JOURNAL OF MECHANICS OF CONTINUA AND MATHEMATICAL SCIENCES
www.journalimcms.org

ISSN (Online) : 2454 -7190 Vol.-20, No.-8, September (2025) pp 115-132 ISSN (Print) 0973-8975

A NOVEL HYBRID MODEL FOR ROBUST CYBER INTRUSION DETECTION IN CLOUD COMPUTING ENVIRONMENTS

Mohamed Loughmari[1], Anass El Affar[2]

[1, 2] Engineering Sciences Laboratory (LSI), Polydisciplinary Faculty of Taza, Sidi Mohamed Ben Abdellah University, Taza, Morocco.

Email : [1]mohamed.loughmari@usmba.ac.ma, [2]anass.elaffar@usmba.ac.ma

Corresponding Author: **Mohamed Loughmari**

https://doi.org/10.26782/jmcms.2025.09.00007

(Received: May 21, 2025; Revised: July 16, 2025; Accepted: August 10, 2025)

Abstract

Security remains one of the most critical concerns in all types and sizes of networks. Among the various strategies and policies designed to protect networks and systems, intrusion detection systems (IDSs) are paramount in identifying and preventing attacks. As security threats evolve, next-generation security solutions are progressively incorporating artificial intelligence (AI) to enhance their effectiveness. Consequently, the building of an effective and intelligent intrusion detection system remains one of the most significant research challenges. This study proposes a novel hybrid IDS model that combines anomaly detection and supervised learning to improve attack detection in Cloud Computing (CC) environments. Our approach utilizes the CICIDS2018 dataset, noted for its large scale, recency, inclusion of diverse real-world attack scenarios, and suitability for CC contexts. Our methodology first employs Isolation Forest for anomaly detection. Then, the anomaly results are added as a new feature to the dataset. Subsequently, the eXtreme Gradient Boosting (XGBoost) model is employed on this enriched dataset. This two-stage hybrid approach enhances the model's learning capabilities and leads to more accurate threat detection. The experimental results indicate that the proposed model achieves superior performance, with high recall, F1-score, precision, and accuracy. Moreover, a comparative analysis with existing literature further confirms these strong results. The findings indicate that combining anomaly detection with supervised learning can provide a more robust approach for enhancing IDS, particularly in demanding environments such as CC.

Keywords: Intrusion Detection System (IDS), Cloud Computing, Hybrid Model, XGBoost, Isolation Forest, Network Security, CICIDS2018

I. Introduction

In recent years, cloud computing (CC) has seen widespread and rapid adoption, fundamentally changing the way organizations manage, store, and access data and applications [IV]. However, among the primary barriers to cloud adoption, security [V] and compliance concerns stand out as the most significant, surpassing other challenges such as resource limitations and technical complexity. Issues such as data security and privacy, effective cloud security management, and timely threat detection and response remain at the forefront of these concerns, as highlighted in the Fortinet 2025 Cloud Report [I].

Intrusion Detection Systems (IDS) are paramount for maintaining the security status of cloud environments by monitoring system activity and network traffic continuously for malicious or suspicious behavior and identifying potential security breaches promptly, ensuring the integrity, confidentiality, and availability of cloud-based resources and data [XIX]. However, conventional IDS, reliant on predefined rules or signatures, often struggles to address dynamic and evolving attack patterns, such as novel or zero-day exploits. This limitation has led to the exploration and implementation of advanced methods, such as machine learning (ML), to strengthen the ability of IDS to function effectively in cloud environments. ML, as a powerful subset of artificial intelligence (AI), offers various techniques for intrusion detection, with anomaly detection and supervised learning being among the most promising in the field of intrusion detection. Anomaly detection is designed to detect departures from standard operational patterns, which can indicate previously unseen attacks, but it often suffers from high false positive rates and may lack the specificity to classify the exact type of detected threat. Supervised learning, on the contrary, entails training models using labeled data to classify network traffic or system behavior as either benign or malicious, but it struggles with detecting novel attacks for which labeled examples are unavailable. Moreover, there is a lack of publicly available, realistically labeled datasets.

To overcome these limitations, this study proposes a hybrid intrusion detection model that integrates Isolation Forest for anomaly detection with XGBoost for classification, leveraging the complementary strengths of both paradigms. Isolation Forest efficiently detects suspicious patterns without relying on labeled data, while XGBoost enhances the classification process, enhancing detection rates and minimizing false positives.

This research contributes to the advancement of intrusion detection in cloud environments by demonstrating the effectiveness of a novel hybrid machine learning approach through the integration of anomaly detection and supervised classification, rigorously evaluated and validated using the entirety of a challenging and realistic dataset to ensure a thorough assessment without data omissions, and through a detailed performance evaluation that further highlights the model's superiority over traditional and state-of-the-art AI-based intrusion detection systems in terms of detection capabilities.

The remainder of this paper is structured as follows: Section 2 reviews the relevant literature and discusses previous studies in the field. Section 3 describes the proposed

hybrid model, detailing its workflow design. Section 4 presents the main experimental results. Section 5 discusses the experimental results, including a comparative analysis with baseline, hybrid alternative method, and state-of-the-art models. Finally, Section 6 concludes the study, addressing key findings, limitations, and future research directions.

II. Related Work

In recent years, many security solutions have started combining AI technologies with security. IDS, as a critical component in safeguarding against emerging threats, has evolved from traditional types, which often rely on a signature database to detect known threats, to more dynamic approaches that employ machine learning and deep learning techniques to handle the growing volume of new and sophisticated threats [II, XX].

Numerous research studies have explored various approaches across different contexts and datasets. However, in this literature review, we highlight a selection of these various studies, with particular emphasis on hybrid solutions designed for cloud computing, and use the same dataset employed in our study.

Yang et al. [XXI] conducted a systematic literature review on anomaly-based network intrusion detection, highlighting the effectiveness of various ML methods in identifying unknown threats. Nonetheless, they also identified challenges such as high false positive rates and the need for heavy computational resources. In [IX], Zhao et al. addressed the limitations of IDSs by proposing a semi-supervised Discriminant Autoencoder (AUE), which combines Denoising Autoencoders with a heuristic method for class separation. As reported in their study, the model achieved 97.90% accuracy and 98.00% for both precision and recall.

Jafarian et al. [XVIII] addressed security challenges in SDN networks by proposing a multi-stage modular approach for anomaly detection and mitigation. Their proposed model shows improved accuracy (98.80%) and reduced false alarm rates (0.38%) compared to existing techniques. Despite the satisfactory performance demonstrated, certain constraints persist, including the limited generalizability of the proposed mechanism and the substantial challenges associated with scalability in centralized SDN architectures.

In the malware detection field, AI impact is also evident through the use of ensemble learning, Mayura et al.[XIV] proposed an ensemble model that combines three base classifiers Sequential model achieved an accuracy rate of 96.20%, distinguishing between malware and benign instances with minimal false alarms. The model's performance was improved through the use of Gray Wolf Optimization (GWO) to select the most relevant features. Despite these findings, their performance may be limited due to the use of relatively simple baseline classifiers.

Hybridization techniques in intrusion detection can be applied in various ways. In [XIII], Sajid et al. proposed a hybrid model combining ML and DL techniques, where feature extraction was performed using Extreme Gradient Boosting (XGBoost) and convolutional neural networks (CNN), while classification was handled by long short-term memory networks (LSTM). Their model was trained for both binary and

multiclass classification tasks using four datasets from different environments. The authors reported achieving high performance metrics across both classification types. However, the effectiveness of their approach was somewhat diluted by the broad focus across multiple datasets. Nevertheless, this choice was justified, as the primary objective was to demonstrate the general usefulness of hybrid approaches in various contexts. Similarly, Jaber and Rehman [III] proposed a hybrid mechanism based on the Fuzzy C-Means Clustering Algorithm with Support Vector Machine (FCM–SVM) to detect the four attack types present in the NSL-KDD dataset. Nevertheless, since their model was trained and tested only on NSL-KDD, its focus was limited to the four predefined attack categories: DoS, Probe, R2L (Remote-to-Local), and U2R (User-to-Root), which may not fully represent the diversity and complexity of real-world threats. Regardless of this, the model achieved accuracy rates of 97.37% or higher across all attack types.

Another relevant study by Qazi et al. [VII] proposed a hybrid intrusion detection framework based on DL using a Convolutional Recurrent Neural Network (CRNN). This framework integrates RNN with CNN, wherein two convolutional layers are followed by multiple RNN layers. The output is then passed into fully connected, flattened, and SoftMax layers, thereby enhancing the model's capacity to detect and classify traffic. Their experiments were conducted on the CICIDS-2018 dataset, achieving an average accuracy of 98.90% along with promising results in terms of F-measure, precision, and recall.

Several studies address security issues in CC, intending to enhance intrusion detection system performance. Alzughaibi and El Khediri developed two models based on Deep Neural Network (DNN). The first model was built on a multi-layer perceptron (MLP) with backpropagation (BP), while the second was trained by an MLP with particle swarm optimization (PSO). These models are employed to address binary and multi-class classification using the updated cybersecurity CSE-CIC-IDS2018 dataset. The best accuracy for binary classification was 98.97%, and for multiclass classification, it was 98.41%, both based on MLP-BP [XVI]. Likewise, the authors in [XX] proposed a DL-based model for CC that utilizes advanced CNNs. To address the issue of class imbalance, the model incorporates the SMOTE technique, while a Dropout layer is integrated to prevent overfitting, leading to effective cyberattack detection. Their method addresses the challenges in analyzing large-scale traffic in CC, achieving over 98.67% in accuracy, precision, and recall metrics. In [X], Attou et al. present an innovative IDS model using DL algorithms, notably Radial Basis Function Neural Network (RBFNN) and Random Forest (RF). RF was used for feature selection, while the RBFNN algorithm was used to detect intrusion in CC. The potential of their proposed IDS model lies in achieving high accuracy using minimal features. However, they used two datasets (Bot-IoT and NSL-KDD) for validation, which may not fully represent the complexities of real-world CC environments. Furthermore, Al-Fawa'reh et al. [XII] proposed a hybrid model integrating Principal Component Analysis (PCA) with a Deep Neural Network (DNN) to identify abnormal network traffic patterns using the CSE-CIC-IDS2018 dataset, achieving enhanced accuracy and reduced complexity. Acknowledging resource limitations, they excluded data from the fourth day (20-02-2018), which

contained high-volume DDoS attacks, limiting the evaluation of their model against those specific attack types.

III. Methods

This section describes the design of our model, as illustrated in Fig. 1. The process begins with data collection, followed by extensive preprocessing to ensure the dataset is clean, robust, and of high quality. This consolidation and refinement process results in a unified, high-quality version suitable for analysis. After performing binary labeling, the data is split for training and testing. Next, our method is applied, which combines anomaly detection using Isolation Forest with a supervised learning model based on XGBoost. In the final stage, we evaluate our model using the most representative metrics.

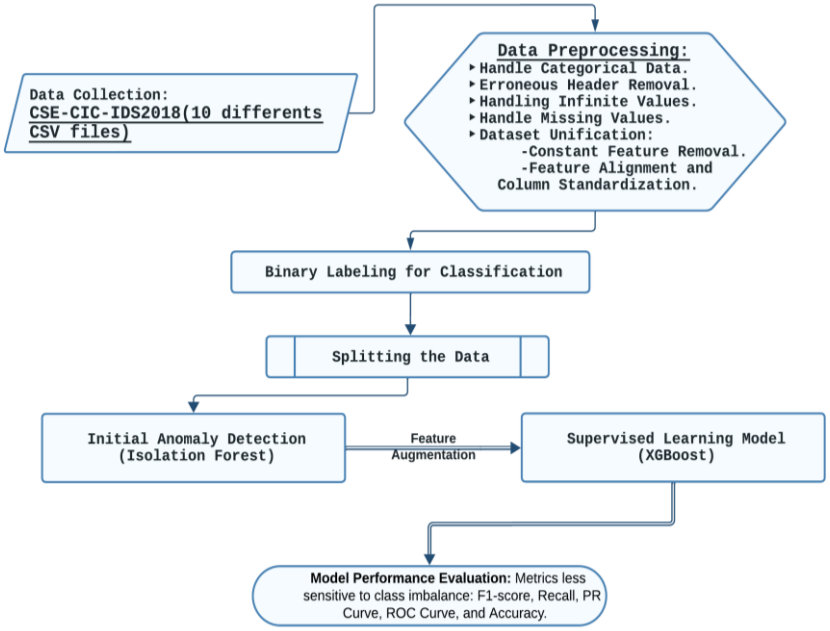

Fig. 1. Schematic design of the proposed system

Dataset

The CSE-CIC-IDS2018 dataset is an essential reference in the field of intrusion detection, developed in collaboration between the Centre for Cybersecurity Excellence (CSE) and the Canadian Institute for Cybersecurity (CIC) [XI]. The dataset was generated on the Amazon Web Services (AWS) cloud environment, which makes it suitable for our research objectives since it provides a realistic and diverse collection of network traffic, including both normal activity and modern cyberattacks. Simulated attacks include brute force, denial of service, infiltration, injection, and bot attacks Fig. 2, with a total of 16,233,002 records distributed across seven scenario days, as detailed in Table 1.

Table 1: CSE-CIC-IDS2018 Network traffic.

Scenario/ Day	Date	Primary Attack Theme(s)	Total Records (Flows)
Day 1	Wed, Feb 14, 2018	Brute Force Attacks	1,048,575
Day 2	Thu, Feb 15, 2018	DoS Attacks	1,048,575
Day 3	Fri, Feb 16, 2018	DDoS LOIC-HTTP, Heartbleed	1,048,575
Day 4	Tue, Feb 20, 2018	DDoS LOIC-UDP & HOIC	7,948,748
Day 5	Wed, Feb 21, 2018	Brute Force (FTP, SSH)	1,048,575
Day 6	Thu, Feb 22, 2018	Web Attacks (Brute Force, XSS, SQLi)	1,048,575
Day 7	Fri, Feb 23, 2018	Web Attacks (Brute Force, XSS, SQLi)	1,048,575
Day 8	Wed, Feb 28, 2018	Infiltration	613,104
Day 9	Thu, Mar 01, 2018	Infiltration	331,125
Day 10	Fri, Mar 02, 2018	Botnet (Ares)	1,048,575

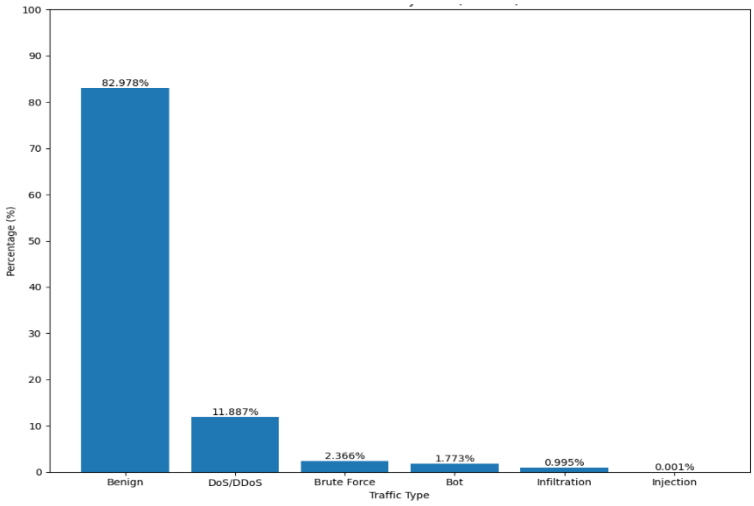

Fig. 2. Traffic distribution by class.

CIC-IDS2018 includes network flows rich in information, and its widespread adoption by the scientific community testifies to its relevance to cybersecurity research.

Preprocessing

The dataset comes in the form of 10 CSV files, and the first data pre-processing step was the clean-up of header lines that weren't parsed correctly and were erroneously included in the dataset.

The next step was handling infinities by dropping them, as they appear only in normal traffic, which is already voluminous, to avoid disruption in the model training or

evaluation process. Following this, we proceed with handling missing values. This process of dataset integrity verification reduces the total records to 16,137,183 records with different numbers of features in each file.

The subsequent operations focus on features to unify the dataset files' shape to make them suitable for combining into a single dataset version, including all the benign and attack traffic; initially, we drop single-value columns, which are non-informative and do not contribute to the model learning. Secondly, we perform feature alignment and column standardization to ensure files have a consistent set of features, making them suitable for combining into one comprehensive dataset.

Following the execution of these pre-processing steps, we encode the target variable into binary values, where normal traffic is labeled as 0 and anomalous traffic as 1. This binary labeling ensures a consistent format for classification and aligns with our objective of detecting intrusions rather than identifying specific attack types, allowing the model to generalize across different malicious activities. Thereafter, the dataset was split using the 80/20 rule, where 80% of the data is used for training the model and the remaining 20% is used for testing to evaluate its performance.

Model Building

Baseline Model:

To establish a meaningful reference for evaluating the proposed hybrid architecture, we implemented two types of baseline models.

***Logistic Regression(LR)*:** In our proposed method, we first established a baseline for comparison. This decision was motivated by the challenge of finding previous research conducted under the same specific conditions as our study, where the entire dataset, including all attack types, was combined into a unified version and subjected to essential preprocessing steps. Given these unique conditions, we selected Logistic Regression (LR) as the baseline model due to its simplicity, interpretability, and widespread use in the intrusion detection field, as well as its role as a benchmark in binary classification tasks [VI]. The model was trained on 80% of the data and tested on the residual 20%, with the default hyperparameters.

***Alternative Hybrid Architecture (K-Means + Logistic Regression)*:** To validate the selection of the proposed IF + XGBoost architecture, we compared it to an alternative hybrid model architecture. In this alternative approach, K-means clustering is combined with LR. The process runs in two steps. First, K-means [XV] is used on the training data to form clusters. Then, for each data point, we measure how far it sits from the nearest cluster center and add that distance as an extra feature in the dataset. Secondly, an LR classifier is trained on this enriched dataset to handle the actual classification.

Anomaly Detection-Based Classification: Isolation Forest

Datasets in the intrusion detection field are typically imbalanced, due to the real-world traffic where normal is predominant, and also because of attack diversity and frequency, where several types of attacks are rarer than others, and certain sophisticated attacks are often challenging to collect. This imbalance causes models

to be biased in favor of the dominant class and struggle to detect less represented classes. It also produces misleading performance metrics. Several strategies have been adopted to address this issue, including resampling techniques, synthetic data generation, and anomaly detection models that focus exclusively on learning the majority class, facilitating the detection of rare classes since they will appear anomalous relative to the Benign majority. Isolation Forest, as a key technique for anomaly detection, is relatively fast and scalable for large datasets like our chosen dataset, and has the capabilities to detect new and previously unseen attacks, provided they deviate significantly from the learned Benign behavior [VIII].

We implemented the Isolation Forest, training it exclusively on the benign samples from the 80% training split. After several trials and adjustments, we set the number of estimators to 100 and used a dynamically determined contamination rate to ensure the model effectively distinguishes between normal and anomalous traffic. The contamination rate, defined as the proportion of anomalies in the dataset, was calculated as:

$$contaminiation = \min\left(0.5, \frac{|\ Anomalous\ Samples\ |}{|\ Total\ Test\ Sample\ |}\right) \quad (1)$$

Supervised Learning-Based Classification: Extreme Gradient Boosting (XGBoost)

XGBoost is a supervised machine learning algorithm that is regarded as a scalable gradient tree boosting system known for its speed and performance in solving classification problems using a minimal amount of resources in comparison to other classification and regression algorithms [XVII]. It employs a boosting methodology, constructing models sequentially, with each subsequent model focusing on correcting the errors made by its predecessor. The iterative refinement employed in this algorithm is instrumental in enhancing the ensemble's predictive capabilities through a systematic and incremental process.

We implemented the XGBoost on 80% of the dataset for training, ensuring exposure to a diverse range of benign and attack samples, and 20% for testing. We used a combination of default and fine-tuned hyperparameters, determined through a trial-and-error approach to balance between performance metrics. The number of estimators was set to 100, the maximum depth to 5, and the learning rate to 0.1.

Hybrid Approach: Combining Anomaly Detection with Supervised Learning

To enhance threat detection, we propose a hybrid approach that leverages both anomaly detection and supervised learning. In the first stage, we applied the Isolation Forest model to identify anomalies in the dataset based on its ability to isolate outliers efficiently. The predictions made by the model, where instances are labeled as normal or anomalous, are then introduced as an additional feature in the dataset. In the second stage, we use this enriched dataset subsequently to train the XGBoost model and benefit from both the raw feature set and the anomaly-based insights provided by the first one. The integration of these two models is intended to improve the overall critical performance metrics by leveraging the strengths of both unsupervised anomaly detection and supervised learning.

Evaluation metrics

Efficiency metrics selection is crucial for evaluating and validating our proposed hybrid model. Therefore, we employed several relevant performance metrics calculated from the confusion matrix (CM), including Recall, Precision, and Accuracy, in addition to ROC and PR Curves.

Confusion Matrix: A visualization component that provides a detailed summary of the model's predictions by displaying true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). This allows us to analyze misclassification patterns and evaluate model performance at a granular level.

Recall: Or sensitivity, a highly significant metric, particularly when dealing with imbalanced datasets, as it ensures the detection of the minority class. Recall computes the proportion of correctly classified anomalous instances out of all actual intrusions.

$$Recall = \frac{TP}{TP + FN} \tag{2}$$

Precision: Measures the correctly classified abnormal(attack) instances among all instances classified as abnormal, and remains a very important metric since false positives can be costly in the cybersecurity field.

$$Precision = \frac{TP}{TP + FP} \tag{3}$$

Accuracy: Represents the overall correctness of the model's predictions, which are defined as:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN} \tag{4}$$

***Receiver Operating Characteristic (ROC)*:** Illustrates the trade-off between the true positive rate (recall) and the false positive rate (FPR) at various classification thresholds. FPR is defined as:

$$FPR = \frac{FP}{FP + TN} \tag{5}$$

ROC provides an area under the curve (AUC) score, which helps assess the model's capacity to discriminate between normal and anomalous traffic.

***Precision-Recall (PR) Curve*:** Since intrusion detection often involves highly imbalanced datasets, we also employ the PR curve as an additional evaluation metric. This curve illustrates the trade-off between precision and recall across different classification thresholds, providing deeper insight into the model's effectiveness.

IV. Results

LR Baseline Model Performance

The results for our baseline model, a Logistic Regression classifier, are presented in the table below:

Table 2: Performance Metrics of the LR Model

	Precision	Recall	F1-Score	Accuracy
Anomaly	0.8566	0.7550	0.8026	0.9368
Benign	0.9509	0.9741	0.9624	
Macro Avg	0.9038	0.8645	0.8825	
Weighted avg	0.9349	0.9368	0.9352	

The results show a high accuracy of 93.68%, with strong performance in detecting benign traffic (Precision = 95.09%, Recall = 97.41%). However, the model exhibits a lower recall for anomaly detection (75.50%), indicating that a significant proportion of anomalous instances are misclassified.

Fig. 3 shows the confusion matrix for the model we examine, including both benign and anomaly classes.

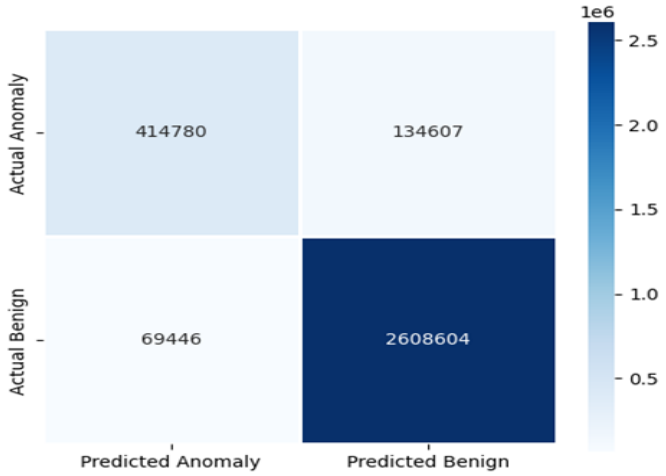

Fig. 3. Confusion Matrix of the LR Model

The CM reveals a significant amount of false negatives, which is critical, and significant false positives, which may result in unnecessary security alerts.

Alternative Hybrid Architecture (K-Means + LR) Performance

The performance of the alternative K-means + Logistic Regression model was evaluated to provide a comparative benchmark. The key performance metrics for this model are presented in Table 3.

Table 3: Performance Metrics of the alternative hybrid model

	Precision	Recall	F1-Score	Accuracy
Anomaly	0.8822	0.7610	0.8171	0.9420
Benign	0.9523	0.9792	0.9656	
Macro Avg	0.9172	0.8701	0.8913	
Weighted avg	0.9404	0.9420	0.9403	

The K-Means + LR hybrid shows slight improvements in both accuracy (94.20% vs. 93.68%) and the F1 score (81.71% vs. 80.26%) for anomaly detection. Recall for anomalies also increased marginally (76.10% vs. 75.50%), indicating that the hybrid approach identifies more malicious instances while maintaining strong performance on benign traffic. These results highlight the benefit of using unsupervised learning as a feature engineering step before classification.

Fig. 4 shows the confusion matrix for the alternative K-Means + LR model, detailing its performance on both benign and anomaly classes.

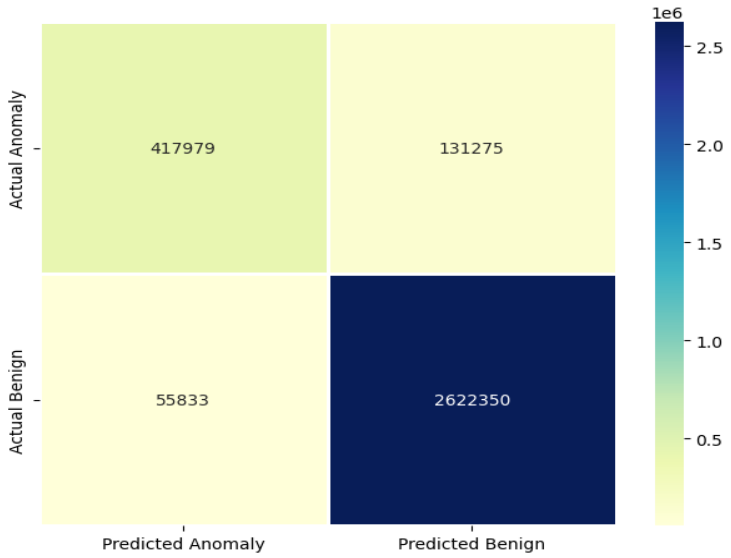

Fig. 4. Confusion Matrix of the K-Means + LR Model

The CM shows a promising decrease in false positives, which leads to fewer false alarms. However, the model's critical weakness persists; It still produces over 131,000 false negatives, leaving a significant volume of threats undetected.

Hybrid Model Performance (Isolation Forest + XGBoost)

Table 4 demonstrates the key efficiency evaluation metrics of our novel method.

Table 4: Hybrid Model Performance (Isolation Forest + XGBoost)

Metric	Value
Precision	0.9895
Recall	0.9995
F1-Score	0.9945
Accuracy	0.9908
False Positive Rate (FPR)	0.0516
False Negative Rate (FNR)	0.0005
ROC-AUC Score	0.9950
PR-AUC Score	0.9989

Compared to the baseline and the alternative hybrid architecture model, our proposed hybrid approach achieves higher precision, recall, and F1-score. This demonstrates its ability to minimize false negatives while maintaining a low false positive rate.

These results highlight the effectiveness of the proposed hybrid approach, particularly in minimizing false negatives, which is a critical consideration in intrusion detection scenarios.

Precision-Recall and ROC Curves

Figures 5 and 6 show the Precision-Recall (PR) curve and the Receiver Operating Characteristic (ROC) curve, respectively, and offer further insight into the model's robustness.

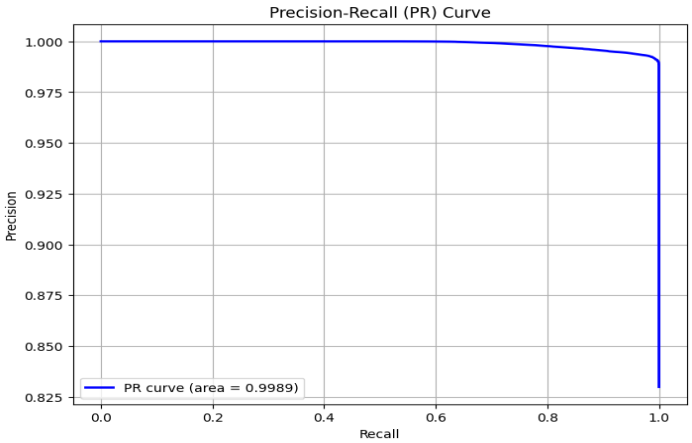

Fig. 5: PR curve of the proposed model

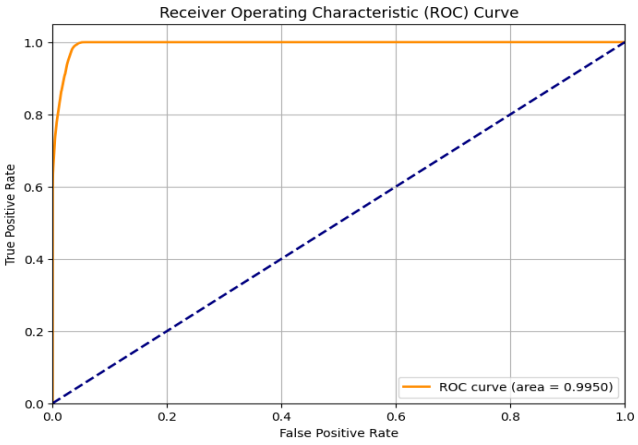

Fig. 6. ROC Curve of the proposed model

The Precision-Recall (PR) curve, presented in Fig. 5, assesses the performance of the classification model across different threshold settings, particularly emphasizing its ability to handle imbalanced datasets. The plot demonstrates a strong performance, maintaining Precision across a wide range of Recall before declining slightly as

Recall approaches 1.0. With an area under the curve equal to 0.9989 as reported, this indicates exceptional model performance in balancing precision and recall, especially for positive class identification. This high AUC confirms that the hybrid model is powerful and reliable at classification, even when positive instances are potentially rare.

The Receiver Operating Characteristic (ROC) curve, shown in Fig. 6, illustrates the relationship between the True Positive Rate (TPR) and False Positive Rate (FPR). The hybrid model achieves an ROC-AUC score of 0.9950, indicating an excellent discriminatory ability between benign and anomalous instances. The curve's proximity to the top-left corner indicates that the model effectively reduces false positives while maintaining a high true positive rate. In comparison to a random classifier (illustrated by a dashed diagonal line), the hybrid model demonstrates significantly superior classification performance.

These improvements indicate that the hybrid approach successfully enhances intrusion detection performance, surpassing the baseline and alternative hybrid models in key evaluation metrics.

V. Discussion

The experimental results obtained demonstrate the viability of our proposed innovative hybrid model for IDS in cloud environments. With its exceptional ability to learn from and process large-scale, high-dimensional, and imbalanced data without manual intervention, the hybrid methodology shows a promising pivotal role in identifying subtle signs of cyberattacks and highlighting its potential in enhancing Intrusion Detection Systems (IDS) performance. This section covers the key findings, practical implications, comparative performance, and potential limitations of the model.

Key Findings and Interpretation

The results demonstrate that the hybrid model excels at detecting intrusion problems. The high recall (0.9995) ensures that almost all types of attacks are detected, and the precision (0.9895) indicates that the model effectively detects anomalies while minimizing false alarms. The low false negative rate (FNR = 0.0005) ensures that very few malicious activities go undetected, which is crucial for real-world intrusion detection applications. Additionally, the false positive rate (FPR = 0.0516) is significantly reduced compared to traditional supervised learning approaches, demonstrating the hybrid model’s ability to differentiate between normal and anomalous traffic more reliably.

The ROC-AUC (0.9950) and PR-AUC (0.9989) values further validate the robustness of the model, suggesting that it maintains high discriminative performance across different classification thresholds. These metrics are particularly important in intrusion detection, where the class distribution is often imbalanced, and maintaining high recall is a priority.

The combination of Isolation Forest (IF) and XGBoost provides the observed performance by leveraging their complementary strengths. IF, as an unsupervised

anomaly detection method, efficiently isolates anomalies using randomly generated decision trees. Because anomalous instances tend to be isolated earlier, they have shorter average path lengths. This feature enhances the model's ability to effectively detect outliers. When integrated with XGBoost, a powerful supervised learning algorithm, the model benefits from improved feature learning and classification capabilities, leading to more accurate predictions. The synergy between these two techniques explains the results achieved, particularly in handling unbalanced data and identifying subtle attack patterns.

To comprehensively evaluate the model's behavior across diverse attack categories, we conducted a per-class analysis using detailed performance metrics (Table 5) alongside a confusion matrix visualization (Fig. 7). This evaluation reveals a distinct behavioral pattern: the model demonstrates proficiency in detecting high-volume, network-centric attacks, but struggles significantly with stealthy, under-represented threats.

Table 5: Per-Class detailed performance metrics

Attack_Type	Recall	Precision	F1-Score	Support
DDOS attack-LOIC-UDP	1.000000	1.0	1.000000	334
DDOS attack-HOIC	1.000000	1.0	1.000000	136869
FTP-BruteForce	1.000000	1.0	1.000000	38810
DoS attacks-Slowloris	1.000000	1.0	1.000000	2226
DoS attacks-SlowHTTPTest	1.000000	1.0	1.000000	27856
DoS attacks-Hulk	1.000000	1.0	1.000000	92445
SSH-Bruteforce	1.000000	1.0	1.000000	37449
DoS attacks-GoldenEye	0.999880	1.0	0.999940	8300
DDoS attacks-LOIC-HTTP	0.999800	1.0	0.999900	115255
Bot	0.999720	1.0	0.999860	57195
Brute Force -XSS	0.902439	1.0	0.948718	41
Brute Force -Web	0.787402	1.0	0.881057	127
SQL Injection	0.266667	1.0	0.421053	15
Infilteration	0.242361	1.0	0.390161	32332

As shown in Table 5, the model achieves nearly perfect recall for volumetric attacks such as DoS/DDoS, service-level brute force (FTP, SSH), and Botnet traffic. It correctly identifies almost 100% of these instances. Interestingly, this strong performance extends even to some classes with limited samples but distinctive signatures, such as Brute Force-XSS (90% recall on 41 samples). This suggests that an attack's pattern can be more critical than its frequency, in contrast to low-and-slow (Infiltration) or application-layer (SQL Injection) attacks that are specifically designed to blend in with normal traffic.

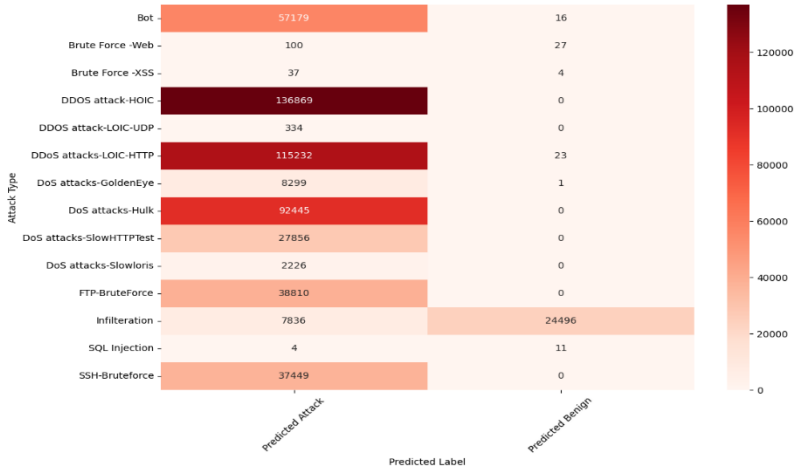

Fig. 7. Per-Attack-Type Confusion Matrix Heatmap

However, as the confusion matrix in Fig. 7 illustrates, the model's primary weakness lies in its inability to detect stealthy attacks that mimic benign traffic. The most significant failure is in identifying 'Infiltration' attacks, where 75% of instances (24,496 out of 32,332) were misclassified as benign. Similarly, the model missed 73% of 'SQL Injection' attacks (11 out of 15). This behavior highlights a classic precision-recall tradeoff; the model's 1.0 precision on these rare classes means it is highly conservative, and when it does flag an attack, it is correct. The consequence, however, is a critically low recall, as the model prefers to misclassify these subtle threats as benign. This analysis underscores the model's limitations and establishes a clear direction for future work focused on enhancing sensitivity to these rare but highly dangerous attack vectors.

Comparative Analysis and Summary of Results

The evaluation results confirm that the proposed hybrid approach significantly enhances intrusion detection performance. To further evaluate our approach and better understand performance gains, we compared our hybrid model with a baseline model, an alternative hybrid architecture, and the reported results from existing state-of-the-art methods, and also to get an impression of how challenging it is to detect attacks in such an environment. Table 6 summarizes the comparative analysis and highlights the key differences.

Table 6: Comparison of our proposed method and other state-of-the-art methods

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
[VII]	98.90	98.64	99.15	99.03
[XVI]	98.97	99.98	98.80	99.38
[IX]	97.90	98.00	98.00	98.00
[XII]	97.93	99.97	97.42	98.68
[XX]	98.67	98.80	99.68	98.79
Baseline	93.68	93.49	93.68	93.52
Alternative Hybrid model	94.20	95.23	97.92	96.56
Our model	99.08	98.95	99.95	99.45

Our proposed hybrid model clearly outperforms both the baseline, the alternative hybrid model and the recent studies summarized in the table above. This comparison demonstrates the effectiveness of integrating anomaly detection techniques, such as Isolation Forest, with supervised learning algorithms such as XGBoost as an advanced architecture. Through this combination, the model can effectively capture complex patterns in cloud network traffic, achieving greater accuracy and robustness. The proposed approach achieves an accuracy of 99,08%, outperforming all other listed models trained on the same CSE-CICIDS2018 dataset. Moreover, it maintains competitive performance across a range of key indicators such as recall, F1 score, and AUC values, making it a highly reliable solution for modern intrusion detection systems. These results highlight the potential of our hybrid method as a scalable and effective alternative for securing cloud computing environments.

VI. Conclusion

This study introduced an innovative hybrid anomaly detection model by integrating Isolation Forest (IF) with XGBoost, combining the distinct benefits of both unsupervised and supervised learning approaches for intrusion detection in a robust, rigorously preprocessed, and unified cybersecurity dataset. The experimental evaluation demonstrates that the proposed IF + XGBoost significantly outperforms traditional supervised approaches in terms of all pertinent metrics, which suggests that pre-filtering anomalies using Isolation Forest before classification enhances detection accuracy and reinforces the importance of hybrid machine learning frameworks in handling class imbalance and improving anomaly detection robustness.

Despite the strong performance achieved in this study, several limitations and challenges remain. These include the necessity for meticulous parameter tuning for both models, the computational overhead due to the two-step learning process, and sensitivity to data distribution shifts. Moreover, while the proposed model achieves excellent performance on frequent, network-centric attacks, its limited sensitivity to application-layer and rare attacks highlights an important direction for future research. Enhancing detection capabilities at these higher layers will be crucial to achieving more comprehensive intrusion detection. Future work will therefore be directed toward improving efficiency and interpretability, while also testing the model on more diverse and dynamic datasets to ensure practical deployment and validate its generalizability in real-world cybersecurity applications.

Conflict of Interest:

There was no relevant conflict of interest regarding this paper.

References

I. “2025-Cloud-Security-Report-Fortinet.”

II. A. Khraisat, I. Gondal, P. Vamplew, and J. Kamruzzaman, "Survey of intrusion detection systems: techniques, datasets and challenges," Cybersecurity, vol. 2, no. 1, p. 20, Jul. 2019. 10.1186/s42400-019-0038-7.

III. A. N. Jaber and S. U. Rehman, "FCM–SVM based intrusion detection system for cloud computing environment," Clust. Comput., vol. 23, no. 4, pp. 3221–3231, Dec. 2020. 10.1007/s10586-020-03082-6.

IV. A. Singh and K. Chatterjee, "Cloud security issues and challenges: A survey," J. Netw. Comput. Appl., vol. 79, pp. 88–115, Feb. 2017. 10.1016/j.jnca.2016.11.027.

V. C. Modi, D. Patel, B. Borisaniya, H. Patel, A. Patel, and M. Rajarajan, "A survey of intrusion detection techniques in Cloud," J. Netw. Comput. Appl., vol. 36, no. 1, pp. 42–57, Jan. 2013. 10.1016/j.jnca.2012.05.003.

VI. E. Besharati, M. Naderan, and E. Namjoo, "LR-HIDS: logistic regression host-based intrusion detection system for cloud environments," J. Ambient Intell. Humaniz. Comput., vol. 10, no. 9, pp. 3669–3692, Sep. 2019. 10.1007/s12652-018-1093-8.

VII. E. U. H. Qazi, M. H. Faheem, and T. Zia, "HDLNIDS: Hybrid Deep-Learning-Based Network Intrusion Detection System," Appl. Sci., vol. 13, no. 8, p. 4921, Apr. 2023. 10.3390/app13084921.

VIII. F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation-Based Anomaly Detection," ACM Trans. Knowl. Discov. Data, vol. 6, no. 1, pp. 1–39, Mar. 2012. 10.1145/2133360.2133363.

IX. F. Zhao, H. Zhang, J. Peng, X. Zhuang, and S.-G. Na, "A semi-self-taught network intrusion detection system," Neural Comput. Appl., vol. 32, no. 23, pp. 17169–17179, Dec. 2020. 10.1007/s00521-020-04914-7.

X. H. Attou et al., "Towards an Intelligent Intrusion Detection System to Detect Malicious Activities in Cloud Computing," Appl. Sci., vol. 13, no. 17, p. 9588, Aug. 2023. 10.3390/app13179588.

XI. I. Sharafaldin, A. Habibi Lashkari, and A. A. Ghorbani, "Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization:," in Proceedings of the 4th International Conference on Information Systems Security and Privacy, Funchal, Madeira, Portugal: SCITEPRESS - Science and Technology Publications, 2018, pp. 108–116. 10.5220/0006639801080116.

XII. M. Al-Fawa'reh, M. Al-Fayoumi, S. Nashwan, and S. Fraihat, "Cyber threat intelligence using PCA-DNN model to detect abnormal network behavior," Egypt. Inform. J., vol. 23, no. 2, pp. 173–185, Jul. 2022. 10.1016/j.eij.2021.12.001.

XIII. M. Sajid et al., "Enhancing intrusion detection: a hybrid machine and deep learning approach," J. Cloud Comput., vol. 13, no. 1, p. 123, Jul. 2024. 10.1186/s13677-024-00685-x.

XIV. Mayura V. Shelke, Jyoti Yogesh Deshmukh, Deepika Amol Ajalkar, and R. B. Dhumal, "A Robust Ensemble Learning Approach for Malware Detection and Classification," J. Adv. Res. Appl. Sci. Eng. Technol., vol. 48, no. 1, pp. 152–167, Jul. 2024. 10.37934/araset.48.1.152167.

XV. P. Fränti and S. Sieranoja, "K-means properties on six clustering benchmark datasets," Appl Intell, vol. 48, no. 12, pp. 4743–4759, Dec. 2018. 10.1007/s10489-018-1238-7

XVI. S. Alzughaibi and S. El Khediri, "A Cloud Intrusion Detection Systems Based on DNN Using Backpropagation and PSO on the CSE-CIC-IDS2018 Dataset," Appl. Sci., vol. 13, no. 4, p. 2276, Feb. 2023. 10.3390/app13042276.

XVII. T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco California USA: ACM, Aug. 2016, pp. 785–794. 10.1145/2939672.2939785.

XVIII. T. Jafarian, A. Ghaffari, A. Seyfollahi, and B. Arasteh, "Detecting and mitigating security anomalies in Software-Defined Networking (SDN) using Gradient-Boosted Trees and Floodlight Controller characteristics," Comput. Stand. Interfaces, vol. 91, p. 103871, Jan. 2025. 10.1016/j.csi.2024.103871.

XIX. T. Sowmya and E. A. Mary Anita, "A comprehensive review of AI based intrusion detection system," Meas. Sens., vol. 28, p. 100827, Aug. 2023. 10.1016/j.measen.2023.100827.

XX. W. H. Aljuaid and S. S. Alshamrani, "A Deep Learning Approach for Intrusion Detection Systems in Cloud Computing Environments," Appl. Sci., vol. 14, no. 13, p. 5381, Jun. 2024. 10.3390/app14135381.

XXI. Z. Yang et al., "A systematic literature review of methods and datasets for anomaly-based network intrusion detection," Comput. Secur., vol. 116, pp. 102675, May 2022. 10.1016/j.cose.2022.102675.