AI-Driven Firewall Log Analysis: Enhancing Threat Detection with Deep Learning Techniques

Yasmine ABOUDRAR, Khalid BOURAGBA, Mohamed OUZZIF
Computer Science and Smart Systems, ESTC, Hassan II, Casablanca, Morocco

***Abstract*—As cyber-attacks get increasingly sophisticated, cybersecurity threats have surged, with 430 million new malware instances identified in 2023 representing a 36% rise compared to 2020 figures in the United States.Traditional firewall defense mechanisms are increasingly restricted. Even though firewalls are the frontline defense mechanism, their reliance on preconfigured rules and signature-based detection leaves them behind in the identification of carefully crafted, dynamic attacks. Furthermore, they generate enormous volumes of logs and hence add high false positive rates, making manual threat analysis a tedious and time-consuming process. In order to counter such issues, we propose an AI-fortified SIEM system using deep learning algorithms for intelligent firewall log analysis. This serves to reduce false positives through event pattern extraction and correlation, allowing for more efficient threat detection. By employing deep neural networks like fully connected, convolutional, and recurrent, our system enhances classification accuracy and optimizes threat detection. We utilize actual firewall logs and benchmarking datasets (UNSW-NB15-training and UNSW-NB15-testing) to assess our system, one for training and the other for testing. Our primary objective is to differentiate between true positive and false positive alarms so that security analysts can respond to cyber threats more effectively. The experimental results demonstrate the effectiveness of our approach in improving threat monitoring and IT security. Besides, they confirm that our learning-based models are better than classical machine learning methods and are therefore a realistic and efficient solution to real-world firewall security.**

Keywords*—*AI-driven SIEM; deep learning; firewall log analysis; threat detection; false positives; cybersecurity

I. Introduction

Based on the development of artificial intelligence (AI) techniques, learning-based techniques for detecting cyber threats have improved significantly, with promising results in most research studies [17]. However, because cyber threats continue to evolve, protection of IT systems against malicious activity is a serious issue. As a result of the increased evolvability of network intrusions and cyberattacks, organizations are compelled to develop robust security features as their top priority for effective defense.

Traditional firewalls are utilized traditionally as the initial line of defense against unauthorized access and cyber attacks [18]. They filter, and route incoming and outgoing network traffic according to security policies. Firewalls may be configured using signature-based systems to detect well-known patterns of attack and enforce access control to deny malicious traffic into the network. In today's cybersecurity architectures, firewalls are augmented with security information and event management (SIEM) systems to collect and process security incidents to help security analysts investigate anomalous behavior and correlate events to determine the likelihood of threats.

Firewalls have limitations in detecting sophisticated cyber attacks. Their increased level of security logs has the tendency to produce overwhelming amounts of data, and this makes it difficult for the analysts to actually identify real threats from numerous false positives. Also, traditional rule-based solutions don't work while detecting new patterns of attacks as they are implemented using pre-defined signatures and policies. To alleviate these challenges, research has gone into AI-based solutions for the analysis of firewall logs, and machine learning is employed to facilitate the detection of threats.

AI-based solutions are able to commit firewall security through the ability to learn from historic threat data and identifying anomalies in real time. Such solutions can help security analysts automate security incidents categorization, priority of alerts, and false positives removal. Learning-based detection models, however, have disadvantages as well. They require labeled datasets in order to effectively train and test models, for instance, but high-quality-labeled data in scale isn't available everywhere. High-quality labeled data enough to use for supervised learning isn't possessed by many commercial SIEM vendors, limiting their value in real-world scenarios.

Second, features derived from security logs across different studies can be non-transferable to the real-world in some cases. Although such datasets are very informative, they can be lacking in terms of reflecting the nature of real-world firewall logs since they have a narrow range of features. In order to promote real-world utility, AI-driven models must be trained and tested with real-world security event data that is sourced from live environments.

Third, anomaly-based detection products, favored by AI-driven security analysis, have the potential to strengthen the detection of unknown threats [19]. However, they generate a large number of false positives, hence overwhelming security teams with unwanted alarms that must be manually verified. This inefficiency increases response time and operational expense. Further, attackers continually update attack strategies, necessitating the AI models to adapt dynamically in order to detect evolving threats. To address these limitations, we propose an AI-based SIEM system that is capable of processing firewall logs using deep learning. Our system aims to distinguish true threats from alarms using event pattern extraction and correlation analysis. Through correlating security events based on concurrency and similarity features, our approach provides useful input data for deep neural networks, which can effectively classify threats.

The major contributions of this work are:

- We propose a method to transform large volumes of firewall logs to structured syslog for scalable security event analysis. Our system learns normal and attack patterns from data collected and considers event frequency and relevance.
- Our AI-fueled SIEM solution, as opposed to traditional sequence-based methods, provides improved input to deep learning algorithms for threat identification with greater precision and less false positives.
- We evaluate our system performance using real-world firewall logs of a security operations center (SOC) compared with existing machine learning methods and establish our method using benchmark dataset such as UNSW-NB15 to demonstrate its efficacy in network security research.

The produced events are subsequently employed as inputs to deep learning models such as fully connected neural networks (FCNN), convolutional neural networks (CNN), and long short-term memory (LSTM) networks. In testing our system using benchmark datasets and actual data, we aim to demonstrate how it is beneficial in enhancing firewall security and safeguarding IT infrastructures against constantly evolving cyber threats. This study is guided by the following central research question:

Can an AI-augmented SIEM system, powered by deep learning, significantly enhance the accuracy and efficiency of threat detection in firewall syslog data compared to traditional rule-based SIEM approaches?

The rest of this study is organized as follows: Section II provides background information on firewall security and AI-driven threat detection, and outlines our system architecture and labeling data approach proposed. Section III illustrates the current literature on AI-driven cyber threat detection. Section IV outlines our research approach utilized, and provides the datasets utilized in our experiments, and outlines the deployment of our AI models. Section V reports the results of evaluation and comparisons with traditional approaches. Section VI presents the discussion and key findings. Section VII details the preprocessing steps applied. Lastly, Section VIII provides our concluding remarks and perspectives for future works.

II. Preliminaries

This section presents, a summary overview of the major concepts related to our study. We begin by demonstrating the major concepts of Firewalls and SIEM systems and introduce deep learning algorithms subsequently. We then describe our big data platform for our AI-based SIEM system.

A. Firewall and SIEM

1) Firewall: As a cornerstone of network security, firewalls play a vital role in protecting systems against such threats [20]. A firewall is a fundamental security appliance that screens and controls incoming and outgoing network traffic based on defined security rules. Firewalls are the first line of protection, which are utilized heavily by organizations to reject unauthorized access and withstand cyber-attacks. However, although firewalls are vital, they are bedeviled by numerous shortcomings when handling modern-day cyber-attacks.

Most traditional firewalls are rule-based and signature-based, hence less capable of handling sophisticated, dynamic, and zero-day attacks. Further, the increasing volume and complexity of network traffic cause problems in accurately distinguishing between good and bad behavior. This is likely to create many false positives, which flood security teams with useless notifications and make true threat detection difficult.

2) SIEM: A Security Information and Event Management (SIEM) system is a valuable part of enterprise security, providing a centralized view of an organization's cyber security stance. Through the collection and analysis of security information from a variety of sources [15], including firewalls, IDS/IPS, and other network security solutions, SIEM identifies cyber threats by identifying patterns and anomalies. They utilize event correlation as well as security policies defined before to attempt discovering possible attacks while also examining the security incidents. Firewalls as perimeter security generate huge volumes of logs which are being devoured by SIEM solutions to process. The logs contain valuable information about network traffic, access attempts, and potentially security incidents. The traditional rule-based threat detection methods, on the other hand, cannot identify benign traffic and real threats, and therefore have a high rate of false positives. This makes manual threat investigation time and labor-intensive for security professionals.

One of the biggest challenges for SIEM-based firewall log analysis is handling the volume and variety of security events [21]. The complexity of modern cyber threats requires detection capabilities that exceed the capabilities of traditional policy-based approaches. As a result, recent SIEM innovations have incorporated artificial intelligence (AI) and machine learning (ML) capabilities to aid in enhancing threat detection and reducing human effort. Machine learning-based SIEM products improve alert prioritization, anomaly detection, and automated response.

Despite such advancements, AI-driven SIEM still has some significant hurdles to cross, e.g., 1) a need for continual analyst monitoring, 2) the unavailability of high-quality labeled datasets to train upon, and 3) continuously changing cyber threats. These shortcomings need to be addressed in an effort to better enhance firewall security in SIEM systems and enhance security operations.

B. Deep Learning Techniques

Deep learning techniques have improved significantly over the past couple of years and are still expanding into numerous other areas beyond traditional machine learning. The techniques rely on mathematical formulations of neural networks based on the human neural network. Two of the most widely used deep neural network models are convolutional neural networks (CNNs) and recurrent neural networks (RNNs). CNNs excel greatly in learning spatial data from data, and for such use like image processing, CNNs are suitable.

Particularly designed for handling spatial data, CNNs leverage the powers of local feature detection and shared parameter

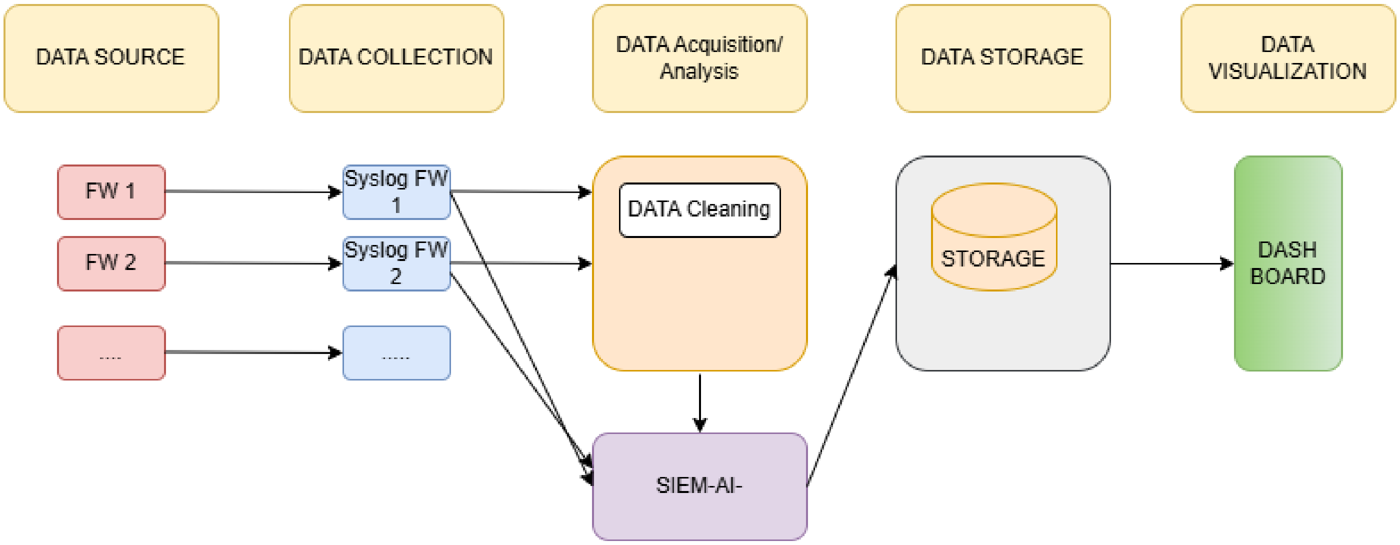

Fig. 1. Proposed architecture SIEM AI based.

setups and therefore have have had their application spread across so many different fields. CNNs have already excelled greatly in applications such as image recognition, biomedicine text mining, and virus classification. In network security logs for an extended duration. Additionally, a platform of this nature can be optimized for data analysis and swift detection of cyber threats.

This is because it has historical data, through which security incidents are investigated and responded to.

To address the aforementioned challenges, we developed a large-scale big data platform leveraging distributed computing technologies to collect, process, store, correlate, and analyze security event logs. As illustrated in Fig. 1, our system architecture is composed of four main components: a data collection module, a data processing unit, a data analysis engine, and a data storage infrastructure. These components operate collaboratively to analyze and respond to cyber threats based on long-term security intelligence.

The Data Cleaning module plays a key role in getting raw syslog data ready for meaningful analysis. By removing noise and inconsistencies, it ensures that only clean, relevant, and high-quality information is passed along. This cleaned data is the only input used by the SIEM-AI module, which helps avoid issues like overfitting that can happen when both raw and processed data are used together. By keeping the process streamlined and focused, this approach preserves the integrity of the data pipeline and leads to more accurate, dependable threat detection powered by AI.

Through big data processing on a large scale, the system can ingest humongous amounts of security events in real-time and process them endlessly. Apart from this, our big data platform is also integrated with AI-based SIEM systems. Through the application of artificial intelligence, the system can better differentiate between actual security alarms and false positives, thus enhancing real-world threat identification and reaction.

III. Related Work

This section explores existing research relevant to AI-driven firewall log analysis, focusing on enhancing threat detection with deep learning techniques. It's divided into two main areas: Deep Learning-Based Intrusion Detection and Real Security Event Analysis.

Latest study [5] highlight how the integration of Next-Gen SIEM solutions with Data Lakes and AI significantly strengthens threat detection and response. The traditional SIEM solutions are limited by poor scalability, high false positives, and delayed processing of data, resulting in slow threat detection and operational inefficiencies. Through the utilization of Data Lakes for real-time big, unstructured data integration and analysis, and the application of sophisticated machine learning algorithms, the system improves at anomaly detection, learning from new threats, and reducing false positives. Not only does this novel approach negate the limitations of the previous SIEM tools, but also enhances detection rates, reduces workload on security teams, and strengthens overall cybersecurity posture of organizations.

A recent study explored how to improve cloud security by deploying a tailored SIEM solution that increases visibility and streamlines incident response. The authors set up a virtual network including VMs, load balancers, Microsoft Defender for Cloud, and a web application firewall (WAF), which helped filter malicious traffic and protect hosted applications. Their system was designed to keep a constant eye on cloud activity, helping to spot threats early while also ensuring that configurations stayed in line with industry standards. The study ultimately shows how combining built-in cloud tools with automation can lead to a more practical, scalable, and effective way to secure cloud infrastructures in the face of growing complexity [4].

A majority of the cyber security research is focused on AI-based anomaly detection in firewalls, where various AI and machine learning methods are proposed for enhancing cyber threat detection [3]. Anomaly detection has been studied in various fields and applications [1].The most challenging task is to detect outliers from normal data, especially when it comes to special attributes and new values.Therefore, the research is aimed at the implementation of Machine Learning and Deep Learning algorithms [1].

Artificial intelligence-based threat detection enhances traditional security controls by detecting advanced threats in real time so that organizations remain ahead of cyber attackers [3]. Deep learning technology is able to rapidly detect unusual behavior and strengthen the robustness of the system by using anomaly detection and log monitoring [7]. Advanced models integrate complex methodologies for mitigating the limitations of traditional methods and thereby improving the accuracy and effectiveness of anomaly detection [7].

Research also investigates WAF (Web Application Firewall) systems based on single and stacked LSTM layers on character sequences of user-input data with the best hyperparameter settings [1]. It employs a semi-supervised approach, trained on real attack payloads and typical HTTP data, with performance indicators expressed in terms of F1 scores [1].

Next-Generation firewalls (NGFWs) that incorporate artificial intelligence (AI) are the focus of further analysis in this study [6], and a comparison of their performances in dealing with modern-day cyber attacks is presented. As conventional firewalls are likely to fall short in detecting and responding to today's advanced attacks, AI offers a promising future ahead. According to critical examination of 20 to 25 authentic sources, the study compares different AI-based firewall models, their technology, and level of accuracy and efficiency. According to machine learning and deep learning principles, the study compares some of the most crucial parameters like detection rates, false positives, and resource utilization. Visual summaries allow for simplicity of presentation of contrast among methodologies, both strengths and weaknesses. By insight into present capabilities as well as limitations, the purpose of this research is to assist more informed decision-making on cybersecurity strategy and provoke yet further innovation around AI-based firewall solutions.

BM25 is a very successful information retrieval algorithm that enhances search effectiveness via precise document ordering from query words [2]. Building upon legacy TF-IDF method, BM25 comes with parameter controls to avoid term frequency saturation and length normalization of documents [2]. There is even a betterment of BM25-similar retrieval by the use of machine learning on learning its input attributes [8].

Although the idea of bringing AI into SIEM systems isn't new, many existing approaches still lean heavily on rule-based methods or conventional machine learning. The problem is, these methods often miss the mark when it comes to spotting complex or previously unseen threats. Our work takes a different path. By using deep learning tailored specifically to the patterns found in cleaned data, we aim to build a smarter, more adaptive system—one that's better equipped to detect modern cyberattacks as they evolve.

A. *Deep Learning-Based Intrusion Detection*

The application of deep learning to intrusion detection has garnered significant attention, offering an alternative to traditional log analysis methods that are often manual, time-consuming, and prone to error. Deep learning's ability to automatically learn features from large datasets makes it a promising avenue for enhancing security.

Several studies have explored deep learning for security log analysis. For instance, one study proposes an intelligent approach for classifying incoming and outgoing firewall traffic packets using deep neural networks [9]. Another research effort highlights the use of machine learning techniques to automatically detect activities recorded in firewall logs, thereby enhancing the security of corporate networks [10]. A comparative analysis of anomaly detection approaches in firewall logs highlights the importance of logs as a source of information and explores methods for generating anomalies to simulate real-world attacker behavior [11]. Additionally, various machine learning models, including K-Nearest Neighbor (KNN), Random Forest (RF), and Deep Neural Network (DNN), have been used to analyze firewall logs and measure performance in terms of accuracy, recall, precision, and F1-score [12].

B. *Real Security Event Analysis*

Analyzing real security events is crucial in cybersecurity. It involves monitoring and analyzing security events across a network. AI and ML are increasingly used to enhance threat detection and automation in this area.

One study presents a novel method for generating anomalies that simulate real attacker actions within the network [11]. Different supervised and unsupervised learning models were compared, revealing that unsupervised learning methods struggle to detect injected anomalies, suggesting they can be integrated into existing firewall logs. Furthermore, multi-class machine learning models can be used to identify the importance of each feature and determine the best actions, using large-scale real-world datasets [12]. Analyzing firewall logs is a significant practice for monitoring network traffic and assessing its impact [13]. One study presents a new framework that utilizes firewall log data to classify incoming data packets as either permitted or forbidden [10]. AdaBoost model can classify incoming data packets as either permitted or forbidden with a remarkable accuracy rate of 99.00% [10].

IV. Methodology

The present study proposes an advanced deep learning approach for the detection of malicious activities through firewall log analysis. To ensure the efficiency and robustness of the model, a rigorous methodology was followed, comprising data acquisition, preprocessing, model architecture design, training, and evaluation phases.

A. *Dataset Description*

The dataset used in this study is the "UNSW-NB15" dataset, a well-known benchmark dataset for evaluating intrusion detection systems (IDS) and firewalls. It contains a mix of normal and malicious network traffic captured in a realistic setting. The dataset consists of "175,341" training instances and "82,332" testing instances, each with "45 attributes", including both numerical and categorical features. The attacks are categorized into nine major classes, including

Fuzzers, Backdoors, DoS, Exploits, Generic, Reconnaissance, Shellcode, Worms, and Analysis.

We selected the UNSW-NB15 dataset, developed by the Australian Center for Cyber Security, due to its clear advantages over older datasets such as NSL-KDD. Featuring a diverse range of contemporary network attack scenarios categorized into nine distinct groups, it provides a highly relevant benchmark for evaluating modern threat detection systems.

It realistically represents enterprise-level network traffic and covers many recent attack types listed in the Common Vulnerabilities and Exposures (CVE) database. The data comes from two simulation sessions—one lasting 16 hours with an attack happening roughly once every second, and another lasting 15 hours with a higher attack rate of about ten per second. Features were extracted using tools like Argus and BRO-IDS, capturing various aspects of the network traffic such as flow characteristics, content details, and timing information. The dataset includes labels for both normal traffic and nine different types of attacks, offering a comprehensive and challenging benchmark for testing detection models [16].

Fig. 2 illustrates the distribution of attack types within the dataset, highlighting the dominance of the Exploits category, followed by Fuzzers and DoS, while rare classes such as Worms and Shellcode are underrepresented [14].

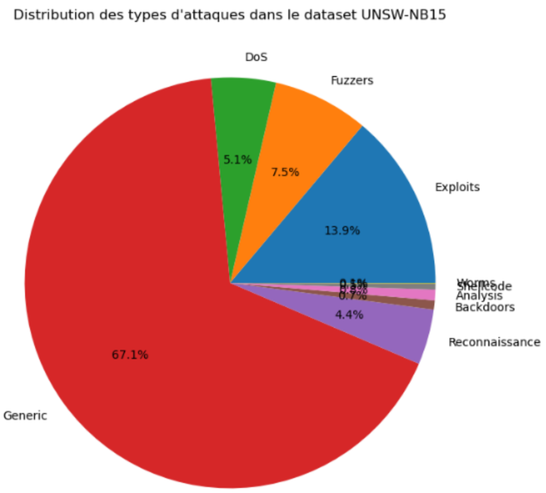

Fig. 2. Distribution of attack types in the UNSW-NB15 dataset.

Feature Selection and Preprocessing :

To enhance model efficiency, we performed feature engineering by:

- Dropping categorical and ID-based features: id, proto, service, state, attack_cat.
- Standardizing numerical features using StandardScaler.
- Encoding the labels using LabelEncoder, converting the binary classification task into a numerical format.
- Handling class imbalance via class weighting to ensure balanced training.

B. Model Selection Rationale

While Recurrent Neural Networks (RNNs), particularly LSTM architectures, are traditionally used for modeling sequential data such as firewall logs, recent studies have shown that one-dimensional Convolutional Neural Networks (1D CNNs) can also perform well in such tasks, especially when speed and scalability are priorities [22]. CNNs can efficiently extract local features from sequential data with significantly lower computational overhead compared to RNNs [23]. In our approach, we selected a 1D CNN due to its fast training time and suitability for real-time log analysis. To address concerns regarding model choice, we also conducted a comparative experiment using an LSTM-based model (see Table I).

C. Model Architecture

This study explores various Deep Learning architectures to improve threat detection. We designed and tested multiple architectures, including Multi-Layer Perceptron (MLP), Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), and hybrid CNN-LSTM models .

Proposed Hybrid CNN-LSTM Model :

The final selected architecture is a hybrid CNN-LSTM model, which leverages:

- CNN layers for spatial feature extraction.
- LSTM layers for capturing sequential dependencies.
- Batch Normalization to improve training stability and convergence.
- Dropout layers to prevent overfitting.

D. Model Structure

The proposed deep learning model follows a hybrid CNN-LSTM architecture, which is composed of the following layers:

- Input Layer: Reshaped input to $(39, 1)$.
- Convolutional Layers:
  - ○ Conv1D (128 filters, kernel size = 3, ReLU activation).
  - ○ Batch Normalization and Dropout (30%).
  - ○ Conv1D (64 filters, kernel size = 3, ReLU activation).
  - ○ Batch Normalization and Dropout (30%).
- LSTM Layer: Extracts temporal dependencies.
- Dense Layers: Fully connected layers for classification.
- Output Layer: Sigmoid activation for binary classification.

Model Architecture

The deep learning model designed for this study is a Deep Neural Network (DNN) composed of several densely connected layers. The architecture is detailed as follows:

- Input layer: 20 neurons corresponding to the selected features.

- Hidden layers:
  - First layer: 64 neurons, ReLU activation, Batch Normalization, Dropout (0.3).
  - Second layer: 32 neurons, ReLU activation, Batch Normalization, Dropout (0.3).
  - Third layer: 16 neurons, ReLU activation, Batch Normalization.
- Output layer: 9 neurons with Softmax activation for multiclass classification.

Training and Hyperparameters :

The models were trained using the following configuration:

- Optimizer: Adam (learning rate = 0.001)
- Loss Function: Binary Cross-Entropy
- Batch Size: 256
- Epochs: 100 (with Early Stopping, patience = 10)
- Validation Split: 20%
- Callbacks: ReduceLROnPlateau and EarlyStopping

E. Performance Metrics

To evaluate model performance, we used the following metrics:

- Accuracy
- Precision, Recall, and F1-score
- Confusion Matrix

V. Results

A. Model Performance Comparison

Table I presents the test accuracy obtained by each evaluated model, including MLP, CNN, LSTM, and the proposed CNN-LSTM architecture :

TABLE I. Test Accuracy Comparison of Different Models

Model	Test Accuracy
MLP	85.04%
CNN	81.95%
LSTM	83.17%
CNN-LSTM (Proposed)	**86.34%**

Fig. 3 illustrates the training progression of the proposed CNN-LSTM model in terms of accuracy over the epochs :

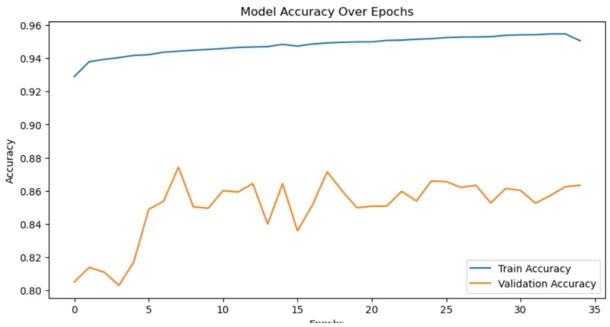

Fig. 3. Accuracy over epochs.

To further evaluate the model's classification performance, Fig. 4 shows the confusion matrix generated from the test dataset :

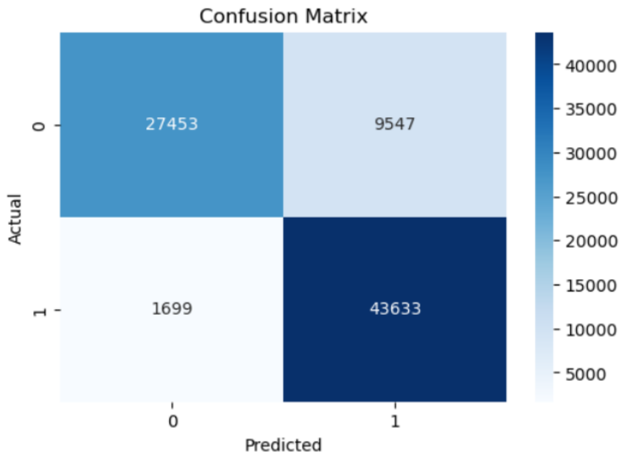

Fig. 4. Confusion matrix.

Additionally, the classification report displayed in Fig. 5 provides detailed metrics including precision, recall, and F1-score for each class.

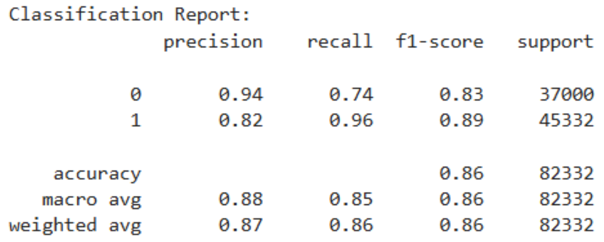

```
Classification Report:
              precision    recall  f1-score   support

           0       0.94      0.74      0.83     37000
           1       0.82      0.96      0.89     45332

    accuracy                           0.86     82332
   macro avg       0.88      0.85      0.86     82332
weighted avg       0.87      0.86      0.86     82332
```

Fig. 5. Classification report.

VI. Discussion

Compared to existing SIEM solutions, our proposed AI-SIEM demonstrates superior performance in detection accuracy, reduced false positives, and adaptability to evolving threats (see Table I). Unlike rule-based systems, our approach can learn complex attack patterns and generalize to novel data.

Table I presents the test accuracy results for the different deep learning models evaluated in our study. The Multi-Layer Perceptron (MLP) achieved a solid performance with an accuracy of 85.04%, while the Convolutional Neural Network (CNN) and the Long Short-Term Memory (LSTM) models reached 81.95% and 83.17%, respectively. Our proposed hybrid model, combining CNN and LSTM architectures, outperformed the others with the highest test accuracy of 86.34%. These results highlight the effectiveness of integrating spatial and temporal feature extraction for improving threat detection in firewall log analysis.

To verify the performance of our AI-SIEM model, we observed its accuracy at training epochs. As shown in Fig. 3, the model enhanced continuously with increasing epochs, and the curve of accuracy converged slowly, indicating steady learning and effective generalization.

In order to analyze the performance of the proposed AI-driven firewall detection model, we compared the confusion matrix, a common evaluation metric for classification issues. The matrix shown in Fig. 4 provides a detailed split of the predictions made by the model versus actual ground truth. It consists of four primary elements: True Positives (TP), where the model correctly identifies an attack; True Negatives (TN), where it correctly identifies normal traffic; False Positives (FP), where innocent traffic is wrongly labeled as malicious; and False Negatives (FN), where actual attacks are not detected. High TPs and TNs with low FP and FN show high detection capacity and reliability of the system. The results from the confusion matrix in this study confirm again that the model not only identifies with very high accuracy but also has a very low false positive rate, which is important to minimize unnecessary alarms in real-world security contexts. These findings vindicate the suitability of utilizing deep learning techniques for intelligent log analysis in future SIEM systems.

In addition to the confusion matrix assessment, a classification report with extensive information was utilized to evaluate the performance of the proposed model. The report in Fig. 5 provides significant metrics such as precision, recall, and F1-score for both classes: normal traffic (class 0) and attacks (class 1). For regular examples, the model was 94% accurate, so most predictions that were labeled as regular were indeed correct. However, the recall for the class was only slightly lower at 74%, so that some regular cases were being incorrectly labeled as attacks. In classifying attacks, though, the model was very good with a recall rate of 96% and precision of 82%, so it will catch most bad activity.

The F1-score, or the balance between precision and recall, was 0.83 for the normal class and 0.89 for the attack class, indicating balanced performance across classes. The overall accuracy of the model on the test set was 86% out of a total of 82,332 samples. Moreover, the macro average F1-score was 0.86, confirming well-balanced detection between the two classes, while the weighted average was also 0.86, confirming the reliability of the model in handling skewed data. The results confirm the effectiveness of the AI-augmented approach towards distinguishing normal and malicious traffic during real-world attacks.

TABLE II. Top 5 Selected Features Based on Relevance Scores

Rank	Feature	Score
1	sbytes	0.91
2	dbytes	0.88
3	ct_state_ttl	0.85
4	ct_src_dport_ltm	0.82
5	ct_dst_sport_ltm	0.80

A. Discussion and Key Findings

To gain a deeper understanding of the model's performance, we carried out a series of evaluation experiments supported by visual metrics. These analyses provide insight not only into the overall accuracy but also into how well the model handles different traffic types and attack categories. The key takeaways are as follows:

- The hybrid CNN-LSTM model achieved the highest accuracy (86.34%), outperforming traditional MLP and LSTM models.
- The confusion matrix shows that the model effectively classifies normal and attack traffic, though some misclassification remains between certain attack categories.
- Precision-Recall curves indicate a strong tradeoff, with high recall ensuring better attack detection.
- The ROC Curve AUC confirms the robustness of our model, achieving a high area under the curve.

VII. Data Preprocessing

To ensure optimal model performance, several preprocessing techniques were applied:

- Data cleaning was performed to remove duplicate records and handle missing values.
- Categorical features were encoded using one-hot encoding.
- Feature selection was conducted via the SelectKBest algorithm, retaining the top 20 features based on their statistical relevance to the target variable. Notable selected features include:
  - sbytes
  - dbytes
  - ct_state_ttl
- Normalization of numerical attributes was performed using MinMaxScaler to harmonize feature scales between 0 and 1.
- The dataset was split into 70% training, 15% validation, and 15% testing subsets to assess generalization.

To boost the model's performance while avoiding unnecessary complexity, we used a feature selection process based on relevance scores. This allowed us to concentrate on the most meaningful variables in the dataset. The top five features identified—sbytes, dbytes, ct_state_ttl, ct_src_dport_ltm, and ct_dst_sport_ltm_are listed in Table II. These features capture essential characteristics of network traffic and played a key role in helping the model accurately distinguish between normal and malicious activity.

A. Advantages of the Proposed Approach

Compared to existing SIEM solutions, our proposed AI-driven architecture demonstrates significant advantages. As shown in Table I, the CNN-LSTM model outperforms traditional approaches in terms of detection accuracy. In addition, it shows reduced false positive rates and stronger adaptability to evolving attack patterns.

Unlike conventional rule-based SIEM systems, which rely on manually defined signatures or heuristics, our model is capable of learning complex relationships within the data. This allows it to detect previously unseen threats and generalize to new types of network behavior. These strengths position the model as a valuable asset for next-generation threat detection in dynamic and large-scale environments.

VIII. Conclusion

In this study, we discussed an AI-enabled SIEM architecture designed to parse firewall logs intelligently and improve cyber threat detection by using deep learning models. Our main goal was to remove false alarms and improve overall threat detection accuracy using three types of models, namely, Fully Connected Neural Networks (FCNN), Convolutional Neural Networks (CNN), and Long Short-Term Memory Networks (LSTM).

To give a comprehensive and realistic evaluation, we trained and tested our models with actual firewall log data as well as the benchmark datasets UNSW-NB15-training and UNSW-NB15-testing. Among the three, the LSTM model performed best, particularly, since it can learn temporal patterns and process sequential data providing higher accuracy and lower false alarms.

By integrating these powerful deep learning techniques into the SIEM solution, we unveiled a wiser and watchful solution that was capable of identifying suspicious behavior in complicated network environments. Our findings reinforce the ability of AI in broadening the functionality of SIEM and paving the way for proactive and better cybersecurity monitoring.

Future work will focus on improving real-time detection capabilities, expanding the system to incorporate additional data sources such as endpoint and operating system logs, and conducting live performance evaluations in enterprise production environments. Furthermore, issues such as model generalizability, dependency on labeled data, and minimizing false positives warrant ongoing attention. Addressing these limitations will be crucial to deploying robust, scalable AI-enhanced SIEM solutions in practice.

References

[1] Toprak, Sezer, and Ali Gökhan Yavuz. ”Web application firewall based on anomaly detection using deep learning.” Acta Infologica 6.2 (2022): 219-244.

[2] Lu, Meng, Catherine Chen, and Carsten Eickhoff. ”Cross-Encoder Rediscovers a Semantic Variant of BM25.” arXiv preprint arXiv:2502.04645 (2025).

[3] Sola, Rajendra Prasad, Nihar Malali, and Praveen Madugula. Cloud Database Security: Integrating Deep Learning and Machine Learning for Threat Detection and Prevention: 0. Notion Press, 2025.

[4] Tuyishime, Emmanuel, et al. ”Enhancing cloud security—proactive threat monitoring and detection using a siem-based approach.” Applied Sciences 13.22 (2023): 12359.

[5] Marri, Rahul, Sriram Varanasi, and Satwik Varma Kalidindi Chaitanya. ”Integrating Next-Generation SIEM with Data Lakes and AI: Advancing Threat Detection and Response.” Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023 3.1 (2024): 446-465.

[6] Ahmadi, Sina. ”Next generation ai-based firewalls: a comparative study.” International Journal of Computer (IJC) 49.1 (2023): 245-262.

[7] Wang, Shuzhan, et al. ”Deep learning-based anomaly detection and log analysis for computer networks.” arXiv preprint arXiv:2407.05639 (2024).

[8] Shetty, Rahul. ”Enhancing Context-Aware Search with Retrieval-Augmented Generation.” Authorea Preprints (2025).

[9] Lillmond, Chandesh, and Geerish Suddul. ”A Deep Neural Network Approach for Analysis of Firewall Log Data.” (2021).

[10] et al. Afrah Fathima, “Cyber security Reinforcement through Firewall Log Analysis and Machine Learning”, IJRITCC, vol. 11, no. 10, pp. 777–782, Nov. 2023.

[11] Comparative Analysis of Anomaly Detection Approaches in Firewall Logs: Integrating Light-Weight Synthesis of Security Logs and Artificially Generated Attack Detection.

[12] Aljabri, Malak, et al. ”Classification of firewall log data using multiclass machine learning models.” Electronics 11.12 (2022): 1851.

[13] AL-Tarawneh, Batool A., and Hani Bani-Salameh. ”Classification of firewall logs actions using machine learning techniques and deep neural network.” AIP Conference Proceedings. Vol. 2979. No. 1. AIP Publishing, 2023.

[14] N. Moustafa and J. Slay, ”UNSW-NB15: a comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set),” 2015 Military Communications and Information Systems Conference (MilCIS), Canberra, ACT, Australia, 2015, pp. 1-6, doi: 10.1109/MilCIS.2015.7348942.

[15] Tiwari, Anamika, et al. ”Implementing Robust Cyber Security Strategies to Protect Small Businesses from Potential Threats in the USA.” Journal of Ecohumanism 4.3 (2025): 322-333.

[16] Li, Yueyang. ”Detection for Malicious Network Traffic Based on Convolutional Neural Networks.”

[17] Conti, Mauro, Tooska Dargahi, and Ali Dehghantanha. ”Cyber threat intelligence: challenges and opportunities.” Cyber threat intelligence (2018): 1-6.

[18] Liang, Junyan, and Yoohwan Kim. ”Evolution of firewalls: Toward securer network using next generation firewall.” 2022 IEEE 12th Annual Computing and Communication Workshop and Conference (CCWC). IEEE, 2022.

[19] Rangaraju, Sakthiswaran. ”Secure by intelligence: enhancing products with AI-driven security measures.” EPH-International Journal of Science And Engineering 9.3 (2023): 36-41.

[20] Kolli, RAJA KUMAR, E. Priyanshi, and S. Gupta. ”Palo Alto Firewalls: Security in Enterprise Networks.” International Journal of Engineering Development and Research, 12 (3), 1-13. rjwave ijedr/viewpaperforall. php? paper= IJE DR200A001 (2024).

[21] Shirazi, Patrick, and Ali Padyab. ”Discerning Challenges of Security Information and Event Management (SIEM) Systems in Large Organizations.” International Symposium on Human Aspects of Information Security and Assurance. Cham: Springer Nature Switzerland, 2024.

[22] Kim, Jiyeon, et al. ”CNN-based network intrusion detection against denial-of-service attacks.” Electronics 9.6 (2020): 916.

[23] Umair, Muhammad Basit, et al. ”A network intrusion detection system using hybrid multilayer deep learning model.” Big data 12.5 (2024): 367-376.