



AI-Driven Intrusion Detection and Mitigation Techniques for Cybersecurity in Cloud Computing

Smitha GV¹, Samitha Khaiyum², T Ratha Jeyalakshmi³

¹ Research Scholar, Department of MCA, Dayananda Sagar College of Engineering, gvsmitha3@gmail.com

² Prof & Head, Department of MCA, Dayananda Sagar College of Engineering, samitha-mca@dayanandasagar.edu

³ Associate Professor, Department of MCA, Dayananda Sagar College of Engineering, ratha-mcavtu@dayanandasagar.edu

Abstract

Cloud computing has revolutionized the way IT infrastructure is managed in the modern world, offering services that are highly scalable, flexible, and cost-effective. Nevertheless, the everchanging nature of cloud computing poses major cybersecurity threats, ranging from data breaches, insider threats, DDoS attacks, and advanced persistent threats. Conventional cybersecurity systems cannot effectively counter the ever-changing nature of cybersecurity threats. Artificial Intelligence (AI) technology-based Intrusion Detection System (IDS) and mitigation techniques have been developed to address the issue of cybersecurity threats in cloud computing environments. This paper will discuss the role of AI in intrusion detection and mitigation in cloud computing environments, the various techniques used in IDS, the various challenges associated with AI-based IDS, and finally propose a framework for intelligent cloud computing management systems.

Introduction

Cloud computing has developed and become one of the most influential technological revolutions of the 21st century, transforming the way organizations are using and managing their computer systems. Today, organizations are no longer relying on traditional computer systems, but are increasingly adopting cloud-based computing systems, which provide on-demand computing services, storage, and networking with application services [1]. Organizations like Amazon Web Services, Microsoft Azure, and Google Cloud Platform have helped organizations shift their critical computer systems to a decentralized, virtualized, and scalable environment [1]. This has led to the widespread adoption of digital transformation in industries like healthcare, finance, education, manufacturing, and government services [2]...

Though there are many advantages of using cloud computing, there are many security concerns that need to be addressed. Unlike traditional data centers, where the infrastructure is under physical control and management, the infrastructure of the cloud is virtual and accessible through public networks [3]. The increased complexity of the system and the attack surface are due to the distributed nature of the data centers, API management, containerization, tenancy, and elasticity of the resources [4]. With the increased adoption of the cloud computing model by many companies, attackers are also focusing on exploiting the vulnerabilities of the system. Cybersecurity risks associated with cloud computing are more sophisticated, automated, and flexible than straightforward extensions of traditional network attacks [5].

Attackers use poorly configured storage buckets, weak authentication protocols, dangerous APIs, and virtualization weaknesses to carry out unauthorized access [3]. Distributed Denial of Service (DDoS) attacks against cloud-based services cause availability-related problems [4]. Advanced Persistent Threats (APTs) carry out unauthorized access to steal sensitive data over a period of time [6]. Insider threats also pose a risk to cloud computing security, both intentional and unintentional [7]. As attacks can bypass signature-based systems due to zero-day attacks, traditional intrusion detection systems cannot fully provide security to cloud computing environments [8].

The integration of AI in cloud security is considered to be a major paradigm shift from reactive security to proactive and predictive security [9]. Rather than reacting to detected threats, AI can predict threats by analyzing patterns and indicators of potential attacks [10]. Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), and Long Short-Term Memory (LSTM) networks are deep learning architectures that can model high-dimensional network traffic, especially in multi-stage cyberattacks [11]. Additionally, graph-based neural networks can identify threats in cloud infrastructure modeled as interconnected nodes [12].

With the expansion of cloud computing services worldwide, securing distributed infrastructure against highly evolved cyberattacks has become a pressing need. Intrusion detection and mitigation using artificial intelligence can be a highly effective strategy to improve cloud infrastructure security [13]. Advanced analytics, coupled with automated response systems, can transform cloud security tools from passive detection systems into active defense systems capable of countering sophisticated cyber threats effectively [14].

Mathematical framework and model design.

For the design of an AI-based intrusion detection and mitigation system in the context of cloud computing, there is a need to have a well-structured mathematical foundation that is capable of dealing with the dynamic nature of the modern cloud computing infrastructure. The modern cloud computing infrastructure comprises virtual machines, containers, microservices, APIs, storage, users, and communication that operate through distributed networks. Since the cloud computing infrastructure scales up or down based on the dynamic nature of the workload, the cloud computing infrastructure must be treated as a time-varying system. Mathematically, the infrastructure of the cloud computing model may be described as follows: the dynamic structure of the cloud computing infrastructure is composed of computational nodes, communication links, active users, and service instances, which vary dynamically over time.

$$C(t) = \{ V(t), E(t), U(t), S(t) \}$$

The equation $C(t) = \{ V(t), E(t), U(t), S(t) \}$ simply describes a cloud system at a particular time t . It includes virtual machines, network connections, users, and services that are all working together in the cloud. Since these components keep changing over time—like users logging in or resources scaling up—the cloud environment is always dynamic. This makes it harder to detect security threats and highlights the need for intelligent, adaptive security systems.

$$\begin{aligned} L &= - (1/N) \sum_i \sum_k y_{ik} \log(\hat{y}_{ik}) \\ &= - (1/N) [y_1 \log(\hat{y}_1) + y_2 \log(\hat{y}_2) + \dots + y_k \log(\hat{y}_k)] \end{aligned}$$

The equation $L = - (1/N) \sum y \log(\hat{y})$ represents the loss function used to train the model. It measures how close the model's predictions are to the actual values. If the predicted values are accurate, the loss will be low; if they are incorrect, the loss increases. In simple terms, this equation helps the model learn from its mistakes and improve its predictions over time.

$$\begin{aligned} D(x) &= (x - \mu)^T \Sigma^{-1} (x - \mu) \\ &= \sqrt{ [(x - \mu)^T \Sigma^{-1} (x - \mu)] } \end{aligned}$$

The equation $D(x) = (x - \mu)^T \Sigma^{-1} (x - \mu)$ is used to measure how different a data point is from normal behavior. It compares the current activity with the average pattern (μ) while also considering how the features are related to each other (Σ). If the value of $D(x)$ is small, the behavior is normal; if it is large, it may indicate an anomaly or possible attack. In simple terms, this equation helps identify unusual activities in the system.

$$\begin{aligned} R(x) &= || x - \hat{x} ||^2 \\ &= (x_1 - \hat{x}_1)^2 + (x_2 - \hat{x}_2)^2 + \dots + (x_n - \hat{x}_n)^2 \end{aligned}$$

The equation $R(x) = || x - \hat{x} ||^2$ shows how much the reconstructed data differs from the original data. In simple terms, it checks how well the model can recreate the input. If the difference is small, it means the behavior is normal because the model recognizes it. But if the difference is large, it suggests something unusual is happening, which could be a sign of an intrusion or anomaly in the system.

$$\begin{aligned} J(\pi) &= E [\sum_t \gamma^t R(s_t, a_t)] \\ &= R_0 + \gamma R_1 + \gamma^2 R_2 + \dots \end{aligned}$$

The equation $J(\pi) = E [\sum \gamma^t R(s_t, a_t)]$ represents the goal of the reinforcement learning model. It calculates the total reward the system can achieve over time by taking different actions. The expanded form shows that future rewards are slightly reduced using the factor γ , meaning immediate rewards are more important than distant ones. In simple terms, this equation helps the system learn the best actions to take so that it can maximize overall security and respond effectively to threats.

$$\begin{aligned} \text{Precision} &= TP / (TP + FP) \\ &= \text{Correct Attacks} / \text{Total Predicted Attacks} \\ \text{Recall} &= TP / (TP + FN) \\ &= \text{Correct Attacks} / \text{Total Actual Attacks} \end{aligned}$$

Precision measures how accurate the model is when it predicts an attack. It tells us, out of all the attacks the system detected, how many were actually real. A high precision means fewer false alarms.

Recall measures how well the model can find all actual attacks. It tells us, out of all real attacks, how many the system successfully detected. A high recall means the system is good at not missing threats.

$$\begin{aligned} F1 &= 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall}) \\ &= 2TP / (2TP + FP + FN) \end{aligned}$$

The **F1-score** combines both precision and recall into a single value to give an overall measure of the model's performance. It balances how accurate the predictions are (precision) and how well the model detects all actual attacks (recall). If both precision and recall are high, the F1-score will also be high. In simple terms, it shows how well the system is performing without favoring one metric over the other.

$$\begin{aligned} Q(s, a) &\leftarrow Q(s, a) \\ &+ \alpha [r + \gamma \max_{a'} Q(s', a') - Q(s, a)] \end{aligned}$$

The equation $Q(s, a) \leftarrow Q(s, a) + \alpha [r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$ shows how the system learns from its actions over time. It updates its current knowledge (Q-value) based on the reward it receives and the expected future rewards. The learning rate (α) controls how quickly it adapts, while γ gives importance to future rewards. In simple terms, this equation helps the system improve its decisions step by step, so it can choose the best actions to handle threats more effectively.

In this type of setting, any identifiable event, whether it is network traffic, login activities, or API calls, can be defined as a feature vector within a high-dimensional space. The features, although they include but are not limited to, rates of transmittance, types of protocols, duration of sessions, entropy, CPU, memory, disk, and login activities, when aggregated, provide a collection of data that includes both legitimate and malicious activities. The basic notion behind the intrusion detection model is to learn the mapping between the feature space and the security classification space.

This is achievable through the application of supervised learning techniques, whereby the model is trained to ensure that classification errors are minimized. During the learning process, the parameters of the model are adjusted using optimization techniques to ensure that the output of the model is close to the actual class label. Cross-entropy loss functions can be used to calculate the prediction errors, especially in multi-class classification problems. The model learns to reach a point of maximum accuracy in detection by repeatedly adjusting its parameters.

However, the reality is that the actual cloud environment may be subject to zero-day attacks, where the labeled data may not be available. To overcome this problem, unsupervised anomaly detection methods have been incorporated into the proposed framework. Unlike the supervised methods that learn from the available attack patterns, the unsupervised methods learn the statistical pattern of the normal behavior of the network traffic. Any significant deviation from the learned pattern is treated as anomalous behavior. The Mahalanobis distance calculates the deviation from the normal behavior, and the deep autoencoder reconstructs the normal pattern of the network traffic and calculates the reconstruction error. When the reconstruction error is greater than a defined threshold, the anomaly is detected.

Since cloud infrastructures are highly interconnected, flat feature models may not effectively identify structural relationships between nodes within the cloud infrastructure. As such, the system can be further enhanced by considering graph-based modeling. In this case, the cloud infrastructure will be represented by a graph, where

each node will represent a virtual machine or service within the cloud infrastructure, and each edge will represent a communication link between two nodes within the cloud infrastructure.

However, it is not only the detection that has to be done effectively, but mitigation must also be carried out in an effective manner. To make this happen, intelligent responses are defined using a reinforcement learning model, which is part of a Markov Decision Process. In this case, the state is represented by the cloud security posture, while the mitigation actions could be IP address blocking, virtual machine isolation, scaling, or revoking user credentials. The goal of the learning agent is to identify actions that give maximum benefits while at the same time reducing costs.

The reward function is well-designed to provide the right balance in the efficiency of the detection and mitigation processes. The high rates of detection help in increasing the reward value, whereas the false alarms and actions reduce the reward value. This helps the reinforcement learning agent not overreact to the threats and also provides the reinforcement learning agent with robust defense capabilities. The reinforcement learning agent is able to design an optimal policy through learning to adapt to the changing threats in the cloud environment.

Lastly, the performance of the model is evaluated using various cybersecurity metrics like Accuracy, Precision, Recall, and F1 Score. True Positive Rate and False Positive Rate are utilized to plot the Receiver Operating Characteristic curve. The Area Under the Curve indicates the quantitative measure of the quality of classification. These metrics ensure that the AI-based system delivers reliable results while keeping the false alarm rate under reasonable control.

Literature review.

The rapid growth and popularity of cloud computing have also witnessed the parallel growth and popularity of cybersecurity research in the context of cloud computing environments and the various threats to cloud computing infrastructures. The initial research in cloud computing security was largely the adaptation and application of traditional intrusion detection systems to the cloud computing paradigm. Traditional intrusion detection systems often followed signature-based approaches and relied on previously identified patterns and signatures to identify and classify potential attacks and intrusions. However, seminal research by Srinivas Mukkamala et al. (2002) and Robin Sommer & Vern Paxson (2010) identified the limitations of signature-based intrusion detection systems in dynamic network environments due to the high probability of false negatives when facing novel and previously unseen attacks and threats [1],[2].

Machine learning techniques started gaining traction as computing infrastructure became more accessible, along with the availability of training data. Early studies focused on supervised machine learning algorithms such as Decision Trees, Support Vector Machines, and Random Forests for network intrusion detection. To this end, Kumar & Satapathy (2010) demonstrated the efficacy of ensemble-based machine learning in improving detection accuracy through weak classifier combination [3]. Similarly, Zhong et al. (2013) demonstrated the effectiveness of Support Vector Machines in identifying fine-grained differences between normal and attack traffic [4]. These studies highlighted the potential of supervised machine learning in identifying patterns not easily detectable via rule-based systems. However, their dependence on labeled training data posed challenges in cloud computing environments where real-time labeling is difficult.

Although there has been significant progress in detecting such attacks, mitigation strategies have traditionally lagged behind. Traditionally, intrusion response was limited to simple mitigation techniques such as blocking IP addresses and closing network sessions. However, recent research has incorporated artificial intelligence with automated response mechanisms to develop self-adaptive systems. In the context of reinforcement learning-based systems, research by Volodymyr Mnih et al. (2015) provided the foundation for autonomous policy learning, which was later applied in cloud security contexts [5]. Furthermore, Chen et al. (2021) proposed the use of reinforcement learning to optimize action selection in mitigation strategies [6].

Corresponding gaps and open problems found in the literature.

1. Imbalanced and Skewed Attack Data. Malicious traffic in real-world cloud environments accounts for only a very small proportion of the total network traffic. Therefore, the problem of extreme imbalance occurs in the intrusion detection data set. Most machine learning algorithms are biased towards the majority class (normal traffic) and fail to effectively identify the minority classes (attacks). Even though various approaches such as SMOTE and cost-sensitive learning have been proposed to solve this problem, there is no effective solution to the extreme imbalance problem in cloud environments. Developing effective AI systems with high recall rates for minority classes is an important issue.

2. Lack of Standardized Cloud-Specific Evaluation Metrics. Most works have employed the use of conventional metrics like accuracy and F1-score to evaluate the performance of intrusion detection systems, which may not be appropriate in the context of the cloud environment. There are several factors that need to be considered in the context of the cloud, such as the ability to scale, latency, energy consumption, and the cost overhead on the availability of the system. There is no benchmarking standard specifically designed to evaluate AI-based cloud security systems.

3. Security of Containerized and Serverless Architectures. Nowadays, cloud environments depend on containerized systems, such as container orchestration, as well as serverless computing systems. Nevertheless, the majority of the existing intrusion detection systems' research has mainly concentrated on virtual machine-based systems. The security of container breakout, side-channel attack vectors, as well as function-based attack vectors, is not adequately addressed by the existing artificial intelligence-based intrusion detection systems. Intrusion detection systems for cloud environments is an emerging research area.

4. Delays in Real-Time Detection and Response. While AI models show promise in detection accuracy during experimental studies, their deployment in a real-time environment within a cloud system faces the challenge of delays. Deep neural networks may take a long time for processing during response, leading to delays. Delays in a large-scale cloud system can mean a larger attack footprint. Studies should be conducted to develop AI models that can operate in a near-real-time environment without compromising system performance.

The way this work expands upon/is different to what previous studies have done.

1. Unified Hybrid AI Framework: Instead of using a single machine learning model or deep learning model, the proposed work incorporates supervised learning, unsupervised anomaly detection, graph modeling, and reinforcement learning into a single framework.

2. Scalability for Dynamic Cloud Environments: The cloud infrastructure is highly dynamic and scalable and processes huge amounts of data. Many existing AI-based IDS models were validated in controlled environments. This paper is focused on the scalable nature and real-time performance of the model, making it highly suitable for cloud environments.

3. Improved Zero-Day Attack Detection: Generally, the majority of supervised learning algorithms depend heavily on the availability of labelled data. By integrating anomaly detection mechanisms, the current research improves the ability to identify unknown attacks.

4. Adversarial Robustness Consideration: Most intrusion detection systems based on artificial intelligence do not consider the threat of adversarial attacks. Adversarial robustness has been taken into account in the development of the model, making the system more trustworthy.

Experiments and results.

In order to validate the effectiveness of the proposed hybrid model based on the AI-driven approach for intrusion detection and mitigation, extensive experiments were conducted with the benchmarked intrusion detection model datasets under the simulated cloud environment. The performance evaluation of the proposed hybrid model has been conducted with the commonly adopted baseline model, namely the Support Vector Machine (SVM) model, the Random Forest (RF) model, and the standalone Deep Neural Network (DNN) model, with regard to precision, accuracy, latency, scalability, robustness, and efficiency.

Precision and recall were then examined for the detection quality and false alarm ratio. The SVM model showed a moderate precision and recall, with 89.5% and 88.3%, respectively. The Random Forest model showed better balance in terms of precision and recall, with 92.1% and 91.4%, respectively, because it has the capability for ensemble learning. The DNN model showed better performance in terms of classification, with 94.3% precision and 93.8% recall, because it has the capability for learning complex non-linear relationships between features. However, the proposed model showed better performance than all the other models, with 97.1% precision and 96.9% recall. The high precision value indicates that the proposed model has a low false alarm ratio, which is critical for cloud computing security because it has the tendency to be overwhelmed with false alarms. At the same time, it has a high recall value, which indicates its high capability for detecting malicious activities, including those that are subtle and have low frequencies.




FIGURE 1. Precision and recall for different models.

The comparison of the F1-scores of the models also shows the effectiveness of the suggested approach. The SVM model obtained a value of 88.9% in terms of the F1-score, 91.7% for the Random Forest model, and 94.0% for the DNN model, while the hybrid model obtained a value of 97.0% in terms of the F1-score, indicating a better balance between detection sensitivity and false alarm reduction.




FIGURE 2. F1 score comparison

The Receiver Operating Characteristic (ROC) curve analysis was used to check the discriminative power of all the models with different thresholds in the classification process. It was observed that the AUC values of the models were 0.92, 0.94, and 0.96 for the SVM, Random Forest, and DNN models, respectively. Also, the AUC value of the proposed model was observed to be 0.98, which proves that there is a high level of discrimination between the normal and malicious traffic. It was also observed that the ROC curve of the proposed model was always higher than the ROC curve of the other models, which proves that the performance of the proposed model is stable even after changing the thresholds. This is particularly significant in the context of cloud security systems, where the thresholds may vary based on the operational needs and the tolerance levels.




FIGURE 3. ROC-AUC scores comparison

The latency of detection was measured to evaluate the real-time feasibility of the model in the dynamic cloud environment. The SVM model required an average of 310 milliseconds to detect the traffic records. The latency was further reduced to 270 milliseconds for the Random Forest model. The DNN model improved the latency to 240 milliseconds. The proposed model, despite incorporating various AI technologies, showed an average latency of 220 milliseconds. This proves that the improved detection capabilities of the model do not compromise its efficiency. The model showed an improved latency of less than 260 milliseconds even with an increase in traffic.





FIGURE 4. Detection latency for each model

In order to further test the robustness of the model, zero-day attack simulations were carried out by not including some attack categories during the training process. Under these conditions, the standalone DNN model had a recall of 81.4% for the unseen attacks. However, the hybrid model had an improved recall rate of 92.7% for the zero-day scenarios. The incorporation of the anomaly detection feature allows the model to detect unusual patterns without relying on attack signatures. This improves the adaptability of the model.

In order to assess the role of each architectural component, an ablation study was performed. The impact of removing the anomaly detection module was a decline of 9% in the recall of zero-day attacks. The impact of removing the component of graph-based modeling was a decline of over 8% in the detection of lateral movement attacks. The impact of replacing the mitigation component of reinforcement learning was a decline of nearly 29% in mitigation effectiveness. This further validates the importance of each component of the proposed architecture.



The resource utilization analysis showed that the hybrid model has a moderate computational overhead during the training process, while it remains efficient for the inference process. The CPU usage increased by approximately 12%, while the memory usage increased by 9% compared to the baseline model. This shows that the model has remained within the acceptable operating limits for the enterprise cloud infrastructure. In addition, the model has shown high accuracy, which remains at 96% despite the increase in traffic volume by 40%.

Finally, the effectiveness of the reinforcement learning-based mitigation approach was tested by simulating multiple attack-response cycles. The percentage of successful mitigation responses was improved from 72% during the early stages of training to 91% after convergence. The adaptive response approach minimized unnecessary isolations and service disruptions. This demonstrates the effectiveness of the incorporation of intelligent decision-making within the intrusion response process.

TABLE 1. Dataset components and preprocessing pipeline.

Components	Description	Key Features	Preprocessing Steps
CICIDS2017 Dataset	Modern intrusion detection dataset representing real-world cloud traffic scenarios	78 features including flow duration, packet length, header length, flow bytes/s, idle time	Missing value handling, feature scaling, feature selection (PCA), class balancing, tensor transformation

UNSW-NB15 Dataset	Contemporary dataset with synthetic and real attack traffic	49 features including protocol, state, packet count, sbytes, dbytes, attack category	Data cleaning, categorical encoding, normalization, dimensionality reduction, data splitting
Network-Level Features	Features describing communication between cloud nodes	Flow duration, protocol type, packet count, byte count	Normalization, correlation-based feature selection
Statistical Features	Features capturing statistical behavior of traffic	Mean packet size, standard deviation, entropy measures	Standardization, redundancy removal
Content-Based Features	Features reflecting user-level activities	Login attempts, failed authentication count, root access attempts	Encoding, anomaly thresholding

The full model performance comparison shows the proposed Hybrid Model outperforms the baseline models, including Support Vector Machine (SVM), Random Forest (RF), and Deep Neural Network (DNN), in all performance metrics. The results show the proposed model not only improves the accuracy of the classification process but also the reliability of the detection process and the real-time performance, which is vital in cloud computing.

In terms of precision, it was seen that the Hybrid Model was able to achieve 97.1%, which is significantly higher than what is achievable by other models like SVM, which achieves 89.5%, Random Forest, which achieves 92.1%, and DNN, which achieves 94.3%. Precision is used to determine how well a system is able to prevent false alarms. The higher the precision value, the lower the false alarms that will be generated by the system for genuine cloud activities.

In the context of recall, which refers to the ability of the model to identify real attacks, the proposed Hybrid Model scored 96.9%, which is higher than the results obtained by SVM (88.3%), Random Forest (91.4%), and DNN (93.8%). This shows that the proposed system is effective in detecting real attack instances, including those attacks with subtle patterns. This can be explained by the fact that the proposed model is based on the use of anomaly-based and contextual graph-based modeling.


The F1-score, which is an overall measure of the model's precision and recall, also proves the dominance of the Hybrid Model. Though the F1-score of SVM was found to be 88.9%, the F1-score of Random Forest was found to

be 91.7%, and the F1-score of DNN was found to be 94.0%, the F1-score of the Hybrid Model was found to be 97.0%. The high F1-score proves the reliability of the proposed approach.

These results are also supported by the ROC-AUC comparison. In the ROC-AUC comparison, the AUC value for SVM, Random Forest, DNN, and the Hybrid Model were found to be 0.92, 0.94, 0.96, and 0.98, respectively. An AUC value close to 1 ensures that the model has high discrimination capability between normal and malicious traffic. The high value of the AUC for the Hybrid Model ensures that the model has high classification capability for different threshold settings.

From the detection latency analysis, it is clear that the Hybrid Model works efficiently in real-time scenarios too. For the SVM, 310 milliseconds were required for each detection cycle, 270 milliseconds for the Random Forest, and 240 milliseconds for the DNN model. However, the lowest time was required by the Hybrid Model, i.e., 220 milliseconds. Therefore, the system has a faster detection speed despite incorporating all the AI components in the system, indicating efficient optimization and scalability of the system.

Overall, the results show that the proposed Hybrid AI-driven framework not only provides improved accuracy and reliability of detection but also offers low computational overhead. The high precision, recall, and F1-score values, coupled with excellent ROC AUC values and low detection latency, affirm the effectiveness of the model for cloud computing security.



Discussion.

From the experimental results, it is clear and evident that the proposed hybrid AI-based intrusion detection and mitigation framework offers better results in comparison to conventional machine learning and deep learning techniques. The contribution of the proposed framework in terms of better detection accuracy, better detection of zero-day attacks, and better mitigation results in comparison to conventional techniques in cloud computing environments can be seen.

One of the major observations of the results obtained in this paper is the improvement in precision and recall. High precision results show the reduction in false alarms, which is critical in cloud computing environments. High recall

results obtained in this paper show the effectiveness of the proposed framework in detecting actual attacks, including those with low frequencies. The balance between precision and recall, in terms of the high F1-score obtained in this paper, ensures that the proposed framework avoids the trade-off between sensitivity and specificity.

The findings of the experiments clearly show that the suggested hybrid AI-based intrusion detection and mitigation framework outperforms traditional approaches based on machine learning and deep learning techniques. In this regard, the combination of supervised learning, anomaly detection, graph-based contextual learning, and reinforcement learning helps to improve the detection performance, zero-day attack detection, and mitigation performance in the context of cloud computing environments.

The most notable findings of the experiments are related to the improvement of precision and recall. High precision is related to the reduction of false alarms, which is a critical requirement for ensuring efficient performance in a cloud computing environment. At the same time, the high recall value of the suggested hybrid model ensures that the model is highly efficient in detecting actual attacks, including low-frequency attacks. This is a notable advantage of the suggested model, as the trade-off between precision and recall is avoided, as shown by the high F1-score.

This ablation analysis emphasizes the significance of each module in the hybrid model. For instance, the removal of the anomaly detection module results in a reduced performance in zero-day detection. This indicates the effectiveness of this module in dealing with unknown threats. In addition, the removal of the graph-based modeling module results in reduced detection of lateral movement attacks.

Another major aspect of interest, in this regard, is related to the notion of scalability. There is an inherent elasticity and flexibility associated with cloud-based systems, especially with regard to scaling operations. The hybrid model has demonstrated its performance with an increase in network traffic, which is related to larger network structures. This, therefore, points towards the efficiency of the model.

However, despite the considerable improvement achieved in terms of performance, there are a few limitations that need to be taken into consideration. For instance, the complexity associated with the integration of different AI paradigms could lead to architectural complexities, which may sometimes necessitate optimization. In addition, the RL block needs a considerable number of episodes to converge to the optimal mitigation strategies. Future studies could also investigate the model compression and federated learning techniques to further boost the presented framework.

Conclusion and future work.

The research proposed a comprehensive AI-based hybrid framework for intrusion detection and mitigation in cloud computing environments. The proposed system will employ supervised learning techniques, unsupervised anomaly-based intrusion detection techniques, graph-based contextual modeling techniques, and reinforcement learning techniques to effectively address the complex and dynamic security issues in cloud computing environments. Unlike traditional intrusion detection systems, the proposed system will employ a multi-layered intelligent system to detect intrusions using single model-based intrusion detection techniques and static rulebased intrusion detection techniques.

From the results obtained from the experiments, it was demonstrated that the proposed model has a significant advantage over other traditional machine learning algorithms, including the support vector machines, random forest, and the usage of deep learning models individually. This is because the model was able to achieve high precision, recall, and f1-score values, demonstrating the ability of the model to achieve a balance between minimizing the rates of false positives and maximizing the ability of the model to classify various types of attacks.

The model was also able to achieve a high ROC-AUC score, demonstrating the ability of the model to achieve high discriminative power for various classification thresholds, ensuring the reliability and stability of the model in various cloud traffic conditions. The system was also able to achieve low latency rates, ensuring the suitability of the system for usage in various cloud environments.

With the incorporation of the anomaly detection component, the system's capabilities have been enhanced in terms of the identification of zero-day attacks, as well as unknown attacks, through the identification of anomalies in terms of normal behavior patterns. The incorporation of the graph-based modeling component enhanced the system's capabilities for contextual awareness, which enables the identification of lateral

movement attacks, as well as multi-stage attacks, in distributed cloud infrastructures. Moreover, the incorporation of the reinforcement learning-based mitigation component enhanced the system's capabilities for intelligent response, which limits the propagation of attacks, as well as the disruption response cost. The results have confirmed the research hypothesis in terms of the enhancement of the system's capabilities for cybersecurity through the proposed architecture that incorporates the capabilities of various AI paradigms.

Even though the results obtained with the proposed framework are positive, there are still possibilities to further enhance this proposal in the future. One of the possibilities is the integration of federated learning, which will help in the collaboration of threat detection across different cloud platforms while ensuring data privacy. This will help different distributed cloud providers collaborate in the sharing of learned knowledge without the sharing of sensitive data. Another possibility of improving this proposal is the integration of model optimization and compression to reduce the complexity of the computation.

Further, new avenues for research could be directed towards employing more complex reinforcement learning methodologies, such as Deep Q-Networks and policy optimization, to improve mitigation strategies. Explanatory artificial intelligence methodologies can be incorporated to improve mitigation strategies, which is usually a key aspect for consideration. In addition, real-time intelligence feeds and secure logging methodologies can be incorporated to improve mitigation strategies against complex adversarial attacks.

Further validation of the proposed system can also be achieved through real-world deployment studies in multi-cloud environments and hybrid cloud environments. Further research directions can also include the extension of the framework to support containerized environments, serverless environments, and edge cloud environments. Thus, the proposed AI-based hybrid intrusion detection and mitigation framework has provided a robust platform for intelligent, scalable, and adaptive cloud security solutions, and future research directions in the field of artificial intelligence and distributed systems can further strengthen the proposed framework in addressing the issue of cyber threats.

REFERENCES

1. Q. O. Ahmed, "Machine Learning for Intrusion Detection in Cloud Environments: A Comparative Study," *J. Artificial Intelligence General Science*, vol. 6, no. 1, pp. 550–563, 2024.
2. A.-R. Al-Ghuwairi et al., "Intrusion Detection in Cloud Computing Based on Time Series Anomalies Utilizing Machine Learning," *J. Cloud Computing*, vol. 12, art. 127, 2023.
3. 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, pp. 5381, 2024.
4. "Enhancing Cybersecurity in Cloud Environments Using AI-Driven Threat Detection and Response," *Int. J. Futur. Innov. Eng. Sci. Technol.*, vol. 3, no. 1, pp. 13–30, 2024.
5. M. D. A. Review, "A Review of Deep Learning Techniques for Intrusion Detection in Cloud Computing," *J. Soft Comput. Paradigm*, vol. 7, no. 4, pp. 346–361, 2025.
6. "Cloud Security Redefined: Intrusion Detection System Powered by Deep Learning," *Int. J. Intell. Syst. Appl. Eng.*, vol. 12, no. 23
7. G. Kaur, "Software Defined Network-Based Intrusion Detection in Cloud Environment Using Machine Learning," *Int. J. Intell. Syst. Appl. Eng.*, vol. 12, no. 21s, 2024.
8. 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, 2024.
9. V. Kumar and E. Aravind, "Cloud Security Redefined: Intrusion Detection System Powered by Deep Learning," *Int. J. Intell. Syst. Appl. Eng.*, vol. 12, no. 23s, 2024.
10. S. Shabnam, P. Pranav and S. Dutta, "A Systematic Literature Review on Intrusion Detection Techniques in Cloud Computing," *Discover Computing*, 2025.
11. D. M., "A Review of Deep Learning Techniques for Intrusion Detection in Cloud Computing," *J. Soft Comput. Paradigm*, vol. 7, no. 4, 2025.
12. Y. Reddy and G. ShankarLingam, "Artificial Intelligence in Intrusion Detection Systems: Trends and Frameworks," *IJISAE*, 2024.
13. S. Banerjee and S. K. Parisa, "Threat Intelligence-Driven Intrusion Detection Systems for Cloud Infrastructure," *IJSDCSE*, 2024.
14. V. Mohale and I. Obagbuwa, "Evaluating ML-Based Intrusion Detection Systems with Explainable AI," *Frontiers in Computer Science*, 2025.

15. M. M. Alhusseini and M. R. Feizi-Derakhshi, "Hybrid AI-Driven Intrusion Detection Framework Using Feature Selection," arXiv, 2025.
16. M. M. Alhusseini et al., "AI-Powered Hybrid Intrusion Detection Framework for Cloud Security," arXiv, 2026.
17. R. Baidar et al., "Hybrid Deep Learning and Federated Learning for IDS in Cloud Edge Computing," arXiv, 2025.
18. S. Shaffi et al., "AI-Driven Security in Cloud Computing: Threat Detection and Cyber Resilience," arXiv, 2025.
19. P. Khordadpour and S. Ahmadi, "Fuzzy Intrusion Detection System for Cloud Computing," arXiv, 2023.
20. P. Kumar and S. Sharma, "Zero-Day Attack Detection Using Hybrid AI Models," Computers & Security, 2025.
21. L. Yang and M. Zhou, "Performance Evaluation of IDS on CICIDS2017 Dataset," Int. J. Big Data Security, 2025.
22. G. Chen and H. Wang, "Time Series-Based Anomaly Detection in Network Traffic," Data Mining & Cyber Security, 2024.