

NeuroDefend: Artificial Intelligence-Powered IOT Cyber Defense an Intelligent Intrusion Detection System for Protecting Smart Device Networks

Janhavi Raut, Prof/Extc, Hemangi Satam

Shree L R Tiwari College of Engineering, Mira road, Maharashtra

Abstract - The widespread connectivity of Internet of Things (IoT) devices in different industries has raised cybersecurity concerns because they are prone to exploitation. This work reports the design of a light-weight, AI-driven intrusion detection system (IDS) that is tailored for IoT networks. The system employs machine learning algorithms for predicting abnormal network traffic patterns, issuing real-time warnings through an easy-to-use web interface. Using datasets like CICIDS 2017 and BoT-IoT, we trained and tested supervised learning models to achieve high detection rates with minimal computational overhead. The system also overcomes device heterogeneity and resource constraints. Our methodology integrates artificial intelligence with popular web technology to produce an adaptive, scalable, and deployable IoT security solution.

1. INTRODUCTION

The Internet of Things (IoT) has transformed industries like healthcare, farming, manufacturing, and home automation through real-time data transfer and automation. While the growth of connected devices has enhanced capabilities, it has also raised the attack surface for cyber attacks. Conventional intrusion detection systems (IDS) are not suitable for IoT because they consume high computational power and are not adaptable. This paper suggests a web-based, AI-powered IDS that can function suitably in IoT environments, providing real-time threat detection as well as easy-to-use monitoring facilities. IoT networks are heterogeneous, containing devices with varying operating systems, protocols of communication, and processing power. IoT networks do not have a centralized management system and are installed in public spaces, thus making them vulnerable to attack. Security in IoT systems is hence specially demanding. Current security solutions, including firewalls and signature-based IDS, prove inadequate because they are static and rely on predetermined rules, which cannot identify unknown or changing threats. To overcome these shortcomings, there is increased interest in integrating IDS with artificial

intelligence and machine learning. These technologies bring the power to learn based on past experiences and recognize patterns of both normal and abnormal behavior. As opposed to conventional solutions, AI-driven systems can evolve over a period of time, and this helps in detecting zero-day attacks and reducing false positives. Further, the addition of a web-based interface improves usability by presenting real-time visual feedback of network activity and alerts. The suggested system utilizes publicly accessible datasets to train an intelligent machine learning model with high precision in detecting network intrusions. The fact that it is built on a Flask-based web dashboard enables simple and effortless interaction, which makes it easily accessible to users without significant technical background. The paper provides a detailed account of the system design, implementation process, performance assessment, and future research with a view to enhancing IoT security via intelligent and scalable solutions

2. LITERATURE SURVEY

Literature on intrusion detection systems (IDS) in IoT environments has dramatically changed with the development of artificial intelligence (AI) and machine learning (ML). Conventional IDS, dependent on rule-based or signature-based detection, are more and more considered insufficient because they cannot detect new attacks and depend on pre-defined rules. With increasingly dynamic and heterogeneous IoT environments, there is a growing requirement for adaptive, intelligent detection techniques that has prompted researchers to investigate ML- and AI-based solutions. Early efforts include Denning (1987) and Axelsson (2000) establishing groundwork models for IDS in traditional computing systems. The models, though efficient for wired networks, prove inefficient in the low resource, highly connected environment of IoT. Later studies have shifted toward the application of

supervised machine learning algorithms like Decision Trees, Naive Bayes, Support Vector Machines (SVM), and Random Forest. For instance, Meidan et al. (2018) showed the effectiveness of Random Forest in device behavior classification on an IoT network with high detection rates and few false positives. Likewise, Doshi et al. (2018) employed supervised learning to detect anomalies in IoT devices and attained competitive accuracy with small feature sets. Unsupervised learning methods have also been investigated for anomaly detection. Autoencoders, K-Means clustering algorithms, and isolation forests are widely used because they can identify zero-day attacks with no labeled datasets. For example, Apruzzese et al. (2019) demonstrated how unsupervised ML could identify nascent attack patterns in dynamically changing network traffic. Such methods, however, are prone to high rates of false alarms and are usually challenging to fine-tune to various IoT environments. Deep learning has indicated potential to enhance detection capacity. Recurrent Neural Networks (RNN) and Long Short-Term Memory (LSTM) models are applied to process time-series data collected from IoT traffic. They are capable of efficiently identifying sequences that indicate anomalous behavior. An example is a study by Diro and Chilamkurti (2018), which applied a deep learning system with LSTM networks trained from the NSL-KDD dataset and showed enhanced detection of Denial-of-Service (DoS) and probing attacks. Regarding datasets, most research studies employ publicly available datasets such as NSL-KDD, CICIDS 2017, BoT-IoT, and IoT-23. Although these datasets are good sources of training and benchmarking, most of them are obsolete or contain limited diversity of real-world IoT traffic. This has been identified in various studies, including those by Ferrag et al. (2020), who highlight the imperative of using real and fresh IoT specific datasets. Federated learning is also a new direction, where training is done jointly across devices without exchanging raw data, thus maintaining privacy and with decreased communication overhead. This is especially suited to IoT, where bandwidth is scarce and privacy of data is of the essence. Notwithstanding great advances, there are issues with real-time detection, computational cost, and explainability. The research by Shone et al. (2018) has emphasized the need for finding a balance between detection performance and computational cost. Explainable AI (XAI) is increasing in popularity to enable decision-makers to know why a particular packet or activity was identified as malicious. Overall, the literature is a dynamic and mature field with the focus on implementing lightweight, scalable, and precise IDS specific to the limitations of IoT settings. This work draws on such findings and suggests an AI-based IDS with real-time visualization and seeks to implement it in practical use on heterogeneous IoT networks.

3. PROPOSED SYSTEM DESIGN

Our system consists of the following:

- IoT Devices: Simulated or actual devices creating network traffic.
- Data Collection Module: Sniffs traffic using applications such as Wireshark or built-in sniffers.
- AI Engine: Supervised machine learning algorithm trained to identify anomalies.
- Web Interface: Built with Flask and JavaScript to display logs and alerts in real-time.
- Alert Mechanism: Alerts users of identified intrusions through UI notifications and optional email.

4. METHODOLOGY

The process includes data collection and preprocessing of IoT traffic, feature selection, machine learning model training and testing, and a responsive web interface implementation. The CICIDS 2017 dataset was chosen for its inclusive attack coverage. Scaling and outlier removal were performed during preprocessing. 80% of the data was used to train a Random Forest classifier with cross-validation. Flask was utilized to develop a real-time dashboard showing alerts and device statuses.

5. RESULTS AND ANALYSIS

The Random Forest algorithm had an accuracy rate of 91%, precision of 89%, recall of 90%, and an F1 score of 89.5%. Real-time latency for alerts was under 500ms. The system effectively identified various attacks such as DoS, port scanning, and data exfiltration. A confusion matrix was employed to cross-validate performance.

6. CHALLENGES

Some of the major challenges are high false positives, resource constraints on IoT devices, and handling high-frequency data streams. These were alleviated by lightweight model design, preprocessing optimizations, and modular system design.

7. CONCLUSION AND FUTURE WORK

This work shows the viability of deploying an AI-driven IDS in IoT settings. Future work will involve federated learning, edge computing, mobile alert apps, and adding explainable AI (XAI) for increased trust and flexibility.

8. DATASET AND FEATURE ENGINEERING FOR IOT INTRUSION DETECTION

The backbone of a successful AI-driven intrusion detection system is good-quality datasets and well designed features. Datasets form the material for training machine learning as well as deep learning

models, whereas feature engineering converts this raw material into useful input that enables such models to differentiate effectively between malicious and benign activity. The most widely used public datasets in IoT intrusion detection are CICIDS 2017, BoT-IoT, UNSW-NB15, and IoT-23. The CICIDS 2017 dataset is based on emulating real-world network traffic and contains a mix of different types of attacks like DoS, brute force, infiltration, and port scanning. BoT-IoT targets botnet based attacks in IoT environments and provides raw as well as preprocessed network flow data. UNSW NB15 includes contemporary normal and attack behavior in a controlled setting, and IoT-23 records traffic from genuine IoT devices such as home appliances and IP cameras, providing realistic behavioral data. These data sets are imperative as they hold labeled instances that distinguish between attack and normal traffic, which is of paramount importance for supervised learning. Impediments exist in the form of class imbalance (where normal traffic greatly exceeds malicious occurrences), noise, and missing values. To overcome these impediments, preprocessing methods like normalization, cleaning, and data augmentation are employed. Normalization provides consistent scaling across features, while data augmentation adds synthetic attack traffic to balance the data set. Feature engineering is also a key step. It involves feature selection and transformation from raw data attributes into the most useful form for the model. Typical features employed include protocol type, packet length, TCP flags, time-to-live (TTL), source and destination ports and IPs, and byte counts. Dimensionality reduction methods such as PCA (Principal Component Analysis) or feature selection methodologies such as Recursive Feature Elimination (RFE) are typically utilized to minimize computation cost and enhance generalization. Domain expertise is important in identifying useful features. For instance, in home automation, an abnormal traffic surge from a low-bandwidth sensor could be a sign of abnormal activity, whereas unusual communication between devices that are not normally communicating could be a sign of a lateral movement attack. Effective feature extraction also enables model explainability and reduces processing time—crucial in IoT settings where computational power is limited. In addition, with federated or distributed learning environments, feature extraction needs to be consistent across devices to guarantee harmonious collaborative training of the model. In summary, both the dataset selection and feature quality have a direct influence on the performance and dependability of AI-powered IDS in IoT networks. In the absence of realistic, well-organized data, even the most advanced models can fall short. Consequently, efforts are presently directed towards developing more varied IoT-specific datasets and improving feature extraction pipelines to facilitate robust and real-time threat detection.

9. PERFORMANCE EVALUATION AND COMPARISON OF AI-POWERED IDS

Performance analysis of AI-driven intrusion detection systems is important to determine how well they can identify threats, particularly in IoT scenarios where false positives and false negatives are important consequences. Common performance measures that are applied in assessing IDS models are precision, recall, F1-score, accuracy, and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). Accuracy refers to the ratio of total positive predictions to total predictions (positive and negative). Although this measure gives an approximate sense of performance, it can be deceptive for imbalanced sets—those with a lot of benign traffic, typical in intrusion detection—since benign traffic will overwhelm the others. Precision (actual positives / total predicted positives) and recall (actual positives / true positives) are more telling in these situations. Precision indicates the number of identified threats which were indeed threats, while recall states the number of actual threats that were detected. The F1-score averages these two into a harmonic mean, providing a well-balanced measurement. AUC-ROC measures the performance of the model for various thresholds and implies how well the model classifies benign from malicious instances. Random Forest classifiers in current research have provided high accuracy (greater than 90%) and low false positive values when applied to data like CICIDS 2017. Deep learning models, including CNNs and RNNs, if properly tuned, tend to outperform conventional ML algorithms in the discovery of intricate patterns but at the expense of longer training time and computational requirements. Hybrid models that combine the strengths of multiple approaches—like CNN-RNN hybrids or Random Forests integrated with unsupervised anomaly detectors—are also becoming popular due to their balanced performance across metrics. A critical consideration in IDS evaluation in IoT environments is how resource-intensive the system is. A model might be good in the lab but lose its charm when deployed in real-world applications on edge devices with low memory or low power. Hence, model evaluation also entails inference time (how fast a decision is taken), memory usage, and energy draw. Another essential evaluation criterion is adversarial robustness, wherein adversaries explicitly make efforts to manipulate input data to circumvent detection. Models trained on diverse and recent datasets are typically more resistant. Real-time performance testing, where simulated attacks are carried out on testbeds or virtual networks, is

frequently used to measure detection latency and accuracy of alerts. Comparison between studies is challenging owing to differences in datasets, attack scenarios, and test setups. Consequently, standardized testbeds and test protocols are being promoted within the research community. Interpretability and explainability are also becoming evaluation criteria, particularly in fields such as healthcare or finance, where it is important to know why a detection occurred. In conclusion, thorough performance evaluation does not only demand statistical measures but also actual world testing and resource consumption analysis. Good IDS should balance accuracy, speed, scalability, and transparency to be practical in real-world IoT deployments.

Comparison Example:

| Model | Accuracy | F1-Score | Notes |
|--------------------|----------|----------|---|
| SVM | 88% | 82% | Poor performance with large IoT datasets |
| Random Forest | 91% | 89.5% | Good balance between speed and accuracy |
| Autoencoder | 82% | 81% | Lightweight but high false alarm rate |
| CNN + RNN (Hybrid) | 93% | 90% | High accuracy but computationally intensive |

Metrics to Evaluate Your Model:

| Metric | Meaning |
|-----------|--|
| Accuracy | Percentage of all correct predictions |
| Precision | Percentage of positive predictions that were correct |
| Recall | Percentage of actual intrusions that were correctly detected |
| F1-score | Harmonic mean of precision and recall |
| AUC-ROC | Model's ability to distinguish between classes |

REFERENCES

- [1] V. T. Nguyen and R. Beuran, "FedMSE: Federated learning for IoT network intrusion detection," Oct. 2024. digital-library.theiet.org
- [2] G. Shen, W. Yang, Z. Chu, J. Fan, D. Niyato, and K.-Y. Lam, "Effective Intrusion Detection in Heterogeneous Internet-of-Things Networks via Ensemble Knowledge Distillation-based Federated Learning," Jan. 2024.
- [3] O. Belarbi, T. Spyridopoulos, E. Anthi, I. Mavromatis, P. Carnelli, and A. Khan, "Federated Deep Learning for Intrusion Detection in IoT Networks," Jun. 2023.
- [4] A. Javeed, M. S. Saeed, M. T. Adil, and P. Kumar, "A federated learning-based zero trust intrusion detection system for Internet of Things," *Ad Hoc Networks*, vol. 162(6), p. 103540, May 2024.
- [5] J. Diro and N. Chilamkurti, "Distributed attack detection scheme using deep learning approach for Internet of Things," *Future Generation Computer Systems*, vol. 82, pp. 761–768, 2023.
- [6] S. Noor et al., "Federated intelligence for intrusion detection and unifying health prediction in IoHT," *IET Conference Proceedings*, Jan. 2024.
- [7] P. Selvam, P. Karthikeyan, S. Manochitra et al., "Federated learning-based hybrid convolutional recurrent neural network for multi-class intrusion detection in IoT networks," *Discover Internet Things*, vol. 5, p. 39, Apr. 2025. (Accepted Mar. 2025)
- [8] V. Lazzarini, H. Tianfield, and V. Charissis, "Federated Learning for IoT Intrusion Detection," *AI*, vol. 4, no. 3, pp. 509–530, Jul. 2023.
- [9] T. Shan et al., "A federated learning approach to network intrusion detection using residual networks in industrial IoT networks," *Journal of Supercomputing*, vol. 80, pp. 18325–18346, May 2024.
- [10] M. Laddi, P. Sonwalkar, S. Allagi, S. Athanikar, B. Nadagoudra, and N. Kishore, "Advanced Cybersecurity Framework for Intrusion Detection Utilizing Federated Machine Learning," *Journal of Information Systems Engineering & Management*, vol. 10, pp. 37s, 2025. (Published Jan. 2025).