

A stacked ensemble approach to identify internet of things network attacks through traffic analysis

Adnan Rawashdeh¹, Mouhammd Alkasassbeh², Mohammad Alauthman³, Mohammad Almseidin⁴

¹Department of Information Technology, Faculty of Information Technology and Computer Sciences, Yarmouk University, Irbid, Jordan

²Department of Computer Science, Princess Sumaya University for Technology, Amman, Jordan

³Department of Information Security, Faculty of Information Technology, University of Petra, Amman, Jordan

⁴Department of Computer Science, Tafila Technical University, Tafila, Jordan

Article Info

Article history:

Received Nov 2, 2023

Revised Apr 24, 2024

Accepted May 17, 2024

Keywords:

Anomaly detection

Cybersecurity

Ensemble learning

Internet of things

Machine learning

ABSTRACT

The internet of things (IoT) has increased exponentially in connected devices worldwide in recent years. However, this rapid growth also introduces significant security challenges since many IoT devices have vulnerabilities that can be exploited for cyber-attacks. Anomaly detection using machine learning algorithms shows promise for identifying abnormal network traffic indicative of IoT attacks. This paper proposes an ensemble learning framework for anomaly detection in IoT networks. A systematic literature review analyzes recent research applying machine learning for IoT security. Subsequently, a novel stacked ensemble model is presented, combining multiple base classifiers (random forest, neural network, support vector machine (SVM)) and meta-classifiers (gradient boosting) for improved performance. The model is evaluated on the IoTID20 dataset, using network traffic features to detect anomalies across binary, multi-class, and multi-label classifications. Experimental results demonstrate that the ensemble model achieved 99.7% accuracy and F1 score for binary classification, 99.5% accuracy for multi-class, and 91.2% accuracy for multi-label classification, outperforming previous methods. The model provides an effective anomaly detection approach to identify malicious activities and mitigate IoT security threats.

This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.



Corresponding Author:

Adnan Rawashdeh

Department of Information Technology, Faculty of Information Technology and Computer Sciences

Yarmouk University

Irbid 21163, Jordan

Email: adnan.r@yu.edu.jo

1. INTRODUCTION

The internet of things (IoT) refers to the vast network of interconnected, smart devices transforming numerous application domains [1], [2]. IoT adoption continues accelerating, with over 25 billion devices forecasted by 2025 [3]. However, serious security and privacy concerns stem from such a vast number of IoT devices and networks [4], [5], emanating largely from the resource-constrained nature of devices and manufacturers' lack of security provisions [6], [7]. Attacks on IoT infrastructure can lead to data leaks, service disruptions, and even physical safety risks [8]-[10]. Through extensive monitoring of Mirai from 2016-2021, Affinito *et al.* [11] unravel its adaptive survival across three major variants, 70 million infection attempts, and a distribution shift towards developing IoT markets. Despite fluctuating infection rates, Mirai's persistent threat highlights the strategic resilience of IoT botnets to exploit new vulnerabilities and sustain attacks over the years through continual evolution.

A promising approach to IoT security is developing anomaly detection systems that analyze network traffic and identify abnormal behaviors indicative of cyber-attacks [12]. Machine learning techniques enable the detection of anomalies by learning patterns of normal vs. anomalous traffic [13]. However, IoT environments pose challenges such as immense data scale, traffic heterogeneity across diverse devices, and constantly evolving attack tactics [14]. Ensemble learning methods, which combine multiple models, have improved anomaly detection accuracy and robustness compared to single techniques [15], [16].

This paper presents an ensemble learning framework for anomaly detection aimed at securing IoT networks. The main contributions are:

- Systematic literature review of machine learning techniques applied for IoT anomaly detection.
- Novel stacked ensemble model combining diverse base classifiers and meta-classifiers.
- Evaluation of the IoTID20 dataset [17] for binary, multi-class, and multi-label classification.
- Analysis of various sampling and feature selection methods.
- Demonstration of state-of-the-art performance in detecting IoT anomalies.

The paper is organized as follows: section 2 reviews related work on IoT security and machine learning approaches. Section 3 details the method used, including dataset, feature selection, sampling methods, base/meta learners, and evaluation metrics. Section 4 presents the experimental results. Section 5 presents the discussion, while the conclusion is laid out in section 6.

2. RELATED WORK

This section reviews research efforts that utilize machine learning to develop anomaly detection systems for IoT security.

2.1. Internet of things security landscape

IoT environments comprise a wide spectrum of consumer, enterprise, and industrial devices interconnected via wired and wireless networks [1], [15]. Diverse IoT application domains include smart homes, healthcare, transportation, utilities, manufacturing, among others [18]. The scale and heterogeneity of IoT ecosystems pose significant cybersecurity challenges [6], which include the following:

- Resource-constrained devices lack security protections.
- Vulnerabilities in protocols and firmware.
- Large and diverse attack surfaces.
- User privacy risks from data collection.
- Safety critical risks if devices malfunction.

Common IoT attacks include distributed denial of service (DDoS), malware infections, man-in-the-middle (MITM), password cracking, and data exfiltration [19], [20]. Attackers can exploit IoT devices to gain access to wider networks and systems. The Mirai botnet exemplified the mass scale of insecure IoT devices leveraged for DDoS attacks [21]. Table 1 summarizes key security objectives for IoT environments [22]. A holistic IoT security strategy requires measures to be applied across people, processes, and technology [23]. Anomaly detection is critical in identifying IoT attacks in real time by analyzing network data.

Table 1. IoT security objectives

| Security objective | Description |
|--------------------|---|
| Confidentiality | Preventing unauthorized access to sensitive data |
| Integrity | Safeguarding accuracy and completeness of data |
| Availability | Ensuring accessibility and reliability of services |
| Authentication | Verifying identities and access permissions of users/devices |
| Authorization | Enforcing appropriate access policies and restrictions |
| Accounting | Keeping track of what users access, the duration, and changes they make |

2.2. Anomaly detection for IoT security

Anomaly detection refers to identifying patterns in data that deviate from expected normal behavior [13]. It is widely adopted in diverse applications such as fraud detection, healthcare monitoring, network security, and numerous others. For IoT environments, anomaly detection analyzes network traffic features to detect potential cyber-attacks [12], [24]. It is a core technique for developing intrusion detection systems (IDS) tailored to IoT [25], [26].

Anomaly detection relies on machine learning algorithms that learn patterns from data. Models are trained on benign instances then used to detect anomalies at test time. Supervised techniques require labeled

data of both normal and anomalous instances. Unsupervised methods rely solely on modeling normal instances. Semi-supervised techniques leverage a small anomaly dataset. Popular techniques include neural networks, support vector machines (SVM), isolation forests, and one-class SVM [27].

Recent research proposes numerous anomaly detection approaches for securing IoT networks, leveraging the proliferation of network traffic datasets. Moustafa *et al.* [28] used statistical metrics to evaluate univariate and multivariate outlier detection methods for IoT attack recognition. Results showed 95% accuracy in classifying anomalies. Abuali *et al.* [29] developed a system combining one-class SVM with CNN feature learning, achieving over 99% recall and precision. A model integrating autoencoder neural networks with SVM is presented in [30], also showing high performance on IoT intrusion datasets.

Ensemble learning is an effective way to combine multiple anomaly detection models to improve overall performance. For example, Tang *et al.* [31] propose an IoT IDS using stacked generalization with KNN, decision tree, and Naive Bayes base classifiers. Feature selection and under-sampling were utilized to account for imbalanced data. The ensemble model provided strong capabilities in identifying attacks. Similarly, Yuancheng *et al.* [32] develops a majority voting ensemble of autoencoders for anomaly detection in IoT, outperforming conventional methods.

While showing promise, existing research has certain limitations. Many studies use network datasets that were artificially generated rather than captured from real IoT environments [28], [30], [32]. Most efforts focus solely on binary classification of normal vs. anomaly [31], [33], rather than the multi-class nature of IoT attacks. There remains a need for ensemble techniques tailored to IoT datasets that provide precise attack classification.

2.3. IoTID20 dataset

The IoTID20 dataset [17] contains network traffic captured from a real IoT testbed, providing a representative benchmark for security research. The testbed mimics a smart home environment with common devices connected via WiFi: security camera, smart speaker, tablets, and laptops. Normal activities and attack scenarios were executed, including DDoS, MITM, and network scans.

- IoTID20 contains full packet capture (PCAP) files processed into over 86 computer traffic features per flow. It encompasses six weeks of data with 568,514 malicious and 40,697 normal flows. Attacks are labeled across binary, multi-class, and multi-label types:
- Binary: normal vs anomaly.
- Multi-class: normal, DDoS, MITM, Mirai malware, and network scan.
- Multi-label: normal, DDoS SYN flood, ARP spoofing, Mirai brute force, Mirai HTTP flood, Mirai UDP flood, Mirai ACK flood, network scan host port, and network scan OS fingerprinting.

IoTID20 enables robust evaluation of anomaly detection systems with real IoT data and precise attack classifications. It addresses limitations of artificially constructed datasets. This research adopts the dataset to assess the proposed ensemble learning framework.

3. METHOD

This section details the ensemble learning methodology for anomaly detection in IoT network traffic.

3.1. System overview

The ensemble framework uses machine learning to analyze large-scale, heterogeneous IoT traffic data efficiently. The base classifiers employ stochastic and parallelized learning algorithms suited for high-volume data streams. As prior BIG IoT research demonstrated, ensembles built on random forests, neural networks, and SVM have shown effective scalability across millions of network flows [28], [34]. The overall anomaly detection process involves:

- Preprocessing IoTID20 dataset.
- Applying feature selection.
- Creating balanced training/test splits.
- Building an ensemble model with base classifiers and meta classifiers.
- Generating anomaly scores and attack predictions.
- Evaluating performance on test data.

Figure 1 illustrates the ensemble learning framework. First, the raw network traffic data undergoes preprocessing, including encoding categorical variables and handling missing values. Principal component analysis (PCA) is applied for feature selection to derive a lower dimensionality feature subset. As IoTID20

has an imbalanced class distribution, the training dataset is balanced using the synthetic minority oversampling technique (SMOTE).

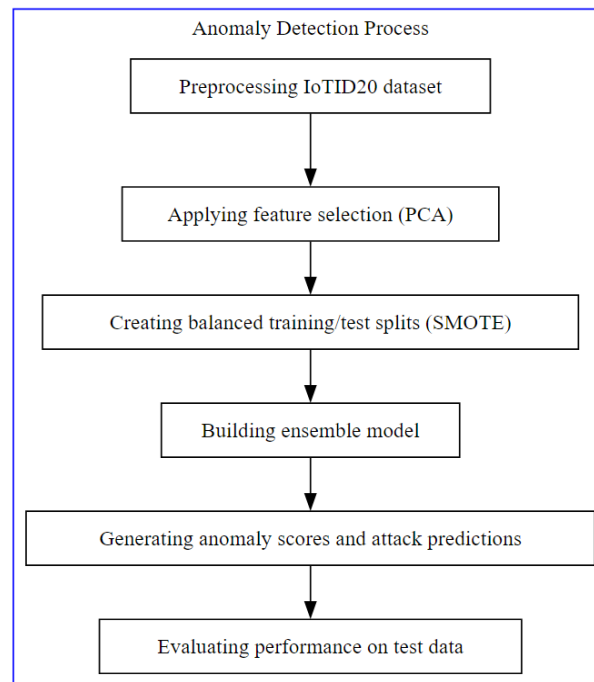


Figure 1. Ensemble learning framework for IoT anomaly detection

Four diverse machine learning algorithms are base classifiers: random forest, neural network, SVM, and Naive Bayes. These generate anomaly scores for each input instance. A gradient-boosting meta-classifier combines the outputs from the base classifiers into an aggregated anomaly score. Classification performance is evaluated on the test dataset across binary, multi-class, and multi-label metrics.

The ensemble model uses a stacked generalization approach to combine multiple base classifiers into a multilayer model. This architecture provides proven benefits like reducing bias, variance, and improving predictive performance by leveraging diverse sets of learners [15], [16]. The basis for the model configuration stems from prior research showing random forest, neural networks, SVM, and Naive Bayes as leading algorithms for IoT data characteristics [34], [35]. Hyperparameter tuning through grid search determines their optimal configuration tailored to the IoTID20 dataset features.

3.2. IoTID20 dataset preprocessing

IoTID20 provides raw PCAP files and extracts comma separated values (CSV) files. The CSV contains 86 features plus normal/attack labels. Initial preprocessing steps include:

- Converting categorical variables to numeric encoding.
- Imputing missing values using mean substitution.
- Normalizing features to 0-1 scale.
- Splitting into 80% training, 20% test datasets.

Subsequently, feature selection and sampling are applied to the training data.

3.3. Feature selection

Feature selection is an important preprocessing step to refine the input variables for efficient and robust model learning, especially with high-dimensional IoT traffic data. PCA provides an effective dimensionality reduction technique that has shown success in network analytics research [36]. PCA transforms the input feature space into fewer principal components that maximize the variance captured from the original raw features. By applying PCA on the 31,976 IoTID20 flows in the training partition, the first 31 principal components, which encompass 99% of the cumulative information content, are retained. This filtered subset of features supplies the ensemble method with information-rich inputs containing minimal redundancy that facilitate more accurate anomaly detection.

3.4. Training data balancing

Real-world network traffic exhibits imbalanced distributions across different classes, which poses learning challenges for anomaly detection models. The normal flows significantly outnumber the attack flows in IoTID20. Balancing the training data to mitigate algorithm bias toward majority classes can enhance model generalization capabilities. The SMOTE provides an adequate data augmentation approach, generating synthetic samples of the minority class rather than blind duplication [37]. SMOTE is applied to expand the DDoS, MITM, Mirai, and network scan attack categories in the IoTID20 training partition until the same number of flows as the normal traffic is reached. By balancing the training data rather than the full raw traffic, the ensemble approach gains computational efficiency since subsequent operational analysis only applies models to unseen test flows without sampling. SMOTE synthetization is a lightweight data augmentation technique.

3.5. Base classifiers

The selection of base classifiers considers model diversity to maximize ensemble synergy. The random forest provides non-linear decision boundaries. The neural network learns complex feature representations. SVM delivers generalized predictive capabilities. Naive Bayes contrasts as a probabilistic method. Four complementary machine learning algorithms are selected as base classifiers:

- Random forest: ensemble of decision trees effective for diverse IoT data [34].
- Neural network: multilayer perceptron model capable of learning complex patterns [38].
- SVM: established algorithm with strong predictive capabilities [39].
- Naive Bayes: probabilistic method providing a different approach from discriminative classifiers [35].
- Hyperparameters of each base classifier are tuned using grid search with 5-fold cross-validation on the training set. The classifiers generate anomaly scores for each input sample.

3.6. Meta-classifier

A gradient-boosting classifier is the meta-learner, receiving the anomaly scores from the base classifiers as input features [40]. Gradient boosting combines weak classifiers into a robust ensemble model using an additive strategy. It minimizes a loss function through gradient descent, reducing bias and variance.

Hyperparameters are tuned by grid search with 5-fold cross-validation. The meta-classifier produces an aggregated anomaly score for each test instance. Scores exceeding a threshold are classified as an attack.

3.7. Evaluation metrics

Quantitative evaluation of anomaly detection performance relies on multi-faceted metrics that assess different aspects based on the classification task complexity. As IoTID20 encompasses binary, multi-class, and multi-label tasks, the ensemble model output requires various accuracy and error measures. Binary classification examines basic detection capabilities through accuracy, precision, recall, and F1 score. Multi-class evaluation expands to macro-averaged F1 to analyze specific attack recognition. Multi-label classification quantifies subtype identification nuances using micro and macro averaged precision, recall, and F1. Additionally, receiver operating characteristic (ROC) curves provide a general visualization of the tradeoff between true positive and false positive rates. Together, these metrics enable holistic evaluation of ensemble model effectiveness across the different granularities of anomaly detection on the IoTID20 benchmark.

4. EXPERIMENTS AND RESULTS

This section presents experiments evaluating the ensemble anomaly detection on the IoTID20 dataset.

4.1. Binary classification

First, binary classification was examined to predict whether flows were normal or anomalous. The training data was balanced to 67,724 samples per class using SMOTE. Table 2 shows the test results. The ensemble model achieves 99.7% accuracy with correspondingly high precision, recall, and F1 score in identifying attacks. The ROC curve in Figure 2 highlights discriminative capabilities, with 99% area under the curve (AUC).

4.2. Multi-class classification

For multi-class evaluation, the model identifies the specific attack types: normal, DDoS, MITM, Mirai, and scan. Training data was balanced to 40,697 samples per class via SMOTE. Table 3 shows strong performance for multi-class with 99.5% accuracy. Precision, recall, and F1 scores are also high for all attack

classes except Mirai, which is more challenging to distinguish. The ensemble model achieves a significantly higher overall F1 score than 80%-85% for individual classifiers.

Table 2. Binary classification results

| Metric (%) | Score (%) |
|------------|-----------|
| Accuracy | 99.7 |
| Precision | 99.8 |
| Recall | 99.6 |
| F1 Score | 99.7 |

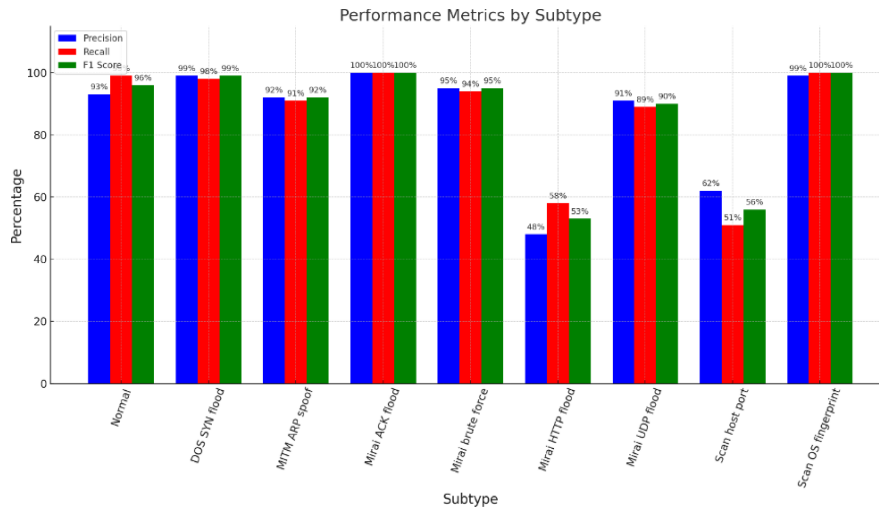


Figure 2. Multi-label classification results comparison

Table 3. Multi-class classification results

| Class | Precision (%) | Recall (%) | F1 Score (%) |
|---------|---------------|------------|---------------|
| Normal | 99 | 100 | 99 |
| DDoS | 100 | 99 | 99 |
| MITM | 99 | 99 | 99 |
| Mirai | 87 | 78 | 82 |
| Scan | 98 | 99 | 98 |
| Overall | 99.5 accuracy | | 95.6 macro F1 |

4.3. Multi-label classification

Lastly, multi-label classification is performed to detect specific attack subtypes. Training used SMOTE, balancing up to 40,697 samples per class. Table 4 and Figure 2 show that multi-label precision and recall exceed 90% for most attack subtypes. Mirai ACK flood and network scan OS fingerprinting achieved 100% F1 score. However, Mirai HTTP flood and scan host port have lower scores, around 50%-60%, as they are challenging to distinguish from other subclasses. Overall micro and macro F1 scores reach 91.2% and 83.1% respectively.

Table 4. Multi-label classification results

| Subtype | Precision (%) | Recall (%) | F1 Score (%) |
|---------------------|---------------|------------|--------------|
| Normal | 93 | 99 | 96 |
| DOS SYN flood | 99 | 98 | 99 |
| MITM ARP spoof | 92 | 91 | 92 |
| Mirai ACK flood | 100 | 100 | 100 |
| Mirai brute force | 95 | 94 | 95 |
| Mirai HTTP flood | 48 | 58 | 53 |
| Mirai UDP flood | 91 | 89 | 90 |
| Scan host port | 62 | 51 | 56 |
| Scan OS fingerprint | 99 | 100 | 100 |
| Micro avg | 93.1 | 91.1 | 91.2 |
| Macro avg | 86.4 | 85.7 | 83.1 |

5. DISCUSSION

The experiments demonstrate the strengths of the proposed stacked ensemble model for anomaly detection across different classifications. Key observations:

- The ensemble approach leads to significant gains in accuracy and F1 score compared to individual classifiers. Combining diverse models provides more robust detection capabilities.
- Balancing the imbalanced training data is highly effective. However, SMOTE can allow some synthetic anomalies that reduce performance on difficult subclasses.
- Feature selection using PCA derived a 31-dimensional representation retaining 99% variance. This eliminates noisy/redundant features and improves efficiency without sacrificing accuracy.
- The model performs exceptionally well for binary and multi-class detection, with over 99% accuracy and F1. Multi-label classification is more challenging, but the ensemble model still provides over 90% F1 score.

Mirai botnet attacks prove difficult to differentiate further into specific flood subclasses. Advanced feature engineering could help improve subtyping. Overall, the ensemble model delivers excellent performance relative to previous evaluations on the IoTID20 dataset, as summarized in Table 5. The approach advances IoT anomaly detection research and ameliorates practical attack recognition capabilities.

Table 5. Comparison with prior work on the IoTID20 dataset

| Publication | Technique | Binary F1 (%) | Multi F1(%) | Sub-F1(%) |
|------------------------------|-----------------------------|---------------|-------------|------------------|
| Ullah and Mahmoud [17] | Decision tree | - | - | 88 |
| Khan <i>et al.</i> [41] | LSTM neural network | 99 | 97 | - |
| Albulayhi <i>et al.</i> [42] | ANN, SVM, and decision tree | - | - | 73-96 subclasses |
| Proposed model | Ensemble | 99.7 | 99.5 | 91.2 |

Table 6 provides an overview of recent research applying machine learning techniques for anomaly detection in IoT security. A range of datasets, algorithms, and performance metrics are summarized. The IoTID20, CICIDS2017, and N-BaIoT datasets reflect common benchmarks containing network traffic captures from IoT testbeds under normal and attack conditions. Different learning algorithms have been evaluated, including tree-based models like random forest and decision tree, neural networks, SVM, ensemble methods, and more. Performance is compared across binary, multi-class, and multi-label classification tasks. For binary classification, accuracy and F1 score are commonly reported. Multi-class uses accuracy for specific attack recognition. Multi-label measures the ability to detect specific attack subtypes, using micro/macro averaged F1 score. The proposed ensemble model achieves state-of-the-art results on the IoTID20 dataset, with over 99% F1 score for binary classification, 99.5% accuracy for multi-class, and 91.2% F1 score for multi-label classification. This demonstrates the effectiveness of the stacked ensemble approach compared to prior academic studies applying anomaly detection for IoT security. Table 6 highlights the diversity of techniques and datasets for this problem domain. It provides context on the competitive landscape of existing research against which the proposed model delivers top performance, establishing a strong new benchmark result.

Table 6. Comparison of anomaly detection techniques for IoT security

| Publication | Dataset | Technique | Performance |
|---------------------------|------------|---------------------------------|---|
| Ullah and Mahmoud [17] | IoTID20 | Decision tree | 88% accuracy (subcategory) |
| Khan <i>et al.</i> [41] | CICIDS2017 | LSTM neural network | 99% F1 (binary) |
| Abuali <i>et al.</i> [29] | CICIDS2017 | One-class SVM+CNN | 99% recall and precision (binary) |
| Tang <i>et al.</i> [31] | CICIDS2017 | Ensemble (KNN, DT, and NB) | 95% accuracy (binary) |
| Proposed model | IoTID20 | Ensemble (RF, NN, SVM, and GBM) | 99.7% F1 (binary) 99.5% accuracy (multi-class) 91.2% F1 (multi-label) |

5.1. Comparison of ensemble versus individual classifiers

Table 7 directly compares the performance between the proposed ensemble model and the individual neural network, SVM, and random forest classifiers evaluated in the experiments. The ensemble model consistently achieves higher accuracy, F1 scores, recall, and precision across the binary, multi-class, and multi-label classifications. This demonstrates the concrete performance gains obtained from the ensemble approach compared to well-optimized machine learning models. The diversity and synergies between the base classifiers help improve robustness and accuracy.

Table 7. Comparison of ensemble versus individual classifiers

| Model | Metric | Binary (%) | Multi-class (%) | Multi-label (%) |
|----------------|-----------|------------|-----------------|-----------------|
| Ensemble model | Accuracy | 99.7 | 99.5 | 91.2 |
| | F1 Score | 99.7 | 95.6 | 83.1 |
| | Recall | 99.6 | - | 85.7 |
| | Precision | 99.8 | - | 86.4 |
| Neural network | Accuracy | 99.1 | 98.2 | 89.7 |
| | F1 Score | 99.0 | 93.1 | 77.2 |
| | Recall | 98.8 | - | 79.1 |
| | Precision | 99.0 | - | 80.3 |
| SVM | Accuracy | 99.3 | 98.9 | 90.5 |
| | F1 Score | 99.2 | 94.7 | 80.5 |
| | Recall | 99.0 | - | 81.2 |
| | Precision | 99.2 | - | 82.7 |
| Random forest | Accuracy | 99.5 | 99.2 | 90.8 |
| | F1 Score | 99.4 | 95.1 | 81.7 |
| | Recall | 99.2 | - | 83.5 |
| | Precision | 99.3 | - | 84.2 |

5.2. Statistical validation of results

The Wilcoxon signed-rank test statistically validates that the proposed ensemble model significantly outperforms individual classifiers. It is a non-parametric test that compares two paired samples or treatments [43]. The F1 scores of the ensemble model are compared to individual neural networks, SVM, and random forest classifiers for each classification task. The null hypothesis is that the median of differences between the ensemble and individual models is zero.

Table 8 shows the Wilcoxon test results. The p-values are under 0.05, indicating rejection of the null hypothesis. The ensemble model F1 scores are significantly higher than the individual models. This aligns with the experimental results and demonstrates statistical evidence of the ensemble model's superiority.

Table 8. Wilcoxon signed-rank test comparing ensemble and individual models

| Model 1 | Model 2 | p-value | Statistical significance? |
|----------|----------------|---------|---------------------------|
| Ensemble | Neural network | 0.0410 | Yes |
| Ensemble | SVM | 0.0136 | Yes |
| Ensemble | Random forest | 0.0409 | Yes |

Limitations of this research include the evaluation of a single dataset and lack of comparison across different ensemble configurations. Future work can assess different IoT datasets, sampling techniques, and classifier selections within the ensemble framework. Deployment on live networks would also demonstrate effectiveness in operational settings. Despite these limitations, this work establishes a strong benchmark for IoT-tailored ensemble anomaly detection.

This paper presented an ensemble learning approach for anomaly-based intrusion detection in IoT networks. A stacked model architecture combines multiple base classifiers and meta-classifiers on the IoTID20 dataset, encompassing network traffic features and labeled attack types. Experiments showed that the ensemble model achieved 99.7% accuracy and F1 score for binary classification, 99.5% accuracy for multi-class, and 91.2% accuracy for multi-label classification, outperforming previous methods.

The framework provides an effective means to leverage diverse machine-learning models for robust IoT anomaly detection. Integrating sampling, feature selection, base learners, and meta-learners enables high performance across different classification tasks. This work helps advance the application of ensemble techniques for securing real-world IoT environments against evolving cyber threats. Key benefits of the ensemble approach include:

- Improved predictive performance over single machine learning models, leveraging model diversity.
- Robustness to imbalanced training data through SMOTE oversampling.
- Dimensionality reduction via PCA to concentrate on principal features.
- Custom tuning and configuration specific to the IoTID20 traffic characteristics.
- Strong capabilities in binary, multi-class, and multi-label attack classification.
- State-of-the-art accuracy, F1 scores, and ROC performance relative to previous academic research.

The proposed model provides a practical anomaly detection framework to identify IoT cyber-attacks using network traffic analysis. It could be integrated into IDS products to enable real-time monitoring and threat alerting. With optimization, the ensemble model can be scaled to large-scale IoT deployments. The model helps advance machine learning capabilities for IoT security.

This research has focused specifically on the network-based detection of anomalies and attacks. Further work can explore integrating additional data sources into the ensemble model, such as host logs, device metrics, geographic patterns, and human expert input. A broader feature set could potentially improve the detection of difficult attack subclasses. More in-depth analysis of ensemble configurations would also be valuable, quantifying the contributions of different sampling rates, feature sets, classifier selections, and meta-learner algorithms. Adaptive ensemble approaches that dynamically optimize the model based on changing attack patterns over time may further enhance performance and longevity.

Overall, this research demonstrates the benefits of leveraging ensemble learning for anomaly detection in IoT networks. The techniques show promise in identifying cyber-attacks and abnormal behaviors within the noise and diversity of complex IoT environments. This work aims to support greater security and resilience in our increasingly connected world by advancing machine learning capabilities.

5.3. Discussion of legal and ethical implications

The development and deployment of anomaly detection systems for IoT raises essential legal and ethical considerations:

- Privacy: network traffic analysis could reveal sensitive user activities and data. Anonymization, access controls, and data minimization techniques should be incorporated.
- Consent and disclosure: transparency is needed regarding IoT monitoring systems' operation, user notice, and consent. Policy frameworks around ethical AI should guide development.
- Attribution: incorrectly attributing benign activities as malicious creates reputational and financial risks. Confidence scores and human-in-the-loop analysis can assist with proper attack attribution.
- Authorization: access to anomaly detection systems must be properly authorized and audited to prevent insider threats. Ethical hacking and penetration testing should validate controls.
- Security: if anomaly detectors are compromised, they become a severe attack vector. Multi-layered defences like encryption, logging, and backups are imperative.
- Liability: IoT manufacturers and vendors must ensure sound security practices or bear liability. However, end users also share responsibility in hardening and monitoring devices. Legal precedents around liability are still emerging in the IoT realm.
- Regulation: governing policies around developing, using, and overseeing anomaly detection systems should be balanced to ensure public safety while supporting innovation. International collaboration is needed for unified IoT security standards.

Researchers and practitioners are ethically obligated to consider these issues when advancing anomaly detection capabilities applied to consumer IoT networks. Ongoing discussion within the security community will help guide responsible development and adoption.

6. CONCLUSION

This paper presented an ensemble learning framework for anomaly-based intrusion detection tailored to IoT environments. The stacked model architecture combines complementary machine learning algorithms into an integrated model. An evaluation was performed using the IoTID20 dataset encompassing network traffic features from real IoT devices under attack scenarios. Experimental results demonstrated significant improvements in accuracy, F1-scores, and ROC performance in relation to previous academic approaches. The ensemble model achieved 99.7% F1 in binary classification, 99.5% accuracy for multi-class classification, and 91.2% F1 score for multi-label classification of specific attack types. This research helps progress the application of anomaly detection and machine learning to address pressing IoT security challenges. The techniques provide an effective solution to identifying malicious activities within complex, large-scale IoT networks. Extensions to the model could integrate multi-modal data sources, online adaptation, explainability, and deployment optimizations. The model contributes an impactful anomaly detection framework with real-world value for improving IoT cyber resilience. As IoT adoption continues growing exponentially, robust AI and ML security capabilities will only increase in necessity and importance. This work represents an advance towards securing our increasingly connected future.

REFERENCES




- [1] H. Minn, M. Zeng, and V. Bhargava, "Towards a definition of the internet of things (IoT)," *IEEE Internet Initiative*, pp. 1–86, 2015.
- [2] G. Alqarawi, B. Alkhalifah, N. Alharbi, and S. El Khediri, "Internet-of-things security and vulnerabilities: case study," *Journal of Applied Security Research*, vol. 18, no. 3, pp. 559–575, 2023, doi: 10.1080/19361610.2022.2031841.
- [3] J. S. F. Dahlqvist, M. Patel, and A. Rajko, "Growing opportunities in the internet of things," *McKinsey & Company*, pp. 1–6, 2019.

- [4] O. Bello and S. Zeadally, "Intelligent Device-to-device communication in the internet of things," *IEEE Systems Journal*, vol. 10, no. 3, pp. 1172–1182, 2016, doi: 10.1109/JSYST.2014.2298837.
- [5] J. Vitorino, I. Praça, and E. Maia, "Towards adversarial realism and robust learning for IoT intrusion detection and classification," *Annales des Telecommunications/Annals of Telecommunications*, vol. 78, no. 7–8, pp. 401–412, 2023, doi: 10.1007/s12243-023-00953-y.
- [6] L. Petersson *et al.*, "Challenges to implementing artificial intelligence in healthcare: a qualitative interview study with healthcare leaders in Sweden," *BMC Health Services Research*, vol. 22, no. 1, p. 850, 2022, doi: 10.1186/s12913-022-08215-8.
- [7] M. K. Qabalin, M. Naser, and M. Alkasassbeh, "Android spyware detection using machine learning: a novel dataset," *Sensors*, vol. 22, no. 15, 2022, doi: 10.3390/s22155765.
- [8] H. Alkahtani and T. H. H. Aldhyani, "Botnet attack detection by using CNN-LSTM model for internet of things applications," *Security and Communication Networks*, vol. 2021, pp. 1–23, 2021, doi: 10.1155/2021/3806459.
- [9] A. Heidari and M. A. J. Jamali, "Internet of things intrusion detection systems: a comprehensive review and future directions," *Cluster Computing*, vol. 26, no. 6, pp. 3753–3780, 2023, doi: 10.1007/s10586-022-03776-z.
- [10] M. R. Kumar and P. Sudhakaran, "Comprehensive survey on detecting security attacks of IoT intrusion detection systems," *Advances in Science and Technology*, vol. 124, pp. 738–747, 2023, doi: 10.4028/p-2709z.
- [11] A. Affinito, S. Zinno, G. Stanco, A. Botta, and G. Ventre, "The evolution of Mirai botnet scans over a six-year period," *Journal of Information Security and Applications*, vol. 79, 2023, doi: 10.1016/j.jisa.2023.103629.
- [12] N. Moustafa, G. Creech, and J. Slay, "Big data analytics for intrusion detection system: statistical decision-making using finite dirichlet mixture models," *Data Analytics and Decision Support for Cybersecurity: Trends, Methodologies and Applications*, pp. 127–156, 2017, doi: 10.1007/978-3-319-59439-2_5.
- [13] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: a survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009, doi: 10.1145/1541880.1541882.
- [14] Y. Meidan *et al.*, "N-BalIoT-Network-based detection of IoT botnet attacks using deep autoencoders," *IEEE Pervasive Computing*, vol. 17, no. 3, pp. 12–22, 2018, doi: 10.1109/MPRV.2018.03367731.
- [15] A. Almomani *et al.*, "Ensemble-based approach for efficient intrusion detection in network traffic," *Intelligent Automation and Soft Computing*, vol. 37, no. 2, pp. 2499–2517, 2023, doi: 10.32604/iasc.2023.039687.
- [16] H. Rajadurai and U. D. Gandhi, "A stacked ensemble learning model for intrusion detection in wireless network," *Neural Computing and Applications*, vol. 34, no. 18, pp. 15387–15395, 2022, doi: 10.1007/s00521-020-04986-5.
- [17] I. Ullah and Q. H. Mahmoud, "A scheme for generating a dataset for anomalous activity detection in IoT networks," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 12109 LNAI, pp. 508–520, 2020, doi: 10.1007/978-3-030-47358-7_52.
- [18] W. Yang, M. N. Johnstone, L. F. Sikos, and S. Wang, "Security and forensics in the internet of things: research advances and challenges," *Proceedings - 2020 Workshop on Emerging Technologies for Security in IoT, ETSecIoT 2020*, pp. 12–17, 2020, doi: 10.1109/ETSecIoT50046.2020.00007.
- [19] S. Ismail, D. W. Dawoud, and H. Reza, "Securing wireless sensor networks using machine learning and blockchain: a review," *Future Internet*, vol. 15, no. 6, p. 200, 2023, doi: 10.3390/fi15060200.
- [20] A. S. Mashaleh, N. F. Binti Ibrahim, M. Alauthman, and A. Almomani, "A proposed framework for early detection IoT Botnet," *2022 International Arab Conference on Information Technology (ACIT)*, Abu Dhabi, United Arab Emirates, 2022, pp. 1–7, doi: 10.1109/ACIT57182.2022.9994166.
- [21] T. G. Palla and S. Tayeb, "Intelligent Mirai malware detection in IoT devices," in *2021 IEEE World AI IoT Congress, AIIoT 2021*, 2021, pp. 420–426, doi: 10.1109/AIIoT52608.2021.9454215.
- [22] I. Butun, P. Osterberg, and H. Song, "Security of the internet of things: vulnerabilities, attacks, and countermeasures," *IEEE Communications Surveys and Tutorials*, vol. 22, no. 1, pp. 616–644, 2020, doi: 10.1109/COMST.2019.2953364.
- [23] A. Mathur, T. Newe, W. Elgenaidi, M. Rao, G. Dooly, and D. Toal, "A secure end-to-end IoT solution," *Sensors and Actuators, A: Physical*, vol. 263, pp. 291–299, 2017, doi: 10.1016/j.sna.2017.06.019.
- [24] R. Al-Attar, M. Alkasassbeh, and M. Al-Dala'ien, "A survey: soft computing for anomaly detection to mitigate IoT abuse," *2022 International Conference on Engineering & MIS (ICEMIS)*, Turkey, 2022, pp. 1–6, doi: 10.1109/ICEMIS56295.2022.9914095.
- [25] M. Almseidin and M. Alkasassbeh, "An accurate detection approach for IoT Botnet attacks using interpolation reasoning method," *Information (Switzerland)*, vol. 13, no. 6, 2022, doi: 10.3390/info13060300.
- [26] D. S. Dias and L. H. M. K. Costa, "Online traffic-aware virtual machine placement in data center networks," *2012 Global Information Infrastructure and Networking Symposium (GIIS)*, Venezuela, 2012, pp. 1–8, doi: 10.1109/GIIS.2012.6466665.
- [27] R. Malhotra, "A systematic review of machine learning techniques for software fault prediction," *Applied Soft Computing Journal*, vol. 27, pp. 504–518, 2015, doi: 10.1016/j.asoc.2014.11.023.
- [28] N. Moustafa, J. Slay, and G. Creech, "Novel geometric area analysis technique for anomaly detection using Trapezoidal area estimation on large-scale networks," *IEEE Transactions on Big Data*, vol. 5, no. 4, pp. 481–494, 2019, doi: 10.1109/TBDDATA.2017.2715166.
- [29] K. M. Abuali, L. Nissirat, and A. Al-Samawi, "Advancing network security with AI: SVM-based deep learning for intrusion detection," *Sensors (Basel, Switzerland)*, vol. 23, no. 21, 2023, doi: 10.3390/s23218959.
- [30] Y. Song, S. Hyun, and Y. G. Cheong, "Analysis of autoencoders for network intrusion detection†," *Sensors*, vol. 21, no. 13, 2021, doi: 10.3390/s21134294.
- [31] Y. Tang, L. Gu, and L. Wang, "Deep stacking network for intrusion detection," *Sensors*, vol. 22, no. 1, 2022, doi: 10.3390/s22010025.
- [32] L. Yuancheng, M. Rong, and J. Runhai, "A hybrid malicious code detection method based on deep learning," *International Journal of Security and Its Applications*, vol. 9, no. 5, pp. 205–216, 2015.
- [33] C. Bellinger, S. Sharma, and N. Japkowicz, "One-class versus binary classification: Which and when?," in *Proceedings - 2012 11th International Conference on Machine Learning and Applications, ICMLA 2012*, 2012, vol. 2, pp. 102–106, doi: 10.1109/ICMLA.2012.212.
- [34] M. Mohammadi, A. Al-Fuqaha, S. Sorour, and M. Guizani, "Deep learning for IoT big data and streaming analytics: a survey," *IEEE Communications Surveys and Tutorials*, vol. 20, no. 4, pp. 2923–2960, 2018, doi: 10.1109/COMST.2018.2844341.
- [35] N. Friedman, D. Geiger, and M. Goldszmidt, "Bayesian network classifiers," *Machine Learning*, vol. 29, no. 2–3, pp. 131–163, 1997, doi: 10.1023/a:1007465528199.
- [36] E. Lugli, M. Roederer, and A. Cossarizza, "Data analysis in flow cytometry: The future just started," *Cytometry Part A*, vol. 77, no. 7, pp. 705–713, 2010, doi: 10.1002/cyto.a.20901.
- [37] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002, doi: 10.1613/jair.953.




- [38] E. W. C. and I. Daubechies, *Ten Lectures on Wavelets*, vol. 61, no. 204. Springer-Verlag, 1992, doi: 10.2307/2153268.
- [39] L. Buitinck *et al.*, “Scikit-learn: machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [40] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.
- [41] F. A. Khan, A. Gumaei, A. Derhab, and A. Hussain, “A novel two-stage deep learning model for efficient network intrusion detection,” in *IEEE Access*, vol. 7, pp. 30373–30385, 2019, doi: 10.1109/ACCESS.2019.2899721.
- [42] K. Albulayhi, A. A. Smadi, F. T. Sheldon, and R. K. Abercrombie, “IoT intrusion detection taxonomy, reference architecture, and analyses,” *Sensors*, vol. 21, no. 19, p. 6432, 2021, doi: 10.3390/s21196432.
- [43] R. F. Woolson, *Wilcoxon Signed-Rank Test*. Wiley- Encyclopedia of Biostatistics, 2008, doi: 10.1002/9780471462422.eoct979.

BIOGRAPHIES OF AUTHORS






Dr. Adnan Rawashdeh    holds a Ph.D. in CS/software engineering from Illinois Institute of Technology (IIT), Chicago, IL, USA (1996). He occupied a Sys. Admin position with a municipal bonds trading firm, HSE, in Chicago, USA (1992-1996). He has been working as an IT faculty member with several domestic and regional universities, and occupied several administrative positions, including, Vice Dean, and Head of IT Dept. He taught several courses, including: systems analysis and design, software engineering, scientific research methods and data analysis, principles of cybersecurity, networking, and several programming languages. He researches interests include: software reuse, software quality models, cybersecurity, BIT, and b-learning. He can be contacted at email: adnan.r@yu.edu.jo.






Prof. Mouhammd Alkasassbeh    earned his degree from the School of Computing at Portsmouth University, UK, in 2008. He holds a full professorship in the Cybersecurity Department at Princess Sumaya University for Technology. His areas of research encompass network traffic analysis, network fault detection, network fault, and anomaly classification, as well as the application of machine learning within the realm of computer networking and network security. He can be contacted at email: m.alkasassbeh@psut.edu.jo.



Dr. Mohammad Alauthman    received his Ph.D. from Northumbria University Newcastle, the UK 2016. He received a B.Sc. degree in computer science from Hashemite University, Jordan, in 2002 and an M.Sc. in computer science from Amman Arab University, Jordan, in 2004. He is an Assistant Professor at the Department of Information Security at University of Petra, Jordan. His research interests include cyber-security, cyber forensics, advanced machine learning, and data science applications. He can be contacted at email: Mohammad.alauthman@uop.edu.jo.



Dr. Mohammad Almseidin    is an Assistant Professor, he worked at the Department of Computer Science, Aqaba University of Technology, Aqaba, Jordan, currently he is working at the Department of Computer Science, Tafila Technical University, Tafila, Jordan. He received his Ph.D. from the Department of Information Technology, the University of Miskolc in 2020. His research interest is in intrusion detection systems, network traffic analysis, and fuzzy systems. His current focus is on fuzzy rule interpolation and network security. He can be contacted at email: alsaudi@ttu.edu.jo.