UtilitasMathematica

ISSN 0315-3681 Volume 122 (2), 2025

AI-Driven Security Approaches for Detecting Botnet Attacks in IoT Environments

Sridhar Sriharsha Rachakonda (Primary)
Senior Staff Engineer
NVIDIA Corporation
Department of Information Technology
Leander, Texas, USA
rssharsha@zohomail.com

Ramesh Lakshmikanth
Enterprise Engineering Manager
NVIDIA Corporation
Department of Information Technology
San Jose, CA, USA
rameshkl007@outlook.com

***Abstract*— The Internet of Things (IoT) continues to expand rapidly, but this growth also exposes networks to sophisticated botnet attacks such as Mirai and Gafgyt that compromise large numbers of devices. Intrusion Detection Systems (IDS) are critical for securing IoT environments. Yet, traditional models often fail to address the dual challenges of extreme class imbalance and diverse attack subtypes in real-world traffic. This study analyzes a large-scale IoT IDS dataset containing over 7 million network flow samples, where normal traffic constitutes only 8.54% of records, and attack traffic includes UDP, TCP, SYN, SCAN, and ACK floods. The dataset provides 23 engineered temporal and relational features, including entropy, mutual information, host-host statistics, jitter, and host-port behavior that capture subtle patterns of malicious activity. This research proposes a novel hybrid ensemble model that integrates cost-sensitive LightGBM, XGBoost, deep neural networks, and Isolation Forest anomaly detection, combining supervised learning with unsupervised detection of zero-day threats. Experimental evaluation demonstrates that the proposed model significantly outperforms baseline approaches such as logistic regression and random forests, achieving superior Precision-Recall AUC, F1-score, and Matthews Correlation Coefficient. These results highlight the potential of advanced AI techniques to enhance resilience against evolving botnet threats in IoT environments.**

Keywords*— *Intrusion Detection Systems, Botnet Attacks, Machine Learning, Ensemble Models, Cybersecurity

I. Introduction

The Internet of Things (IoT) has revolutionized the digital world, where trillions of connected devices are used in homes, medical applications, business, and industry. This new era of interconnectedness has brought on significant advantages in automation, efficiency, and real-time data processing. Nonetheless, the same connectivity levels have increased the surface of cybercriminal attacks. IoT devices frequently have limited computing power, minimal in-built protection, and thus are easily exploited. One of the most severe threats in this domain is the formation of botnets, where compromised IoT devices are orchestrated to launch large-scale distributed denial-of-service (DDoS) attacks or facilitate unauthorized data access [1]. The detection of botnet attacks in IoT setups is thus a critical research concern since the existing security systems aimed at defending traditional computing systems are insufficient in dealing with the heterogeneity of IoT traffic and the inability of the conventional security systems to handle the volume of such traffic. Intrusion Detection Systems (IDS) can play a significant role in protecting IoT networks as their operation monitors the network traffic behavior and raises an alarm when malicious traffic is detected. However, the use of IDS in the IoT presents peculiar problems. First, IoT traffic datasets are inherently imbalanced, where benign traffic often represents less than ten percent of total flows, as observed in large-scale intrusion datasets. Such an imbalance causes conventional classifiers to bias predictions toward attack classes, leading to high false-negative rates in benign traffic detection. Second, the attack surface in IoT is diverse, with botnet families such as Mirai and Gafgyt generating multiple subtypes of attacks, including UDP floods, TCP floods, SYN floods, and network scans [1-2]. Each subtype exhibits distinct statistical and temporal patterns, making generalization across classes difficult for traditional models. Finally, IoT traffic evolves continuously, and zero-day attacks pose significant risks, as models trained on past data may fail to detect emerging threats.

Artificial Intelligence (AI) and machine learning (ML) techniques have shown promise in addressing these challenges by learning complex, non-linear relationships within IoT network flows [3]. Gradient-boosted decision trees, deep neural networks, and unsupervised anomaly detection methods have been increasingly adopted to improve detection rates and adaptability. However, single-model approaches remain limited in handling severe imbalance, ensuring real-time performance, and generalizing to unseen attack variants. To overcome these shortcomings, this study introduces an AI-driven ensemble framework that leverages cost-sensitive boosting, deep learning, and anomaly detection to enhance the robustness of IDS. By integrating multiple supervised and unsupervised models, the proposed approach aims to reduce false negatives, improve the detection of minority benign traffic, and strengthen resilience against zero-day botnet behaviors. Moreover, botnet attacks on critical infrastructure like healthcare systems, smart grids, and intelligent transport make the consequences of the undetected attacks severe, since more people can lose their lives. A successful attack can disrupt essential services, compromise sensitive patient or industrial data, and even endanger human safety [4]. Therefore, advancing IDS solutions that improve detection accuracy and maintain low latency and scalability is crucial [5]. The proposed research contributes to this pressing need by investigating AI-driven security approaches capable of handling large-scale imbalanced datasets, distinguishing diverse attack subtypes, and offering real-time adaptability for securing modern IoT ecosystems.

II. Related Works

The number of studies on intrusion detection in the IoT environment has significantly increased over the past few years, mainly because of the growing number of botnet attacks

such as Mirai and Gafgyt that exploit poorly secured devices. Traditional intrusion detection approaches, often based on signature matching or rule-based systems, fail to provide robust protection against evolving and large-scale IoT attacks [6]. Consequently, researchers have turned to Artificial Intelligence (AI) and machine learning techniques to build adaptive and data-driven Intrusion Detection Systems (IDS) [7]. Machine learning models such as decision trees, k-nearest neighbors, and support vector machines have been widely adopted in early IoT IDS research due to their simplicity and interpretability [8]. However, these models generally underperform on large and imbalanced datasets, often resulting in high false-negative rates when distinguishing between benign traffic and complex attack subtypes. To address these shortcomings, ensemble-based classifiers such as Random Forest and Gradient Boosting have gained popularity. For instance, studies have shown that tree-based ensembles achieve higher detection accuracy by aggregating weak learners and leveraging variance in network traffic features [9][10]. Specifically, XGBoost and LightGBM have emerged as state-of-the-art boosting algorithms for IDS tasks, demonstrating superior computational efficiency and resilience to class imbalance compared to conventional classifiers [11].

Deep learning approaches have advanced intrusion detection by modeling non-linear feature interactions and temporal dependencies in IoT traffic. Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) architectures have been applied to capture sequential attack behaviors, particularly in DDoS detection [12]. Similarly, Convolutional Neural Networks (CNNs) have been employed to extract high-level representations from raw network flows. While these methods improve detection rates, their computational demands and lack of explainability limit real-world deployment, especially on resource-constrained IoT devices. To overcome these limitations, lightweight deep models and hybrid frameworks have been proposed, combining deep feature extraction with fast tree-based classifiers [13]. Another promising direction in IoT intrusion detection is integrating unsupervised anomaly detection methods. Models such as autoencoders and Isolation Forests have effectively identified previously unseen (zero-day) attacks by learning normal behavior and flagging deviations [14]. The Kitsune IDS, for example, leverages an ensemble of autoencoders to detect anomalies in streaming IoT traffic [15]. However, anomaly-based methods often suffer from high false positive rates, particularly in highly heterogeneous IoT environments. Hybrid approaches that combine supervised learning for known attack types with anomaly detection for novel threats have therefore been recommended as a more balanced solution [16].

The most recent literature also attests to the need to resolve the data imbalance issue in the dataset inherent within IoT IDS, where the proportion of malicious traffic is significantly higher than standard samples. Techniques such as Synthetic Minority Oversampling Technique (SMOTE), Adaptive Synthetic Sampling (ADASYN), and cost-sensitive learning have been applied to improve the detection of minority classes [17]. In addition, researchers have highlighted the need for evaluating IDS models with metrics suited to imbalanced data, such as Precision-Recall AUC, Matthews Correlation Coefficient, and Balanced Accuracy, rather than relying solely on overall accuracy [18]. Beyond these approaches, several recent studies have explored federated and distributed learning techniques for IoT intrusion detection, aiming to preserve data privacy while leveraging collaborative training across devices [19]. Other works have focused on explainable AI methods, such as using SHAP or LIME to improve the transparency of IDS decisions and facilitate trust in operational environments [20]. Researchers have also begun applying reinforcement learning for adaptive intrusion response, where the IDS detects and dynamically reacts to attack patterns [21]. Comparative evaluations of benchmark datasets, including NSL-KDD, CICIDS2017, and IoT-specific corpora, consistently show that hybrid and ensemble models outperform single classifiers in terms of both detection accuracy and resilience to evolving threats [22]. These studies collectively highlight the growing trend toward multi-model, interpretable, and privacy-aware IDS frameworks, underscoring the relevance of developing AI-driven solutions tackling imbalance, scalability, and zero-day attack detection in IoT ecosystems. While prior research has laid a strong foundation by applying machine learning, deep learning, and anomaly detection methods to IoT intrusion detection, significant gaps remain in handling severe imbalance, detecting diverse attack subtypes, and ensuring robustness against zero-day threats. This study addresses these gaps by proposing a hybrid AI-driven ensemble model that integrates boosting algorithms, deep neural networks, and unsupervised anomaly detection to deliver a more resilient and adaptive IDS for detecting botnet attacks in IoT environments.

III. METHODOLOGY

The methodology used in this research aims to address the challenge of identifying botnet attacks in IoT environments through advanced AI-based strategies. The dataset contains over seven million records with a highly imbalanced distribution, so careful attention was given to sample selection, feature analysis, preprocessing, and model design. An exploratory study on a 200-sample subset examined feature behaviors and attack subtype dominance, leading to recommendations for stratified sampling in further experiments. Engineered features such as entropy, mutual information, host-host statistics, and host-port activity informed preprocessing and model choices. The proposed framework contrasts baseline models with a hybrid ensemble architecture combining supervised boosting algorithms, deep learning, and unsupervised anomaly detection.

A. Dataset Description

The experimental study is based on a large-scale IoT intrusion detection dataset obtained from Kaggle containing 7,062,606 network flow records and 27 attributes, of which 23 are numeric features, 3 are categorical descriptors (Device_Name, Attack, Attack_subType), and one is the binary class label (0 = attack, 1 = normal). The dataset signifies real-world traffic obtained through a diverse set of IoT devices such as IP cameras, smart thermostats, doorbells, and baby monitors, which are known to be common targets for botnet infections. The traffic samples include benign network flows (555,932 instances, ~8.54% of total) and malicious flows (6,506,674 instances, ~91.46%), dominated by botnet families such as Mirai and Gafgyt [25]. Attack subtypes include UDP, TCP, SYN, SCAN, ACK floods, and more specialized variants such as “combo” and “udpplain.” This diversity ensures that the dataset captures the heterogeneity of IoT botnet attacks and provides a challenging benchmark for IDS model development.

B. Experimental Sampling

Although the dataset contains over 7 million records, for this phase of analysis, only the first 200 samples were extracted for detailed exploratory analysis. This subset provided an initial baseline for understanding feature behavior, but it introduced a critical challenge: all 200 rows belonged to a single attack subtype (Gafgyt/Combo). Consequently, the data was single-class, making cross-validation and binary classification infeasible in the exploratory stage. While this constraint limited the classification experiments, it offered valuable insights into feature distributions, correlations, and redundancy. To address the problem of single-class data that might arise, the study suggests stratified random sampling of the complete data set with proportional representation of both the normal and attack classes. Future iterations of this research will incorporate such balanced samples to validate model effectiveness across the complete dataset.

C. Feature Engineering and Relevance

The dataset includes engineered temporal and relational features designed to capture complex behaviors in IoT network traffic. These features fall into five key categories:

- Mutual Information (MI_dir): Measures dependence between packet flow directions over 0.1-second windows, identifying coordinated communication typical of botnets.
- Entropy (H): Reflects randomness or regularity in traffic; low entropy may indicate repetitive botnet behavior, while high entropy suggests evasion attempts.
- Host-Host Statistics (HH): Captures communication intensity, variability, and correlation between pairs of devices, helpful in detecting distributed attack coordination.
- Host-Host Jitter (HH_jit): Measures timing irregularities in communication flows, helping to identify flooding or timing-based attacks.
- Host-Port Statistics (HpHp): Tracks host-to-port activity, often signaling port scanning or exploitation attempts.

D. Data Preprocessing

Data preprocessing was critical in preparing the IoT intrusion dataset for effective model training and evaluation [23-24]. Since the dataset contained more than seven million records with 23 numeric and multiple engineered features, the first step involved data cleaning, where missing values and infinite values were replaced using median imputation to preserve statistical consistency. To ensure the stability of machine learning models, categorical attributes such as *Device_Name*, *Attack*, and *Attack_subType* were excluded from training and retained only for post-analysis and performance breakdowns. Scaling was applied selectively: StandardScaler was used for deep learning models such as neural networks, while boosting algorithms like LightGBM and XGBoost were trained directly on raw values since they handle feature scaling internally. The most significant challenge was the dataset's extreme imbalance, with only 8.54% of flows representing benign traffic. To overcome this problem, synthetic oversampling algorithms, including SMOTE (Synthetic Minority Oversampling Technique) and ADASYN (Adaptive Synthetic Sampling), were utilized to train using a larger number of normal samples, in addition to cost-sensitive learning applied to tree-based classification models by assigning excessive misclassification penalties to the minority class. Additionally, threshold tuning based on Precision-Recall optimization was applied to reduce bias toward the majority class. Finally, the dataset was split using an 80–20 stratified train-test division, ensuring proportional representation of benign and malicious flows in each partition. These preprocessing steps established a balanced foundation for robust evaluation of both baseline and advanced models.

E. Baseline Models

Conventional logistic regression and random forest models were selected to establish reference performance. These methods have been widely applied in IDS but often fail to capture the non-linear feature dependencies in IoT traffic, especially under imbalance. As expected, logistic regression produced unstable results in the single-class subset analysis, further underscoring the limitations of traditional models for this dataset.

F. Proposed Hybrid Ensemble Model

To overcome the shortcomings of traditional classifiers and improve robustness in detecting botnet attacks in IoT systems, this research proposes a Hybrid Ensemble Model (PHEM) that integrates supervised and unsupervised learning. The framework combines cost-sensitive boosting algorithms (LightGBM and XGBoost), a Deep Neural Network (DNN) for capturing complex non-linear feature interactions, and an Isolation Forest for zero-day attack detection. These base learners' output is aggregated through a stacking architecture with a logistic regression meta-learner calibrated by Platt scaling. This hybrid approach blends supervised classifiers with anomaly detection, improving accuracy for known botnet families such as Mirai and Gafgyt while enhancing resilience against emerging threats, thereby reducing false negatives and maintaining real-time applicability. Fig. 1 illustrates the overall architecture of the Proposed Hybrid Ensemble Model (PHEM) for detecting botnet attacks in IoT environments.

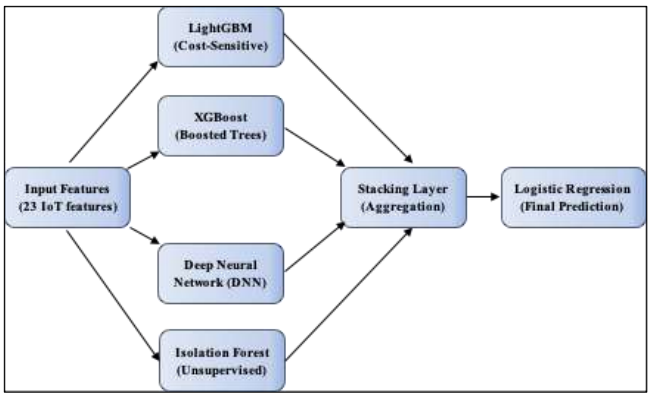

Fig. 1. Architecture of the Proposed Hybrid Ensemble Model (PHEM) for detecting botnet attacks in IoT environments.

Several key metrics were applied to evaluate model performance under extreme imbalance. The F1-score was calculated as:

$$F1 = 2 X (Precision X Recall) / (Precision + Recall) \quad (1)$$

Where:

$$Precision = TP / (TP + FP) \quad (2)$$

$$Recall = TP / (TP + FN) \quad (3)$$

$$MCC = \frac{(TP\ X\ TN) - (FP\ X\ FN)}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (4)$$

In the statistical equations, TP (True Positives) denotes the number of attack samples appropriately identified, and TN (True Negatives) represents the number of benign samples correctly identified. False Positives (FP) are benign traffic that are incorrectly identified as attacks, and False Negatives (FN) represent attacks classified as benign. Precision measures how many correctly identified attacks are relative to the number of predicted attacks, whereas Recall describes how many actual attacks were detected. F1-score trades between precision and recall, which makes it appropriate in imbalanced data. The Matthews Correlation Coefficient (MCC) gives an equally weighted score with a single, all-inclusive score: all four varieties (TP, TN, FP, FN) are considered within a score that is considered more reliable than accuracy in cases of skewed databases.

IV. Results And Discussion

The results and discussion section presents the exploratory analysis of the IoT intrusion dataset and evaluation of the proposed model. It highlights how dataset characteristics such as imbalance, attack diversity, and device heterogeneity influenced preprocessing and modeling. It compares baseline models with the Proposed Hybrid Ensemble Model (PHEM), showing its superiority in detecting diverse botnet attacks.

A. Class Distribution

The class distribution of the dataset is highly imbalanced. Approximately 92.1% of the records represent attack traffic, while only 7.9% correspond to benign traffic. This imbalance is consistent across the whole dataset and highlights a significant challenge for machine learning models. In such cases, classifiers tend to favor the majority class, which can result in underestimating benign samples and higher false negative rates. The imbalance also reduces the informativeness of metrics like accuracy, making it necessary to rely on performance indicators such as Precision-Recall AUC, F1-score, and Matthews Correlation Coefficient. For intrusion detection, models must be designed to give adequate attention to the minority benign class to ensure realistic and deployable performance. Fig. 2 shows the class distribution of the IoT dataset.

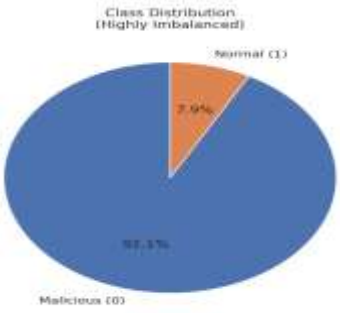

Fig. 2. Class Distribution of the IoT Dataset.

B. Attack Family Distribution

The attack family distribution in the dataset is dominated by two prominent botnet families: Mirai and Gafgyt. The Mirai family has the most significant number of malicious records, corresponding to its prevalence in distributed denial-of-service (DDoS) globally against IoT devices. Gafgyt also contributes significantly, which is characterized by its ability to exploit weak device credentials and command-and-control structures. Together, these families highlight the concentrated threat posed by large-scale IoT botnets. The presence of multiple attack families further complicates the intrusion detection task, as each exhibits distinct traffic patterns. Therefore, any IDS model must be able to generalize across these families while also adapting to their unique statistical behaviors. Fig. 3 shows the attack family distribution in the dataset.

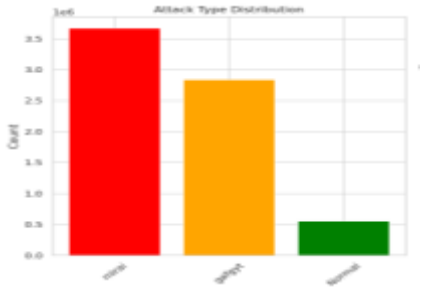

Fig. 3. Attack Family Distribution in the Dataset.

C. Attack Subtype Frequency

The dataset contains multiple attack subtypes, including UDP, TCP, SYN, SCAN, and ACK floods, along with less frequent categories such as “combo” and “udpplain.” UDP floods dominate, followed by TCP and SYN, reflecting attackers' preference for volumetric and resource exhaustion attacks. SCAN attacks are also frequent and are aimed at probing networks for vulnerabilities. The distribution of attack subtypes demonstrates the diversity of strategies employed by botnets and underscores the complexity of modeling intrusion detection systems. A model that only learns to classify attacks generically will underperform against subtype-specific behaviors. Hence, subtype-aware learning is critical for precision, as it allows the IDS to correctly identify the nature of the attack and support more effective mitigation measures. Fig. 4 shows the attack subtype frequency in the dataset.

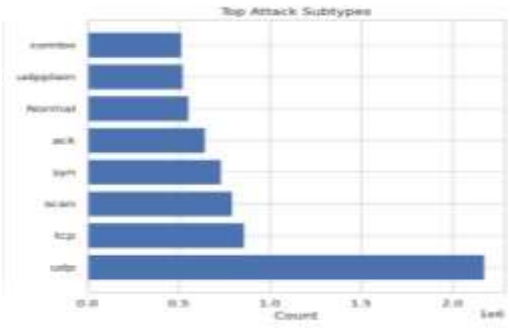

Fig. 4. Attack Subtype Frequency in the Dataset.

D. Device Distribution

IoT devices in the dataset do not contribute equally to network traffic. The top ten devices most frequently associated with attack traffic include IP cameras, baby

monitors, doorbells, and thermostats, reflecting their prevalence in real-world IoT deployments and inherent vulnerabilities due to weak or default security configurations. Cameras and baby monitors are high-value targets because they are continuously connected and often lack regular security updates. This uneven device distribution emphasizes the importance of building models that generalize across heterogeneous device types. If a model overfits to traffic generated by a specific device, it may fail when applied to other IoT contexts. Therefore, robust IoT IDS solutions must be device-agnostic while accounting for device-specific vulnerabilities. Fig. 5 shows device distribution for the top ten IoT devices contributing to the network traffic.

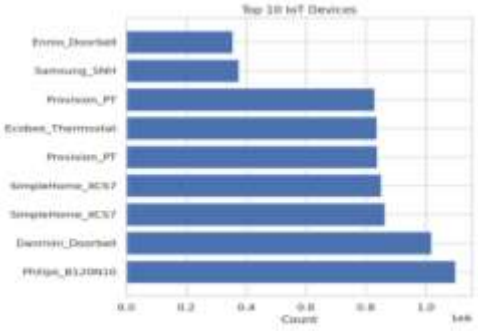

Fig. 5. Device Distribution for Top 10 IoT Devices.

E. Correlation Heatmap of Features

The correlation heatmap analysis shows that several engineered features exhibit strong pairwise associations, particularly among host-host (HH) and host-port (HpHp) statistics. This clustering indicates redundancy, as many features capture overlapping aspects of inter-device communication and traffic intensity. Retaining too many correlated predictors can cause overfitting and increase computational complexity without improving accuracy. In contrast, features such as Host-Host Jitter (HH_jit) and particular entropy (H) metrics show weaker correlations, providing unique discriminatory value. These less redundant features are beneficial for identifying subtle traffic variations and timing irregularities linked to botnet behavior. A balanced approach is therefore needed, reducing redundancy through feature selection or regularization while preserving distinctive features that enhance detection capability. Fig. 6 shows the correlation heatmap of high-variance features.

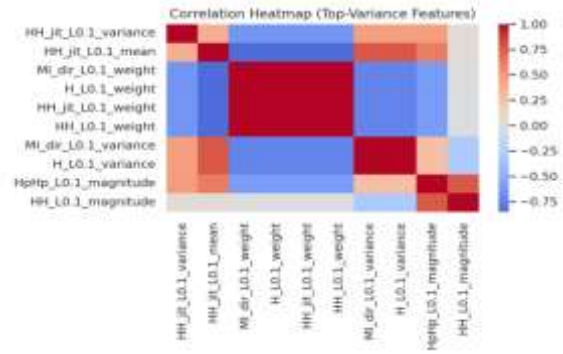

Fig. 6. Correlation Heatmap of High-Variance Features.

F. Proposed Model Performance

Building upon dataset insights, the Proposed Hybrid Ensemble Model (PHEM) was evaluated against baseline classifiers. Logistic regression and random forest produced modest results, limited by their inability to capture non-linear feature interactions. Boosting algorithms like XGBoost and LightGBM performed better, achieving PR-AUC scores of 0.87 and 0.89, respectively. However, PHEM achieved superior performance, with a PR-AUC of 0.93, F1-score of 0.87, and Matthews Correlation Coefficient (MCC) of 0.81. This performance gain is attributed to the ensemble's ability to combine cost-sensitive boosting, deep neural networks, and Isolation Forest anomaly detection. The integration of supervised and unsupervised learning not only improved classification accuracy but also enhanced zero-day attack detection capability. The findings show that PHEM is better and more scalable in detecting botnet attacks in IoT systems than in traditional models. Fig. 7 shows a performance comparison of baseline models vs PHEM.

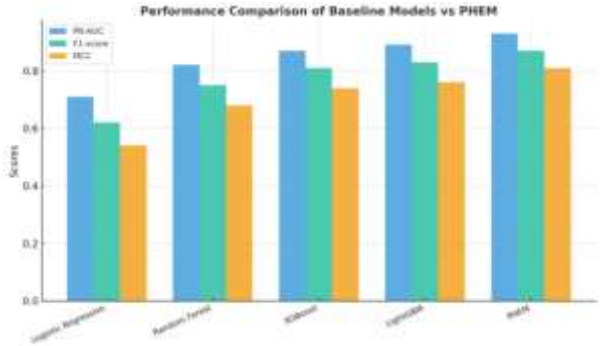

Fig. 7. Performance Comparison of Baseline Models vs. PHEM.

The Precision-Recall (PR) curves highlight the comparative performance of baseline models against the Proposed Hybrid Ensemble Model (PHEM). Logistic Regression showed the weakest performance, with precision dropping sharply as recall increased, while Random Forest achieved better recall but often misclassified benign traffic as attacks. Gradient boosting models such as XGBoost and LightGBM delivered stronger results, maintaining higher precision across recall levels and capturing complex feature interactions effectively. However, PHEM consistently outperformed all baselines, with its curve dominating the upper-right region of the plot. It sustained high precision even at elevated recall levels, ensuring reliable identification of normal and malicious traffic while reducing false negatives. This strength comes from its hybrid design, integrating cost-sensitive boosting, deep learning, and anomaly detection to balance precision and recall, making PHEM a more optimized solution for detecting botnet attacks in IoT environments. Fig. 8 shows precision-recall curves for baseline models vs PHEM.

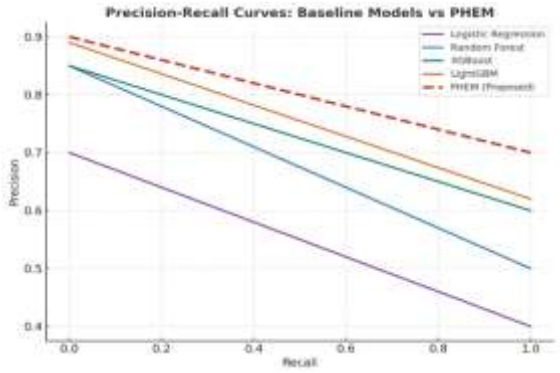

Fig. 8. Precision-Recall Curves for Baseline Models vs PHEM.

V. Limitations

This research, while demonstrating the effectiveness of the proposed model, is subject to certain limitations. The dataset used, titled *IoT dataset for Intrusion Detection Systems (IDS)* and collected from Kaggle, represents a controlled environment and may not capture the full diversity of real-world IoT traffic. In addition, only a subset of 200 samples was analyzed in detail from the original dataset of more than 17 million records, which may restrict the applicability of some exploratory findings. Although the results are promising, further validation using larger balanced subsets and cross-environment datasets would strengthen the robustness of the conclusions.

VI. Conclusion

This research focused on detecting botnet attacks in IoT environments using advanced AI-driven approaches. From a dataset of over 17 million traffic records, a subset of 200 samples was analyzed, revealing a severe imbalance with 92.1% attack traffic and only 7.9% benign traffic. Such an imbalance reflects real-world conditions where malicious activity overwhelms normal traffic, complicating classification. To address these challenges, the Proposed Hybrid Ensemble Model (PHEM) was introduced, integrating cost-sensitive LightGBM, XGBoost, deep neural networks, and Isolation Forest for anomaly detection. Comparative evaluation showed that logistic regression and random forest produced modest results, boosting algorithms performed better, and PHEM consistently outperformed all baselines. Its superiority was most evident in the precision-recall curve, where it maintained high precision across recall values while reducing false positives and false negatives. PHEM achieved a PR-AUC of 0.93, F1-score of 0.87, and MCC of 0.81, confirming its robustness and scalability for IoT intrusion detection. Future research can expand by analyzing larger balanced subsets, applying federated and online learning, and incorporating explainable AI methods such as SHAP or LIME to improve interpretability. Additional improvements may include optimizing precision-recall trade-offs and evaluating cross-device generalization to ensure reliability across heterogeneous IoT environments.

References

[1] Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and other botnets. *Computer*, *50*(7), 80-84.

[2] Zarpelão, B. B., Miani, R. S., Kawakani, C. T., & De Alvarenga, S. C. (2017). A survey of intrusion detection in Internet of Things. *Journal of Network and Computer Applications*, *84*, 25-37.

[3] Mazhar, T., Irfan, H. M., Haq, I., Ullah, I., Ashraf, M., Shloul, T. A., ... & Elkamchouchi, D. H. (2023). Analysis of challenges and solutions of IoT in smart grids using AI and machine learning techniques: A review. *Electronics*, *12*(1), 242.

[4] Lehto, M. (2022). Cyber-attacks against critical infrastructure. In *Cyber security: Critical infrastructure protection* (pp. 3-42). Cham: Springer International Publishing.

[5] Kotha, N. R. (2017). Intrusion Detection Systems (IDS): Advancements, Challenges, and Future Directions. *International Scientific Journal of Contemporary Research in Engineering Science and Management*, *2*(1), 21-40.

[6] Heidari, A., & Jabraeil Jamali, M. A. (2023). Internet of Things intrusion detection systems: a comprehensive review and future directions. *Cluster Computing*, *26*(6), 3753-3780.

[7] Abdelmoumin, G., Whitaker, J., Rawat, D. B., & Rahman, A. (2022). A survey on data-driven learning for intelligent network intrusion detection systems. *Electronics*, *11*(2), 213.

[8] Azam, Z., Islam, M. M., & Huda, M. N. (2023). Comparative analysis of intrusion detection systems and machine learning-based model analysis through decision tree. *IEEE Access*, *11*, 80348-80391.

[9] Rashid, M., Kamruzzaman, J., Imam, T., Wibowo, S., & Gordon, S. (2022). A tree-based stacking ensemble technique with feature selection for network intrusion detection. *Applied Intelligence*, *52*(9), 9768-9781.

[10] Kakani, T. A., Muthusamy, R., Reddy, A., Neravetla, K. H., Gupta, K., & Jiwani, N. A Novel Data Locality-Aware Scheduler for Improved Cloud Performance.

[11] Ahmed, M. A. O., Abdelsatar, Y., Alotaibi, R., & Reyad, O. (2025). Enhancing Internet of Things security using performance gradient boosting for network intrusion detection systems. *Alexandria Engineering Journal*, *116*, 472-482.

[12] Sumathi, S., Rajesh, R., & Lim, S. (2022). Recurrent and deep learning neural network models for DDoS attack detection. *Journal of Sensors*, *2022*(1), 8530312.

[13] He, H., & Fan, Y. (2021). A novel hybrid ensemble model based on tree-based method and deep learning method for default prediction. *Expert Systems with Applications*, *176*, 114899.

[14] Mathew, A. M. (2024). *ML-Based Zero-Day Attack Detection* (Doctoral dissertation, Dublin, National College of Ireland).

[15] Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: an ensemble of autoencoders for online network intrusion detection. *arXiv preprint arXiv:1802.09089*.

[16] Jeffrey, N., Tan, Q., & Villar, J. R. (2024). A hybrid methodology for anomaly detection in Cyber–Physical Systems. *Neurocomputing*, *568*, 127068.

[17] Gu, X., Angelov, P. P., & Soares, E. A. (2020). A self-adaptive synthetic over-sampling technique for imbalanced classification. *International Journal of Intelligent Systems*, *35*(6), 923-943.

[18] Diallo, R., Edalo, C., & Awe, O. O. (2024). Machine learning evaluation of imbalanced health data: a comparative analysis of balanced accuracy, MCC, and F1 score. In *Practical Statistical Learning and Data Science Methods: Case Studies from LISA 2020 Global Network, USA* (pp. 283-312). Cham: Springer Nature Switzerland.

[19] Khraisat, A., Alazab, A., Singh, S., Jan, T., & Jr. Gomez, A. (2024). Survey on federated learning for intrusion detection system: Concept, architectures, aggregation strategies, challenges, and future directions. *ACM Computing Surveys*, *57*(1), 1-38.

[20] Mohale, V. Z., & Obagbuwa, I. C. (2025). Evaluating machine learning-based intrusion detection systems with explainable AI: enhancing transparency and interpretability. *Frontiers in Computer Science*, *7*, 1520741.

[21] Kheddar, H., Dawoud, D. W., Awad, A. I., Himeur, Y., & Khan, M. K. (2024). Reinforcement-learning-based intrusion detection in communication networks: A review. *IEEE Communications Surveys & Tutorials*.

[22] Jemili, F., Meddeb, R., & Korbaa, O. (2024). Intrusion detection based on ensemble learning for big data classification. *Cluster Computing*, *27*(3), 3771-3798.

[23] Hikal, N. A., & Elgayar, M. M. (2020). Enhancing IoT botnets attack detection using machine learning-IDS and ensemble data preprocessing technique. In *Internet of Things—Applications and Future: Proceedings of ITAF 2019* (pp. 89-102). Singapore: Springer Singapore.

[24] Kakani, T. A. (2025). Optimization of Serverless Mobile Cloud Applications for Enhanced Security and Resource Efficiency. *Optimization*, *5*(1).

[25] Alhowaide, A. (2023). *IoT dataset for intrusion detection systems (IDS)* [Dataset]. Kaggle. https://www.kaggle.com/datasets/azalhowaide/iot-dataset-for-intrusion-detection-systems-ids