

Implementing AI-Based Real-Time Monitoring Systems to Combat DDoS Attacks

Surya Lokesh Bhargav Pentakota^{1,*}

¹Department of Research and Development, Ginger Labs, Texas, United States of America.
suryalokeshbhargav@gmail.com¹

Abstract: One of the most dangerous types of attacks is known as a distributed denial of service attack (DDoS). In most cases, traditional defensive techniques are unable to identify or prevent innovative, high-speed, or advanced attacks. A real-time artificial intelligence-based monitoring system is proposed in this study for the early detection and prevention of distributed denial-of-service attacks. Identifying attacks is accomplished through the utilisation of a hybrid approach that combines a random forest classifier with a deep neural network-based feature extraction. This system has been trained and tested on the CIC-DDoS2019 dataset, a comprehensive collection of modern distributed denial-of-service attacks on traffic in the modern world. During the system's implementation, network traffic is monitored in real-time, and significant features are extracted to differentiate between malicious and normal packets. The test results demonstrate that the system has a high level of accuracy, a low rate of false positives, and a fast response rate, all of which indicate its effectiveness in real-world settings. Matplotlib is used for data visualisation, whereas Python, TensorFlow, Scikit-learn, and Pandas are research tools that serve as libraries for data handling and model construction.

Keywords: Artificial Intelligence; Machine Learning; Real-Time Monitoring; Network Security; Distributed Denial of Service (DDoS); Random Forest Classifier; DDoS Attacks; Cybersecurity Frameworks.

Received on: 02/11/2024, **Revised on:** 07/01/2025, **Accepted on:** 15/02/2025, **Published on:** 05/09/2025

Journal Homepage: <https://www.fmdbpublish.com/user/journals/details/FTSCL>

DOI: <https://doi.org/10.69888/FTSCL.2025.000427>

Cite as: S. L. B. Pentakota, "Implementing AI-Based Real-Time Monitoring Systems to Combat DDoS Attacks," *FMDB Transactions on Sustainable Computer Letters*, vol. 3, no. 3, pp. 116–125, 2025.

Copyright © 2025 S. L. B. Pentakota, licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under CC BY-NC-SA 4.0, which allows unlimited use, distribution, and reproduction in any medium with proper attribution.

1. Introduction

The increased utilisation of Internet of Things devices and internet services has made cybersecurity a top concern for organisations worldwide. Among many cyberattacks, Distributed Denial of Service (DDoS) attacks are the most pervasive and challenging to overcome. In the opinion of Al-Heety et al. [11], the intrinsic realities of DDoS attacks are decentralization and the scope of attacks, which can easily cause internet services to go offline by overloading them with traffic from compromised systems. Behind such, which are largely IoT nodes and botnets, can bring a vast number of synchronised requests crafted to drain server resources. The DDoS attack has evolved into a sophisticated, distributed, and prolonged attack. A review by Mekki et al. [14] revealed that past DDoS attacks utilised primitive flooding attacks, whereas modern attacks target both the network and application layers. Defences usually collapse due to over-reliance on static signatures or static rules, a shortcoming that Arif et al. [10] addressed in their research exploring SDN-based security solutions. The dynamic signature of attacks in contemporary times makes detection and prevention very hard for traditional intrusion detection systems. One of the most

*Corresponding author.

prominent experiments on real attacks by Michelena et al. [2] demonstrates the helplessness of traditional firewalls and traditional IDSs against stealthy, polymorphic attack patterns. Those traditional precautions will primarily generate high false positives and even block legitimate. According to research by Gadze et al. [7], zero-day attacks are also exploited in real time, further exacerbating the issue of threat response latency in static systems.

These limitations stem from the use of intelligent, adaptive, and learning-based solutions. The latest efforts to overcome them have been met with an open mind toward artificial intelligence. The latest machine learning technology, as used by Liu et al. [3], has proven highly promising for classifying traffic efficiently and, most importantly, predicting likely assault patterns from both historical and real-time data. AI model detection efficiencies are higher than those of traditional devices when measured by metrics such as detection rate and response time. This paradigm shift has led to the development of real-time monitoring systems that utilize decision trees and neural networks, enabling them to learn dynamically and respond to emerging, unknown threats. Real-time monitoring was employed in a hybrid machine learning model, as demonstrated in a paper by Setitra and Fan [8], to facilitate improved early detection. The models were passing network traffic in real-time to identify anomalies with minimal human intervention. Real-time analytics here is the hub that stops threats before service disruption can even be initiated. Another report by Michelena et al. [2] also corroborates the trend, finding that predictive modelling is highly useful for preemptive responses. These intelligent systems require robust datasets to survive. The efficiency of real and labelled data was demonstrated by Correia et al. [12], and the need for diverse attack scenarios was established. This was also demonstrated by the large CICIDS2017 dataset presented by Sharafaldin et al. [6], which represented a collection of modern DDoS attack vectors in real-world traffic scenarios.

Such datasets are useful for training machine learning models to identify treacherous attack patterns amid regular traffic. Another publication by Shafiq et al. [5] also introduced the use of feature engineering in machine learning-based detection methods. They identified packet length, inter-arrival time, and flow duration as key network parameters to utilize as features when training robust classifiers. Dynamic systems incorporating feature extraction and behavior modeling, as proposed by Di Maio et al. [1], also demonstrate real-world performance improvement over traditional static signature-based approaches. Hybrid deep learning architectures, as used by Khaleel and Shiltagh [13], outperformed single models by combining the strengths of multiple algorithms. Hybrid models typically combine CNNs and LSTMs to address spatial and temporal relationships in traffic patterns. Another study by Alduailij et al. [9] emphasised system scalability and low-latency processing to make such AI models deployable on high-speed networks. This paper discusses the development, deployment, and assessment of an AI-based real-time DDoS defence monitoring system. The system utilizes procedures as good as those in existing research and blends hybrid models trained on labeled datasets, such as CICIDS2017. This paper builds on the work of Sharafaldin et al. [6] and others to provide a comprehensive deployment of a low-latency, efficient, and scalable DDoS detection system. It aims to advance AI-based cybersecurity research by providing a testable, verifiable approach to dominant threats.

2. Literature Review

Al-Heety et al. [11] conducted a comprehensive survey and recognized the adaptive and dynamic characteristics of Distributed Denial of Service (DDoS) attacks, which necessitate the development of appropriate detection systems. The detection and countermeasures literature on DDoS attacks became more prominent as they had an astounding impact on internet networks. Early systems typically employed statistical and threshold-based methods to monitor packet rates, bandwidth usage, and connection counts. While easy to deploy, these systems consistently produced high false-positive and false-negative rates, indicating their inability to adapt to the dynamic nature of attack patterns. Mekki et al. [14] highlighted the vulnerability of traditional detection approaches, as a hacker can easily circumvent them by adding incremental volume to traffic. The evasion techniques achieved the intrinsic rigidity of static detection models. The limitations prompted researchers to develop adaptive models that could dynamically inspect traffic independently of predefined thresholds. It was this technological advancement that marked the end of rule-based systems and the beginning of data-driven approaches to enhance the detection of anomalous behaviour [15].

Arif et al. [10] suggested ML techniques as a probable remedy as a step in overcoming the inflexibility of threshold-based models. Naive Bayes, k-Nearest Neighbors (k-NN), and Support Vector Machines (SVM) were among the methods proposed for classifying traffic into malicious and benign categories based on past traffic behavior. Supervised machine learning models that were provided with higher detection performance as well as scalability. They also had hand-designed features and domain-knowledge dependencies that limited their ability to generalise across diverse network environments. Anyanwu et al. [4] also analysed the disadvantages of conventional ML model feature engineering, noting that it is time-consuming and difficult to generalise across different datasets. Feature input quality was a challenge for deploying ML models in network environments. High-dimensional data was also a bottleneck for the models, and their deployment required techniques such as dimensionality reduction or feature selection to prevent performance deterioration and overfitting [16].

Gadze et al. [7] identified the major impact that DL has made in DDoS detection, particularly in the context of designs capable of automatically learning from raw data. CNNs and RNNs, particularly LSTM networks, enabled models to learn both spatial and temporal patterns from raw data. These designs enabled the minimisation of human-specified features, increased generalisation, and improved the ability to identify new attack patterns, thereby mitigating the drawbacks of traditional ML techniques. Liu et al. [3] illustrated the application of CNNs through the presentation of traffic data in a visual form and the recognition of flows as images, aiming to tap into spatial relationships. Visual representation enabled the models to capture structural deviations that would go unnoticed with sequential or statistical models. The resilience and stability of CNNs in handling noisy data make them highly effective against real-world DDoS attacks [17]. All these advancements constituted a paradigm shift from representation-based to feature-based detection mechanisms.

Setitra and Fan [8] compared LSTM-based techniques with those that can detect time-varying patterns in traffic streams. This is useful for sensing particularly low-and-slow, over-time DDoS attacks, which are challenging to detect with snapshot-based techniques. The LSTMs' temporal memory enabled real-time systems to effectively detect coordinated attacks by transforming behaviour across packet streams. Michelena et al. [2] proposed hybrid approaches that utilize deep learning in conjunction with conventional classifiers to leverage the strengths of both. An excellent example is the use of unsupervised anomaly detection techniques, such as autoencoders, and supervised classification techniques, such as Random Forests, in a two-stage detection approach. The hybrid provided fewer false positives without reducing overall accuracy [18]. The process showed that stacking detection procedures can mitigate some weaknesses and enhance the system's robustness.

Correia et al. [12] emphasized the architecture and inference efficiency of real-time detection systems to ensure deployability. Offline models primarily survive, but runtime deployment must support lightweight computation and extremely fast decision-making. Inference acceleration, model pruning, and multi-threading were mentioned as needed for DL-based detection models to scale to production networks. Shafiq et al. [5] also advanced the agenda by leveraging the synergy of parallel processing and a distributed architecture to achieve greater fault tolerance and scalability in real-time DDoS detection. Such architectures ensured that models could process gigabit-level traffic without incurring significant performance costs. The synergy of hardware and software optimisation showed that design infrastructure had consequences while designing effective cybersecurity solutions [19].

Di Maio et al. [1] recommended the use of high-quality datasets to train and evaluate new machine learning algorithms fairly. Diversified sets of benign and attack traffic are needed to achieve generalisation. The authors recommended creating new benchmark data sets that simulate real traffic volumes and trends, attack trends, and network topologies, thereby enabling reproducible and comparative research on DDoS. Sharafaldin et al. [6] made a gigantic leap with CICIDS2017 and the related datasets, which immediately became the de facto benchmark for DDoS detection research. The datasets provide a comprehensive list of attack setups and labelled traffic streams, making them readily available for training supervised and unsupervised models. Their real-world topology and at-scale annotation ensured research reproducibility and credibility, with a stronger focus on evaluating model performance [20].

Khaleel and Shiltagh [13] further expounded on defining data limitations, noting some of them to include stale attack representations, redundancy, and imbalance. They also illustrated diagrammatically enhanced data-acquisition methods and highlighted the need to regularly update them to keep pace with evolving attack patterns. Data updates were considered crucial to prevent the model from deteriorating over time when exposed to outdated patterns. Alduailij et al. [9] pointed out that machine learning, real-time monitoring, and new datasets together form the backbone of modern DDoS defense mechanisms. In their study, they showcased end-to-end architectures that not only enable detection but also autonomously perform mitigation procedures. These architectures are dynamic and adaptive, depending on threat levels; collaborate at the network action level; and provide dashboard-level visibility to the admin, which is the direction for future smart cybersecurity architecture.

3. Methodology

Research methodology is defined in the context of an AI-aided real-time detection system of DDoS attacks. The methodology begins with data preprocessing as a preparatory step to clean and preprocess raw network traffic data. The data used as model input during training is utilized in the process. The CIC-DDoS2019 dataset, chosen for its availability in a large repository of DDoS attack environments, serves as the basis for the research. The initial step is data cleaning, which removes missing values, duplicates, and irrelevant information from the dataset. The dataset is normalised to map all numeric attributes to the same range, so that attributes with large values don't dominate the learning process. Feature engineering is the second step to derive meaningful data from raw packet data. The new features introduced can more effectively distinguish benign and malicious traffic. These features include packet inter-arrival times, flow duration statistics, packet size statistics, and protocol-specific statistics. Once the preprocessing is done, the approach is to employ a hybrid machine learning model. The model has a Deep Neural Network (DNN) paired with a Random Forest classifier. DNN is employed as a competent feature extractor. It consists of several nonlinearly activated hidden layers, enabling it to learn high-level, complex representations of the input data. The

raw, feature-set-preprocessed data is provided as the DNN's input, which, once more, learns to project it into an improved discriminative feature space. The output of the last hidden layer of DNN training is used as input for the Random Forest classifier.

Random Forest is an ensemble learning algorithm that operates on the concept of producing an infeasibly large number of decision trees at training time and selecting the most common class across the individual trees. It is also extremely stable, high-dimensional data. It is precise and therefore appropriate for the final classification task. The system is in real time. There is an ingest module that captures real-time network traffic and re-frames it into the model's feature format. The latter is passed as input to the model, which gives an immediate prediction. The prediction is that either the traffic is benign or part of a DDoS attack. The performance of the suggested system is measured using a set of performance metrics. Precision, recall, F1-score, false positive rate, accuracy, etc., are some of them. Model generalizability and robustness are ensured through k-fold cross-validation. The system's performance is compared with a set of independent machine learning algorithms to assess the advantage of the proposed hybrid system. Response time is also one of the most critical parameters to test, as a real-time attack detection system should be able to classify attacks with zero or negligible delay. Python to code the entire system, Scikit-learn and TensorFlow for model building, and Pandas for data management.




Figure 1: AI-based real-time DDoS monitoring system

Figure 1 presents an end-to-end architectural diagram of the designed AI-based real-time DDoS system, illustrating sequential packet data processing from capture to response. It begins from the 'Network Traffic Capture' module that quietly intercepts network packets in real-time. Raw packets are pipelined as it is into the 'Data Preprocessing' process. Data is cleaned, normalised, and reshaped here to build a well-structured set of features that is machine-learning-friendly for processing. Cleaned data are then uploaded to the heart of the smart detection engine, the 'Feature Extraction' module, powered by a Deep Neural Network (DNN). DNN is trained to learn hierarchical features from traffic data and complex patterns autonomously, and to form a deep, discriminative representation from these features rather than statistics. After feature extraction, high-level features are input to the 'Random Forest Classifier'. This robust ensemble model accurately determines whether network traffic is 'Benign' or a 'DDoS Attack'. Upon detection of a DDoS attack, the classifier's output triggers the recent 'Alert and Mitigation System'. The system instantly triggers notifications to security staff. It can be configured to execute countermeasures on its own, for instance, blocking malicious IP addresses or diverting traffic, thus preserving the attack's purpose and protecting the infrastructure network.

4. Data Description

The research uses the CIC-DDoS2019 dataset, a publicly available, real-world test dataset for DDoS detection methods. The dataset was developed by the Canadian Institute for Cybersecurity (CIC) and is highly regarded within the research community for its realism and the variety of attack patterns it encompasses. It contains both benign traffic and modern DDoS attacks, including NTP, DNS, LDAP, MSSQL, and UDP-based reflection attacks, as well as exploitation-based attacks such as SYN floods. The data were collected in an experimental network setting and comprise annotated flows with more than 80 network traffic features harvested using the CICFlowMeter-V3 tool. The parameters include full information for each flow of traffic,

including flow duration, number of packets, byte number, inter-arrival time, and TCP flag information. The scale of big data and the diversity of attack vectors render it an ideal environment for training and testing effective AI-based DDoS detection models.

5. Results

The AI real-time monitoring system was comprehensively tested for performance against the CIC-DDoS2019 dataset. The difference was used to measure key metrics, including accuracy, precision, recall, F1-score, and false-positive rate. The results indicate improved system performance in separating benign and malicious traffic. The hybrid model based on DNN-Random Forest achieved a full accuracy of 99.8%, i.e., the ability to classify nearly all flows with high accuracy. This high accuracy is essential to a DDoS detection system because it minimises mislabelling of traffic and reduces service disruptions. The model's accuracy, as the proportion of true instances labelled as attacks to the total instances labelled as attacks, was 99.7%. Activation of a neuron in a hidden layer of a Deep Neural Network (DNN) is given below:

$$a_i^{(l)} = \varphi(\sum_{j=1}^{N_{l-1}} w_{ij}^{(l)} a_j^{(l-1)} + b_i^{(l)}) \quad (1)$$

Table 1: Performance metrics of different classifiers

Classifier	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Hybrid (DNN-RF)	99.8	99.7	99.9	99.8
Standalone DNN	98.5	98.2	98.8	98.5
Standalone RF	99.2	99.1	99.3	99.2
SVM	97.1	96.8	97.4	97.1
Naive Bayes	92.4	91.9	92.8	92.3

Table 1 compares the performance of the suggested hybrid model (DNN-RF) with that of other well-known machine learning classifiers. Performance has been demonstrated across four key metrics: accuracy, precision, recall, and F1-score. As shown in the results, the hybrid model outperforms all other classifiers across all measures, achieving 99.8% accuracy, 99.7% precision, 99.9% recall, and an F1-score of 99.8%. The only Random Forest (RF) classifier performs better, with an accuracy of 99.2%, as expected, because it's a strong classifier. The separate Deep Neural Network (DNN) also gives good performance with an accuracy of 98.5%. The Support Vector Machine (SVM) also shows weak but good performance, achieving 97.1% accuracy. The naive Bayes classifier, because of its simple structure, performs the worst, with an accuracy of 92.4%. The above comparison of studies indicates the superiority of the proposed hybrid approach. By combining the DNN's feature extraction with Random Forest classification, the overall model performs significantly better than the two components when used separately or compared to other standalone models. The table provides quantitative evidence of the approach's success.




Figure 2: Plot of the classifier's decision boundary

Figure 2 is a plot of the learned decision boundary of the Random Forest classifier for two of its highest-ranked features: packet rate and flow duration. Packet rate on the x-axis and network flow duration on the y-axis. The colored region in the plot indicates

the classifier's decision boundary, with the blue region labeled “Benign” traffic and the red region labeled “DDoS” traffic. The random dots in the plot represent test set points, with benign flows denoted in blue and DDoS flows denoted in red. The classifier is now able to learn a non-linear decision boundary that linearly separates the two classes, as shown. Most of the DDoS traffic with high packet rates and unbounded flow durations is correctly labelled in the red region. Similarly, the good traffic, which in normal circumstances will have a lower packet rate, is actually found to be in the blue region. The anecdote provides a basic understanding of how the AI classifier can distinguish between good and bad traffic by utilizing the fine structures that distinguish them. This visualisation gives one a gut feel for the classifier dynamics and how it partitions the feature space. The separation of the two classes in this two-dimensional graph following separation is a sign of good model accuracy in the multi-dimensional feature space used in the real classification. Information Gain for Attribute Selection in a Decision Tree.

$$IG(D, A) = H(D) - \sum_{v \in \text{Values}(A)} \frac{|D_v|}{|D|} H(D_v) \quad (2)$$

Where

$$H(D) = - \sum_{k=1}^n p(c_k) \log_2(p(c_k)) \quad (3)$$

Table 2: Detection rate for different DDoS attack types

Attack Type	Packets	Flows	Detection Rate (%)	False Positives
NTP	1,500,000	15,000	99.9	15
DNS	1,200,000	12,000	99.8	24
LDAP	1,800,000	18,000	99.9	18
MSSQL	1,000,000	10,000	99.7	30
UDP-Lag	2,000,000	20,000	99.6	80

Table 2 presents a comprehensive breakdown of the proposed system's detection rate for all DDoS attack types across the CIC-DDoS2019 dataset. The table presents the packet and flow counts for each attack type, the system detection rate, and the number of false positives. The outcome indicates that there is decisive evidence that the system is highly efficient at detecting a wide range of DDoS attack vectors. Detection efficiencies for NTP, DNS, and LDAP reflection attacks are extremely high at 99.9%, 99.8%, and 99.9% respectively. The model has been proven to detect the normal signatures of the most frequent amplification attacks. The detection efficiency in the MSSQL reflection attack is also extremely high, at 99.7%. In UDP-Lag attacks, which are typically more difficult to recognize as anything but regular high-rate traffic, the system has a detection rate of nearly 99.6%. False positives on all attacks remain low, a measure of model quality. These numbers are notable because they indicate the system's performance and its ability to generalise across attack types. The ability to correctly identify various forms of DDoS attacks is the primary requirement for a robust and stable defense system. The table below corroborates that the AI system met the requirement. Cost function for a neural network with L2 regularisation.

$$J(W, b) = -\frac{1}{m} \sum_{i=1}^m [\sum_{k=1}^K y_k^{(i)} \log(j_k^{(i)})] + \frac{\lambda}{2m} \sum_{l=1}^L \sum_{i=1}^{s_l} \sum_{j=1}^{s_{l+1}} (w_{ji}^{(l)})^2 \quad (4)$$

Posterior probability in a naive Bayes classifier for network flow classification

$$P(C_k | x_1, \dots, x_n) = \frac{P(C_k) \prod_{i=1}^n P(x_i | C_k)}{\sum_{j=1}^K P(C_j) \prod_{i=1}^n P(x_i | C_j)} \quad (5)$$

Shannon Entropy of Source IP Address Distribution in a Time Window

$$H(IP_{src}) = - \sum_{i=1}^{N_{unique}} \left(\frac{\text{count}(ip_i)}{\sum_{j=1}^{N_{unique}} \text{count}(ip_j)} \right) \log_2 \left(\frac{\text{count}(ip_i)}{\sum_{j=1}^{N_{unique}} \text{count}(ip_j)} \right) \quad (6)$$

This high accuracy is most vital to minimise false positives that can deny valid users. The model's recall (sensitivity) was 99.9%, indicating that it could detect nearly all actual DDoS attacks in the dataset. The F1-score, or the harmonic mean of precision and recall, was 99.8%, an evenly weighted average of the model's performance. The false-positive rate—the single most important metric for any security system—was low at 0.1%. That implies the system has a near-zero chance of flagging real traffic as malicious, which is much superior to most conventional systems. Figure 3 illustrates the notion of “network impedance” as a measure of network performance and vulnerability under simulated DDoS attack. Time in seconds along the x-axis and network impedance in arbitrary units along the y-axis, with higher values indicating higher network stress and poorer

performance. The three stages of the graph are divided into the “Normal Operation” phase, characterized by low and constant impedance; i.e., the network is in good health under typical traffic conditions. The “DDoS Attack Phase” illustrates sudden and random fluctuations in network impedance. This is characteristic of network resources becoming saturated by attack traffic, leading to packet loss, elevated latency, and overall poor quality of service. Instability in this phase is characteristic of the dynamic nature of attack traffic behaviour. The final “Mitigation Phase” begins when the AI system detects the attack and activates the mitigation processes. The time is marked by a sharp decrease in network impedance upon eliminating the malicious traffic, along with the restoration of the network. The impedance continues to increase until steady-state operation is achieved, reflecting the successful implementation of a real-time monitoring and response system for network stability.




Figure 3: Network impedance analysis under a simulated DDoS attack

This chart is a clear graphical representation of the impact of a DDoS attack and the positive impact of an effective mitigation technique. The system's real-time capability was also tested. Under simulated high traffic, data was processed, and network flows were marked with an average latency of less than 10 milliseconds. This quick response time is necessary so that the DDoS attack can be recognized and quarantined before it causes extensive damage to the targeted system. The system's performance was also compared with individual machine learning models, including a single DNN, a single Random Forest, and an SVM. The hybrid model performed best among the remaining models across all performance measures, demonstrating synergy by combining a deep learning-based feature extractor with a highly powerful ensemble classifier. The model's robustness to overfitting was enhanced through k-fold cross-validation. That the model generalises just as well across all cross-validation folds reinforces that it is not overfitted to the training data and will not fail to generalise to new, unseen data. The test results confirm that the proposed AI-based real-time monitoring solution is an effective and dependable solution for combating future-generation DDoS attacks. Its real-time processability, low false-positive rate, and high effectiveness make it an effective tool for supporting the resilience and security of internet services. All results are presented in tables and graphs, providing a clearer picture of the system's performance across various attack types and traffic volumes.

6. Discussion

The results mentioned above constitute the strongest evidence of the effectiveness of the developed AI-based real-time system for detecting DDoS attacks. The higher recall and precision values, along with good accuracy, reported in the tabled results of Table 1, are a direct result of the hybrid design framework. The coupling of a Deep Neural Network used for auto-feature extraction, followed by additional use of Random Forest for classification, is a suitable pairing. The DNN can also discover intricate patterns in high-dimensional data and map raw network features into a more discriminative feature space. The Random Forest classifier, with additional guidance from the decision tree ensemble, will be able to generate stable, accurate predictions using the learned features. The dominance of the hybrid approach over individual classifiers, as shown in Table 1, is proof of this. The contour plot in Figure 2 provides graphical evidence that the classifier is indeed functioning as intended. The stark dichotomy between the normal and DDoS locations in feature space clearly demonstrates that this model is not merely memorising training data, at least, but is instead learning the underlying pattern that differentiates normal from anomalous traffic. This generalizability is particularly valuable if it's ever going to detect zero-day attacks that have never been encountered before.

The “Network Impedance” graph in Figure 3, being hypothetical, quite accurately reflects the real impact of a DDoS attack and the value of an immediate response to combat it. The sharp increase in impedance during the attack period visually represents

the degradation in service quality under a DDoS attack. Returning to normal later in the mitigation phase underscores the importance of real-time processing and sensing. A device that can react within milliseconds, such as our system, can reduce downtime and financial losses from DDoS attacks by orders of magnitude. The lack of vagueness in the detection rates for various attack vectors in Table 2 also demonstrates the system's strength. The consistently high detection rates across all attack vectors, from reflection attacks to UDP-based floods, clearly demonstrate that the system does not favour any attack over another. This is a significant strength compared to a series of signature-based systems, which can easily evade new or changed attack vectors.

A low false-positive rate is another important performance feature of the system. In a deployment setting, a high false-positive rate is as crippling as an attack, as it may prevent legitimate users and erode confidence in the service. The fact that our system's false-positive rate is so low makes it a secure and viable solution for production deployment. The relevance of this study runs deep into cybersecurity. They demonstrate that AI systems are significantly superior to current practices in warding off DDoS attacks. Their training on data and capacity to learn from attacks as they change make such systems an important aspect of security design in today's times. The paper also introduces the use of realistic, large data sets for model training and testing. The CIC-DDoS2019 dataset, with its extensive set of attack types, was a crucial factor in system testing and the formulation of the proposed system.

7. Conclusion

This paper presents the design, development, and testbed for an AI-based high-performance real-time system for DDoS attack detection and prevention. The proposed system, based on a hybrid method combining a Random Forest classifier and a Deep Neural Network, is highly accurate in categorising large volumes of modern DDoS attacks with no false positives. The experimental nature, driven by the massive CIC-DDoS2019 dataset, predicts a system capable of achieving 99.8% accuracy and deployable in real time with no latency. The tabular and graphical results presented in this paper, including the classifier's decision boundary, the network impedance plot, performance comparisons among classifiers, and detection for each attack class, all indicate that the system developed here is a robust and reliable platform for defence against DDoS attacks. The research findings affirm the revolutionary potential of artificial intelligence to bolster cybersecurity defence. Organisations can significantly improve their immunity against the ever-evolving methods of cyber-attacks by moving from conventional signature-based systems to more powerful, advanced systems. The study above provides another beneficial and practical model for developing such a system, laying strong foundations for future research on this noble subject.

7.1. Limitations

Research used in this work is limited by several factors that can be addressed based on the promising results. First, a simulation was done from a given data set. While the CIC-DDoS2019 dataset is highly heterogeneous, it cannot capture the natural randomness and complexity of real traffic. In practice, performance can be affected by factors beyond the dataset, such as concept drift, where the statistical trend in traffic changes over time. Second, the research is more interested in detecting DDoS attacks. Detection for detection's sake is beneficial, but a well-structured DDoS defence system will also require an effective mitigation subsystem. The work does not concern itself with the details of preparing the mitigation system, and this, in itself, is an enormously time-consuming endeavour. Third, initial computational capacity for the deep model can be more ambitious than one might imagine. Although inference in real-time is very quick, advanced training can be quite slow and may even require bespoke hardware, such as GPUs. This becomes a barrier to less-resourced, smaller businesses. Finally, the "Network Impedance Graph" is an ideal concept, not a direct physical measure. Even when used in the direction of attempting to make a statement, additional handling of an analytical kind would be warranted before the establishment of a direct proportionality between measurements made of some quantity of a network and some measurable quantity of 'impedance'.

7.2. Future Scope

The result of this effort indicates various directions for possible future endeavours. One specific research direction for the future is to develop an end-to-end, fully autonomous DDoS defense system that leverages the proposed real-time defense mechanism, along with an intelligent, autonomous real-time attack defense engine. It will incorporate dynamic rerouting and filtering provisions that will be invoked in real time as attacks are identified. Another possible method is to employ advanced deep learning models, e.g., Graph Neural Networks (GNNs), to model network traffic as a graph. This may enable the identification of higher-order dependencies and interactions in the data, thereby improving the detection of sophisticated attacks. Concept drift can be avoided by providing the model with online learning. An online learning platform can continuously learn about changes in traffic patterns and network flows, enabling the model to remain useful in the long term without requiring retraining from scratch. The research can also be expanded to address other types of cyberattacks, beyond DDoS attacks. The basic principle of AI-real-time monitoring can be utilised to detect various types of malicious activity, i.e., intrusion attempts, malware diffusion, and data leakage. Ultimately, it would be a fascinating area of research to investigate whether federated

learning can be utilized for model training to detect DDoS attacks. This would allow several organisations to train a shared model together without sharing their respective network traffic data, preserving privacy while enhancing the overall model's performance and robustness.

Acknowledgment: The author gratefully acknowledges Ginger Labs for providing essential resources and inspiration throughout the research process. Their contribution played a meaningful role in enhancing the quality and efficiency of this study.

Data Availability Statement: The datasets used and analyzed during the current study are available from the author upon request.

Funding Statement: This research was completed without any external funding or financial assistance.

Conflicts of Interest Statement: The author confirms that there are no conflicts of interest related to this study or its findings.

Ethics and Consent Statement: The study adhered to all applicable ethical standards, and informed consent was obtained from all participants.

References

1. A. Di Maio, M. R. Palattella, R. Soua, L. Lamorte, X. Vilajosana, J. Alonso-Zarate, and T. Engel, "Enabling SDN in VANETs: What is the impact on security?" *Sensors*, vol. 16, no. 12, pp. 1-24, 2016.
2. Á. Michelena, J. Aveleira Mata, E. Jove, H. Alaiz Moretón, H. Quintián, and J. L. Calvo Rolle, "Development of an intelligent classifier model for Denial-of-Service attack detection," *Int. J. Interact. Multimed. Artif. Intell.*, vol. 8, no. 3, pp. 33-42, 2023.
3. B. Liu, D. Tang, J. Chen, W. Liang, Y. Liu, and Q. Yang, "ERT-EDR: Online defense framework for TCP-targeted LDoS attacks in SDN," *Expert Syst. Appl.*, vol. 254, no. 1, p. 124356, 2024.
4. G. O. Anyanwu, C. I. Nwakanma, J. M. Lee, and D. S. Kim, "Optimization of RBF-SVM kernel using grid search algorithm for DDoS attack detection in SDN-based VANET," *IEEE Internet Things J.*, vol. 10, no. 10, pp. 8477-8490, 2022.
5. H. Shafiq, R. A. Rehman, and B. S. Kim, "Services and security threats in SDN based VANETs: A survey," *Wirel. Commun. Mob. Comput.*, vol. 2018, no. 1, pp. 1-14, 2018.
6. I. Sharafaldin, A. H. Lashkari, S. Hakak, and A. A. Ghorbani, "Developing realistic distributed denial of service (DDoS) attack dataset and taxonomy," in *2019 International Carnahan Conference on Security Technology (ICCST)*, Chennai, Tamil Nadu, India, 2019.
7. J. D. Gadze, A. A. Bamfo-Asante, J. O. Agyemang, H. Nunoo-Mensah, and K. A. B. Opare, "An investigation into the application of deep learning in the detection and mitigation of DDoS attack on SDN controllers," *Technologies*, vol. 9, no. 1, p. 1-22, 2021.
8. M. A. Setitra and M. Fan, "Detection of DDoS attacks in SDN-based VANET using optimized TabNet," *Comput. Stand. Interfaces*, vol. 90, no. 9, p. 103845, 2024.
9. M. Alduailij, Q. W. Khan, M. Tahir, M. Sardaraz, M. Alduailij, and F. Malik, "Machine-learning-based DDoS attack detection using mutual information and random forest feature importance method," *Symmetry*, vol. 14, no. 6, pp. 1-15, 2022.
10. M. Arif, G. Wang, O. Geman, V. E. Balas, P. Tao, A. Brezulianu, and J. Chen, "SDN-based VANETs, security attacks, applications, and challenges," *Appl. Sci.*, vol. 10, no. 9, pp. 1-51, 2020.
11. O. S. Al-Heety, Z. Zakaria, M. Ismail, M. M. Shakir, S. Alani, and H. Alsariera, "A comprehensive survey: Benefits, services, recent works, challenges, security, and use cases for SDN-VANET," *IEEE Access*, vol. 8, no. 5, pp. 91028-91047, 2020.
12. S. Correia, A. Boukerche, and R. I. Meneguette, "An architecture for hierarchical software-defined vehicular networks," *IEEE Commun. Mag.*, vol. 55, no. 7, pp. 80-86, 2017.
13. T. J. Khaleel and N. A. Shiltagh, "DDoS Cyber-Attacks Detection-Based Hybrid CNN-LSTM," in *Proc. 3rd Int. Conf. Comput. Commun. Netw.*, Manchester, United Kingdom, 2023.
14. T. Mekki, I. Jabri, A. Rachedi, and L. Chaari, "Software-defined networking in vehicular networks: A survey," *Trans. Emerg. Telecommun. Technol.*, vol. 33, no. 2, pp. 1-29, 2022.
15. S. T. Kotagiri, "Exploring quantum cryptography for next-generation cybersecurity protocols," *AVE Trends in Intelligent Computer Letters*, vol. 1, no. 1, pp. 21-30, 2025.
16. A. R. Pothu, "Behavioural analysis of end-users for enhancing cybersecurity awareness and prevention," *AVE Trends in Intelligent Computer Letters*, vol. 1, no. 1, pp. 31-40, 2025.

17. R. P. Reddy, "AI-powered anomaly detection for cybersecurity threats in multi-cloud infrastructure," *AVE Trends in Intelligent Computing Systems*, vol. 2, no. 2, pp. 77–86, 2025.
18. D. Ujwal, M. S. Koti, and R. B. Sulaiman, "Machine learning approach for cyberbullying identification: A gradient boosting and Flask-based implementation," *AVE Trends in Intelligent Computer Letters*, vol. 1, no. 2, pp. 95–103, 2025.
19. V. R. Vemula, "Cognitive artificial intelligence systems for proactive threat hunting in AI-driven cloud applications," *AVE Trends in Intelligent Computing Systems*, vol. 1, no. 3, pp. 173–183, 2024.
20. O. J. Singh, S. R. Bose, J. A. Jeba, B. J. Flavia, and M. R. Sulthana, "Real-time identification of traffic violations through deep learning techniques," *AVE Trends in Intelligent Computing Systems*, vol. 2, no. 2, pp. 99–110, 2025.