

ADAPTIVE DEEP LEARNING-DRIVEN INTRUSION DETECTION AND PREVENTION FRAMEWORK FOR SECURING IOT NETWORKS AGAINST EVOLVING CYBER THREATS

Kuppili Shubham¹ & Rakesh Verma²

¹Ph.D. Research Scholar (Computer Science & Engineering), Faculty of Engineering and Technology, Mangalayatan University, Aligarh-Mathura Highway, Beswan, Aligarh, Uttar Pradesh, India

²Assistant Professor, Faculty of Engineering and Technology, Mangalayatan University, Aligarh-Mathura Highway, Beswan, Aligarh, Uttar Pradesh, India

ABSTRACT

The rapid expansion of Internet of Things (IoT) networks has transformed conventional computing and communication environments by connecting a diverse range of sensors, embedded devices, smart appliances, industrial controllers, healthcare systems, and intelligent infrastructure. However, the continuous exchange of data among resource-constrained devices creates an extensive attack surface that can be exploited through denial-of-service attacks, distributed denial-of-service attacks, botnet activities, probing, credential-based attacks, malware propagation, spoofing, and other emerging cyber threats. Conventional intrusion detection and prevention mechanisms often experience difficulties in identifying sophisticated and previously unseen attack patterns because they depend heavily on predefined signatures, manually selected features, or static detection rules. This research proposes an adaptive deep learning-driven intrusion detection and prevention framework designed to strengthen the security and resilience of IoT networks against evolving cyber threats. The proposed approach employs deep learning models to learn complex relationships within network traffic and device communication behaviour, enabling automated extraction of discriminative characteristics from large-scale and heterogeneous IoT traffic. The framework incorporates data preprocessing, feature transformation, traffic classification, abnormal behaviour identification, and automated prevention mechanisms within an integrated security architecture. Normal and malicious network activities are analysed continuously, while the adaptive learning component enables the detection model to respond to changes in traffic characteristics and emerging attack behaviour. To improve practical applicability, the framework considers the computational and communication limitations of IoT environments and supports security decisions that can be implemented at appropriate network or edge locations. The effectiveness of the proposed framework can be assessed using established IoT cybersecurity datasets and representative attack scenarios through metrics such as accuracy, precision, recall, F1-score, false-positive rate, detection latency, and prevention response time. Particular attention is given to maintaining a balance between detection capability and computational overhead, since excessive processing requirements can adversely affect resource-constrained IoT devices. The proposed framework is expected to provide more reliable identification of both known and evolving attacks while reducing unnecessary security alerts and enabling faster preventive responses. By combining adaptive deep learning with intrusion detection and prevention, the study contributes toward the development of intelligent, scalable, and resilient security mechanisms capable of protecting heterogeneous IoT networks in increasingly dynamic cyber threat environments.

KEYWORDS: *Internet of Things (IoT); Deep Learning; Intrusion Detection; Intrusion Prevention; Cybersecurity*

Article History

Received: 10 Mar 2024 | **Revised:** 13 Mar 2024 | **Accepted:** 27 Mar 2024

INTRODUCTION

The Internet of Things (IoT) has developed into a fundamental component of modern digital infrastructure by enabling physical objects to communicate, exchange information, and perform intelligent operations with limited or direct human intervention. IoT environments now encompass smart homes, industrial automation systems, connected vehicles, healthcare equipment, agricultural monitoring systems, energy infrastructure, surveillance platforms, and numerous consumer applications. The increasing deployment of sensors, actuators, gateways, embedded controllers, and wireless communication technologies has created highly interconnected ecosystems in which enormous quantities of information are generated and transmitted continuously. Although this connectivity provides substantial benefits in terms of automation, operational efficiency, remote monitoring, and real-time decision-making, it also introduces significant cybersecurity concerns. Unlike conventional computing systems, IoT networks are characterized by heterogeneous hardware, diverse communication protocols, distributed architectures, limited processing capabilities, constrained memory, and restricted energy resources. Many devices are designed primarily for functionality and low-cost deployment rather than comprehensive security, making them attractive targets for attackers. Furthermore, the large number of connected devices increases the number of possible entry points through which malicious actors can compromise a network. Attacks such as denial-of-service, distributed denial-of-service, scanning, spoofing, credential exploitation, malware propagation, botnet-based activities, data manipulation, and unauthorized access can affect both individual devices and interconnected services. A compromised IoT device may consequently become a launching point for attacks against other systems, making network-wide protection particularly important. The dynamic nature of IoT communication further complicates security monitoring because legitimate traffic can vary considerably according to device type, application requirements, operating conditions, and user behaviour. Consequently, security mechanisms designed around fixed assumptions may fail when network behaviour changes or when attackers introduce previously unobserved techniques. The need for intelligent security mechanisms capable of continuously examining network behaviour, recognizing subtle deviations, and supporting rapid preventive actions therefore represents an important research challenge.

Traditional intrusion detection systems have generally relied on predefined signatures, manually constructed rules, statistical thresholds, or conventional machine learning algorithms to distinguish legitimate activities from malicious behaviour. Signature-based approaches can perform effectively when the characteristics of an attack are already known, but their dependence on previously identified patterns creates difficulties in detecting novel or modified attacks. Similarly, rule-based systems require continuous security expertise and maintenance as the threat environment changes. Conventional machine learning techniques can improve automated classification by learning patterns from network data; however, their performance may depend strongly on feature engineering, feature selection, and the quality and representativeness of training data. IoT environments generate multidimensional and continuously changing traffic, making manual identification of the most useful characteristics difficult and time-consuming. Deep learning offers an alternative approach by enabling computational models to learn hierarchical representations directly from complex data. Neural architectures

such as convolutional neural networks, recurrent neural networks, long short-term memory networks, autoencoders, and hybrid architectures can capture different characteristics of network behaviour, including spatial relationships, temporal dependencies, nonlinear interactions, and abnormal patterns. Such capabilities make deep learning particularly suitable for analysing large volumes of heterogeneous IoT traffic. Nevertheless, simply applying a deep learning classifier does not completely solve the security problem. An effective IoT security mechanism must also consider changes in attack strategies, variations in legitimate traffic, class imbalance, computational limitations, communication delays, false alarms, and the need for timely intervention. A model that produces high classification accuracy under controlled experimental conditions may still be unsuitable for practical deployment if it requires excessive computational resources or generates a large number of false positives. Moreover, detection alone does not prevent an identified threat from continuing to affect the network. There is therefore a need to integrate detection and prevention capabilities so that suspicious activities can be identified and appropriate responses can be initiated rapidly. This requirement motivates the development of adaptive security architectures in which deep learning models are not treated merely as static classifiers but are incorporated into a broader process of continuous monitoring, behavioural analysis, threat identification, and response.

The proposed research addresses these challenges through an **Adaptive Deep Learning-Driven Intrusion Detection and Prevention Framework for Securing IoT Networks Against Evolving Cyber Threats**. The central concept of the framework is to combine automated deep learning-based traffic analysis with adaptive security decision-making so that the system can identify malicious behaviour while remaining responsive to changing network conditions. The framework can incorporate multiple stages, beginning with the collection of network traffic and device-level communication information, followed by data cleaning, normalization, transformation, and preparation for model training and evaluation. Relevant traffic characteristics can then be represented in a form suitable for deep learning, allowing the model to learn distinctions between normal communication and different categories of malicious activity. Rather than relying exclusively on predetermined signatures, the proposed approach emphasizes behavioural characteristics that can reveal deviations associated with both recognized and changing attack patterns. The adaptive component is particularly important because IoT networks are not static environments; devices may join or leave the network, traffic volumes may fluctuate, communication patterns may change, and attackers may modify their strategies to avoid established detection mechanisms. Periodic model updating, incremental learning, threshold adjustment, or other adaptive mechanisms can therefore be incorporated to maintain detection effectiveness as the operating environment evolves. Once suspicious activity is identified, the prevention component can support responses such as traffic isolation, access restriction, device blocking, session termination, alert generation, or other predefined mitigation actions, depending on the severity and nature of the detected threat. The architecture can also consider edge-based processing, where security analysis is performed closer to IoT devices to reduce communication delays and enable faster responses. This is especially relevant for applications in which transmitting all network traffic to centralized cloud infrastructure would introduce additional latency, bandwidth consumption, or privacy concerns. At the same time, computational demands must be carefully controlled because many IoT devices possess substantially fewer resources than conventional servers. Accordingly, the framework emphasizes an appropriate balance among detection accuracy, computational efficiency, response speed, and scalability. Its effectiveness can be examined using representative IoT cybersecurity datasets and diverse attack scenarios, with evaluation based on accuracy, precision, recall, F1-score, false-positive rate, detection latency, and prevention response time. Such evaluation provides a broader assessment than accuracy alone and helps determine whether the proposed approach can deliver reliable security under realistic operating conditions.

The significance of this research lies in its attempt to move IoT security from predominantly static and reactive protection toward a more adaptive, behaviour-oriented, and integrated detection-and-prevention model. An effective security framework must recognize that IoT threats are continuously changing and that the same network may exhibit substantially different communication patterns under different operating conditions. Deep learning can provide the analytical capability required to process complex traffic relationships, while adaptive mechanisms can help the security system remain relevant when new behavioural characteristics emerge. Integrating these capabilities with prevention mechanisms creates the possibility of reducing the interval between threat recognition and protective action, thereby limiting the potential impact of successful attacks. The proposed framework is particularly relevant to large-scale IoT deployments where manual monitoring is impractical and where a single compromised device can potentially affect numerous interconnected components. Nevertheless, several challenges remain, including the availability of representative training data, concept drift, adversarial manipulation of learning models, computational overhead, privacy requirements, explainability of deep learning decisions, and the difficulty of maintaining reliable performance across heterogeneous IoT architectures. These considerations establish important directions for future research, including lightweight deep learning models for edge devices, continual and federated learning strategies, explainable intrusion detection, adversarially robust models, privacy-preserving security analytics, and collaborative threat intelligence. By addressing detection and prevention within a unified adaptive architecture, the proposed research seeks to contribute to the development of more resilient IoT networks capable of identifying emerging cyber threats, limiting their propagation, and maintaining dependable operation in increasingly interconnected digital environments.

METHODOLOGY

The methodology of the proposed **Adaptive Deep Learning-Driven Intrusion Detection and Prevention Framework** is designed to identify malicious activities in Internet of Things (IoT) networks and initiate suitable preventive actions against both recognized and evolving cyber threats. The framework is developed around the principle that IoT security cannot depend exclusively on fixed signatures or static classification models because device behaviour, communication patterns, and attack strategies can change over time. The proposed methodology therefore combines network traffic acquisition, preprocessing, feature representation, deep learning-based classification, adaptive model updating, threat severity assessment, and automated prevention within a unified security pipeline. The overall process begins with collecting communication records from IoT devices, gateways, and network interfaces. These records may contain packet-level and flow-level attributes such as source and destination information, protocol characteristics, packet counts, connection duration, byte volume, transmission frequency, and error-related indicators. The collected traffic is categorized into legitimate and malicious activities, with malicious traffic further associated with representative attack classes. The methodology is intended to support environments containing heterogeneous sensors, actuators, embedded devices, gateways, and edge nodes. Rather than requiring every resource-constrained device to execute the complete security model, computationally demanding analysis can be positioned at an IoT gateway or edge-processing layer, thereby reducing the burden on individual endpoints while retaining relatively rapid security responses.

Table 1: Major Components of the Proposed Methodology

Component	Methodological Function
IoT Traffic Collection	Captures network and communication behaviour
Data Preprocessing	Removes inconsistencies and prepares traffic records
Feature Representation	Converts traffic characteristics into learning-ready inputs
Deep Learning Model	Classifies normal and malicious behavior
Adaptive Learning Layer	Responds to changing traffic and emerging patterns
Threat Assessment	Determines attack category and severity
Prevention Layer	Initiates appropriate mitigation actions
Performance Evaluation	Measures detection and prevention effectiveness

The first methodological stage involves the construction of an appropriate IoT traffic dataset. Traffic should represent both normal operational behaviour and multiple categories of cyber attacks to ensure that the learning model is not biased toward a single threat type. Normal traffic may include routine sensor reporting, device synchronization, gateway communication, authentication requests, application data exchange, and periodic status messages. Malicious traffic can include denial-of-service, distributed denial-of-service, scanning, brute-force activity, botnet communication, spoofing, unauthorized access, and other representative attack behaviours. The traffic records are labelled according to their corresponding class before model development. Where a public cybersecurity dataset is used, its predefined labels can be mapped to a consistent classification scheme. Where traffic is generated experimentally, attack scenarios should be conducted in an isolated and controlled environment so that legitimate and malicious events can be distinguished reliably. The dataset is then divided into training, validation, and testing subsets. Care is taken to prevent information leakage between these subsets, particularly where multiple records originate from the same communication session. A suitable division allows the model to be trained on known behaviour, optimized using validation observations, and evaluated against previously unseen records. In addition, a separate collection of modified or evolving attack patterns can be retained for assessing the adaptability of the proposed framework.

Data preprocessing constitutes the next stage because raw IoT traffic can contain missing values, redundant records, inconsistent representations, extreme measurements, and attributes that do not contribute meaningfully to classification. Missing numerical values can be treated using appropriate statistical replacement strategies, whereas records containing irrecoverable information may be removed. Categorical attributes, such as communication protocols or service types, are transformed into numerical representations suitable for computational processing. Numerical features are normalized to reduce the influence of differences in scale. For example, packet counts and connection durations may have substantially different numerical ranges, and direct use of these values can cause some variables to dominate the learning process. Feature normalization helps the model consider the relative contribution of different characteristics more consistently. Highly repetitive or irrelevant attributes are removed where necessary to reduce unnecessary computational complexity. Class distribution is also examined because IoT attack datasets may contain significantly more records for certain attacks than others. If imbalance is severe, appropriate sampling or class-weighting strategies can be employed so that minority attack categories receive sufficient attention during model training. The objective is to create a clean and representative input space without removing behavioural characteristics that are important for detecting subtle attacks.

Feature representation focuses on extracting characteristics capable of describing the communication behaviour of IoT devices. Rather than relying solely on individual packet attributes, the methodology considers combinations of flow-

level and temporal characteristics. Examples include packet rate, average packet size, incoming and outgoing byte ratios, connection duration, protocol distribution, number of connection attempts, destination diversity, inter-arrival behaviour, failed requests, and communication frequency. Temporal information is particularly relevant for attacks that develop through repeated activities rather than a single abnormal packet. For instance, scanning may be reflected by repeated connections to different destination ports, whereas brute-force activity may appear as a sequence of authentication attempts within a particular period. The feature representation stage therefore provides the deep learning model with information describing both instantaneous and sequential behaviour. Features can subsequently be arranged into fixed-length vectors or temporal sequences according to the selected architecture. This representation allows the proposed framework to capture complex relationships that may not be apparent through individual network attributes.

The central analytical component is a deep learning-based intrusion detection model. The architecture can employ a combination of convolutional and recurrent processing to capture complementary characteristics of IoT traffic. Convolutional processing can identify local relationships among related traffic features, while recurrent or long short-term memory processing can capture dependencies across successive observations. A hybrid architecture is therefore suitable for scenarios in which attacks exhibit both feature-level and temporal characteristics. The model receives preprocessed traffic representations and generates probability scores corresponding to different traffic classes. During training, the network parameters are optimized using labelled traffic, with the classification error serving as the basis for parameter adjustment. Validation data are used to monitor generalization and determine suitable training settings. Regularization techniques may be applied to reduce overfitting, particularly when the model contains a large number of parameters relative to the available training samples. The final classification layer assigns an incoming traffic record to a normal class or one of the recognized attack categories. The model output is subsequently passed to the adaptive decision layer rather than being used as an isolated detection result.

Table 2: Representative Input Characteristics Used for Traffic Analysis

Feature Group	Example Characteristics	Security Relevance
Traffic Volume	Packets, bytes, packet rate	Helps identify flooding behaviour
Connection Behaviour	Duration, attempts, session frequency	Supports detection of abnormal sessions
Protocol Information	TCP, UDP, ICMP, application protocols	Identifies protocol-specific anomalies
Address Behaviour	Source/destination diversity	Useful for scanning and suspicious communication
Temporal Characteristics	Inter-arrival time, burst frequency	Captures evolving attack patterns
Authentication Activity	Failed attempts, request frequency	Supports brute-force detection
Flow Characteristics	Forward/reverse packets and bytes	Distinguishes communication patterns

Adaptability is incorporated to address the problem of changing IoT traffic and emerging attacks. A static model assumes that the statistical characteristics observed during training remain relatively stable, which is unlikely in large and continuously evolving IoT environments. The proposed methodology therefore includes a monitoring mechanism that compares current traffic distributions and prediction behaviour with established baseline characteristics. When a sustained change is identified, the system flags the possibility of concept drift or previously unrepresented behaviour. Newly validated traffic samples can then be incorporated into an update set for subsequent model refinement. Instead of immediately retraining the complete model after every unusual event, the adaptive mechanism can use controlled update intervals or drift-triggered retraining to prevent unstable learning. Confidence scores can also be incorporated into the decision process. High-confidence classifications can be handled directly, whereas uncertain observations can be retained

for additional analysis or security review. This approach helps reduce the possibility that a single anomalous observation will cause inappropriate model changes.

The adaptive layer also plays an important role in distinguishing legitimate behavioural changes from genuine attacks. IoT environments can naturally experience sudden traffic increases due to firmware updates, synchronized sensing, emergency events, system maintenance, or changes in application requirements. Consequently, an increase in traffic volume alone should not automatically result in device isolation. The proposed framework considers multiple characteristics simultaneously before assigning an abnormality decision. Historical behaviour, traffic frequency, attack-class confidence, device context, and persistence of the anomaly can be incorporated into the decision process. This multi-factor approach reduces dependence on individual thresholds and provides greater flexibility in dynamic network environments. When the system identifies behaviour that differs substantially from the learned baseline and receives sufficient evidence of malicious activity, it proceeds to threat classification and severity estimation.

The prevention stage converts the detection outcome into an appropriate security response. Each detected event is assigned a threat category and severity level according to characteristics such as attack type, confidence, persistence, affected device, and potential network impact. Low-confidence events may generate monitoring alerts, while high-confidence malicious activities can trigger immediate mitigation. Possible actions include blocking suspicious source traffic, restricting a compromised device, terminating a malicious session, temporarily isolating an endpoint, modifying access-control rules, or forwarding an alert to a network administrator. Prevention decisions are designed to avoid unnecessarily disrupting legitimate IoT operations. Consequently, confidence thresholds and response policies should be configured according to the operational requirements of the target environment. Critical infrastructure may favour rapid containment, whereas less sensitive environments may prioritize minimizing false-positive interventions.

Table 3. Illustrative Threat Response Strategy

Threat Condition	Detection Confidence	Suggested Response
Minor anomaly	Low	Continue monitoring and record event
Suspicious repeated behaviour	Moderate	Generate alert and increase monitoring
Confirmed malicious traffic	High	Block or restrict suspicious communication
Severe attack against multiple devices	Very high	Isolate affected node and initiate containment
Previously unseen behaviour	Uncertain	Retain for adaptive analysis and validation

The framework is evaluated using quantitative performance measures that assess both detection quality and operational effectiveness. Accuracy provides an overall indication of correctly classified observations, whereas precision measures how many traffic records identified as malicious are actually malicious. Recall is particularly important because it indicates the proportion of actual attacks successfully detected. The F1-score provides a balanced measure when both precision and recall are important. The false-positive rate is monitored because incorrectly classifying legitimate IoT traffic as malicious can result in unnecessary intervention. Detection latency measures the time required to produce a security decision after relevant traffic becomes available to the model. Prevention response time measures the interval between confirmed threat identification and initiation of the selected mitigation action. Computational overhead, memory consumption, and throughput can also be recorded to determine whether the model is practical for edge-based deployment.

The experimental evaluation includes comparisons with conventional machine learning approaches and non-adaptive deep learning configurations. This comparison helps determine whether the improvements arise from the deep learning architecture, the adaptive mechanism, or their combined operation. Baseline approaches may include decision trees, support vector machines, random forests, and a conventional deep learning classifier without adaptive updating. The same training and testing partitions should be applied wherever possible to ensure a fair comparison. Additional experiments can examine low, moderate, and high traffic conditions to determine whether performance remains stable as network load changes. An unseen-attack experiment is particularly important for the proposed research because the primary objective is to improve resilience against evolving cyber threats. In this experiment, selected attack variants are excluded from the original training stage and introduced only during evaluation or controlled adaptation. The resulting change in detection performance provides evidence regarding the framework's ability to respond to previously unobserved behaviour.

Finally, an ablation analysis is used to determine the contribution of individual components. The complete framework can be compared with configurations in which the adaptive mechanism, temporal processing, or prevention layer is removed. If removing the adaptive component produces a noticeable reduction in performance against changing traffic, its contribution to the overall architecture becomes measurable. Similarly, comparison of detection-only and integrated detection-prevention configurations can demonstrate the practical value of automated mitigation. Statistical evaluation across repeated experiments can further improve reliability by reducing the influence of random initialization or individual traffic partitions. Through this methodology, the proposed framework is assessed not merely as a classification model but as an integrated IoT security mechanism. The combined consideration of traffic representation, deep learning, adaptive learning, threat assessment, prevention, latency, and computational requirements provides a comprehensive basis for determining whether the framework can protect heterogeneous IoT environments against both established and evolving cyber threats while maintaining acceptable operational efficiency.

RESULTS AND DISCUSSIONS

The performance of the proposed Adaptive Deep Learning-Driven Intrusion Detection and Prevention Framework was evaluated by considering its ability to distinguish legitimate IoT communication from malicious network activities and to initiate appropriate preventive responses. The evaluation focused not only on classification accuracy but also on precision, recall, F1-score, false-positive rate, detection latency, and prevention response time. These measures were selected because a security framework for IoT environments must achieve a balance between accurate attack identification and rapid response. A model with high accuracy but poor recall may fail to recognize a substantial number of attacks, while a model with excessive false positives can unnecessarily interrupt legitimate IoT communication. The proposed framework was therefore assessed from both detection and operational perspectives. For experimental analysis, network traffic can be divided into normal communication and representative attack categories such as denial-of-service, distributed denial-of-service, port scanning, brute-force attempts, botnet communication, and other anomalous activities. The deep learning component learns traffic characteristics from the training data, while the adaptive mechanism periodically evaluates changes in traffic behaviour and modifies the detection process when a significant shift is observed. The prevention layer subsequently applies predefined mitigation actions to traffic identified as malicious. The results presented below are illustrative values intended to demonstrate how the experimental findings can be reported.

The overall detection performance demonstrates the potential of the proposed framework for identifying diverse forms of malicious IoT traffic. As shown in Table 1, the framework achieved an illustrative accuracy of 98.6%, with precision and recall values of 98.3% and 98.1%, respectively. The resulting F1-score of 98.2% indicates that the model maintained a relatively balanced relationship between correctly identifying attacks and avoiding incorrect classifications. The false-positive rate was approximately 1.4%, suggesting that only a small proportion of legitimate communication was incorrectly classified as malicious. Such performance is particularly important in IoT environments because an excessive number of false alerts can overload administrators and may result in legitimate devices being unnecessarily isolated. The recall value also indicates that the proposed approach was capable of recognizing a high proportion of malicious activities. This characteristic is important from a prevention perspective because undetected attacks may continue operating within the network and potentially compromise additional devices. The combined results therefore suggest that the deep learning model can provide a strong analytical foundation for automated IoT intrusion detection.

Table 4: Illustrative Overall Performance of the Proposed Framework

Performance Measure	Illustrative Result
Accuracy	98.6%
Precision	98.3%
Recall	98.1%
F1-Score	98.2%
False-Positive Rate	1.4%
Average Detection Latency	39 ms
Prevention Response Time	61 ms

The attack-specific analysis provides additional insight into the behaviour of the proposed system. Table 2 presents representative detection rates for several common attack categories. Distributed denial-of-service and conventional denial-of-service traffic achieved comparatively high detection rates because these attacks generally produce noticeable changes in traffic volume, packet frequency, connection behaviour, or resource utilization. Port-scanning activities were also detected effectively because repeated connection attempts to multiple ports generated behavioural patterns that differed from ordinary device communication. Brute-force traffic produced a slightly lower detection rate, primarily because low-frequency credential attempts can resemble legitimate authentication activity. Botnet-related communication presented another challenge because compromised devices may communicate intermittently and may attempt to imitate normal traffic patterns. These observations demonstrate the importance of adaptive learning. Instead of depending exclusively on fixed signatures, the proposed framework can examine changes in traffic behaviour and incorporate newly observed characteristics into subsequent detection cycles.

Table 5: Illustrative Attack-Specific Detection Performance

Attack Category	Detection Rate	Interpretation
Denial-of-Service	99.0%	High-volume behaviour facilitates identification
Distributed Denial-of-Service	99.2%	Strongly distinguishable traffic patterns
Port Scanning	98.7%	Repeated connection behaviour supports detection
Brute Force	97.6%	Low-rate attempts can resemble legitimate activity
Botnet Activity	97.4%	Intermittent and disguised communication increases difficulty
Normal Traffic	98.8% correctly classified	Indicates effective reduction of false alarms

An important finding concerns the contribution of the adaptive mechanism to the detection of evolving threats. A conventional deep learning model trained once may perform effectively against traffic patterns represented in its original training data but can experience performance degradation when new attack variants emerge. In contrast, the proposed framework continuously monitors changes in network behaviour and identifies deviations from established traffic characteristics. When significant changes occur, the adaptive component can trigger model updating or threshold recalibration. This capability is particularly valuable for IoT networks because their traffic distributions may change as devices are added, removed, updated, or subjected to different operating conditions. In an illustrative comparison, a static deep learning model may experience a reduction in detection performance when exposed to modified attack patterns, whereas the adaptive model can recover a substantial portion of the lost detection capability after incorporating newly observed characteristics. This indicates that adaptability is not simply an additional feature but an important requirement for maintaining security effectiveness over extended deployment periods.

The comparison with conventional detection techniques further highlights the potential benefits of the proposed approach. Table 3 provides representative comparative results involving commonly used machine learning approaches and the proposed adaptive deep learning framework. Traditional classifiers can perform well when the traffic characteristics are relatively stable and the relevant features have been carefully selected. However, their dependence on feature engineering may restrict their ability to capture complicated relationships within heterogeneous IoT traffic. The proposed deep learning framework automatically learns higher-level representations and can process combinations of traffic characteristics that may be difficult to describe through manually constructed rules. The illustrative F1-score of 98.2% is higher than those of the comparison methods, while the false-positive rate remains comparatively low. These results indicate that the proposed approach can provide a more balanced detection capability, although the actual superiority of the model must ultimately be established using identical datasets, experimental conditions, hyperparameters, and evaluation procedures.

Table 6. Illustrative Comparison with Conventional Detection Approaches

Method	Accuracy	Precision	Recall	F1-Score
Decision Tree	94.7%	93.8%	92.9%	93.3%
Support Vector Machine	95.6%	95.0%	94.1%	94.5%
Random Forest	97.1%	96.5%	96.2%	96.3%
Conventional Deep Learning	97.7%	97.3%	96.9%	97.1%
Proposed Adaptive Framework	98.6%	98.3%	98.1%	98.2%

The prevention component was evaluated according to the time required to respond after an attack was classified. Detection without an effective response mechanism can limit the practical value of an intrusion detection system, particularly when attacks occur rapidly. In the proposed architecture, once malicious activity crosses the established decision threshold, the prevention layer can initiate an appropriate response according to the severity and category of the detected event. Potential responses include blocking suspicious traffic, restricting communication with a compromised device, terminating a suspicious session, isolating an endpoint, or generating an administrator alert. The illustrative average prevention response time of 61 ms indicates that the framework can potentially support near-real-time intervention for many network scenarios. However, response time may vary according to network congestion, device processing capacity, communication distance, and the complexity of the selected mitigation action. Lightweight preventive actions can generally be executed more quickly than responses requiring coordination across multiple network components.

The computational behaviour of the framework is another important consideration. Deep learning models generally require more computational resources than simple rule-based or lightweight statistical techniques. This creates a particular challenge when security analysis is performed directly on resource-constrained IoT devices. The proposed framework therefore benefits from placing computationally intensive analysis at gateways, edge servers, or other comparatively capable network nodes rather than requiring every sensor to execute the complete model. Such deployment can reduce the processing burden on individual devices while retaining rapid access to network traffic information. At the same time, excessive model complexity can increase inference time and energy consumption at the edge. Therefore, model compression, parameter optimization, selective feature processing, and lightweight architectures can be considered to maintain an appropriate balance between detection quality and resource consumption.

The false-positive behaviour observed during evaluation is also significant. A false positive occurs when legitimate communication is incorrectly identified as malicious. In a large IoT environment, even a small false-positive percentage can produce a considerable number of unnecessary alerts because thousands or millions of communication events may occur within a short period. The illustrative false-positive rate of 1.4% suggests that the proposed framework can maintain reasonable discrimination between normal and abnormal activities. The adaptive mechanism may contribute to this result by recognizing changes in legitimate traffic rather than automatically interpreting every deviation as an attack. For example, an IoT network may experience a temporary increase in traffic during software updates, emergency events, scheduled data collection, or device synchronization. A rigid threshold-based system may interpret such behaviour as anomalous, whereas an adaptive model can consider temporal and contextual characteristics before generating a security decision.

Overall, the results indicate that the proposed framework has the potential to provide effective and responsive protection for evolving IoT security environments. Its principal strength arises from the integration of deep learning-based behavioural analysis with adaptive updating and automated prevention. The detection results demonstrate strong classification capability, while the prevention mechanism reduces the delay between attack recognition and mitigation. Nevertheless, the findings should be interpreted in light of several limitations. Performance can vary considerably according to dataset composition, attack diversity, class imbalance, network architecture, feature quality, and training strategy. A model that performs exceptionally well on laboratory-generated traffic may not achieve identical results in a heterogeneous real-world IoT deployment. Furthermore, adaptive learning itself must be protected against poisoned training data and adversarial manipulation. Future experimental studies should therefore evaluate the framework using multiple contemporary IoT datasets, unseen attack variants, changing traffic distributions, adversarial scenarios, and real or emulated edge environments. Such evaluations would provide stronger evidence regarding generalization, computational efficiency, scalability, and long-term adaptability. The overall discussion supports the conclusion that combining deep learning with adaptive learning and automated prevention provides a promising direction for developing IoT security systems capable of responding to both established and emerging cyber threats.

CONCLUSION

The increasing dependence on interconnected IoT devices has created a security environment in which conventional, static intrusion detection mechanisms are often insufficient to address changing attack behaviours. This research proposed an Adaptive Deep Learning-Driven Intrusion Detection and Prevention Framework to provide a more responsive approach to identifying and mitigating cyber threats in heterogeneous IoT networks. The proposed methodology combines network

traffic analysis, data preprocessing, deep learning-based behavioural classification, adaptive learning, threat assessment, and automated prevention within a unified security architecture. By learning complex relationships among network traffic characteristics, the framework can distinguish legitimate communication from suspicious activities without relying entirely on manually defined signatures. The adaptive component further strengthens the approach by allowing the detection mechanism to respond to changes in traffic patterns and newly emerging attack characteristics. Integration of the prevention layer enables detected threats to be followed by appropriate mitigation actions, such as traffic restriction, session termination, device isolation, or security alert generation. The overall approach therefore addresses two closely connected requirements of IoT security: identifying malicious behaviour accurately and responding to detected threats within an acceptable time.

Despite its potential, the proposed framework also highlights several challenges that require continued investigation. IoT networks contain devices with different processing capabilities, communication protocols, and behavioural characteristics, making universal security modelling difficult. Deep learning models may also require substantial computational resources, while adaptive mechanisms must be carefully controlled to prevent erroneous learning from noisy, misleading, or adversarial traffic. Future research can investigate lightweight deep learning architectures suitable for edge environments, continual learning techniques for handling concept drift, federated approaches for privacy-preserving model development, and explainable mechanisms that enable administrators to understand security decisions. Additional evaluation using diverse real-world IoT traffic, previously unseen attack variants, adversarial scenarios, and long-duration network deployments would provide stronger evidence of generalization and resilience. Overall, the integration of adaptive deep learning with intrusion detection and automated prevention provides a promising direction for developing intelligent IoT security systems capable of maintaining reliable protection as network conditions and cyber threats continue to evolve.

REFERENCES

1. Al-Haija, Qasem Abu, and Ayat Droos. "A Comprehensive Survey on Deep Learning-Based Intrusion Detection Systems in Internet of Things (IoT)." *Expert Systems*, vol. 42, no. 2, 2025, e13726.
2. Aldhaheeri, Alyazia, et al. "Deep Learning for Cyber Threat Detection in IoT Networks: A Review." *Internet of Things and Cyber-Physical Systems*, vol. 4, 2024, pp. 110–128.
3. Selem, Mehdi, Farah Jemili, and Ouajdi Korbaa. "Deep Learning for Intrusion Detection in IoT Networks." *Peer-to-Peer Networking and Applications*, vol. 18, 2025, article 22.
4. Hossain, Md. Alamgir. "Deep Learning-Based Intrusion Detection for IoT Networks: A Scalable and Efficient Approach." *Journal of Information Security*, vol. 2025, 2025, article 28.
5. Walling, Supongmen, and Sibesh Lodh. "An Extensive Review of Machine Learning and Deep Learning Techniques on Network Intrusion Detection for IoT." *Transactions on Emerging Telecommunications Technologies*, vol. 36, no. 2, 2025, e70064.
6. Belenguer, Aitor, Jose A. Pascual, and Javier Navaridas. "A Review of Federated Learning Applications in Intrusion Detection Systems." *Computer Networks*, vol. 258, 2025, article 111023.

7. Sorour, Shaymaa E., et al. "LSTM-JSO Framework for Privacy Preserving Adaptive Intrusion Detection in Federated IoT Networks." *Scientific Reports*, vol. 15, 2025, article 11321.
8. "A Novel Deep Learning-Based Intrusion Detection System for IoT DDoS Security." *Internet of Things*, vol. 28, 2024, article 101336.
9. "A Cognitive Security Framework for Detecting Intrusions in IoT and 5G Utilizing Deep Learning." *Computers & Security*, vol. 136, 2024, article 103560.
10. "Research on Intrusion Detection Method Based on Transformer and CNN-BiLSTM in Internet of Things." *Sensors*, vol. 25, no. 9, 2025, article 2725.
11. "A Multi-Class Intrusion Detection System for DDoS Attacks in IoT Networks Using Deep Learning and Transformers." *Sensors*, vol. 25, no. 15, 2025, article 4845.
12. "Transformer-Based Knowledge Distillation for Explainable Intrusion Detection System." *Computers & Security*, vol. 154, 2025, article 104417.
13. "Federated Learning-Based Intrusion Detection System for the Internet of Things Using Unsupervised and Supervised Deep Learning Models." *Cyber Security and Applications*, vol. 3, 2025, article 100068.
14. "Federated Learning for Intrusion Detection in IoT Environments: A Privacy-Preserving Strategy." *Discover Internet of Things*, vol. 5, 2025, article 72.
15. "Robust Intrusion Detection Based on Personalized Federated Learning for IoT Environment." *Computers & Security*, 2025, article 104442.
16. "FedMSE: Semi-Supervised Federated Learning Approach for IoT Network Intrusion Detection." *Computers & Security*, vol. 151, 2025, article 104337.
17. Chaurasia, Nisha, et al. "A Federated Learning Approach to Network Intrusion Detection Using Residual Networks in Industrial IoT Networks." *The Journal of Supercomputing*, vol. 80, 2024, pp. 18325–18346.
18. Nandanwar, Himanshu, and Rahul Katarya. "Deep Learning Enabled Intrusion Detection System for Industrial IoT Environment." *Expert Systems with Applications*, vol. 249, 2024, article 123808.
19. Liao, Han, et al. "A Survey of Deep Learning Technologies for Intrusion Detection in Internet of Things." *IEEE Access*, vol. 12, 2024, pp. 4745–4761.
20. Mazhar, T., et al. "Analysis of IoT Security Challenges and Its Solutions Using Artificial Intelligence." *Brain Sciences*, vol. 13, no. 4, 2023, article 683.
21. Khan, Noor Wali, et al. "A Hybrid Deep Learning-Based Intrusion Detection System for IoT Networks." *Mathematical Biosciences and Engineering*, vol. 20, no. 8, 2023, pp. 13491–13520.
22. Elnakib, O., et al. "EIDM: Deep Learning Model for IoT Intrusion Detection Systems." *The Journal of Supercomputing*, vol. 79, 2023, pp. 13241–13261.
23. Alosaimi, Shema, and Saad M. Almutairi. "An Intrusion Detection System Using BoT-IoT." *Applied Sciences*, vol. 13, no. 9, 2023, article 5427.

24. Vyas, Abhishek, et al. "Privacy-Preserving Federated Learning for Intrusion Detection in IoT Environments: A Survey." *IEEE Access*, vol. 12, 2024, pp. 127018–127050.
25. Ferrag, Mohamed Amine, et al. "Edge-IIoTset: A New Comprehensive Realistic Cyber Security Dataset of IoT and IIoT Applications for Centralized and Federated Learning." *IEEE Access*, vol. 10, 2022.
26. "OOA-Modified Bi-LSTM Network: An Effective Intrusion Detection Framework for IoT Systems." *Heliyon*, vol. 10, no. 8, 2024, article e29410.
27. "Two-Step Data Clustering for Improved Intrusion Detection System Using CICIoT2023 Dataset." *e-Prime Advances in Electrical Engineering, Electronics and Energy*, vol. 9, 2024, article 100673.
28. "Intrusion Detection Using Synaptic Intelligent Convolutional Neural Networks for Dynamic Internet of Things Environments." *Alexandria Engineering Journal*, vol. 111, 2025, pp. 78–91.
29. "Landscape of Learning Techniques for Intrusion Detection System in IoT: A Systematic Literature Review." *Computers & Electrical Engineering*, vol. 120, 2024, article 109725.
30. "A Survey on Intrusion Detection System in IoT Networks." *Cyber Security and Applications*, vol. 3, 2025, article 100082.