



INTERNATIONAL JOURNAL OF TRENDS IN EMERGING RESEARCH AND DEVELOPMENT

INTERNATIONAL JOURNAL OF TRENDS IN EMERGING RESEARCH AND DEVELOPMENT

Volume 3; Issue 4; 2025; Page No. 477-482

Received: 19-05-2025
Accepted: 28-06-2025
Published: 22-07-2025

Ai-Powered Intrusion Detection Systems to Improve Cloud-Based Infrastructure Cybersecurity

¹Sreedhar Korubilli and ²Dr. Shashank Naidu

¹Research Scholar, YBN University, Ranchi, Jharkhand, India

²Associate Professor, YBN University, Ranchi, Jharkhand, India

DOI: <https://doi.org/10.5281/zenodo.21918371>

Corresponding Author: Sreedhar Korubilli

Abstract

An enormous surge in cybercrimes has made cybersecurity a top concern on a global scale, impacting vital companies and infrastructures all over the globe. Attacks against cloud platforms, IoT systems, and industrial control systems have increased dramatically, with a focus on India. The increasing complexity of security breaches is highlighted by a thorough examination of cybercrime data by area, industry, and time period, which highlights the critical need for cybersecurity frameworks driven by artificial intelligence. These worrying tendencies are brought to light in this work, which further proves that sophisticated threat detection systems are required. Our extensive user survey included a wide range of sectors, from information technology and banking to healthcare and critical infrastructure, and its purpose was to shed light on the practical consequences of cybersecurity risks. In order to assess the present security picture, identify existing gaps, and identify adoption barriers of AI-driven solutions, the poll solicited responses from cybersecurity experts and IT stakeholders.

Keywords: Cybersecurity, Cloud Computing, Intrusion Detection System (IDS), Artificial Intelligence (AI), Deep Learning

Introduction

Organizations' data and application management has been revolutionized by cloud computing. According to Mell and Grance of the National Institute of Standards and Technology (NIST), "cloud computing" is a paradigm that allows users to easily access, through the internet, pools of configurable computing resources such as networks, servers, storage resources, applications, and services with little requirement for the physical characteristics of the IT infrastructure (2011). As a result, many businesses have grown more rapidly and shown that they can adapt to market conditions by cutting costs in a number of operational areas. There has been meteoric rise in the use of cloud computing in recent years. Over the next four years, the global market for public cloud services is projected to increase by 23%, according to Gartner's tech research. as compared to \$270 billion in 2020, 1% in the scenario developed for 2021 results in a total of \$332.3 billion.

Based on these results, it's evident that cloud computing is

transforming the face of enterprises as cloud services converge across diverse industries. Oligopoly persists in the cloud computing industry, with a small number of dominant players providing innovative and widely used services. Amazon Web Services (AWS), Microsoft Azure, Google Cloud, and Oracle Cloud are the four main cloud service providers. From the early days of cloud computing and its underlying principles to the widespread use of AI and ML, these companies have been instrumental in propelling cloud solutions and services forward.

The typical price tag for a data breach in hybrid cloud settings was \$4.45 million in 2023, as reported by IBM. Attacks on supply chains, vulnerabilities in cloud-native apps, and advanced persistent threats (APTs) are all examples of new dangers. Additional vulnerabilities may be exploited by hackers via linked systems due to the proliferation of edge computing, IoT, and multicloud methods. Striking a balance between robust security measures and a smooth user experience is still a big

problem. To be effective, systems must identify threats automatically, implement intelligent access control, and use adaptive encryption. Furthermore, cloud infrastructures are required to adhere to strict data protection standards by rules including GDPR, HIPAA, CCPA, and FISMA. In order to be agile and in compliance with these rules, organizations need to do frequent audits and use real-time monitoring.




Fig 1: Workloads running in cloud

Core Threats in Cloud Environments

- **Data Breaches and Loss:** Insecure APIs, stolen passwords, or improperly configured storage (such as AWS S3 buckets) might lead to the exposure of sensitive data. Without proper backup procedures, data might potentially be irretrievably lost due to ransomware and accidental deletion.
- **Insider Threats:** Credentials may be misused by malicious or negligent insiders with lawful access to steal or harm data. Not constantly keeping an eye on access records and user activity increases this danger.
- **Service Hijacking and Account Compromise:** Through techniques like phishing and keyloggers, attackers may obtain access to a system and impersonate legitimate users in order to escalate their privileges.
- **Insecure Interfaces and APIs:** Api reliance is high on cloud platforms. These interfaces provide a security risk because they enable attackers to circumvent authentication and get unauthorized access in the absence of adequate restrictions.
- **Advanced Persistent Threats (APT):** Advanced persistent threats (APTs) are a kind of cyberattack that may go undetected for an extended period of time, with the goal of collecting data or interfering with system functions.

Literature Review

Businesses have quickly put in place robust security measures after adopting cloud computing to back their structures, processes, procedures, and operations. To protect themselves against cyber assaults and data leaks, businesses are looking for ways to strengthen their digital environments. Cyber assaults are more common and more sophisticated than ever before. When it comes to finding, safeguarding, and stopping cyber assaults, it is a great component of cloud environment security. Yusuf, Shefiu.

(2025) ^[1].

Cloud computing's quick development has brought new security concerns, calling for sophisticated defenses to counteract cyberattacks. Complex attack vectors are outpacing the capabilities of traditional intrusion detection systems (IDS), necessitating smart and adaptive solutions. Improving cloud security through the use of AI Intrusion Detection Systems (AI-IDS) is the focus of this study. Smith, Jordan & Kevin, Emily. (2025) ^[2].

The creation and assessment of sophisticated machine learning models for cloud-based intrusion detection is the subject of this article. We look at the results of traditional GNNs, CNNs, LSTMs, Isolation Forests, and Transformer-based Spatio-Temporal Graph Neural Networks (ST-GNNs) on three separate datasets: NSL-KDD, CICIDS2017, and our own synthetic dataset. Wang, Fei & Xie, Sanshan. (2025) ^[3].

By allowing enterprises to disperse workloads across several cloud service providers, multi-cloud environments have drastically changed current IT infrastructures. Heterogeneous platforms, dynamic resource provisioning, and dispersed network borders all contribute to complicated security issues, even while this design improves scalability, flexibility, and resilience.

The increasing complexity and unpredictability of cyberattacks has made intrusion systems driven by AI an absolute need. Through studying attack patterns, adapting to new behaviors, and discovering previously undisclosed weaknesses, contemporary artificial intelligence systems enhance intrusion detection, as examined in this study. Akhilesh & Patel, Hemant & Gautam, Chandra Shekhar. (2025) ^[4].

Protecting these settings has becoming more difficult as cloud computing keeps taking over the world of digital infrastructure. By continuously scanning network traffic for signs of intrusion, Intrusion Detection Systems (IDS) hosted in the cloud provide an extra line of protection. Unfortunately, there are a lot of obstacles to overcome when it comes to optimizing and implementing IDS in cloud systems. Micheal, Dave. (2025) ^[5].

By improving the capacity to identify and react to complex cyber threats, Intrusion Detection Systems (IDS) driven by Artificial Intelligence (AI) have transformed cybersecurity. In order to monitor network traffic, detect abnormalities, and forecast possible assaults in real time, these systems use advanced techniques like deep learning, machine learning, and natural language processing.

Research Methodology

This study employs an experimentally-driven, systematic research methodology to evaluate the effectiveness of machine learning models for network intrusion detection. The methodology is structured to ensure a rigorous and reproducible investigation of how various algorithms perform when detecting abnormal traffic patterns in benchmark datasets.

Cybercrime is expected to cause damages of up to \$10.5 trillion annually by 2025, highlighting the need for improved intrusion detection and prevention systems (IDPS). Traditional signature-based intrusion detection systems (IDS) fail to identify novel attacks without known patterns. In contrast, machine learning algorithms can

analyze large volumes of network traffic data to identify subtle anomalies and enhance detection rates continually. This study aims to investigate AI-based intrusion detection models capable of recognizing abnormal network behavior with high accuracy. It involves both conceptual and experimental research, proposing a framework based on literature and testing machine learning models with publicly available datasets (e.g., NSL-KDD, UNSW-NB15) under controlled conditions. The research seeks to enhance IDS effectiveness in identifying malicious traffic patterns across various attack types and environments.

Research Design

This study adopts a quantitative experimental research design focused on evaluating the effectiveness of various machine learning models for network intrusion detection. The research is structured in two main phases: (1) conceptual development of an AI-based intrusion detection framework, and (2) empirical testing of machine learning algorithms on benchmark intrusion detection datasets. The study does not involve implementation in a live production environment, but rather conducts experiments under controlled settings using publicly available datasets.

Data Collection and Description

Training and testing datasets must be diverse, relevant, and of high quality in order for machine learning algorithms to perform well in intrusion detection. We used four well-known cybersecurity datasets-CICIDS2017, UNSW-NB15, NSL-KDD, and ToN_IoT-that are accessible to the public for our investigation. These were chosen to allow for the comparison of both legacy and future network environments, including networks based on the Internet of Things.

Results

Method of the Development Assignment

Conducting a battery of tests on four industry-standard intrusion detection data sets was a part of the development

assignment. In order to get comparable findings, the same procedure was used for all data sets. In the beginning, we did data preprocessing and EDA (exploratory data analysis). Data preparation included encoding category attributes, cleansing the data (such as removing missing values), and scaling numerical attributes as needed. To illustrate the point, one-hot encoding or label encoding were used to transform categorical field data (such as protocol types or service names) into numerical type. For the purpose of training Support Vector Machines (SVMs) and Multi-Layer Perceptrons (MLPs), continuous features were normalized using methods like min-max scaling or standardization. The models were trained using data that was already class-imbalanced and performance measures that were selected to address this issue, rather than using oversampling or augmented data.

Several measures were used to assess distinct elements of effectiveness while evaluating the model's performance. The ratio of occurrences that were properly categorized was used as an easy way to quantify overall accuracy, which is known as accuracy. The accuracy score alone might be deceiving when dealing with unbalanced data, thus we also presented the F1-score (the harmonic mean of recall and precision) for each model. In the multi-class instance, we focused on class-wise F1 and weighted-average F1. The effectiveness of minority attack classes might be seen from this. To examine the distribution of misclassifications across classes, a confusion matrix was created for each classifier. These matrices are particularly helpful in a multi-class scenario for seeing which types of assaults are regularly misclassified. To evaluate the trade-off between the true positive rate and the false positive rate in a binary classification issue, we constructed Receiver Operating Characteristic (ROC) plots and computed the Area Under the Curve (AUC). With a large area under the curve (AUC), the model is quite good at differentiating between benign and malicious data at different thresholds. Figure 2 shows a confusion matrix for one data set's multi-class categorization as an example.




Fig 2: Confusion matrices for multi-class classification on the ToN_IoT dataset using (a) Linear SVM, (b) MLP Classifier, and (c) Random Forest

Each matrix illustrates the classifier's ability to distinguish among ten attack types. Most models show high precision along the diagonal, with minor misclassifications observed in classes 3 and 4.

The appendix contains the Jupyter Notebooks that were used to arrange the analysis of each dataset during the development process. While adhering to the general

consistent methodology, this modular approach made sure that data pretreatment, model training, and assessment were customized to each dataset's individual properties. Jupyter Notebooks also made it easier to do in-depth EDA, which helped with things like seeing feature significance and class distributions graphically, which aided in understanding the findings.

Development Assignment Data and Collection

This study used four publicly available intrusion detection datasets: CICIDS2017, UNSW-NB15, NSL-KDD, and ToN_IoT. These datasets were chosen because they provide a wide range of assessment criteria for AI-based intrusion detection system models, including different time frames, network environments, and attack kinds. The datasets have

been extensively used in cybersecurity research and are all annotated, allowing for the accountability of model performance to well established baselines. The number of records, features, and attack classifications are summarized in Table 1, which also provides a summary of each dataset. The following is a synopsis of each dataset, including its collection methods and the types of attack traffic it includes.

Table 1: The intrusion detection datasets used in this study.

Dataset	Year	Total Records	Features	Attack Classes (multi-class)
CICIDS2017	2017	~2,830,000 Flows	80 (traffic features)	1 normal + ~13 attack types (Brute Force, DoS, DDoS, PortScan, Botnet, Web attacks, Infiltration, etc.)
UNSW-NB15	2015	2,540,044	49 features	1 normal + 9 attack families (filegdp, inppqun, jmfkehgkxes) (Fuzzers, Analysis, Backdoor, DoS, Exploits, Generic, Reconnaissance, Shellcode, Worms)
NSL-KDD	2009	148,517 connections (125,973 train, 22,544 test)	41 features	1 normal + 4 attack categories (DoS, Probe, R2L, U2R; multiple sub-attack types in each category)
ToN_IoT	2020	Millions of network flows (subset used for training/testing)	42 features (network traffic subset)	1 normal + multiple attack types (DoS, DDoS, scanning, ransomware, etc.)
NSL-KDD	2009	148,517 connections (125,973 train, 22,544 test)	41 features	1 normal + 4 attack categories (DoS, Probe, R2L, U2R; multi-ple sub-attack types in each category)
ToN_IoT	2020	Millions of network flows (subset used for training/testing)	42 features (network traffic subset)	1 normal + multiple attack types (DoS, DDoS, scanning, ransomware, etc.)

Analysis of Data

CICIDS2017 Binary Classification (Normal vs Attack)

The four models demonstrated remarkable accuracy in detecting attacks on CICIDS2017, indicating that distinguishing between legitimate and malicious traffic in this dataset is a breeze. The Random Forest classifier outperformed the others, demonstrating near-perfect recognition of attack flows and normal flows with a binary classification accuracy of almost 99.9 percent. Similar to the Random Forest, the MLP neural network achieved remarkable results, approaching 99.7 percent accuracy. The SVM achieved an accuracy of around 99.3%, which is somewhat lower. This might be because to scaling constraints imposed by the whole data set's complexity or a small number of edge instances that it incorrectly categorized. With an accuracy rate of around 98.6%-a high rate-but obviously more errors than the non-linear ones, Logistic Regression was the weakest of the four on this dataset, but it was still good.

Multi-class Classification (Multiple Attack Types)

The performance degrades while differentiating between different sorts of attacks on CICIDS2017, but the models maintain their excellent accuracy. Once again, in the multi-class example, Random Forest produced the best overall accuracy, coming in at about 98.7 percent. Almost all flows (almost 99 out of 100) were correctly identified as benign or as the right sort of assault. Afterwards, the MLP achieved an accuracy of 97.0%, and SVM followed with 96.1%. The accuracy of Logistic Regression plummeted to around 93.2% in multi-class.

UNSW-NB15 Binary Classification (Normal vs Attack)

Despite the models' claims of great accuracy, distinguishing between attack and routine traffic on UNSW-NB15 proved a modest challenge. On the test set, Random Forest achieved a binary classification accuracy of almost 95.0%, demonstrating its superior performance once again. This

indicates that RF was successfully determining the maliciousness of a connection 95% of the time. Additionally, the MLP was performing well, with an accuracy rate of around 92.0%. SVM achieved over 90% accuracy, whereas Logistic Regression achieved around 88% accuracy. The larger degree of overlap between benign and attack patterns in UNSW-NB15 is the reason why these numbers are lower than those in CICIDS2017.

Multi-class Classification (Multiple Attack Families)

With nine distinct attack families and an unequal distribution of classes, UNSW-NB15 presents a formidable multi-class problem. In terms of overall multi-class accuracy, the Random Forest classifier managed to reach about 89.4 percent on the test set. With 10 classifications (normal included), this is a respectable level of accuracy, suggesting that most traffic cases were appropriately classified. Among the three methods tested, MLP achieved an accuracy of around 85.7%, SVM approximately 82.1%, and Logistic Regression about 80.0%. As before, these numbers represent a roughly ranked order of models, with LR being the worst and RF being the best.

NSL-KDD Binary Classification (Normal vs Attack)

Although it was lower in NSL-KDD compared to more current datasets, the binary classification was still good. When it came to differentiating between legitimate and malicious traffic on the test set, the Random Forest achieved an accuracy of around 85.0%. This means that it classified 85 percent of links as either safe or harmful. The 83–84% accuracy rates of the MLP and SVM were comparable. With an accuracy of around 80.0%, Logistic Regression trailed. The ~99% of CICIDS2017 and the ~95% of UNSW-NB15 are somewhat higher, but these numbers are still acceptable.

Multi-class Classification (Multiple Attack Categories)

The five classes used to assess NSL-KDD in the multi-class scenario were Normal, DoS, Probe, R2L, and U2R, with the

latter four representing attack types. On the test set, Random Forest achieved an overall accuracy of around 80.2% for the classes. With an accuracy of around 75.0%, MLP, SVM, and LR were all above average. Now that the focus is on accurately marking the sort of assault, it is expected that accuracy would decline. The Random Forest confusion matrix for NSL-KDD. If we apply Random Forest to the NSL-KDD confusion matrix, we will get a jumble of results. The Denial-of-Service (DoS) category, which include attacks such as SYN flood and smurf, performed very well when tested using RF. T

ToN_IoT Binary Classification (Normal vs Attack)

The models achieved respectable accuracy on the ToN_IoT binary classification test (identifying malicious or benign traffic), but at a somewhat lower level than on CICIDS2017. On the test set, the Random Forest achieved an accuracy of about 94.5 percent. This indicates that it accurately classified about 95% of the flows as either harmless or harmful. The accuracy rates for the MLP, SVM, and Logistic Regression were about 90%, 88.0%, and 85.0%, respectively. Because of the somewhat wider difference between RF and the remainder than with previous sets, it is possible that a decision tree ensemble may detect more complicated patterns in the ToN_IoT dataset than a linear model or support vector machine (SVM).

Multi-class Classification (Multiple Attack Types)

Different types of attacks (beyond the usual) were distinguished using ToN_IoT's multi-class categorization. To name a few types of attacks, the dataset description states that ToN_IoT is vulnerable to distributed denial of service (DDoS), scanning, ransomware, and perhaps data exfiltration or injection attacks. Across all classes, the Random Forest achieved an accuracy of around 95.0% in our trial. This is a remarkable result, showing that the model gets the class of 95% of flows exactly right. Similar to their binary performance ranking, although somewhat lower in the multi-class situation, the MLP neural network obtained about 90% accuracy, the SVM around 88%, and Logistic Regression around 85.0. It is highly probable that the Normal class and large attack classes, such as DoS and DDoS, are clearly separated in the Random Forest's confusion matrix on ToN_IoT (which is conceptually similar to the ones mentioned earlier).

The ToN_IoT analysis: applying AI models to intrusion data from IoT networks shows great promise, with top-tier models attaining impressive detection and classification rates. To tackle the complexity of IoT data, one must be meticulous when selecting and configuring models, since there are slight performance decreases while moving from binary to multi-class and across multiple models.

Since Random Forest is so dominant across all datasets, including ToN_IoT, ensemble approaches should be considered the gold standard for intrusion detection issues.

Discussion

Overarchingly, our findings across all four datasets demonstrate that AI models, particularly Random Forest and deep learning techniques, are capable of accurately detecting known attack types. But they also point out that addressing class disparity is a constant struggle. Despite

calls for more robust model explainability, resampling methods, and hybrid learning strategies in the literature, models failed miserably in every dataset when it came to representing underreaching threats. These results support the core theoretical claims of this thesis, which state that AI may significantly improve IDS performance, but that a multi-layered, complex strategy is necessary to guarantee efficacy against varied and ever-changing attack scenarios. From logistic regression and support vector machines (SVMs) to decision trees, XGBoost, Random Forest, and artificial neural networks (ANNs), we put a plethora of ML models through their paces in our experiments to see how well they detected intrusions. We have made an effort to demonstrate the current theoretical discussions in IDS research by contrasting standard algorithms with deep learning approaches.

Conclusion

This study supports the idea that intrusion detection systems powered by AI can effectively identify suspicious patterns in network traffic that might indicate a cyberattack. We showed that machine learning models can outperform existing detection approaches in terms of accuracy and detection rates by designing and evaluating them on four standard intrusion detection data sets. More specifically, our models virtually always accurately detected common attack techniques including port scans, brute-force, and DoS/DDoS floods. The findings back up the main premise, which states that advanced signs of compromise may be used to train machine learning algorithms to successfully detect intrusions automatically. This allowed us to accomplish the goals of this thesis, which were to train and evaluate ML models for intrusion detection, compare their performance (binary vs. multi-class, across techniques and datasets), and research their resilience and generalizability.

References

1. Yusuf S. Developing AI-powered intrusion detection system for cloud infrastructure. 2025.
2. Smith J, Kevin E. AI-powered intrusion detection systems for next-generation cloud. 2025.
3. Wang F, Xie S. Cybersecurity in cloud computing: AI-driven intrusion detection and mitigation strategies. IEEE Access. 2025. doi:10.1109/ACCESS.2025.3580569.
4. Wao A, Patel H, Gautam CS. AI-powered driven intrusion systems in cyber security and zero-day attack detection. International Journal of Scientific Research in Engineering and Management. 2025;9:1-9. doi:10.55041/IJSREM54733.
5. Micheal D. Cloud-based intrusion detection systems: challenges and best practices. 2025.
6. Cate M. AI-powered intrusion detection systems: challenges and opportunities. 2025.
7. Johnson A, Townsend R, McCoy D. AI-based cybersecurity frameworks for enterprise cloud systems. 2025.
8. Hamilton B, Onami S. Artificial intelligence-driven intrusion detection systems for next-generation cyber threats. 2023.
9. Chauhan G. AI-driven intrusion detection systems: enhancing cybersecurity with machine learning

- algorithms. 2019;7:131-139.
10. Modi N. AI-powered intrusion detection using CNN-LSTM for cloud and edge networks: a hybrid deep learning approach. 2025. doi:10.21203/rs.3.rs-6928225/v1.
 11. Desai A, Kowalski L, Al-Mutairi Y, Wei C, Hughes V, Paul C, et al. AI-powered intrusion detection systems in cloud environments. 2025.
 12. Motala B, Ngubane SN, Yedwa L. Systematic review artificial intelligence in cybersecurity: AI-driven threat detection and response systems. 2025. doi:10.5281/zenodo.17372293.
 13. Lindsey O, College W, Mohammed O. AI-driven cybersecurity for cloud-based data systems: anomaly detection, threat mitigation, and performance optimization. 2025. doi:10.13140/RG.2.2.18277.10721.
 14. Tarafdar R. AI-powered cybersecurity threat detection in cloud environments. *International Journal of Computer Engineering & Technology*. 2025;16:3858-3869. doi:10.34218/IJCET_16_01_266.
 15. Carter M, Hayes J, Miller R, Thompson E, James C. AI-based intrusion detection for multi-cloud workloads. 2025.
 16. Brian J, Alexander D. AI-powered intrusion detection systems for next-gen cloud. 2023.
 17. Gudimetla S, Kotha N. AI-powered threat detection in cloud environments. *Turkish Journal of Computer and Mathematics Education*. 2018;9:638-642. doi:10.61841/turcomat.v9i1.14730.

Creative Commons (CC) License

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license. This license permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.