classification. *Comput. Ind. Eng.* **2005**, *48*, 395–408. [CrossRef]
36. Bounsiar, A.; Madden, M.G. One-class Support Vector Machines revisited. In Proceedings of the 2014 International Conference on Information Science & Applications (ICISA), Seoul, Republic of Korea, 6–9 May 2014; pp. 1–6. [CrossRef]
37. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [CrossRef]
38. Jordan, M.I.; Mitchell, T.M. Machine learning: Trends, perspectives, and prospects. *Science* **2015**, *349*, 255–260. [CrossRef]
39. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016.
40. Tavallaee, M.; Bagheri, E.; Lu, W.; Ghorbani, A.A. A detailed analysis of the KDD Cup 99 data set. In Proceedings of the 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications, Ottawa, ON, Canada, 8–10 July 2009; pp. 1–6.
41. Sharafaldin, I.; Lashkari, A.H.; Ghorbani, A.A. Toward generating a new intrusion detection dataset and intrusion traffic characterization. In Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP), Funchal-Madeira, Portugal, 22–24 January 2018; pp. 108–116.
42. Kuhn, M.; Johnson, K. *Applied Predictive Modeling*; Springer: New York, NY, USA, 2013.
43. Mittelstadt, B.D.; Allo, P.; Taddeo, M.; Wachter, S.; Floridi, L. The ethics of algorithms: Mapping the debate. *Big Data Soc.* **2016**, *3*, 2053951716679679. [CrossRef]
44. Boulle, M.; Huynh, T. OpenML: An R package to connect to the machine learning platform. *J. Mach. Learn. Res.* **2020**, *21*, 1–5.
45. Koroniotis, N.; Moustafa, N.; Sitnikova, E.; Turnbull, B. Towards the development of realistic botnet dataset in the Internet of Things for network forensic analytics: Bot-IoT dataset. *Future Gener. Comput. Syst.* **2019**, *100*, 779–796. [CrossRef]
46. Ahmed, M.; Byreddy, S.; Nutakki, A.; Sikos, L.F.; Haskell-Dowland, P. *ECU-IoHT: A Dataset for Internet of Healthcare Things*; Edith Cowan University: Joondalup, Australia, 2021. [CrossRef]
47. Areia, J.; Afonso Bispo, I.; Santos, L.; de Costa, R.L.C. IoMT-TrafficData: Dataset and tools for benchmarking intrusion detection in Internet of Medical Things. *IEEE Access* **2024**, *12*, 115370–115379. [CrossRef]
48. Hussan, M.I.; Reddy, G.V.; Anitha, P.T.; Kanagaraj, A.; Naresh, P. DDoS attack detection in IoT environment using optimized Elman recurrent neural networks based on chaotic bacterial colony optimization. *Clust. Comput.* **2024**, *27*, 4469–4490. [CrossRef]
49. Kouekam, M.; Khennou, F. Advancing IoMT security with privacy-preserving federated learning techniques. In Proceedings of the 13th International Symposium on Digital Forensics and Security (ISDFS), Boston, MA, USA, 24–25 April 2025; pp. 1–6. [CrossRef]