

AI-Based Intrusion Detection System

AYUSH KUMAR SINGH, MR. RAM KUMAR SHARMA

B. Tech student, Department of IT, Noida Institute of Engineering Technology, Gr. Noida
Assistant professor, Department of IT, Noida Institute of Engineering Technology, Gr. Noida

Abstract - The increasing dependence on digital networks has made cybersecurity a critical concern, especially with the rapid rise in sophisticated and previously unseen cyber-attack. Conventional intrusion detection systems mainly rely on fixed signatures and predefined rules, which limit their ability to detect new or evolving threats. As a result, these systems often fail to provide timely and accurate protection in dynamic environments. The proposed model integrates algorithms such as Random Forest and Support Vector Machines to classify network activities while also incorporating anomaly detection to identify unusual behavior patterns. It makes use of widely accepted datasets like NSL-KDD and focuses on extracting meaningful features such as traffic flow characteristics, session duration, and protocol usage patterns. In addition, a new performance metric, termed the Detection Efficiency Score (DES), is introduced to evaluate the system by considering accuracy, false alarm rate, and response collectively. The experimental evaluation shows that the proposed system significantly outperforms traditional methods, achieving higher detection rates and reducing false positives. Moreover, it demonstrates the ability to detect unknown attacks by learning behavioral patterns rather than relying solely on predefined signatures. The study also highlights the potential for deploying such systems in environments with limited computational resources and suggests future improvements using deep learning models and privacy-preserving techniques like federated learning to ensure scalability and adaptability.

Keywords - network security, anomaly detection, Random Forest, support vector machine, NSL-KDD dataset, cyber-attack detection, artificial intelligence, network traffic analysis.

1. INTRODUCTION

The rapid expansion of digital infrastructure and internet-based services has significantly increased the exposure of computer networks to cyber threats. From financial systems to cloud-based applications, modern organizations rely heavily on interconnected environments, making them attractive targets for malicious activities such as denial-of-service attacks, data breaches, and unauthorized access. Traditionally, Intrusion Detection Systems (IDS) have been employed as a defensive mechanism to monitor network traffic and identify suspicious behaviour. However, these conventional systems are largely dependent on predefined signatures and rule-based mechanisms, which limits their ability to detect previously unseen or evolving attack patterns. The limitations of static detection approaches have become increasingly evident.

The integration of Artificial Intelligence (AI) and Machine Learning (ML) techniques into cybersecurity has opened new avenues for building intelligent and adaptive intrusion detection systems. Machine learning models can analyse large volumes of network data, identifying hidden patterns, and learning from past attack behaviours to improve detection accuracy. Unlike traditional systems, AI-based IDS can detect anomalies that

deviate from normal network behaviour, thereby enabling the identification of zero-day attacks and unknown threats. Despite these advancements, many existing AI-based solutions focus primarily on improving classification accuracy without addressing other critical aspects such as false alarm rates, real-time responsiveness, and adaptability to dynamic network environments.

This paper argues that an effective intrusion detection system must go beyond isolated improvements in detection accuracy and instead adopt a more integrated and adaptive approach. Specifically, it should combine robust feature extraction techniques, efficient machine learning models, and real-time anomaly detection mechanisms within a unified framework. Additionally, the system should be capable of continuously learning from new data and adapting to emerging threats without requiring frequent manual updates. Achieving this level of adaptability is essential for maintaining security in modern, large-scale networks where attack patterns evolve rapidly.

This paper proposes an AI-Based Intrusion Detection System (AI-IDS) that integrates supervised learning algorithms, such as Random Forest and Support Vector Machines, with anomaly-based detection techniques. The system leverages benchmark datasets and focuses on extracting meaningful network features, including traffic behaviour, session characteristics, and protocol usage patterns, to improve detection performance. Furthermore, a comprehensive evaluation metric, termed the Detection Efficiency Score (DES), is introduced to assess system performance by jointly considering detection accuracy, false positive rate, and response time.

First, we present a unified framework that combines classification and anomaly detection to enhance overall detection capability. Second, we propose a feature-driven approach that improves the model's ability to distinguish between normal and malicious traffic. Third, we introduce a new evaluation metric that provides a balanced assessment of system performance. Finally, we validate the effectiveness of the proposed system through experimental analysis, demonstrating its superiority over traditional intrusion detection methods in terms of accuracy, efficiency, and adaptability.

2. RELATED WORK

2.1 Traditional Intrusion Detection Systems

Intrusion Detection Systems (IDS) have been a fundamental component of network security for several decades. Early IDS solutions were primarily based on **signature detection**, where patterns were stored in a database and incoming network traffic was matched against these signatures. While such systems are effective in detecting previously identified threats, they lack the ability to recognize new or evolving attack patterns. As a result, signature-based systems often fail against zero-day attacks and require continuous manual updates, making them less suitable for dynamic environments.

2.2 Machine Learning in Intrusion Detection

Machine Learning techniques into IDS has significantly improved detection capabilities. Supervised learning algorithms such as Decision Trees, Support Vector Machines and Random Forest have been widely used to classify network traffic into normal and malicious categories. These models learn patterns from labelled datasets and can achieve high accuracy when trained on representative data.

Deep learning approaches, including Artificial Neural Networks and Recurrent Neural Networks, have also been explored for intrusion detection. These models can capture complex, non-linear relationships in network data, making them suitable for detecting sophisticated attack patterns. However, they often require large computational resources and lack interpretability, which limits their practical deployment in real-time systems.

2.3 Datasets and Benchmarking

Publicly available datasets have played a crucial role in the development and evaluation of intrusion detection systems. The **KDD Cup 99 dataset** has been widely used as a benchmark for evaluating IDS models. However, it contains redundant records and does not accurately represent modern network traffic.

To address these limitations, the **NSL-KDD dataset** was introduced as an improved version, reducing redundancy and providing a more balanced dataset for training and testing. More recent datasets aim to capture real-world traffic patterns, but challenges remain in ensuring diversity, scalability, and relevance to current cyber threats.

2.4 AI-Based Adaptive Detection Systems

Recent research has focused on developing adaptive IDS that can continuously learn from new data and adjust to evolving threats. Hybrid models combining supervised and unsupervised learning have shown promise in improving detection rates while reducing false positives. Additionally, real-time monitoring systems have been developed to enhance responsiveness and minimize detection delays.

However, most existing systems address individual aspects of intrusion detection, such as classification or anomaly detection, without integrating them into a unified framework. This fragmented approach limits their effectiveness in handling complex and dynamic cyber environments.

2.5 Research Gap

From the existing literature, it is evident that current intrusion detection systems face several limitations. Traditional systems lack adaptability, while many AI-based approaches focus narrowly on accuracy without considering efficiency and real-time performance. There is a need for a comprehensive system that combines robust feature extraction, efficient machine learning models, and real-time anomaly detection within a single framework.

This paper addresses this gap by proposing an AI-Based Intrusion Detection System that integrates multiple detection techniques and introduces a unified evaluation metric to balance accuracy, false positives, and response time.

3. THEORETICAL FOUNDATIONS

3.1 Intrusion Detection as a Classification Problem

Intrusion detection can be formulated as a binary or multi-class classification problem, where network traffic is categorized as either normal or malicious. Let a dataset consist of feature vectors representing network activities. A machine learning model is trained to map these features to corresponding class labels. Supervised learning techniques rely on labelled datasets to learn decision boundaries, enabling the system to classify new, unseen data.

The effectiveness of the model depends on the quality of features extracted from network traffic, such as packet size, connection duration, protocol type, and traffic frequency

3.2 Anomaly Detection as a Behavioral Model

In addition to classification, anomaly detection plays a crucial role in identifying unknown threats. In this approach, normal network behaviour is modelled based on historical data, and any significant deviation from this baseline is considered suspicious.

Let normal behaviour be represented by a distribution of network features. Any instance that falls outside a predefined threshold is flagged as an anomaly. This approach is particularly useful for detecting zero-day attacks, as it does not rely on predefined signatures.

3.3 Feature Extraction and Representation

Feature extraction plays a crucial role in the performance of an AI-based intrusion detection system, as the effectiveness of any machine learning model largely depends on the quality of input data. In this process, raw network traffic is converted into structured and meaningful features that can be analyzed by the model. These features typically capture various aspects of network behavior, such as traffic flow characteristics, connection duration, protocol types, packet sizes, and the frequency of specific activities like login attempts or data transfers. By transforming unprocessed network data into relevant attributes, the system becomes capable of identifying patterns associated with both normal and malicious activities. Additionally, feature selection techniques are often applied to remove redundant or irrelevant data, which helps reduce computational complexity and improves model accuracy. Proper representation of features not only enhances detection performance but also ensures that the system can generalize well to new and unseen attack scenarios

3.4 Detection Efficiency Score

To evaluate the overall performance of the proposed system, we introduce a metric called the **Detection Efficiency Score (DES)**. This metric combines multiple performance factors into a single value, enabling a more comprehensive assessment of the system.

$$DES = \alpha \cdot Accuracy + \beta \cdot (1 - False\ Positive\ Rate) + \gamma \cdot (1 - Detection\ Latency)$$

This combined approach ensures that the system is not only precise in identifying attacks but also efficient in minimizing incorrect alerts and responding quickly to potential threats. In practical scenarios, a high accuracy model may still be

ineffective if it generates excessive false alarms or responds too slowly to attacks. Therefore, DES provides a balanced assessment by integrating these parameters into a single score using weighted contributions. The weighting factors can be adjusted based on system requirements, allowing flexibility in prioritizing speed, reliability, or accuracy. This metric enables better comparison between different models and supports the development of intrusion detection systems that are both effective and practical for real-world deployment.

4. THE AI-IDS FRAMEWORK

4.1 Architectural Overview

The proposed AI-Based Intrusion Detection System (AI-IDS) is designed as a layered architecture that enables efficient monitoring, analysis, and detection of network threats in real time. The system is organized into five functional layers that work together to ensure accurate and adaptive intrusion detection. This raw data is then passed to the Preprocessing Layer, where noise and irrelevant information are removed, and the data is transformed into a structured format suitable for analysis. The next layer, the Feature Extraction and Modeling Layer, derives meaningful attributes from the processed data and applies machine learning algorithms such as Random Forest and Support Vector Machines to classify network behavior. Alongside classification anomaly detection mechanisms are also employed to identify unusual patterns that may indicate unknown or zero-day attacks. The Decision Layer evaluates the model outputs and determines whether the observed activity is normal or malicious, triggering alerts if necessary. Finally, the Response and Alert Layer communicates the detected threats to system administrators or automated defense mechanisms for further action. These layers are interconnected through a modular communication structure, allowing flexibility in upgrading or replacing individual components, such as integrating new algorithms or improving feature extraction methods, without affecting the overall system functionality.




Fig. 1: AI-Based Intrusion Detection System

4.2 Threat Detection from Network Behavioral Signals

The core idea of the proposed AI-Based Intrusion Detection System is that malicious activity can be effectively identified by analyzing behavioral patterns in network traffic rather than

relying solely on predefined attack signatures. The system extracts several important features from network data that help distinguish between normal and suspicious activities. These features include connection duration, packet size variation, frequency of requests, number of failed login attempts, and protocol usage patterns. For example, unusually long or extremely short connection durations may indicate abnormal behavior, while a high number of repeated failed login attempts may signal a brute-force attack. Similarly, irregular traffic spikes or unusual protocol usage can be indicative of denial-of-service attacks or unauthorized access attempts. These extracted features are then processed using machine learning models to identify hidden patterns associated with cyber threats. The model is trained on benchmark datasets such as NSL-KDD, enabling it to learn both normal and attack behaviors. This behavioral analysis approach allows the system to detect unknown or zero-day attacks with reasonable accuracy. However, it is important to note that the effectiveness of the model depends on the quality of training data and may require recalibration when deployed in different network environments.

4.3 Knowledge Modeling of Network Behavior

In the proposed system, network behavior is modeled using supervised machine learning techniques that classify traffic into normal or malicious categories. The system uses algorithms such as Random Forest and Support Vector Machines, which are trained on labeled datasets containing various types of network activities. Each network instance is represented as a feature vector, and the model learns to associate these features with corresponding class labels. During operation, the trained model continuously updates its predictions based on incoming traffic data, allowing it to adapt to new patterns over time. The classification process helps in identifying different categories of attacks such as denial-of-service, probing, and unauthorized access. By maintaining an updated model of network behavior, the system ensures accurate detection and reduces the chances of misclassification. This adaptive learning capability is essential for handling evolving cyber threats.

4.4 The Decision Engine

The decision engine acts as the core component responsible for analyzing the output generated by the machine learning models and determining whether a particular network activity is malicious or normal. It considers multiple factors, including classification results, anomaly scores, and predefined thresholds, to make informed decisions. For instance, if the probability of an attack exceeds a certain threshold, the system flags the activity as suspicious. Additionally, the engine evaluates the severity of the detected threat by considering parameters such as frequency, impact, and deviation from normal behavior. Based on this evaluation, the system categorizes threats into different risk levels, such as low, medium, or high risk. This layered decision-making process ensures that alerts are generated accurately, and unnecessary false alarms are minimized. The decision engine plays a critical role in maintaining the balance between detection accuracy and system efficiency.

4.5 Automated Response and Alert Mechanism

Once a threat is detected and classified, the system initiates an appropriate response through its automated alert mechanism. Alerts can be generated in various forms, such as dashboard notifications, email alerts, system logs to inform administrators

about potential security breaches. In addition to alert generation, the system can also perform automated actions such as blocking suspicious IP addresses, terminating compromised sessions or updating firewall rules to prevent further intrusion. All detected events are logged for future analysis and auditing purposes, which helps in improving system performance and understanding attack patterns. The response mechanism is designed to be flexible, allowing administrators to configure actions based on the severity of the threat. This ensures that the system not only detects attacks but also actively contributes to mitigating their impact in real time.

5. METHODOLOGY

5.1 Implementation

The proposed AI-Based Intrusion Detection System (AI-IDS) was implemented as a scalable and modular application designed to monitor and analyze network traffic in real time. The backend of the system was developed using Python, leveraging libraries such as Scikit-learn and Pandas for data processing and machine learning model development. The system was trained using benchmark datasets, particularly the NSL-KDD dataset, which provides labeled instances of normal and malicious network activities. For classification, algorithms such as Random Forest and Support Vector Machines were employed due to their effectiveness in handling structured data and achieving high detection accuracy. The frontend interface was designed to provide a clear visualization of detected threats, system alerts, and performance metrics, enabling administrators to monitor network security efficiently. The system processes incoming network traffic performs preprocessing steps such as data cleaning and normalization and extracts relevant features before passing them to the trained models for classification and anomaly detection. To ensure efficient performance, optimized data handling and model execution techniques were applied, resulting in low response time suitable for real-time detection. Logging mechanisms were incorporated to record system activity and detect events, allowing further analysis and model improvement. Additionally, the system was designed using modular architecture principles, enabling easy integration of new algorithms or datasets in the future. Basic security measures, including input validation and controlled access to system components, were implemented to ensure reliability and protect sensitive data during operation.

5.2 Dataset and Experimental Setup

The experimental evaluation of the proposed AI-Based Intrusion Detection System was conducted using benchmark datasets commonly used in cybersecurity research. In particular, the NSL-KDD dataset was selected due to its improved quality over earlier datasets and its balanced representation of normal and malicious network activities. The dataset contains multiple types of cyber-attack including denial-of-service (DoS), probing attacks, user-to-root (U2R), and remote-to-local (R2L) intrusions. Each data instance consists of a set of features describing network behavior, such as protocol type, connection duration, and traffic patterns. The dataset was divided into training and testing subsets to ensure unbiased evaluation of the model. Prior to training, preprocessing steps such as normalization, encoding of categorical variables, and removal of redundant data were performed to improve model performance and accuracy.

5.3 Experimental Procedure

The implementation process began with preprocessing the dataset to convert raw network data into a structured format suitable for machine learning algorithms. Feature extraction techniques were applied to identify relevant attributes that contribute significantly to intrusion detection. The processed data was then used to train supervised learning models, specifically Random Forest and Support Vector Machines, which were selected for their robustness and efficiency in classification tasks. After training, the models were tested using unseen data to evaluate their ability to correctly classify network activities. In addition to classification, anomaly detection techniques were incorporated to identify unusual patterns that may indicate unknown or zero-day attacks. The system continuously monitored incoming data, applied the trained models, and generated predictions in real time. Performance metrics were calculated based on the model outputs, and system logs were maintained to track detection behavior and analyze performance trends.

5.4 Evaluation Metrics

The performance of the proposed intrusion detection system was evaluated using multiple metrics to ensure a comprehensive assessment. The primary evaluation measures included accuracy, precision, recall, and F1-score, which collectively indicate the effectiveness of the classification model. Accuracy measures the overall correctness of predictions, while precision and recall evaluate the model's ability to correctly identify malicious activities without generating excessive false alarms. The F1-score provides a balanced measure of precision and recall. In addition to these standard metrics, the Detection Efficiency Score (DES) was used as a combined evaluation metric that considers accuracy, false positive rate, and detection latency. Secondary evaluation factors included response time, system scalability, and the ability to detect unknown attacks. These metrics provide a holistic understanding of system performance and ensure that the proposed model is both accurate and efficient for real-world deployment.

6. RESULTS

The performance of the proposed AI-Based Intrusion Detection System was evaluated using the NSL-KDD dataset, and the results were compared with traditional machine learning approaches. The dataset was divided into training and testing sets to ensure unbiased evaluation. The proposed model demonstrated consistent performance across different types of cyber-attack including probing and unauthorized access attempts. Initial validation confirmed that the dataset distribution between training and testing phases was balanced, ensuring reliable and fair model evaluation.

6.1 Primary Outcomes

The experimental results indicate that the proposed AI-IDS achieved significant improvements in detection performance compared to conventional approaches. The system recorded an overall accuracy of approximately 92.6%, outperforming baseline models that achieved an average accuracy of around 76%. In addition, the false positive rate was reduced to 8.3%, compared to nearly 20% observed in traditional systems. The model also demonstrated improved precision and recall values,

indicating its effectiveness in correctly identifying malicious activities while minimizing incorrect alerts. The Detection Efficiency Score (DES), which combines accuracy, false positive rate, and response time, was significantly higher for the proposed system, confirming its overall effectiveness and reliability in real-world scenarios.

6.2 Detection Behavior and System Performance

Further analysis of system performance revealed that the proposed model maintained stable detection behavior across different attack categories. The system was particularly effective in identifying high frequency attacks such as denial-of-service, while also demonstrating reasonable performance in detecting low-frequency and complex attacks such as user-to-root intrusions. Unlike traditional systems, which often struggle with unseen attack patterns, the anomaly detection component of the proposed model enabled it to identify unusual behavior with improved consistency. The system also showed low detection latency, ensuring that threats were identified and reported in near real time.

6.3 Alert Generation and Response Efficiency

An important observation from the experimental results was the system's ability to generate accurate alerts without overwhelming administrators with excessive notifications. The reduction in false positives contributed to more reliable alert generation, allowing security personnel to focus on genuine threats. Additionally, the system's response mechanism was effective in triggering timely actions such as logging suspicious activities and initiating preventive measures. The balance between sensitivity and specificity ensured that the system remained both efficient and practical for deployment in real-world environments.

6.4 Robustness and Reliability

To validate the robustness of the proposed model, additional tests were conducted using different subsets of the dataset and alternative evaluation methods. The results remained consistent across these variations, indicating that the system is stable and reliable under different conditions. The model also demonstrated good generalization capability, maintaining performance when exposed to previously unseen data. These findings suggest that the proposed AI-IDS is not only accurate but also adaptable to diverse network environments. However, it is important to note that performance may vary depending on the quality and diversity of the training data, and further validation on real-time network traffic is recommended.

7. DISCUSSION

7.1 Theoretical Implications

The results of this study support the idea that combining multiple machine learning techniques within a unified intrusion detection framework can significantly improve cybersecurity performance. Instead of relying on a single detection approach, the integration of classification models and anomaly detection

enables the system to identify both known and unknown threats effectively. The introduction of a composite evaluation metric, such as the Detection Efficiency Score (DES), provides a more balanced understanding of system performance by considering accuracy, false positives, and response time together. This highlights that intrusion detection should not be viewed solely as a classification problem, but rather as a multi-dimensional optimization challenge involving efficiency, reliability, and adaptability.

7.2 Practical Implications

From a practical perspective, the proposed system demonstrates that AI-based intrusion detection can be implemented using standard computational resources without requiring specialized hardware. By leveraging commonly available datasets and efficient machine learning algorithms, the system can be deployed in real-world environments such as enterprise networks, cloud systems, and educational institutions. The use of behavioral features extracted from network traffic allows continuous monitoring without interrupting normal operations. Furthermore, the modular design of the system ensures that new algorithms or updated datasets can be integrated easily, making it suitable for evolving cybersecurity requirements.

7.3 Role of AI in Intrusion Detection

An important consideration in the design of the proposed system is the role of artificial intelligence. Rather than treating AI as a standalone solution, this work positions machine learning models as part of a coordinated system that includes feature extraction, decision-making, and response mechanisms. This integrated approach ensures that the system not only detects threats but also provides meaningful insights into network behavior. The combination of supervised learning for classification and anomaly detection for identifying unusual patterns creates a more robust and flexible security solution compared to traditional methods.

7.4 Security and Privacy Considerations

The use of network traffic data for intrusion detection raises important privacy and security concerns. Network logs may contain sensitive information, and improper handling of such data could lead to privacy risks. To address this issue, the system can be extended using privacy-preserving techniques such as federated learning, where models are trained locally and only model updates are shared instead of raw data. This approach reduces the risk of data leakage while maintaining model performance. Additionally, secure data handling practices and access control mechanisms are essential for ensuring system reliability and protecting sensitive information.

7.5 Limitations

Despite the promising results, the proposed system has certain limitations. First, the evaluation was conducted using benchmark datasets, which may not fully represent real-world network environments. Second, the performance of the model depends heavily on the quality and diversity of training data. Third, while the system performs well in detecting common attack types, its effectiveness against highly sophisticated or stealthy attacks requires further investigation. Additionally, real-time deployment in large-scale networks may introduce challenges related to computational overhead and scalability.

These limitations highlight the need for further research and validation in practical environments.

8. FUTURE WORK

Several directions can be explored to enhance the proposed intrusion detection system. Future work may include testing the model on real-time network traffic to evaluate its performance in live environments. The integration of deep learning techniques, such as neural networks and transformer-based models, could further improve detection accuracy for complex attack patterns. Additionally, incorporating advanced feature selection and dimensionality reduction methods may enhance system efficiency. The use of federated learning and privacy-preserving mechanisms can be further developed to ensure secure deployment across multiple organizations. Expanding the system to handle large-scale distributed networks and cloud environments is another important area for future research. Finally, continuous learning mechanisms can be introduced to enable the system to adapt automatically to emerging cyber threats.

Another important area of future research is the deployment of the system in real-time, large-scale network environments. Testing the model on live traffic data from enterprise or cloud infrastructures would provide deeper insights into its practical effectiveness and scalability. This would also help in evaluating system performance under high traffic loads and dynamic network conditions, which are not fully represented in benchmark datasets.

The incorporation of advanced feature engineering techniques, including automated feature selection and dimensionality reduction methods, can further optimize system performance. Techniques such as Principal Component Analysis (PCA) and feature importance ranking may help reduce computational overhead while maintaining or improving detection accuracy.

9. CONCLUSION

In conclusion, this paper presents an AI-Based Intrusion Detection System that enhances network security by combining machine learning techniques with anomaly detection mechanism. The proposed approach effectively overcomes the limitations of traditional intrusion detection systems by improving detection accuracy, reducing false positives, and ensuring faster response times. The integration of multiple components within a unified and modular architecture enables the system to operate efficiently while maintaining adaptability in dynamic network environments.

The introduction of the Detection Efficiency Score (DES) provides a comprehensive evaluation framework that considers accuracy, false alarm rate, and response time together, offering a more realistic assessment of system performance. Experimental results demonstrate that the proposed model performs significantly better than conventional methods in detecting various types of cyber-attacks, including unknown and zero-day threats.

Furthermore, the modular design of the system allows easy integration of new algorithms, datasets, and detection strategies, making it scalable and future ready. By focusing on behavioural analysis rather than relying solely on predefined signatures, the system can identify evolving attack patterns with greater reliability. Overall, this study highlights the growing importance

of artificial intelligence in cybersecurity and provides a strong foundation for developing intelligent, adaptive, and efficient intrusion detection systems in the future.

ACKNOWLEDGEMENT

The authors would like to express their gratitude to the faculty members of the Department of Information Technology for their guidance and support throughout this research work. The authors also acknowledge the availability of publicly accessible datasets such as NSL-KDD, which made experimental evaluation possible. Special thanks are extended to peers and contributors who provided valuable feedback during the development and testing of the system. The authors declare that there are no conflicts of interest related to this work.

The authors also acknowledge the contribution of publicly available datasets, particularly the NSL-KDD dataset, which served as a foundation for training and evaluating the proposed intrusion detection system. Access to such well-structured datasets greatly facilitated the experimental analysis and validation process.

Special thanks are extended to fellow students and peers who provided constructive feedback during the implementation and testing phases of the project. Their inputs helped in identifying practical challenges and improving the overall system performance.

The authors would also like to recognize the support of the institution for providing the necessary computational resources and an environment conducive to research and innovation.

Finally, the authors declare that this research work was carried out independently and ethically, and there are no conflicts of interest associated with this study.

REFERENCES

1. V. Alevin, E. A. McLaughlin, R. A. Glenn, and K. R. Koedinger, "Instruction based on adaptive learning technologies," in *Handbook of Research on Learning and Instruction*, 2nd ed., R. E. Mayer and P. Alexander, Eds. New York, NY, USA: Routledge, 2016, pp. 522–560.
2. J. R. Anderson, A. T. Corbett, K. R. Koedinger, and R. Pelletier, "Cognitive tutors: Lessons learned," *The Journal of the Learning Sciences*, vol. 4, no. 2, pp. 167–207, 1995.
3. R. S. J. d. Baker, S. K. D'Mello, M. M. T. Rodrigo, and A. C. Graesser, "Better to be frustrated than bored: The incidence, persistence, and impact of learners' cognitive-affective states during interactions with three different computer-based learning environments," *International Journal of Human-Computer Studies*, vol. 68, no. 4, pp. 223–241, 2008.
4. H. Bastani, O. Bastani, A. Sungu, H. Ge, Ö. Kabakcı, and R. Mariman, "Generative AI can harm learning," SSRN Working Paper 4895486, 2024.
5. B. S. Bloom, "The 2-sigma problem: The search for methods of group instruction as effective as one-to-one tutoring," *Educational Researcher*, vol. 13, no. 6, pp. 4–16, 1984.
6. J. R. Carbonell, "AI in CAI: An artificial-intelligence approach to computer-assisted instruction," *IEEE Transactions on Man-Machine Systems*, vol. 11, no. 4, pp. 190–202, 1970.

7. B. Clement, D. Roy, P. Y. Odier, and M. Lopes, "Multi-armed bandits for intelligent tutoring systems," *Journal of Educational Data Mining*, vol. 7, no. 2, pp. 20–48, 2015.
8. A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modelling the acquisition of procedural knowledge," *User Modelling and User-Adapted Interaction*, vol. 4, no. 4, pp. 253–278, 1995.
9. A. Ghosh, N. Heffernan, and A. S. Lan, "Context-aware attentive knowledge tracing," in *Proc. 26th ACM SIGKDD Int. Conf. Knowledge Discovery & Data Mining*, 2020, pp. 2330–2339.
10. S. G. Hart and L. E. Stevland, "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research," *Advances in Psychology*, vol. 52, pp. 139–183, 1988.