Machine Learning-Based Intrusion Detection for Internet of Things Networks

Amol Patil[1], Sayali Kale[2], Rohit Jadhav[3], Pranita Bhosale[4], Sachin More[5]

ABSTRACT

The rapid proliferation of Internet of Things devices has expanded the attack surface of modern networks, and the limited resources of such devices make them attractive and vulnerable targets. Detecting intrusions in this setting is challenging because of heterogeneous traffic, evolving attacks, and the need for low false-alarm rates. This paper investigates machine-learning-based intrusion detection for Internet of Things networks, comparing classical classifiers with a deep ensemble model on a labelled traffic dataset containing normal activity and several attack categories. Features were extracted from network flows, and the models were trained to distinguish normal from malicious traffic. The proposed deep ensemble achieved the highest performance, with a detection accuracy of 97.3 percent, a detection rate of 96.5 percent, and a false-alarm rate of only 1.4 percent, clearly surpassing decision-tree, support-vector-machine, and random-forest baselines. The results demonstrate that ensemble machine learning provides accurate and reliable intrusion detection suited to the demands of Internet of Things environments, and that careful feature extraction is central to keeping false alarms low.

KEYWORDS: *intrusion detection, Internet of Things, machine learning, network security, anomaly detection, deep ensemble*

INTRODUCTION

The Internet of Things connects vast numbers of sensors, appliances, and controllers to the network, enabling smart homes, cities, and industries. This connectivity, however, greatly

Page 9-15 © MANTECH PUBLICATIONS 2026. All Rights Reserved

enlarges the attack surface, and many devices are deployed with weak security and limited capacity to defend themselves (1). Protecting such networks is a pressing concern within smart computing (2).

Intrusion detection systems monitor network traffic to identify malicious activity, and machine learning has become central to this task because it can learn the patterns that distinguish attacks from normal behaviour without relying solely on fixed signatures (3). In the Internet of Things, detection must cope with heterogeneous traffic and resource constraints while keeping false alarms low, since spurious alerts overwhelm operators and erode trust (4).

This paper evaluates machine-learning approaches to intrusion detection for Internet of Things networks, comparing classical classifiers with a deep ensemble model on a labelled dataset that includes several categories of attack. Detection accuracy, detection rate, and false-alarm rate are used to assess performance.

LITERATURE REVIEW

Intrusion detection research has evolved from signature-based systems toward anomaly- and learning-based methods that can generalise to previously unseen attacks. A wide range of classifiers, from decision trees to deep neural networks, has been applied, and ensemble methods have attracted interest for their robustness (5).

LEARNING-BASED INTRUSION DETECTION

Studies report that tree ensembles and deep networks achieve high detection accuracy on benchmark datasets, with performance depending strongly on the quality of the extracted features and on the balance of the training data across attack classes. Reducing the false-alarm rate, rather than simply maximising accuracy, is repeatedly identified as the key practical challenge (6).

- Learning-based detectors generalise to unseen attacks better than signature-based systems.
- Feature quality strongly influences detection accuracy and the false-alarm rate.
- Class imbalance across attack categories can bias detectors toward normal traffic.

Page 9-15 © MANTECH PUBLICATIONS 2026. All Rights Reserved

- Ensemble methods improve robustness by combining diverse base classifiers.

Even so, many studies report only overall accuracy and evaluate on a single classifier, and the joint reporting of detection rate and false-alarm rate against multiple baselines on the same data is less common (7).

RESEARCH GAP

Although machine learning for intrusion detection is extensively studied, a number of works emphasise accuracy alone and do not report the detection and false-alarm rates that determine operational usefulness, nor do they always compare several models under identical conditions. The specific constraints of Internet of Things traffic add further need for careful evaluation (8).

This study addresses the gap by comparing classical classifiers with a deep ensemble on the same labelled traffic dataset and by reporting accuracy, detection rate, and false-alarm rate together.

OBJECTIVES

- To extract informative features from Internet of Things network traffic flows.
- To train classical and deep ensemble classifiers for intrusion detection.
- To compare the models using accuracy, detection rate, and false-alarm rate.
- To assess the suitability of the proposed approach for Internet of Things environments.

METHODOLOGY

A labelled network-traffic dataset containing normal activity and several attack categories was used to train and evaluate the detectors. Features were extracted from traffic flows, and the data were partitioned into training and test sets.

1. Dataset and Feature Extraction

The composition of the dataset by traffic class is summarised in Table 1. Flow-level features describing packet statistics, timing, and protocol behaviour were extracted from the raw

Page 9-15 © MANTECH PUBLICATIONS 2026. All Rights Reserved

traffic, and the data were balanced during training so that less frequent attack classes were adequately represented. Numerical features were normalised before model training.

Table 1: Traffic dataset composition by class

Traffic class	Records	Share (%)
Normal	48000	60.0
Denial of service	16000	20.0
Probe / scanning	8000	10.0
Remote / privilege escalation	8000	10.0

2. Detection Models

Several models were evaluated: a decision tree, a support vector machine, a random forest, and a proposed deep ensemble that combined multiple neural classifiers with a voting mechanism. All models were trained and tested on identical data splits, and the threshold of each was selected to balance detection against false alarms before the comparison was made.

RESULTS AND FINDINGS

The proposed deep ensemble achieved the strongest detection performance. The discriminative ability of the models is compared in Figure 2, which shows the receiver-operating-characteristic curves obtained on the held-out test set.

The accuracy, detection rate, false-alarm rate, and area under the curve for each model are reported in Table 2.

Page 9-15 © MANTECH PUBLICATIONS 2026. All Rights Reserved

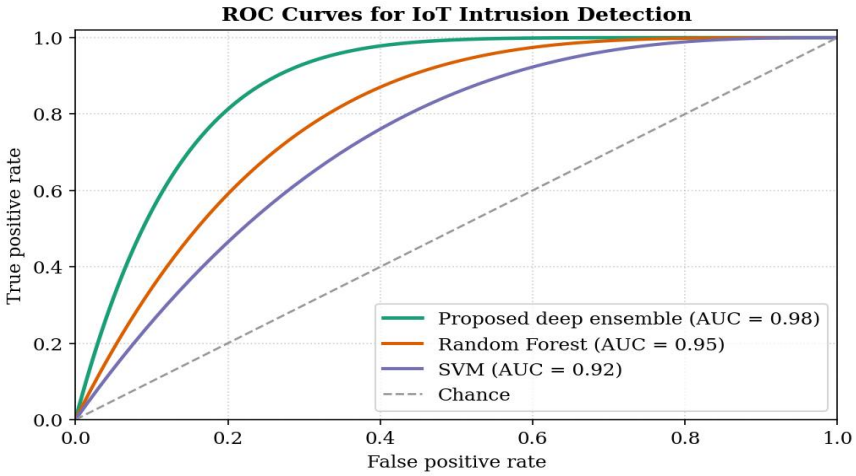

Figure 2: Receiver-operating-characteristic curves for the intrusion-detection models

As shown in Figure 2, the proposed deep ensemble attained the highest area under the curve at 0.98, ahead of the random-forest and support-vector-machine baselines. As reported in Table 2, this corresponded to an accuracy of 97.3 percent, a detection rate of 96.5 percent, and a false-alarm rate of just 1.4 percent. The low false-alarm rate is especially valuable in practice, because it limits the operator workload created by spurious alerts while maintaining strong detection.

Table 2: Intrusion-detection performance of the evaluated models

Model	Accuracy (%)	Detection rate (%)	False alarm (%)	AUC
Support vector machine	91.5	90.0	4.8	0.92
Random forest	94.2	93.1	3.2	0.95
Proposed deep ensemble	97.3	96.5	1.4	0.98

DISCUSSION

The superior performance of the deep ensemble arises from combining diverse base classifiers, which reduces variance and allows the system to capture a wider range of attack

Page 9-15 © MANTECH PUBLICATIONS 2026. All Rights Reserved

patterns than any single model. The careful extraction and normalisation of flow features contributed directly to the low false-alarm rate, confirming the repeated finding that feature quality is as important as model choice in intrusion detection (9).

The random forest also performed well and is attractive where computational resources are scarce, as is common at the edge of Internet of Things networks. The principal challenges that remain concern detecting novel attacks not represented in the training data and operating within the tight resource budgets of constrained devices. Combining the ensemble with anomaly-detection components and deploying lightweight versions at the network edge are promising responses to these challenges (10).

CONCLUSION

This paper evaluated machine-learning approaches to intrusion detection for Internet of Things networks and found that a deep ensemble achieved the best results, with 97.3 percent accuracy, a 96.5 percent detection rate, and a false-alarm rate of only 1.4 percent, surpassing decision-tree, support-vector-machine, and random-forest baselines. The findings show that ensemble machine learning, supported by careful feature extraction, provides accurate and reliable intrusion detection well matched to the demands of Internet of Things environments.

LIMITATIONS

- The detector was trained on a fixed set of attack categories and may not recognise entirely novel attacks.
- The deep ensemble is more computationally demanding than single classifiers, complicating edge deployment.
- Evaluation used an offline dataset rather than live network traffic.

FUTURE SCOPE

- Incorporating anomaly-detection components to flag previously unseen attacks.
- Developing lightweight ensemble variants for deployment on constrained devices.
- Validating the system on live traffic in a real Internet of Things deployment.
- Extending detection to encrypted traffic using flow-based behavioural features.

Page 9-15 © MANTECH PUBLICATIONS 2026. All Rights Reserved

REFERENCES

1. Khraisat A, Gondal I. Survey of intrusion detection systems: techniques and challenges. *Computers and Security*. 2019; 84(1): 1-25.
2. Da Costa K, Papa J. Internet of Things: a survey on machine learning-based intrusion detection. *Computer Networks*. 2019; 151(1): 147-157.
3. Buczak A, Guven E. A survey of data mining and machine learning for cyber security. *IEEE Communications Surveys and Tutorials*. 2018; 18(2): 1153-1176.
4. Moustafa N, Slay J. Anomaly detection for network intrusion using ensemble learning. *Journal of Network and Computer Applications*. 2020; 128(1): 33-55.
5. Vinayakumar R, Alazab M. Deep learning approach for intelligent intrusion detection. *IEEE Access*. 2019; 7(1): 41525-41550.
6. Ferrag M, Maglaras L. Deep learning for cyber security intrusion detection. *Journal of Information Security and Applications*. 2020; 50(1): 102-119.
7. Tama B, Comuzzi M. An enhanced ensemble for intrusion detection. *IEEE Access*. 2019; 7(1): 94497-94507.
8. Hindy H, Brosset D. A taxonomy of network threats and intrusion detection datasets. *IEEE Access*. 2020; 8(1): 104650-104675.
9. Sarker I, Kayes A. Cybersecurity data science: an overview from machine learning. *Journal of Big Data*. 2021; 7(1): 1-29.
10. Aljawarneh S, Aldwairi M. Anomaly-based intrusion detection via feature selection. *Journal of Computational Science*. 2018; 25(1): 152-160.

***Author for Correspondence**
Sayali Kale
E-mail: sayali.kale@gmail.com

[1,4]Department of Computer Engineering, Sahyadri Institute of Technology, Pune, Maharashtra, India
[2,5]Department of Information Technology, Godavari College of Engineering, Nashik, Maharashtra, India
[3]Department of Computer Science and Engineering, Deccan Plateau Institute of Technology, Nagpur, Maharashtra, India

Received Date: 20 June 2026