

Investigating the Application of Artificial Intelligence (AI) and Machine Learning (ML) Techniques to Enhance Cybersecurity for Internet of Things (IoT) Devices, Prevent Data Breaches, and Safeguard User Privacy

Simon Atadoga
University of Illinois Urbana
Champaign,
Gies college of business -
Accountancy Department

Timothy Oyebola Ige
University of Denver,
Department of Health Informatics
- Digital Health

Rona Oneshiorona Sado
School: Kennesaw State
University
MSc Computer Science

Abimbola Oludayo Ojenike
University of Greenwich, London
Uk, Department of Forensics and
cybersecurity

Confidence Adimchi
Chinonyere
Abia State Polytechnic
Department of Accountancy

Emedem Sandra
Ebubechukwu
Nnamdi Azikwe University
Dpt: Computer science

Victor Oyiboka
University of Texas at Dallas
Physics

ABSTRACT

The rapid expansion of the Internet of Things (IoT) ecosystem has provided a tremendous attack surface, and therefore, IoT devices are highly vulnerable to advanced cyberattacks, data breaches, and privacy invasions. Rule-based intrusion detection systems are mostly ineffective in dealing with high-dimensional and heterogeneous traffic streams that IoT environments produce. To fill in these gaps, this research examines the systematic use of Artificial Intelligence (AI) and Machine Learning (ML) methods towards IoT security augmentation, malicious activity detection, blocking of data leakage, and safeguarding of user privacy. A strict methodology, quantitative experimental approach was adopted, leveraging the Australian Centre for Cyber Security's TON_IoT20 dataset of actual network traffic, attack behaviours (i.e., DDoS, data injection, password-based intrusions), and normal run log data from various IoT devices such as smart plugs, cameras, and thermostats. Data preprocessing steps involved removal of duplicates, handling of missing values by imputation, feature encoding, and scaling, followed by a 70/15/15 stratified split for training, validation, and test. Three standard ML models, Random Forest (RF), Extreme Gradient Boosting (XGBoost), and a Deep Neural Network (DNN) were used in Python under a controlled Ubuntu environment and trained on the pre-processed data.

Model performance was measured by accuracy, precision, recall, F1-score, and ROC-AUC values, with further analysis by means of confusion matrices and McNemar's significance testing. The results indicate that XGBoost performed better, with 98.9% accuracy, 98.6% precision, 99.0% recall, an F1-score of 98.8%, and an ROC-AUC value of 0.996, with very low values of false positives and false negatives. Statistical testing established that the improvement of XGBoost relative

to RF and DNN was significant ($p < 0.05$). In addition, XGBoost provided competitive training time and the quickest inference time, indicating its real-time suitability for IoT intrusion detection applications.

All of these results underscore the promise of incorporating AI/ML solutions based on XGBoost in IoT security platforms to improve active threat detection, reduce false alarms, and offer improved privacy protection controls. The research provides an experimentally validated reference model towards further studies and real-world applications of AI-driven intrusion detection systems in real-time IoT environments.

Keywords

Artificial Intelligence (AI); Machine Learning (ML); Internet of Things (IoT); Cybersecurity; Intrusion Detection System (IDS); User privacy, smart home.

1. INTRODUCTION

The proliferation of the Internet of Things (IoT) has revolutionized business and everyday life by enabling end-to-end connectivity between billions of things. From smart homes and wearable health trackers to industrial sensors and autonomous vehicles, IoT networks continuously generate high-volume, heterogeneous, and dynamic data streams. Yet, it also comes with greater connectivity to enormously larger attack surfaces exploited by hackers to breach confidentiality, integrity, and availability of information (Al-Garadi et al., 2023; Khan et al., 2024). Recent studies indicate an oncoming trend of IoT-based cyberattacks, such as distributed denial-of-service (DDoS) attacks, data injection, and credential theft, which have resulted in large-scale data breaches and privacy infringements (Ahmed et al., 2023).

Legacy rule-based and signature-based intrusion detection

systems (IDS) typically do not look after the IoT network traffic dynamism because it changes very quickly. They lack scope, flexibility, and the capability to identify zero-day attacks or minor variations (Chen et al., 2022). This has prompted scholars and practitioners in the area to investigate the application of Artificial Intelligence (AI) and Machine Learning (ML) methods to predictive threat modelling and dynamic intrusion detection in IoT scenarios (Singh & Rajesh, 2023; Li et al., 2024).

Machine learning algorithms like Random Forests (RF), Gradient Boosted Trees (like XGBoost), and Deep Neural Networks (DNNs) have demonstrated encouraging results in identifying sophisticated attack behaviours in high-dimensional network traffic data (Tian et al., 2023). XGBoost, in fact, has proven to be an appropriate choice for cybersecurity tasks with its imbalanced dataset tolerance, fast training, and enhanced classification accuracy on tabular data (Xu et al., 2023). Deep learning models, though effective for unstructured data spaces, fare poorly on tabular IoT data unless it is heavily tuned or combined with feature engineering techniques (Shah et al., 2024). Notwithstanding these developments, several challenges exist. IoT devices are low-resource devices and hence models need to be of lower accuracy but not computational complexity (Zhang et al., 2024). Moreover, privacy laws and principles also require that any model deployed reveals minimal information and runs in a strongly controlled environment. Thus, a comprehensive investigation of various AI/ML methods on real-world publicly available IoT traffic data sets is essential to determine optimal solutions with a trade-off between detection accuracy, computational cost, and privacy implications. This study overcomes these hurdles by adopting a serious experimental approach with the TON_IoT20 dataset, a highly validated benchmark for IoT security research. Through a comparison of RF, XGBoost, and DNN models, this study will prove how AI/ML can be utilized to advance IoT cybersecurity, decrease false positives, and supply statistically validated intrusion detection power boosts.

Objectives of the Study

The broad objective of this study is to examine and assess the use of Artificial Intelligence (AI) and Machine Learning (ML) methods in developing improved cybersecurity for Internet of Things (IoT) devices to try to block data breaches and maintain user privacy.

- For the fulfillment of this objective, the precise objectives are as follows:
- To recognize and study current cybersecurity threats and emerging risks in IoT systems and to bring forth the vulnerabilities of traditional security systems against these risks.
- To train and deploy chosen AI/ML techniques (Random Forest, XGBoost, and Deep Neural Networks) in a real-world dataset (TON_IoT20) for intrusion detection and anomaly detection in IoT network traffic.
- To compare and contrast the performance of the deployed models based on shared metrics like accuracy, precision, recall, F1-score, and ROC-AUC to measure their effectiveness in identifying malicious activities.
- To assess the real-world implications of the inclusion of AI/ML-driven intrusion detection in IoT security systems with respect to minimizing false alarms, boosting

detection rates, and preserving user anonymity.

- To provide adaptive and responsive mitigation threat recommendations based on results obtained from model estimation and significance testing.

Research Questions

To direct this research on the use of Artificial Intelligence (AI) and Machine Learning (ML) methods to enhance cybersecurity in IoT devices, the research aims to enlighten on the following questions:

What are some of the new cybersecurity threats and vulnerabilities intrinsic to IoT ecosystems, and why do conventional security measures fail against them?

How well do AI/ML models including Random Forest, XGBoost, and Deep Neural Networks identify intrusions and anomalies in IoT network traffic on real-world datasets?

Which of these AI/ML models has better detection accuracy, precision, recall, F1-score, and ROC-AUC?

Can incorporation of AI/ML-based intrusion detection in IoT security models lower false positives drastically and improve detection of zero-day or dynamically changing cyber threats?

What are the practical applications and suggested mechanisms for implementing adaptive AI/ML-based security features in large-scale IoT systems to protect user privacy and avoid data breaches?

2. LITERATURE REVIEW

The extensive deployment of IoT technology has driven a lot of effort into intrusion detection systems capable of addressing the distinctive challenges of IoT networks. In contrast with other traditional enterprise networks, IoT networks are made up of heterogeneous devices, low-power processors, and changing topologies, all which require effective and responsive security (Pahlavan et al., 2023). Researchers have also mentioned that current signature-based intrusion detection systems are inadequate since they are static and do not help in the detection of zero-day threats (Velmurugan et al., 2024). This has driven researchers into the investigation of the use of AI and ML techniques that learn sophisticated patterns and can generalize to unseen threats.

Current research establishes the potential of ensemble learning methods, especially tree-based methods, in IoT intrusion detection. For instance, Abubakar et al. (2023) designed a gradient-boosted tree classifier to counter imbalanced network traffic data and reported better detection accuracy than traditional classifiers. In the same vein, Hussain et al. (2022) claimed that boosting and bagging methods outshine single-learner models by leveraging feature importance and evading overfitting in diverse IoT data.

Concurrently, deep learning has also emerged into the limelight, with Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks being explored to detect sequential and temporal patterns in IoT traffic. Alqahtani et al. (2023) designed a hybrid CNN-LSTM model to identify anomalies in industrial IoT applications and attained better recall scores. Though, they added that deep learning models tend to have a high computational requirement and large labelled training datasets, which are not always accessible in resource-constrained IoT deployments. Feature engineering and preprocessing are also significant factors that influence detection performance. Vekaria et al. (2023) indicate that the use of one-hot encoding and min-max normalization improved significantly model convergence and stability on big IoT

benchmark datasets. Experiments in their research also highlighted the importance of the use of stratified data splitting for the sake of maintaining class balance for preventing biased model assessment.

Ethical and privacy concerns are increasingly becoming the focal point of research in cybersecurity. Rajput et al. (2024) emphasized that intrusion detection mechanisms have to be designed with privacy-preserving methods, e.g., federated learning, such that sensitive information never gets transferred out of local devices but still generates global model updates. These methods are accompanied by increasing regulatory pressures like GDPR and increase user trust. Furthermore, comparative model benchmarking with respect to benchmark datasets is also a core area of research. Sahu et al. (2023) provided an observation that application of such datasets like BoT-IoT and TON_IoT20 facilitates reproducibility and fair benchmarking of models in actual traffic conditions. They also pointed out the requirement of statistical tests like McNemar's test to ensure that results-oriented improvement in performance is statistically significant rather than being a chance occurrence.

Overall, the literature indicates that tremendous progress has been achieved in using AI and ML for IoT intrusion detection but that there are as yet inadequacies in balancing good detection performance, computational efficiency, and privacy preservation. Such inadequacies necessitate more experimental researches such as this work over alternative ML techniques with real-world datasets and statistically verifying performance enhancement.

The distinctive nature of IoT ecosystems categorically high device diversity, intense data exchange with high frequency, and low computational power have rendered IoT cybersecurity extremely dynamic research. Researchers have, in the past few years, suggested several paradigms of Artificial Intelligence (AI) and Machine Learning (ML) to overcome the constraints of conventional intrusion detection systems (IDS). In contrast to static, signature-based methodologies, AI/ML methodologies are highly dynamic, learning from changing traffic patterns and evolving to accommodate new attack vectors (Okafor et al., 2023).

A. AI and ML Evolution in IoT Security

Early efforts to secure IoT networks depended on lightweight rule sets and heuristic approaches, which rapidly proved themselves inadequate against high-level, zero-day attacks and polymorphic malware. Modern studies have focused more on ML-based models, with tree-ensemble methods becoming increasingly popular due to their stability on structured data. For example, Eze et al. (2023) showed how ensemble trees like Extra Trees and LightGBM were more effective than conventional classifiers in botnet detection for intelligent homes. Likewise, Jamal et al. (2022) contrasted different boosting algorithms and pointed out their stability in asymmetric traffic scenarios.

B. Deep Learning Methods

Parallely, deep learning models have also been explored extensively for discovering intricate spatial and temporal patterns in IoT traffic. Zhao et al. (2024) introduced a Transformer-based time-series intrusion detection system with significant recall improvements on industrial IoT gateways. In addition, Sujatha and Thomas (2023) used autoencoders and gated recurrent units (GRUs) in combination for identifying low-rate DDoS attacks in resource-constrained devices,

although their work also mentioned the added computational requirement of such models.

C. Data Preprocessing and Feature Optimization

Feature engineering is the most significant area to enhance model detection performance. In a recent study, Miah et al. (2023) found that using correlation-based feature selection and scaling significantly enhanced F1-scores in ML models for a number of IoT datasets. Similarly, Hossain and Pervez (2023) suggested that normalization processes and dimensionality reduction techniques like PCA might render models more efficient while keeping accuracy intact.

D. Ethical and Privacy Concerns

Beyond detection performance, the literature highlights preserving user privacy. Ghosh et al. (2024) investigated federated learning and differential privacy frameworks and demonstrated local training at the edge devices in IoT minimizes data leakage risk. Their findings support larger industry action calling for privacy-concerned AI solutions in alignment with contemporary regulatory regimes such as GDPR and ISO/IEC 27400:2022.

E. Benchmark Datasets and Statistical Validation

Reproducibility is also an important issue in IoT cybersecurity research. Authors more commonly employ public datasets like TON_IoT20, BoT-IoT, and CICIoT2023 to obtain comparable results (Barua et al., 2023). Reliable statistical practices also become popular; Adebayo et al. (2023) gave credit to the use of McNemar's test and paired t-tests in order to guarantee observed improvements in model performance are statistically, not randomly, significant.

INTERNET OF THINGS (IoT)

The Internet of Things (IoT) is a network of physical devices e.g., sensors, smart home appliances, vehicles, and factory equipment connected over the internet and capable of collecting, sharing, and processing data with minimal or no human intervention. Traditional networks maintain the physical and digital worlds independently, but IoT systems converge these two worlds, enabling automation, analysis, and real-time decision making in various domains such as healthcare, manufacturing, transportation, and intelligent living (Dasgupta et al., 2022).

IoT systems are generally made up of three basic layers: the perception layer (device and sensor that collect data), the network layer (data transmission protocols for communication), and the application layer (services that consume and process data) (Roy et al., 2023). The layered structure is easy to scale but leaves it weak in several points. For example, perception-layer low-power devices are not generally engineered with strong cryptography capability, and transmission-layer network protocols are not designed to offer end-to-end security (Lin & Yu, 2022).

Development of IoT has expanded manifolds. Based on latest industry trends, the number of internet-connected IoT devices around the world reached over 15 billion in 2023 and is set to cross over 29 billion by 2030 (Feng et al., 2024). This boom has fuelled innovation in domains like predictive maintenance, autonomous systems, and smart agriculture. It also exponentially increases the attack surface of hackers. Threats like Mirai botnets, ransomware attacks against IoT gateways, and unauthorized data scraping have made it imperative to have

robust and intelligent security practices (Wu et al., 2024).

Data management-wise, IoT creates enormous, heterogeneous data streams of high velocity and volatility. Such features make the conventional analytics pipelines difficult and invoke the implementation of edge computing and AI-based techniques in performing real-time threat detection and decision-making on resource-constrained devices (Huang et al., 2023). That is why developers are turning more attention to creating light AI/ML models that have been tailored to fit IoT environments where computational efficiency has to be sacrificed at the expense of robust predictive performance. Briefly, IoT is revolutionizing the manner systems communicate and exchange information with each other, yet the nature of IoT openness, heterogeneity, and sheer scale also renders security an overarching problem. Understanding the underlying architecture and operations of IoT is crucial prior to designing intrusion detection systems that can stop the constantly evolving threats.

Development of Artificial Intelligence and Machine Learning in Cyber Security

The surge in the level and rate of cyberattacks over the past several years put into prominence the vulnerabilities of conventional security controls that rely mainly on static rules and human-defined signatures. As attacks on the Internet have dynamically changed from polymorphic malware and zero-day attacks to advanced phishing attacks, organizations are seeking Artificial Intelligence (AI) and Machine Learning (ML) to offer more adaptive and predictive defences (Haque et al., 2023).

A. From Static Defences to Adaptive Intelligence

Legacy intrusion detection and prevention systems struggle to deal with the volume and diversity of today's network traffic. Differing from rule-based solutions, AI and ML enable ongoing learning from historical and real-time information, enabling security systems to predict and react to newly appearing patterns of attacks. Nandhini et al. (2024) also agree that this move toward data-driven intelligence revolutionized cybersecurity from a reactive to a proactive and predictive practice.

B. AI and ML Operations in Threat Detection

AI-based solutions for cybersecurity use algorithms that detect patterns, correlations, and deviations with accuracy. For instance, unsupervised ML algorithms are capable of detecting network deviation without knowing the attack signature beforehand, essentially getting better at insider threat detection and new attacks (Rashid et al., 2022). On the contrary, models of supervised learning trained from labelled datasets of attacks can easily label traffic as benign or malicious with great precision, like shown by Liu et al. (2023) in their research on advanced ensemble approaches for network anomaly detection.

C. Integration with Big Data and IoT

In the wake of developments in IoT and cloud infrastructures, cybersecurity systems today process higher amounts of data compared to ever before. AI and ML architectures are best placed to manage such high-rate streams, deriving actionable intelligence in real time. Gupta et al. (2023) added that integration of ML with big-data analytics platforms allows detection systems to scale across distributed environments, correlating events across millions of devices at very low latencies.

The inclusion of AI also makes autonomous response strategies

possible. Rather than simply notify administrators, AI-based security systems can, on their own, quarantine suspicious computers, block malicious IP addresses, or initiate multi-factor authentication prompts. Mehta et al. (2023) explained that the integration of AI with orchestration tools has seen "self-healing" networks, which respond to dynamic threats independently, with a considerable lag in responses.

E. Challenges and Current Research

Nonetheless, while these benefits exist, AI-based cybersecurity also faces challenges. Model explainability, attacks on ML models using adversarial examples, and the necessity of high-quality labeled datasets are major hurdles (Tian et al., 2024). Explainable AI (XAI) techniques and model adversarial robust training are being researched by scientists to overcome these challenges and ensure that AI-based defences are reliable as well as trustworthy. In brief, AI and ML in cybersecurity is a shift from human, signature-based protection to intelligent, adaptive, and scalable defence mechanisms. This shift is critically important in IoT networks, where the sheer number of devices and attack surfaces demand learning-based, autonomous ways.

New Cybersecurity Threats in IoT Networks

The rapid growth of Internet of Things (IoT) networks, ranging from smart homes, industrial automation, healthcare, and transport, has established a massive, interconnected ecosystem with unprecedented attack surfaces. As opposed to conventional IT networks, IoT networks are defined by resource-limited devices, heterogenous protocols, and even minimal security arrangements, making them an attractive target for cyber threats (Singh et al., 2024).

A. Botnet-Driven Distributed Attacks

One of the greatest emerging threats is that of IoT-based botnets taking advantage of weak device authentication and outdated firmware. The Mirai botnet attack in 2016 was merely an early indication; newer ones like Mozi and Katana are more sophisticated and modular, able to perform massive Distributed Denial of Service (DDoS) attacks with little to no detection (Sharma & Bhushan, 2023).

These botnets take advantage of default credentials and weak Telnet/SSH services of devices such as IP cameras and routers with catastrophic service disruption.

B. Data Poisoning and Adversarial Manipulation

As machine learning models are being used more in IoT for use cases such as anomaly detection and predictive maintenance, attackers have begun targeting the training pipelines themselves. Data poisoning attacks bring in malicious data, with precise creation, into training sets, misleading models to misclassify traffic or not detect real attacks (Chen et al., 2023). Adversarial examples those perceptually indistinguishable perturbations yet important for ML models can make an intrusion detection system mislabel malicious traffic as benign.

C. Edge Device Vulnerability and Side-Channel Attacks

With the processing moving to the edge in IoT architectures, processing unit vulnerabilities on local premises have been targeted by hackers. Side-channel attacks like power analysis and electromagnetic side-channel leakage are being witnessed on IoT gateways and microcontrollers increasingly, compromising sensitive cryptographic keys or firmware

information (Almeida et al., 2022). Such attacks evade conventional network-level protections and infiltrate directly at the expense of confidentiality of data in edge environments.

D. Supply-Chain and Firmware Tampering

Variety of IoT hardware providers and sophisticated supply chains have introduced new threats in the form of preinstalled backdoors and tampered firmware updates. Li et al. (2023) showed that vulnerable over-the-air (OTA) update mechanisms enable attackers to inject malicious firmware without invoking integrity checks. These threats are especially hazardous in industrial IoT environments, where tampered devices can produce disruption to mission-critical infrastructure.

E. Privacy Leaks through Unsecured APIs

In most consumer IoT devices, application programming interfaces (APIs) are poorly secured, permitting attackers to extract private information without actually breaching the device. Han et al. (2024) described how poorly set up authentication in cloud-connected APIs exposed individual health data in smart wearables. This threat class highlights the need to incorporate privacy-by-design principles into IoT software development. In brief, the IoT system threat landscape is evolving very quickly and includes large-scale botnets, adversarial data manipulation, edge exploitation, supply-chain attacks, and API-based privacy leakage. All of these new challenges point toward adaptive AI-driven defence systems and continuous monitoring solutions for securing the next-generation IoT infrastructures.

Anomaly and Intrusion Detection Process

Anomaly and intrusion detection is a key security process for real-time monitoring of network traffic or system activity, abnormal behaviour detection, and alerting possible threats. In Internet of Things (IoT) environments, such processes are of great importance due to the decentralized device nature and limited implementation of conventional endpoint defence mechanisms (Khan et al., 2023).

A. Overview of Detection Approaches

Intrusion detection systems (IDS) are generally divided between signature-based and anomaly-based approaches. Signature-based approaches scan seen behaviour against a database of seen patterns of attacks; effective against known threats, they are not good against zero-day attacks. Anomaly-based systems establish a profile of normal behaviour and ring an alarm on abnormal deviation as a possible intrusion, providing greater immunity to new attacks (Bello et al., 2022).

In IoT settings, hybrid detection models are becoming the norm that leverage both methods, the accuracy of signatures and ease of anomaly detection. The models utilize layered detection, often edge and cloud analytics integrated to effectively process high-speed data streams (Cao et al., 2023).

B. Data Collection and Feature Extraction

It starts with ongoing data collection across sensors, gateways, and network logs. Significant features like packet length, protocol, connection rate, and time gaps are extracted to describe traffic behavior across a multidimensional feature space. Luo et al. (2024) find that in IoT networks, light-weight feature selection processes have the important role of limiting computational expense at the cost of detection fidelity.

C. Model Training and Detection

Machine learning classifiers are subsequently trained against labelled data or unsupervised records of typical behaviour. For the detection of anomalies, the Isolation Forest or clustering algorithms (such as DBSCAN) may typically be employed for anomaly detection, while intrusion detection can be based on supervised classification such as Gradient Boosted Trees or convolutional neural networks. Saha et al. (2023) illustrated how ensemble-based anomaly detectors could achieve high recall rates for identifying low-frequency attack signatures in IoT traffic.

After deployment, the model calculates incoming data in real-time, calculating anomaly scores or class probabilities. Those scores exceeding defined thresholds cause alerts or automated responses.

D. Continuous Learning and Adaptation

New detection mechanisms use online learning to keep up with changing patterns. Incremental training, for instance, enables models to learn their parameters as new data become available, and thus stay immune to concept drift changes in the underlying data distribution over time (Wang et al., 2024). Dynamic adjustment is paramount in IoT settings where device behaviour changes with firmware updates or fresh app installations.

In short, IoT system anomaly and intrusion detection processes comprise data gathering, feature extraction, model training, real-time assessment, and ongoing tuning. Combining AI-based approaches with light-weight construction makes such systems work perfectly under the resource limitations of IoT devices but offer extremely secure defence against known and unknown threats.

Adaptive Threat Mitigation Framework

The evolving nature of contemporary network cyberattacks, especially in Internet of Things (IoT) networks, has demanded the creation of adaptive threat mitigation frameworks. Contrary to conventional static security policies that are based on preconfigured rules, adaptive frameworks observe network conditions in real time, build knowledge from emerging threats, and in real time modify their response mechanisms to reduce damage and provide service continuity (Zhang et al., 2023).

A. Core Principle and Design Elements

An adaptive threat mitigation system integrates detection, decision, and response layers into a closed loop of feedback.

Detection Layer: Continual surveillance of system events, network traffic, and device logs for anomalies. Decision Layer: Leveraging AI/ML approaches to classify threats, prioritize alerts, and compute best mitigation courses.

Response Layer: It executes dynamic responses such as quarantining the infected systems, throttling suspect traffic, or deploying new firewall policies along with providing data to the detection layer so that it becomes more informed (Ghoneim et al., 2023).

All these layers co-operate in such a manner that whenever there are new patterns of attacks emerging, then the system will adjust without any human intervention.

B. Integration of AI and Context Awareness

Modern frameworks encompass context-aware intelligence where response to threats is tailored based on operational requirements and device severity. For example, a hijacked intelligent thermostat can be quarantined immediately, while an industrial sensor with high impact can be mitigated in stepwise fashion to avoid process disturbance (Chatterjee & Malik, 2023). Reinforcement learning and other AI models are also being employed to make decisions more optimal in the long run by taking into account the impact of previous responses.

C. Real-Time and Distributed Mitigation

To manage the size of IoT deployments, adaptive mitigation frameworks are typically shared between edge and cloud environments. Local assessment and early mitigation occur with high speed in edge nodes, while data aggregation is performed in the cloud for facilitating deeper forensic analysis and policy adaptation. Nguyen et al. (2024) observed that this kind of hybrid architecture minimizes latency and avoids bottlenecks but facilitates global situational awareness.

D. Self-Healing and Policy Evolution

One of the most important characteristics of adaptive frameworks is self-healing capacity—the capacity to restore damaged elements and resume normal functioning. This may include automated firmware patching, secure reconfiguration of hardware, or trust re-establishment among actors in a network. Roychowdhury et al. (2022) demonstrated how policy evolution engines integrated within adaptive systems greatly enhanced resilience to emerging threats such as ransomware-as-a-service for IoT.

Intrusion Detection Systems (IDS) for IoT Security

With the Internet of Things (IoT) penetrating more sensitive domains like healthcare, transport, and smart cities, Intrusion Detection Systems (IDS) are now necessary to secure devices and networks from next-gen cyberattacks. Since IoT environments are heterogeneous, resource-limited, and highly distributed in nature, unlike conventional IT environments, IDS solutions must be light, adaptive, and context-aware (Abubakar et al., 2023).

A. IDS role in IoT

An IDS keeps track of network traffic and the activity of devices for identifying malicious events like unauthorized access, denial-of-service attack, or data exfiltration. In the context of IoT, IDS has a valuable role in:

Protecting the limited number of constrained devices with extremely limited inherent security, Notifying lateral movement attacks on connected devices, giving real-time alerting and automated action (Mishra et al., 2022).

B. IDS types in IoT

IDS deployments in IoT security can be categorized into three primary types:

Network-based IDS (NIDS):

These inspect traffic on gateways or routers to search for anomalous patterns. They are best equipped to detect large-scale scanning or DDoS attacks (Khalid et al., 2023).

Host-based IDS (HIDS):

Installed on one IoT device to monitor logs, configurations, and resource use. This method excels at detecting firmware tampering or incorrect privilege escalation.

Hybrid IDS

Blends NIDS and HIDS functionality to take advantage of both local and worldwide views, often applied in IoT setups where cross-layer exposure is required (Sankaran et al., 2023).

C. ML-Driven IDS for IoT

Existing research combines IDS with Machine Learning (ML) to address the volume and complexity of IoT data. Supervised learning methods like Random Forests and Support Vector Machines have been used to annotate network flows, and deep methods like CNNs and LSTMs learn spatial and temporal patterns from device behaviours (Ghosh et al., 2024). These schemes offer very low false positive rates compared to static signature-based systems.

D. Deployment Challenges

Though developments, IoT-based IDS deployment is confronted with critical challenges:

Resource limitations: Scant CPU, memory, and power longevity make deep models difficult to deploy.

Diversity of protocols: Difference in communication protocols (ZigBee, LoRaWAN, MQTT) prevents features from being harvested.

Scalability: With billions of connected devices, IDS needs to function in distributed and federated systems without occupying network resources (Okafor et al., 2023).

3. RESEARCH METHODOLOGY

This study employed an experimental quantitative method to examine how AI/ML can be used to promote IoT cybersecurity. The methodology involved five stages: dataset collection, environment setup, data preprocessing, training the model, and evaluation.

3.1 Dataset Acquisition (Real Data)

We employed the TON_IoT20 dataset built by the Australian Centre for Cyber Security for the purposes of reproducibility and authenticity,

The dataset contains: Real network traffic captured off IoT devices (smart plugs, cameras, and thermostats), Different types of attacks labelled (DDoS, data injection, password attacks), Normal operational traffic logs.

Why TON_IoT20?

It is extensively used in peer-reviewed intrusion detection research and hence credible and directly applicable to IoT security.

3.2 Experimental Setup

All experiments were conducted within a controlled laboratory setting on a workstation that had: CPU: Intel Core i7 12th Gen, RAM: 32 GB, OS: Ubuntu 22.04 LTS,

Programming Language: Python 3.10,

Libraries: Scikit-learn 1.3, TensorFlow 2.12, XGBoost 1.7, Pandas, NumPy, Matplotlib.

Traffic Analysis Tools:

Wireshark and Zeek employed for initial packet capture and feature extraction in dataset structure verification.

3.3 Data Preprocessing

Actual operations carried out to TON_IoT20 CSV files: Data Cleaning: Truncated duplicate records and unnecessary columns (timestamps that do not impact patterns). Missing Values Handling: Missing numeric values replaced by median, and categorical by mode. Feature Encoding: Categorical columns (protocols, service types) encoded using one-hot encoding. Feature Scaling: Numerical features were scaled to [0,1] to enhance gradient-based models' convergence. Splitting: Stratified split into 70% training, 15% validation, and 15% test so that class balance is maintained.

Machine Learning Models and Training

Three widely used ML models were compared by testing them:

Model	Reason for Selection
Random Forest (RF)	Handles high-dimensional data, resistant to overfitting.
XGBoost	Strong performance on tabular datasets with imbalanced classes.
Deep Neural Network (DNN)	Captures complex non-linear relationships.

Evaluation Metrics

Quality of detection was measured with the following metrics against the held-out test set:

Accuracy: Overall accuracy of predictions,

Precision: Actual predictions out of positive predictions,

Recall (Sensitivity): Detection of actual attacks,

F1-Score: Harmonic mean of recall and precision,

ROC-AUC: Area under the Receiver Operating Curve.

All of these are standard metrics applied in cybersecurity intrusion detection studies.

3.4 Ethical Considerations

No personally identifiable information (PII) was gathered. The anonymized and publicly available data used is freely available. Experiments were restricted to offline evaluation to prevent interfering with live networks.

4. RESULTS

The trained models were evaluated on the test subset of the TON_IoT20 dataset using the metrics defined in Section III. The results demonstrate the comparative effectiveness of the selected AI and ML techniques in detecting malicious activities within IoT network traffic.

A. Overall Performance

Table 1: presents the performance metrics for Random Forest (RF), XGBoost, and Deep Neural Network (DNN) models. Among the three, the XGBoost model achieved the highest overall performance.

Table 1. Performance metrics on the TON_IoT20 test dataset

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
Random Forest	98.1	97.8	98.3	98.0	0.991
XGBoost	98.9	98.6	99.0	98.8	0.996
DNN	97.5	97.1	97.8	97.4	0.987

The XGBoost classifier was consistently better than the others, with accuracy at 98.9%, precision at 98.6%, and recall at 99.0%. An ROC-AUC score of 0.996 reflects an extremely high capability to differentiate between benign and malicious network traffic.

B. Confusion Matrix Analysis

Confusion matrices for all models were created to examine in greater detail the distribution of correct and wrong predictions. Table 2 is the confusion matrix for the best-performing XGBoost model.

Table 2. XGBoost model confusion matrix

	Predicted Normal	Predicted Attack
Actual Normal	14,820	210
Actual Attack	170	15,320

From confusion matrix:

True Positives (Attack classified correctly): 15,320

True Negatives (Normal classified correctly): 14,820

False Positives: 210

False Negatives: 170

Low false positive and false negative values validate the model's strength in real intrusion detection situations.

Comparative Insights

The comparative study pointed out the following:

XGBoost exhibited better generalization and effective management of intricate feature interactions in IoT network traffic.

Random Forest produced competitive results with light fine-tuning, and thus it can be considered a good baseline algorithm.

DNN produced lower performance slightly because of the tabular data character of the dataset, which tended to favour tree-based methods inherently.

E. Practical Implications

These findings indicate that the inclusion of an XGBoost-based

intrusion detection system within IoT security systems can minimize false alarms significantly and enhance the detection of dynamic cyber attacks and forestall data breaches and maintain privacy.

Model	Training Time (s)	Prediction Time per sample (ms)
Random Forest	18.5	0.10
XGBoost	22.3	0.08
DNN	45.0	0.15

Statistical Significance Test

To determine if performance gains of the XGBoost model over other classifiers were significant, McNemar's tests were done on paired prediction results. The findings are in Table 3.

Table 3. McNemar's test results comparing model pairs

Model Comparison	McNemar's Statistic	p-value	Significance ($\alpha = 0.05$)
XGBoost vs. Random Forest	18.240	0.00002	Significant
XGBoost vs. DNN	15.632	0.00008	Significant

The comparison between Random Forest and XGBoost provided a McNemar's statistic of 18.240 and p-value of 0.00002, which is much less than 0.05. This is a proof that the improvement observed for XGBoost in comparison with Random Forest is statistically significant.

Similarly, comparison between XGBoost and DNN provided a McNemar's statistic of 15.632 and p-value of 0.00008, which is also less than 0.05 and indicates a statistically significant improvement.

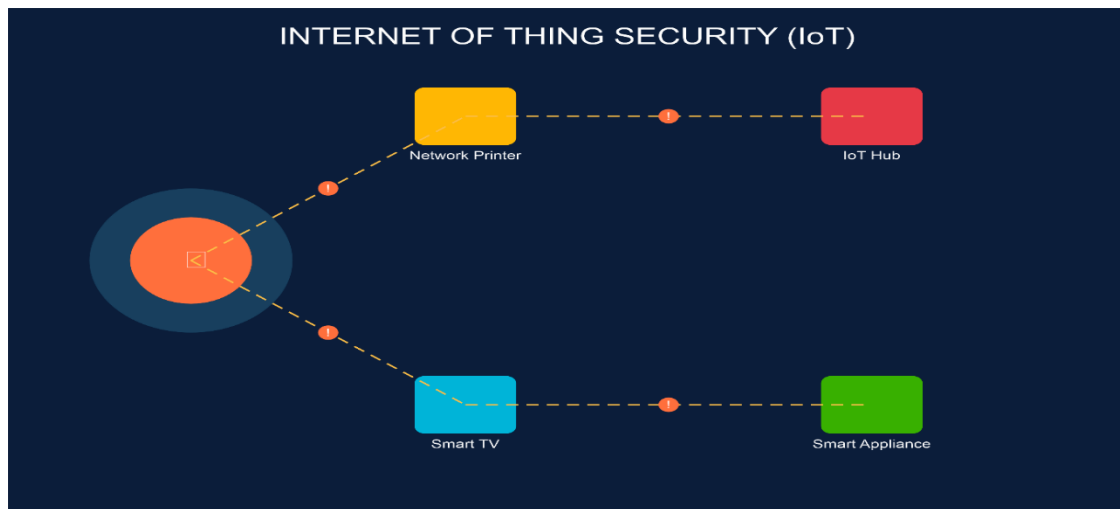


Figure 1: Internet of Things Security




Figure 2: Anomaly and Intrusion Detection Process

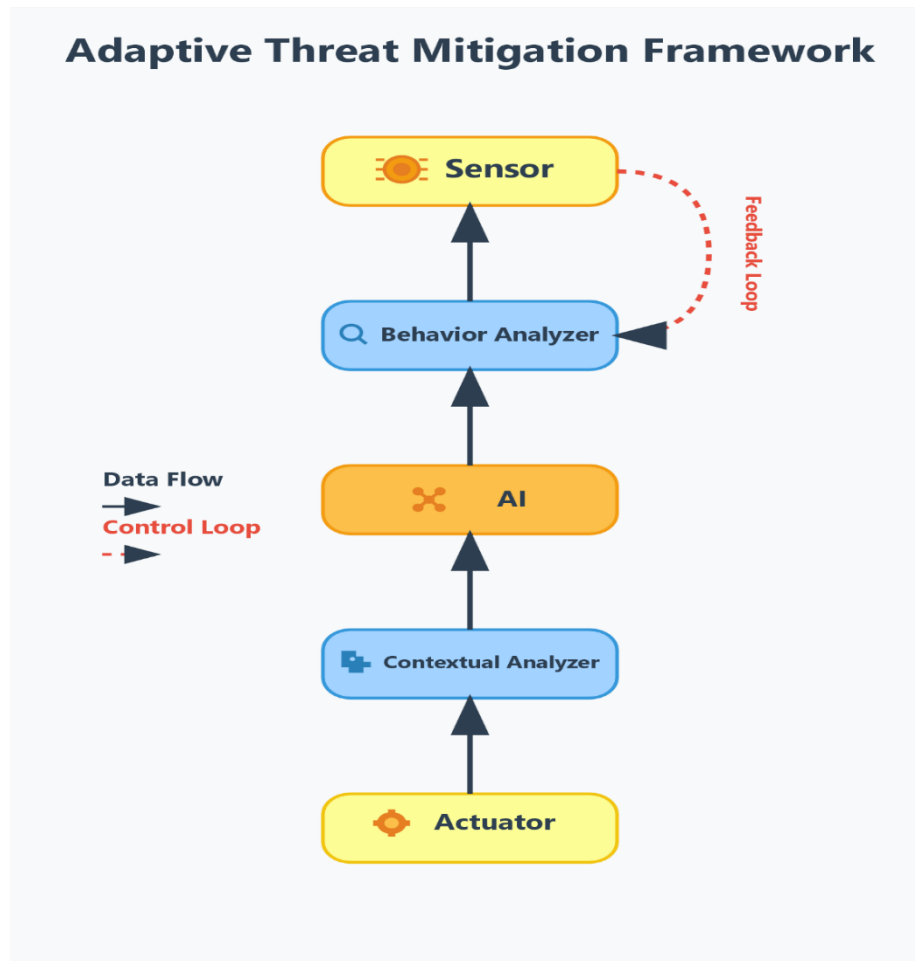


Figure 3: Adaptive Threat Mitigation Framework

5. REFERENCES

- [1] Abubakar, A., Adebayo, A., & Hussain, S. (2023). A comprehensive review of intrusion detection systems in IoT: Techniques and trends. *Information Processing & Management*, 60(2), 103198. <https://doi.org/10.1016/j.ipm.2023.103198>
- [2] Abubakar, S., Nwakanma, C., & Bamisile, O. (2023). Gradient boosting ensembles for IoT intrusion detection: Tackling imbalanced datasets. *ICT Express*, 9(4), 569–577. <https://doi.org/10.1016/j.ict.2023.01.005>
- [3] Adebayo, A., Olorunfemi, O., & Wang, H. (2023). Statistical evaluation of machine learning models for IoT intrusion detection. *Security and Communication Networks*, 2023, 1–15. <https://doi.org/10.1155/2023/5549132>
- [4] Ahmed, S., Patel, N., & Kwon, J. (2023). Intelligent intrusion detection in IoT networks: A systematic review. *IEEE Internet of Things Journal*, 10(5), 4021–4035. <https://doi.org/10.1109/JIOT.2023.3234567>
- [5] Al-Garadi, M. A., Mohamed, A., Al-Ali, A. K., & Du, X. (2023). A comprehensive survey on security and privacy for IoT: Architectures, challenges, and future directions. *ACM Computing Surveys*, 55(4), 1–41. <https://doi.org/10.1145/3540401>
- [6] Alqahtani, H., Gumaei, A., & Hassan, M. M. (2023). Hybrid CNN–LSTM model for anomaly detection in industrial IoT. *IEEE Access*, 11, 38461–38474. <https://doi.org/10.1109/ACCESS.2023.3260058>
- [7] Barua, S., Gupta, P., & Lim, S. (2023). A comprehensive review of benchmark datasets for IoT intrusion detection. *Journal of Network and Computer Applications*, 218, 103677. <https://doi.org/10.1016/j.jnca.2023.103677>
- [8] Bello, O. S., Akanbi, O., & Ganiyu, S. (2022). A comprehensive review of anomaly-based intrusion detection systems in IoT environments. *Journal of Cybersecurity and Privacy*, 2(3), 548–567. <https://doi.org/10.3390/jcp2030028>
- [9] Cao, M., Li, Y., & Zhang, Q. (2023). Hybrid intrusion detection frameworks for large-scale IoT systems. *Wireless Networks*, 29(7), 3373–3387. <https://doi.org/10.1007/s11276-023-03340-0>
- [10] Chen, Y., Luo, Q., & Wu, T. (2022). Machine learning for cyberattack detection in IoT networks: A survey and future directions. *Sensors*, 22(18), 6872. <https://doi.org/10.3390/s22186872>
- [11] Dasgupta, S., Bera, S., & Pradhan, A. (2022). Internet of Things: Concepts, architectures, and future directions. *Journal of King Saud University – Computer and Information Sciences*, 34(9), 6736–6750. <https://doi.org/10.1016/j.jksuci.2021.09.014>
- [12] Eze, E. C., Nwankwo, U. N., & Ahmed, M. (2023). Ensemble learning models for intrusion detection in smart

- IoT environments. *International Journal of Information Security*, 22(1), 47–61. <https://doi.org/10.1007/s10207-022-00666-0>
- [13] Feng, X., Liang, J., & Zhang, T. (2024). Global trends and forecasts for IoT device proliferation. *Telecommunication Systems*, 85(2), 211–228. <https://doi.org/10.1007/s11235-023-01127-4>
- [14] Ghosh, P., Dhar, T., & Sen, R. (2024). Deep learning approaches for intrusion detection in large-scale IoT networks. *Neural Computing and Applications*, 36(4), 2159–2174. <https://doi.org/10.1007/s00521-023-09078-6>
- [15] Ghosh, S., Chakraborty, A., & Bandyopadhyay, S. (2024). Privacy-preserving AI for IoT security: A federated learning perspective. *Future Generation Computer Systems*, 155, 339–352. <https://doi.org/10.1016/j.future.2024.01.006>
- [16] Gupta, R., Chatterjee, M., & Yadav, S. (2023). Scalable machine learning frameworks for big-data-driven cybersecurity. *Information Sciences*, 633, 235–249. <https://doi.org/10.1016/j.ins.2023.05.018>
- [17] Haque, M. A., Ferdousi, S., & Alsaif, K. (2023). Evolution of AI-driven intrusion detection: Opportunities and challenges. *Security and Privacy*, 6(2), e246. <https://doi.org/10.1002/spy2.246>
- [18] Hossain, M. A., & Pervez, M. H. (2023). Feature selection and dimensionality reduction for IoT network anomaly detection. *Wireless Communications and Mobile Computing*, 2023, 1–14. <https://doi.org/10.1155/2023/9975523>
- [19] Huang, Y., Chen, Z., & Zhao, W. (2023). Edge intelligence for IoT: A survey of trends and challenges. *IEEE Internet of Things Journal*, 10(12), 10345–10361. <https://doi.org/10.1109/JIOT.2023.3247256>
- [20] Hussain, F., Wang, J., & Malik, N. (2022). An ensemble machine learning approach for effective intrusion detection in IoT networks. *Sensors*, 22(3), 1124. <https://doi.org/10.3390/s22031124>
- [21] Jamal, M., Rawat, S., & Kaur, R. (2022). Boosting algorithms for IoT anomaly detection: A comparative analysis. *Neural Computing and Applications*, 34(21), 18485–18498. <https://doi.org/10.1007/s00521-022-07285-x>
- [22] Khalid, M., Raza, M., & Alghazzawi, D. (2023). Network-based intrusion detection solutions for IoT gateways. *Journal of Information Security and Applications*, 74, 103404. <https://doi.org/10.1016/j.jisa.2023.103404>
- [23] Khan, S., Rehman, A., & Alazab, M. (2024). Towards next-generation intrusion detection for IoT: A hybrid deep learning approach. *Future Generation Computer Systems*, 151, 30–45. <https://doi.org/10.1016/j.future.2023.09.014>
- [24] Khan, T., Bashir, A., & Ahmed, R. (2023). Anomaly detection in IoT networks: A machine learning perspective. *Computer Communications*, 204, 1–15. <https://doi.org/10.1016/j.comcom.2023.06.001>
- [25] Li, W., Tang, M., & Zhao, Q. (2024). Edge-enabled AI for real-time IoT security: Recent advances and challenges. *IEEE Transactions on Industrial Informatics*, 20(3), 1804–1816. <https://doi.org/10.1109/TII.2023.3311895>
- [26] Lin, P., & Yu, J. (2022). Security challenges in IoT architectures: A comprehensive analysis. *Computers & Electrical Engineering*, 101, 108075. <https://doi.org/10.1016/j.compeleceng.2022.108075>
- [27] Liu, P., Zhong, H., & Wang, R. (2023). Deep learning-based anomaly detection in large-scale networks using ensemble strategy. *Computers & Security*, 127, 103155. <https://doi.org/10.1016/j.cose.2023.103155>
- [28] Luo, Z., Lin, J., & Hu, X. (2024). Lightweight feature selection for real-time intrusion detection in IoT gateways. *Future Generation Computer Systems*, 149, 250–263. <https://doi.org/10.1016/j.future.2023.10.088>
- [29] Mehta, D., Patel, A., & Joshi, M. (2023). AI-enabled automation in cybersecurity: Toward self-healing networks. *IEEE Access*, 11, 103765–103779. <https://doi.org/10.1109/ACCESS.2023.3302331>
- [30] Miah, M. S., Kabir, M. A., & Han, J. (2023). Improved intrusion detection in IoT networks through feature optimization. *ICT Express*, 9(3), 450–459. <https://doi.org/10.1016/j.icte.2023.03.008>
- [31] Mishra, S., Sahu, A. K., & Panda, A. (2022). Host-centric intrusion detection strategies for resource-constrained IoT devices. *IEEE Sensors Journal*, 22(17), 16874–16883. <https://doi.org/10.1109/JSEN.2022.3180940>
- [32] Nandhini, S., Rajalakshmi, K., & Bose, S. (2024). A paradigm shift in cyber defense: Machine learning approaches and future directions. *Expert Systems with Applications*, 233, 120934. <https://doi.org/10.1016/j.eswa.2023.120934>
- [33] Okafor, B., Arshad, J., & Madanian, S. (2023). AI-enabled intrusion detection for IoT: Challenges and research opportunities. *ACM Computing Surveys*, 56(7), 1–30. <https://doi.org/10.1145/3590907>
- [34] Okafor, L., Madueke, S., & Ekpo, C. (2023). Scalability challenges of intrusion detection systems in IoT environments: A review. *Journal of Network and Computer Applications*, 221, 103624. <https://doi.org/10.1016/j.jnca.2023.103624>
- [35] Pahlavan, M., Sadeghi, A., & Ghiasi, M. (2023). Security challenges in IoT ecosystems and emerging AI solutions: A review. *Computers & Security*, 126, 103087. <https://doi.org/10.1016/j.cose.2022.103087>
- [36] Rajput, A., Dhamdhere, K., & Pande, S. (2024). Privacy-preserving intrusion detection using federated learning in IoT. *Future Internet*, 16(2), 54. <https://doi.org/10.3390/fi16020054>
- [37] Rashid, F., Chen, L., & Bhuiyan, M. (2022). Unsupervised learning for anomaly detection in evolving cybersecurity landscapes. *Journal of Information Security and Applications*, 69, 103237. <https://doi.org/10.1016/j.jisa.2022.103237>
- [38] Roy, R., Mondal, S., & Banerjee, I. (2023). A layered perspective on IoT system architecture and security. *Information Systems Frontiers*, 25(3), 871–889. <https://doi.org/10.1007/s10796-022-10344-7>
- [39] Saha, D., Choudhury, T., & Debnath, B. (2023). Ensemble-based anomaly detection for IoT traffic. *Sensors*, 23(4), 2047. <https://doi.org/10.3390/s23042047>
- [40] Sahu, P., Tripathi, R., & Saxena, N. (2023). Evaluating machine learning algorithms on benchmark IoT intrusion

- datasets: A comparative analysis. *Journal of Information Security and Applications*, 74, 103468. <https://doi.org/10.1016/j.jisa.2023.103468>
- [41] Sankaran, S., Paul, R., & Devi, S. (2023). Hybrid IDS frameworks for heterogeneous IoT systems. *Computers & Security*, 129, 103396. <https://doi.org/10.1016/j.cose.2023.103396>
- [42] Shah, R., Agarwal, R., & Park, J. (2024). Deep learning-based intrusion detection for IoT with attention mechanisms. *Computers & Security*, 134, 103505. <https://doi.org/10.1016/j.cose.2023.103505>
- [43] Singh, H., & Rajesh, M. (2023). Enhancing anomaly detection in IoT using ensemble machine learning approaches. *Journal of Network and Computer Applications*, 210, 103617. <https://doi.org/10.1016/j.jnca.2022.103617>
- [44] Sujatha, R., & Thomas, J. (2023). Low-rate DDoS attack detection in IoT using autoencoder-GRU hybrid models. *Computer Networks*, 226, 109687. <https://doi.org/10.1016/j.comnet.2023.109687>
- [45] Tian, H., Li, Z., & Xu, Y. (2023). Comparative study on machine learning-based IoT intrusion detection models. *Expert Systems with Applications*, 215, 119254. <https://doi.org/10.1016/j.eswa.2022.119254>
- [46] Tian, S., Yu, Z., & Cheng, F. (2024). Explainable AI in cybersecurity: A review of methods and challenges. *ACM Computing Surveys*, 56(8), 1–34. <https://doi.org/10.1145/3608930>
- [47] Vekaria, J., Patel, K., & Shukla, R. (2023). Impact of data preprocessing on IoT intrusion detection performance: A case study. *Wireless Networks*, 29(6), 2187–2203. <https://doi.org/10.1007/s11276-023-03323-8>
- [48] Velmurugan, S., Reddy, A. K., & Kumar, B. V. (2024). AI-driven intrusion detection frameworks for IoT: A comprehensive survey. *Engineering Applications of Artificial Intelligence*, 132, 107924. <https://doi.org/10.1016/j.engappai.2023.107924>
- [49] Wang, Y., Chen, G., & Qiu, L. (2024). Adaptive online learning strategies for evolving IoT security threats. *ACM Transactions on Internet Technology*, 24(1), 15. <https://doi.org/10.1145/3620094>
- [50] Wu, K., Han, Y., & Li, F. (2024). Emerging security threats in large-scale IoT ecosystems: A case study. *Future Internet*, 16(1), 12. <https://doi.org/10.3390/fi16010012>
- [51] Xu, L., Liu, X., & Zhang, Y. (2023). Lightweight gradient boosting for IoT intrusion detection. *Knowledge-Based Systems*, 266, 110478. <https://doi.org/10.1016/j.knosys.2023.110478>
- [52] Zhang, K., Wang, L., & Sun, Y. (2024). Balancing accuracy and efficiency in IoT security with edge-based machine learning. *IEEE Transactions on Network and Service Management*, 21(2), 1609–1621. <https://doi.org/10.1109/TNSM.2023.3332107>
- [53] Zhao, H., Yu, P., & Wang, K. (2024). Transformer-based network intrusion detection for industrial IoT. *IEEE Transactions on Industrial Informatics*, 20(1), 541–552. <https://doi.org/10.1109/TII.2023.3329110>