Detecting DDoS Attacks in IoT Healthcare

Eurydice Tracy Makena and Sara Sutton
Grand Valley State University, Allendale, MI, USA

makenae@mail.gvsu.edu
suttosar@gvsu.edu

Abstract: The increased adoption of IoT devices in healthcare domains has significantly increased the attack surface, leaving critical infrastructure vulnerable to attacks such as Distributed Denial of Service (DDoS). Our work investigates the role of feature selection and temporal dependency modelling in detecting MQTT (Message Queuing Telemetry Transport) DDoS attacks using the CICIoMT 2024 dataset. We compare two approaches; using the Naïve Bayes model which assumes flow independence and Long Short-Term Memory (LSTM) model, which captures sequential dependencies in the raw network flows. Feature selection was also performed to reduce 84 raw features to 31 informative features. Our preliminary results show that the LSTM model outperformed naïve bayes by leveraging the temporal patterns that are characteristic of DDoS traffic. This study evaluates the importance of guided feature selection and temporal dependency modelling in DDoS detection in IoT healthcare networks.

Keywords: DDoS detection, IoT healthcare, MQTT, Long short-term memory (LSTM), Intrusion detection

1. Introduction

Internet of Things (IoT) devices continue to be widely adopted across various domains, and it is predicted to rise to over 40 billion devices by 2030. One such domain is in healthcare where IoT devices such as patient monitors, infusion pumps, and wearables interact to enable real-time healthcare delivery. With this adoption comes a wide range of cybersecurity challenges as discussed by (ElSayed et al., 2025) in their paper. This connectivity introduces critical security risk, IoMT devices often run outdated firmware, use weak authentication mechanisms, and run on limited computational resources making it challenging to implement security measures. This makes them easy targets for botnet recruitment and large-scale DDoS attacks that bring down organizations. (Censinet.com, 2026) estimated that the cost of downtime from DDoS attacks is over $100,000 per hour. Service interruptions in healthcare environments directly impact patient safety. Therefore, robust defences must be prioritized to ensure continuity of services.

Among the protocols commonly exploited in IoMT environments, MQTT (Message Queuing Telemetry Transport) is particularly vulnerable due to its lightweight publish/subscribe model, weak authentication mechanisms, and unencrypted message transmission. Attackers exploit MQTT through Connect and Publish Flood attacks that overwhelm brokers and clients with excessive traffic, overwhelming legitimate device communications. In their paper, (Lakshminarayana et al., 2025) evaluated the security vulnerabilities of MQTT v5.0 protocol that made them susceptible to attacks. They found that the protocol's support for large, unbounded properties such as connect properties and user-defined metadata enables attackers to craft oversized or numerous malicious packets that exhaust broker resources like CPU and memory.

Most traditional intrusion detection methods are signature based and rely on static rules to flag suspicious activities. However, as attacks become more sophisticated, these methods struggle to detect varying attack patterns. While machine learning has also advanced beyond signature-based detection, existing methods mostly treat the network flows as independent events and fail to capture the temporal relationships between the flow attributes. Our work addresses two fundamental research questions in DDoS detection, are there specific features that most effectively discriminate between benign and attack traffic and does modelling short-term flow sequences provide quantifiable benefits over independent flow classification in detecting MQTT-based DDoS attacks. To address these questions, we make the following contributions: We perform feature selection through Random Forest importance scoring and Information Gain to identify a smaller, informative feature subset. We then evaluate two contrasting models, the Naive Bayes classifier, which treats each flow independently; and an LSTM model, which captures temporal dependencies across consecutive flows. We demonstrate that temporal modelling achieves better results in detection accuracy and false positive reduction on the MQTT Connect and Publish Flood traffic from the CiCIoMT 2024 dataset.

2. Related Works

With the rise of cyberattacks on IoMT devices, such as DDoS, researchers have made efforts to develop detection algorithms and mitigation techniques.

(Clark et al., 2024) leveraged machine learning models for anomaly detection to identify abnormal activity in IoMT systems with a goal of detecting DoS and ARP spoofing attacks. Their research focuses on Bluetooth, Wi-Fi, and MQTT enabled devices where they trained different supervised and unsupervised machine learning models to classify network traffic as either malicious or benign. The models were trained on the CIC IoMT 2024 dataset, and the findings were that the decision tree model performed best for detecting DoS attacks on Bluetooth devices, while the random forest model performed best for detecting ARP spoofing attacks on Wi-Fi and MQTT enabled devices.

(Hnamte et al., 2022) implemented a Naïve Bayes classifier for DDoS detection in IoT networks. They selected and optimized features using statistical techniques, specifically the correlation coefficient and Dice similarity coefficient. A Naïve bayes classifier was trained to accurately distinguish between malicious and benign traffic while reducing false positives. Their model achieved a high accuracy of 89.7% and true positive rates of about 90%.

(Kirubavathi et al., 2024) proposed a two-model based approach to anomaly detection trained on the CICIoMT 2024 dataset. The models used were Autoencoder for identifying errors through reconstruction errors and TabNet for utilizing attention mechanisms on tabular data. The Autoencoder had an input layer and an output layer of the same shape that attempted to reconstruct the traffic. If the resultant error was beyond a predetermined threshold, then it implied that some unusual activity was occurring. The TabNet used attention mechanisms to identify anomalies from the output of the model.

(Berguiga et al., 2025) presented a hybrid deep learning intrusion detection system for IoMT devices. The methodology combined two neural network architectures; Convolutional Neural Network (CNN) and Long Short-Term Memory (LSTM) neural network trained on the IoTID20 and the Edge-IoTset datasets. The datasets contain samples categorized into different attacks, including DoS and DDoS attacks. The CNN model performed feature extraction on the data while the LSTM model was used for sequence data prediction. The hybrid model achieved a high accuracy of 99.92%.

(Zabeehullah et al., 2025) also suggested a hybrid CNN-LTSM model for DDoS detection in IoMT. They used the CNN model for spatial feature extraction, and the LTSM module was trained on time-series data for network traffic features, augmenting their pipeline with SHAP-based explainability. They used the CICDDoS2019 and IoT Healthcare Security datasets, achieving a high accuracy of 99.59%, and providing interpretable feature attributions.

(Bhutia et al., 2022) proposed a three-phase methodology using just three machine learning algorithms: the KNN classifier, the Logistic Regression classifier, and SVM (Linear and RBF), for DDoS detection at the source, which is often bots. Their dataset of choice was traffic data collected from various commercial IoT devices that were exploited in the MIRAI and BASHLITE DDoS attacks. They extracted 23 features for classification of traffic as either benign or junk, and their proposed models achieve a detection accuracy of 99%.

(Phan et al., 2024), feature selection techniques of IoT attacks (Information Gain, Recursive Feature Elimination, Random Forest, XGBoost, and Logistic Regression) for machine learning detection on the CIC IoT 2023 dataset. Their results established that Random Forest and Information Gain consistently achieved high and stable accuracy, supporting our decision to use the two methods for our feature selection.

Prior works have advanced feature engineering and machine learning methods for IoMT DDoS detection. While hybrid models combine spatial and temporal learning, our work isolates and investigates the contribution of temporal learning by comparing a probabilistic classifier with a temporal sequence model using the same selected feature set.

3. Design And Methodology

3.1 Dataset

We use the CiCIoMT 2024 dataset (Dadkhah et al., 2024), generated by the Canadian Institute of Cybersecurity by executing different cyberattacks against an IoMT testbed of 40 devices across multiple protocols. Overall, the dataset consists of attacks categorized into different attacks such as DDoS, DoS, Recon, MQTT, and spoofing. Our study focuses on MQTT DDoS traffic, specifically Connect Flood and Publish Flood variants, against benign MQTT traffic to create a binary classification problem.

The CiCFlowMeter tool is a network traffic flow generator and analyser which generates bidirectional flows. A flow consists of multiple packets that are defined by the same five-tuple (source IP, destination IP, source port,

destination port, protocol). Using the tool, we extracted bidirectional network flow statistics from the dataset PCAP files, with each flow record containing 84 statistical features characterizing packet-level behaviour including duration, packet counts, byte volumes, inter-arrival times, and flag counts.

3.2 Preprocessing and Feature Selection

To ensure consistency, non-numerical identifiers were excluded from the dataset and infinite, missing and NaN values were handled through dropping. Normalization was performed by applying MinMax scaling which converts values to a range of (0,1). CiCIoMT datasets are usually known to suffer from data imbalance, with the attack data being greater than the benign data. To address the class imbalance, we applied under sampling to the majority class during training data preparation, maintaining a balanced training distribution.

With so many features, dimensionality reduction is important to reduce the model complexity and remove noisy features that degrade model performance. Previous work by (Phan et al., 2024), compared multiple feature selection techniques for machine learning detection of IoT attacks. Their results showed high and stable performance in Random Forest and Information Gain, which we employed for feature reduction to result in a smaller, informative feature set.

Information Gain works by quantifying how much information a given feature contributes to correctly classifying an instance.$IG(D, A) = H(D) - H(D|A)$ where IG(D,A) is the information gain of a feature A on a dataset D. Features with high IG provide significant discriminative information and are prioritized for retention.

Random Forest evaluates the relative importance of each feature by evaluating how much the feature contributes to reducing the Gini impurity in the decision tree. If a feature results in a higher decrease in impurity, it is considered more important since it improves the model's ability to classify correctly.

Both methods independently ranked the 84 extracted features, and features were selected based on consistent high rankings across both methods as well as domain knowledge. Information Gain ranking showed Fwd Seg Size Min, and TCP flag counts (SYN, RST, and FIN) as the most discriminative features for distinguishing the attacks. These features reflect the handshake and connection termination patterns characteristic of MQTT Connect and Publish attacks. Other connection related features such as FWD Init Win Bytes and Bwd Header Length also ranked highly since they captured early-stage resource negotiation that would noticeably differ between benign and DDoS traffic. The Random Forest rankings corroborated these findings and where feature rankings differed significantly, domain knowledge was used to evaluate the relevance. This process resulted in a set of 31 features that effectively characterize network traffic and provide strong indicators of DDoS attacks. These include flow duration, bidirectional packet and byte counts, inter-arrival time statistics (mean, standard deviation, maximum), active and idle time metrics, header length statistics, and TCP flag counts. The selected features are summarized in Table 1.

Table 1: Selected Features

Category	Features
TCP Flag Counts	SYN Flag Count, ACK Flag Count, PSH Flag Count, RST Flag Count, FIN Flag Count
Flow-level metadata	Protocol, Flow Duration, Flow Bytes/s, Flow Packets/s, Down/Up Ratio
Packet Counts	Total Fwd Packet, Total Bwd Packets
Packet size statistics	Packet Length Min, Packet Length Mean, Packet Length Max, Packet Length Std, Packet Length Variance, Average Packet Size, Fwd Segment Size Avg, Fwd Seg Size Min
Header lengths	Fwd Header Length, Bwd Header Length
Initial Window Sizes	FWD Init Win Bytes, Bwd Init Win Bytes
Subflow bytes	Subflow Fwd Bytes, Subflow Bwd Bytes
Inter-arrival times	Fwd IAT Mean, Fwd IAT Total, Bwd IAT Mean, Bwd IAT Total, Flow IAT Mean

4. Model Architecture and Training

4.1 Model 1: Naïve Bayes

The Naïve Bayes model is a probabilistic classifier that assumes each feature contributes independently to the final decision. This assumption directly contrasts with the temporal dependency modelling of the LSTM model, making it a suitable control for assessing the contribution of sequential learning.

After the preprocessing step, each flow was represented as a vector of the selected features. We then applied a Gaussian Naïve Bayes classifier, which assumes that the feature values follow a normal distribution. During training, the model estimates the mean and variance of each feature separately for the benign and DDoS classes. For each new flow instance, the model computes the likelihood of the observed feature values under each class and assigns the class with a higher probability.

By treating each flow as an independent instance, the classification decisions were based purely on the statistical characteristics of that individual flow, without considering temporal relationships between consecutive flows.

4.2 Model 2: Long Short-Term Memory (LSTM)

Each network flow in our dataset includes a timestamp column indicating when each network activity was recorded. For temporal context, the flows were sorted chronologically based on their timestamps. Consecutive flows were then grouped into non-overlapping fixed-length sequences of 36 flows, allowing the model to observe how the traffic features evolve over short time periods rather than observing the network flows in isolation.

DDoS flooding attacks exhibit characteristic patterns of inter-arrival times and packet rates that evolve over time, whereas benign MQTT traffic tends to exhibit varying and unpredictable patterns. Since these behavioural differences evolve over time, modelling temporal order provides meaningful information for network intrusion detection.

The LSTM layer consists of 64 units feeding into a GlobalAveragePooling1D layer which aggregates the temporal representations across all timesteps. This is followed by two fully connected dense layers of 32 and 16 units respectively, each with Rectified Linear Unit (ReLU) activation and a dropout rate of 0.3. Dropout and recurrent dropout are applied to reduce overfitting. The output layer is a single dense unit with sigmoid activation for binary classification. The Adam optimizer and binary cross entropy loss function were used. The model architecture is summarized in Figure 1.

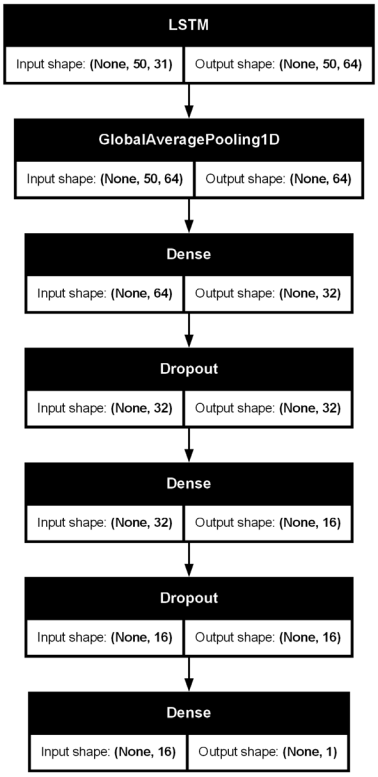

Figure 1: Model Architecture

Unlike Naïve Bayes, the LSTM retains information across timesteps, allowing it to model temporal dependencies within the flow data.

5. Experimental Results

Performance was evaluated on an unseen test set that was not used at any stage of the training process, using accuracy, precision, recall, F1-score, and the confusion matrix to assess false positive and false negative rates.

A summary comparison of the performance between the two models is shown in Table 2.

Table 2: Model Performance

Model	Accuracy	Precision		Recall		F1-Score	
		Benign	DDoS	Benign	DDoS	Benign	DDoS
Naïve Bayes	0.8990	0.99	0.86	0.69	0.99	0.82	0.92
LSTM	0.9585	0.97	0.95	0.90	0.98	0.94	0.97

Tables 3 and 4 show the normalized confusion matrices.

Table 3 Bayes Confusion Matrix

	Predicted Benign	Predicted DDoS
Actual Benign	75.3%	24.7%
Actual DDoS	0.09%	99.91%

Table 4: LSTM Confusion Matrix

	Predicted Benign	Predicted DDoS
Actual Benign	100%	0%
Actual DDoS	0.013%	99.987%

The LSTM model achieved higher overall accuracy (95.85%) compared to Naïve Bayes (89.90%). The high accuracy scores on both models demonstrate that the selected features contain sufficient discriminative information for detection. While both models demonstrated strong recall for DDoS traffic, the notable differences emerged in the benign traffic classification. The 69% score on benign traffic in Naïve Bayes indicates a high false positive rate (24.7%), where a lot of benign traffic flows were misclassified as DDoS. In contrast, LSTM improved the benign recall to 90% indicating that the temporal context helped distinguish sustained attack behaviour from legitimate short bursts of activity. In isolation, individual benign flows can resemble short attack bursts because legitimate MQTT clients frequently establish connections with brokers and publish messages in succession. This generates packet rates, TCP flag patterns and segment sizes that are similar to the early stages of a Connect or Publish attack. Without temporal context, a classifier may struggle to distinguish between a legitimate connection event and the start of a volumetric attack.

These results indicate that single flow features may not be enough to rule out false positives. The LSTM's temporal context allows it to evaluate whether a suspicious flow is part of an ongoing flood or a behavioural variation, therefore reducing false positives.

6. Conclusion and Future Works

This study examined the viability of feature selection and temporal modelling in detecting MQTT-based DDoS attacks within IoMT networks using the CICIoMT 2024 dataset. We demonstrated that guided dimensionality reduction results in concise, discriminative features that support both traditional machine learning and deep

learning classifiers. Our results show that by leveraging sequential flow dependencies, the LSTM model was able to significantly outperform the Naïve Bayes model.

Our main finding is that temporal modelling is meaningful for DDoS detection. This is because DDoS attacks usually occur as huge and repetitive volumes of normal looking traffic over short intervals of time. This is easier to distinguish from legitimate communication when modelled as sequences rather than independent instances.

Future works may include expanding beyond MQTT DDoS attacks. The CiCIoMT 2024 dataset also contains traffic data for other attack types, such as TCP SYN floods and spoofing attacks, which we did not include in our work. It would be worth exploring whether the same model trained on MQTT data can be generalised for other attacks since the network behaviour differs for each. Additionally, our current solution performs binary classification between benign and DDoS traffic. Extending this to multi-class classification to distinguish among specific attack types would provide more granular and informative detection. This would be particularly useful in critical sectors such as the healthcare industry, where accurately detecting the type of attack can inform more effective incident response efforts. Future work could also explore lightweight alternative models that could achieve comparable performance at a lower computational cost due to the resource constraints in IoMT devices. This would make deployment on constrained edge devices more feasible.

Ethics Declaration: Ethical clearance was not required for this research since the work used publicly available datasets and did not include human subjects or any sensitive information.

AI Declaration: AI tools were used only for minor formatting assistance. The intellectual contribution, methodology, and content of this paper are solely the work of the authors.

References

Berguiga, A., Harchay, A. and Massaoudi, A. (2025). HIDS-IoMT: A Deep Learning-Based Intelligent Intrusion Detection System for the Internet of Medical Things. IEEE Access, 13, pp.32863–32882. doi:https://doi.org/10.1109/access.2025.3543127.

Bhutia, N. T., Verma, H., Chauhan, N., & Awasthi, L. K. (2022). DDoS Attacks Detection in 'Internet of Medical Things' Using Machine Learning Techniques. *2022 IEEE Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI)*, 1–6. https://doi.org/10.1109/IATMSI56455.2022.10119428

Censinet.com. (2026). Ultimate Guide to DDoS in Healthcare. [online] Available at: https://censinet.com/perspectives/ultimate-guide-ddos-healthcare [Accessed 12 May 2026].

Clark, T., Rahouti, M., & Xiong, K. (2024). Machine Learning-Based Detection for Cyber Attacks in Internet of Medical Things Devices. *2024 IEEE MIT Undergraduate Research Technology Conference (URTC)*, 1–5. https://doi.org/10.1109/URTC65039.2024.10937650

Dadkhah, S., Neto, E. C. P., Ferreira, R., Molokwu, R. C., Sadeghi, S., & Ghorbani, A. A. (2024). CICIoMT2024: A benchmark dataset for multi-protocol security assessment in IoMT. *Internet of Things*, *28*, 101351. https://doi.org/10.1016/j.iot.2024.101351

ElSayed, Z., Abdelgawad, A. and Elsayed, N., 2025, April. Cybersecurity and frequent cyber attacks on IoT devices in healthcare: Issues and solutions. In 2025 IEEE 4th International Conference on Computing and Machine Intelligence (ICMI) (pp. 1-7). IEEE. doi:https://doi.org/10.1109/icmi65310.2025.11141075.

Hnamte, V., Balram, G., Priyanka, V. and Kolluru, V., 2022. Implementation of Naive Bayes classifier for reducing DDoS attacks in IoT networks. Journal of Algebraic Statistics, 13(2), pp.2749-2757.

Kirubavathi G, Sivakumar, S. and Manickam, S. (2024). Optimized Hybrid Approach for Anomaly Detection of DDoS and Network Attacks in IoMT Systems using Autoencoders and TabNet. [online] pp.1–7. doi:https://doi.org/10.1109/icmnwc63764.2024.10872358.

Phan, V. A., Jerabek, J., & Malina, L. (2024). Comparison of Multiple Feature Selection Techniques for Machine Learning-Based Detection of IoT Attacks. *Proceedings of the 19th International Conference on Availability, Reliability and Security*, 1–10. https://doi.org/10.1145/3664476.3670440

S. Lakshminarayana and P. Santhi Thilagam, "Next-Generation DDoS Attacks on IoT Deployments: Targeting the Advanced Features of MQTT v5.0 Protocol," in IEEE Internet of Things Journal, vol. 12, no. 12, pp. 22090-22109, 15 June15, 2025, doi: 10.1109/JIOT.2025.3549784.

Zabeehullah, Arif, F., Mazhar Ul Haq, Q., Ali Khan, N., Ud Din, I., Almogren, A., Ali Khan, M., Alsaleh, O., & Guizani, M. (2025). Hybrid CNN-LSTM Model for DDoS Attack Detection in Internet of Things-Based Healthcare Industry 5.0. *IEEE Internet of Things Journal*, *12*(22), 46075–46082. https://doi.org/10.1109/JIOT.2025.3569610