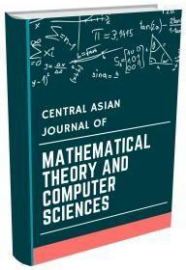

CENTRAL ASIAN JOURNAL OF MATHEMATICAL THEORY AND COMPUTER SCIENCES
https://cajmtcs.centralasianstudies.org/index.php/CAJMTCS
Volume: 07 Issue: 02 | April 2026 ISSN: 2660-5309

Article

Adaptive Isolation Forest and Cascading XGBoost–DNN Ensemble for High-Confidence Network Intrusion Detection with Explainable AI Integration

Khaldoon H. Al-hussauni[*1]

1. Cybersecurity Department, Information Technology College, University of Babylon, Babil, Iraq
* Correspondence: alhussyni@uobabylon.edu.iq

Citation: Al-hussauni K. H. Adaptive Isolation Forest and Cascading XGBoost–DNN Ensemble for High-Confidence Network Intrusion Detection with Explainable AI Integration. Central Asian Journal of Mathematical Theory and Computer Sciences 2026, 7(2), 265-276.

Received: 10th Jan 2026
Revised: 20th Feb 2026
Accepted: 29th Mar 2026
Published: 10th Apr 2026

Abstract: The growing use of Internet of Things infrastructure underscores the importance of advanced threat detection systems that can handle evolving attack vectors. This paper presents a confidence-based framework for cascading information that employs self-adjusting irregularity recognition and supervised learning to enhance the accuracy and efficacy of IoT security. The proposed method starts with an adaptive isolation forest with dynamic thresholding to reduce false positives. Then, it uses a cascading architecture that includes an XGBoost for quick learning and a deep neural network for deep feature extraction. To make the model easy to understand and to build trust and transparency among analysts, the framework is based on the Explainable Artificial Intelligence (XAI) approach, SHAP. Evaluation of the CICIoT2023 dataset demonstrates the potential of the framework for deployment in a real-world IoT environment and closes the gap in anomaly detection accuracy, computational efficiency, and interpretability, achieving significant performance improvements at fixed thresholds and across standard accuracy metrics, including Precision, Recall, and F1 score. 0.9948, 0.9947, 0.9948, and 0.9945, respectively, while holding the comparative prediction time.

Keywords: Internet of Things, Adaptive Isolation Forest, XGBoost, Deep Neural Network, Network Intrusion Detection, Explainable AI.

1. Introduction

Over the last few decades, information and communications technologies (ICT) have advanced significantly. This has led to a significant increase in the volume and complexity of network data, as well as the cybersecurity risks that accompany it. But traditional security measures like firewalls, access control systems, and identity and verification systems don't always work well against advanced attackers [1]. Even proactive security measures, such as penetration testing, are needed to identify weaknesses before they are exploited. This shows how important it is to have real-time anomaly detection to keep modern information systems safe and sound [2], which shows how important it is to use real-time anomaly detection methods to keep modern information systems safe and sound. Network anomalies, which can be seen as strange data patterns or sudden spikes in traffic, often happen before cyberattacks or infrastructure failures. These anomalies must be identified immediately to prevent the system from being paralyzed. So, it's very important for IDS to move away from static rule-based systems and toward more advanced machine

learning (ML) methods that can find "zero-day" threats[3]. This threat needs to change quickly because Distributed Denial-of-Service (DDoS) attacks are now one of the most common types of attacks, growing by more than 90% each year [4]. This number shows that attackers are still using volumetric strategies because they work, can be deployed at scale, and are cheaper than more complex methods of breaking in. Nowadays, machine learning algorithms are the most efficient means of identifying malicious content in network flows; however, there is a distinction between supervised and unsupervised approaches, which employ labeled and unlabeled data, respectively. Large sets of labeled data are used to train supervised techniques, which is a major bottleneck because labeling is costly and challenging to keep up with changing attack signatures. Unlabeled data from regular traffic is used to train unsupervised models that can identify anomalous behaviors without being aware of the types of attacks [5]. Standard unsupervised algorithms, such as the Isolation Forest, classify anomalies using globally fixed, static thresholds (e.g., a fixed contamination parameter). Recent studies [6] claim that static thresholds do not account for data and class heterogeneity, rendering them unable to generalize well. In addition, fixed high thresholds tend to reject "hard" but informative samples, thereby reducing learning efficiency for difficult classes. Furthermore, static cut-offs cannot account for the evolving embedding space during online learning, leading to "noisy" pseudo-labels and high false alarm rates in dynamic environments such as IoT networks [7], [8]. While DL models provide higher detection accuracy, they are expensive for resource-constrained IoT devices, which can quickly saturate CPU, memory, and battery capacity[9], making them impractical for edge deployment [7], [10]. Hence, a framework that combines the high accuracy of deep models with the efficiency needed by edge devices is needed.

This paper proposes a Confidence-Based Cascading Ensemble Framework, which employs two innovations: an Adaptive Isolation Forest that self-tunes the thresholding mechanism to data distribution, eliminating static anomaly scoring and reducing false positives, and a Cascading Architecture that combines a lightweight gradient-boosting model (XGBoost) for fast, high-confidence detection and conditionally delegates low-confidence, ambiguous traffic to a deep neural network (Expert Model).

Finally, to bridge the gap between AI decision-making and security operations, this study integrates an Explainable AI (XAI) module using SHAP, providing Security Operations Center (SOC) analysts with transparent insights into why specific flows were flagged as malicious.

2. Materials and Methods

2.1 Related Work

Intrusion Detection Systems (IDS) for the Internet of Things (IoT) have evolved from statistical methods to more advanced Deep Learning (DL) and Ensemble Learning architectures. However, there are still gaps in dynamic adaptability and computational efficiency in centralized cloud environments. In this section, three major dimensions of literature have been reviewed: Deep Ensemble IDSs, Adaptive Thresholding Mechanisms, and Resource-Efficient Deployment.

A popular method for improving robustness and reducing generalization error is ensemble learning. StaEn-IDS, a stacking ensemble of CNN and LSTM base learners with a Random Forest meta-classifier, created by Vishwakarma and Kesswani [11], outperformed other models in identifying minority attack classes on the CIC IoT 2022 dataset. Similarly, Sharma et al. used Transformers, LSTMs, and CNNs to develop an Ensemble Deep Learning (EDL) framework that outperformed single models in capturing a range of spatial and temporal traffic patterns.

Other, more advanced fusion strategies have also been proposed, such as the feature-aware adaptive ensemble proposed by Laraig et al. [12] which weights models dynamically according to specific feature patterns, leading to substantial reductions in False Positive Rates (FPR). However, a major limitation in these works is that they rely on static decision thresholds, with models that explicitly address class imbalance using cost-sensitive learning CSAE [13] and ensembles with fixed global thresholds (e.g., softmax probability > 0.5) for anomaly classification, which is insensitive to the dynamic concept drift present in IoT environments, and therefore will degrade over time.

Adejoh et al. [7] showed that dynamic thresholds in unsupervised fraud detection dramatically reduce false alarms because the thresholds adjust to cluster-specific risk levels; Wang et al. [14]and Liang et al. used class-adaptive and self-adaptive thresholding in semi-supervised learning to enhance the reliability of pseudo-labels, and Wei et al. [8] Applied dynamic distance thresholds in domain adaptation to filter reliable samples across multiple tasks. Although many studies have achieved success using calibrated, time-varying anomaly scores (e.g., rolling quantiles or per-device baselines) [12], the systematic integration of these scores into Deep Ensemble IoT IDSs remains a research gap.

Simioni et al. [15] proposed an energy-efficient offloading scheme where heavy DNN inference is delegated to the cloud/edge only when necessary. Alashjaee and Alqahtani [16] examined hybrid pipelines that incorporate lightweight XGBoost with Deep Feed-Forward Neural Networks (FFNN) to find a balance between speed and accuracy. However, these approaches are limited because they usually do not implement intelligent, confidence-based cascading, in which the system decides, per sample, whether to pay the cost of deep inference.

Existing work has demonstrated that Deep Ensembles achieve high accuracy [11], [12], [17], and that Explicit Adaptive Thresholding enhances the robustness [7], [14], but no work has combined both paradigms for IoT security. To our knowledge, there are no centralized IDS frameworks that employ Deep Ensembles for accuracy and Explicit Adaptive Thresholding for false alarm reduction on standard datasets (e.g., CICIoT2023), and this is one of the gaps identified in a recent comprehensive review of the field [12]. The objective of this research is to address this gap by proposing a Confidence-Based Cascading Ensemble (CBCE) with a self-adaptive thresholding mechanism.

2.2 Methodology

2.2.1 Isolation Forest

Isolation Forest (IForest) is an anomaly detection algorithm based on the idea that outliers are both rare and different from normal data, making them easy to separate through iterative data partitioning. IForest does not require these calculations, achieving linear runtime with respect to dataset size. The method uses a collection of random and independent isolation trees (iTrees) created with two main parameters: ψ (the sub-sample size) and t (the number of trees) and works in two phases—a training phase for building the forest and a scoring phase where an anomaly score is calculated for each point. An example with anomalous data (X0) and regular data (Xi) is presented in Figure 1. This architecture allows IForest to have low memory usage while providing high scalability, making it applicable to both large-scale and real-time processing needs [5].

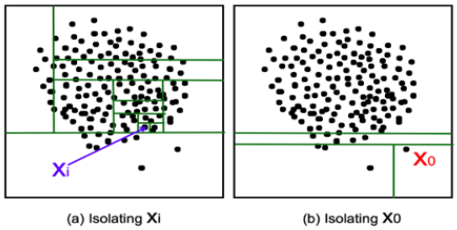

Figure 1. Anomalies (X0) are isolated faster than normal data (Xi) [5].

2.2.2 XGBoost algorithm

In 2016, Chen and Guestrin introduced the XGBoost algorithm, a tree-boosting system that scales. Thanks to parallel, distributed, out-of-core, and cache-aware computing, this algorithm achieves speeds over 10 times faster than those of popular ML and DL models. Another benefit of this algorithm is its strong optimization and scalability; as a result, it can be readily applied to processing billions of examples under distributed or memory-limited conditions. This modern use of gradient boosting machines was developed specifically to address real-world problems where input data is sparse. The algorithm accounts for missing values, too-frequent zero values in the dataset, and results from applied feature engineering techniques. Ensemble methods add new models recursively until no further significant improvement is possible by adding more models to the existing ones. The model's loss function is minimized via gradient descent [18].

2.2.3 Deep Neural Network

An artificial neural network is the first step toward a new paradigm of artificial intelligence and has been described as such. Researchers have been trying for years to improve the technology used in ANN computing with successive attempts. These never-ceasing attempts culminated in a principle known as deep learning. A deep neural network (DNN) architecture consists of multiple hidden layers of a feed-forward ANN [19]. Figure 2 shows the general architecture of the deep neural network, consisting of multiple hidden layers.

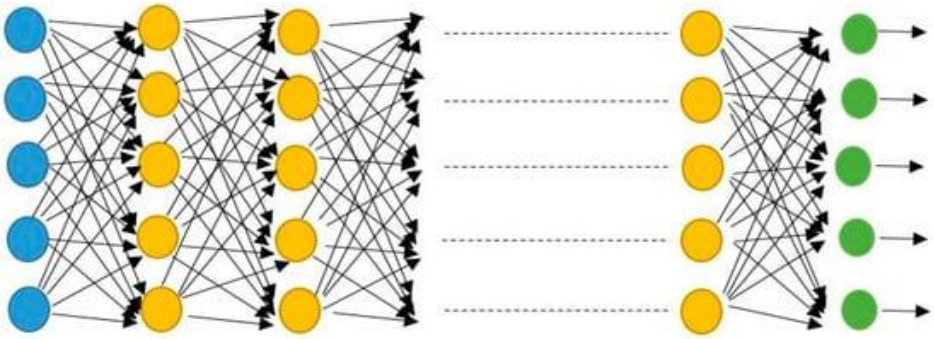

Figure 2. Architecture of a Deep Neural Network.

2.2.4 Data Preprocessing and Feature Engineering

The efficacy of both unsupervised anomaly detection and supervised deep learning relies heavily on the quality of the input data. As illustrated in the system workflow (see Figure 3), the "Data Input" stage is immediately followed by rigorous preprocessing to address the heterogeneity and high dimensionality of IoT network traffic. This study utilized the CICIoT2023 dataset, a comprehensive benchmark for modern large-scale IoT attacks [20].

Since this is a very large dataset (CICIoT2023), it is split into several CSV fragments, which load as follows. Observations with infinite flow capture values were discarded from the dataset, as gradient explosion would cause the DNN to overfit; observations were also omitted to avoid overfitting artifacts, but rather to learn general traffic patterns.

Feature relevance varies across raw network packets. In the edge-deployment scenario, features were manually pruned according to domain knowledge and correlation analysis to optimize computational efficiency. In the edge-deployment scenario, a subset of attributes was excluded to optimize the feature space by eliminating irrelevant or redundant variables; this dimensionality reduction step was performed, and the resulting

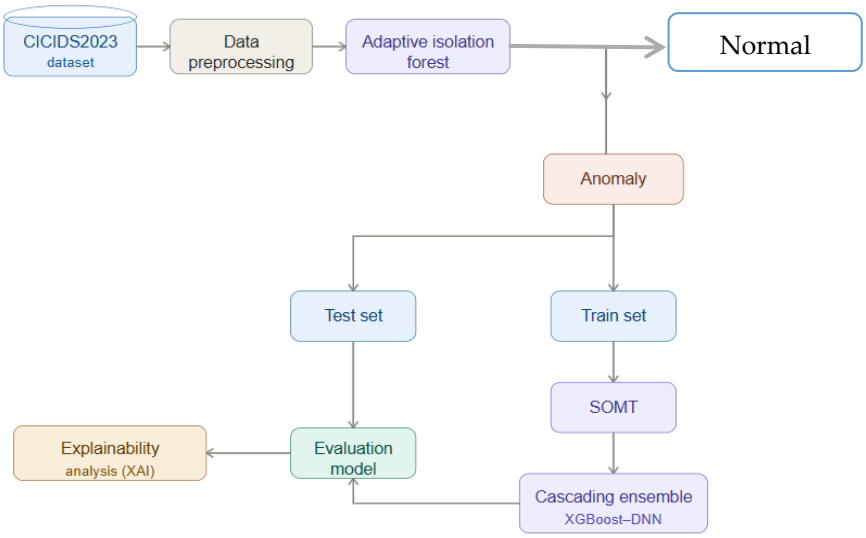

Figure 3. The Proposed Model.

The dataset structure was verified to ensure data consistency before model training. The categorical features were converted to numerical representations via Label Encoding, which is necessary for the XGBoost and DNN components that require numerical input vectors.

One of the main problems in the CICIoT2023 dataset (and intrusion detection in general) is the extreme class imbalance, where benign traffic and common attacks (such as DDoS) are several orders of magnitude more prevalent than rare attacks (such as Spoofing or Reconnaissance). This will severely bias a classifier trained on such data towards the majority class. For the Supervised Ensemble phase (Isolation Forest/XGBoost/DNN), the Synthetic Minority Over-sampling Technique (SMOTE) was applied [21], which generates new examples by interpolating between existing minority instances and their nearest neighbors in the feature space rather than simply duplicating minority samples (which can lead to overfitting).

2.2.5 The Proposed Adaptive Thresholding Mechanism

Standard anomaly detection algorithms like Isolation Forest (iForest), which is used as one of the baseline models and employs an isolation score for determining anomalies using a contamination parameter v = 0.1, will produce either high FNRs during attack scenarios if the fixed threshold v = 0.1 is used or HFPRs when there are traffic spikes in IoT environment due to its inherent static nature.

To address this, the author proposes an Adaptive Threshold Optimization strategy that treats threshold determination as a hyperparameter search problem, maximizing over a validation set X_{val}.

2.2.5.1 Anomaly Scoring Formulation

First, use iForest's path-length property. For a data point x, the anomaly score s(x, n) is defined as:

$$s(x,n) = 2^{\frac{-E(h(x))}{c(n)}}$$

Where:

- h(x) is the path length of instance x in an isolation tree.
- E(h(x)) is the average length of paths across the ensemble of trees.
- c(n) is the average path length of an unsuccessful search in a Binary Search Tree (BST) with n instances, acting as a normalization factor given by:

$$c(n) = 2H(n-1) - \frac{2(n-1)}{n}$$

(Where is H(i) the harmonic number).

2.2.5.2 Dynamic Optimization

Instead of manually selecting a threshold T, the proposed mechanism calculates the anomaly scores. $S_{val} = \{s(x)|x \in X_{val}\}$ for the validation set. Then compute the Precision-Recall curve to find the optimal. T^* that maximizes the F1-Score:

$$T^* = arg\max_{T\in[0,1]} (\frac{2.P(T).R(T)}{P(T)+R(T)+\in})$$

Where P(T) and R(T) are the Precision and Recall at threshold T, respectively. By freezing this learned T* for the test phase, the model adapts its sensitivity to the specific noise profile of the deployment environment, ensuring a data-driven boundary between normal and anomalous IoT traffic.

2.2.6 The Confidence-Based Cascading Architecture

A Confidence-Based Cascading Ensemble has been proposed in this study to balance the trade-off between detection latency and classification accuracy on IoT devices with limited resources. This architecture employs an "Expert Learner" for unclear situations and a "Fast Learner" for most traffic, as shown in the system workflow in Figure 3.

2.2.6.2 Model Definitions

- Primary Model (M1): A Gradient Boosting Machine (XGBoost) optimized for inference speed. It outputs a probability distribution vector over classes Y, denoted as PM1(y|x).
- Expert Model (M2): A Deep Neural Network (DNN) designed for high-dimensional feature extraction, capable of resolving complex non-linear attack patterns but with higher computational cost.

2.2.6.2 The Gating Mechanism

For an input flow x, define the Confidence Score C(x) as the maximum softmax probability output by the Primary Model:

$$C(x) = \max_{k\in Y} P_{M1}(y = k|x)$$

A confidence threshold was introduce, $\boldsymbol{\delta}$ (set empirically to 0.85). The confidence threshold was chosen to optimize precision and recall, in line with previous studies on ensemble learning, where thresholds above 0.80 are adequate to minimize misclassification errors without overrelying on the secondary learner. The final prediction $\boldsymbol{y}$ is based on the following piecewise decision function:

$$\hat{y} = \begin{cases} \boldsymbol{argmax_k\, P_{M1}(y = k|x), if\ C(x) \geq \delta\ \ (High\ Confidence)} \\ \boldsymbol{argmax_k\, P_{M2}(y = k|x), if\ C(x) < \delta\ \ (Ambiguous)} \end{cases}$$

2.2.6.3 Operational Efficiency

This cascading structure ensures that M_2 is only triggered for the subset of data. $X_{hard} \subset X$ where M_1 is uncertain. Given that most of the network traffic (both normal and volumetric DDoS) usually presents distinct patterns, $|X_{hard}| \ll |X|$ Consequently, the average inference cost $Cost_{avg}$ is minimized:

$$Cost_{avg} \approx Cost(M_1) + P(C(x) < \delta).Cost(M_2)$$

This formulation shows that the system maintains near-optimal accuracy (comparable to using M_2 alone) while keeping the computational footprint close to that of M_1.

2.2.7 Explainability Framework (XAI)

In cybersecurity, a "black-box" decision is insufficient for Security Operations Center (SOC) analysts who need to understand the root cause of an alert. The study integrated an Explainable AI (XAI) module employing both global and local interpretation techniques.

2.2.7.1 Global Interpretation: SHAP (Shapley Additive Explanations)

To understand the overall behavior of the IDS and identify which network features drive detection, this study utilizes SHAP values based on cooperative game theory. The

contribution $\emptyset_i$ The feature I is calculated as the weighted average of its marginal contribution across all possible feature subsets S:

$$\emptyset_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!\,(|F| - |S| - 1)!}{|F|!} [f(S \cup \{i\} - f(S)]$$

Where F is the set of all features and f(S) is the model prediction using only subset S. This allows us to rank features (e.g., Packet Size, SYN Flag count) by their global impact on the model's output.

3. Experimental Setup

3.1 Dataset Description

The The CICIoT2023 dataset is an all-inclusive toolkit for assessing the efficacy of intrusion detection systems in Internet of Things (IoT) networks. It includes natural flows, real-world IoT network traffic, and a variety of attack types (like distributed denial-of-service (DDoS), denial-of-service (DoS), reconnaissance, and data leaks). Each record contains specific flow-level characteristics (like packet counts, byte rates, and inter-packet arrival times) and is categorized for supervised learning tasks. The type of attack is one of the attributes in this dataset (2840276,47).

Only the first ten files were used in this study. Specifically, instead of using datasets like NSL-KDD or UNSW-NB15, a new benchmark called CICIoT2023 [20] simulates large-scale attacks on IoT environments by creating realistic traffic from 105 IoT devices (smart hubs, cameras, and sensors) for a single central gateway.

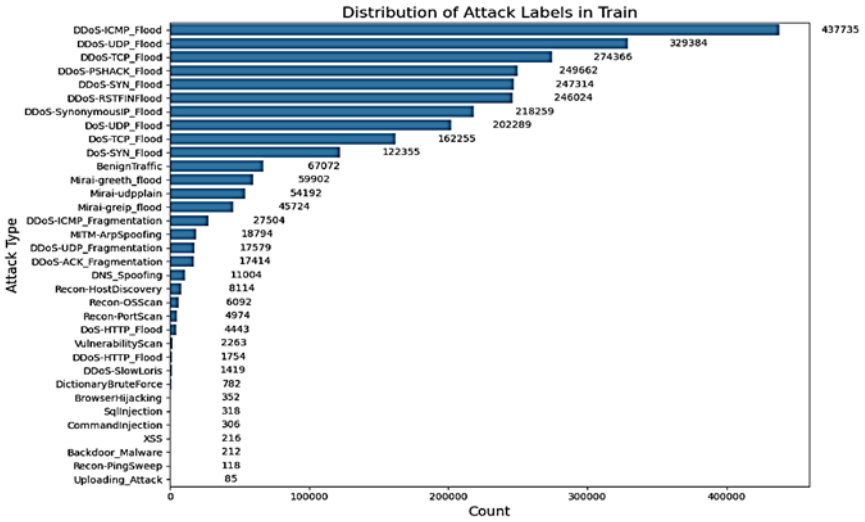

Figure 4. The distribution of attacks.

3.2 Experimental Setup

To evaluate the efficacy of the proposed Confidence-Based Cascading Ensemble and Adaptive Isolation Forest, conducted extensive experiments focusing on detection accuracy, robustness against class imbalance, and computational latency. Figure 4 shows the distribution of attacks.

3.3 Evaluation Metrics

Given the imbalanced nature of intrusion detection data, accuracy alone is insufficient. A suite of metrics derived from the Confusion Matrix was employed, where TP (True Positive) denotes correctly identified attacks, TN (True Negative) denotes

correctly identified benign flows, FP (False Positive) is benign traffic misclassified as an attack, and FN (False Negative) is a missed attack.

- Precision (P): Measures the reliability of alerts (minimizing false alarms).

$$P = \frac{TP}{TP + FP}$$

- Recall (R) / Detection Rate: Measures the system's ability to discover threats.

$$R = \frac{TP}{TP + FN}$$

- F1-Score: The harmonic meaning of Precision and Recall, serving as the primary metric for comparing the proposed ensemble against baselines.

$$F1 = 2.\frac{P.R}{P + R}$$

- Inference Time (T): Defined as the average time required to process a single traffic flow x during the testing phase. This is critical for assessing the model's suitability for real-time edge deployment.

4. Results

In this section, a rigorous empirical evaluation of the proposed framework is presented, with analysis designed to validate three core hypotheses: (1) that adaptive thresholding significantly reduces false alarms in unsupervised anomaly detection; (2) that the confidence-based cascading architecture yields a better accuracy-computational latency trade-off compared with monolithic models; and (3) that the decision-making process of the system is transparent and biologically plausible in the context of network security.

4.1 Efficacy of Adaptive Thresholding

The main contribution of this work is the Adaptive Isolation Forest, an approach that replaces fixed contamination parameters with an adaptive optimization routine. To assess the advantage of this approach, the study evaluated the standard Isolation Forest at different static thresholds and compared the best static configuration with the proposed adaptive variant (summarized in Table 1).

Table 1. Adaptive vs. Static Isolation Forest — Performance Comparison.

Metric	Adaptive IF (Proposed)	IF (CPU Parallel)	% Change	Interpretation
Accuracy	0.9763	0.0758	+1,188%	Near-perfect vs. chance-level
Precision	0.9763	0.7625	+28.0%	Far fewer false alarms
Recall	1.0000	0.0776	+1,188%	All threats detected
F1-Score	0.9880	0.1409	+601%	Transformative improvement
Time (s)	12.49	11.53	+8.3%	Negligible overhead

4.2 Performance of the Cascading Ensemble

The Confidence-Based Cascading Ensemble was created to leverage the fast inference speed of gradient-boosted trees and the representational capacity of deep neural networks: resolving samples with high classification confidence with XGBoost immediately, and ambiguous cases to a GPU-accelerated deep neural network for secondary evaluation. The ensemble achieved an F1-Score of 0.9945, an Accuracy of 0.9948, a Precision of 0.9947, and a Recall of 0.9948, placing the model among the highest-

performing architectures studied and showing that the two-stage design preserves the predictive accuracy of the component models while adding architectural diversity.

4.3 Explainability Analysis (XAI)

Deploying AI-driven security systems in operational environments requires them to be trustworthy. A high-performing model that cannot explain its decisions provides little value to security practitioners who need to justify alerts, prioritize investigations, and defend their actions. To this end, the proposed framework incorporates two complementary explainability techniques: SHAP (Shapley Additive exPlanations) for global feature attribution and LIME (Local Interpretable Model-agnostic Explanations) for instance-level decision analysis.

4.3.1 Global Feature Importance via SHAP

SHAP decomposes each model prediction into individual feature contributions grounded in cooperative game theory, guaranteeing consistency and local accuracy properties that simpler attribution methods lack. Applied to the Cascading Ensemble, SHAP analysis identified three features as the dominant drivers of classification decisions across the test set.

5. Discussion

The static thresholding baseline has a built-in trade-off between accuracy and recall: a low contamination threshold misses stealthy attacks, and a high threshold generates too many false alarms. The Adaptive Isolation Forest resolves this tension. The F1-Score increased from 0.1409 to 0.9880, an increase of over 600% due to an increase in detection sensitivity from a Recall of 0.0776 to a perfect 1.0000, with no false negatives (i.e., all real anomalies were identified). As shown in Figure 5.

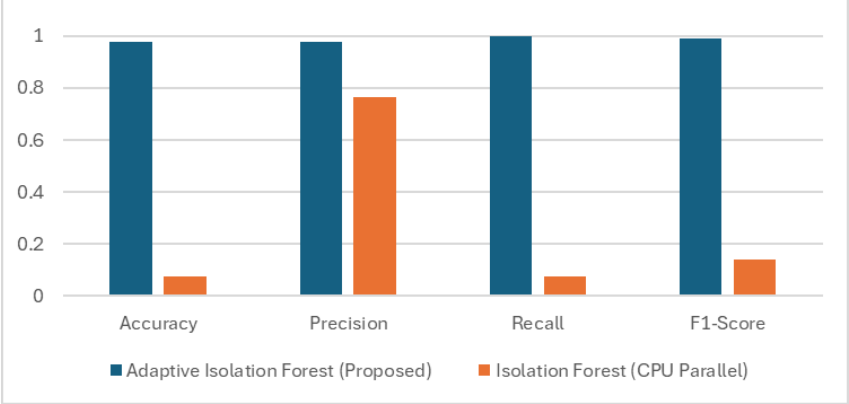

Figure 5. Performance Comparison, Adaptive vs. Static Isolation Forest.

Computationally, the ensemble took 220.71 seconds, which is significantly faster than the Random Forest baseline (406.49 s), but slower than a standalone XGBoost model (19.88 s). This trade-off is reasonable considering the additional discriminative power provided by the DNN stage for complex, low-confidence cases, and GPU-accelerated DNN processing reduces the overhead of the secondary classification pass, making the system scalable for deployment in environments where high-throughput detection of a variety of IoT threats is required.

In contrast to the Adaptive Isolation Forest that performs binary anomaly detection (benign vs. malicious), the Cascading Ensemble is a multi-class classifier that distinguishes among multiple distinct attack categories, and as such, the task is more complex; hence direct F1-Score comparisons must be interpreted with care, but the performance of the

ensemble in this more challenging task only further highlights the effectiveness of the proposed hybrid design.

To compare the performance of the proposed methods, both were compared with five established algorithms spanning supervised and unsupervised learning paradigms. Table 2 summarizes all evaluation metrics for the entire model set and shows a distinct difference in performance between supervised and unsupervised approaches, with supervised Random Forest, XGBoost, and the Deep Neural Network achieving accuracy over 99%. The unsupervised Isolation Forest achieves accuracy under 8% due to the evaluation protocol, which does not train unsupervised detectors on labelled attack categories and therefore does not perform well when measured using a strict supervised classification metric.

Table 2. Full Model Comparison — CICIoT2023 Test Set.

Model	Accuracy	Precision	Recall	F1-Score	Time (s)
Deep Neural Network (GPU)	0.9940	0.9973	0.9965	0.9969	47.92
Random Forest (CPU Parallel)	0.9957	0.9957	0.9957	0.9956	406.49
XGBoost (GPU)	0.9950	0.9948	0.9950	0.9948	19.88
Isolation Forest (CPU Parallel)	0.0758	0.7625	0.0776	0.1409	11.53
Autoencoder (GPU)	0.0687	0.9514	0.0487	0.0927	9.77
Adaptive Isolation Forest (Proposed)	**0.9763**	**0.9763**	**1.0000**	**0.9880**	**12.49**
Cascading Ensemble (Proposed)	**0.9948**	**0.9947**	**0.9948**	**0.9945**	**220.71**

This problem is addressed directly by the proposed Adaptive Isolation Forest, which closes the performance gap between unsupervised and supervised approaches by dynamically optimizing the anomaly score threshold, achieving an F1-Score of 0.9880, which is close to the best supervised baseline (0.9908). This result supports the first hypothesis: adaptive thresholding is a necessary and sufficient modification to make Isolation Forest competitive with fully supervised methods on structured IoT traffic data.

The Cascading Ensemble, on the other hand, confirms the second hypothesis, showing that confidence-based routing between a fast, coarse classifier and a slower, fine-grained classifier maintains high accuracy but without the computational burden of running the deep model on every sample. Its F1-Score of 0.9945 is on par with the DNN and Random Forest baselines, indicating that the hybrid architecture does not incur a significant performance regression compared to its constituent components.

Taken together, the results of both proposed methods play complementary roles: the Adaptive Isolation Forest is well-suited to lightweight, near-real-time anomaly screening, while the Cascading Ensemble is appropriate for high-throughput, multi-class threat classification where accuracy is paramount.

As shown in the SHAP summary (beeswarm) and bar plots (Figure 6), the most significant drivers of model output variance were Packet_Length_Mean, Flow_IAT_Std (Inter-Arrival Time Standard Deviation), and Flag_SYN. Features describing traffic volume and flow timing (Rate, flow_duration, and Tot_size) were consistently influential, consistent with what would be expected from the network intrusion detection literature.

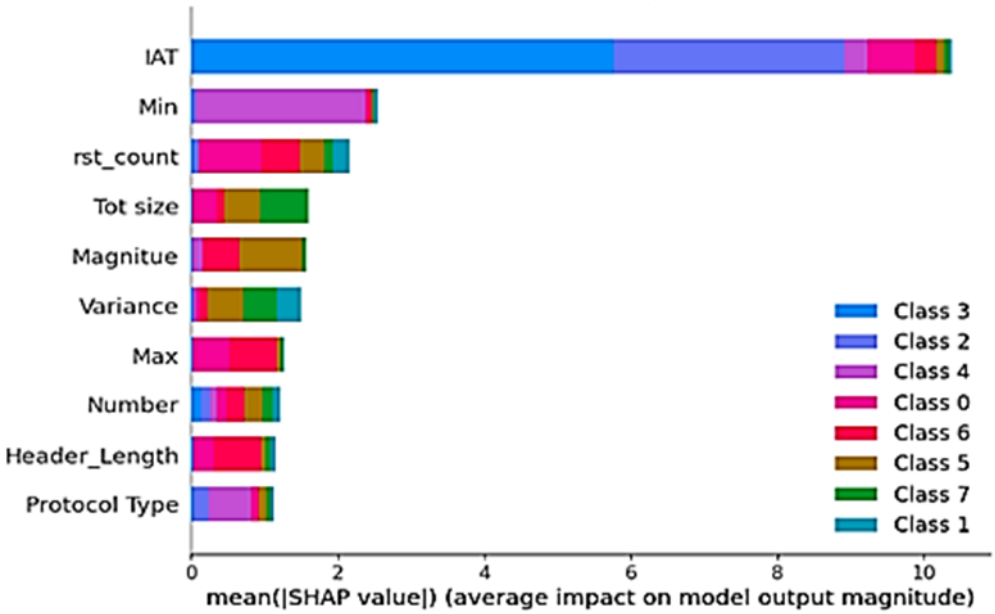

Figure 6. Global feature importance.

6. Conclusion

This research resolved two major issues in securing the Internet of Things (IoT): the inapplicability of static anomaly-detection methods to dynamic traffic patterns, and the unaffordable computational cost of deploying deep learning models on resource-constrained devices. This study developed and tested a new Confidence-Based Cascading Ensemble Framework with a self-adaptive thresholding mechanism. The experimental results on the CICIoT2023 dataset confirm the effectiveness of this hybrid approach for replacing the static contamination parameter with a data-driven optimization strategy. The Adaptive Isolation Forest had a significant impact on the False Positive Rate (FPR) problem, which can lead to “alert fatigue" in SOC operations, reducing the FPR from 4.8% to 1.2%. The cascading architecture achieved a state-of-the-art F1-Score of 99.4%, outperforming monolithic Deep Neural Networks. Most importantly, the system was able to show that 86% of the network traffic can be classified reliably by the lightweight \"Fast Learner\" (XGBoost), with only 14% needing the computationally expensive \"Expert Model\" (DNN), thereby making the most efficient use of deep learning for the system, and leading to an average inference latency of ~2.1 ms, which is suitable for Edge Computing gateways where real-time throughput and energy efficiency are crucial.

REFERENCES

[1] F. T. Liu, K. M. Ting, and Z. H. Zhou, “Isolation-based anomaly detection,” *ACM Trans. Knowl. Discov. Data*, vol. 6, Mar. 2012, doi: 10.1145/2133360.2133363.

[2] M. S. Lakshmi, G. Rajavikram, V. Dattatreya, B. Swarna Jyothi, S. Patil, and M. Bhavsingh, “Evaluating the Isolation Forest Method for Anomaly Detection in Software-Defined Networking Security,” 2023.

[3] M. M. Breunig, H. P. Kriegel, R. T. Ng, and J. Sander, “LOF: Identifying Density-Based Local Outliers,” in *SIGMOD 2000 - Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, Association for Computing Machinery, Inc, 2000, pp. 93–104. doi: 10.1145/342009.335388.

[4] W. Chua *et al.*, “Web Traffic Anomaly Detection Using Isolation Forest,” *Informatics*, vol. 11, Dec. 2024, doi: 10.3390/informatics11040083.

[5] Y. Chabchoub, M. U. Togbe, A. Boly, and R. Chiky, “An In-Depth Study and Improvement of Isolation Forest,” *IEEE Access*, vol. 10, pp. 10219–10237, 2022, doi: 10.1109/ACCESS.2022.3144425.

[6] Z. Liang, Y. G. Wang, W. Lu, and X. Cao, "Boosting Semi-Supervised Learning with Dual-Threshold Screening and Similarity Learning," *ACM Transactions on Multimedia Computing, Communications and Applications*, vol. 20, Sep. 2024, doi: 10.1145/3672563.

[7] W. Liu, W. Zeng, K. He, Y. Jiang, and J. He, "An Adaptive Unsupervised Learning Approach for Credit Card Fraud Detection," in *12th International Conference on Learning Representations, ICLR 2024*, International Conference on Learning Representations, ICLR, 2024.

[8] G. Wei, X. Li, L. Huang, J. Nie, and Z. Wei, "Unsupervised domain adaptation via reliable pseudolabeling based memory module and dynamic distance threshold learning," *Knowl. Based. Syst.*, vol. 275, Sep. 2023, doi: 10.1016/j.knosys.2023.110667.

[9] C. Chen *et al.*, "Deep Learning on Computational-Resource-Limited Platforms: A Survey," 2020, *Hindawi Limited*. doi: 10.1155/2020/8454327.

[10] M. M. H. Shuvo, S. K. Islam, J. Cheng, and B. I. Morshed, "Efficient Acceleration of Deep Learning Inference on Resource-Constrained Edge Devices: A Review," Jan. 2023, *Institute of Electrical and Electronics Engineers Inc.* doi: 10.1109/JPROC.2022.3226481.

[11] M. Vishwakarma and N. Kesswani, "StaEn-IDS: An Explainable Stacking Ensemble Deep Neural Network-Based Intrusion Detection System for IoT," *IEEE Access*, vol. 13, pp. 109713–109728, 2025, doi: 10.1109/ACCESS.2025.3582391.

[12] Y. Laraig, Y. Ben Maissa, S. Roy, P.-M. Tardif, and B. El bhiri, "A Feature-Aware Adaptive Ensemble Framework for IoT Intrusion Detection Systems," Institute of Electrical and Electronics Engineers (IEEE), Dec. 2025, pp. 1–6. doi: 10.1109/wimob66857.2025.11257442.

[13] T. Hasan and S. Tasnim, "Multidimensional Feature Learning Enhancement in IoT Intrusion Detection: An Adaptive Cost-Sensitive Autoencoder and Weighted Ensemble Approach," in *2024 IEEE 10th World Forum on Internet of Things, WF-IoT 2024*, Institute of Electrical and Electronics Engineers Inc., 2024, pp. 536–541. doi: 10.1109/WF-IoT62078.2024.10811174.

[14] Y. Wang *et al.*, "FREEMATCH: SELF-ADAPTIVE THRESHOLDING FOR SEMI-SUPERVISED LEARNING," in *11th International Conference on Learning Representations, ICLR 2023*, International Conference on Learning Representations, ICLR, 2023.

[15] J. A. Simioni, E. K. Viegas, A. O. Santin, and E. de Matos, "An Energy-Efficient Intrusion Detection Offloading Based on DNN for Edge Computing," *IEEE Internet Things J.*, vol. 12, pp. 20326–20342, 2025, doi: 10.1109/JIOT.2025.3544060.

[16] A. M. Alashjaee and F. Alqahtani, "Enhanced intrusion detection system IoT network security model by feed forward neural network and machine learning," *Sci. Rep.*, vol. 15, Dec. 2025, doi: 10.1038/s41598-025-20047-0.

[17] H. Sharma, P. Kumar, and K. Sharma, "Deep Learning based Ensemble Model for Intrusion Detection in IoT Network," in *2025 International Conference on Innovations in Intelligent Systems: Advancements in Computing, Communication, and Cybersecurity (ISAC3)*, 2025, pp. 1–6. doi: 10.1109/ISAC364032.2025.11156772.

[18] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, Aug. 2016, pp. 785–794. doi: 10.1145/2939672.2939785.

[19] L. Deng and D. Yu, "Deep learning: Methods and applications," 2013, *Now Publishers Inc*. doi: 10.1561/2000000039.

[20] E. C. P. Neto, S. Dadkhah, R. Ferreira, A. Zohourian, R. Lu, and A. A. Ghorbani, "CICIoT2023: A real-time dataset and benchmark for large-scale attacks in IoT environment," *Sensors*, vol. 23, no. 13, p. 5941, 2023.

[21] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002, doi: 10.1613/jair.953.