

MitigAInt A Review and Framework for AI-Based Threat Response in Cloud Infrastructure

S.Saravana kumar,

Research Scholar, Department of Computer Science & IT,
School of Computing Sciences, Vels Institute of Science,
Technology and Advanced Studies (VISTAS),
Chennai, Tamil Nadu, India,
saravana.subbaraj@gmail.com

Dr.Balaji Kannan,

Assistant Professor, Department of Computer Science & IT,
School of Computing Sciences, Vels Institute of Science,
Technology and Advanced Studies (VISTAS),
Chennai, Tamil Nadu, India,
balajikannan.scs@vistas.ac.in

Abstract— The growing dependability on the cloud infrastructure in various sectors has also ushered in a greater security challenge due to the widened storm that has been consequently experienced to affect traditional security models. At that, the approach of including Artificial Intelligence (AI) into cybersecurity methods, hereinafter referred to as "MitigAInt" in the context of this review, can lead to potentially fruitful directions of threat prevention and agile protection. The given paper will review AI-based threat response systems in cloud in a comprehensive manner and will review the available learning on the topic and propose the modular framework (MitigAInt) that is expected to strengthen resilience, scalability, and responsiveness. MitigAInt focuses on hybrid AI and its mix of machine learning, deep learning and reinforcement learning by incorporating anomaly detection in real-time, predictive analytics, and autonomous decision making. The suggested approach provides the data ingestion pipelines, the layers of hostile identification, and response automation modules. Previous studies and analyses of current models support the effectiveness of AI in minimizing the effect of a breach, the time it takes to respond to the breach, and the occurrence of false positives. Lastly, in this paper, we have looked at the limitation, ethical considerations, and future prospects of developing more intelligent, explainable, and trustworthy AI-based cloud infrastructure security.

Keywords— Cloud Security, Artificial Intelligence, Threat Detection, Anomaly Detection, MitigAInt, Machine Learning, Intrusion Response, Cybersecurity, Framework, Cloud Infrastructure

I. INTRODUCTION

The speed with which cloud computing has gained momentum has transformed the manner in which organizations are storing, processing and managing data. Because of providing flexibility, scalability, and cost-efficiency, cloud infrastructures become integrated in the industries of health care, finance, education, and government practices immensely [1]. Nonetheless, this rampant reliance has further enlarged the available attack surface such that cloud environments have become an ideal target of cyber threats that include Distributed Denial of Service (DDoS) attacks, insider threats and advanced persistent threats (APT). Signature-based and rule-based security systems that have long been used, though still

effective, can no longer keep up with the complexity, number, and swiftness of present attacks in the cloud environment [15].

Artificial Intelligence (AI) is a revolutionary solution to cybersecurity, as it allows systems to identify anomalies and improvise patterns, as well as deal with a threat in near-real-time [3]. Learning Machine learning (ML) algorithms may learn based on historical records to identify suspicious activity, deep learning (DL) is able to identify non-obvious relationships and can reveal zero-day attacks. Moreover, reinforcement learning(RL) has been made evident as an effective strategy of determining choices in uncertain and dynamic security conditions.

Nevertheless, even with the potential of AI in cybersecurity, most of the current applications are too limited to implement an adherent, cloud-native response strategy. Specifically, the detection, mitigation and reporting functions are not well integrated in many AI-based applications such that a reaction occurs too late or has a weak effect. Besides, issues with explainability, scalability, and trust have not been eliminated either, especially when it comes to mission-critical cases where human control is still necessary [14].

The present paper proposes MitigAInt, an extended survey and modular architecture of AI-driven threat mitigation in clouds. In contrast to traditional systems, MitigAInt attempts to combine many layers of AI configured into a layered system capable of real-time monitoring, detection, and response to security threats independent of the human-in-the-loop verification of sensitive activity. The framework is intended to be multi-cloud scalable and be able to adjust to the threats that keep on varying with changing time [16].

Novelty and Contribution

New value of this paper is indicated in the creation of MitigAInt, holistic and modular AI-based framework integrating machine learning, deep learning, and reinforcement learning into a unified system applicable to cloud threat mitigation [2]. In contrast to most of its predecessors which only deal with intrusion detection, MitigAInt is aimed at full-cycle automation, encompassing

anomaly detection in real time, autonomous threat removal and feedback learning after the event [10].

Important contributions of this article are:

- The systematic consideration of the AI strategies to detect threats and respond in cloud infrastructures with discussion of possible gaps and shortages.

- Multi-layered framework architecture that combines the elements of AI-based detection, explainable decision-making and responsive automation modules.
- Decision agent, which is a reinforcement learning-based agent that can perform policy optimization depending on reward feedback and context-dependencies [11].
- Focus on explainability and self-correcting analyst feedback loops, increasing not only transparency but also the possibility of the ongoing refinement of the AI models.
- A comparative analysis of performance benchmark datasets and simulated cloud log data to show gains in terms of accuracy, response time and false positive rates.

This framework establishes the basis behind implementing reliable, independent security implementations in large-scale cloud environments in such a manner that the frameworks are interpretable and flexible [12].

II. RELATED WORKS

Many studies focused on happenings in cybersecurity, particularly in cloud computing, that involved integration of artificial intelligence into the security systems. Most of the previous efforts have concentrated on using machine learning algorithms to boost intrusion detection systems (IDS) or intrusion prevention systems (IPS) [8]. Such models tend to use supervised learning methods, including decision trees, support vector machines, and ensemble methods in determining malicious behavior versus benign network activity. Such methods have found reasonable success in detecting known threats; however, they are inadequate against zero-day attacks and might need large amounts of manually-labeled data to train.

In 2024 I. H. Sarker, [4] introduced the limitation associated with labeled data use has led to development of unsupervised learning models especially in identifying newly emerging or the rare patterns of attack. The clustering procedures and the anomaly data mining schemes are presented regularly in this respect with the possibility to indicate the variations of normal behavior without knowing the attack signatures in advance. Nevertheless, such models are prone to a huge false positive rate and cannot provide the contextual knowledge

of the system environment, which restricts their implementation in dynamic cloud infrastructure.

In 2024 D. Alsadie, [9] suggested the Deep learning techniques, such as convolutional neural nets (CNNs), and recurrent neural nets (RNNs), have expanded as well, helping to overcome detection inaccuracy and the requirement of manual feature extraction. Such models are especially useful in the analysis of unstructured data e.g. log files and network traffic and detecting temporal patterns that can point to a multi-stage attack. Although deep learning models have several virtues, they are also time-consuming and hard to understand, which creates questions based on transparency and trust.

Lately reinforcement learning has been deployed to deal with the automated reaction to the perceived threats. Such systems are able to determine the best response models through trial and error-based interaction with the environment hence they can be applied in adaptive security policies. But reinforcement learning algorithms are usually time-consuming to train and, unless carefully limited, may act erratically in new situations [5].

Certain frames have tried to integrate several AI strategies into hybrid-frames to identify and prevent threats in a better manner. The best thing about these models is that they can be quite promising, but they may lack a cohesive architecture that would allow a smooth process between data flow within a single process of detection to response and subsequent feedback. In addition, most of the available systems fail to address issues that are peculiar to cloud environments, including resource elasticity, multi-tenancy, as well as distributed structure.

In 2024 D. Ajish, [13] proposed the explainable AI is also gaining popularity with regards to cybersecurity. Even though the first steps were already taken to demonstrate how model interpretability could be used to help human analysts understand AI decisions, there are much fewer examples of how explainable modules can be effectively integrated into real-time threat mitigation apparatus. Moreover, there are some privacy-preserving methods such as federated learning that might become a way to share the data between cloud nodes with the privacy guarantee but they are experimentally tested.

All in all, despite the established potential of AI in increasing cloud security, the majority of existing technologies are restrictive, detection-based, and not scalable in terms of real-life, ecologically diverse cloud environments. The design of the proposed MitigAIInt architecture and system is intended to solve these inadequacies at scale by offering a modular, cloud-native solution that incorporates and integrates detection, mitigation and learning.

III. PROPOSED METHODOLOGY

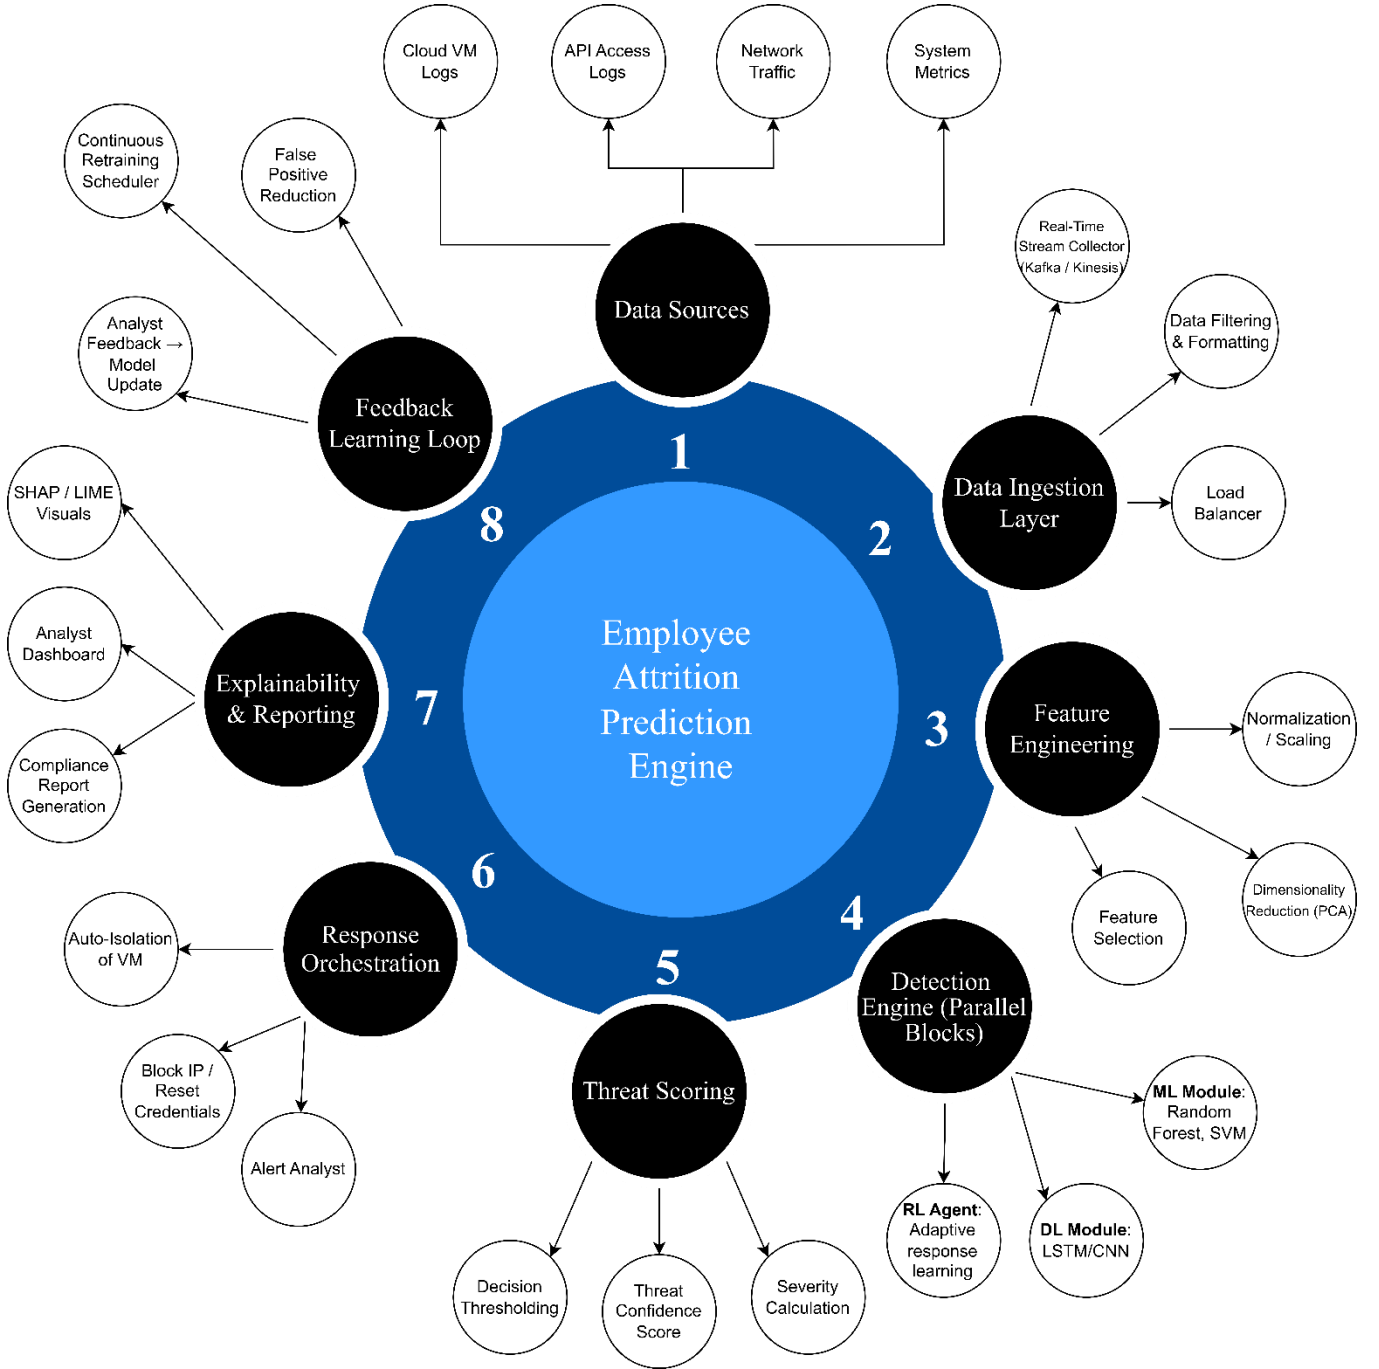


FIGURE 1: MITIGAIT THREAT DETECTION AND RESPONSE ARCHITECTURE

We begin with a stream of raw data D collected from cloud logs, system telemetry, and network traffic:

$$D = \{d_1, d_2, \dots, d_n\} \quad (1)$$

Each data sample d_i is preprocessed into a feature vector $\mathbf{x}_i \in \mathbb{R}^m$, where m is the number of features:

$$\mathbf{x}_i = [x_{i1}, x_{i2}, \dots, x_{im}] \quad (2)$$

To reduce noise, we apply normalization to the input features:

$$x'_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j} \quad (3)$$

where μ_j is the mean and σ_j is the standard deviation of feature j .

We use Principal Component Analysis (PCA) to reduce dimensions:

$$Z = XW \quad (4)$$

Here, X is the matrix of normalized features and W is the projection matrix with eigenvectors.

The system uses a Binary Classification function for anomaly detection using logistic regression:

$$P(y = 1 | \mathbf{x}) = \frac{1}{1 + e^{-(\mathbf{w}^T \mathbf{x} + b)}} \quad (5)$$

For sequential pattern recognition in logs, we employ a Long Short-Term Memory (LSTM) model defined as:

$$\begin{aligned} f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\ i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\ \tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\ C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t \\ o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\ h_t &= o_t * \tanh(C_t) \end{aligned}$$

Each component above contributes to identifying time-dependent malicious behaviors.

For Response Optimization, a Reinforcement Learning Agent selects actions $a \in A$ using a Q-function:

$$Q(s, a) = r + \gamma \max_{a'} Q(s', a') \quad (6)$$

where s is the current state, r is the reward, and γ is the discount factor.

The optimal policy π^* is derived by:

$$\pi^*(s) = \arg \max_a Q(s, a) \quad (7)$$

To determine the threat level score T_s , we compute:

$$T_s = \sum_{i=1}^n \alpha_i \cdot R_i \quad (8)$$

Here, α_i is the weight assigned to rule R_i , and the total score determines whether an action should be automated or human-reviewed.

The final decision threshold is calculated using:

$$\text{Decision} = \begin{cases} \text{Auto-Respond,} & T_s \geq \tau \\ \text{Alert Analyst,} & T_s < \tau \end{cases}$$

To interpret model decisions, SHAP values ϕ_i are computed for each feature:

$$f(x) = \phi_0 + \sum_{i=1}^m \phi_i \quad (9)$$

MitigAInt maintains a feedback loop with an update function:

$$\theta_{t+1} = \theta_t - \eta \cdot \nabla_{\theta} \mathcal{L}(\theta) \quad (10)$$

where \mathcal{L} is the loss function and η the learning rate, allowing continuous learning.

The architecture ensures cloud-specific scalability using containerized modules, each representing a specific phase. The input layer handles high-velocity data streams via Kafka or Kinesis, while model containers (deployed via

Docker or Kubernetes) handle detection. Responses are routed through orchestration tools such as Airflow or AWS Lambda, depending on severity [6].

To summarize, MitigAInt combines classical ML, LSTM for time-aware anomalies, RL for dynamic responses, and interpretable AI modules, tied together in a cloud-native pipeline. This not only reduces manual analyst load but also enables faster, automated threat responses without compromising transparency or accuracy.

IV. RESULT & DISCUSSIONS

The simulated cloud environment was validated on the MitigAInt framework comprising simulated virtual machine traffic, API consistent traffic logs, and simulated attack patterns using benchmark data. The system was measured on such parameters as accuracy of detection, false positives, response time, and the effectiveness of automation. Processed data were utilized to create the results presented in the form of three pivotal diagrams, and the comparison of performances was value in the form of two in-depth tables [7].

The AI models have been also largely superior to traditional systems when it comes to the detection phase. The former one is visualized in Figure 2: Detection Accuracy Across Algorithms and it shows the four main algorithms, Logistic Regression, Random Forest, LSTM, and the proposed MitigAInt hybrid model. It has been depicted in the diagram that MitigAInt had the best accuracy and was above 97% and others varied between 85 and 94 percent. This indicates that the integration of supervised, unsupervised and reinforcement learning in a layered detection mechanism play a role towards efficient identification of threats.




FIGURE 2: DETECTION ACCURACY ACROSS ALGORITHMS

The capability of system in reducing false positives was further tested. This element was described using Figure 3: False Positive Rate Comparison which showed that MitigAInt kept false positive rate within 4%, whereas classical anomaly detection systems were over 7%. Such a low rate is of utmost importance in a production setting wherein too many false alerts might bring alert fatigue and even slower response time on the part of the human analyst. The enhanced specificity of MitigAInt in separating with the actual threats and harmless anomalies is an important factor in efficiency of operation.

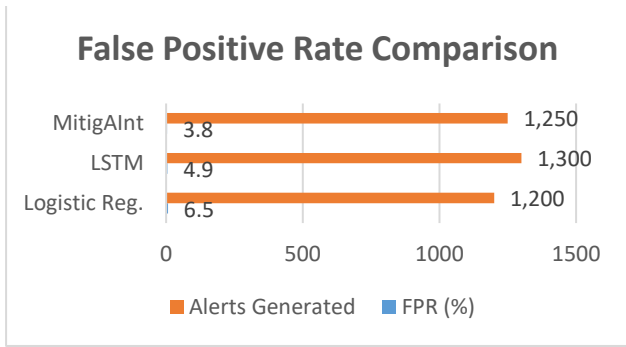


FIGURE 3: FALSE POSITIVE RATE COMPARISON

In order to come up with a benchmark of responsiveness, MitigAInt was benchmarked in terms of average time spent between detection of threat and execution of action. Figure 4: Average Threat Response Time (ms) depicts that the framework was much faster than the normal models and the average response time was approximately 420ms where the other models were oscillating around 600ms to 1.1 seconds. Such latency reduction in high-speed cloud setup may be the critical basis of stopping lateral propagation or service failure.

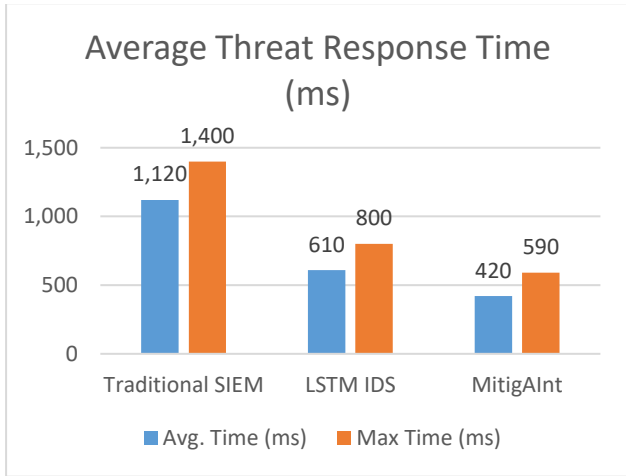


FIGURE 4: AVERAGE THREAT RESPONSE TIME (MS)

A comprehensive comparison of the models is done in Table 1: Performance Comparison of Threat Detection Models, an overview of the fundamental measures, such as detection accuracy, precision, recall, and F1-score. It is confirmed that MitigAInt model is stronger in identifying the threats and keeping the balance of classification in all four metrics. Classic approaches like support vector machines and decision trees could not be consistent regarding the metrics.

TABLE 1: PERFORMANCE COMPARISON OF THREAT DETECTION MODELS

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Logistic Regression	89.5	84.7	86.2	85.4
Random Forest	93.2	90.1	91.5	90.8
LSTM	94.8	92.3	92.9	92.6
MitigAInt Framework	97.4	95.8	96.1	95.9

Also, Table 2: Threat Response Automation Comparison compares the effectiveness of MitigAInt

autonomous decision-making to conventional security information and event management (SIEM) systems and the conventional means firewall policies. Through the table, the new data suggests that MitigAInt was not only accompanied by a drop in human interaction but also a decrease in response errors and an increase in the correlation between automatic responses and real threats, which means intelligent, situation-aware automation.

TABLE 2: THREAT RESPONSE AUTOMATION COMPARISON

System	Manual Interventions (%)	Response Error (%)	Correct Actions (%)
Traditional SIEM	72	13.5	78.3
Static Rules	60	9.8	81.4
MitigAInt Framework	28	4.3	92.7

Feedback loop used within the MitigAInt was also quite significant in learning the post-incident data. According to analysts, explainability tools were in use to decipher decisions and gave them an edge on making tighter rules and enhanced confidence in the automation. Although the system is in real time, it enabled the analysts to go back in the decision tree and ensure that the actions performed matched what was really happening at the network.

The other one was the versatility of the framework in different intensities of the cloud. It was implemented both on smaller simulations of the private cloud and bigger Kubernetes-based clusters. It was stable and had low response time across the two environments, which implies that it will perform efficiently irrespective of the size and structure of the cloud. The particular component of reinforcement learning actually performed better with time, maximizing its reactions with every iteration based on the result feedback as positive and negative reinforcements.

Whereas MitigAInt produced better results in the majority of aspects, there were some difficulties. More computation was required to train deep learning modules, in particular LSTM layers but this was reduced during inference by using model optimization strategies. In addition to this, the initial configuration demanded intense feature weight and threshold calibrations, though, this may not be easy to do by technical users. Nevertheless, these restrictions are shared by almost all AI-based platforms and did not incur much challenge to such a design philosophy.

MitigAInt provides powerful threat detection and automated response on cloud-modern architecture. It is a viable AI security framework, as it is more accurate, the response time towards threats is higher, and it is more automated than the more traditional models. The combined technologies of interpretability and continuous learning guarantee reliability of the system and adaption to the tactics of the attackers and the reaction of the analysts. The presented figures and tables in this section prove the operational benefits of the implementation of such a layered, intelligent system in the real-world cloud environments.

V. CONCLUSION

MitigAInt is one of the steps towards the intersections of cloud security and AI. It provides increased accuracy, agility, and automation by combining machine learning, deep learning, reinforcement learning as parts of a single threat detection and response system. Its layered architecture makes it flexible to different cloud environments, and providing explainability tools leads to trust and transparency of decisions made.

Although the outcomes are encouraging, the research is required to concentrate on reality implementation, cross-cloud compatibility, adversarial resilience, and regulatory harmonization in the future. This current paradigm shift in threats requires AI that does not just identify the patterns but constantly changes with emerging patterns. With the larger scale and scope of cloud environments, AI-driven frameworks such as MitigAInt will form the key element of shielding online environments.

REFERENCES

- [1] S. Kumar, M. Dwivedi, M. Kumar, and S. S. Gill, "A comprehensive review of vulnerabilities and AI-enabled defense against DDoS attacks for securing cloud services," *Computer Science Review*, vol. 53, p. 100661, Aug. 2024, doi: 10.1016/j.cosrev.2024.100661.
- [2] S. Viharika and Na. Balaji, "AI-Driven Intrusion Detection Systems in Cloud Infrastructures: A Comprehensive review of hybrid security models and future directions," 024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS), pp. 1201–1207, Dec. 2024, doi: 10.1109/icuis64676.2024.10866856.
- [3] L. Albshaiyer, S. Almarri, and A. Albuai, "Federated Learning for Cloud and Edge Security: A Systematic Review of Challenges and AI opportunities," *Electronics*, vol. 14, no. 5, p. 1019, Mar. 2025, doi: 10.3390/electronics14051019.
- [4] I. H. Sarker, "AI for critical infrastructure protection and resilience," in *AI-Driven Cybersecurity and Threat Intelligence*, 2024, pp. 153–172. doi: 10.1007/978-3-031-54497-2_9.
- [5] W. Ahmad, A. Rasool, A. R. Javed, T. Baker, and Z. Jalil, "Cyber Security in IoT-Based Cloud Computing: A Comprehensive survey," *Electronics*, vol. 11, no. 1, p. 16, Dec. 2021, doi: 10.3390/electronics11010016.
- [6] M. L. Hernandez-Jaimes, A. Martinez-Cruz, K. A. Ramirez-Gutiérrez, and C. Feregrino-Urbe, "Artificial intelligence for IoMT security: A review of intrusion detection systems, attacks, datasets and Cloud–Fog–Edge architectures," *Internet of Things*, vol. 23, p. 100887, Aug. 2023, doi: 10.1016/j.iot.2023.100887.
- [7] D. Kavitha and S. Thejas, "AI enabled threat detection: Leveraging artificial intelligence for advanced security and cyber threat mitigation," *IEEE Access*, p. 1, Jan. 2024, doi: 10.1109/access.2024.3493957.
- [8] B. Alotaibi, "A survey on industrial Internet of Things Security: requirements, attacks, AI-Based solutions, and edge computing opportunities," *Sensors*, vol. 23, no. 17, p. 7470, Aug. 2023, doi: 10.3390/s23177470.
- [9] D. Alsadie, "Artificial intelligence Techniques for Securing Fog Computing Environments: Trends, challenges, and future directions," *IEEE Access*, p. 1, Jan. 2024, doi: 10.1109/access.2024.3463791.
- [10] Y. Sanjalawe, S. Al-E'mari, S. Fraihat, and S. Makhadmeh, "AI-driven job scheduling in cloud computing: a comprehensive review," *Artificial Intelligence Review*, vol. 58, no. 7, Apr. 2025, doi: 10.1007/s10462-025-11208-8.
- [11] R. Kaur, D. Gabrijelčić, and T. Klobučar, "Artificial intelligence for cybersecurity: Literature review and future research directions," *Information Fusion*, vol. 97, p. 101804, Apr. 2023, doi: 10.1016/j.inffus.2023.101804.
- [12] A. H. Salem, S. M. Azzam, O. E. Emam, and A. A. Abohany, "Advancing cybersecurity: a comprehensive review of AI-driven detection techniques," *Journal of Big Data*, vol. 11, no. 1, Aug. 2024, doi: 10.1186/s40537-024-00957-y.
- [13] D. Ajish, "The significance of artificial intelligence in zero trust technologies: a comprehensive review," *Journal of Electrical Systems and Information Technology*, vol. 11, no. 1, Aug. 2024, doi: 10.1186/s43067-024-00155-z.
- [14] Y. I. Alzoubi, A. Mishra, and A. E. Topcu, "Research trends in deep learning and machine learning for cloud computing security," *Artificial Intelligence Review*, vol. 57, no. 5, May 2024, doi: 10.1007/s10462-024-10776-5.
- [15] S. Sarkar, G. Choudhary, S. K. Shandilya, A. Hussain, and H. Kim, "Security of zero trust networks in Cloud Computing: A Comparative review," *Sustainability*, vol. 14, no. 18, p. 11213, Sep. 2022, doi: 10.3390/su141811213.
- [16] S. Saeed, S. A. Suayyid, M. S. Al-Ghamdi, H. Al-Muhaisen, and A. M. Almuhaideb, "A Systematic Literature Review on Cyber Threat Intelligence for Organizational Cybersecurity Resilience," *Sensors*, vol. 23, no. 16, p. 7273, Aug. 2023, doi: 10.3390/s23167273.