Intelligent Detection and Analysis of Sophisticated Cyber Threats Using Real-World Network Traffic

Juli Dutta, Jeffrey Berg, Qingzhong Liu
Department of Computer Science
Sam Houston State University, Huntsville, TX, USA
Email: {jxd148, jdb118, qxl005}@shsu.edu

***Abstract*—This paper presents an intelligent machine-learning framework for detecting cyberattacks using real-world network traffic from the CIC-IDS2017 dataset. In addition to classical flow-based modeling with Random Forest, we incorporate two modern baselines—XGBoost and a one-dimensional Convolutional Neural Network (1D-CNN)—to evaluate the effectiveness of ensemble learning and deep learning approaches for intrusion detection. Our comparative analysis demonstrates that while the Random Forest classifier achieves the highest overall performance (macro-F1 ≈ 0.98), XGBoost performs competitively (macro-F1 ≈ 0.95), confirming the strength of tree-ensemble methods on tabular flow-feature data. In contrast, the 1D-CNN model shows substantially lower performance (macro-F1 ≈ 0.55), highlighting the limitations of deep learning when applied to non-sequential statistical features. Cross-day evaluation further reveals that mixed-attack training substantially improves generalization across diverse threat types, including DDoS, infiltration, and brute-force attacks. These findings provide practical insights into the selection of effective AI-based IDS architectures and contribute to a deeper understanding of how different model families perform on real-world flow-based intrusion detection tasks.**

***Index Terms*—Intrusion Detection, Random Forest, XGBoost, Deep Learning, Network Forensics**

I. Introduction

The proliferation of interconnected devices and online services has led to an unprecedented increase in both the volume and complexity of cyber threats [1], [2]. Modern cyber attacks—ranging from advanced persistent threats (APT) and botnets to command-and-control (C2) infrastructure and lateral movement—often bypass traditional signature-based detection systems [3]. As a result, the cybersecurity community is moving towards more intelligent, data-driven approaches to threat detection, leveraging real-world network traffic to build robust and adaptable defense mechanisms [1], [4].

One effective method involves the analysis of packet capture (PCAP) datasets from real network environments, which provide a detailed record of communication events and behaviors [5]. Publicly available datasets such as CTU-13 [6], CIC-IDS2017 [7], and MAWI Working Group traces serve as critical benchmarks for modeling and evaluating detection systems under realistic threat scenarios [8].

In this paper, we focus on the detection and analysis of sophisticated cyber threats using the CIC-IDS2017 dataset, which includes diverse attack types such as botnet traffic, web intrusions, and denial-of-service (DoS) attacks [7]. We employ intelligent detection techniques—including flow-based feature extraction, supervised machine learning models, and behavioral analysis—to identify malicious activities. Additionally, we propose a defense strategy based on the insights derived from threat attribution and behavioral profiling.

This work provides a comparative evaluation of classical machine-learning models and modern AI-based intrusion detection approaches. In addition to Random Forest, we incorporate XGBoost and a 1D-CNN baseline to assess model performance on flow-based network traffic. Our results show that gradient-boosting methods (XGBoost) achieve accuracy comparable to Random Forest, while deep learning models underperform on tabular flow-feature data. These findings offer practical insights for selecting effective AI-based IDS architectures and highlight the limitations of deep-learning approaches when applied to non-sequential statistical flow representations.

II. Literature Review

The evolution of cyber threats has necessitated the development of intelligent detection systems that go beyond traditional intrusion detection systems (IDS). These systems incorporate machine learning (ML), deep learning (DL), and advanced data analytics to enhance the detection of complex and evasive threats in network traffic.

1) Machine Learning for Intrusion Detection: Several researchers have applied ML techniques for network-based intrusion detection. Buczak and Guven [1] provided a comprehensive review of data mining and ML techniques used for cybersecurity, highlighting the promise of supervised and unsupervised models for anomaly detection. Ring et al. [8] examined the effectiveness of flow-based ML algorithms in detecting network intrusions using labeled datasets.

2) Use of Public Datasets: Real-world datasets like CIC-IDS2017 and CTU-13 have become common in validating intrusion detection frameworks. García et al. [6] introduced the CTU-13 dataset as a collection of botnet scenarios that include normal, background, and malicious traffic. Earlier datasets such as NSL-KDD [9] have also been extensively used, although Tavallaee et al. [10] highlighted limitations of KDD CUP'99 and motivated newer datasets like CIC-IDS2017 that better reflect real-world traffic complexities.

3) Traffic Characterization Approaches: Recent studies emphasize characterizing network traffic using temporal and

statistical flow features rather than relying solely on payload inspection. Lashkari et al. [11] explored time-based features to differentiate Tor anonymized traffic from regular web traffic, showing that metadata such as flow duration, inter-arrival times, and packet size distributions can reveal critical behavioral patterns even when payloads are encrypted. Such characterization is essential with the growing prevalence of encrypted communications that hinder traditional deep-packet inspection.

4) Deep Learning Approaches: Deep learning models have shown strong performance in some IDS applications, but classical ensemble techniques like Random Forest remain attractive for their interpretability and speed. Kim et al. [12] explored LSTM and CNN models for real-time detection of DDoS attacks and malware communication. Wang et al. [13] showed that 1D CNNs can classify encrypted network traffic without payload inspection. Shone et al. [14] demonstrated that deep autoencoders can significantly improve IDS performance by learning traffic representations.

5) Taxonomy of IDS Techniques: Hindy et al. [2] proposed a comprehensive taxonomy of IDS designs, categorizing systems by data source (host-based vs. network-based), detection methodology (signature-based, anomaly-based, hybrid), and deployment architecture (centralized vs. distributed). Their survey highlights challenges such as handling encrypted traffic, minimizing false positives, and achieving real-time detection.

6) Threat Attribution and Defense Modeling: The integration of threat attribution techniques with detection systems is gaining attention. MITRE's ATT&CK framework [15] has been used to align detected behaviors with known adversary tactics. Tools such as Zeek and Suricata facilitate extraction of enriched metadata from PCAP files to support contextual analysis and incident response.

These works collectively demonstrate the feasibility of using intelligent systems to detect advanced cyber threats. However, many existing studies either focus on a narrow threat type or fail to integrate real-time feature engineering, threat attribution, and defense proposal into a cohesive system. Our research addresses this gap by combining intelligent detection with end-to-end analysis and defense planning based on real-world traffic.

III. Methodology

For this paper, we utilized the CIC-IDS2017 dataset [7], a widely used benchmark for evaluating intrusion detection systems. The dataset comprises five days of captured network traffic (Monday through Friday), each representing different types of attack scenarios. Rather than working directly with PCAP files, we downloaded the corresponding CSV files made available on the official CIC website, which include pre-extracted network flow features using CICFlowMeter. These CSVs served as the foundation for both training and evaluating our machine learning models. Figure 2 illustrates the overall system pipeline, beginning with data cleaning and filtering, followed by feature selection, model training, and evaluation.

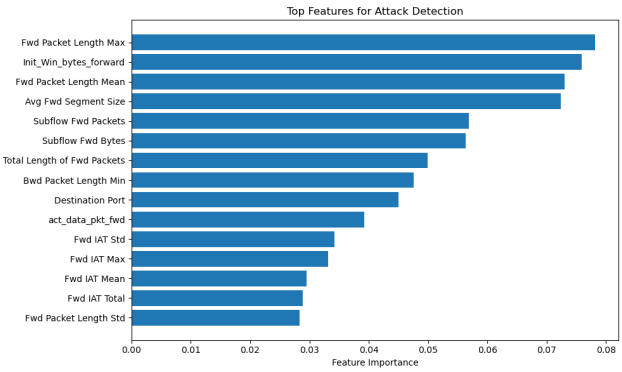

Fig. 1. Confusion matrix showing classification results of the model on DDoS data.

A. DDoS Attack Detection

We initially focused on the Friday DDoS dataset, which contains labeled instances of both benign and Distributed Denial-of-Service (DDoS) traffic. Each entry includes 78 network flow features, such as packet size statistics, TCP flag counts, and inter-arrival times. Using these, we trained a Random Forest classifier to differentiate between benign and malicious traffic. Feature importance analysis was performed after training, allowing us to identify the top 15 most influential features contributing to detection accuracy. These features included metrics such as flow duration, total forward packets, and average length of forward packets.

To evaluate generalization beyond DDoS detection, we extended the training data by incorporating CSV files from Wednesday and Thursday, which contain web-based attacks, infiltration attempts, and other network intrusions. This enabled us to build a more robust model capable of identifying a wider variety of attack types under a binary classification scheme (benign vs. attack). The *Label* column in each CSV was used as ground truth. We mapped `BENIGN` labels to 0 and all other attack types to 1, framing the problem as a supervised binary classification task.

Model performance was evaluated using confusion matrices and classification reports, with a focus on precision, recall, and F1-score to assess detection capability.

B. Multi-Attack Generalization

When tested against a broader set of attack types, the DDoS-only model exhibited high precision but poor recall for non-DDoS attacks, highlighting its specialization to flooding patterns. We then created a mixed model that incorporates attacks from multiple days (including web-based and infiltration attacks) to ensure coverage across diverse threat types. After retraining on the combined datasets, the mixed model achieved improved generalization, correctly classifying a wider variety of malicious traffic while maintaining strong performance on DDoS flows.

C. Cross-Day Evaluation

After creating models for DDoS-only and mixed attacks, we tested these models on other days such as Monday and Tuesday. When evaluated on the Monday CSV file, which

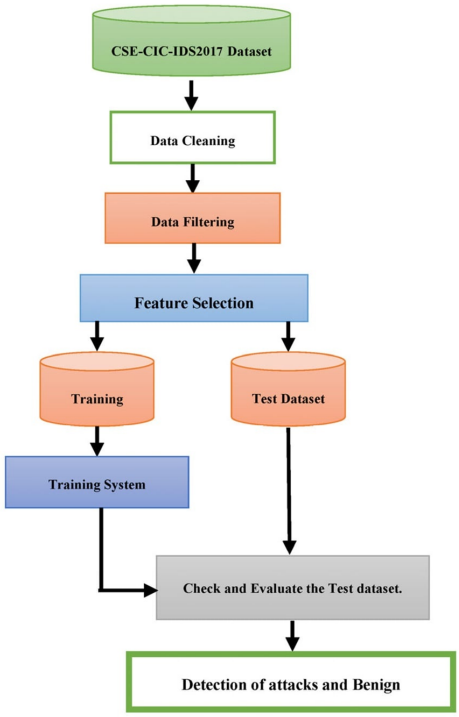

Fig. 2. Workflow of the proposed detection pipeline.

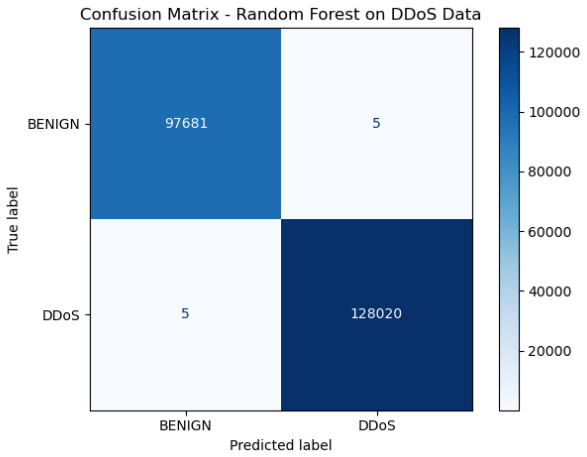

Fig. 3. Confusion matrix showing classification results of the mixed model.

contains only benign traffic, the model did not detect any attacks, consistent with the dataset's ground truth. For the Friday DDoS subset, the confusion matrix in Fig. 1 shows near-perfect separation between benign and DDoS flows. For Tuesday, we ran the mixed model and it was able to predict which attacks were present in the CSV file along with benign packets (Table I).

Tables II and III compare the performance of the Random Forest (RF) classifier and a Neural Network (NN) under imbalanced and balanced training conditions.

D. Evaluation Metrics

To evaluate classification models for cyber threat detection, we used several metrics derived from confusion matrices.

TABLE I
MULTICLASS CLASSIFICATION REPORT SHOWING PRECISION, RECALL, F1-SCORE, AND SUPPORT PER CLASS FOR TUESDAY.

Class	Precision	Recall	F1-Score	Support
BENIGN	1.00	1.00	1.00	431,813
DDoS	0.00	0.00	0.00	0
DoS GoldenEye	0.00	0.00	0.00	0
DoS Hulk	0.00	0.00	0.00	0
DoS Slowhttptest	0.00	0.00	0.00	0
DoS Slowloris	0.00	0.00	0.00	0
FTP-Patator	1.00	1.00	1.00	7,935
Heartbleed	0.00	0.00	0.00	0
SSH-Patator	1.00	1.00	1.00	5,897
Web Attack Brute Force	0.00	0.00	0.00	0
Web Attack SQL Injection	0.00	0.00	0.00	0
Web Attack XSS	0.00	0.00	0.00	0
micro avg	1.00	1.00	1.00	445,645
macro avg	0.25	0.25	0.25	445,645
weighted avg	1.00	1.00	1.00	445,645

TABLE II
CLASSIFIER PERFORMANCE COMPARISON (WEIGHTED METRICS).

Classifier	Accuracy	F1-Score	Recall
Random Forest	96.4%	96.2%	96.7%
Neural Network (Imbalanced)	74.0%	63.0%	74.0%

Accuracy,

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \tag{1}$$

provides a general measure of correctness but can be misleading in imbalanced datasets. Precision and recall are defined as

$$\text{Precision} = \frac{TP}{TP + FP}, \tag{2}$$

$$\text{Recall} = \frac{TP}{TP + FN}, \tag{3}$$

where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively. The F1-score,

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}, \tag{4}$$

offers a balanced assessment of precision and recall. Macro-averaged F1 treats all classes equally,

$$F1_{\text{macro}} = \frac{1}{N} \sum_{i=1}^{N} F1_i, \tag{5}$$

while weighted F1 reflects class distribution,

$$F1_{\text{weighted}} = \sum_{i=1}^{N} w_i F1_i, \quad w_i = \frac{n_i}{\sum_j n_j}. \tag{6}$$

For the neural network classifier, training was guided by minimizing the cross-entropy loss,

$$L = -\frac{1}{N} \sum_{i=1}^{N} \sum_{c=1}^{C} y_{i,c} \log(\hat{y}_{i,c}), \tag{7}$$

and feature importance for tree-based models was quantified using information gain,

$$IG(Y, X) = H(Y) - H(Y \mid X). \tag{8}$$

These metrics collectively provide a robust framework for assessing both overall and class-specific effectiveness of the detection models.

TABLE III
CLASSIFIER COMPARISON: RF VS NN (BALANCED).

Metric	RF	NN
Accuracy	0.9996	0.9974
F1 (Macro Avg)	0.9975	0.9806
F1 (Weighted)	0.9996	0.9974
F1 – BENIGN	0.9989	0.9935
F1 – BRUTE FORCE	0.9998	0.9866
F1 – DDoS	0.9999	0.9997

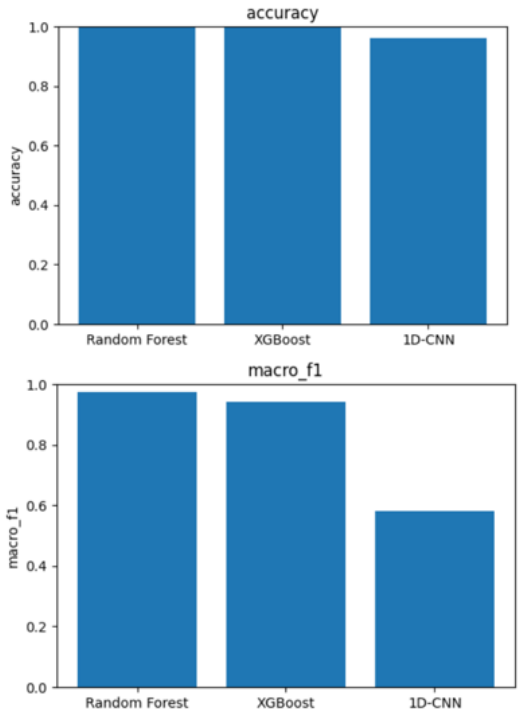

Fig. 4. Overall accuracy comparison of RF, XGBoost, and 1D-CNN.

E. Feature Analysis and Baseline Models

Feature analysis was performed to identify which network flow attributes contributed most to the Random Forest model's performance on the DDoS dataset. The most important features included flow duration, total forward/backward packets, total length of forward/backward packets, flow bytes per second, and TCP flag counts (e.g., SYN, ACK). DDoS traffic typically exhibits extremely high packet and byte rates, short-lived unidirectional flows, and abnormal flag distributions compared to benign traffic. Using only the top 10 features yielded nearly identical performance to the full feature set, indicating that a small subset captures most predictive power.

XGBoost: XGBoost [16] is a state-of-the-art gradient boosting framework widely used for tabular data. We trained XGBoost with grid-searched hyperparameters (learning rate, maximum tree depth, subsampling ratio) and evaluated it on the same train/test split as Random Forest, providing a strong contemporary baseline.

1D-CNN: The 1D-CNN architecture follows established deep-learning approaches for network intrusion detection [17]. We implemented a lightweight 1D-CNN with two convolutional layers, one max-pooling layer, and a fully connected layer. Flow-feature vectors were standardized and reshaped into 1D sequences; training used the Adam optimizer (learning rate = 0.001), batch size 64, and a fixed number of epochs. All models were trained on identical splits to ensure fair comparison.

TABLE IV
PERFORMANCE COMPARISON – DDOS-ONLY MODEL VS MIXED ATTACK MODEL.

Model	Precision	Recall	F1-Score
DDoS-Only Model	1.00	1.00	1.00
Mixed Attack Model	1.00	0.84	0.91

IV. RESULTS

Our experimental results demonstrate the effectiveness of the Random Forest-based detection approach on real-world network traffic from CIC-IDS2017. Initially, the classifier trained exclusively on Friday DDoS traffic achieved near-perfect performance, with precision, recall, and F1-scores approaching 1.0 for DDoS flows. As shown in Fig. 1, the confusion matrix indicates virtually no false negatives or false positives when tested against DDoS-specific flows.

Table IV compares the detection performance of the specialized DDoS model and the generalized mixed attack model. The DDoS-only model achieved perfect classification on specialized traffic (precision, recall, and F1-score of 1.00), but lacked versatility when applied to heterogeneous attack types. In contrast, the mixed attack model maintained high precision (1.00) while significantly improving recall (0.84) and F1-score (0.91) across diverse threats, demonstrating better generalization.

When evaluated against datasets containing diverse attacks such as infiltration, web-based intrusions, and brute force attempts, the DDoS-trained model exhibited significant performance degradation (recall dropping to 0.34), highlighting the limitations of using a model trained on a single attack type for broader detection tasks. The mixed model, however, demonstrated improved versatility across attack classes. Cross-day evaluation on the Monday dataset, which contained only benign traffic, yielded no false alarms, indicating a low false-positive rate.

A. Comparative Evaluation of RF, XGBoost, and 1D-CNN

Fig. 4 shows that Random Forest achieved a macro-F1 of approximately 0.98, while XGBoost reached approximately 0.95, confirming that tree-based methods remain the strongest performers for flow-based IDS tasks. The 1D-CNN model attained a macro-F1 around 0.55, indicating that deep learning models are less effective when limited to statistical flow features rather than temporal packet sequences.

The per-class F1 comparison in Fig. 5 reveals that both Random Forest and XGBoost consistently achieve F1-scores above 0.90 across most attack classes, including FTP-Patator, SSH-Patator, and DoS variants. CNN performance varies significantly per class, excelling only in a subset of attacks with strong feature-level patterns. Overall, tree-ensemble models provide stronger generalization than convolutional architectures for real-world flow-based IDS.

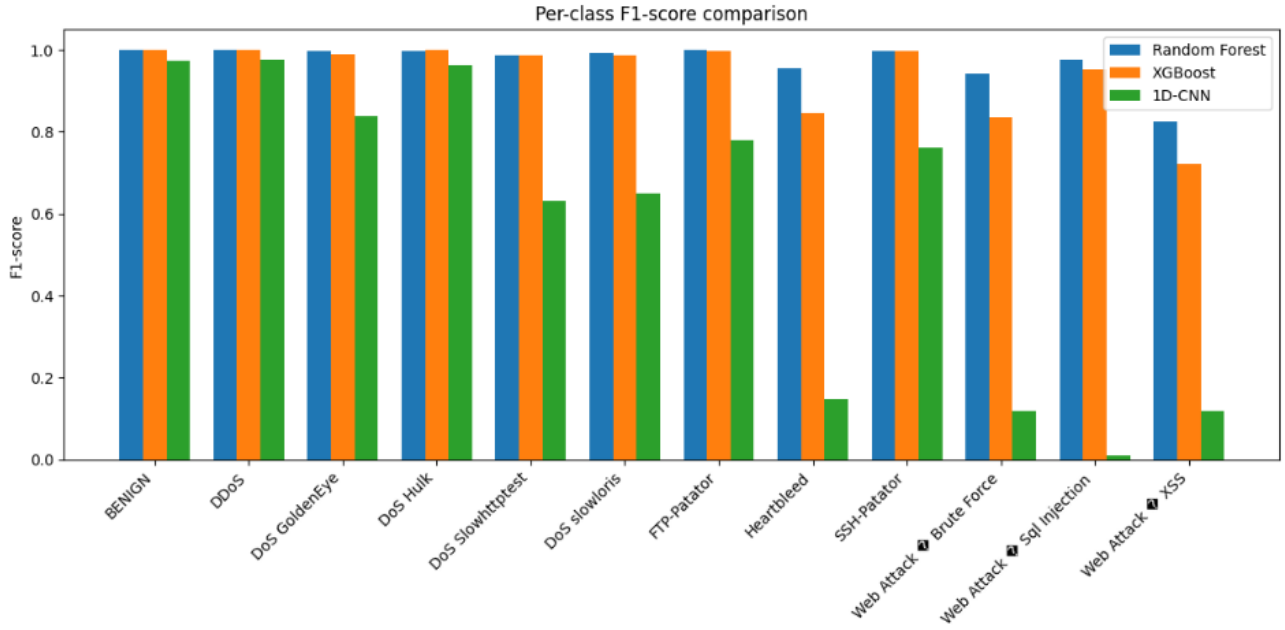

Fig. 5. Per-class F1-score comparison across all attacks.

A key finding of this study is that deep learning does not automatically outperform classical ML for IDS tasks based on flow-level features. RF and XGBoost outperform the 1D-CNN by a wide margin, underscoring the importance of choosing architectures aligned with feature structure.

V. Conclusion and Future Work

We proposed an intelligent detection framework based on flow-level feature extraction and machine-learning modeling for the detection of sophisticated cyber threats in real-world network traffic. Using the CIC-IDS2017 dataset, we showed that a Random Forest classifier trained on a mixed set of attack scenarios achieves high detection performance across multiple threat types, including DDoS, infiltration, and web-based intrusions.

Our experiments highlight that while specialized models excel at detecting targeted attack patterns, they suffer when generalized to diverse threat landscapes. Incorporating multiple attack scenarios during training significantly enhances the model's ability to detect heterogeneous threats without compromising precision. The inclusion of modern baselines such as XGBoost and 1D-CNN further validates that tree-ensemble models remain the most effective for flow-based IDS.

Future work will explore more advanced deep learning approaches, such as recurrent neural networks (RNNs) and CNNs operating on raw packet or flow sequences, to capture temporal and payload patterns more effectively. We also plan to investigate online learning techniques for continuous adaptation to evolving threats and to extend the framework to semi-supervised settings that leverage unlabeled network traffic for model enhancement, ultimately advancing toward real-time and adaptive intrusion detection systems.

References

[1] J. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 2, pp. 1153–1176, 2016.

[2] H. Hindy *et al.*, "A taxonomy and survey of intrusion detection system design techniques," *Comput. Commun.*, vol. 161, pp. 1–20, 2020.

[3] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE S&P*, 2010, pp. 305–316.

[4] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.

[5] P. García-Teodoro, J. Díaz-Verdejo, G. Maciá-Fernández, and E. Vázquez, "Anomaly-based network intrusion detection: Techniques, systems and challenges," *Comput. Secur.*, vol. 28, no. 1–2, pp. 18–28, 2009.

[6] S. García, M. Grill, J. Stiborek, and A. Zunino, "An empirical comparison of botnet detection methods," *Comput. Secur.*, vol. 45, pp. 100–123, 2014.

[7] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in *Proc. ICISSP*, 2018.

[8] M. Ring *et al.*, "A survey of network-based intrusion detection data sets," *Comput. Secur.*, vol. 86, pp. 147–167, 2019.

[9] L. Dhanabal and S. Shantharajah, "A study on NSL-KDD dataset for intrusion detection system based on classification algorithms," *IJCSIT*, vol. 7, no. 6, pp. 231–238, 2015.

[10] M. Tavallaee *et al.*, "A detailed analysis of the KDD CUP 99 dataset," in *Proc. IEEE CISDA*, 2009.

[11] A. Lashkari *et al.*, "Characterization of Tor traffic using time-based features," in *Proc. ICISSP*, 2017.

[12] S. Kim, Y. Kim, and Y. Kim, "Efficient detection of DDoS attacks using deep learning models," *Electronics*, vol. 9, no. 6, p. 916, 2020.

[13] W. Wang *et al.*, "End-to-end encrypted traffic classification with 1D CNNs," in *Proc. IEEE ICNP*, 2017.

[14] T. Shone *et al.*, "A deep learning approach to network intrusion detection," *IEEE TETCI*, vol. 2, no. 1, pp. 41–50, 2018.

[15] MITRE ATT&CK Framework. [Online]. Available: https://attack.mitre.org

[16] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. ACM SIGKDD*, 2016, pp. 785–794.

[17] H. Kim *et al.*, "LSTM-CNN deep learning model for network intrusion detection," *IEEE Access*, vol. 8, pp. 202502–202511, 2020.