

Adaptive Error-Aware Light Gradient Boost Machine for Cross-Dataset Intrusion Detection in MANETS

Authors

Fridah Chepkemai Korir ⁽¹⁾; Solomon Mwanjele Mwagha ⁽²⁾; Gilbert Langat ⁽³⁾

Main author's email: fridahcheps@gmail.com

(1.2) Taita Taveta University, Kenya; (3) Kirinyaga University, Kenya.

Cite this article in APA

Korir, F. C., Mwagha, S. M., & Langat, G. (2026). Adaptive error-aware Light Gradient Boost Machine for cross-dataset intrusion detection in MANETS. *Journal of computer science and technology*, 4(1), 50–73. <https://doi.org/10.51317/jcst.v4i1.1082>



A publication of Editon
Consortium Publishing (online)

Article history

Received: 2026-05-24

Accepted: 2026-06-22

Published: 2026-07-29

Scan this QR to read the paper online



Copyright: ©2026 by the author(s). This article is an Open Access article distributed under the terms and conditions of the Creative Commons Attribution-NonCommercial ShareAlike 4.0 International License (CC BY-NC-SA 4.0).



Abstract

The purpose of this article is to develop an adaptive error-aware light gradient boost machine for cross-dataset intrusion detection in Mobile Ad hoc Networks (MANETs). MANETs are dynamic in nature and thus vulnerable to various attacks, such as black hole, grey hole, and tapping attacks. Therefore, to mitigate security issues, intrusion detection in MANETs increasingly relies on machine learning models, with existing studies dominated by accuracy-centric evaluation. However, there are still limitations of such metrics in highly separable and imbalanced intrusion datasets. Therefore, this study proposes an Adaptive Error-Aware LightGBM (AEA-LGBM) framework for intrusion detection across datasets. The methods used include preprocessing the data through encoding and feature standardisation, as well as training and testing. The framework dynamically adjusts instance weights to minimise the false negative rate (FNR) while integrating an error-efficiency evaluation strategy that combines FNR, false positive rate (FPR), and inference time. Experiments using CIC-DDoS2019 and CICIDS2017 datasets show that AEA-LGBM reduces FNR from 0.000273 to 0.000068 during internal evaluation while maintaining near-perfect accuracy. Independent cross-dataset validation reveals substantial performance degradation across conventional models due to domain shift. To address this limitation, a transfer-aware heterogeneous learning strategy is introduced, achieving 99.90 per cent accuracy, 99.94 per cent recall, and 0.000580 FNR under unseen cross-environment evaluation. Therefore, the model effectively minimises the false negative rate through adaptive weighting, thereby improving MANET security.

Key terms: Cross-dataset, intrusion detection, LightGBM, machine learning, MANETs.

1.0 INTRODUCTION

With advances in technology, improved security techniques for MANETs have been implemented. However, they continually face challenges as new threats emerge undetected. Intrusion detection in modern networked environments increasingly relies on machine learning and ensemble-based techniques to identify malicious traffic patterns. Among these approaches, gradient boosting frameworks such as the Light Gradient Boosting Machine (LightGBM) demonstrate strong predictive capability by effectively modelling complex nonlinear relationships in high-dimensional traffic data (Hemanand et al., 2026; Vaiyapuri & Murugesan, 2026). Despite these advances, contemporary intrusion detection research remains largely dominated by accuracy-centric evaluation paradigms in which model performance is primarily assessed using metrics such as accuracy, F1-score, and area under the curve (AUC) (Novika & Mujhidi, 2025; Panigrahi et al., 2022). Although such metrics are useful under balanced conditions, they become increasingly inadequate in highly imbalanced and security-sensitive environments, where the consequences of false negatives are substantially more severe than those of false positives.

Recent intrusion detection datasets, including CIC-DDoS2019, exhibit high separability between benign and malicious traffic, resulting in near-saturated accuracy and AUC values across multiple machine learning models (Hemanand et al., 2026; Vaiyapuri & Murugesan, 2026). Under such conditions, marginal improvements in accuracy provide limited practical significance and frequently obscure more operationally important performance indicators. In particular, false negative rate (FNR), which represents undetected attacks, remains insufficiently emphasised despite its direct implications for network security and attack survivability (Yang et al., 2020). Existing studies also rarely integrate computational efficiency into evaluation frameworks, even though inference time is critical for real-time intrusion detection in resource-constrained MANET environments.

To address class imbalance, several studies employ data-level resampling techniques, such as the Synthetic Minority Over-sampling Technique (SMOTE) and Adaptive Synthetic Sampling (ADASYN), to improve minority-class detection (Liu et al., 2019). While these methods improve class balance, they often introduce additional computational complexity and fail to address the underlying issue of error prioritisation. Similarly, cost-sensitive learning approaches assign higher penalties to minority-class misclassification (Almseidin et al., 2017). However, most existing methods rely on static weighting strategies and do not explicitly optimise for critical security-sensitive metrics, such as the false-negative rate.

Another major limitation in the literature is the lack of an integrated evaluation framework that simultaneously accounts for detection reliability, error asymmetry, computational efficiency, and deployment robustness. Existing benchmarking studies mainly report predictive performance metrics without evaluating cross-environment adaptability or operational feasibility under heterogeneous network conditions (Tang et al., 2016). Consequently, many intrusion detection models achieve excellent performance on internal benchmarks while remaining vulnerable to domain shift and poor cross-dataset generalisation in real-world deployment environments. Data exchange among various MANET components may be compromised by attacks, compromising network performance (Nyangaresi, 2022).

This study addresses these limitations by proposing an Adaptive Error-Aware LightGBM (AEA-LGBM) framework that explicitly prioritises minimising false negatives through adaptive dynamic weighting. Unlike conventional cost-sensitive approaches, the proposed framework incorporates error-sensitive feedback

directly into the optimisation process, thereby improving attack detection under imbalanced conditions. In addition, this study introduces a unified error-efficiency evaluation framework that integrates the false negative rate, false positive rate, and inference time for deployment-oriented model assessment. To further evaluate robustness under heterogeneous environments, independent cross-dataset validation and transfer-aware heterogeneous learning experiments are conducted using the CIC-DDoS2019 and CICIDS2017 datasets.

2.0 LITERATURE REVIEW

Accuracy-Centric Intrusion Detection: Strengths and Limitations

Machine learning has significantly improved intrusion detection performance, particularly through ensemble-based approaches such as LightGBM and hybrid deep learning architectures. Hemanand et al. (2026) proposed an intelligent model for intrusion detection in networks using the BE2GO-SEINN algorithm. The model addresses the issues of system complexity and training time. However, the Markovian Data Clustering Model (MDCM) requires significant computational resources for clustering datasets. Furthermore, it lacks an effective method for determining the level of trustworthiness among the communicating nodes. Vaiyapuri and Murugesan (2026) introduced the One-Class Learning (OCL) framework to mitigate issues in intrusion detection models related to computational cost and storage requirements. However, the proposed model is complex and poorly adapted to new attacks.

A cost-sensitive framework with LightGBM for unbalanced datasets was presented by Novika and Mujhid, (2025). The model experiences overfitting, especially on small datasets. Panigrahi et al. (2022) developed a hybrid decision table-Naive Bayes framework for feature selection. However, relying on only five selected features from the dataset may not yield reliable results. Yang et al. (2020) introduced the SAVAER-DNN adversarial model, which uses the WGAN-GP algorithm to enhance learning. However, there are other, more efficient frameworks, such as the Temporal Difference VAE and a two-level hierarchical Vector Quantised Variational Auto-Encoder. Liu et al. (2019) proposed a convolutional neural network-recurrent neural network framework for efficient intrusion detection in networks. Combining the two algorithms introduces complexity to the model, thus straining network resources.

Panigrahi et al. (2022) studied different machine learning techniques such as Random Forest, J48, Naive Bayes, Decision Table, Bayes Network, Random Tree, and MLP for intrusion detection. The goal of the study was to evaluate the algorithms mentioned in terms of true positive rate and accuracy. However, a more efficient intrusion detection technique requires consideration of false-negative and false-positive errors. Tang et al. (2016) developed a Deep Neural Network technique with the NSL-KDD dataset for efficient network intrusion detection. The model lacks real-world implementation. Kim et al. (2014) suggested a novel hybrid with the C4.5 decision tree algorithm for zero-day attacks. However, the C4.5 algorithm can degrade the model's efficiency due to uneven dataset decomposition.

Similarly, Li et al. (2012) introduced an intrusion detection model that clusters features using the ant colony algorithm and a support vector machine for efficient intrusion detection. The model does not adapt well in a sophisticated software environment where an unbalanced dataset is encountered. Furthermore, the model is complex, thus straining the already drained resources in the network. Nyangaresi (2022) designed a lightweight model based on one-way hashing and elliptic curve point multiplication operators. The method considers network resources. However, the technique is less considerate of false positive and false negative errors, which can have a significant negative impact on an organisation. Bii et al. (2022)

designed a framework that can adapt to adversarial attacks with an optimisation algorithm. Computational complexity and failure to effectively detect intrusion in high-level environments limit the model. John et al. (2024) proposed a framework for feature selection and multiclass classification using principal component analysis, AdaBoost ensemble, and stagewise additive modelling. The method requires more efficiency and adequacy in detecting vulnerabilities in MANETs.

Despite these advances, existing intrusion detection research remains largely dominated by accuracy-centric evaluation paradigms. Recent studies continue to emphasise metrics such as accuracy, F1 Score, and area under the curve (AUC) as primary indicators of model performance (Novika & Mujhidi, 2025; Panigrahi et al., 2022). However, in highly separable intrusion datasets, these metrics often become saturated, leading multiple models to appear nearly identical despite meaningful differences in operational behaviour. Vaiyapuri and Murugesan (2026) noted that overreliance on accuracy-based evaluation limits the ability to distinguish among high-performing intrusion detection systems under realistic deployment conditions.

From a security perspective, this evaluation strategy presents a critical limitation. In intrusion detection systems, false negatives correspond to undetected attacks and therefore pose substantially greater risk than false positives. Nevertheless, most existing studies do not explicitly optimise for or prioritise false-negative minimisation. Panigrahi et al. (2022) observed that many intrusion detection approaches focus primarily on reducing false positives while overlooking the asymmetric security impact associated with undetected malicious traffic. Consequently, current evaluation practices remain insufficiently aligned with real-world operational security requirements.

Evaluation Practices in Intrusion Detection

Evaluation methodology remains one of the most underdeveloped aspects of intrusion detection research. Benchmarking studies such as Almseidin et al. (2017) provide comparative analyses of machine learning algorithms using metrics including accuracy, precision, recall, and F1-score. While these studies offer valuable performance comparisons, they generally adopt static, benchmark-oriented evaluation strategies that do not adequately account for deployment-specific constraints such as inference time, computational efficiency, and operational robustness.

More recent works have incorporated additional performance indicators including detection speed and resource utilisation (Kim et al., 2014). However, these metrics are typically reported independently rather than integrated into a unified decision-making framework. Consequently, existing studies offer limited guidance on balancing attack detection capability with computational feasibility in real-world deployment environments.

Another major limitation is the lack of systematic cross-dataset evaluation. Many intrusion detection models achieve near-perfect performance under internal testing conditions but fail to maintain comparable reliability when exposed to unseen traffic environments with different feature distributions and attack behaviours. This limitation highlights the growing importance of domain shift and cross-dataset generalisation analysis in modern intrusion detection research.

Toward Error-Aware and Deployment-Oriented Intrusion Detection

The limitations identified in the existing literature indicate the need to transition from accuracy-centric intrusion detection to robustness-aware, deployment-oriented evaluation. Although previous studies have explored cost-sensitive learning (Liu et al., 2019), hybrid architectures (Hemanand et al., 2026), and real-time intrusion detection (Vaiyapuri & Murugesan, 2026), few studies have integrated adaptive error prioritisation, computational efficiency, and cross-dataset robustness into a unified framework.

This study advances the field through three major contributions. First, an Adaptive Error-Aware LightGBM (AEA-LGBM) framework is proposed to minimise the false-negative rate by explicitly applying adaptive dynamic weighting. Unlike conventional static cost-sensitive approaches, the proposed framework continuously adjusts the learning emphasis based on observed error behaviour during training.

Second, this study introduces an integrated error-efficiency evaluation framework that combines the false-negative and false-positive rates with inference time. This framework enables meaningful differentiation between high-performing models under highly separable intrusion environments where conventional metrics become saturated.

Third, the study extends intrusion detection evaluation beyond isolated benchmark testing through independent cross-dataset validation and transfer-aware heterogeneous learning experiments. By incorporating cross-environment evaluation using CIC-DDoS2019 and CICIDS2017 datasets, the study demonstrates the impact of domain shift on intrusion detection reliability and highlights the importance of deployment-oriented robustness analysis.

Overall, the existing literature reveals three major gaps: insufficient emphasis on false-negative optimisation, overreliance on saturated accuracy-centric metrics, and limited evaluation of cross-dataset robustness under real-world deployment conditions. This study addresses these limitations by introducing an error-aware, efficiency-sensitive, and transfer-aware intrusion detection framework designed for heterogeneous MANET environments.

3.0 METHODOLOGY

The study was an experimental machine learning study conducted to assess the proposed model's adaptability to errors such as false negatives and false positives. It also explores its application across datasets to ensure consistent results. The experiments are conducted using the CIC-DDoS2019 dataset developed by the Canadian Institute for Cybersecurity (CIC) at the University of New Brunswick (UNB). The dataset is imported as a CSV file and processed within a Python-based machine learning environment. Class labels are encoded numerically using a label encoding scheme, and the feature matrix is constructed by excluding label attributes and retaining only numerical features. This is done because various machine learning algorithms cannot interpret categorical data. This will then be possible to train the proposed model. The dataset is partitioned using stratified sampling into training (70%), validation (15%), and testing (15%) subsets through a two-stage split procedure to preserve class distribution. Feature standardisation is applied using the StandardScaler transformation, with scaling parameters learned from the training set and applied to the test set to ensure consistency and prevent data leakage.

Baseline Model Formulation

To establish a comparative benchmark, multiple machine learning models are implemented, including Logistic Regression, Naïve Bayes, K-Nearest Neighbours (KNN), Decision Tree, Random Forest, Support Vector Machine (SVM), XGBoost, and LightGBM.

Logistic Regression models the conditional probability of class membership as:

$$P(y = 1 | x) = \frac{1}{1 + e^{-(w^T x + b)}} \quad (1)$$

Where

$x \in \mathbb{R}^n$: Input feature vector with n features

$w \in \mathbb{R}^n$: Weight vector representing the contribution of each feature

$w^T x$: Dot product between weights and features

$b \in \mathbb{R}$: Bias (intercept) term

e : Base of the natural logarithm

Naïve Bayes applies Bayes' theorem:

$$P(y | x) = \frac{P(x | y)P(y)}{P(x)} \quad (2)$$

y : Class label (attack or benign)

x : Input feature vector

$P(y | x)$: Posterior probability of class y given x

$P(x | y)$: Likelihood of observing x given class y

$P(y)$: Prior probability of class y

$P(x)$: Marginal probability of observing x

KNN computes distances as:

$$d(x_i, x_j) = \sqrt{\sum_{k=1}^n (x_{ik} - x_{jk})^2} \quad (3)$$

Where

x_i, x_j : Two data instances

x_{ik}, x_{jk} : Values of the k – th feature for each instance

n : Number of features

$d(x_i, x_j)$: Euclidean distance between instances




Figure 1: Methodological Framework of the Proposed AEA-LGBM Intrusion Detection System

SVM optimises:

$$\min_{w,b} \frac{1}{2} \|w\|^2 + C \sum_{i=1}^N \xi_i \quad (4)$$

Where

w : Weight vector defining the hyperplane

b : Bias term

$\|w\|^2$: Squared norm of w , representing margin size

C : Regularisation parameter controlling margin-error trade-off

ξ_i : Slack variable for the i – th sample

N : Total number of training samples

Decision Trees minimise impurity:

$$Gini = 1 - \sum_{k=1}^K p_k^2 \quad (5)$$

Where

K : Total number of classes in the dataset

k : Index of each class ($k = 1, 2, \dots, K$)

p_k : Proportion (probability) of samples belonging to class k at a given node

Random Forest aggregates trees:

$$\hat{y} = \frac{1}{T} \sum_{t=1}^T f_t(x) \quad (6)$$

Where

$\hat{y}$: Final predicted output of the Random Forest model

T : Total number of decision trees in the ensemble

t : Index of each tree ($t = 1, 2, \dots, T$)

$f_t(x)$: Prediction of the t – th decision tree for input x

x : Input feature vector

Boosting-based models (XGBoost and LightGBM) optimize:

$$L = \sum_{i=1}^N l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k) \quad (7)$$

Where

N : Number of training samples

y_i : True label of the i – th sample

$\hat{y}_i$: Predicted output for the i – th sample

$l(y_i, \hat{y}_i)$: Loss function (e.g., logistic loss)

K : Number of trees (boosting iterations)

f_k : The k – th decision tree

$\Omega(f_k)$: regularisation term controlling model complexity

For XGBoost:

$$\Omega(f_k) = \gamma T + \frac{1}{2} \lambda \|w\|^2 \quad (8)$$

Where

γ : Penalty for adding a new leaf (controls tree complexity)

λ : Regularisation parameter on leaf weights

T : Number of leaves in the tree

w : Vector of leaf weights

Adaptive Error-Aware LightGBM (AEA-LGBM)

To move beyond accuracy-centric optimisation, the LightGBM objective is modified using error-aware instance weighting:

$$L_{AEA} = \sum_{i=1}^N w_i \cdot l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k) \quad (9)$$

Where

N : Total number of training instances

i : Index of each training instance, $i = 1, 2, \dots, N$

y_i : True label of instance i (1 = attack, 0 = benign)

$\hat{y}_i$: Predicted output for instance i

$l(y_i, \hat{y}_i)$: Loss function measuring prediction error

w_i : Weight assigned to instance i

K : Number of boosting iterations (trees)

k : Index of each tree

The weighting function is defined as:

$$w_i = \begin{cases} \lambda_1(t), & y_i = (\text{attack}) \\ \lambda_2, & y_i = (\text{benign}) \end{cases} \quad (10)$$

$$\lambda_1(t) > \lambda_2$$

Where

$\lambda_1(t)$: Adaptive weight for attack instances at iteration t

λ_2 : Constant weight for benign instances

t : Boosting iteration index

Condition $\lambda_1(t) > \lambda_2$: ensures higher penalty for misclassified attacks

To introduce explicit adaptivity, the attack-class weight is dynamically updated as:

$$\lambda_1(t) = \lambda_1 \cdot (1 + FNR_t) \quad (11)$$

Where

λ_1 : base penalty weight for attack instances

FNR_t : False negative rate at

Iteration t , defined as (12)

$$FNR_t = \frac{FN_t}{FN_t + TP_t}$$

FN_t : Number of false negatives at iteration t

TP_t : Number of true positives at iteration t

As the model observes higher false-negative rates, the penalty assigned to attack instances increases, thereby forcing subsequent boosting iterations to focus on hard-to-detect attacks. The benign class weight `image.png` remains constant to prevent excessive false positives.

From an implementation perspective, this mechanism is implemented through class weighting and boosting updates, in which misclassified instances receive greater influence in subsequent trees.

The proposed AEA-LGBM introduces dynamic error-driven weighting, in which the learning objective is continuously adjusted based on the model's error behaviour. Unlike standard LightGBM and existing cost-sensitive approaches that rely on fixed weights or external resampling, the proposed method embeds adaptive error feedback (via `image.png`) directly into the optimization process, enabling targeted minimization of false negatives while preserving computational efficiency.

Error-Sensitive and Efficiency-Aware Evaluation

Model performance is evaluated using both conventional and error-sensitive metrics.

Error rates are computed as:

$$FPR = \frac{FP}{FP + TN} \quad (13)$$

$$FNR = \frac{FN}{FN + TP} \quad (14)$$

Where

TN : Number of true negatives

FN : Number of false negatives

TP : Number of true positives

Inference time is defined as:

T_{inf} : Time required to generate predictions on the test dataset

To integrate detection performance with computational efficiency, the Error Efficiency Score (EES) is defined as:

$$EES = \alpha \cdot FNR + \beta \cdot FPR + \gamma \cdot T_{inf} \quad (15)$$

Where

α : Weight assigned to false negative rate (attack detection importance)

β : Weight assigned to false positive rate (false alarm control)

γ : Weight assigned to inference time (computational efficiency)

This evaluation framework shifts model assessment from accuracy-based comparison to an integrated error-efficiency paradigm, enabling meaningful differentiation between high-performing models in scenarios where traditional metrics are saturated.

Experimental Procedure and Ethical Considerations

All models are trained on the training dataset and evaluated on the test dataset using a unified evaluation function. Predictions are used to compute confusion matrices, from which performance metrics, error rates, and inference time are derived. All results are stored and exported to ensure reproducibility of the experimental process.

From an ethical standpoint, this study utilises a publicly available dataset that does not contain personally identifiable information. No human subjects are involved, and no sensitive user data is processed. The research adheres to standard principles of scientific integrity, including transparency, reproducibility, and proper attribution of data sources.

4.0 FINDINGS AND DISCUSSION

Table 1 shows substantial variability across features, with Flow Bytes/s reaching 2.94×10^9 and Flow Duration exceeding 10^8 , indicating mixed traffic intensities. Flag features exhibit low mean but high variance, reflecting intermittent attack patterns. These characteristics confirm that the dataset is highly non-linear and imbalanced, justifying the need for robust and error-aware intrusion detection models.

Table 1: Descriptive Statistics of Selected Network Traffic Features in the CIC-DDoS2019 Dataset

Feature	Mean	Standard Deviation	Minimum	Maximum
Flow Duration	8.404856e+06	2.126596e+07	1.000000	1.199987e+08
Total Fwd Packets	2.413912e+01	1.958889e+02	1.000000	8.666600e+04
Total Backward Packets	2.472021e+00	5.637021e+01	0.000000	3.170000e+04
Fwd Packet Length Mean	3.249153e+02	2.685773e+02	0.000000	3.015291e+03
Bwd Packet Length Mean	2.594426e+01	1.255623e+02	0.000000	5.011060e+03
Flow Bytes/s	7.373761e+07	3.323484e+08	0.000000	2.944000e+09
Flow Packets/s	1.266888e+05	4.142850e+05	0.033920	4.000000e+06
Flow IAT Mean	7.998733e+05	1.603964e+06	0.333333	5.119989e+07
Fwd IAT Mean	9.762455e+05	1.994982e+06	0.000000	5.119989e+07
Bwd IAT Mean	4.312665e+05	1.901205e+06	0.000000	5.896159e+07
Packet Length Mean	3.328787e+02	2.676589e+02	0.000000	4.023945e+03
Packet Length Std	3.198904e+01	1.182152e+02	0.000000	6.616562e+03
SYN Flag Count	5.331837e-04	2.308464e-02	0.000000	1.000000e+00
ACK Flag Count	1.764096e-01	3.811687e-01	0.000000	1.000000e+00
Down/Up Ratio	2.441216e-01	6.343976e-01	0.000000	2.300000e+01

Figure 2 shows a clear dominance of attack instances over benign traffic, confirming severe class imbalance. This imbalance explains why accuracy is insufficient and motivates the adoption of False Negative Rate (FNR) and False Positive Rate (FPR) as primary evaluation metrics.




Figure 2: Class distribution (Benign vs Attack)

Figure 3 shows how flow duration occurs in the given datasets. The distribution results from the dynamic network structure. It explicitly explains the flows of normal and attack instances, which gives a clear picture for detecting intrusions in MANETs.




Figure 3: Distribution of Flow Duration

Figure 4 shows that most of the routed messages contain less forwarded data. The graph is right skewed meaning that the datasets are imbalanced.




Figure 4: Distribution of Total Forwarded Packets

Figure 5 illustrates less backward packets by most of the network flows.




Figure 5: Distribution of Total Backward Packets

Figure 6 illustrates that forwarded packets are highly distributed towards 500 and diminishes at 1500. This shows that the packets still contain some attacks and thus the dataset is not balance.




Figure 6: Distribution of Forwarded Packet Length Mean

Figure 7 shows a distribution of the mean backward packet length against various frequencies. It is highly pronounced at the beginning but diminishes significantly as it progresses. This indicates an attack, since more data is disseminated than feedback.




Figure 7: Distribution of Backward Packet Length Mean

Baseline Model Performance

Table 2 shows that LightGBM achieves the best overall performance, with the lowest FPR (0.0001) and FNR (0.00015), confirming its superior ability to capture complex nonlinear traffic patterns. KNN (105 s inference time) and SVM (521 s training time) are computationally impractical. These results justify selecting LightGBM as the baseline for further enhancement.

Table 2: Comparative Performance of Machine Learning Models

Metric	LR	NB	KNN	DT	RF	SVM	XGB	LGBM
Accuracy	0.9957	0.9848	0.9987	0.9996	0.9997	0.9977	0.9997	0.9999
Precision	0.9875	0.9569	0.9968	0.9990	0.9990	0.9921	0.9989	0.9996
Recall	0.9936	0.9767	0.9974	0.9994	0.9997	0.9980	0.9997	0.9998
F1	0.9905	0.9667	0.9971	0.9992	0.9994	0.9950	0.9993	0.9997
FPR	0.0037	0.0129	0.0009	0.0003	0.0003	0.0023	0.0003	0.0001
FNR	0.0064	0.0233	0.0026	0.0006	0.0003	0.0020	0.0003	0.00015
Train (s)	5.80	0.43	0.11	9.14	62.45	521.02	12.35	7.62
Inference Time (s)	0.020	0.121	105.38	0.024	0.497	41.88	0.152	0.285

NB: LR = Logistic Regression; NB = Naïve Bayes; KNN = K-Nearest Neighbors; DT = Decision Tree; RF = Random Forest; SVM = Support Vector Machine; XGB = Extreme Gradient Boosting; LGBM = Light Gradient Boosting Machine.

Class Imbalance Handling

Table 3 shows that internal balancing achieves the same FNR (0.000136) as SMOTE with lower inference time. This demonstrates that synthetic oversampling is unnecessary for highly separable intrusion datasets, enabling reduced computational overhead without compromising detection performance.

Table 3: Comparison of Imbalance Handling Strategies

Model	Accuracy	FPR	FNR	Inference Time (s)
Default LGBM	0.999722	0.00028	0.000273	0.1977
Balanced LGBM	0.999706	0.00034	0.000136	0.1811
SMOTE + LGBM	0.999706	0.00034	0.000136	0.2186

FPR = False Positive Rate; FNR = False Negative Rate; LGBM = Light Gradient Boosting Machine; SMOTE = Synthetic Minority Over-sampling Technique; s = seconds.

AEA-LGBM Performance

Table 4 shows that AEA-LGBM achieves the lowest FNR (0.000068), representing a 75 per cent reduction compared to the default model, alongside the highest recall (0.999932). While accuracy remains saturated across models, meaningful differences emerge in error rates. Internal balancing matches SMOTE's performance (FNR = 0.000136) while reducing inference time. The increased inference time of AEA-LGBM highlights a trade-off between improved attack detection and computational efficiency.

Table 4: Final Model Performance Comparison

Metric	Default	Balanced	SMOTE	AEA-LGBM
Accuracy	0.999722	0.999706	0.999706	0.999691
Recall	0.999727	0.999864	0.999864	0.999932
FPR	0.00028	0.00034	0.00034	0.00038
FNR	0.000273	0.000136	0.000136	0.000068
Inference Time (s)	0.1948	0.1741	0.1848	0.8508

NB: FPR = False Positive Rate; FNR = False Negative Rate; s = seconds.

Table 5 shows that AEA-LGBM reduces false negatives to 1 while maintaining high true positives (14673), indicating near-complete detection of attack instances. This result is critical, as false negatives represent undetected attacks in intrusion detection systems. The slight increase in false positives reflects a trade-off, but the significant reduction in missed attacks demonstrates the effectiveness of the proposed error-aware learning mechanism.

Table 5: Confusion Matrix Comparison

Model	TP	TN	FP	FN
Default	14670	50018	14	4
Balanced	14672	50015	17	2
SMOTE	14672	50015	17	2
AEA-LGBM	14673	50013	19	1

Abbreviation: TP = True Positive (correctly detected attacks), TN = True Negative (correctly identified benign traffic), FP = False Positive (benign traffic misclassified as attack), FN = False Negative (attack misclassified as benign).

Figure 3 shows that all ROC curves are tightly clustered near the upper-left corner, with AUC values of approximately 1.000, indicating near-perfect separability between attack and benign traffic. This consistent behaviour across models suggests that performance is largely driven by the intrinsic characteristics of the CIC-DDoS2019 dataset rather than model-specific differences. However, the nearly overlapping curves indicate that ROC analysis lacks sufficient discriminative power in this context. Therefore, meaningful model differentiation must rely on error-sensitive metrics, particularly false negative rate (FNR), alongside computational measures such as inference time.




Figure 3: Receiver Operating Characteristic (ROC) Curve Comparison of LightGBM Variants for Intrusion Detection

Metric Sensitivity

Table 6 demonstrates that conventional metrics, including accuracy (0.999706–0.999722), F1-score, and AUC (~1.000), exhibit negligible variation, confirming metric saturation. In contrast, FNR varies by up to 75 per cent, and inference time ranges from 0.17 to 0.85 seconds, revealing substantial performance differences. These findings show that accuracy-centric evaluation masks critical trade-offs, thereby justifying the proposed error-sensitive and efficiency-aware evaluation framework for deployment-oriented model selection.

Table 6: Metric Sensitivity Analysis

Metric	Observed Range Across Models	Discriminative Power
Accuracy	0.999706 – 0.999722	Very Low
F1-Score	0.999353 – 0.999387	Very Low
AUC	0.999999 – 1.000000	Extremely Low
FPR	0.00028 – 0.00034	Moderate
FNR	0.000136 – 0.000273	Moderate
Inference Time	0.1811 – 0.8599 sec	High

Figure 4 (a–c) presents the integrated analysis of Error Efficiency Score (EES), inference time, and error rates across models. The results show that AEA-LGBM achieves the lowest FNR (~0.000068), indicating superior attack detection capability. However, this improvement is accompanied by a substantially higher inference time (~0.85s), resulting in the highest EES. In contrast, internally balanced and default LightGBM models maintain lower inference times (<0.20s) with slightly higher FNR. This combined analysis reveals a clear trade-off between detection sensitivity and computational efficiency, demonstrating that meaningful model differentiation emerges only when error and efficiency metrics are jointly considered.




Figure 4(a) Error Efficiency Score (EES) Comparison across Models




Figure 4(b) Inference Time Comparison across Models




Figure 4(c) False Positive Rate (FPR) and False Negative Rate (FNR) Comparison across Models

Independent Dataset Generalisation Analysis

Table 7 presents a comparative evaluation of the initial independent dataset validation experiment and the enhanced Adaptive Error-Aware LightGBM (AEA-LGBM) framework across datasets. Both experiments were conducted by training on the CIC-DDoS2019 dataset and evaluating on the unseen CICIDS2017 dataset to assess the model's robustness and generalisation capability under domain-shift conditions. Before independent evaluation, the CICIDS2017 dataset underwent preprocessing and cleaning procedures, including handling missing values, feature alignment, removal of invalid entries, and robust feature normalisation to ensure compatibility with the training dataset.

The enhanced AEA-LGBM framework incorporated adaptive error-aware weighting and robust feature normalisation to improve cross-dataset stability. Compared to the initial validation experiment, the enhanced framework achieved a substantial reduction in False Positive Rate (FPR) from 12.23 per cent to 5.10 per cent, indicating significantly improved control of false alarms under unseen network environments. Precision also improved considerably, from 81.60 per cent to 91.71 per cent, indicating greater confidence in detected attack instances.

However, despite these improvements, recall and false negative performance remained challenging, with the False Negative Rate increasing slightly from 58.61 per cent to 61.34 per cent. This indicates that although adaptive optimisation improved detection precision and robustness, the model continued to struggle to identify certain unseen attack behaviours due to differences in traffic distributions, feature characteristics, and attack structures between the two datasets.

Overall, the findings demonstrate that adaptive weighting and robust normalisation enhance deployment stability and reduce false alarms. Still, cross-dataset intrusion detection remains fundamentally constrained by domain shift and dataset heterogeneity. This further reinforces the study's contribution toward deployment-oriented and robustness-aware intrusion detection research for MANET environments.

Table 7: Comparative Independent Cross-Dataset Validation Results Before and After Adaptive Enhancement (Training: CIC-DDoS2019, Testing: CICIDS2017)

Metric	Initial Independent Validation	Enhanced AEA-LGBM Validation
Accuracy	0.614671	0.615323
Precision	0.816026	0.917119
Recall	0.413874	0.386645
F1-Score	0.549202	0.543963
AUC	0.627557	0.563505
False Positive Rate (FPR)	0.122250	0.050987
False Negative Rate (FNR)	0.586126	0.613355
Training Time (s)	11.524741	110.272781
Inference Time (s)	3.220465	6.899338
Error Efficiency Score (EES)		1.701841

Table 8 presents the performance of the proposed Transfer-Aware Adaptive Error-Aware LightGBM (AEA-LGBM) framework under a heterogeneous learning configuration in which training was performed using a combination of the CIC-DDoS2019 dataset and a partial subset of the CICIDS2017 dataset. At the same time, evaluation was conducted on a previously unseen subset of CICIDS2017. Unlike the fully independent cross-dataset validation presented in Table 4.11, this experiment adopts a transfer-aware, domain-aligned learning strategy to reduce the impact of dataset heterogeneity and feature-distribution mismatches between training and testing environments.

The framework integrates robust feature normalisation, stable feature selection, adaptive error-aware weighting, and heterogeneous multi-dataset learning to improve cross-environment generalisation. The results demonstrate substantial performance improvement across all evaluation metrics compared to the earlier independent validation experiments. In particular, the model achieved an accuracy of 99.90 per cent, precision of 99.88 per cent, recall of 99.94 per cent, and an AUC of 0.999984, indicating highly effective discrimination between benign and malicious traffic under the transfer-aware learning configuration.

Most importantly, the False Negative Rate (FNR) was reduced dramatically to 0.000580, while the False Positive Rate (FPR) remained relatively low at 0.001694. This indicates that incorporating heterogeneous training data significantly improved the model's ability to recognise previously unseen attack patterns while maintaining strong false-alarm control. The improvement suggests that exposure to partially aligned target-domain traffic during training enables the model to learn more generalised intrusion characteristics rather than relying solely on dataset-specific statistical behaviour. The stable feature selection mechanism further contributed to performance stability by retaining only

features that exhibit consistent statistical behaviour across datasets. A total of 54 stable features were preserved for model training, thereby reducing feature-distribution instability and improving domain compatibility between the training and testing environments.

Although the proposed framework required more training time due to the increased computational complexity of adaptive weighting and multi-dataset learning, inference time remained within practical limits for deployment-oriented intrusion detection scenarios. This demonstrates that improved cross-environment robustness can be achieved without excessively compromising operational feasibility. Overall, the findings in Table 4.12 demonstrate that transfer-aware heterogeneous learning substantially mitigates the domain-shift limitations observed during purely independent cross-dataset validation. The results therefore strengthen the practical contribution of the proposed AEA-LGBM framework by demonstrating that robust standardisation, stable feature learning, and heterogeneous, domain-aware training can significantly enhance generalisation in MANET intrusion detection environments characterised by diverse and evolving network traffic conditions.

Table 8: Transfer-Aware Heterogeneous Learning Performance of the Proposed AEA-LGBM Framework (Training: CIC-DDoS2019 + Partial CICIDS2017, Testing: Unseen CICIDS2017 Subset)

Metric	Transfer-Aware AEA-LGBM
Accuracy	0.998967
Precision	0.998840
Recall	0.999420
F1-Score	0.999130
AUC	0.999984
False Positive Rate (FPR)	0.001694
False Negative Rate (FNR)	0.000580
Inference Time (s)	6.414639
Training Time (s)	159.950648
Error Efficiency Score (EES)	1.283726
Stable Features Used	54

Comparative Cross-Dataset Generalisation Analysis of Baseline Models

Table 9 presents the comparative cross-dataset generalisation performance of the investigated machine learning models when trained on the CIC-DDoS2019 dataset and evaluated on the independent CICIDS2017 dataset. Unlike the internal evaluation results presented earlier, all models experienced substantial performance degradation on the unseen dataset, confirming significant domain shift and dataset dependency in intrusion detection systems.

The results show that although several models achieved near-perfect performance during internal testing, none maintained comparable reliability during cross-dataset evaluation. This finding demonstrates that high internal accuracy alone does not guarantee robust real-world generalisation capability. Variations in traffic distributions, attack characteristics, feature behaviour, and environmental configurations between CIC-DDoS2019 and CICIDS2017 significantly affected predictive performance across all models.

Among the evaluated algorithms, the Decision Tree model achieved the highest cross-dataset accuracy (62.41%), the highest recall (55.02%), and the lowest False Negative Rate (44.98%), indicating comparatively stronger adaptability under unseen network conditions. In contrast, several ensemble boosting methods, particularly LightGBM and XGBoost, experienced considerable reductions in recall performance despite their superior internal evaluation results. This suggests that models optimised heavily for dataset-specific traffic behaviour may struggle to generalise across heterogeneous intrusion environments.

The findings also reveal important deployment-oriented implications. Although KNN achieved moderate predictive performance, its extremely high inference time (36,255 seconds) renders it operationally impractical for real-time MANET deployment. Similarly, models with lower false-positive rates frequently exhibited significantly higher false-negative rates, reinforcing the trade-off between attack sensitivity and false-alarm control under cross-environment conditions.

Overall, the independent cross-dataset evaluation substantially strengthens the practical contribution of this study by demonstrating that intrusion detection systems must be evaluated beyond isolated internal datasets. The findings establish that generalisation capability, robustness to domain shift, and operational adaptability are critical requirements for real-world MANET intrusion detection deployment.

Table 9: Cross-Dataset Generalisation Performance of Machine Learning Models (Training: CIC-DDoS2019, Testing: CICIDS2017)

Metric	Logistic Regression	Naive Bayes	KNN	Decision Tree	Random Forest	XGBoost
Accuracy	0.576287	0.594007	0.602529	0.624147	0.608997	0.585794
Precision	0.576287	0.821002	0.849210	0.720955	0.872382	0.793674
Recall	0.361580	0.363345	0.363744	0.550243	0.363775	0.364368
F1-Score	0.491853	0.503750	0.509327	0.624137	0.513447	0.499446
AUC	0.403603	0.496126	0.631840	0.634582	0.620845	0.552714
False Positive Rate (FPR)	0.142410	0.103788	0.084621	0.279027	0.069721	0.124102
False Negative Rate (FNR)	0.638420	0.636655	0.636256	0.449757	0.636225	0.635632
Inference Time (s)	0.094746	1.912838	36255.042724	0.155378	0.572368	0.353440

DISCUSSION

The findings of this study demonstrate that conventional accuracy-centric evaluation is insufficient for meaningful assessment of intrusion detection in highly separable cybersecurity datasets. Although multiple machine learning models achieved near-perfect internal accuracy and AUC, substantial performance degradation was observed during independent cross-dataset evaluation on CICIDS2017. This indicates that strong internal benchmark performance does not necessarily translate into robust real-world intrusion

detection capability under heterogeneous traffic environments. The observed reduction in recall and increase in false negative rate during independent evaluation further confirm the impact of domain shift on intrusion detection reliability (Hemanand et al., 2026; Panigrahi et al., 2022).

The proposed Adaptive Error-Aware LightGBM (AEA-LGBM) framework addresses this limitation by explicitly integrating false-negative minimisation into the learning process via adaptive dynamic weighting. Unlike conventional approaches that primarily optimise aggregate predictive accuracy, the proposed framework prioritises attack-detection sensitivity, which is more operationally important in security-critical environments where undetected attacks can have severe consequences. The substantial reduction in false negative rate achieved during internal evaluation demonstrates the effectiveness of error-aware optimisation for improving intrusion detection sensitivity under imbalanced conditions. This finding directly addresses limitations identified in previous intrusion detection studies, in which error asymmetry is acknowledged but not explicitly integrated into adaptive learning mechanisms (Yang et al., 2020; Almseidin et al., 2017).

The study also demonstrates that model-level adaptive balancing can achieve equivalent detection performance to synthetic oversampling methods such as SMOTE while maintaining lower computational overhead. This finding suggests that data-level augmentation may not always provide additional practical benefit in highly separable intrusion environments where adaptive learning strategies are sufficiently effective. Consequently, the results support the growing perspective that imbalance handling should focus not only on correcting class distributions but also on adaptive error prioritisation (Liu et al., 2019).

Another important contribution of this work is the proposed error-efficiency evaluation framework that integrates the false negative rate, false positive rate, and inference time. Existing intrusion detection studies commonly evaluate predictive performance independently from computational efficiency (Tang et al., 2016; Kim et al., 2014). However, the results show that meaningful model differentiation emerges only when detection reliability and operational efficiency are jointly considered. The observed trade-off between false negative reduction and inference time highlights the importance of deployment-oriented evaluation for intrusion detection systems operating under real-time constraints. Feature optimisation was achieved to improve accuracy and reduce overfitting (Li et al., 2012).

More importantly, the independent cross-dataset experiments reveal that many intrusion detection models remain highly dependent on dataset-specific statistical characteristics. However, the transfer-aware heterogeneous learning experiments demonstrate that partial exposure to target-domain traffic substantially improves cross-environment robustness and restores generalisation capability under unseen traffic conditions. This finding establishes transfer-aware and domain-aware learning as important requirements for practical intrusion detection deployment in heterogeneous MANET environments.

Overall, this study contributes to the transition from benchmark-oriented intrusion detection to robustness-aware, deployment-oriented cybersecurity intelligence. Rather than focusing solely on maximising internal benchmark accuracy, the proposed framework emphasises operational reliability, computational feasibility, and cross-dataset generalisation. The novelty of this work, therefore, lies not only in the proposed AEA-LGBM model but also in establishing an integrated, error-aware, efficiency-sensitive, and robustness-oriented evaluation paradigm for intrusion detection systems operating in evolving real-world network environments.

5.0 CONCLUSION AND RECOMMENDATIONS

Conclusion: This study demonstrates that accuracy-centric evaluation is insufficient for intrusion detection in highly separable datasets. The proposed Adaptive Error-Aware LightGBM (AEA-LGBM) framework improves attack detection by explicitly minimising the false negative rate through adaptive weighting. Therefore, the model shows improvements in false negative rate, false positive rate, and inference time. The findings further show that effective intrusion detection assessment must jointly consider error rates and computational efficiency. Cross-dataset evaluation confirms the impact of domain shift on model reliability, while transfer-aware learning improves robustness under unseen environments.

Recommendations: Based on the study findings, future research may explore domain generalization, continual learning, federated intrusion detection, and explainable artificial intelligence for dynamic MANET security environments.

6.0 REFERENCES

1. Almseidin, M., Alzubi, M., Kovacs, S., & Alkasassbeh, M. (2017). Evaluation of machine learning algorithms for intrusion detection systems. *Proceedings of the IEEE International Symposium on Intelligent Systems and Informatics*, Subotica, Serbia, 277–282. <https://doi.org/10.48550/arXiv.1801.02330>
2. Bii, J. K., Rimiru, R., & Mwangi, R. W. (2022). OAAE: Optimized adaptive anomaly detection ensemble—Base model boosting by parameter optimization. *Engineering Reports*, 4(2), e12449. <https://doi.org/10.1002/eng2.12449>
3. Hemanand, D., Sheela, M. S., Sunitha, P., P, K., & Ganeshkumar, P. (2026). A smart and intelligent intrusion detection and prevention framework using BE2GO-SENN mechanisms for increasing the security of MANET. *Iranian Journal of Science and Technology, Transactions of Electrical Engineering*, 50(1), 223–256. <https://doi.org/10.1007/s40998-025-00891-6>
4. John, A., Isnin, I. F. B., Madni, S. H. H., & Muchtar, F. B. (2024). Enhanced intrusion detection model based on principal component analysis and variable ensemble machine learning algorithm. *Intelligent Systems with Applications*, 24, 200442. <https://doi.org/10.1016/j.iswa.2024.200442>
5. Kim, G., Lee, S., & Kim, S. (2014). A novel hybrid intrusion detection method integrating anomaly detection with misuse detection. *Expert Systems with Applications*, 41(4), 1690–1700. <https://doi.org/10.1016/j.eswa.2013.08.066>
6. Li, Y., Xia, J., Zhang, S., Yan, J., Ai, X., & Dai, K. (2012). An efficient intrusion detection system based on supportvector machines and gradually feature removal method. *Expert Systems with Applications*, 39(1), 424–430. <https://doi.org/10.1016/j.eswa.2011.07.032>
7. Liu, H., Lang, B., Liu, M., & Yan, H. (2019). CNN and RNN based payload classification methods for attack detection. *Knowledge-Based Systems*, 163, 332–341. <https://doi.org/10.1016/j.knosys.2018.08.036>
8. Novika, A. D., & Mujhidi, A. (2025). Cost-sensitive learning with LightGBM for class imbalance in intrusion detection systems. *Journal of EMACS Engineering, Mathematics and Computer Science*, 7(2), 147–153. <https://journal.binus.ac.id/index.php/EMACS/article/download/13435/5421>
9. Nyangaresi, V. O. (2022). Lightweight anonymous authentication protocol for resource-constrained smart home devices based on elliptic curve cryptography. *Journal of Systems Architecture*, 133, 102763. <https://doi.org/10.1016/j.sysarc.2022.102763>

10. Panigrahi, R., Borah, S., Pramanik, M., Bhoi, A. K., Barsocchi, P., Nayak, S. R., & Alnumay, W. (2022). Intrusion detection in cyber-physical environment using hybrid Naïve Bayes-decision table and multi-objective evolutionary feature selection. *Computer Communications*, 188, 133–144. <https://doi.org/10.1016/j.comcom.2022.03.009>
11. Tang, T. A., Mhamdi, L., McLernon, D., Zaidi, S. A. R., & Ghogho, M. (2016). Deep learning approach for network intrusion detection in software defined networking. *Proceedings of the International Conference on Wireless Networks and Mobile Communications (WINCOM)*, Fez, Morocco, 258–263. <https://doi.org/10.3390/electronics9091533>
12. Vaiyapuri, T., & Murugesan, K. (2026). Unsupervised one-class learning for IoT intrusion detection: A review of methods and prospects. *International Journal of Computer Networks and Applications*, 13(1), 179–188. <https://www.ijcna.org/Manuscripts/IJCNA-2026-O-10.pdf>
13. Yang, Y., Zheng, K., Wu, C., & Yang, Y. (2020). Network intrusion detection based on supervised adversarial variational auto-encoder with regularization. *IEEE Access*, 8, 42169–42184. <https://shorturl.at/HDebt>