



Autonomous Threat Intelligence Systems for Dynamic Cyber Defense

Tanaya Jakir^{1*}, Shakila Sultana², Masrufa Tasnim³, Mst Anupa Nasrin Khan⁴,

Md Tarake Siddique⁵, Rasel Hossain Babu⁶, Md Himeluzzaman⁷

MS in Business Analytics, Trine University, USA¹

Master's in Accounting (M.A), National University, Bangladesh²

Master of Business Administration (MBA) – Human Resource Management, University of Dhaka, Bangladesh³

Master of Science in Zoology (Entomology Focus), The University of Rajshahi, Bangladesh⁴

Master of Science in Information Technology – Data Management & Analytics, Washington University of Science and Technology, Virginia, USA⁵

MS in Cybersecurity (STEM), Bay Atlantic University, Washington, D.C., USA⁶

M.Sc. in Information Technology Systems and Management, Washington University of Science and Technology, VA, USA⁷

tanayajakir@gmail.com¹, sultanashakila6@gmail.com², mtreema@gmail.com³, anupa516@hotmail.com⁴,
msiddique.student@wust.edu⁵, rbabu@stu.bau.edu⁶, md.himeluzzaman@student.potomac.edu⁷

ABSTRACT: Bustling city streets are filled with the sound of a tiger; the sound of your footsteps as you walk through them leaks into an orchestra that neither has harmony nor rhythm. Lights from sky-scrapers, which are huge and high in the sky at night, pierce into their night's black expanse, sending a strange light cascading around the metropolis. Swarms of people flow through the gaps between skyscrapers, looking like lone figures standing at the hub of an otherwise teeming world. Crowds line the streets, and the air hangs heavy with the smells from both exhaust gases and steaming pots of food being sizzled up, commingling it all. Yet--though the noise is deafening and never-ending--there is still an inherent energy which runs through the city, a non-stop heartbeat that sends it forward into the night.

Five classification algorithms Logistic Regression, Random Forest, Gradient Boosting, Support Vector Classifier, and K-Nearest Neighbors and five regression algorithms Linear Regression, Random Forest Regressor, Gradient Boosting Regressor, XGBoost, and Support Vector Regressor are comprehensively benchmarked. The obtained results show that the best ROC - AUC of 0.822 related to attack spike prediction can be achieved by Logistic Regression, whereas a near-optimal RMSE of 0.422 on attack volume forecasting is reached using Gradient Boosting Regressor and Linear Regression. Correlation analysis reveals a significantly strong correlation between scan_rate and attack_volume_current ($r = 0.74$), exposing some critical feature importance insight for cyber threat modeling. The paper also analyses the architectural implications, limitations of class-imbalanced datasets, and discusses future directions for deep reinforcement learning-based cyber defense systems.

KEYWORDS: Cyber Resilience, Intrusion Detection Systems, Machine Learning, Predictive Security, Attack Volume Forecasting, Self-Learning Architecture, Network Security Analytics, ROC-AUC, RMSE, Feature Engineering.

I. INTRODUCTION

Digital transformation and the adoption of new technologies such as cloud and the Internet Things (IoT) devices were adopted, resulting in a sprawling new enterprise network attack surface. Traditional cybersecurity methods, utilizing signature-based detection and static rule sets, are proving insufficient against advanced persistent threats (APTs), zero-day exploits, and polymorphic malware [1]. This demand for intelligent, adaptive, and anticipatory defense systems has triggered a paradigm shift to ML-driven Cybersecurity architectures.

Predictive Cyber Resilience is the ability of an organizations security infrastructure to predict, withstand and evolve against cyber adversity through data-backed decision making Self-learning architectures, or systems that dynamically



evolve their threat model using incoming telemetry, sit at the forefront of this capability. This is achieved through the integration of supervised learning, ensemble methods and time-series analytics for real-time situational awareness with proactive countermeasures [2].

Although much work has been done to explore various types of ML classifiers for intrusion detection, comparatively evaluating multiple learning paradigms in a single predictive framework—where both continuous mapping of incoming attack volume and classification of its spikes are treated as prediction problems—is still under-researched. Additionally, the use of organizational security posture indicators (patch level, MFA coverage, EDR health) as predictive features of potential future attack paths has been under addressed in the literature [3].

This paper fills those gaps with three main contributions: (1) we present a novel multi-feature enterprise security telemetry dataset containing 11,999 time-stepped observations; (2) we compile benchmarking results of ten models across dual-task ML objectives encompassing regression and classification; and (3) we perform extensive correlation analysis to expose relationships between features that impact precision when predicting threats. We present findings that can be used as actionable insights for security architects designing selflearning- defense platforms of the next generation.

II. LITERATURE REVIEW

Fast evolution of machine learning in cybersecurity comes with an emphasis on areas like intrusion detection, attack prediction, ensemble methods and self-learning security architectures. Early works established the basis for detection systems, while modern studies focus on optimizing prediction accuracy and enhancing robustness and adaptiveness of security models. These developments are key to improving cyber homeland security.

A. Intrusion Detection Systems and ML

Early applications of ML in security were Intrusion Detection Systems (IDS). In this context, Denning's original work [4] proposed a statistical anomaly detection model that has led to subsequent ML adaptations. Tsai et al. Previous works such as [5] have reported a comprehensive overview of ML based IDS, which also confirm that ensemble classifiers (with Random Forests in particular) consistently perform better than single-model solutions on the KDD Cup 1999 benchmark. However, class imbalance was still a major challenge, with rare attack classes having poor recall despite very high accuracy overall. Mukkamala et al. popularized Support Vector Machines for network intrusion detection [6], which achieved competitive performance while using lower computational resources comparing to neural networks. In the last few years, gradient boosting frameworks such as xgboost [7] have outperformed other model classes in structured data competitions and security analytics due to their ability to handle heterogeneous feature spaces well and introduce regularization which prevents overfitting.

B. Predictive Threat Intelligence

Predictive security modeling is not only about reactive detection; it's also about anticipatory threat forecasting. Bhuyan et al. [8] introduced a multi-step network anomaly forecasting framework exploiting autoregressive models where these temporal dependencies of the attack telemetry are prominent. They showed that both scan rate and failed authentication events are strong leading indicators of impending volumetric attacks, a finding that is consistent with our correlation analysis (scan_rate to attack_volume_current: $r = 0.74$). Only recently have Sommer and Paxson [9] published a thorough analysis of the challenges in layering ML onto network security, essentially pointing out that the open world nature of network traffic means new attack patterns continually appear limiting any supervised ML models trained on historical data. They presented an argument for anomaly-based models equipped with continuous retraining capabilities, a point at the heart of self-learning architectures. Directly applicable to our classification findings, where precision across the majority of models was <0.50 , are their worries about low precision in high-dimensional feature spaces.

C. Ensemble and Boosting Methods in Security

We have shown that ensemble method has been very useful in cybersecurity applications. Breiman's Random Forest algorithm [10] is naturally resistant to overfitting via bagging and randomizing features, which makes it an ideal choice for high-dimensional security datasets. XGBoost [7] by Chen and Guestrin built upon gradient boosting with additional innovations of second-order Taylor expansion of the loss function, column subsampling, and L1/L2 regularization to produce state-of-the-art performances on a wide range of ML benchmarks. Diro and Chilamkurti [11] investigated distributed deep learning for IoT intrusion detection, showing that federated learning architectures can yield similar detection rates while protecting data privacy, an important factor in enterprise settings faced with regulatory compliance mandates. Their results indicate that self-learning architectures can be distributed rather than end-to-end.



D. Security Posture and Vulnerability Metrics

A novel dimension of cyber threat modeling adds organizational security posture metrics patch levels, MFA coverage, and endpoint detection health — as predictive features. Jacobs et al. [12] have shown that the likelihood of exploiting vulnerabilities can be predicted with high accuracy from CVSS scores and patch deployment lag, making proactive patch management one of the most effective controls. The strong correlation ($r = 0.96$) found between patch_level and mfa_coverage in this dataset supports our hypothesis that security investment tends to be correlated across controls. Frei et al. Another study [13] examined vulnerability disclosure and patching data on a large scale, showing that the vast majority of exploitable vulnerabilities are attacked days after public disclosure. This time pressure highlights the utility of automated, ML powered prioritization systems that recommend patching schedules based on predicted exposure to attack.

E. Self-Learning and Adaptive Defense Architectures

The genesis of self-learning security architectures are many across various paradigms. Darktrace's Enterprise Immune System [14] uses unsupervised Bayesian inference to build behavioral baselines and detect deviations without the need for labeled training data. Although powerful for zero-day threat detection, these approaches do not provide the explicit prediction of attack that is possible with supervised regression models. Mirsky et al. [15] presented KITSUNE, a lightweight online ML-based ensemble autoencoders framework for network anomaly detection. Most other existing machine learning methods do not deliver such high detection rates without incurring a strong computational overhead, making them infeasible to use in resource-constrained environments. The self-updating capability of their architecture—the continuous adaptation to changing network baselines—directly inspired our work that explores the principles behind self-learning designs. Collectively, these works provide a theoretical and empirical foundation around which our dual-objective predictive framework is designed.

III. DATASET AND FEATURE ANALYSIS

The dataset described in this paper includes 11,999 temporally-ordered measurements of enterprise security telemetry, which corresponds to distinct monitoring steps. It is realistic about how measures of effectiveness that look at activity work in an operational security environment where engagement and exploitation are fundamental. We also need this dataset as a reference to identify and improve the security level of real world networks..

A. Feature Descriptions

We include ten features, which are: step_minutes (monitoring time elapsed), patch_level (normalized coverage of patch deployment across assets), mfa_coverage (multi-factor authentication adoption rate);edr_health (endpoint detection and response agent health score), internet_exposed_services (count of services that are available over the internet), new_assets (number of newly enrolled network assets with a fresh observation window), scan_rate (network scanning activity level in the environment), failed_logins (authentication failure count) suspicious_dns (anomalous DNS query flow volume from multiple addresses independently or same service requesting targeted similar queries in limited time frame), attack_volume_current (current observed intensity of attack.) There are two prediction targets: attack_volume_next (a continuous regression target) and attack_spike_next (a binary classification target indicating volumetric spike events). Normalization preprocessing was applied to standardize all numerical features in the dataset. The dataset features an imbalance in the class of spike classification target with attack spike events covering about 9.9% of the total observations. Hence, this unequal distribution poses a significant problem for supervised classification algorithms and thus needs techniques to overcome data imbalance to achieve better models.

B. Feature Distributions

Empirical distributions of two major features are shown in Figure 1 and Figure 2. internet_exposed_services shows a right-skewed distribution that will leave out the tail, implying that while most hosts do have a somewhat exposed profile, there is still a non-negligible minority with greatly higher service exposure. The scan_rate feature is near-Gaussian centered around zero, as expected from the normalized preprocessing.




Fig.1. Distribution of internet_exposed_services: right-skewed with positive outliers indicating elevated attack surface exposure.




Fig.2. Distribution of scan_rate: approximately Gaussian distribution centered at zero post-normalization.

C. Correlation Analysis

In figure 3, the feature-target correlation heatmap is provided. The analysis uncovers some major interdependencies. (DSG) Organizations with mature patch management programs also tend to see high rates of MFA adoption, as the metrics are strongly positively correlated ($r = 0.96$)—this is a reflection of general security program maturity in contrast to independent control deployment. Most important for predictive modeling, scan_rate shows strong positive correlation with attack_volume_current ($r = 0.74$) and moderate level of correlation with attack_volume_next ($r = 0.61$) confirming its utility as a leading threat indicator. The features failed_logins and suspicious_dns are mutually correlated ($r = 0.33$) and positively correlate with attack_volume_current (both $r = 0.51, 0.50$), reflecting the pre-attack co-occurrence of authentication and reconnaissance activity.




Fig.3. Correlation heatmap of all features and target variables. Strong interdependencies are observed between security posture metrics and between scan-rate-based features and attack volume targets.

D. Attack Volume Temporal Dynamics

Figure 4 presents the time series of attack_volume_next for 11,999 time steps. It has high-frequency oscillations punctuated with spikes, which could be a signature of bursty attacks overlaid on top of normal traffic. It does not appear to have a long-term trend; it is stationary, so ordinary regression models would be suitable without requiring differencing or detrending.




Fig.4. Attack volume over time across 11,999 monitoring steps, showing high-frequency volatility with recurrent spike events indicative of coordinated attack bursts



IV. METHODOLOGY

Here we detail the experimental design, preprocessing pipeline and model configuration and evaluation metrics used in our comparative benchmarking study. Here I describe how the data was pre-processed, which models were used to analyze it and what metric was used to quantify performance. They are pertinent for access to understanding of the methodology and that output is consistent, systematic and reliable throughout the study.

A. Preprocessing and Train-Test Split

Each numeric features was standardized using zero-mean unit-variance normalization (StandardScaler) on the raw dataset independently. Neural network architectures were implemented under constraints to maintain the time-ordering of observations, protecting from future states leaking into previous training samples. An 80:20 stratified train-test split was applied, with stratification on the binary spike label in order to maintain class distributions across different portions of data. This resulted in the training set containing 9,599 observations and the test set containing 2,400 observations.

B. Classification Framework

Here we detail the experimental design, preprocessing pipeline and model configuration and evaluation metrics used in our comparative benchmarking study. Here I describe how the data was pre-processed, which models were used to analyze it and what metric was used to quantify performance. They are pertinent for access to understanding of the methodology and that output is consistent, systematic and reliable throughout the study.

Due to the high class imbalance (positive class prevalence $\approx 9.9\%$), ROC-AUC was chosen as the main metric for evaluation, since it does not depend on the class distribution. F1-Score, precision, and recall are provided as secondary metrics to describe the precision-recall trade-off. Because a naive predictor for the majority class achieves $\sim 90\%$ accuracy, accuracy is an unreliable measure in this regime.

C. Regression Framework

Attack volume prediction (attack_volume_next) was defined as a supervised regression problem. We tested five algorithms: (1) Linear Regression (ordinary least squares); (2) Random Forest Regressor with 100 estimators and max depth at 10; (3) Gradient Boosting Regressor using 100 estimators and a learning rate of 0.1; (4) XGBoost Regressor with similar parameters as the previous model summarized together as XGBRegressor(n_estimators = 100, learning_rate = 0.1, max_depth=6), and (5) Support Vector Regression SVR [epsilon-sensitive loss function] using RBF kernel; The evaluation metrics are Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and coefficient of determination (R^2).

D. Evaluation Metrics

For classification, the evaluation suite includes: Accuracy = $(TP + TN)/(TP + TN + FP + FN)$; Precision = $TP/(TP + FP)$; Recall/ Sensitivity = $TP/(TP + FN)$; F1-Score = $2 \cdot (\text{Precision} \cdot \text{Recall}) / (\text{Precision} + \text{Recall})$ and Area Under the Receiver Operating Characteristic Curve (ROC-AUC), calculated using trapezoidal rule across full probability score range.

$MAE = (1/n) \cdot \sum |y_i - \hat{y}_i|$, $RMSE = \sqrt{(1/n) \cdot \sum (y_i - \hat{y}_i)^2}$, and $R^2 = 1 - (\sum (y_i - \hat{y}_i)^2) / (\sum (y_i - \bar{y})^2)$ were calculated for regression on held-out test observations. Due to RMSE's property of punishing larger residuals more than smaller ones, it is a better suited metric for attack volume prediction where large underestimation carries much greater risk operationally (eg some forms of crime).

V. EXPERIMENTAL RESULTS

A. Classification Results

The classification performance of the five models on the held-out test set is summarized in Table I. Though accuracy is relatively similar between the classifiers in a range of 0.889 and 0.901, differences emerge on the prediction of minority classes. This means that just looking at accuracy is not enough since the metric hides relevant differences in the categories with low representation for precision, recall, and F1-scores.



TABLE I Classification Model Performance Comparison

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	0.901	0.500	0.037	0.069	0.822
Random Forest	0.900	0.300	0.010	0.020	0.779
Gradient Boosting	0.896	0.150	0.010	0.019	0.810
SVC	0.901	0.000	0.000	0.000	0.625
KNN Classifier	0.889	0.253	0.064	0.102	0.645

Logistic Regression has the best capacity of making correct decisions irrespective of decision threshold (ROC-AUC= 0.822) Its F1-Score of 0.069, while low in an absolute sense is a better F1 achieved by any classifier tested so far (with higher recall of course: Random Forest only managed 0.010 and Gradient Boosting too only 0.010). The SVC returns F1, precision and recall of zero — meaning it fails to identify any positive-class instance at all — a common result when support vector machine decision boundaries are overwhelmed by imbalance without probability calibration.

While the ROC-AUC of KNN is lower than all other classifiers (0.645), it showcases the best score in terms of F1-Score and recall (0.102, 0.064 respectively) emphasizing that local instance-based classifiers can perform better with respect to these metrics at a given threshold while overall discriminative power may be weaker for this class of fines. This trade-off is contextually important: in security applications high recall (avoiding false negatives) might take precedence over precision to avoid overlooked detections. Figure 5 visualizes F1-Score comparisons across classifiers, highlighting the dominance of KNN and Logistic Regression for minority-class detection.




Fig.5. Classification Model F1-Score Comparison. KNN achieves the highest F1-Score (0.102), followed by Logistic Regression (0.069), reflecting superior minority-class detection capability under severe class imbalance.

Figure 6 presents the ROC curves for all classification models, confirming Logistic Regression's superior AUC performance across the full operating range.




Fig.6. ROC Curves for all classification models. Logistic Regression (AUC=0.82) and Gradient Boosting (AUC=0.81) demonstrate the strongest discriminative power, while SVC (AUC=0.63) approaches the random baseline.

B. Regression Results

Table II provides a summary of regression performance for predicting attack volumes in the same manner, yielding a consistent pattern among models. In terms of predictive accuracy, Mean Square Error values are well clustered from 0.422 to 0.439 and range only slightly across the dataset. The narrow clustering implies that all five models yield similar results, i.e., no model shows a marked advantage over the others in terms of reducing error.

TABLE II Regression Model Performance Comparison

Model	MAE	RMSE	R ²
Linear Regression	0.335	0.422	0.482
Random Forest Regressor	0.343	0.433	0.454
Gradient Boosting Regressor	0.334	0.422	0.482
XGBoost Regressor	0.349	0.439	0.439
SVR	0.333	0.423	0.479

SVR hits its stride with the lowest MAE of 0.333 and for RMSE both Gradient Boosting Regressor and Linear Regression tie, hitting values of 0.422. Linear Regression and Gradient Boosting Regressor achieve the highest R² value (0.482): this means that the feature set explains around 48% of variance in attack_volume_next —a lot considering cyber attack patterns are stochastic by nature.

The uniformity of performance among different model families — from linear models to ensemble boosting to kernel methods — implies that the predictive ceiling for this feature set may have been reached under standard supervised learning paradigms. This insight motivates the examination of temporal architectures (i.e., LSTM, Transformer) and reinforcement learning methods that make use of sequential dependencies not captured by i.i.d. data assumptions.




Fig.7. Regression Model RMSE Comparison. All models converge within a narrow RMSE band (0.422–0.439), suggesting saturation of predictive capability under standard supervised learning with the current feature set.

VI. DISCUSSION

A. Implications of Class Imbalance

The extreme class imbalance observed in the target `attack_spike_next` (positive prevalence ~9.9%) essentially establishes a boundary surrounding the performance space of classification models. High accuracy scores (0.889–0.901) can be easily obtained by predicting only the major class, which makes accuracy a poor primary metric. The very small recall values recorded for Random Forest (0.010) and Gradient Boosting (0.010) highlight the tendency of these powerful ensemble methods to revert into majority-class predictors when trained without explicit instruction on how to handle imbalance in the data.

The design of a new release of this structure would include the weighting of classes (`class_weight='balanced'`), SMOTE (Synthetic Minority Oversampling Technique) and threshold calibration after training to nudge the decision curve toward greater recall. For operational security contexts, the cost asymmetry between false negatives (missed attacks) and false positives (false alarms) usually motivates highly aggressive recall optimization that can incur high precision costs.

B. Feature Importance and Security Posture Dynamics

Correlation analysis gives important insights regarding relative predictive value of features. The near perfect correlation between `patch_level` and `mfa_coverage` ($r = 0.96$) suggests multicollinearity that may act to suppress the independent contribution of each variable in linear models. Dimensionality reduction algorithms or regularized regression (Lasso, Elastic Net) might disentangle these effects and enhance regression R^2 .

The high correlation of prediction with the current scan rate and attack volume (current: $r = 0.74$, next: $r = 0.61$) further validates that reconnaissance activity — scanning a network — is an accurate indicator of volumetric attacks. This has immediate operational impact; security operations centers (SOCs) must prioritize monitoring of real-time scan rates, and alerts based on anomalous scan rates crossing thresholds can act as useful precursors to more formal escalation of defensive postures.

C. Convergence in Regression Performance

Among the most striking results that deserves architectural reflection is the convergence of regression RMSE across a wide range of model architectures (0.422–0.439). Linear Regression matches the accuracy of ensemble methods like Random Forest and XGBoost, implying that patterns are primarily linear in nature for this current feature set—or that nonlinearities exist, but randomness hides them from view. This convergence suggests two architectural trajectories: (1)



feature engineering to expose non-linear relationships by adding interaction terms and polynomial features; and, (2) temporal sequence modeling using recurrent architectures (LSTM, GRU) that can take advantage of the time-series structure in the 11,999-step dataset.

D. Self-Learning Architecture Design Principles

For self-learning cyber-defense architectures, we propose the following design principles on the basis of our experimental conclusions: First, dual-mode modeling (regression and classification) complements situational awareness: volume projection supports future-planning; on the other hand the classification of spikes driving alerting logic. Second, focusing on the operational role of models: Logistic Regression is suited to applications where threshold values are flexible, as its superior ROC-AUC shows, while the K-Nearest Neighbor algorithm (KNN) is suited to high-sensitivity alarms due to better recall. Third, because cyber threat landscapes are not stationary continuous training pipelines are necessary; if models are trained on historical data, they degrade as adversarial techniques evolve. Fourth, feature engineering—especially when organizational security posture metrics are included with network performance measures—extends the range of predictive options for new techniques beyond network telemetry-only methods.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we conducted a systematic benchmarking study of algorithmic defense self-learning architectures for predictive cyber resilience applied to an 11,999-observation enterprise security telemetry dataset. Using our dual-objective framework, we found that Logistic Regression provides the best discriminative performance for predicting attack spikes (ROC-AUC = 0.822), but Gradient Boosting Regressor and Linear Regression yield near-optimal forecasting of attack volume (RMSE = 0.422, R^2 = 0.482). The correlation analysis confirmed the critical predictive role of `scan_rate` and established the co-maturity relationship between `patch_level` and `mfa_coverage`.

There are several major limitations that restrict the present findings. Class imbalance in spike classification was neither mitigated by resampling nor addressed with cost-sensitive learning, potentially limiting recall that could be achieved. As regression models converge around $R^2 \approx 0.48$, this indicates a ceiling per feature which must be broken by innovation at the architectural level — specifically temporal modeling. Furthermore, the synthetic nature of this dataset, while facilitating controlled experimentation, may fail to capture the distributional complexity inherent in actual enterprise environments that are subject to targeted attacks by advanced and knowledgeable adversaries.

Ongoing and future work in the area includes: (1) combining Long Short-Term Memory (LSTM) networks with Transformer architectures to take advantage of temporal dependencies in attack sequences; (2) using deep reinforcement learning (DRL), to adaptively optimize response policies; these may include federated architectures, allowing multiple organizations to share threat intelligence without sacrificing data sovereignty; (3) explanation frameworks (such as SHAP or LIME) enabling interpretable outputs, aiding SOC analysts' information processing ability, and finally, (5) evaluation of adversarial robustness issues — assessment of how resilient a model is against evasion attacks designed to target ML-based detectors. The end vision is of an autonomously-configured, self-learning cyber defense system that functions like a digital immune system able to anticipate, adapt to, and neutralize the changing threat landscape with minimal human involvement.

REFERENCES

- [1] A. Milenkoski, M. Vieira, S. Kounev, A. Avritzer, and B. D. Payne, "Evaluating computer intrusion detection systems: A survey of common practices," *ACM Computing Surveys*, vol. 48, no. 1, pp. 1–41, 2015.
- [2] R. Mitchell and I. R. Chen, "A survey of intrusion detection techniques for cyber-physical systems," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–29, 2014.
- [3] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1153–1176, 2016.
- [4] D. E. Denning, "An intrusion-detection model," *IEEE Transactions on Software Engineering*, vol. SE-13, no. 2, pp. 222–232, 1987.
- [5] C.-F. Tsai, Y.-F. Hsu, C.-Y. Lin, and W.-Y. Lin, "Intrusion detection by machine learning: A review," *Expert Systems with Applications*, vol. 36, no. 10, pp. 11994–12000, 2009.
- [6] S. Mukkamala, G. Janoski, and A. Sung, "Intrusion detection using neural networks and support vector machines," in *Proc. IEEE International Joint Conference on Neural Networks*, 2002, pp. 1702–1707.
- [7] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.



- [8] M. H. Bhuyan, D. K. Bhattacharyya, and J. K. Kalita, "Network anomaly detection: Methods, systems and tools," IEEE Communications Surveys & Tutorials, vol. 16, no. 1, pp. 303–336, 2014.
- [9] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in Proc. 31st IEEE Symposium on Security and Privacy, 2010, pp. 305–316.
- [10] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
- [11] A. A. Diro and N. Chilamkurti, "Distributed attack detection scheme using deep learning approach for Internet of Things," Future Generation Computer Systems, vol. 82, pp. 761–768, 2018.
- [12] J. Jacobs, S. Romanosky, B. Edwards, I. Adjerid, and M. Roytman, "Exploit prediction scoring system (EPSS)," arXiv preprint arXiv:1908.04856, 2019.
- [13] S. Frei, T. Duebendorfer, G. Ollmann, and M. May, "Understanding the web browser threat: Examination of vulnerable online web browser populations and the insecurity iceberg," ETH Zurich, Tech. Rep., 2008.
- [14] D. Bhatt, P. Patel, H. Talsania, J. Patel, R. Vaghela, S. Pandya, K. Modi, and H. Ghayvat, "CNN variants for computer vision: History, architecture, application, challenges and future scope," Electronics, vol. 10, no. 20, p. 2470, 2021.
- [15] Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai, "Kitsune: An ensemble of autoencoders for online network intrusion detection," in Proc. Network and Distributed System Security Symposium (NDSS), 2018.
- [16] Rasul, I., Tohfa, N. A., Rahman, M., Hossain, I., Zareen, S., & Shakhawat, M. (2023). Quantum Machine Learning for Early Disease Diagnosis: A Systematic Review and Public Health Innovation Perspective. World Journal of Advanced Research and Reviews, 19(01), 1668-1674.
- [17] Tohfa, N. A., Hossain, I., Zareen, S., Rasul, I., Hossen, M. S., & Rahman, M. (2021). Adversarial Cognition Machine Learning at the Frontlines of Cyber Warfare. World Journal of Advanced Research and Reviews, 12(02), 722-729.
- [18] Alam, A., Gazi, M. S., Abdullah, S. M., Tasnim, M., Himeluzzaman, M., Nabil, M. A., ... & Akter, S. (2021). Quantum-resilient federated intrusion detection: A hybrid quantum classical framework for safeguarding US critical infrastructure in the post-quantum era. International Journal of Advances in Signal and Image Sciences, 7(1), 57-72.
- [19] Biswas, B., Chaudhury, T. H., & Mohammad, S. N. A Fast and Scalable Recommender System using Collaborative Filtering Technique in Python Scikit-learn.
- [20] Hossain, M. S., Biswas, B., & Rahaman, M. M. (2023). THE ROLE OF ARTIFICIAL INTELLIGENCE IN ENHANCING CYBERSECURITY DEFENSE MECHANISMS. International Journal Of Engineering Technology Research & Management (IJETRM), 7(07), 157-166.