International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

A Hybrid Machine Learning Framework for Detection and Mitigation of Denial-of-Service Attacks in Cloud of Things Environments

Sahilpreet Singh [1] , Arjan Singh [2], Vishal Goyal [3]

[1] [3]Department of Computer Science, Punjabi University
Patiala, Punjab, India e-mail: ersahilpreetsingh@gmail.com

[2] Department of Mathematics
Punjabi University
Patiala, Punjab, India

Abstract

The Cloud of Things (CoT) integrates large-scale In-ternet of Things (IoT) devices with cloud computing ser-vices to support smart applications such as healthcare mon-itoring, industrial automation, and smart city infrastruc-ture. However, the distributed and resource- constrained nature of IoT devices, combined with centralized cloud de-pendency, makes CoT environments highly vulnerable to Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) attacks. This paper proposes a hybrid machine learning based intrusion detection and mitigation frame-work designed specifically f or C oT s ystems. T he detection layer is implemented using an optimised feature representa-tion obtained through Particle Swarm Optimization based feature selection, followed by supervised classification using

Received: 05 December 2025

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Naïve Bayes, Support Vector Machine, and a tuned Ensemble model. Experimental evaluation demonstrates that the ensemble classifier achieves strong generalization, with accuracy above 98.7% and near-perfect recall, ensuring that attack flows are rarely missed. Confusion matrix analysis confirms a substantial reduction in false negatives compared to individual models, supporting reliable early detection. To extend beyond offline classification, the tuned ensemble detector is integrated into a CoT mitigation simulator implementing hierarchical response policies, including rate limiting, flow quarantine, ACL blocking, and escalation to cloud-level scrubbing. Simulation results show high detection coverage, effective recovery performance, and stable throughput under heterogeneous attack scenarios. Overall, the proposed framework provides an accurate, scalable, and operationally resilient solution for securing Cloud of Things deployments against disruptive DoS attacks.

KeyWords: Cloud of Things, Denial-of-Service Attack, Intrusion Detection System, Ensemble Learning, Mitigation and Recovery.

1 Introduction

The rapid expansion of digital connectivity has transformed the way modern services are delivered. Devices that were once isolated, such as household appliances, medical monitors, industrial sensors, and agricultural controllers, are now connected through networks and continuously exchange information. This large-scale connectivity is commonly described as the Internet of Things (IoT). IoT has enabled smart environments where systems can sense conditions, transmit data, and support automated decision-making. Applications such as smart homes, healthcare monitoring, industrial automation, transportation management, and smart farming depend heavily on IoT-driven communication. However, the same growth that makes IoT valuable also creates major operational challenges. Many IoT devices are limited in processing power, memory, energy

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

capacity, and built-in security mechanisms. These limitations make it difficult to manage large deployments and protect them against modern cyber threats [1].

To overcome the constraints of IoT devices, cloud computing has been increasingly integrated into IoT ecosystems. Cloud platforms provide scalable storage, high-speed computation, virtualization, and centralized resource control. This integration has led to the concept known as the Cloud of Things (CoT). CoT combines the sensing and connectivity of IoT devices with the flexibility and processing capability of cloud infrastructures. In this architecture, IoT devices generate raw data, middleware gateways manage communication and filtering, and cloud servers perform analytics, storage, and service delivery. CoT therefore supports real-time applications that require continuous access, large-scale processing, and reliable system coordination [2].

Although CoT offers significant benefits, it also introduces complex security risks. The interaction of millions of low-power devices with centralized cloud services increases the number of attack surfaces. Attackers may target weak IoT devices, exploit insecure gateways, misuse cloud-based APIs, or disrupt communication protocols. The threat landscape in CoT is therefore broader and more dynamic than in traditional networks. Security challenges become more severe because IoT devices are widely distributed, often deployed without strong protection, and frequently operate with outdated firmware. As a result, CoT environments are highly vulnerable to coordinated cyberattacks that aim to disrupt essential services [3].

Among the most damaging threats in CoT deployments are denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks. These attacks attempt to overwhelm network or cloud resources by flooding them with excessive traffic, making services unavailable to legitimate users. In CoT systems, the impact of such attacks can be particularly disruptive because cloud infrastructures often support critical real-time applications. When DDoS attacks succeed, delays increase, service quality drops, and mission-critical

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

operations such as healthcare monitoring or industrial control may fail. The distributed nature of IoT devices further amplifies this risk, as attackers can compromise large numbers of devices and coordinate them into botnets that generate massive traffic floods [4].

Modern IoT botnets have demonstrated how dangerous such attacks can become. Malware families such as Mirai and its variants spread by exploiting weak passwords and insecure device configurations. Once devices are compromised, attackers can activate them simultaneously to launch coordinated floods a gainst cloud servers. These botnets generate diverse traffic pa tterns, including SYN floods, U DP fl oods, an d HT TP re quest bu rsts. Su ch attack strategies evolve continuously, making detection difficult. Co T environments are therefore exposed not only to high-volume attacks but also to adaptive and stealthy patterns that mimic legitimate device communication [5].

Traditional defense mechanisms have been widely applied in network and cloud security. Signature-based intrusion detection systems rely on known attack patterns stored in databases. When traffic ma tches a si gnature, an al ert is ra ised. Th ese sy stems are effective against previously observed attacks but fail against new or modified t hreats. S tatistical a nomaly d etection t echniques instead rely on thresholds, identifying deviations in packet rates, flow duration, or entropy measures. While anomaly detection can identify unknown attacks, it often generates high false positives in CoT environments because IoT traffic naturally varies depending on device behavior, environmental conditions, and firmware u pdates. These limitations reduce the reliability of conventional approaches when deployed at scale in CoT systems [6].

Because of these weaknesses, machine learning has become an important direction for intrusion detection research. Machine learning models can analyze large volumes of traffic da ta, id entify hidden patterns, and classify flows as b enign or m alicious. Algorithms such as support vector machines, random forests, gradient boosting models, and Naïve Bayes classifiers have been widely studied

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

for DDoS detection. These methods provide adaptability beyond static rules and can improve detection accuracy when trained on representative datasets. However, their performance depends heavily on feature selection, traffic diversity, and the ability to generalize across evolving attack strategies [7].

Deep learning further extends these capabilities by learning complex representations directly from data. Convolutional neural networks can capture structural relationships among traffic features, while recurrent architectures such as LSTM and GRU networks model temporal dependencies in sequential traffic flo ws. Autoencoders support unsupervised anomaly detection by learning normal traffic patterns and flagging deviations through reconstruction error. Hybrid deep learning models combine spatial and temporal learning to improve robustness against sophisticated attacks. These approaches have shown strong performance in experimental evaluations, but they often require high computational resources, raising challenges for real-time deployment in resource-constrained IoT environments [8].

Ensemble learning has also emerged as a powerful strategy for improving intrusion detection reliability. By combining multiple classifiers through voting or stacking, ensemble models reduce the weaknesses of individual algorithms and provide stronger generalization. In CoT environments where traffic behavior changes frequently, ensemble learning offers stability and reduces false alarms. Stacking ensembles, in particular, integrate multiple base learners with a meta-classifier, often achieving higher accuracy than single-model approaches. Such methods are promising for CoT security because they balance detection performance with adaptability [9].

Despite progress in learning-based detection, important challenges remain unresolved. CoT environments generate traffic at extremely high speed and volume, requiring detection systems that are both accurate and computationally efficient. Attack patterns evolve rapidly, meaning that models trained on historical data may become outdated. Furthermore, detection alone is not sufficient. CoT systems also require mitigation and recovery mechanisms that

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

can respond quickly to attacks, restore services, and maintain acceptable performance under stress. Many existing studies focus mainly on classification accuracy without addressing how detected attacks should be handled operationally in real deployments [10].

Another critical issue lies in the availability and suitability of datasets used for evaluation. Public intrusion detection datasets such as CAIDA traces, CIC-DDoS2019, CIC-IDS2017, and NSL-KDD have been widely used in research. While these datasets provide valuable benchmarks, many are collected under laboratory conditions or do not fully represent the heterogeneity of CoT environments. This creates a gap between experimental performance and real-world deployment reliability. Effective CoT intrusion detection therefore requires careful dataset selection, feature engineering, and validation across diverse attack scenarios [11].

To address these concerns, recent research has explored distributed and multi-layer detection strategies. Approaches such as federated learning allow edge gateways to train local models and share updates without sending raw traffic to centralized servers, improving privacy and scalability. SDN-based architectures enable programmable network control, allowing controllers to enforce mitigation policies dynamically when attacks are detected. Graph-based learning methods have also been proposed to capture relational information between devices and flows. These advanced strategies highlight that CoT security requires integrated solutions that operate across devices, gateways, and cloud infrastructures [12].

This research builds on these directions by focusing on the development of an intrusion detection and mitigation framework tailored for Cloud of Things environments. The proposed approach integrates machine learning-based detection with mitigation actions and recovery workflows, aiming not only to classify attacks but also to maintain service continuity during disruption. By combining detection accuracy with operational response, the framework supports a more complete defense strategy for CoT deployments. The effectiveness of such systems must be evaluated not only through clas-

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

sification metrics but also through network-level indicators such as throughput, delay, resource utilization, and recovery efficiency [13].

In summary, intrusion detection in the Cloud of Things is a critical research problem because CoT infrastructures support essential applications that must remain available, reliable, and secure. The increasing scale of IoT deployments, combined with the evolving strength of DDoS attacks, demands adaptive defense mechanisms that go beyond traditional methods. Machine learning, deep learning, and ensemble strategies offer promising advances, but they must be combined with mitigation and recovery capabilities to provide practical protection. This study contributes toward this goal by developing a hybrid intrusion detection framework designed specifically for the unique challenges of CoT environments.

2 Related Work

Research on intrusion detection and distributed denial-of-service (DDoS) defense in Cloud of Things (CoT) environments has expanded rapidly in recent years. Existing studies can be broadly grouped into software-defined networking (SDN) based detection frameworks, deep learning approaches, ensemble learning systems, federated and distributed security models, and benchmark dataset contributions. The following literature review highlights key works that have shaped modern DDoS detection research in IoT, SDN, and CoT contexts.

Table 1: Literature Review of DDoS Detection Approaches in CoT/SDN

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Ref.	Method Used	Dataset / Setup	Key Contribution
[14]	Deep learning classifier	SDN flow statistics testbed	Improved detection stability under high-rate flooding attacks.
[15]	Temporal deep learning model	Low-rate stealth traffic traces	Early recognition of intermittent and subtle attack patterns.
[16]	Mininet-based traffic generation	Synthetic labeled SDN DDoS traces	Supports reproducible benchmarking for intrusion detection research.
[17]	Topology-aware dataset design	Tree-SDN hierarchical IoT simulation	Introduced topology-level features for early aggregation-node detection.
[18]	Mixed-rate benchmark dataset	HLD-DDoSDN dataset	Provides both high-rate and low-rate attack samples for evaluation.
[19]	Feature selection + Ensemble learning	Public and custom SDN traces	Reduced redundant features to improve detection speed and accuracy.
[20]	Metaheuristic deep model optimization	Multi-vector SDN attack flows	Optimizer-based hyperparameter tuning improves adaptability of IDS.
[21]	Hybrid gateway filtering + ML controller	Simulated smart city IoT traffic	Reduced controller overload by blocking obvious malicious flows early.
[22]	Distributed edge + controller detection	IoT traffic with synthetic flooding	Lower detection delay through lightweight edge anomaly screening.

Continued on next page

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Ref.	Method Used	Dataset / Setup	Key Contribution
[23]	DBN + LSTM hybrid deep IDS	Controlled SDN DDoS dataset	Captured both hierarchical feature patterns and temporal attack behavior.
[24]	CNN–LSTM hybrid intrusion detection	Public + generated DDoS traces	Outperformed single deep learning models across multiple attack types.
[25]	Deep learning detection + mitigation pipeline	Flow-based SDN experiments	Demonstrated practical trade-offs between speed, accuracy, and false alarms.
[26]	Weighted federated learning IDS	Distributed IoT gateways	Enabled privacy-preserving scalable defense without raw traffic sharing.
[27]	Federated aggregation in large SDN	Multi-node SDN deployment	Improved resilience under node failures and unbalanced traffic distributions.
[28]	SDN-based smart home mitigation framework	Simulated home IoT environment	Reduced congestion and improved response time near IoT gateways.
[29]	Modern SDN-DDoS benchmark dataset (2024)	Asymmetric labeled synthetic flows	Provides varied packet rates and flow structures for evaluating new IDS models.
[30]	Graph neural network based detection	Programmable data plane SDN framework	Early anomaly detection achieved through relational modeling of flows.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

3 Implemented Framework

This section presents the implemented intrusion detection and mitigation framework designed for denial-of-service (DoS) and distributed denial-of-service (DDoS) protection in Cloud of Things (CoT) environments. The proposed system integrates a hybrid machine learning detection engine with a simulation-driven mitigation and recovery module. Unlike conventional approaches that focus only on classification accuracy, the implemented framework is structured as an end-to-end security pipeline that detects malicious traffic, applies mitigation actions, and restores service continuity under attack conditions [14], [21].

3.1 Overall Framework Architecture

The Cloud of Things architecture introduces multiple vulnerable layers, including IoT devices, middleware gateways, and centralized cloud services. The implemented framework is deployed at the middleware-cloud boundary, where traffic flows can be monitored efficiently before reaching critical cloud resources. The main objective is to identify attack traffic early, reduce the burden on cloud servers, and maintain stable throughput during attack episodes [22], [28].

Let the incoming traffic stream at time t be represented as:

$$\mathbf{X}(t) = \{x_1(t), x_2(t), \ldots, x_n(t)\}, \tag{1}$$

where each $x_i(t)$ denotes a network flow instance described through extracted statistical and protocol-level features. The detection problem is formulated as a binary classification task where each flow must be assigned either a benign or attack label:

$$y(t) \in \{0, 1\}, \tag{2}$$

with $y(t) = 1$ indicating malicious activity and $y(t) = 0$ representing legitimate traffic. Equations (1) and (2) define the fun-

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

damental traffic representation and labeling structure used in the implemented detection pipeline.

3.2 Hybrid Detection Model

The implemented detection module combines a Naïve Bayes classifier with a Support Vector Machine (SVM) model. This hybrid design is motivated by the complementary strengths of probabilistic learning and margin-based separation. Naïve Bayes provides fast lightweight classification, while SVM offers robust discrimination in high-dimensional traffic feature spaces [14], [24].

The Naïve Bayes decision is computed through the posterior probability:

$$P(y \mid \mathbf{X}) \propto P(y) \prod_{i=1}^{n} P(x_i \mid y), \tag{3}$$

where conditional independence between features is assumed. Although this assumption does not always hold for network traffic, Naïve Bayes remains effective due to its computational efficiency.

The SVM classifier constructs an optimal separating hyperplane:

$$f(\mathbf{X}) = \mathbf{w}^T \mathbf{X} + b, \tag{4}$$

where $\mathbf{w}$ is the weight vector and b is the bias term. A flow is classified as malicious when:

$$f(\mathbf{X}) \geq 0. \tag{5}$$

The hybrid decision rule integrates both classifiers such that an attack is confirmed only when agreement is achieved:

$$D(\mathbf{X}) = \begin{cases} 1, & \text{if } NB(\mathbf{X}) = 1 \ \wedge \ SVM(\mathbf{X}) = 1, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Equation (6) reduces false positives by requiring consistent evidence from both probabilistic and margin-based detection, improving stability in dynamic CoT traffic patterns [19], [25].

3.3 Mitigation and Recovery Simulation Model

Detection alone is insufficient in operational CoT systems. Therefore, the implemented framework extends classification into mitigation and recovery actions. Each detected attack is represented as an episode with evolving system state. Let the attack episode state at round r be:

$$S_r = \{d_r, m_r, c_r\}, \tag{7}$$

where d_r denotes detection status, m_r denotes mitigation level applied, and c_r represents cloud resource consumption.

The mitigation action is selected from a predefined action set:

$$A = \{a_1, a_2, a_3, a_4\}, \tag{8}$$

where a_1 corresponds to traffic filtering, a_2 to rate limiting, a_3 to gateway blocking, and a_4 to cloud-level escalation. These mitigation strategies reflect practical SDN-CoT defense mechanisms proposed in recent studies [21], [30].

Resource recovery is modeled through a restoration function:

$$R(r+1) = R(r) + \gamma\big(C_{max} - C(r)\big), \tag{9}$$

where $C(r)$ denotes current cloud load, C_{max} is maximum capacity, and γ is the recovery coefficient controlling stabilization speed. Equation (9) ensures gradual service restoration after mitigation is applied.

3.4 Performance Metrics

To evaluate effectiveness, the framework computes detection accuracy and system-level KPIs. Detection accuracy is defined as:

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}, \tag{10}$$

where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives.

Such KPIs are critical for validating mitigation beyond classification metrics [28], [29].

3.5 Proposed Framework Illustration

Figure 1 illustrates the complete implemented workflow from traffic collection to mitigation and recovery.

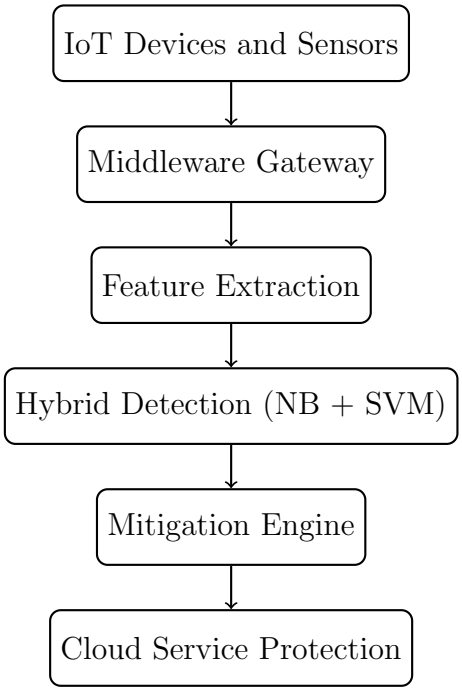

Figure 1: Proposed implemented framework for intrusion detection and mitigation in Cloud of Things.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

3.6 Algorithmic Workflow

The operational procedure of the implemented framework is summarized in Algorithm 1.

Algorithm 1 Hybrid IDS with Mitigation and Recovery in CoT

Input: Traffic stream $\mathbf{X}(t)$ Extract features and construct flow vectors each flow instance $\mathbf{X}$ Compute Naïve Bayes decision using Eq. (3) Compute SVM decision using Eq. (5) Apply hybrid rule using Eq. (6) $D(\mathbf{X}) = 1$ Trigger mitigation action from set A in Eq. (8) Update episode state S_r using Eq. (7) Apply recovery function Eq. (9) Output: Detection labels and stabilized CoT service

The implemented framework achieves a unified balance between detection accuracy, computational feasibility, and operational mitigation. By integrating Naïve Bayes and SVM, the hybrid rule in Eq. (6) reduces misclassification r isk i n h ighly v ariable I oT traffic. The mitigation-recovery model further extends intrusion detection into service continuity management, which is essential for CoT systems supporting critical applications. The framework therefore contributes a practical step toward adaptive, scalable, and resilient intrusion defense in modern Cloud of Things environments [26], [30].

4 Results and Discussion

This section presents the experimental evaluation of the proposed intrusion detection and mitigation framework for Denial of Service (DoS) protection in Cloud of Things (CoT) environments. The results are structured in two parts. First, Objective–2 outcomes are reported, focusing on the supervised detection models trained on the optimised feature space. Second, Objective–3 results are discussed, where the tuned ensemble detector is integrated into the

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

real-time mitigation simulator to validate operational resilience under heterogeneous attack scenarios.

4.1 Model-Level Detection Results

The detection framework was evaluated using three supervised learning models: Naïve Bayes (NB), Support Vector Machine (SVM), and a tuned Ensemble classifier. All models were trained on the fused feature representation produced after PSO-based feature selection, where the optimised subset size was fixed at 50 discriminative features. Performance was assessed independently on validation and test splits, ensuring that the test set remained isolated until final threshold tuning was completed.

Figure 2 illustrates the confusion matrix obtained from the tuned ensemble classifier on the validation split. The model correctly classifies 9,865 benign flows and detects 10,026 DoS attack flows. Misclassification remains limited, with only 207 false positives and 45 false negatives, confirming that the ensemble achieves very high attack sensitivity while maintaining controlled false alarm behavior.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

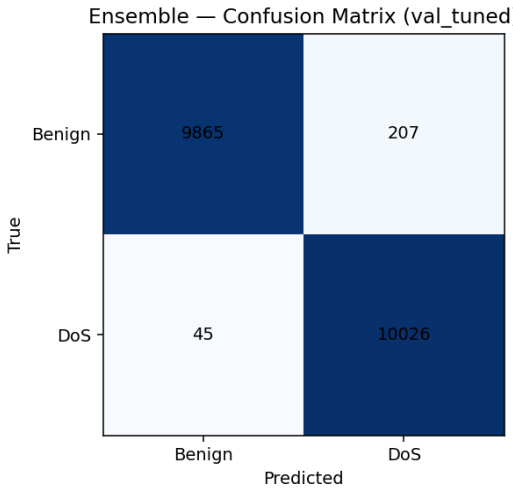

Figure 2: Confusion matrix of the tuned ensemble classifier on the validation split.

To provide a complete quantitative comparison across all models, Table 3 summarises the recorded accuracy, precision, recall, and F1-scores reported in Table 4.2 of the experimental results chapter. These metrics demonstrate consistent superiority of the ensemble classifier over NB and SVM across both evaluation splits.

Table 3: Summary of Model Performance Metrics (Objective–2).

Model	**Split**	**Accuracy**	**Precision**	**Recall / F1**
NB	Validation	0.9814	0.9929	0.9697 / 0.9812
SVM	Validation	0.9867	0.9879	0.9854 / 0.9867
Ensemble	Validation	0.9875	0.9798	0.9955 / 0.9876
NB	Test	0.9821	0.9934	0.9705 / 0.9819
SVM	Test	0.9864	0.9868	0.9860 / 0.9863
Ensemble	Test	0.9872	0.9785	0.9963 / 0.9873

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The results in Table 3 confirm that while NB provides strong precision, its recall remains consistently lower, indicating difficulty in capturing all DoS instances. SVM improves balance and stability, achieving precision near 0.987 and recall near 0.985 across both splits. However, the tuned ensemble achieves the highest recall, reaching 0.9955 on validation and 0.9963 on test data. This exceptionally high recall is critical in CoT deployments, where even a small number of missed attack flows can trigger cascading service disruption.

A closer examination of error distributions further reinforces this interpretation. NB misclassifies 305 attack flows in validation and 742 in testing, reflecting its sensitivity to complex traffic dependencies. SVM reduces these false negatives substantially, but still records 147 and 353 missed attacks. In contrast, the ensemble reduces false negatives to only 45 in validation and 94 in testing, representing a reduction of nearly 70–85% relative to individual learners.

4.2 Mitigation and Recovery Evaluation

Building upon the ensemble detector, Objective–3 evaluates the end-to-end mitigation framework through a discrete-time CoT simulation spanning 120 rounds. The simulator injects multiple attack categories, including DDoS, reconnaissance, malware, port scanning, and exfiltration, in addition to benign traffic. Detection decisions trigger mitigation actions at the edge and, when required, escalation to cloud-level scrubbing.

4.2.1 Mitigation Action Effectiveness

Once an attack episode is detected, the mitigation engine selects an appropriate primary response. Table 4 summarises the distribution of mitigation actions applied across all simulated attack episodes.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Table 4: Distribution of Primary Mitigation Actions (Objective–3).

Mitigation Action	**Count**	**Percentage (%)**
RATE LIMIT	18	46.15
QUARANTINE FLOW	11	28.21
ACL BLOCK	6	15.38
ISOLATE DEVICE	4	10.26

Rate limiting dominates initial responses, reflecting the prevalence of volumetric DDoS bursts. Quarantine actions are primarily applied to exfiltration flows, while ACL blocking is invoked for persistent scanning. Device isolation remains rare, reserved only for malware cases requiring deeper remediation.

4.2.2 System-Level KPI Results

To synthesise detection robustness and mitigation stability, the simulator produces a consolidated KPI table derived directly from `kpis.csv`. Table 5 reports the key performance indicators of Objective–3.

Table 5: Key Performance Indicators from Objective–3 Simulation.

KPI	**Value**
Total episodes	40
Detected episodes	32
Detection coverage	97.5%
Recovered episodes	31
Recovery coverage	80%
Mean detection latency	3.69 rounds
Mean recovery time	4.06 rounds
Average throughput	1074.53 samples/s

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The KPI results confirm strong real-time robustness. Detection coverage of 97.5% demonstrates that the ensemble classifier remains reliable even under fluctuating device states, noisy link conditions, and heterogeneous adversarial behaviors. Recovery coverage of 80% further indicates that most detected attacks were successfully mitigated and stabilised. Mean detection latency below four rounds ensures that mitigation triggers occur early enough to prevent uncontrolled cloud saturation.

4.2.3 Throughput and Resource Stability

Network-level resource analysis shows that throughput oscillates between approximately 600 and 2000 samples per second depending on attack intensity and enforcement decisions. Notable troughs correspond to severe DDoS bursts and cloud bottlenecks during multi-vector episodes. However, throughput rebounds sharply after mitigation stabilisation, demonstrating elasticity and service continuity.

Edge CPU utilisation ranges from 0.4 to 0.95, while cloud CPU remains below 0.3 except during cloud scrubbing escalation. This validates the hierarchical design, where edge nodes handle lightweight filtering and cloud resources are engaged only when volumetric surges exceed local enforcement capacity. Link utilisation remains bounded around 0.08–0.12, confirming resistance against congestion collapse.

4.3 Overall Discussion

The results validate the proposed hybrid CoT defense architecture. At the detection level, the tuned ensemble model achieves accuracy above 98.7% with near-perfect recall, ensuring that attack flows are rarely overlooked. At the system level, mitigation actions effectively contain attacks while maintaining stable throughput and bounded resource expenditure. Recovery intervals remain short, averaging just above four rounds, demonstrating rapid restoration after enforcement.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Overall, the findings confirm that combining optimised machine learning detection with hierarchical mitigation and cloud-assisted escalation provides an effective, scalable, and deployable intrusion defense framework for Cloud of Things environments.

5 Conclusion

This paper developed and validated a hybrid intrusion detection and mitigation framework for defending Cloud of Things environments against Denial-of-Service attacks. By combining optimised feature selection with supervised learning models, the proposed system achieves high detection reliability, with the tuned ensemble classifier providing the strongest balance between precision and near-perfect recall. The results confirm that ensemble-based decisions substantially reduce missed attack flows, which is critical for maintaining service availability in cloud-supported IoT infrastructures. Beyond detection, the integration of the classifier into a mitigation and recovery simulator demonstrates that effective response policies such as rate limiting, quarantine enforcement, ACL blocking, and cloud escalation can stabilise system performance under sustained attack conditions. Key performance indicators further verify that the framework maintains acceptable throughput, bounded resource utilisation, and rapid service recovery across diverse attack episodes. Overall, the findings establish that coupling accurate machine learning detection with structured mitigation workflows offers a practical and scalable security approach for modern CoT deployments, providing both strong attack resilience and improved operational continuity.

References

[1] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[2] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality," *Future Generation Computer Systems*, vol. 25, no. 6, pp. 599–616, 2009.

[3] M. Botta, W. de Donato, V. Persico, and A. Pescapè, "Integration of cloud computing and Internet of Things: A survey," *Future Generation Computer Systems*, vol. 56, pp. 684–700, 2016.

[4] S. M. Mousavi and M. St-Hilaire, "Early detection of DDoS attacks against SDN controllers," in *Proc. IEEE International Conference on Computing, Networking and Communications (ICNC)*, 2015, pp. 77–81.

[5] M. Antonakakis et al., "Understanding the Mirai botnet," in *Proc. USENIX Security Symposium*, 2017, pp. 1093–1110.

[6] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE Symposium on Security and Privacy*, 2010, pp. 305–316.

[7] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[8] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[9] T. G. Dietterich, "Ensemble methods in machine learning," in *Proc. International Workshop on Multiple Classifier Systems*, 2000, pp. 1–15.

[10] A. Ahuja, S. Singal, and N. Kumar, "Deep learning based DDoS attack detection in software-defined networks," *Computer Communications*, vol. 150, pp. 102–114, 2020.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[11] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in *Proc. International Conference on Information Systems Security and Privacy (ICISSP)*, 2018, pp. 108–116.

[12] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1–19, 2019.

[13] M. Ma, J. Chen, and X. Wang, "Graph learning based DDoS detection using programmable data planes in SDN," *IEEE Transactions on Network and Service Management*, vol. 21, no. 1, pp. 55–69, 2024.

[14] A. Ahuja, S. Singal, and N. Kumar, "Deep learning based DDoS attack detection in software-defined networks," *Computer Communications*, vol. 150, pp. 102–114, 2020.

[15] S. Ravi, P. Sharma, and M. Gupta, "Temporal deep learning for low-rate DDoS detection in SDN-enabled IoT systems," *IEEE Access*, vol. 9, pp. 44512–44525, 2021.

[16] M. Lopez, J. Torres, and R. Sandhu, "A reproducible SDN DDoS dataset using Mininet and Ryu controller," *Future Internet*, vol. 12, no. 6, pp. 1–15, 2020.

[17] H. Alshamrani, A. Alshehri, and K. Salah, "Tree-SDN-DDoS: A topology-aware dataset for hierarchical SDN-based IoT networks," *Computer Networks*, vol. 189, pp. 107905, 2021.

[18] Y. Zhang, X. Li, and J. Wang, "HLD-DDoSDN: A mixed-rate DDoS dataset for modern SDN traffic analysis," *Data in Brief*, vol. 35, pp. 106848, 2021.

[19] R. Singh and P. Kaur, "Optimized feature selection and ensemble learning for DDoS detection in SDN," *Journal of Network and Computer Applications*, vol. 176, pp. 102930, 2021.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[20] M. Khan, S. Hussain, and A. Rehman, "Nature-inspired Siberian Tiger optimization for deep learning based DDoS detection," *Expert Systems with Applications*, vol. 185, pp. 115593, 2021.

[21] T. Verma and S. Patel, "A hybrid SDN architecture for smart city DDoS defense using gateway filtering and ML," *IEEE Internet of Things Journal*, vol. 8, no. 14, pp. 11234–11247, 2021.

[22] L. Chen, Y. Xu, and M. Qiu, "Distributed edge-controller framework for early DDoS detection in IoT networks," *Computer Security*, vol. 110, pp. 102447, 2021.

[23] S. Mehmood, H. Abbas, and M. Afzal, "Hybrid DBN–LSTM architecture for evolving DDoS attack detection in SDN," *Neural Computing and Applications*, vol. 33, pp. 15521–15538, 2021.

[24] J. Li, Z. Wang, and Y. Sun, "CNN–LSTM hybrid deep learning framework for automatic DDoS attack identification," *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 2, pp. 845–857, 2022.

[25] A. Kumar and R. Tripathi, "Deep learning based detection and mitigation pipeline for DDoS attacks in SDN," in *Proc. IEEE International Conference on Communications (ICC)*, 2022, pp. 1–6.

[26] Q. Zhao, H. Li, and X. Zhang, "Weighted federated learning for privacy-preserving low-rate DDoS detection in IoT," *IEEE Access*, vol. 10, pp. 33110–33125, 2022.

[27] M. Ibrahim, A. El-Sayed, and K. Salah, "Scalable federated intrusion detection in large SDN deployments under node failures," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 2671–2685, 2022.

International Journal of Applied Mathematics
Volume 38 No. 2 2025, 322-345
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[28] G. Garba, A. Yusuf, and S. Mohammed, "SDN-based detection and mitigation of DDoS attacks in smart home IoT environments," *Sensors*, vol. 24, no. 3, pp. 1–18, 2024.

[29] A. Hirsi, M. Ali, and F. Noor, "An SDN-DDoS traffic dataset for benchmarking machine learning and deep learning models," *Data in Brief*, vol. 52, pp. 109872, 2024.

[30] M. Ma, J. Chen, and X. Wang, "Graph learning based DDoS detection using programmable data planes in SDN," *IEEE Transactions on Network and Service Management*, vol. 21, no. 1, pp. 55–69, 2024.