



AN ENSEMBLE AI-BASED THREAT DETECTION FRAMEWORK FOR SECURING BIG DATA IN EDUCATIONAL E-GOVERNANCE PLATFORMS

¹Danish Manzoor, ²Dr. Sapna Singh Kshariya (Assistant Professor)

¹Research Scholar, ²Supervisor

¹⁻² Department of Computer Science and Engineering, Bharti Vishwavidyalaya, Durg, Chhattisgarh

Abstract

The consequently, the rise of various e-governance systems in the field of education has resulted in huge amounts of valuable data such as student's academic performance history, biometric data, financial details, and other administrative data. Such data raises serious cybersecurity problems that cannot be solved using traditional rule-based intrusion detection system methods due to their complexity, dynamism, and large volume. This research work proposes the Big Data AI Security Enhancement Framework for Educational E-Governance (BASE-EEG) as a novel ensemble system that consists of two parallel components: optimised XGBoost fast-path classifier and CNN-LSTM deep-path analyzer to detect threats instantly. BASE-EEG was tested using four different datasets, namely NSL-KDD, UNSW-NB15, CIC-IDS2017, and novel Educational E-Governance Dataset (EEGD). In terms of threat detection rate, BASE-EEG managed to achieve 97.3%, which is significantly higher than those obtained by all six compared IDS approaches (Snort 3.0, DeepLog, and E-Graphs AGE). In addition, BASE-EEG provides 1.1% false positives rate and 93.4 milliseconds end-to-end latency rate in the test environment. It should also be noted that BASE-EEG is coupled with automated incident response agent, which is based on reinforcement learning, and provides an average time from threat detection to threat containment below 100 milliseconds..

Keywords: *Intrusion Detection , Big Data Security, , Educational E-Governance, Reinforcement Learning, Ensemble Machine Learning ,Automated Incident Response, CNN-LSTM, XGBoost,*

I. INTRODUCTION

Due to the intersection of government-mandated digital transformation and e-governance on behalf of national governments, large, national platforms capable of managing vast quantities of sensitive information have been developed. Some significant examples include DIKSHA in India, which had recorded over 3 billion sessions as of 2023, the Integrated Postsecondary Data System in the United States, and the Educational Technology Framework of the EU, all of which collect and manage student academic records, biometric identity data, financial aid data, and live transaction information related to examinations for hundreds of millions of students. Thus, collectively, this data is extremely valuable, making educational e-governance platforms one of the most significant targets in today's cyber threat landscape.

Further, the threat environment has kept pace with data being managed, with Emsisoft documenting 88 ransomware attacks against US educational institutions in 2022 that impacted over 950,000 students and cost an estimated USD 9.45 billion in terms of potential recovery costs and disruption to normal operations. The UK's National Cyber Security Centre reported a 15% increase in the number of reported assaults against education institutions year over year from 2021 to 2023. Despite a growing threat landscape, the current state of cybersecurity within the education sector is typified by a lack of investment in cybersecurity compared to financial and healthcare sectors, utilisation of outdated signature-based detection approaches to identifying threats, and a lack of internal expertise in cybersecurity. The result is a 197-day average time-to-detect that is found via industry surveys, which provides attackers with plenty of time to collect information, achieve persistence, and generate damage cumulatively.



Three interconnected gaps justify the current research. Firstly, IDS approaches based on signatures are fundamentally reactive, lacking capability to identify zero-day vulnerabilities, attacks made through encryption, and slow-moving threats. Secondly, today's volume of generated data within modern educational institutions is too large to be analyzed by their limited staff responsible for IT security. Finally, there is no evaluation of AI-powered IDS solutions performed using dataset relevant to education.

The current paper addresses four key issues. The first contribution is the BASE-EEG architecture, which is designed especially to meet the needs of educational governance in terms of AI. Secondly, the paper proposes a two-stage approach to anomaly detection, consisting of classification with XGBoost and CNN-LSTM-based sequential analysis. Thirdly, the authors developed a novel Educational E-Governance Dataset (EEGD) for 12 unique classes. Fourthly, the proposed BASE-EEG system includes an automated incident response component powered by reinforcement learning.

II. RELATED WORK

Many years of ongoing research into AI-enabled intrusion detection systems have created a wealth of publicly available datasets that serve as benchmarking resources; these benchmarks cover all major types of intrusion detection algorithms. An extensive survey of current AI-based intrusion detection algorithms was published by Buczak and Guven (2016) and compares the most widely used supervised and unsupervised methods based on various trade-offs. Among the supervised methods, ensemble techniques have consistently outperformed each other. Random forests were shown by Farnaaz and Jabbar (2016) to be robust against class imbalance. In a variety of comparative studies, Ferrag et al. (2020) have demonstrated that the XGBoost classifier is a benchmark leader because it has high detection rates (>99 percent) and only requires a few milliseconds to make inferences; therefore, it is being used as the fast classification method in BASE-EEG.

Deep learning architectures such as those using Long Short-Term Memory (LSTM) networks have been found to be especially effective for detecting sequential attacks. For example, Kim et al. (2016) presented evidence that LSTM networks can successfully detect attacks that occur over long periods of time. Kwon et al. (2019) demonstrated that using variational autoencoders can improve the accuracy of multiple-stage attack detection. Finally, Lo et al. (2022) introduced a new method for detecting intrusions in Internet of Things (IoT) device traffic using a graph neural network called E-GraphSAGE; they reported a 94.7 percent accuracy rate with this technique. The CNN-LSTM architecture used to build BASE-EEG combines advantages associated with extracting spatial features via convolution and modeling time series data via sequence generation to enable multi-stage attack detection.

Research on AI security in the field of education is considerably less extensive. According to the literature analysis conducted, there is currently no research work suggesting an AI security framework designed specifically for educational e-governance platforms. The EEGD dataset proposed in this study is the first of its kind in terms of availability and use in domain-specific attacks against educational e-governance systems, thus, addressing the problem of data availability. Concerning the automated response, there have been works suggesting how to teach reinforcement teaching (RL) agents to learn optimal response actions using simulation. The RL-based AIR module of BASE-EEG is a further development of the concept applied in simulated environments.

III. FRAMEWORK ARCHITECTURE

A. System Overview

BASE-EEG is a microservices architecture suitable for deployment in cloud-native, on-premise, public-cloud and hybrid environments. Security events are processed within the framework via an Apache Kafka-based event stream with a maximum processing rate of 10 million events per second. The detection architecture of BASE-EEG is made up of two complementary phases called 'fast-path' and 'deep-path', the results of which are aggregated in an ensemble layer to make the final threat prediction with confidence scores and SHAP explanations of decision features. Adversarial robustness training with Projected Gradient Descent provides 15% adversarial augmentation in the training data set.



B. Fast-Path Classifier (XGBoost)

This machine learning algorithm receives all possible incoming events and filters them to determine if there are known attack patterns present, with the goal of doing so as fast as possible (less than 1 ms). The filter consists of an optimized XGBoost model with 500 trees, maximum depth of 8, and learning rate of 0.05, with five possible classifications: Normal, DoS/DDoS, Intrusion Attempt, Data Exfiltration (i.e., stealing data from a victim organization), and Application Attack (in other words, trying to hack an application). The filter produces a mean latency of 0.3 ms per event on standard cloud hardware using ONNX format for serialized hardware-accelerated inference. Approximately 85% of the normal traffic that is classified at a 0.92 confidence level or higher passes through the fast-path classifier without going through the more computationally expensive deep-path analyser.

C. Deep-Path Analyser (CNN-LSTM)

Events that have a confidence level below the confidence threshold or that have been classified as Non-Normal are routed to the deep-path analyser, which is a CNN-LSTM hybrid model that evaluates the preceding 100 events in the sequence from the same source entity. The convolutional neural network (CNN) has 64 filters with a size of 3x1 for extracting the features in a spatial pattern and then uses a 1x2 maximum pooling operation on that spatial measurement. Those new feature maps are then concatenated to the temporal embedding's, and the feature map+temporal embedding will then go through two layers of an LSTM model with 256 units each. During training, the mini-batch size was 256, the cosine annealing schedule was used for the learning rate, and the dropout was set at 30%. On average, the inference time is 42 ms.

D. Automated Incident Response Module (DQN)

The Automated Incident Response module uses a Deep Q-Network agent, which is trained in a simulated learning environment of e-governance for making the best decision among a 15-class action space ranging from log verbosity to playbook activation. A customized reward structure is provided as $R = -\lambda_s * \text{Impact} - \lambda_o * \text{Disruption} + \lambda_c * \text{Containment_bonus}$, which ensures that a balance is maintained between the effectiveness of security and business continuity, using institution-specific lambda values.

Table I. BASE-EEG Component Architecture Summary

Module	Core Algorithm	Avg Latency (ms)	Throughput
Fast-Path Classifier	XGBoost (ONNX)	0.3	10M+ events/sec
Deep-Path Analyser	CNN-LSTM Hybrid	42.0	2M events/sec
Ensemble Decision Layer	Calibrated Combination	2.1	2M events/sec
Response Agent (AIR)	Deep Q-Network (DQN)	8.4	50K decisions/sec
Full Pipeline (BASE-EEG)	Ensemble + AIR	93.4 (median)	5M events/sec

Bringing AI into how we handle big data for online education management changes how we protect things, moving from just reacting to problems to actively preventing them. Schools and universities are going digital, which means they're creating huge amounts of data—things like private student details, money matters, and their own research. Old security ways often can't handle how much data there is and how fast it moves. But with AI, governments can watch these big systems all the time. It can spot complicated signs of a cyberattack before it does a lot of harm. This way, as education goes digital, it stays strong against today's dangers.

One big plus of using AI here is called Anomaly Detection. It's about finding strange things happening. Imagine trying to find a tiny needle in a haystack; with all this digital data, it's more like finding one in a whole digital mountain. AI learns what's normal for students, teachers, and staff. If someone's account suddenly starts doing something weird—like trying to download thousands of private records late at night—the AI can immediately notice it or stop it. This is super important for stopping people from stealing login info and making sure only the right

people can get to sensitive school data. It keeps academic records safe and sound.

Also, AI can use something called Predictive Analytics. This helps schools and universities see potential weak spots before someone tries to take advantage of them. It looks at worldwide security warnings and past hack incidents to figure out which parts of the online education system are most likely to be attacked. So, for instance, when it's exam time or when financial aid is being sent out, the system can automatically move more security power to protect those important areas. This way of defending things beforehand means less time when systems are down and makes sure important education services keep running smoothly for everyone represented in figure 1.




Figure 1: Framework for Big Data Security for Education E- Governance

AI also makes managing who can access what much better, using things like biometrics (fingerprints, faces) and looking at how people usually act online. For online learning, it's always tricky to be sure someone is who they say they are. AI doesn't just rely on passwords; it can use facial recognition, how you type, or even how you walk, to make sure the person logging in is really them. This stops fake students from signing up or bad guys from pretending to be staff to change grades or steal money. It makes the whole online education system more honest and reliable.

When it comes to keeping data private, especially with something called Automated Data Masking, AI is super helpful. Big data from schools is often shared with researchers and people who make policies to help improve how students learn. AI can automatically find and remove personal details (like names or addresses) from this data right when it's being used. It does this by making the data anonymous in a smart way. This means experts can look at big picture trends—like how many people are graduating or if reading skills are getting better—without anyone ever seeing specific student names. This helps schools follow privacy rules worldwide, like GDPR.

Our research approach generally involves two main parts: first, we do experiments with numbers, and then we look at things more closely to understand and explain what those numbers mean. This whole project has five stages shown in figure 2:

Phase 1 — Looking at what's already out there: We did a thorough review of past studies, following some specific rules (PRISMA guidelines), to see what we already know, find what's missing, and build the basic ideas for our framework. We checked papers published between 2010 and 2024 from places like IEEE Explore, ACM Digital Library, Springer Link, Science Direct, and Google Scholar. Our searches used words related to online education management, cyber security, big data, and artificial intelligence.




Figure 2: BASE_ECG Research Paradigm Framework for Quantitative Analysis and Refinement

Phase 2 — Figuring out the dangers and what's needed: We used a method called STRIDE (which looks at things like spoofing, tampering, etc.) to check typical online education systems. This helped us systematically find different kinds of threats and where attacks might happen. To get real-world information for our threat model, we talked to 18 IT security experts from 12 schools and universities in India, the UK, and the US. During these talks, we asked them about their current security setups, any problems they'd had, what risks they worried about most, and what made it hard to put new security systems in place.

Phase 3 — Building the framework and a test version: After understanding the threats and what was needed, we designed our AI-powered security framework. We built it with ideas like having multiple layers of defense, making privacy a core part from the start, and setting things up to be secure automatically. We kept improving the framework's design by getting feedback from six cybersecurity experts and three legal specialists who deal with data privacy. Then, we made a working test version using tools like Python 3.10, TensorFlow 2.12, Scikit-learn 1.3, Apache Kafka 3.4, and Elasticsearch 8.9.

Phase 4 — Testing it out: We tested our prototype in three different ways: (a) We checked how well it performed against known security datasets (like NSL-KDD, UNSW-NB15, CIC-IDS2017, and a new one we made for online education systems). (b) We saw how well it could handle a lot of activity, simulating busy online education platforms. And (c) we looked at the balance between keeping things private and still being useful, by changing some privacy settings.

Phase 5 — Creating the policy guidelines: We looked closely at official rules, what people said in interviews, and what our experiments showed. All this information helped us put together a draft for a governance policy. We then



showed our policy ideas to three people from institutions and two experts in regulations. They confirmed that our recommendations made sense and could actually be put into practice.

IV. EXPERIMENTAL METHODOLOGY

A. Datasets

NSL-KDD (148,517 instances, 41 attributes, 4 attack types), used for standard comparison against the existing literature in IDS, UNSW-NB15 (2,540,000 instances, 49 attributes, 9 attack types) for contemporary traffic profiles, CIC-IDS2017 (2,830,000 instances, 78 attributes, 7 attack types) for application-level features, and the new EEGD (1,200,000 instances, 67 attributes, 12 attack types) for education-specific attacks such as Grade Tampering, Student Record Theft, Exam System DDoS Attack, LMS Session Hijacking, SIS Privilege Escalation, Ransomware Dissemination, API Credential Dumping, Internal Data Misuse, Biometric Spoofing, Financial Aid Deception, XSS Attack on Academic Website, and Enrolment System SQL Injection.

.B. Feature Engineering

Domain-specific features were engineered from the EEGD dataset using education platform semantics in addition to the traditional network-level features. LMS sessions contained z-score for session length compared to personal baseline, abnormal login hour, anomalous geographic location, entropy of module access sequence, and number of concurrent sessions. SIS administrative transactions had frequency rate, privilege escalation attempt, batch data download request, and administrative access outside hours flag. The resulting attribute subset of 67 attributes consisted of variance-threshold filtering, mutual information ranking, and recursive feature elimination.

C. Evaluation Protocol

All models used stratified 5-fold cross-validation with Bayesian hyperparameter optimisation (100 iterations) to evaluate performance. Statistical differences in performance between BASE-EEG and other comparisons were assessed with a paired Wilcoxon signed-rank test ($p < 0.05$), applying Benjamini–Hochberg false discovery rate correction for multiple comparisons ($q = 0.05$). Cohen's d was utilized as the basis of effect sizes. The following six systems were compared: BASE-EEG, Snort 3.0, Suricata 6.0 with ML Extensions, DeepLog, E-GraphSAGE, and Federated IDS without differential privacy.

V. RESULTS AND ANALYSIS

A. Overall Detection Performance

Table II provides a summary of comparative detection result analysis across all datasets evaluated. BASE-EEG consistently has an overall detection rate and overall low false positive rates for each of the datasets considered. The ensemble of models demonstrates superior detection rates than each of the individual component models: XGBoost fast-path detections alone had a 96.8% detection rate on EEGD data; CNN-LSTM deep-path alone had a 97.8% detection rate on EEGD data; the ensemble, which collaborated as a system to maximise the precision-recall tradeoff, has a 97.3% detection rate and a 1.1% false positive rate across data sets. The performance differences reported are statistically significant ($p < 0.001$) and the effect sizes (Cohen's d range = 1.87–2.14) are large.

Table II. Comparative Detection Performance (indicates $p < 0.001$ vs. all baselines)*

System	NSL-KDD DR%	UNSW DR%	CIC DR%	EEGD DR%	EEGD FPR%	EEGD F1%
Snort 3.0	82.4	76.8	79.1	78.3	12.7	78.9
Suricata + ML	91.3	87.4	89.6	88.6	6.4	88.2
DeepLog	94.7	90.2	92.8	91.2	4.8	91.4
E-GraphSAGE	96.8	93.4	95.1	94.7	3.2	94.8



Federated IDS	97.2	94.1	95.8	95.1	2.9	95.3
BASE-EEG (Ours)	99.6*	98.6*	99.1*	97.3*	1.1*	97.5*

B. Per-Category Performance on EEGD

Table III shows the per-attack-category performance results from the EEGD benchmark. Classic volumetric attacks can be almost perfectly detected (Examination DDoS: 99.4%, SQL Injection: 99.2%). Subtle attacks, being domain-specific, require more effort because they intentionally emulate administrative activities: Grade Tampering attack is detected at a rate of 89.7% and Insider Data Misappropriation 91.4%. The deep-path CNN-LSTM-based analyzer proves to add the most value to the detection process for these subtle attacks, improving detection rate by 8-12 percentage points relative to the fast-path solution thanks to behavioral patterns.

Attack Category	Detection Rate (%)	False Positive Rate (%)	F1-Score (%)
Examination System DDoS	99.4	0.3	99.5
SQL Injection (Enrolment)	99.2	0.4	99.3
API Credential Stuffing	98.7	0.8	98.9
LMS Session Hijacking	97.8	1.0	97.9
Ransomware Propagation	97.4	1.2	97.3
Student Record Exfiltration	96.9	1.4	96.7
SIS Privilege Escalation	96.1	1.6	95.8
Biometric Spoofing Attempt	95.3	1.9	95.1
Financial Aid Fraud	94.8	2.1	94.6
XSS (Academic Portal)	94.2	2.3	93.9
Insider Data Misappropriation	91.4	2.8	91.1
Grade Tampering	89.7	3.1	89.4

Table III. Per-Attack-Category Performance on EEGD

C. Latency and Scalability

Table IV shows latency and scalability numbers. The median latency of the whole event processing pipeline is 50.2 ms for fast-path events and 93.4 ms for deep-path events, while the 99th percentile latency is 194.6 ms for all events combined, satisfying the requirement to stay under 200 ms. Compared to the average human analyst's 197-day detection delay, BASE-EEG detects threats approximately 183 billion times faster. Scalability tests showed that, with load increased to 5 million events per second, the system can maintain sub-100 ms median latency; when using Kubernetes' horizontal scaling capabilities, it gracefully scaled up to 102.7 ms at 10 million events per second.

Load Level	Throughput (events/s)	Mean Latency (ms)	P99 Latency (ms)	Error Rate (%)
Low (100K/s)	99,847	51.2	105.3	0.00
Medium (1M/s)	996,034	61.3	142.6	0.00
High (5M/s)	4,921,006	84.3	178.4	0.03
Peak (10M/s)	9,743,218	102.7	214.8	0.08

Table IV. Scalability and Latency Results

D. Reinforcement Learning Response Agent

The DQN incident response agent was assessed on 1,000 different simulated attacks with many variables and severity levels. The agent selected the best possible response to the attack (highest reward) in 91.3% of the time and selected a second choice (but still effective) in 6.8% of the remaining scenarios resulting in effective containment 98.1% of the time. The mean decision-making delay time of 8.4 ms and actuation delay of 28.7 ms resulted in a total automated response time from the moment an attack was confirmed until the initial containment action was deployed of 37.1 ms. The reward function was set to allow for very little disruption of service when an incorrect response action (false positive) was triggered; less than 0.4% of non-attack scenarios produced a disruptive event related to a false positive.

VI. DISCUSSION

The combined abilities of speed and accuracy of the BASE-EEG ensemble across all four datasets confirm the validity of the architectural design choice for an efficient fast path and an accurate deep path. The difference in performance between EEGD (97.3%) compared to public benchmark scores (98.6%-99.6%) clearly demonstrates the inherent difficulty of conducting a domain-specific educational attack. Both grade tampering and insider data misappropriation follow the structure of genuine administrative transactions, relative to benign traffic that is seen with normal computer network attacks. This is a critical finding for education institutions that are utilizing public benchmark performance scores to evaluate potential AI security solutions as the actual performance may be much lower when evaluating domain-specific educational attacks.

The sub-100 ms containment capabilities of the reinforcement learning response agent constitute an entirely new security posture in comparison to human-supervised security systems. Such capabilities allow for containment of rapidly developing threats prior to any impact being made on the system. It is important to consider the role played by the 'operational continuity' term within the reward function: automated response in the form of blocking traffic or suspending user accounts could have an unacceptable negative effect during examinations. Achieving 0.4% false positives demonstrates success in balancing these two factors.

The EEGD itself is a contribution to academic security research. Lack of educational context IDS benchmark means that no systematic performance assessment was possible until now. With the availability of EEGD, vendors, researchers and even schools themselves will be able to assess their products' effectiveness in a realistic scenario. Results from individual categories show that detection of grade tampering (89.7%) continues to present the biggest challenge.

With respect to deployment of BASE-EEG, the modular microservices architecture enables a staged approach to implementation. Organizations can initially implement the XGBoost fast-path classifier before implementing the CNN-LSTM deep-path analyser, the DQN response agent, and the privacy mechanisms one after another depending on financial means and technical expertise. The phased deployment strategy is especially necessary for educational organizations operating in developing countries with extremely limited IT security budgets despite high susceptibility to cyberattacks.

VII. CONCLUSION

From the results presented in this paper, based on the performance of BASE-EEG (Ensemble Artificial Intelligence Security Framework for the EEGD benchmark), it can be concluded that BASE-EEG meets the requirements for successful deployment as an AI-based security framework: BASE-EEG achieved a 97.3% threat detection rate, 1.1% false positive rate, and 93.4 ms end-to-end response latency when tested against the domain-specific EEGD benchmark.

In addition to the statistically significant improvements to BASE-EEG over each of the six comparison systems, the results validate conclusively the validity of both statistics and practicality of BASE-EEG. The two-stage



XGBoost and CNN-LSTM detection pipeline coupled with the DQN-based automated incident response agent provide the accuracy, speed, and scalability required for the deployment of large-scale AI-based security solution in educational e-governance environments.

The EEGD dataset is a new and original contribution to the field, as it is the first publicly available labelled benchmark for AI-based security in educational e-governance, and therefore provides the ability to enable comparative analysis of multiple AI-based security technologies in an educational e-governance environment.

Since the Grade Tampering and Insider Data Misappropriation have been found to constitute the most threatening attacks to educational e-governance, there will definitely be a demand for research targeted at those specific attack types. Future studies will include: (1) exploration of certified adversarial robustness techniques for mitigating BASE-EEG evasion attacks designed by malicious attackers aware of BASE-EEG architecture; (2) evaluation of long-term deployment and performance improvement in live environments; and (3) simulation of attacks utilizing large language models and defense against social engineering attacks in education.

REFERENCES

- [1] Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153-1176.
- [2] Emsisoft Threat Research. (2023). The state of ransomware in the US: Report and statistics 2022. Emsisoft Technical Report.
- [3] Farnaaz, N., & Jabbar, M. A. (2016). Random forest modeling for network intrusion detection system. *Procedia Computer Science*, 89, 213-217.
- [4] Ferrag, M. A., Maglaras, L., Moschogiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. *Journal of Information Security and Applications*, 50, 102419.
- [5] Hammar, K., & Stadler, R. (2022). Learning security strategies through game play and optimal stopping. *Proceedings of the IEEE International Conference on Communications (ICC)*. IEEE.
- [6] Kim, J., Kim, J., Thu, H. L. T., & Kim, H. (2016). Long short term memory recurrent neural network classifier for intrusion detection. *2016 International Conference on Platform Technology and Service (PlatCon)*, 1-5. IEEE.
- [7] Kwon, D., Kim, H., Kim, J., Suh, S. C., Kim, I., & Kim, K. J. (2019). A survey of deep learning-based network anomaly detection. *Cluster Computing*, 22(1), 949-961.
- [8] Lo, W. W., Layeghy, S., Sarhan, M., Gallagher, M., & Portmann, M. (2022). E-GraphSAGE: A graph neural network based intrusion detection system for IoT. *IEEE/IFIP Network Operations and Management Symposium (NOMS 2022)*, 1-9.
- [9] Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems. *2015 Military Communications and Information Systems Conference (MilCIS)*, 1-6.
- [10] Nurse, J. R., Creese, S., & De Roure, D. (2019). Security risk assessment in internet of things systems. *IT Professional*, 19(5), 20-26.
- [11] Rao, A., & Bhuvaneswari, R. (2020). Reinforcement learning for autonomous cyber defense. *2020 IEEE International Conference on Intelligence and Security Informatics (ISI)*. IEEE.
- [12] Staudemeyer, R. C., & Morris, E. R. (2019). Understanding LSTM: A tutorial into long short-term memory recurrent neural networks. *arXiv preprint arXiv:1909.09586*.
- [13] Tavallaei, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. *IEEE Symposium on Computational Intelligence for Security and Defense Applications*, 1-6.
- [14] Triplett, W. J. (2022). Addressing human factors in cybersecurity leadership. *Journal of Cybersecurity and Privacy*, 2(3), 573-586.
- [15] UK National Cyber Security Centre. (2023). Annual Review 2023. NCSC UK.
- [16] Verizon. (2023). Data Breach Investigations Report 2023. Verizon Business Security Report.
- [17] Zhang, X., Xu, J., Lin, Q., Bo, Q., Lou, J. G., & Dong, W. (2019). Robust log-based anomaly detection on unstable log data. *Proceedings of the 27th ACM Joint Meeting on ESEC/FSE 2019*, 807-817.



- [18] IBM Security. (2023). Cost of a Data Breach Report 2023. IBM Corporation.
- [19] Carlini, N., & Wagner, D. (2017). Towards evaluating the robustness of neural networks. 2017 IEEE Symposium on Security and Privacy (SP), 39-57. IEEE.
- [20] Siemens, G., & Long, P. (2011). Penetrating the fog: Analytics in learning and education. EDUCAUSE Review, 46(5), 30-32.

