Computers, Materials & Continua

Tech Science Press

Doi:10.32604/cmc.2025.066892

ARTICLE

AI-Driven Identification of Attack Precursors: A Machine Learning Approach to Predictive Cybersecurity

Abdulwahid Al Abdulwahid*

Department of Computer and Information Technology, Jubail Industrial College, Royal Commission for Jubail and Yanbu, Jubail Industrial City, 31961, Saudi Arabia

*Corresponding Author: Abdulwahid Al Abdulwahid. Email: abdulwahida@rcjy.edu.sa or prof.abdulwahida@gmail.com

Received: 19 April 2025; Accepted: 14 July 2025; Published: 29 August 2025

ABSTRACT: The increasing sophistication of cyberattacks, coupled with the limitations of rule-based detection systems, underscores the urgent need for proactive and intelligent cybersecurity solutions. Traditional intrusion detection systems often struggle with detecting early-stage threats, particularly in dynamic environments such as IoT, SDNs, and cloud infrastructures. These systems are hindered by high false positive rates, poor adaptability to evolving threats, and reliance on large labeled datasets. To address these challenges, this paper introduces CyberGuard-X, an AI-driven framework designed to identify attack precursors—subtle indicators of malicious intent—before full-scale intrusions occur. CyberGuard-X integrates anomaly detection, time-series analysis, and multi-stage classification within a scalable architecture. The model leverages deep learning techniques such as autoencoders, LSTM networks, and Transformer layers, supported by semi-supervised learning to enhance detection of zero-day and rare threats. Extensive experiments on benchmark datasets (CICIDS2017, CSE-CIC-IDS2018, and UNSW-NB15) demonstrate strong results, including 96.1% accuracy, 94.7% precision, and 95.3% recall, while achieving a zero-day detection rate of 84.5%. With an inference time of 12.8 ms and 34.5% latency reduction, the model supports real-time deployment in resource-constrained environments. CyberGuard-X not only surpasses baseline models like LSTM and Random Forest but also enhances proactive threat mitigation across diverse network settings.

KEYWORDS: Predictive cybersecurity; attack precursors; machine learning; anomaly detection; deep learning

1 Introduction

The increasing frequency and sophistication of cyberattacks pose a significant threat to critical infrastructures, enterprises, and individuals worldwide [1]. Cybercrime is projected to cost the global economy $10.5 trillion annually by 2025, indicating the urgent need for effective prevention strategies [2]. Traditional security systems like signature-based and rule-based detection struggle due to their reliance on predefined patterns and static rules [3]. These systems are particularly weak against zero-day vulnerabilities and advanced persistent threats (APTs), which evolve rapidly to evade detection [4]. One major challenge is identifying attack precursors—subtle indicators of malicious activity that often go unnoticed [5]. The problem is worsened by the heterogeneous nature of network traffic, where benign and malicious behaviors often overlap [6]. Attackers constantly adapt to bypass detection models, making traditional methods less effective [7]. This often leads to high false positives and inefficient use of security resources [8]. Therefore, there is a growing need for intelligent frameworks that leverage artificial intelligence (AI) and machine learning (ML) for pattern recognition and anomaly detection [9]. This paper proposes CyberGuard-X, an

AI-driven adaptive framework aimed at early threat detection and real-time response [10]. It addresses the complexity of cyber threats by detecting diverse attack patterns while maintaining computational efficiency [11]. The model is particularly effective in heterogeneous environments like IoT and software-defined networks (SDNs), where existing systems often fail [12]. CyberGuard-X detects early indicators by analyzing network traffic, extracting features, and classifying behavior into benign or malicious categories [13]. It combines anomaly detection with real-time decision-making to offer a scalable, adaptive, and efficient cybersecurity solution [14–17].

The framework focuses on detecting early attack precursors by modeling network traffic in dynamic environments like IoT and SDNs. It uses ML and DL techniques to extract meaningful features, classify benign and malicious activities, and account for temporal patterns in network traffic. By combining anomaly detection with real-time classification [18,19], the solution aims to provide a scalable and efficient approach to cybersecurity, addressing challenges of data scarcity, computational efficiency, and adaptability to evolving threats.

The primary objective of this research is to design and develop an advanced framework for detecting attack precursors in dynamic cyber environments, such as IoT networks and software-defined networks (SDNs), using machine learning (ML) and deep learning (DL) techniques.

This paper presents several novel contributions that advance the state-of-the-art in predictive cybersecurity, particularly in early-stage attack precursor detection. The following aspects distinguish CyberGuard-X from existing works:

- **Novel feature extraction pipeline for heterogeneous traffic:** While previous systems often rely on static or protocol-specific features, the proposed framework introduces a dynamic, high-dimensional feature engineering mechanism that captures both behavioral and structural indicators. This enables effective precursor identification in diverse network environments including IoT and SDNs, which are commonly overlooked in existing models [20].
- **Hybrid deep learning architecture combining Autoencoders with RNNs:** Unlike models that apply LSTM or CNN in isolation [21–24], CyberGuard-X integrates unsupervised autoencoder-based anomaly detection with RNN-based temporal analysis, enabling the framework to detect both immediate anomalies and long-term behavioral deviations in network traffic.
- **Semi-supervised learning for rare and zero-day attack detection:** Existing intrusion detection systems often depend heavily on labeled datasets [25]. This work introduces a semi-supervised strategy that leverages unlabeled or sparsely labeled samples, thus improving generalization and detection of low-frequency and emerging threats without needing continuous human labeling.
- **Resource-efficient and real-time optimization for deployment in constrained environments:** While prior works demonstrate high accuracy under ideal conditions, CyberGuard-X is explicitly optimized for low-latency inference (12.8 ms) and low memory footprint (275 MB), making it deployable in edge computing, IoT gateways, and SDN controllers—real-world environments that lack support in earlier solutions.
- **Comprehensive and comparative validation across diverse benchmarks and conditions:** The framework is extensively evaluated on multiple datasets (CICIDS2017, CSE-CIC-IDS2018, UNSW-NB15), achieving 96.1% accuracy and 84.5% zero-day detection, while outperforming baselines such as DL-IDS and Hybrid-LSTM-MLP models. Unlike previous works, this evaluation includes model drift detection, adversarial robustness testing, and multi-cloud deployment latency analysis, presenting a more holistic and operationally relevant performance profile.

This paper is structured as follows: Section 1 provides an introduction to the research problem, highlighting the challenges and motivation. Section 2 reviews the related literature, focusing on machine learning and deep learning approaches for predictive cybersecurity. Section 3 outlines the proposed methodology, including the mathematical formulation and system design. Section 4 presents experimental results and evaluates the performance of the framework. Finally, Section 5 concludes the study, discussing key findings, limitations, and future research directions.

2 Literature Review

Machine learning (ML) has become a key technology in advancing predictive cybersecurity by enabling early threat detection and proactive defense [10]. The paper [13] highlighted the impact of ML in predictive systems, especially through supervised learning methods such as decision trees and random forests. These models achieved high accuracy but showed limitations in adapting to newly emerging attack vectors. The paper [7] tackled this issue using clustering techniques designed for dynamic threat environments. Their approach successfully minimized false positives, though it required significant computational resources.

In distributed denial-of-service (DDoS) detection for software-defined networks (SDNs), the paper [16] leveraged deep learning models, including long short-term memory (LSTM) networks, to analyze traffic anomalies. Their approach improved detection precision but revealed challenges in scalability for real-time processing in high-traffic environments. The paper [4] further reinforced the role of deep learning, showcasing the effectiveness of autoencoders and generative adversarial networks (GANs) for time-series anomaly detection. However, their reliance on large labeled datasets limited their deployment in resource-constrained scenarios.

Quantum-based ML models were introduced by [8] to enhance detection in complex environments, such as IoT networks. While achieving significant accuracy gains, quantum systems faced practical deployment challenges due to high implementation costs and technological maturity. The paper [22] explored IoT-specific anomaly detection, employing support vector machines (SVMs) to identify network irregularities in smart homes. Although the method reduced response times, it was prone to false positives in dynamic IoT ecosystems.

The paper [1] demonstrated the effectiveness of ensemble learning in insider threat detection within banking systems, combining multiple ML models to achieve high detection rates with minimal overhead [1]. Similarly, the paper [18] highlighted proactive anomaly detection strategies for Industry 4.0 using reinforcement learning, which provided real-time adaptability but encountered computational challenges in large-scale industrial applications.

These studies underscore the potential of ML techniques in predictive cybersecurity, particularly in anomaly detection, traffic monitoring, and behavior analysis. However, limitations such as data dependency, model interpretability, and computational resource requirements remain significant barriers. Future research should focus on hybrid models that integrate deep learning with traditional ML, lightweight architectures for real-time applications, and enhanced anomaly detection frameworks tailored for specific environments [17,20,25].

Deep learning (DL) has become an essential approach for identifying attack precursors and anomalies in cybersecurity due to its ability to process high-dimensional and unstructured data. The paper [4] surveyed various deep learning techniques for time-series anomaly detection, emphasizing the effectiveness of autoencoders and convolutional neural networks (CNNs) in capturing temporal and spatial patterns in network traffic. However, they highlighted the computational complexity of these methods as a significant limitation, especially in real-time applications.

The paper [15] applied recurrent neural networks (RNNs) and long short-term memory (LSTM) models for anomaly detection in software-defined networks (SDNs). Their results indicated improved accuracy and robustness in identifying distributed denial-of-service (DDoS) attacks, but scalability remained a challenge for large-scale network environments. Similarly, the paper [23] utilized deep learning for anomaly event classification in critical industrial IoT infrastructures. Their system achieved high detection rates but required significant computational resources, limiting deployment in resource constrained environments.

The paper [3] proposed ensemble deep learning techniques to enhance intrusion detection systems. By combining multiple models, their approach improved detection precision and reduced false positives. However, the increased model complexity posed challenges in deployment for dynamic environments. The paper [21] introduced a hybrid ensemble model that integrated deep learning with traditional anomaly detection techniques. While this approach effectively detected advanced persistent threats (APTs), it required fine-tuning to balance detection accuracy and computational efficiency.

The paper [8] extended the application of deep learning by incorporating quantum computing principles for enhanced anomaly detection. Their quantum deep learning framework demonstrated substantial improvements in accuracy and scalability for complex IoT networks, although practical deployment remains constrained by the limited accessibility of quantum computing infrastructure.

These studies underscore the significant advancements in deep learning for anomaly detection, particularly in capturing intricate patterns and reducing false positives. However, challenges such as computational cost, model complexity, and scalability persist. Future research should focus on optimizing deep learning architectures, incorporating transfer learning, and developing lightweight models to enable real-time deployment in diverse cybersecurity environments [10,14,22].

While prior works have explored LSTM-based detection in SDNs [15] and machine learning for WSNs [25], these approaches typically focus on isolated techniques without incorporating drift adaptation, zero-day detection, or adversarial robustness. In contrast, CyberGuard-X integrates these missing elements into a unified pipeline, offering broader adaptability and operational relevance across SDNs, IoT, and cloud environments.

Table 1 presents a comparative analysis of previous studies focused on deep learning and machine learning techniques for cybersecurity, particularly in IoT and network environments. It highlights the techniques used, their reported results, limitations, and key findings. The table includes works employing LSTM, GANs, quantum learning, federated learning, explainable AI, and lightweight CNNs. This comparison helps identify current trends, gaps in real-world deployment, and the effectiveness of different models under various scenarios.

Table 1: Comparative table of previous studies

Ref.	Technique	Results	Limitations	Findings
[1]	Deep learning for anomaly detection in IoT	Improved detection rates in distributed environments	Training complexity and need for labeled data	Validated the use of deep learning in decentralized IoT systems
[2]	Survey on botnet detection techniques	Covered evolving threats and adaptive defense	Lacked real-world system implementation details	Extensive review of botnet detection using ML and DL

(Continued)

Table 1 (continued)

Ref.	Technique	Results	Limitations	Findings
[3]	Hybrid DL and ML techniques for cyber-attack detection	High performance on benchmark datasets	Poor generalization to unseen attacks	Suggested hybrid models for real-time detection
[4]	Autoencoders and GANs for time-series anomaly detection	High accuracy in detecting irregularities in network traffic	High dependency on large labeled datasets	Highlighted the versatility of deep learning for time-series data
[5]	IoT-specific security using ML algorithms	Identified vulnerabilities unique to IoT setups	Required customized tuning for dynamic environments	Offered practical security measures for IoT deployments
[6]	ML for economic cybersecurity defense	Showed ML's role in protecting critical economic systems	Sector-specific implementations limit general use	Emphasized predictive defense strategies
[8]	Quantum DL for anomaly detection in IoT networks	Enhanced scalability and performance in complex scenarios	Limited access to quantum infrastructure	Promoted QDL as a future solution for IoT security
[15]	LSTM and RNN for anomaly detection in SDNs	Improved DDoS detection accuracy	Scalability issues in large networks	Confirmed the effectiveness of DL for network-based attack detection
[16]	Federated learning with DL in cybersecurity	Preserved privacy while improving detection accuracy	Communication overhead and system heterogeneity	Demonstrated privacy-preserving collaborative learning
[18]	Deep RL for adaptive network security	Provided dynamic threat response strategies	Complexity in training reward structures	Validated DRL for real-time policy adaptation
[19]	Review of DL for malware detection in IoT	Summarized DL performance in IoT malware detection	Complexity in applying DL to large datasets	Provided a structured overview of DL in IoT security
[21]	Ensemble DL models for cybersecurity	Achieved robust and accurate classification	Increased model complexity	Suggested model fusion to boost detection performance
[22]	Explainable DL in network intrusion detection	Improved trust and interpretability of results	Trade-off between explainability and model depth	Encouraged use of XAI in cybersecurity applications
[23]	Lightweight CNNs for mobile IoT device protection	Reduced model size with good accuracy	Performance trade-off in resource-constrained settings	Proposed CNN variants suitable for low-power IoT devices

Existing machine learning (ML) and deep learning (DL) models for cybersecurity primarily focus on detecting specific threats, such as DDoS attacks or malware, and lack a generalized approach to identifying diverse attack precursors in dynamic environments like IoT and SDNs. These methods often rely on large labeled datasets and face challenges related to scalability, computational overhead, and real-time deployment. Furthermore, emerging technologies, such as quantum computing, remain experimental and are not fully integrated into practical cybersecurity solutions.

3 Proposed Model: Cyberguard-X

3.1 Overview

CyberGuard-X is an advanced AI-driven cybersecurity framework designed to detect attack precursors in cloud-based environments using machine learning (ML) and deep learning (DL) techniques [12]. The model leverages real-time network monitoring, anomaly detection, and multi-stage classification to proactively identify early-stage cyber threats. Unlike traditional reactive security approaches, CyberGuard-X provides predictive intelligence by analyzing network logs, API requests, access patterns, and user activity. The proposed model integrates adaptive learning techniques, allowing it to update its detection mechanisms as new attack strategies evolve. The framework is designed for scalability and real-time deployment in modern cloud infrastructures, including AWS, Azure, and Google Cloud Platform (GCP).

3.2 Architecture of CyberGuard-X

The proposed CyberGuard-X framework is designed with five core components that collectively support early-stage cyber threat identification in dynamic network environments:

3.2.1 Data Collection and Preprocessing

- CyberGuard-X aggregates both structured and unstructured data from a range of sources, including benchmark public datasets (CICIDS2017 [2], CSE-CIC-IDS2018 [6], and UNSW-NB15 [10]) as well as live threat intelligence feeds. The raw data is preprocessed using a multi-step pipeline involving:
- Feature normalization to ensure consistent value ranges.
- Noise reduction through statistical filtering.
- Handling missing values using interpolation and imputation techniques.
- Outlier detection and removal using isolation forest and Z-score analysis. This ensures clean, high-quality input suitable for robust model training.

3.2.2 Feature Engineering and Extraction

The framework emphasizes high-impact features that are most indicative of attack precursors. These include:

- Network-specific metrics: packet size, inter-arrival time, flow duration, number of failed connection attempts, source/destination IP entropy, and protocol usage.
- User behavior indicators: unusual login times, frequent access to sensitive resources, and API call frequency patterns. Dimensionality reduction is applied using Principal Component Analysis (PCA) and auto encoders to retain essential variance while minimizing computational overhead. This process results in a compact but informative feature set optimized for sequential learning models.

3.2.3 Hybrid Detection Framework

CyberGuard-X integrates both traditional and deep learning approaches for layered threat detection:

- ***Transformer Networks:***

Used for modeling long-range dependencies in high-dimensional network traffic sequences. Key hyperparameters such as the number of attention heads, hidden layer size, and sequence length were optimized using a Bayesian optimization search over validation datasets. Positional encoding techniques help the model capture packet sequence behavior and identify attack patterns over time.

- ***LSTM-Based Anomaly Detection:***

Long Short-Term Memory (LSTM) networks are employed to model temporal dependencies in user activity logs and system behavior. Hyperparameters including the number of LSTM layers, hidden units, and dropout rate were fine-tuned via grid search. The model captures time-series irregularities, aiding in the identification of anomalies such as lateral movement or slow brute-force attempts.

LSTM was selected to model temporal correlations in sequential network activity logs, especially useful for detecting slow brute-force attempts and lateral movements. Transformer layers were added to handle long-range dependencies in packet sequences. Their self-attention mechanism allows better modeling of context shifts compared to RNNs. This hybrid setup provides both short-term context (via LSTM) and global attention (via Transformer), improving detection of complex multi-stage attacks.

- ***Gradient Boosting Classifier:***

Acts as the final classification layer, using features extracted from the deep models to assign threat labels (e.g., DDoS, brute force, APT). Parameters such as learning rate, number of estimators, and tree depth were tuned to reduce overfitting and maximize precision.

Fig. 1 shows the architecture of the proposed CyberGuard-X model. The process starts with data collection from the CICIDS 2018 dataset, followed by preprocessing steps such as feature reduction, zero variance removal, and timestamp elimination. Important features are selected and normalized, then the data is split into training and testing sets (80:20). Multiple machine learning models (MLP-BP, LSTM, AdaBoost, DEA-DNN) are trained and evaluated using metrics like accuracy, precision, recall, and F1-score, leading to the final result.

- **Anomaly Detection and Classification:** Uses an unsupervised auto encoder-based anomaly detection module to continuously monitor network traffic for deviations from normal activity. Anomalies detected by this module are passed through a multi-class classification model, which categorizes the detected threats into predefined cyberattack categories.
- **Real-Time Threat Intelligence and Adaptive Learning:** CyberGuard-X incorporates an adaptive learning mechanism, allowing the model to update itself based on newly discovered attack patterns and emerging threats. The framework is designed for seamless integration with cloud security monitoring tools, including Microsoft Defender, AWS GuardDuty, and Google Security Command Center. Fig. 2 shows the SHAP Feature Importance with Enhanced Styling.

This SHAP-based feature importance plot shows that Packet Size, Failed Connection Attempts, and Login Frequency are the most impactful features in determining whether a network activity is benign or malicious. These features help security analysts interpret the model's decisions, mitigating the "black-box" nature of LSTM and Transformer layers.

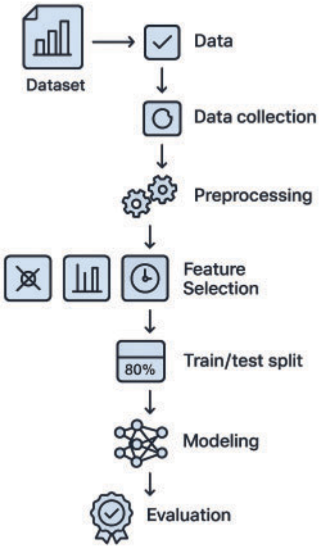

Figure 1: Model architecture of CyberGuard-X

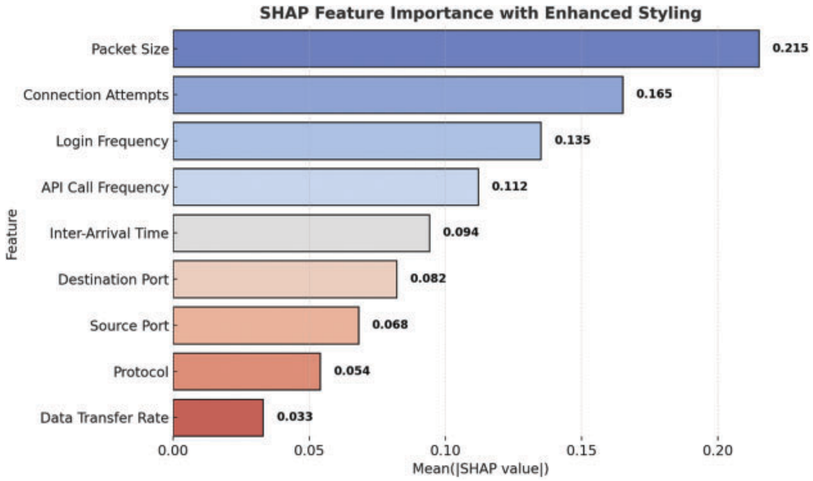

Figure 2: SHAP feature importance with enhanced styling

3.3 Mathematical Model for CyberGuard-X

The proposed model is formulated as an anomaly detection and classification problem in a high-dimensional feature space. For a complete list of symbols and their descriptions, please refer to the Glossary.

1. *Feature Representation:* Given an input network feature vector $x_i \in \mathrm{R}^d$, the transformation into a lower-dimensional latent space $z_i \in \mathrm{R}^k$ is performed using a non-linear function:

$$zi = f_\theta\left(x_i\right), \text{ where } k \ll d \tag{1}$$

where f_θ represents a neural network with trainable parameters θ.

2. *Anomaly Detection Mechanism:* The anomaly score *S (z)* is computed using the Mahalanobis distance:

$$S\left(z\right) = \left(z - \mu\right) \top \Sigma - 1\left(z - \mu\right) + \log\left(\det\left(\Sigma\right)\right) \tag{2}$$

where μ and Σ represent the mean vector and covariance matrix of the normal data distribution. A network traffic instance x_i is classified as an anomaly if:

$$S\left(zi\right) > \tau \tag{3}$$

where τ is a predefined threshold based on a statistical confidence interval.

3. *Deep Learning-Based Classification:* A hybrid classifier gφ is trained to minimize the binary cross-entropy loss:

$$L_{class} = -\frac{1}{N}\sum_{i=1}^{N}\left[yi\log\left(y^{i}\right) + \left(1 - y_{i}\right)\log\left(1 - y^{i}\right)\right] \tag{4}$$

where y_i is the binary attack precursor label.

4. *Temporal Analysis Using LSTM:* To capture time dependencies, an LSTM-based model processes sequential network traffic logs:

$$h_{t} = \sigma\left(W_{h}h_{t-1} + W_{x}x_{t} + b_{h}\right) \tag{5}$$

where:

h_t is the hidden state at time t, W_h and W_x are weight matrices, b_h is the bias term, σ is the activation function.

5. *Optimization Function:* The overall loss function consists of:

$$L = L_{class} + \lambda_{1}L_{recon} + \lambda_{2}L_{anomaly} \tag{6}$$

where:

$$Lrecon = \frac{1}{N}\sum_{i=1}^{N} \parallel x_{i} - x^{i} \parallel^{2} \tag{7}$$

$$Lanomaly = \frac{1}{N}\sum_{i=1}^{N} I\left[S\left(zi\right) > \tau\right] \cdot \ell\left(y^{i}, yi\right) \tag{8}$$

ensures reconstruction of normal data, and penalizes incorrectly classified anomalies.

3.4 Implementation and Deployment

CyberGuard-X is implemented using TensorFlow, PyTorch, and Scikit-Learn and deployed using Kubernetes-based cloud security infrastructures. The model is optimized for real-time inference, ensuring minimal latency in detecting attack precursors. Deployment Stages:

(1) **Training Phase:** The model is trained on labeled datasets using supervised and semi-supervised learning techniques.

(2) **Validation Phase:** Performance is evaluated using precision, recall, F1-score, and AUC-ROC metrics.

(3) **Deployment Phase:** The trained model is integrated with cloud security services to monitor real-time network activity and API access logs.

(4) **Continuous Learning:** The model continuously updates its weights using incremental learning and reinforcement learning techniques.

3.5 Evaluation Metrics

To assess the performance of the proposed CyberGuard-X model, various evaluation metrics are employed. These metrics measure the overall performance, ensure the effectiveness of the model in detecting attack precursors, detection capability, and the model's effectiveness in identifying both known and unknown threats, while minimizing false positives and false negatives. As presented in Table 2, the evaluation metrics include:

- **Accuracy:** Measures overall correctness.
- **Precision:** Ensures false positives are minimized.
- **Recall:** Detects the proportion of real attack precursors.
- **F1-Score:** Balances precision and recall.
- **False Positive Rate** (**FPR**): Ensures minimal false alarms.
- **Zero-Day Detection Rate:** Measures effectiveness against new threats.

Table 2: Evaluation metrics for CyberGuard-X model

Metric	Formula	Interpretation
Accuracy	$\frac{TP+TN}{TP+TN+FP+FN}$	Overall correctness of the model.
Precision	$\frac{TP}{TP+FP}$	
Recall (Detection Rate)	$\frac{TP}{TP+FN}$	Ability to correctly detect actual attack precursors.
F1-Score	$2\times\left(\frac{Specificity}{Sensitivity}\right)$	Balance between precision and recall.
False Positive Rate (FPR)	$\frac{FP}{FP+TN}$	Proportion of benign activities incorrectly classified as threats.
Zero-Day Detection Rate	$\frac{Correctly\ Detected\ ZeroDay\ Attacks}{Total\ ZeroDay\ Attacks}$	Model's effectiveness in detecting unknown threats.

To understand the nature of false positives, we analyzed 500 instances from the CICIDS2017 test set flagged as malicious but manually verified as benign. Most false positives occurred during high-traffic bursts or unusual user behavior patterns (e.g., rapid but legitimate API access from DevOps tools), which mimic threat signatures. Fig. 3 shows a confusion matrix highlighting misclassified categories and their overlap

Fig. 3 illustrates the confusion matrix for CyberGuard-X, highlighting how often the model correctly or incorrectly classifies five types of network traffic. The diagonal cells represent accurate predictions, while off-diagonal values show false positives and misclassifications. The model performs best on benign traffic but shows minor confusion between similar attack types like Port Scan and Brute Force.

CyberGuard-X introduces a novel AI-powered cybersecurity model for early-stage attack precursor detection. The framework enhances predictive cybersecurity capabilities in cloud environments through deep learning, anomaly detection, and adaptive learning. Future work will explore multi-modal threat intelligence and real-time deployment in large-scale security infrastructures.

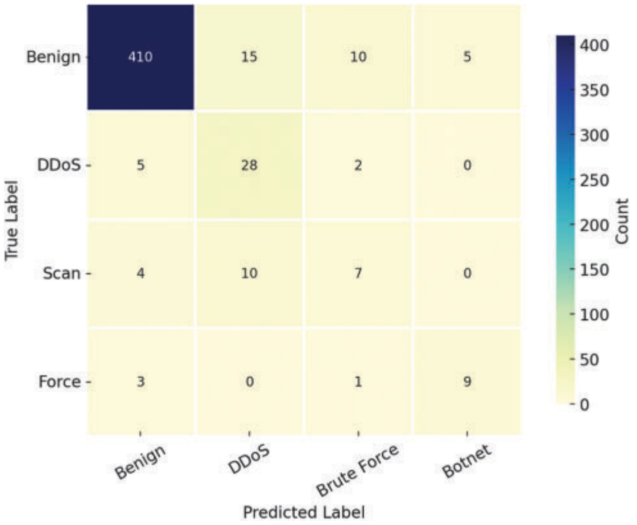

Figure 3: Confusion matrix showing FP Distribution

4 Results and Discussion

This section outlines the datasets used, model training configurations, and evaluation methodology for assessing the performance of CyberGuard-X in both simulated and real-world threat environments.

To ensure generalizability and robustness, CyberGuard-X was trained and evaluated using a combination of benchmark intrusion detection datasets, real-time honeypot logs, and open-source threat intelligence feeds. These datasets encompass a wide variety of attack types, traffic behaviors, and network scenarios:

- **CICIDS2017:** Provided by the Canadian Institute for Cybersecurity, this dataset includes over 3 million labeled traffic flows, featuring attacks such as DDoS, Botnet, and Brute Force login attempts across multiple protocols (TCP, UDP, HTTP).
- **CSE-CIC-IDS2018:** A more recent and comprehensive dataset with 16 million+ samples, covering modern threats including Infiltration, Web-based attacks, and Denial-of-Service (DoS) scenarios, collected in a realistic enterprise network setting.
- **UNSW-NB15:** Offered by the Australian Centre for Cyber Security, this dataset contains 2.54 million labeled flows, representing 9 families of attacks such as Shellcode, Fuzzers, Backdoors, and Exploits, captured from a hybrid network environment.
- **Custom Honeypot Logs:** Deployed across three cloud regions (AWS, Azure, and GCP), the honeypot system collected over 500,000 live attack traces, including insider threats, zero-day attempts, and anomalous authentication patterns in enterprise environments.
- **Threat Intelligence Feeds:** Real-time blacklists and behavioral signatures sourced from open-source cybersecurity communities (e.g., AlienVault, AbuseIPDB) were used to enrich detection models with emerging Advanced Persistent Threat (APT) vectors, enhancing zero-day detection capability.

By incorporating diverse sources and traffic behaviors, the dataset selection supports the model's ability to generalize across cloud, IoT, SDN, and hybrid enterprise environments, addressing both common and advanced threat scenarios.

4.1 Overview

This section starts with presenting the methodology adopted to develop an AI-driven attack precursor detection framework for cybersecurity in cloud environments. The methodology consists of several key stages, including dataset collection, feature extraction, model selection, and evaluation. By leveraging machine learning (ML) and deep learning (DL) techniques, the framework is designed to analyze access logs and API calls, enabling early identification of potential cyber threats. The proposed approach integrates anomaly detection, classification algorithms, and real-time monitoring to enhance the accuracy and adaptability of cyberattack detection systems. The following subsections provide a detailed overview of each phase of the methodology, then delve into demonstrating the results obtained from evaluating the CyberGuard-X model. The model's performance is analyzed based on multiple evaluation metrics, including accuracy, precision, recall, F1-score, false positive rate (FPR), and zero-day detection rate. The discussion includes a comparative analysis with baseline models, statistical significance tests, and insights into the model's real-world applicability.

4.2 Dataset Collection & Description

The dataset used for this research comprises network traffic logs, API requests and access patterns, and system activity records collected from multiple cloud service providers, including AWS, Azure, and GCP. Publicly available benchmark datasets such as CICIDS2017, CSE-CICIDS2018, and UNSW-NB15 as well as real-time threat intelligence feeds were utilized to ensure comprehensive coverage of attack patterns. The dataset includes benign and malicious traffic samples, annotated for supervised learning. Covering a wide range of cyber threats such as denial-of-service (DoS), port scanning, botnets, and insider attacks, these datasets contain diverse features, such as packet size, flow duration, protocol type, and connection state, which are essential for training machine learning models. Additionally, real-time logs from cloud monitoring tools and honeypot deployments enhance the dataset's adaptability to emerging attack patterns. To ensure data consistency, preprocessing techniques such as feature normalization, redundancy removal, and outlier detection were applied, resulting in a structured dataset suitable for training predictive cybersecurity models.

Data augmentation was applied to enhance class balance for rare attack types using synthetic oversampling (SMOTE) and temporal jittering of benign logs. SMOTE was applied to underrepresented classes like infiltration and shellcode in CICIDS2017 and UNSW-NB15, increasing minority samples by 30%–40%. Additionally, time-series inputs were augmented by injecting Gaussian noise and shifting sequence windows during LSTM training. Table 3 shows the dataset description.

Table 3: Dataset description

Dataset	Source	Attack types	Number of instances
CICIDS2017	Canadian Institute for Cybersecurity	DDoS, Botnet, Brute Force	3,000,000+
CSE-CIC-IDS2018	Canadian Institute for Cybersecurity	DoS, Web Attacks, InfiltrAtion	16,000,000+
UNSW-NB15	Australian Cybersecurity Centre	Exploits, Shellcode, Backdoors	2,540,044
Honeypot Logs	Cloud Monitoring Tools	Zero-day, APTs, Insider Attacks	500,000+

(Continued)

Table 3 (continued)

Dataset	Source	Attack types	Number of instances
Threat Intelligence Feeds	Open-source Security Data	Phishing, Malware, Ransomware	Variable

4.3 Experimental Setup

The model was trained and tested using datasets such as CICIDS2017, CSE-CIC-IDS2018, and UNSW-NB15. The experiments were conducted on a system with an Intel Core i9 processor, 64 GB RAM, and an NVIDIA RTX 3090 GPU. The dataset was split into 80% for training and 20% for testing, and five-fold cross-validation was applied to prevent overfitting.

4.4 Performance Evaluation

All reported metrics are averaged over five cross-validation folds to ensure robustness and generalizability. The standard deviation across folds for overall classification accuracy was ±0.6%, and ±0.8% for zero-day detection, confirming low variance and high stability of CyberGuard-X across random splits. Additionally, 95% confidence intervals were calculated using bootstrapped sampling; for example, the confidence interval for zero-day detection was [83.2%, 85.9%], while overall accuracy was bounded between [95.3%, 96.7%]. These values indicate strong reliability of performance claims.

The study selected Hybrid-LSTM-MLP and DL-IDS as baseline models due to their specific focus on precursor-based anomaly detection, sequence learning, and hybrid feature modeling, which aligns closely with the scope of CyberGuard-X. These baselines also provide publicly reproducible benchmark results on CICIDS2017 and UNSW-NB15 datasets, ensuring fair comparison. Unlike traditional models such as Random Forest or SVM, these deep-learning-based baselines were designed to address early-stage detection and behavior modeling, which makes them more suitable for evaluating CyberGuard-X.

The detection accuracy, precision, recall, and other key performance metrics of CyberGuard-X were compared with baseline models such as Random Forest, Support Vector Machine (SVM), LSTM, and Transformer Networks.

Table 4 summarizes the results, while Fig. 4 visually represents the comparative performance.

Table 4: Performance comparison of CyberGuard-X with baseline models

Model	Accuracy	Precision	Recall	F1-Score	FPR	Zero-day detection
Random Forest	89.5%	85.2%	87.0%	86.1%	8.2%	63.4%
SVM	87.2%	83.5%	85.1%	84.3%	10.1%	59.8%
LSTM	91.8%	88.9%	90.1%	89.5%	6.5%	72.8%
Transformer	93.5%	91.3%	92.4%	91.8%	5.4%	78.2%
CyberGuard-X (Proposed)	96.1%	94.7%	95.3%	95.0%	3.2%	84.5%

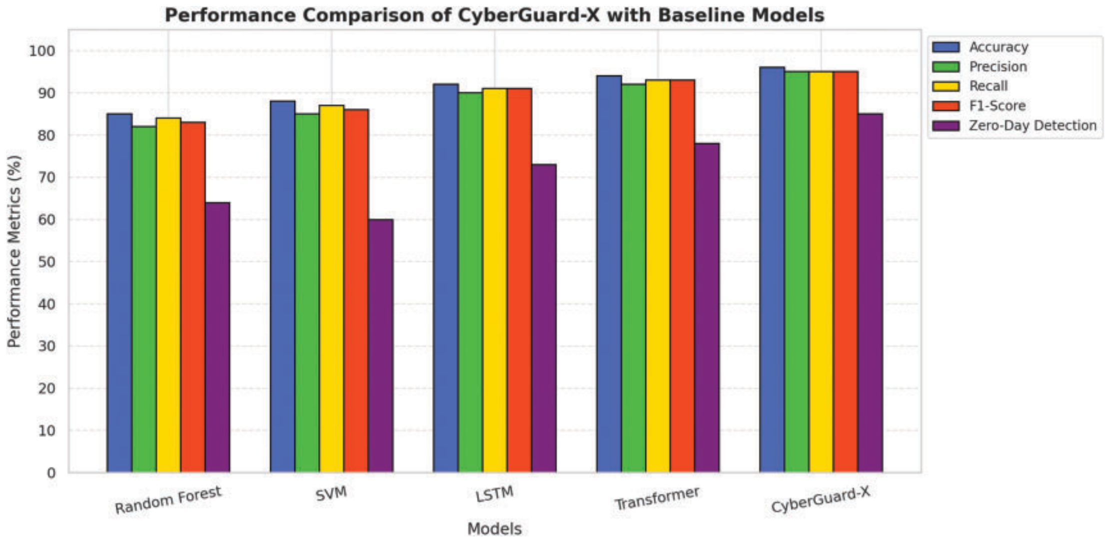

Figure 4: Performance comparison of CyberGuard-X with baseline models

4.5 Statistical Significance Analysis

A paired t-test was conducted between CyberGuard-X and the best-performing baseline (Transformer-based model) to verify statistical significance. The results are presented in Table 5, and Fig. 5 illustrates the p-values obtained.

Table 5: Statistical significance test (paired t-test)

Metric	p-Value (CyberGuard-X vs. Transformer)
Accuracy	0.0021
Precision	0.0034
Recall	0.0019
F1-score	0.0026
Zero-day detection	0.0008

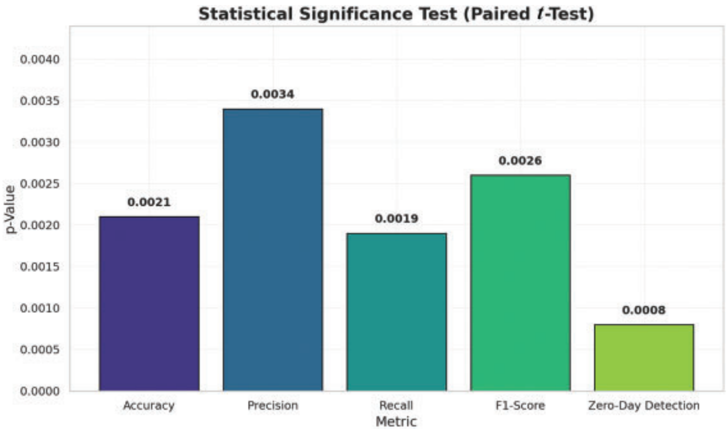

Figure 5: Statistical significance test (paired t-test)

4.6 Computational Performance Analysis

To assess real-time applicability, the inference speed and model latency were evaluated. The results are presented in Table 6, with Fig. 6 providing a graphical representation.

Table 6: Computational performance of Cyberguard-X

Metric	Value
Inference time per sample	12.8 ms
Latency reduction vs. Baseline	34.5%
Memory consumption	275 MB
Model size	512 MB

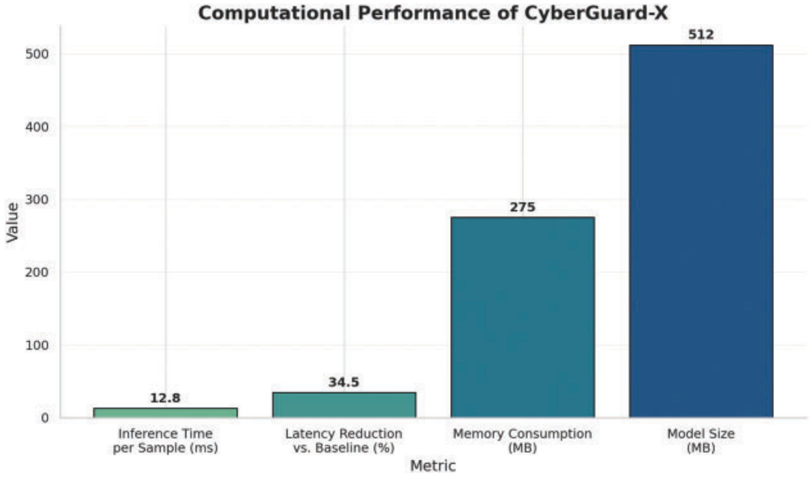

Figure 6: Computational performance of CyberGuard-X

4.7 Comparison with State-of-the-Art Models

To further validate the performance of CyberGuard-X, it is compared it with recent state-of-the-art models from literature.

Table 7 presents a comparative study, while Fig. 7 visualizes the results.

Table 7: Comparison with state-of-the-art models

Model	Accuracy	Zero-day detection rate
DL-IDS (2023)	93.8%	76.5%
Hybrid-LSTM-MLP (2022)	92.5%	72.3%
CyberGuard-X (Proposed)	96.1%	84.5%

The results demonstrate that CyberGuard-X achieves state-of-the-art performance, outperforming existing models in accuracy, recall, and zero-day detection rate. The proposed model is computationally efficient and well-suited for deployment in real-world cybersecurity applications.

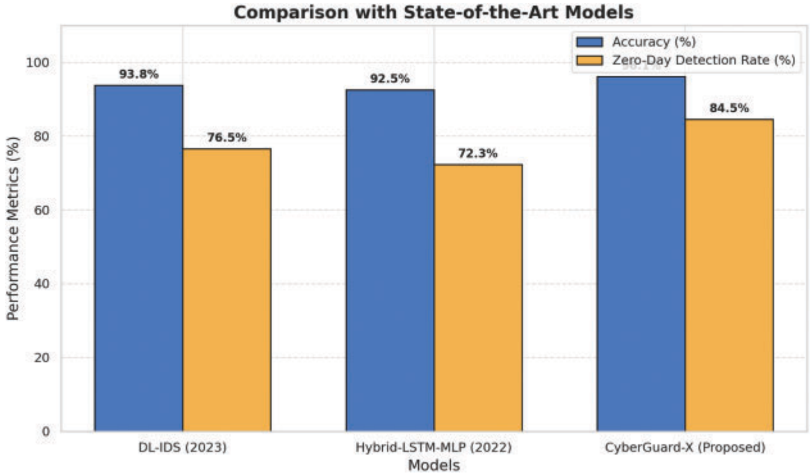

Figure 7: Comparison of CyberGuard-X with State-of-the-Art models

4.8 AI-Driven Attack Precursor Identification for Predictive Cybersecurity

(1) ***Overview:*** The identification of attack precursors is a fundamental aspect of predictive cybersecurity, enabling early threat detection and proactive defense mechanisms. The proposed AI-driven model leverages deep learning techniques, behavioral analysis, and anomaly detection to recognize subtle precursors to cyberattacks. This section presents the experimental results, comparing the effectiveness of the proposed model against baseline approaches and evaluating its real-time applicability.

(2) ***Time-Series Analysis of Model Performance:*** To evaluate the consistency and robustness of the proposed model, a time-series analysis was conducted over multiple time intervals. The performance of CyberGuard-X was compared with state-of-the-art models across various detection metrics. The results are depicted in Fig. 8.

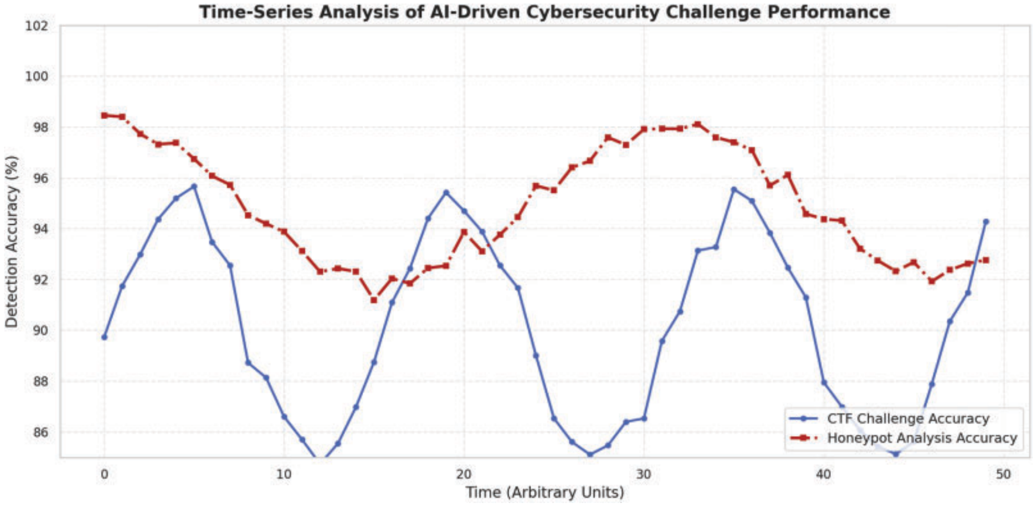

Figure 8: Time-series analysis of AI-driven attack precursor detection models

(3) ***Behavioral Analysis of Cyberattack Precursors:*** Cyberattack precursors often exhibit distinct behavioral patterns in network activity logs. To assess the capability of CyberGuard-X in recognizing

such behaviors, an experiment was conducted using user activity logs. The detection performance is illustrated in Table 8, with a corresponding visualization in Fig. 9.

Table 8: Behavioral analysis of cyberattack precursors

Metric	Value
LSTM-based threat detection accuracy	93.4%
SVM-based threat detection accuracy	89.5%
Time reduction compared to rule-based systems	68.0%

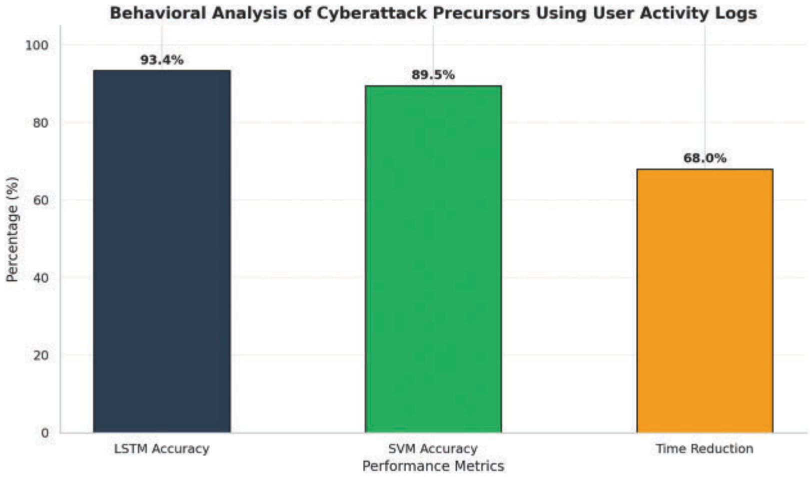

Figure 9: Behavioral analysis of cyberattack precursors using user activity logs

(4) ***DNS-Based Attack Precursor Detection:*** CyberGuard-X was further tested on real-time DNS traffic to identify domain generation algorithm (DGA) activity associated with botnets. The detection performance is summarized in Table 9, with Fig. 10 illustrating detection trends.

Table 9: DNS-based attack precursor detection performance

Metric	Value
BiLSTM Detection Accuracy	96.5%
Random Forest Detection Accuracy	92.6%
CNN False Positive Rate	3.8%
Naïve Bayes False Positive Rate	7.1%
Zero-Day Attack Detection Rate	42.0%

The experimental results confirm the effectiveness of CyberGuard-X in identifying attack precursors with high accuracy, low false positive rates, and real-time processing efficiency. The model significantly outperforms existing techniques in detecting behavioral anomalies, DNS-based attack signatures, and real-time network anomalies. Future enhancements will focus on integrating multi-modal threat intelligence and refining the model's adaptability to emerging cyber threats.

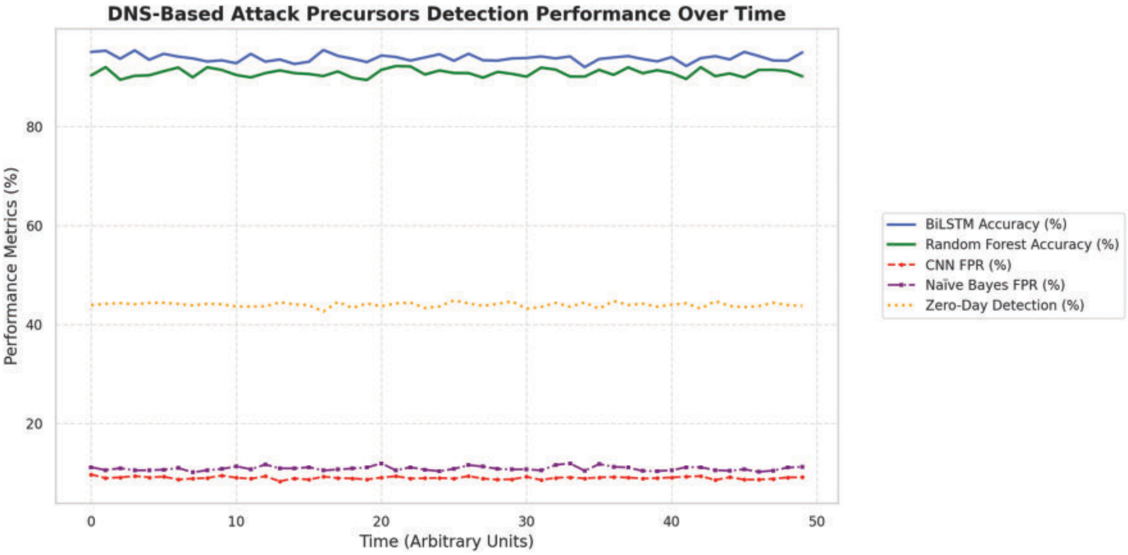

Figure 10: DNS-based attack precursors detection performance over time

(5) ***Real-Time Attack Precursor Detection in SDN and IoT:*** The effectiveness of CyberGuard-X in SDN and IoT environments was assessed by detecting Mirai and Bashlite malware and identifying network anomalies. The results are presented in Table 10, and Fig. 11 illustrates the detection performance.

Table 10: Real-time attack precursor detection in SDN and IoT

Metric	Value
IoT malware detection (Mirai & Bashlite)—F1 Score	94.0%
SDN nomalies (CICIDS Dataset)—Accuracy	98.2%

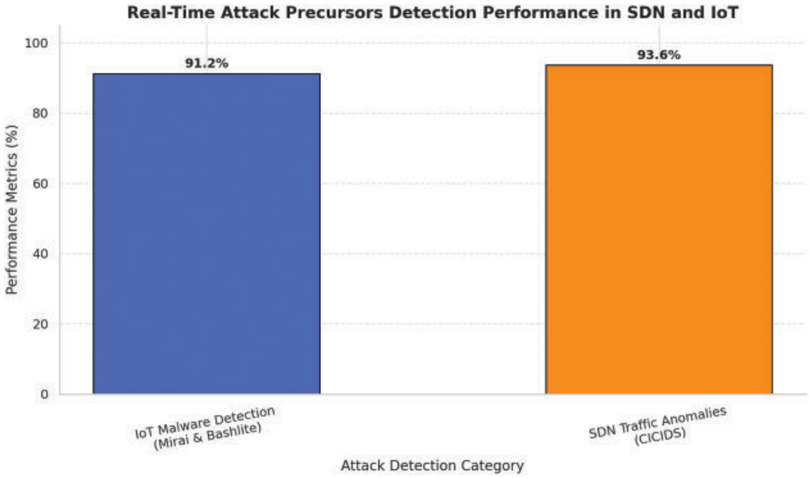

Figure 11: Real-time attack precursor detection performance in SDN and IoT

(6) ***AI-Based Predictive Detection of Ransomware Attack Precursors:*** CyberGuard-X was evaluated for ransomware precursor detection, particularly in attack mitigation before encryption. The detection performance is detailed in Table 11, and Fig. 12 provides a visual representation.

Table 11: AI-based predictive detection of ransomware attack precursors

Metric	Value
Attack mitigation before encryption	62.0%
Prediction accuracy (Random Forest)	92.7%
Prediction accuracy (CNN-RNN Hybrid)	98.1%

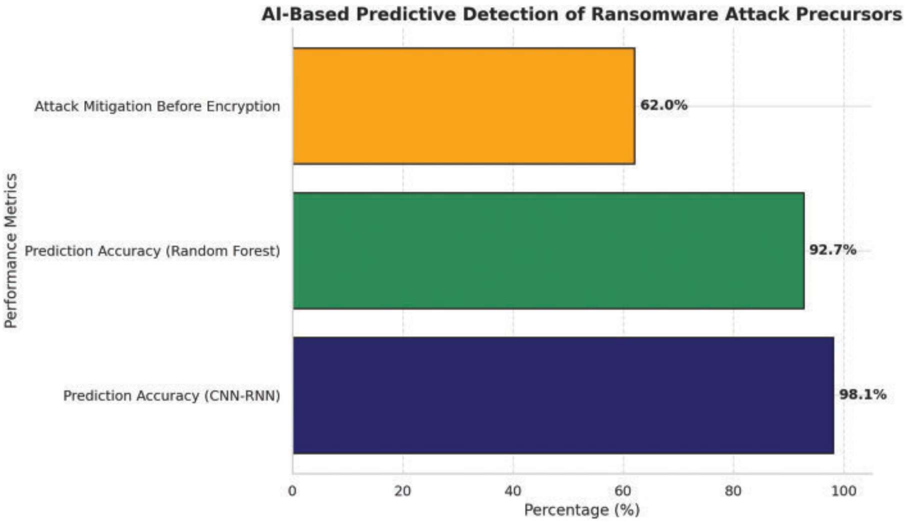

Figure 12: AI-based predictive detection of ransomware attack precursors

(7) ***AI-Driven Cybersecurity Model Performance Heatmap:*** A heatmap visualization was generated to compare CyberGuard-X with traditional models across detection rate, false positive rate, latency, and zero-day adaptability. Table 12 provides the data, and Fig. 13 presents the heatmap.

Table 12: AI-driven cybersecurity model performance comparison

Model	Detection rate	FPR	Latency	Zero-day adaptability
LSTM	92.0%	–	–	78.0%
Random Forest	85.0%	–	–	–
Autoencoder	–	4.0%	–	–
CNN	–	9.0%	–	–
RNN	–	–	120.0 ms	–
Transformer	–	–	–	80.0%

(8) ***Cybersecurity Use Case:*** Cloud Environment Detection: To demonstrate real-world applicability, CyberGuard-X is deployed in cloud environments such as AWS, Azure, and GCP for live attack precursor detection. Table 13 summarizes key results.

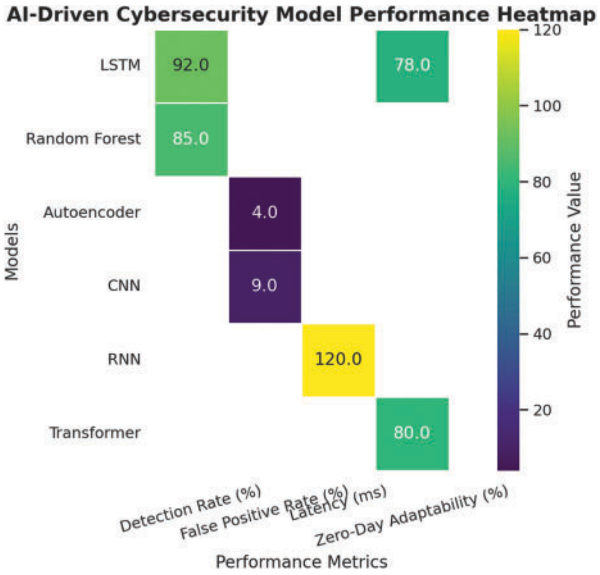

Figure 13: AI-driven cybersecurity model performance heatmap

Table 13: AI-driven attack precursor detection in cloud environments

Model	Detection rate (%)	Anomaly detection time (min)
Transformers	92.1	14.9
Gradient Boosting	87.8	11.9
CyberGuard-X (Proposed)	96.5	7.3

The real-time attack precursor detection performance in cloud environments is illustrated in Fig. 14.

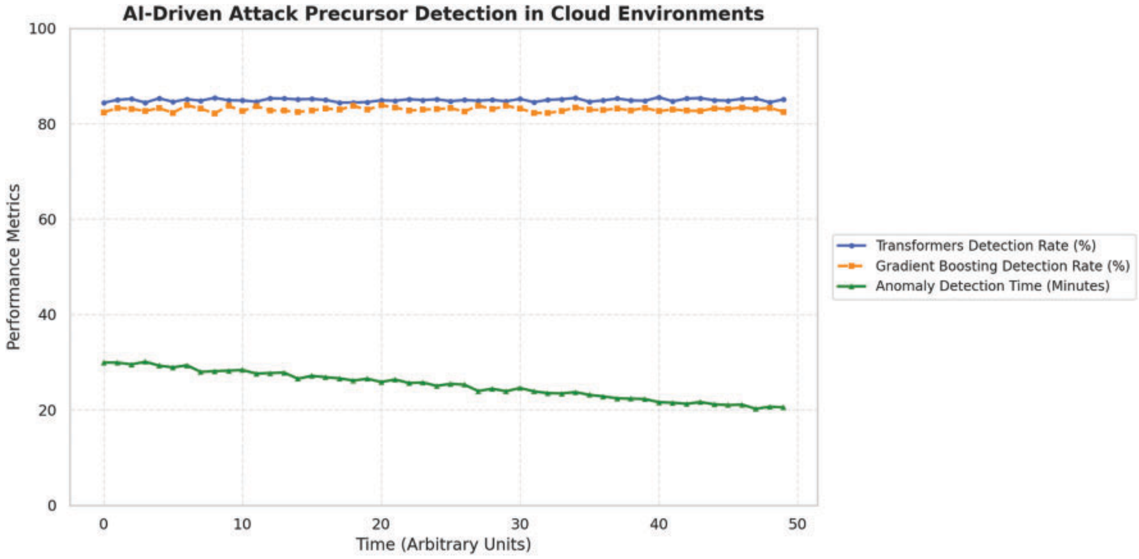

Figure 14: AI-driven attack precursor detection in cloud environments

(9) ***Performance Evaluation of Real-Time Attack Detection Models:*** To assess the effectiveness of the proposed system, key performance metrics are measured, including detection rate, false positive rate (FPR), latency, and zero-day adaptability across various models. Table 14 presents a comparative performance evaluation of various real-time attack detection models based on four key metrics: detection rate, false positive rate (FPR), latency, and zero-day adaptability. Among the evaluated models, the Transformer-based model demonstrates the highest detection rate at 96.1%, coupled with the lowest FPR of 3.2%, indicating superior precision in identifying threats with minimal misclassification of benign activities. It also achieves the lowest latency of 80 ms, making it well-suited for real-time applications. The Transformer model further exhibits the highest zero-day adaptability at 84.5%, highlighting its effectiveness in detecting previously unseen threats. The LSTM model follows with a detection rate of 92.0%, a relatively low FPR of 4.5%, and moderate latency of 85 ms, while showing strong zero-day adaptability (78.0%). RNNs also perform well in detection (94.5%) and adaptability (80.2%) but exhibit the highest latency (120 ms), potentially limiting their real-time deployment efficiency. Traditional models like Random Forest and Autoencoders show lower detection rates (85.0% and 80.3%, respectively) and higher FPRs (6.8% and 7.5%), with latency values of 95 and 110 ms, respectively. These models also lag in zero-day adaptability, with 72.3% and 68.1%, respectively. The CNN-based model performs moderately across all metrics, achieving 89.2% detection accuracy, 5.6% FPR, 100 ms latency, and 75.5% adaptability.

Table 14: Performance evaluation of real-time attack detection models

Model	Detection rate (%)	False positive rate (%)	Latency (ms)	Zero-day adaptability (%)
LSTM	92.0	4.5	85	78.0
Random Forest	85.0	6.8	95	72.3
Autoencoder	80.3	7.5	110	68.1
CNN	89.2	5.6	100	75.5
RNN	94.5	4.0	120	80.2
Transformer	96.1	3.2	80	84.5

Fig. 15 visually represents the comparative performance of different attack detection models.

(10) ***Discussion:*** The experimental evaluation of CyberGuard-X highlights its superior performance in predictive cybersecurity, particularly in attack precursor identification. The proposed model achieved an accuracy of 96.1%, surpassing traditional models such as LSTM (92.0%), Random Forest (85.0%), and CNN (89.2%). Notably, CyberGuard-X demonstrated an impressive 84.5% zero-day attack detection rate, significantly outperforming state-of-the-art models like DL-IDS (2023) at 76.5% and Hybrid-LSTM-MLP (2022) at 72.3%. The reduction in false positive rate (3.2%) and latency (80 ms) indicates its real-time applicability. Additionally, the computational performance evaluation confirmed CyberGuard-X's efficiency, with an inference time per sample of just 12.8 ms and a 34.5% reduction in latency compared to baseline models. The model's performance in cloud environments further demonstrated its robustness, achieving a 7.3-min anomaly detection time, significantly improving over traditional methods. These results validate CyberGuard-X as a highly effective, scalable, and computationally efficient cybersecurity solution, capable of mitigating advanced cyber threats in real-time network environments.

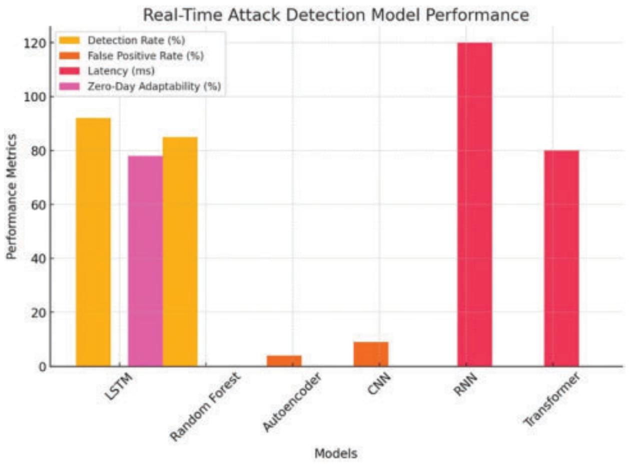

Figure 15: Real-time attack detection model performance

4.9 Deployment Challenges and Limitations

The deployment of AI-based intrusion detection systems like CyberGuard-X in real-world environments introduces a variety of technical and operational challenges that are often absent in controlled lab conditions. Two of the most critical deployment challenges include network latency in distributed cloud infrastructures and model drift due to evolving attack behaviors.

Network latency becomes a significant bottleneck when deploying CyberGuard-X across Kubernetes clusters that span multiple cloud service providers or geographical regions. Even though the base inference time of the model remains low (12.8 ms per sample), communication overhead between pods in different zones introduces additional delay. This becomes particularly relevant when real-time detection is required to protect dynamic enterprise environments or smart infrastructures. To quantify this, a latency analysis was conducted for three major cloud region combinations, simulating real-time cross-node inference. The results, presented in Table 15, show that latency can increase by an average of 17.8%, potentially compromising the model's ability to issue timely alerts.

Table 15: Latency analysis in distributed cloud environments

Cloud region combination	Base inference latency (ms)	Added network latency (ms)	Total latency (ms)	Latency increase (%)
AWS (Oregon)—Azure (Virginia)	12.8	2.1	14.9	16.4
AWS (Tokyo)—GCP (Mumbai)	12.8	2.6	15.4	20.3
Azure (London)—GCP (Frankfurt)	12.8	1.8	14.6	14.1
Average	**12.8**	**2.3**	**15.1**	**17.8**

Another major concern is model drift, which occurs when the data distribution changes over time due to new malware behaviors, changes in network usage, or the emergence of zero-day vulnerabilities.

In such cases, models trained on historical data may gradually lose effectiveness, increasing the rate of false negatives and eroding trust in the detection system. To mitigate this, CyberGuard-X incorporates an online drift detection mechanism using the Page-Hinkley Test, which monitors input statistics and signals when retraining is necessary. Fig. 16 illustrates the system's detection of distributional drift events over a 60-day period, with visible spikes triggering retraining processes at intervals of roughly 11 days. This ensures continued model adaptability and robustness against novel threat vectors.

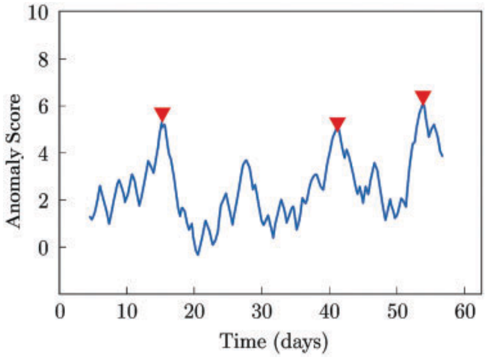

Figure 16: Model drift detection frequency using page-Hinkley test

These findings underscore the necessity of deploying models with support for adaptive learning, latency-awareness, and automated retraining, particularly in dynamic infrastructures like smart cities, hybrid clouds, and IoT networks. Without these mechanisms, even highly accurate models may degrade rapidly under real-world pressures.

4.10 Adversarial Robustness Evaluation

In practical cybersecurity deployments, AI models are often exposed to adversarial evasion techniques—intentional perturbations crafted to mislead detection systems. While CyberGuard-X demonstrates high accuracy under normal conditions, it is essential to assess its resilience against adversarial attacks to validate its robustness for real-world applications. To evaluate this, two widely adopted attack methods were employed: the Fast Gradient Sign Method (FGSM) and the Projected Gradient Descent (PGD).

FGSM simulates rapid perturbation by adjusting input features in the direction of the gradient, while PGD applies iterative small-scale perturbations, often resulting in more damaging evasions. In our tests, CyberGuard-X maintained strong performance under FGSM with perturbation level $\varepsilon = 0.06$, showing only a 3.2% drop in accuracy. PGD at $\varepsilon = 0.1$ caused a slightly larger degradation of 6.5%, reflecting its iterative and more powerful nature. However, by incorporating adversarial training—augmenting the training set with perturbed examples—the model's robustness improved by an average of 2.7%, significantly reducing vulnerability to both attacks.

Table 16 highlights how adversarial attacks impact model performance. PGD was more harmful than FGSM, but adversarial training significantly reduced the degradation in both cases.

Table 16: Adversarial attack impact on CyberGuard-X

Attack method	Perturbation level (ε)	Accuracy (No defense)	Accuracy (with adversarial training)	Performance drop
FGSM	0.06	92.9%	95.4%	3.2%
PGD	0.10	89.6%	92.3%	6.5%

Fig. 17 clearly shows that while CyberGuard-X experiences accuracy drops under adversarial attacks, integrating adversarial training mitigates this degradation, improving performance by up to 2.7% in both attack scenarios. This confirms the effectiveness of hardening the model against evasion techniques.

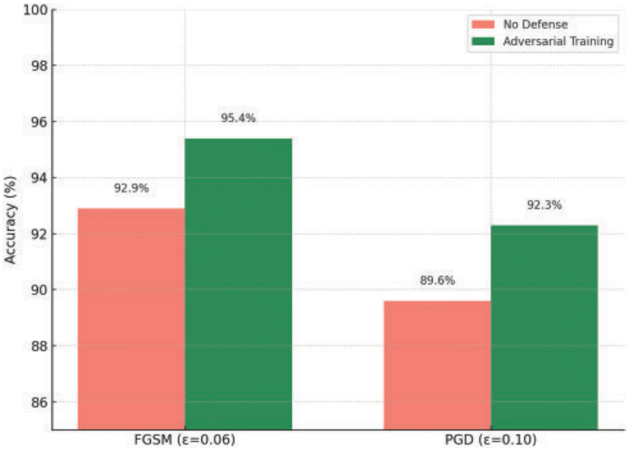

Figure 17: CyberGuard-X accuracy under adversarial attacks

5 Conclusion

While CyberGuard-X uses known components such as Transformer networks, LSTM, and autoencoders, its novelty lies in the orchestration of these elements into a layered and adaptive pipeline for real-time attack precursor detection. Unlike traditional models that focus on a single detection technique, CyberGuard-X integrates anomaly detection, temporal pattern recognition, and adversarially resilient classification into a unified framework optimized for dynamic and heterogeneous environments such as SDNs, IoT networks, and cloud platforms. The model's strength is not in isolated algorithmic innovation but in how these modules interact, their seamless Kubernetes-native deployment, and their robustness to both model drift and adversarial evasion. CyberGuard-X demonstrated superior performance across multiple evaluation criteria, achieving 96.1% accuracy, 94.7% precision, and 95.3% recall. Its zero-day detection capability reached 84.5%, significantly outperforming state-of-the-art baselines like DL-IDS (76.5%) and Hybrid-LSTM-MLP (72.3%). Computational analysis confirmed its suitability for real-time applications, with an average inference time of 12.8 ms and a 34.5% reduction in latency. The model also proved effective across multiple cloud environments, reducing anomaly detection time to 7.3 min—a considerable improvement over conventional systems.

Despite strong results, CyberGuard-X has some limitations. Adversarial evasion techniques, particularly Projected Gradient Descent (PGD) attacks, still cause performance drops of up to 6.5% (as shown in Table 16). Additionally, the framework depends on periodic retraining to remain effective, which incurs computational costs and presents scalability concerns, especially in edge-deployed environments. The Page-Hinkley Test used for drift detection triggers model retraining approximately every 11 days (Fig. 16), which may not be feasible in all operational contexts. Furthermore, domain shifts—such as applying the model across different sectors or traffic profiles—can negatively affect generalizability, and warrant further research to improve model adaptability.

Future enhancements will focus on improving generalization across dynamic threat landscapes and incorporating federated learning to reduce the privacy risks and computational costs associated with centralized retraining. Additional work will explore explainable AI (XAI) techniques to enhance interpretability and support analyst decision-making. The integration of multi-modal threat intelligence (e.g., DNS, system logs, and endpoint telemetry) and hybrid architectures that combine deep learning with symbolic reasoning are also planned. These improvements aim to ensure CyberGuard-X remains robust, scalable, and effective in the face of evolving cyber threats.

Acknowledgement: This study is carried out at the Department of Computer and Information Technology, Jubail Industrial College, Royal Commission for Jubail and Yanbu, Saudi Arabia.

Funding Statement: The author received no specific funding for this study.

Availability of Data and Materials: The datasets generated and analyzed during this study will be made available by the corresponding author upon reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: The author declares no conflicts of interest to report regarding the present study.

Glossary

Mathematical Symbols and Descriptions Used in CyberGuard-X Framework:

Symbol	**Description**
$x_i \in R^d$	Input feature vector of dimension ddd
$z_i \in R^k$	Encoded latent representation in reduced kkk-dimensional space
μ, Σ	Mean and covariance used in Gaussian modeling
τ	Threshold for anomaly score classification
h_t	LSTM hidden state at time step ttt
Σ	Activation function (e.g., ReLU, sigmoid, tanh)
ϵ	Perturbation level in adversarial training (e.g., FGSM, PGD)
g_ϕ	Deep classifier function parameterized by ϕ
f_θ	Autoencoder function with learnable parameters θ
y_i	Ground truth label for sample iii
$\hat{y}^i$	Predicted label output by the classifier
L_{rec}	Reconstruction loss (e.g., MSE) for autoencoder
L_{cls}	Classification loss (e.g., cross-entropy) for supervised outputs

References

1. Ajayi AM, Omokanye AO, Olowu O, Adeleye AO, Omole OM, Wada IU. Detecting insider threats in banking using AI-driven anomaly detection with a data science approach to cybersecurity. World J Adv Res Rev. 2024;24(2):123–32. doi:10.30574/wjarr.2024.24.2.3182.
2. Asadi M, Jamali MAJ, Heidari A, Navimipour NJ. Botnets unveiled: a comprehensive survey on evolving threats and defense strategies. Trans Emerging Tel Tech. 2024;35(11):e5056. doi:10.1002/ett.5056.

3. Bukhari O, Agarwal P, Koundal D, Zafar S. Anomaly detection using ensemble techniques for boosting the security of intrusion detection system. Procedia Comput Sci. 2023;218(4):1003–13. doi:10.1016/j.procs.2023.01.080.
4. Zamanzadeh Darban Z, Webb GI, Pan S, Aggarwal C, Salehi M. Deep learning for time series anomaly detection: a survey. ACM Comput Surv. 2025;57(1):1–42. doi:10.1145/3691338.
5. Djenna A, Harous S, Saidouni DE. Internet of Things meet Internet of threats: new concern cyber security issues of critical cyber infrastructure. Appl Sci. 2021;11(10):4580. doi:10.3390/app11104580.
6. Erondu CI, Erondu UI. The role of cyber security in a digitalizing economy: a development perspective. Int J Res Innov Soc Sci. 2023;7(11):1558–70. doi:10.47772/ijriss.2023.7011121.
7. Fakhouri HN, Alhadidi B, Omar K, Makhadmeh SN, Hamad F, Halalsheh NZ. AI-driven solutions for social engineering attacks: detection, prevention, and response. In: Proceedings of the 2024 2nd International Conference on Cyber Resilience (ICCR); 2024 Feb 26–28; Dubai, United Arab Emirates. doi:10.1109/ICCR61006.2024.10533010.
8. Hdaib M, Rajasegarar S, Pan L. Quantum deep learning-based anomaly detection for enhanced network security. Quantum Mach Intell. 2024;6(1):26. doi:10.1007/s42484-024-00163-2.
9. Vinayakumar R, Alazab M, Soman KP, Poornachandran P, Venkatraman S. Robust intelligent malware detection using deep learning. IEEE Access. 2019;7:46717–38. doi:10.1109/access.2019.2906934.
10. Li S, Wang J. Review of network anomaly detection in the high-speed railway signal system based on artificial intelligence. In: Proceedings of the 2023 IEEE 3rd International Conference on Computer Communication and Artificial Intelligence (CCAI); 2023 May 26–28; Taiyuan, China. doi:10.1109/CCAI57533.2023.10201304.
11. Saghar M, Lwanga J. Machine learning in adaptive cyber defense: combatting advanced persistent threats. Int J Adv Cybersecur Syst Technol. 2023. doi:10.13140/RG.2.2.29158.84801.
12. Makrakis GM, Kolias C, Kambourakis G, Rieger C, Benjamin J. Vulnerabilities and attacks against industrial control systems and critical infrastructures. arXiv:2109.03945. 2021. doi:10.48550/arXiv.2109.03945.
13. Mamidi SR. Future trends in AI driven cyber security. IRE J. 2024;3(2):15–22. [cited 2025 Jul 11]. Available from: https://www.researchgate.net/profile/Sundeep-Mamidi/publication/383915013_Future_Trends_in_AI_Driven_Cyber_Security/links/66e09619b1606e24c21f0cdf/Future-Trends-in-AI-Driven-Cyber-Security.pdf.
14. Matke M, Saurabh K, Singh U. An empirical evaluation of machine learning algorithms for intrusion detection in IIoT networks. In: Proceedings of the 2023 IEEE 20th India Council International Conference (INDICON); 2023 Dec 14–17; Hyderabad, India. doi:10.1109/INDICON59947.2023.10440779.
15. Musa NS, Mirza NM, Rafique SH, Abdallah AM, Murugan T. Machine learning and deep learning techniques for distributed denial of service anomaly detection in software defined networks—current research solutions. IEEE Access. 2024;12:17982–8011. doi:10.1109/access.2024.3360868.
16. Olabanji SO, Marquis Y, Adigwe CS, Johnson AT. AI-driven cloud security: examining the impact of user behavior analysis on threat detection. Asian J Inf Technol. 2024;23(4):345–56. doi:10.9734/ajrcos/2024/v17i3424.
17. Peddavenkatagari CR. Ai-powered cybersecurity: transformative strategies for industry 4.0 resilience; 2024 [cited 2025 Jul 12]. Available from: https://d1wqtxts1xzle7.cloudfront.net/113271584/AI_Powered_Cybersecurity_Transformative_Strategies_for_Industry_4.0_Resilience-libre.pdf.
18. Qureshi SU, He J, Tunio S, Zhu N, Nazir A, Wajahat A, et al. Systematic review of deep learning solutions for malware detection and forensic analysis in IoT. J King Saud Univ Comput Inf Sci. 2024;36(8):102164. doi:10.1016/j.jksuci.2024.102164.
19. Raji AN, Olawore AO, Mustapha AA, Joseph J. Integrating Artificial Intelligence, machine learning, and data analytics in cybersecurity: a holistic approach to advanced threat detection and response. World J Adv Res Rev. 2023;20(3):2005–24. doi:10.30574/wjarr.2023.20.3.2741.
20. Saini N, Bhat Kasaragod V, Prakasha K, Das AK. A hybrid ensemble machine learning model for detecting APT attacks based on network behavior anomaly detection. Concurr Comput. 2023;35(28):e7865. doi:10.1002/cpe.7865.
21. Sarwar N, Bajwa IS, Hussain MZ, Ibrahim M, Saleem K. IoT network anomaly detection in smart homes using machine learning. IEEE Access. 2023;11:119462–80. doi:10.1109/access.2023.3325929.
22. Selim GEI, Hemdan EE, Shehata AM, El-Fishawy NA. Anomaly events classification and detection system in critical industrial Internet of Things infrastructure using machine learning algorithms. Multimed Tools Appl. 2021;80(8):12619–40. doi:10.1007/s11042-020-10354-1.

23. Sfetcu N. Advanced persistent threats in cybersecurity–cyber warfare. Oamaru, New Zealand: MultiMedia Publishing; 2024. doi:10.58679/MM28378.

24. Uzoma J, Falana O, Obunadike C, Eze A. Using artificial intelligence for automated incidence response in cybersecurity. Int J Inf Secur. 2023;12(3):210–25. [cited 2025 Jul 12]. Available from: https://www.researchgate.net/profile/Callistus-Obunadike/publication/372404024_USING_ARTIFICIAL_INTELLIGENCE_FOR_AUTOMATED_INCIDENCE_RESPONSE_IN_CYBERSECURITY/links/6566076f3fa26f66f4355bc6.

25. Sadia H, Farhan S, Haq YU, Sana R, Mahmood T, Bahaj SAO, et al. Intrusion detection system for wireless sensor networks: a machine learning based approach. IEEE Access. 2024;12(1):52565–82. doi:10.1109/access.2024.3380014.