

Innovative trust and privacy solutions for IOT devices in 5G environments

Ahmed Naeem Jasim^a, Hussein M. Jebur^b, Mohammed Yousif Arabi^c

^aDhi Qar Education Directorate. Open College of Education, Shatrah Branch, Iraq. Email: ahmedziady17@gmail.com

^b Education Directorate of Thi-Qar, Ministry of Education, Iraq. Email: cbook586@gmail.com

^c College of Technical Engineering, National University of Science and Technology, Dhi Qar, Iraq. Email: mohammed.y.arabi@nust.edu.iq

ARTICLE INFO

Article history:

Received: 18 /03/2026

Revised form: 04 /05/2026

Accepted : 07 /05/2026

Available online: 30 /09/2026

Keywords:

5G-IOT security, IDS Technology, cyberattack detection, Machine learning ensemble, XGBoost

ABSTRACT

This research study presents an ensemble-based intrusion detection framework for fifth-generation (5G)-driven Internet of Things (IoT) devices. It combines XGBoost, Random Forest, SVM, KNN and Logistic Regression with weighted soft voting in the model as well. The framework has been validated against the CSE-CICIDS2018 dataset, which encompasses different, state-of-the-art attack patterns associated with 5G-IoT systems. Normalisation, feature selection, and missing value imputation are performed in a preprocessing pipeline to improve generalisation. K-fold cross-validation ensures the robustness of the model; statistical testing confirms improvements in performance. The ensemble obtains an accuracy of 99.15%, an F1-score of 99.14% and an AUC of 1.00, outperforming each classifier alone. Privacy awareness is through flow-based features without inspecting Raw payloads. However, future work remains to be done on real-world validation and more intentional trust mechanisms.

MSC..

<https://doi.org/10.29304/jqcm.2026.18.32749>

1. Introduction

The intersection of the Internet of Things and 5G mobile networks is driving a revolution in the digital ecosystem, which will allow future self-supported scenarios like teleclinical diagnostics, cyber-physical factories, and robotised transport. The 5G platform establishes the operating window required by high-priority, low-latency IoT constellations such as those with ultra-low latencies and high data rates, along with megascale device populations [1-2]. However, this mix also makes the innovation happening among them a source of more security and privacy vulnerabilities. The layered and distributional aspects of the networks and the ever-growing number of endpoints and signal stacks enlarge the attack surface [3-4].

The 5G platform enables the operating range required for priority low-latency IoT constellations, including ultralow latency and a high number of I/Os at a megabits-per-second data rate and megascale device population [1-2]. However, this blend of innovation is also a recipe for an increase in security and privacy vulnerabilities. The multi-layer and distributional characteristics of 5G, as well as the constantly increasing amount of endpoints and signal stacks, also widen the attack surface [3-4].

*Corresponding author: Mohammed Yousif Arabi

Email addresses: mohammed.y.arabi@nust.edu.iq

Communicated by 'sub editor'

One limitation that has hindered the full-scale integration of distributed ledger and machine learning detection algorithms is insufficient maturity to process transactions in real-time, because the latency introduced by DLT transactions is still making it difficult. In the context of IoT networks for smart city infrastructures, anomaly detection has been investigated by Al Farsi, Ahmed et al. [6], who evaluated XGBoost and SVM classifiers using the CICIDS2017 benchmark. They find that their SVM model performs better than its gradient-boosted version. Nevertheless, the CICIDS2017 dataset was generated before 5G networks had been widely rolled out.

Thus, that model does not describe the traffic condition that the new generation architecture currently points to. In addition, ensemble-based classifiers or algorithms with provable sensing data privacy at the network level were not taken into account in the study.

Yuan et al. [7] proposed a random forest fog-layer intrusion-detection model, which has tuned the ensemble to the stringent latency requirements of 5G IoT deployments by aligning KNN classifiers jointly with logistic regression. However, the architecture had a high number of false positives in heavy traffic. It had an inadequate device behaviour audit schema because it did not have a heightened trust model to verify the integrity of the connected endpoint conduit's ecosystem.

Rayyis, Anfal Musadaq, et al. [8] proposed a deep ensemble model consisting of LSTM and CNN. The architecture made it possible to enhance privacy by anonymising the extracted vector features. It was with high precision of classification from the data. However, the computational cost of the solution is high, such that it cannot be deployed in low-end IoT devices.

These results demonstrate that there still exists a persistent void in the literature; practically all systems are inappropriate for deployment within end-to-end secure 5G scenarios, as their main goal towards attaining accuracy does not align with basic tenets of trust and privacy. Note that, interestingly enough, to the best of our knowledge, this corpus is also significantly devoid of investigation into the use of light yet robust ensemble classifiers (i.e., majority-voting principle-based) capable of providing a realistic balance between scalability and computational parsimony on one hand and prediction power on the other; hence, being deployable in practical infrastructures.

Even though machine learning and deep learning have been extensively used to perform IoT intrusion detection since their veracity of use in traditional schemes, there are still several associated challenges from machine learning/deep learning that need significant research attention in order to address them for 5G-enabled systems. Support Vector Machines (SVMs) with good classifier capabilities and no huge scalability or ability to adapt to 5G-IoT traffic are considered traditional models. Federated Learning (FL) enhances privacy by decentralising learning but comes with communication overhead, convergence issues on non-IID data and system complexity, making it impractical for resource-constrained devices. Although high accuracy is achieved using CNN and LSTM methods of a deep learning model, there is a high computational cost. This research presents a lightweight ensemble via soft voting for real-time 5G-IoT frameworks, which allows accurate as well as scalable intrusion detection systems with privacy preservation features.

The main contributions of this Study are the proposed privacy- and trust-preserving architecture, including secure data preprocessing, selective feature extraction, adaptive trust scoring, and robust anomaly detection. It evaluates five classifiers (XGBoost, Random Forest, KNN, Logistic Regression, and SVM) experimentally on a dataset with varying cyberattacks. Finally, it brings forward a new soft-voting ensemble engine and then exceeds the state-of-the-art SVM model directly with an accuracy of 99.15%, an F1 score of 99.14%, and an AUC of 1.00, respectively. SVM-only: 96.40% accuracy and 1.00 AUC, respectively. The work provides a benchmark reference for subsequent 5G-IoT cybersecurity research.

Despite the considerable progress achieved in machine learning-based network intrusion detection, the existing paradigms have suffered from two major drawbacks: i) a lack of lightweight ensemble architectures suitable for resource-constrained IoT devices, and ii) a general under-representation of privacy-aware compositions. Thus, a vital void still exists in designing efficient, scalable and privacy-aware IDSs, especially for 5G-enabled IoT ecosystems.

2. Dataset and Preprocessing

In this paper, the cybersecurity research community adopts the CSE-CICIDS2018 dataset as a standard [9]. It embeds an extensive set of modern attack features in actual user data, emulating various types of production networks. The data was gathered in a testbed during a test performed and developed to simulate aspects of the working conditions, including ad hoc user profiles and carefully prepared raid pathways. The dataset is especially useful for training and testing security models targeting the IoT end-devices in the emerging 5G ecosystem, as it intentionally encompasses traffic that is benign as well as malicious.

The dataset consists of about 16 million network flow records over multiple days, including benign and attack traffic. The dataset is imbalanced, as benign traffic constitutes more than 83% of the total data, and attack classes (e.g., DDoS, brute-force, and botnet) take up the remaining proportion. The dataset is chosen because it mimics the current traffic generation environment, containing modern attack scenarios and high-volume network behaviour, aligning it with 5G-IoT conditions that encompass high throughput per unit area, heterogeneous device types, and the nature of accessing networks during dynamic traffic states.

All attack techniques and DoSpack resources are tested against a broad range of actions: simulations like port scans, login attempts, DDoS with different strategies such as flooding, and the most challenging RAPE logic to mimic sophisticated multistaged infiltrations. Each of those attacks is relatively short and appears to increasingly target a specific vulnerability in IoT devices connected to 5G networks. A full set of flow-based features is provided for every entry, which encompasses timings (such as duration), overall size in packets or bytes, Inter-Arrival Times (IATs), and statistical sets to discriminate between benign and malicious activities.

The dataset applies a wide range of modern attack techniques in real user traffic, reproducing the diversity of production networks. Data was collected in a testbed experiment implemented to simulate aspects of the operational environment, including ad hoc user profiles and pre-configured (though sensibly constructed) raid paths. It is beneficial for use in constructing and training security models targeting IoT end devices in 5G networks, as this dataset intentionally includes benign traffic as well as malicious traffic.

The attack forms and DOSpack resources have all been proven to be wide-spectrum actions: scans, port scans, simulation attempts to log in (incl. dictionary attacks), and DDoS attacks of the aggressive types, such as flooding. The most challenging type of RAPE to replicate is a successful multi-staged attack in a DF invasion. Each of these attacks is fairly short-lived, and the campaign seems to be exploiting a specific weakness in IoT gadgets that are running on 5G networks.

3. Proposed Framework

To enhance the privacy as well as trust assurances of the 5G-integrated IoT devices. The structure of the proposed model is illustrated based on ensemble methods. The first step is to collect and preprocess the flows generated by the IoT terminals. The reference set of the traffic statistics stream is used in this preliminary study. In this stage, the initial packet captures are cleansed, converted, transformed, and recast so that the derived feature set genuinely and succinctly corresponds to the operational domain of the network.

In this work, a federated learning model by combining multiple base classifiers (e.g., XGBoost, RF, and RF feature sets) is proposed in the field of IoT security. The integrated model is proposed to identify the security anomalies and malfeasance among IOT networks, especially those caused by 5G network extension. The fusion ensemble unit is specially made for fusing the diverse signals and raising the decision to a higher level. Meanwhile, the base classifier produces suspicious probabilistic scores. The ensemble highlights the divergently accurate predictions and mitigates the idiosyncrasies of models by employing the artificial fusion strategies – weighted voting [3], soft score averaging and majority voting [5], respectively. The evaluations on the different test datasets of the performances of all matrices – accuracy, precision, recall, and F1-score – outperformed isolated classifiers. Therefore, the resulting system is a flexible, robust, and scalable intrusion detection architecture design that is capable of evolving as operational requirements change in 5G, transforming IOT ecosystems.




Fig. 1- Overview of Machine Learning Models Used for Intrusion Detection in 5G-IoT Environments

3.1 Machine Learning Models

XGBoost

XGBoost (eXtreme Gradient Boosting) [10] is a high-performance gradient boosting library. English Abstract In this strategy, various decision trees are built, and their results are added into consideration step by step. The objective of this algorithm is to optimise the differentiable loss function with L1 and L2 norm regularisation terms implemented in parallel during training in order to combat overfitting [11]. Its architecture builds on the out-of-core computing and sparsity-exploiting algorithms in breaks; it can even deal with extremely large-scale, high-dimensional data sets without sacrificing prediction performance.

Random Forest (RF)

The RF is a type of ensemble learning algorithm that uses multiple decision trees trained on bootstrapped training sets [12]. It combines the mean prediction of many trees and increases stability. The power of RFs to exhibit complex, non-univariate patterns without losing explanatory value stems from the bootstrapped observation sampling and feature subsampling [12-13], or “feature-bagging”, at each split that they perform.

K-Nearest Neighbours (KNN)

A naïve approach of labelling an instance by the label that most frequently appears in its K nearest points is good when the data is locally constructed. This distortion-inducing effect on classification accuracy is anticipated with distance-based classifiers, as they suffer from unnormalised features, irrelevant dimensions, and noisy data [14-15].

Logistic Regression (LR)

LR is a linear discriminative model between two classes. It is effectively taking output from an arbitrary binary classification model and converting it to a real number. This value is then fed into the logistic sigmoid function to produce a probability (between 0 and 1) [16]. The approach is formulated as computationally efficient, interpretable, and represented naturally in low dimensions. Regarding LR, the most possible operation is that which considers it can still be estimated by effectively a Bound: a hyperplane separating in the projection feature space [17].

Support Vector Machine (SVM)

SVM is a very well-known method that aims to generate a hyperplane that maximises margins between classes. Non-linear classification problems can be handled using kernel functions. Its margin-based regularisation enables it to avoid overfitting and is well-adapted to high-dimensional feature spaces in which few instances are available for training [18-19].

Ensemble learning Technique.

The basic concept of ensemble learning is to form a consensus from predictions made by multiple base classifiers, which are usually expected to be dissimilar, such as majority voting or weighted averaging [20]. This technique exploits the complementary nature of classifiers and, to reduce singleton model bias, derives a combination model that is more accurate and more reliable than any single classifier in Figure 2. Combined context-based nonverbal behaviour classifier: The combiner can be built from different sources of hypothesis generation processes.

3.2 System Setup

3.2.1 Feature Extraction

The detailed diagnostic traffic metrics were extracted, tapping on selected suspicious candidate time intervals from the preprocessed network records, to achieve a dual goal of enhancing detection accuracy and minimising computational overhead. More specifically, the flow duration accumulator, the inter-packet arrival time histogram, and the byte length histogram were chosen to implement these metrics, which turned out to play a vital role in perceiving the intrinsic nature of network traffic behaviour. Compression of the feature space shortened the database redundancy and irrelevant variability, which made the normal/abnormal behaviour distinction in the 5G-IoT environment glaring as well as laid a taproot for a functioning intrusion detection system.

3.2.2 Implementation

The model design was set up to preserve trained models against novel and previously unseen network conditions to guide how to train, deploy, and evaluate the intrusion detection architecture. The empirical evaluation was powered by the use of the CICIDS database since it comprises genuine network traffic traces and encompasses a broad and comprehensive range of attack types, e.g., botnet command-and-control dialogues, brute-force credential scanning, and distributed denial-of-service attacks. The multidimensional source of the attack spectrum makes it a critical benchmark for assessing the cybersecurity readiness of the systems. More concretely, the training data had to pass through an extensive preprocessing pipeline that involved continuous-valued feature normalisation by z-scores, redundant and uninformative feature dimension removal via recursive feature elimination, and missing data fill-in by k-nearest neighbour imputation. A wrapper-type feature selection algorithm was used to discriminate the most informative features, namely octet-length histograms, interarrival time differentials, and session duration, and their concatenated contribution emerged as a contribution of the dimensionality-cut stance that not only curtailed the algorithmic complexity but also enhanced detection accuracy.

Moreover, just like the normal 80%-20% split, we split it into the training and test sets. This approach is a common practice in machine learning to establish a training corpus that is large enough for the model to adequately converge while containing an echo-free holdout sample upon which to perform unbiased evaluations of generalisation by the model. The training corpus, constituting 80 % of the data, is employed for the initialisation of a model parameter.

The remaining 20% was used as new data and provided a strict standard for comparing prediction performance. To ensure robustness, K-fold cross-validation ($k = 5$ or $k = 10$) was used to reduce model bias and prevent overfitting. This would then guarantee that each model was trained and tested on multiple data partitions and, therefore, provided a more robust estimate of predictive variance. The hyperparameter optimisation stage for the key tuning parameters, the number of trees in the random forest estimator, the scale of regularisation strength used in logistic regression, the type of kernel applied with SVM, and the setup parameters for the extreme gradient boosting ensemble were executed through a grid set or a random search.

The input packet captures and netflow records are cleaned; feature information is extracted directly from them, thus rendering a set of scalars and feature vectors ready for use in supervised learning paradigms. The latter layer follows this as the joint learning of hyperparameters and discriminative model calibration. The feature space separation is done by each tuner in a “smart” manner: support vector machines aim to optimise the hyperplane (possibly very high-

dimensional hyperplanes) that separates the training examples with maximum margin; ensembles of trees, such as gradient boosting and random forests, build a hierarchical tree-based or rule-based partition. The last layer is some voting or stacking of the outputs, here tuned to average the outputs in all lead scores and return one confident label that says which points are most likely to be represented. Such abstract layered construction enables iteratively tuning local modules and interpretable decision-making paths while being scalable to large or even versatile network contexts.

The 20% data not involved in training the above classifiers was held out to test this classifier. Accuracy, precision, recall, F1-score, and AUC-ROC, the major performance measures, were also calculated in this subset. The ensemble approach reported the best results (AUC = 1.00, accuracy = 99.15%), and the benign traffic was almost completely separated from the malicious one. The AUC scores of 0.98 also further validate the good discrimination abilities for support vector (linear and all) and KNN classifiers. This is where the ensemble model excels over all other models since both the false positive and false negative overall incident rates are extremely low, as deduced from confusion matrix analysis. This is critical, particularly when it comes to 5G-IoT applications where security trust and service continuity must be established through a dependability-driven approach. Therefore, to some extent, based on the robustness of the potential test case or practical application under academic need, our simple approach provides reliable and reproducible results; The strict condition from the training and test phase leads to a high classification performance.

BEGIN XGBoost Algorithm INPUT: Training dataset (X_{train}, y_{train}), Test dataset (X_{test}, y_{test}) PREPROCESS: - Normalise/standardise features - Handle missing values - Encode labels (Benign=0, Malicious=1) INITIALIZE model with hyperparameters: - $n_{estimators} = 150$ - $max_depth = 6$ - $learning_rate = 0.1$ - $subsample = 0.9$ - $colsample_bytree = 0.8$ - $reg_lambda = 1.0$	- $random_state = 42$ - $use_label_encoder = False$ - $eval_metric = 'logloss'$ TRAIN model on X_{train}, y_{train}: FOR each boosting round: Compute gradient and Hessian Build a decision tree to minimise loss Apply regularization (L1/L2) Update ensemble prediction PREDICT on X_{test}: $y_{pred} = model.predict(X_{test})$ $y_{proba} = model.predict_proba(X_{test})[:, 1]$ OUTPUT: Accuracy, Precision, Recall, F1, AUC END
--	---

Algorithm-1. Pseudo-code of XGB Regression algorithm

BEGIN Random Forest Algorithm INPUT: Training dataset (X_train, y_train), Test dataset (X_test, y_test) PREPROCESS: - Standardise features (optional) - Handle missing values - Encode labels INITIALIZE model with hyperparameters: - n_estimators = 200 - max_depth = 10 - min_samples_split = 5 - min_samples_leaf = 2	- class_weight = 'balanced' - random_state = 42 FOR each tree in n_estimators: - Sample data with replacement (bootstrap) - Select a random subset of features at each split - Grow a decision tree without pruning COMBINE predictions via majority voting PREDICT on X_test: y_pred = majority_vote(trees' predictions) y_proba = average_probability_across_trees OUTPUT: Accuracy, Precision, Recall, F1, AUC END
--	---

Algorithm-2. Pseudo-code of the Random Forest algorithm

BEGIN KNN Algorithm INPUT: Training dataset (X_train, y_train), Test dataset (X_test, y_test) PREPROCESS: - Standardise features using StandardScaler - Encode labels SET hyperparameters: - n_neighbors = 5 - metric = 'minkowski', p = 2 (Euclidean distance)	FOR each test sample in X_test: - Compute distance to all training samples - Find k nearest neighbours - Assign class by majority vote - Compute probability as proportion of votes PREDICT: y_pred = predicted classes y_proba = predicted probabilities OUTPUT: Accuracy, Precision, Recall, F1, AUC END
--	--

Algorithm-3. Pseudo-code of the KNN algorithm

BEGIN Logistic Regression Algorithm INPUT: Training dataset (X_train, y_train), Test dataset (X_test, y_test) PREPROCESS: - Standardise features - Encode labels INITIALIZE model with hyperparameters: - C = 1.0 (inverse regularisation strength) - penalty = 'l2' - solver = 'liblinear' - max_iter = 1000	TRAIN: - Optimise weights using logistic loss + L2 penalty - Use gradient descent to minimise the cost function - class_weight = 'balanced' - random_state = 42 PREDICT: y_pred = sigmoid(dot_product(weights, X_test)) > 0.5 y_proba = sigmoid(output) OUTPUT: Accuracy, Precision, Recall, F1, AUC END
---	--

Algorithm-4. Pseudo-code of the Logistic Regression algorithm

BEGIN SVM Algorithm INPUT: Training dataset (X_train, y_train),	- class_weight = 'balanced' - random_state = 42
--	--

Test dataset (X_test, y_test) PREPROCESS: - Standardise features - Encode labels INITIALIZE model with hyperparameters: - kernel = 'rbf' - C = 10.0 - gamma = 'scale' - probability = Tru	TRAIN: - Find optimal hyperplane maximising margin - Use kernel trick for nonlinear separation - Solve the dual optimisation problem PREDICT: y_pred = model.predict(X_test) y_proba = model.predict_proba(X_test)[: , 1] OUTPUT: Accuracy, Precision, Recall, F1, AUC END
---	--

Algorithm-5. Pseudo-code of the SVM algorithm

BEGIN Ensemble (Soft Voting) Algorithm INPUT: Training dataset (X_train, y_train), Test dataset (X_test, y_test) PREPROCESS: - Standardise features (required for SVM & KNN) - Encode labels DEFINE base models: - model1 = XGBoost (as defined above) - model2 = SVM (as defined above) - model3 = Logistic Regression (as defined above)	INITIALIZE VotingClassifier: - estimators = [('xgb', model1), ('svm', model2), ('lr', model3)] - voting = 'soft' // average predicted probabilities - weights = [1, 1, 1] (optional) TRAIN ensemble on X_train, y_train PREDICT: y_proba_avg = (p_xgb + p_svm + p_lr) / 3 y_pred = argmax(y_proba_avg) OUTPUT: Accuracy, Precision, Recall, F1, AUC END
---	--

Algorithm-6. Pseudo-code of the Ensemble algorithm

3.3 Statistical Analysis

Statistical significance was estimated using the Wilcoxon signed-rank test for matched samples of performance differences between our ensemble model and baseline classifiers. The ensemble model demonstrated statistically significant improvements ($p < 0.05$) in all metrics relative to each of the individual models separately across accuracy and F1-score.

4. Results

4.1 Performance Metrics

The key performance indicators are listed in Table 2 and used to measure the accuracy of classification models.

Table 1- Performance Metrics Definitions and Formulas

Metric	Definition	Equation
Accuracy [21]	Calculate the existing accuracy of the model making the prediction against actual and predicted class labels, i.e., how well our model has predicted the output.	$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$

Metric	Definition	Equation
Precision [22]	Indicates the frequency that true positive samples are actually taken as negative. False positive/(false negative + false positive). The accuracy ratio of optimistic prediction is correct.	$\text{Precision} = \frac{TP}{TP+FP}$
Recall [23]	Assesses the capacity of the model to identify all relevant instances as it computes the percentage of true positives detected among all positive instances.	$\text{Recall} = \frac{TP}{TP+FN}$
F1-score [24]	Since precision and recall are not balanced measures, their harmonic mean is used as a metric that balances the false positives and false negatives.	$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$

Where:

- TP = True Positives
- TN = True Negatives
- FP = False Positives
- FN = False Negatives

4.2 Hyperparameters

We report the chosen important hyperparameters for the five heterogeneous classifiers that we employed in our framework in Table 1.

Table 2- Model Hyperparameters and Settings

Model	Key hyperparameter
XGBoost	n_estimators=150,max_depth=6,learning_rate=0.1,subsampling=0.9,colsample_bytree=0.8,reg_lambda=1.0
Random Forest (RF)	n_estimators=200,max_depth=10,min_samples_split=5,min_samples_leaf=2,class_weight='balanced'
KNN	n_neighbors=5,metric='minkowski',p=2,weights='uniform'
Logistic Regression (LR)	C=1.0,penalty='l2',solver='liblinear',max_iter=1000,class_weight='balanced'

SVM	kernel='rbf',C=10.0,gamma='scale',probability=True,class_weight='balanced
Ensemble (Soft Voting(	voting='soft', models = [XGBoost, SVM, LR], weights = uniform

These parameters were carefully tuned using grid search and random search strategies to retain the best accuracy in detections, with false positives kept at a minimum level, and to keep computational feasibility on resource-limited IoT endpoints. The classifiers are XGBoost, RF, KNN, LR, and SVM, respectively, with their own advantages to form the ensemble. The accuracy of the ensemble can be understood in terms of the oracle number of estimators, maximum tree size, learning rate regularisation strength, and neighbourhood size (demonstrated as 99.15% with AUC = 1 for the ensemble). Both parameters are set to very small values, which is beneficial for the system to achieve higher sensitivities and defences against different attack strategies in a dynamic 5G-IoT environment.

4.3 Experimental Setup

It can be seen from Table 3 that, uniformly, the ensemble model achieved an accuracy of 99.15% (95% CI: 98.9%–99.3%), confirming the stability and reliability of the results and enhanced precision, recall, and F1-score measurements. On the held-out test set, SVMs also outperform any other individual model setups based on accuracy and precision, with an accuracy of 96.40%, a precision of 96.52%, and a recall of 96.12% against all taxonomy single classifiers except for the F1 score, where full taxonomies were much higher (i.e., 97.83%). These outcomes reveal that the model is able to differentiate between clean and malicious packets, with a conservative balance in order not to raise false positive alarms but to detect actual attacks. Accordingly, the SVM achieved AUC = 1.00 with an available ROC curve; that means class separability over the whole range of probability threshold settings. Reliable anomaly detection under the upcoming scenarios of 5G-IoT deployments relies heavily on this attribute.

The KNN technique was found to be successful in minimising the number of false alarms and achieved an accuracy of 94.20% and a precision of 96.01%. In contrast, its recall rate showed 90.89%. Although this is acceptable, it is still lower than the current state-of-the-art methods. This fact suggests that, although the KNN is optimised for malignant packet generation, it still fails to capture a significant amount of true attack patterns.

With XGBoost, a score of 94.31% on the F1 was also generated, with a balanced precision of 94.63 and recall of 94.00; so much for the accuracy value of 94.90%. The latter dealt with the complicated feature interactions in network traffic and could generalise well. Its high classification success rate, especially in fast and reliable classification-based systems, qualifies it for intrusion detection.

The Random Forest achieved an F1 score of 91.40% and 92.40% accuracy on test data. The precision of RF and 89.78% recall also indicate that the subtler attacks that exist in the web traffic cannot be captured. RF: RF is also a tree-based model, but the model is an ensemble one, which might have its own diversity and may avoid overfitting. This suboptimal behaviour is likely caused by either a non-harmful drive/feature subset for differentiating the most informative dimensions or the “homogeneous” classification with all trees voting in an ensemble.

Table 3- Performance Comparison of Machine Learning Models in Intrusion Detection

Model	Accuracy	Precision	Recal	F1-score
XGBoost	94.90%	94.63%	94.00%	94.31%

RF	92.40%	93.09%	89.78%	91.40%
KNN	94.20%	96.01%	90.89%	93.38%
LR	82.20%	83.05%	80.00%	81.50%
SVM	96.40%	96.52%	96.12%	96.32%
Ensembel	99.15%	99.09%	99.19%	99.14%

Also, the Ensemble method of averaging the predicted probabilities from 5 different classifiers was very high in all performance measures with AUC = 1.00, accuracy = 99.15%, precision = 99.09%, recall = 99.19%, and F1-score = 99.14%.

These excellent results indicate that ensemble methods can exploit complementary behaviours of weakly correlated learners. The ensemble cancels out the weaknesses of each model. That is, it counterbalances the logistic regression's failure to model complex decision boundaries and the random forest's lower recall, as in failing to include enough discriminatory information, leading to a denser and stronger mapping of the input feature space into output classes.

The ensemble model generates a confusion matrix that captures the classification performance on the dataset, as illustrated in Figure 2. The matrix provides a detailed analysis of the algorithm's capacity to accurately classify both benign and adversarial flows by enumerating the counts of true positives, true negatives, false positives, and false negatives. The matrix's distribution within the binary classification framework demonstrates a distinct dominance of true positives and true negatives, suggesting that nearly every attack instance is identified, while benign traffic is as precisely assigned to the non-attack class. The false-positive and false-negative rates that result are negligible, suggesting that the model rarely misidentifies legitimate traffic as malicious. This minimises the occurrence of false alarms that could compromise operational continuity, while simultaneously achieving a high coverage of authentic attack vectors, thereby improving the system's overall resilience. Furthermore, the superior accuracy, precision, recall, and F1-score reported in Table 2 are further validated by the distinct diagonal density of counts that extends from the upper-left to the lower-right extremities of the confusion matrix. The precision rate of 99.09% is supported by the nearly nonexistent number of false positives. In comparison, the exemplary recall of 99.19% is substantiated by the similarly low rate of false negatives. Collectively, these metrics characterise the detector's performance as both resilient and equilibrated.

Figures 2–4 collectively demonstrate the model's strong classification performance, stability, and confidence separation between benign and malicious traffic. The successful integration into operational 5G-IoT networks will rely on these low false positive and negative rates known from commanding precision and recall. Its high threshold recall rate will guarantee that all positive attacks can be almost completely detected, ensuring a strong and reliable confidence in the security of the infrastructure. The lack of false alarms has, in the meantime, removed a great burden from system maintenance and operation, liberating operators from being troubled by unnecessary alarms. Moreover, the better instance rankability of ensemble learning is indicated by its non-zero area under the curve, between which there exists a threshold for distinguishing good from bad traffic, no matter what fixed value we use for practical detection.




Fig. 2- Confusion matrices for the models




Fig. 3- The Area Under the Curve (AUC) for the models

The histogram presents two probability density curves: the former corresponds to benign traffic (centred at zero); the latter is based on attacks (concentrated at one). The y-axis shows the absolute number of cases per confidence score, or normalised density. On the other hand, the x-axis represents how likely an event is to be

classified as an attack. The nearly ideal disjointness of these curves (with only a small overlap of cones at or near 0.5) reflects the confidence and accuracy levels of the model in classification. With such a clear cut-off, the false negatives in the attacks are minimised, and regular traffic can be classified relatively better, thus effectively avoiding the uncertain predictions.

The model's efficacy in reducing classification uncertainty and its strong generalisation capacity are emphasised by the significant separation in the distribution of confidence scores. These attributes are essential for the timely detection of security vulnerabilities in 5G-IoT ecosystems, where the relentless reduction of undetected assaults and false-alarm rates is of the utmost importance [25-27]. The quantitative findings summarised in Table 3 and illustrated in Figures 2 and 3 are intended to be reinforced by the accompanying histogram, which corroborates the assertion that the recommended architecture produces elevated confidence scores for every categorised instance while simultaneously enhancing aggregate classification accuracy.

A significant separation of confidence scores in this distribution indicates that the model is effective at reducing classification uncertainty and a more quantitative level of performance generalisation. These factors bear great significance in the early warning of security threats in IoT scenarios, since the asymptotical attenuation of missed attack rate and false alarm ratio is increasingly demanded [25-27]. The histogram visualisation supports the numerical results shown in Table 3 and Figures 2-3, i.e., that, without deteriorating the classification of each categorisation sample (each having high confidence scores), the proposed architecture enhances overall classification accuracy.




Fig. 4- Comparison of performance metrics across models

5. Discussion

The rapid proliferation of the Internet of Things (IoT) in fifth-generation (5G) networks will provide new opportunities for real-time, intelligent, and scalable applications such as smart cities, industrial automation, healthcare, or transportation. However, the mentioned fusion has brought in highly challenging security and privacy problems, which require efficient, adaptive, and low-cost IDSes. In this paper, we introduced a novel soft-voting ensemble model of XGBoost, Random Forest, Support Vector Machine (SVM), K-Nearest Neighbours (KNN), and Logistic Regression to address these challenges in 5G-IoT. Our method also achieved a remarkable accuracy of 99.15%, an F1-score of 99.14%, and AUC ($A = 1:0$) for the CSE-CICIDS2018 dataset, which are competitive enough to beat single classifiers and some recent state-of-the-art results reported in the existing

literature. This near-perfect AUC (1.00) can be an indication of overfitting due to the characteristics of your dataset. On the contrary, high validation scores across folds, in conjunction with cross-validation results, indicate good generalisation ability, but more real-world experiments are needed to confirm that.

Performance comparison Table 4 compares the performance of the recent proposed state-of-the-art machine learning and deep learning-based models used for intrusion detection in 5G-IoT-enabled environments with the ensemble model presented in this paper. A synthesis detailing the most relevant methodological aspects—from federated learning with differential privacy to hybrid deep architectures—and respective classification accuracies on different data sets is presented in the table. It is also noted that the soft-voting ensemble method outperforms previous existing methods and achieves an accuracy of 99.15% on the CSE-CICIDS2018 dataset, much higher than all other known works (i.e., from 94 to 98.42%). These superior results demonstrate the possibility of ensembles with heterogeneous classifiers—XGBoost, Random Forest, SVM, KNN, and LR—with a dynamic privacy- and trust-aware design that also adapts to the evolving threat landscape in 5G-IoT ecosystems. The ablation study further concludes that ensemble-based techniques offer a more robust, scalable, and low-computation approach for practical deployment over secure IoT infrastructures.

Table 4- Performance Comparison of Machine Learning Models for 5G-IoT Intrusion Detection

Study	Model	Method	Accuracy
[28]	Neural Network (NN) with Extended Kalman Filter (EKF)	EKF optimises NN learning by reducing estimation errors and improving convergence for intrusion detection in IoT networks.	94%
[29]	Federated Learning (FL) with Differential Privacy (DP)	Collaborative training of IDS on distributed industrial IoT devices without sharing raw data to ensure privacy and data security.	95.59%
[30]	Federated Learning (FL)	Collaborative intrusion detection model training across decentralised CPS clients, preserving privacy without sharing raw data.	96.5%
[31]	Hybrid deep learning model (MIX_LSTM) integrating LSTM and XGBoost	Uses XGBoost for feature selection and LSTM with an optimised loss function to handle imbalanced data and improve anomaly detection in IIoT.	98.42%
[32]	Multiple ML classifiers (KNN, XGBoost, Decision Tree, LDA, QDA)	Comparison of classifiers for IDS in intelligent city networks, with KNN performing best in detecting intrusions efficiently.	98%
[33]	Combined Deep Q-Learning (CDQL) with Spider Monkey Optimiser (SMO)	SMO selects optimal features; CDQL continuously monitors and detects anomalies in IoT smart city data, improving detection accuracy.	95%
[34]	Logistic Boosting, Random Forest, SVM	Comparative evaluation of these machine learning models for anomaly detection in IoT-driven factories.	Logistic Boosting: 96.6% RF: 94.3% SVM: 93.8%
Our	Ensembl	Security of 5G-IoT intrusion detection.	99.15%

Over the last few years, there has been increased research work for utilising machine learning and deep learning techniques in 5G-IoT systems. For instance, Senevirathna et al. [28] presented EKF-based DeepCaS, a 94% accurate intrusion detection in IoT, which eliminates estimation error and enhances convergence, and introduced a neural network for IoT intrusions along with the extended Kalman filter. Although their approach is useful for stabilising the training dynamics, it has two major drawbacks: 1) neural networks are computationally expensive, and 2) there is no explicit mention of trust or privacy preservation, which indeed becomes a must in the demand of 5G IoT environments. Similarly, Kulkarni et al. [29] proposed an FL-based DP scheme for industrial IoT and achieved an accuracy of up to 95.59%. Their method is privacy-preserving by not sharing raw data but suffers from high communication costs and non-convergence when the underlying distributions are non-IID, as in many heterogeneous IoT environments.

More recently, Ruzafa-Alcázar et al. [30] designed a privacy-preserving FL-based intrusion detection system (IDS) for CPS with 96% accuracy and 5%. Despite the fact that their architecture is built for decentralisation, they do not use ensemble techniques for bias and variance reduction that are used to guard an organisation from a multitude of different attack vectors. In contrast, Hazman et al. [31] proposed a mixed LSTM to select and balance a feature set under the IoT, for which cooperation of not only LSTM but also XGBoost was reported, as well as 98.42% accuracy that was considered state-of-the-art compared with our referred works. However, such model (sequential) and feature (high-dimensional) types are used to build their solution, so it is not suitable for resource-constrained IoT edge devices because expensive memory and computation overhead are required.

To remedy these weaknesses, we propose a simple yet effective fusion of strong layers from different classifiers with our ensemble technique. Our model can achieve arbitrary discriminative performance without such a strong decomposition, which is absent in deep architectures. The soft-voting mechanism presents the probabilistic predictions of individual base learners, which likely leads to fine-grained decision boundaries reflecting subtle variations in attacking profiles. This is ideal in 5G-IoT use cases where the traffic type is highly dynamic and latency needs and device capabilities differ. In addition, our preprocessing workflow—comprising z-score normalisation, recursive feature elimination, and k-NN imputation—is capable of preserving high data quality while facilitating scalability.

Remarkably, all of these excellent performers are still outperformed by our results. Aly and Behiry [34], for example, compared logistic boosting, random forest, and SVM as well to anomaly detection in IoT-enabled factories; logistic boosting obtained the best tracking performance at 96.6% accuracy, 2.55% points lower than our squad! Similarly, Shibu et al. [33] designed the CDQL-SMO model for smart city IoT security, which achieved only 95% accuracy. These differences do indeed prove the combined strength of the ensemble in utilising and blending complementary strengths: nonlinear capture in interactions for XGBoost, margin optimisation in complex pictures for SVM, and locally detecting trouble with KNN -- also by spreading individual deficiencies like overfitting (RF), miserable recall (LR), or distance shrinkage (KNN).

Moreover, our model considers implicit trust and privacy. Since the model operates on flow-based anonymised features, such as inter-arrival times, packet sizes, and session durations (not raw payloads), we do not compromise user privacy. In addition, the high precision (99.09%) and recall values (99.19%) remove false alarms, as well as missed attacks, which would be essential requirements for maintaining mission-critical 5G-IoT applications such as self-driving cars or remote surgery operations. The close to zero false positives also cut alert fatigue among security operators, which is a well-acknowledged problem in monitoring systems at scale.

For present purposes, it should be noted that the ultimate backend implementation is not one of a centralised system but rather a federated one (although this entails future work). Such a transformation could preserve accuracy and offer offloading of computation to the edge-only nodes – ensuring scalability as well as protection against point failures via redundancy and compliance with data sovereignty legislation. Using secure aggregation protocols and local differential privacy may increase confidentiality, but without compromising model utility, this is being explored in current research [29-30].

We intend to investigate in the future how the proposed ensemble-based anomaly-detection approach can be applied within a communication-efficient federated-learning algorithm, CREION_MLDL19. One possibility is to outsource model refinement from the data owners, because, on one hand, this would help in reducing sensitive data aggregation risk into a single site, and on the other hand, it becomes affordable by various types of devices

and deployments in 5G-IoT networks. With the help of appropriate privacy-preserving primitives (e.g., secure multiparty computation and differential privacy), we design the distributed training process such that, while the sensitive telemetry data is involved in the distributed training rounds, this data does not directly appear in any of these rounds. It would make the device identity private and promote a shared model.

While the model proposed shows great performance on benchmark datasets, there has been no real-world validation of this in 5G-IoT deployments so far. Future work will be deployed in live network environments and assessed over real-time traffic conditions

6. Conclusion

This paper presents a novel intrusion detection scheme on 5G-IoT for privacy and trust. Our method surpasses other existing state-of-the-art methods by making use of a soft-voting ensemble of five different classifiers and obtaining an AUC score of 1.00 and an accuracy rate of 99.15% on the CSE-CICIDS2018 dataset. The design also makes heavy use of on-the-fly trust scoring and privacy-preserving data conditioning techniques. All of these work together to detect advanced persistent threats efficiently and prevent the attack tools from being played among massive IoT ecosystems with ubiquitous 5G coverage at a large scale and to resist trust erosion and information leakage over a huge number and wide range of IoT deployments. Despite the overall centralised design and its lack of balance with scalability, the reduced computation of ensembles against DNNs may open a possibility in using them at low-capacity end nodes. Attention is warranted.

References

- [1] K. Anabi Hilary, et al., The four-C framework for high capacity ultra-low latency in 5G networks: A review, *Energies*, 12(18) (2019) 3449.
- [2] W. Mazurczyk, et al., Challenges and novel solutions for 5G network security, privacy and trust, *IEEE Wireless Communications*, 27(4) (2020) 6–7.
- [3] A. S. George, 5G-enabled digital transformation: Mapping the landscape of possibilities and problems, *Partners Universal Innovative Research Publication*, 2(3) (2024) 01–37.
- [4] S. Sicari, A. Rizzardi and A. Coen-Porisini, 5G in the Internet of Things era: An overview on security and privacy challenges, *Computer Networks*, 179 (2020) 107345.
- [5] M. Wazid, et al., Security in 5G-enabled Internet of Things communication: Issues, challenges, and future research roadmap, *IEEE Access*, 9 (2020) 4466–4489.
- [6] A. Al Farsi, et al., Comparative performance evaluation of machine learning algorithms for cyber intrusion detection, (2024).
- [7] Q. Yuan and Y. Lai, Towards efficient information retrieval in Internet of Things environments via machine learning approaches, *Journal of The Institution of Engineers (India): Series B*, 106(1) (2025) 363–386.
- [8] A. M. Rayyis, et al., Optimizing 5G resource allocation with attention-based CNN-BiLSTM and squeeze-and-excitation architecture, *Frontiers in Communications and Networks*, 6 (2025) 1629347.
- [9] W. Chimphee and S. Chimphee, Hyperparameters optimization XGBoost for network intrusion detection using CSE-CIC-IDS 2018 dataset, *IAES Int. J. Artif. Intell.*, 13(1) (2024) 817.
- [10] T. Chen, et al., XGBoost: Extreme gradient boosting, *R package version*, 1.4 (2015) 1–4.
- [11] J. Ma, et al., Application of the XGBoost machine learning method in PM2.5 prediction: A case study of Shanghai, *Aerosol and Air Quality Research*, 20(1) (2020) 128–138.
- [12] P. Probst, M. N. Wright and A.-L. Boulesteix, Hyperparameters and tuning strategies for random forest, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(3) (2019) e1301.
- [13] H. A. Salman, A. Kalakech and A. Steiti, Random forest algorithm overview, *Babylonian Journal of Machine Learning*, (2024) 69–79.
- [14] P. Cunningham and S. J. Delany, K-nearest neighbour classifiers: A tutorial, *ACM Computing Surveys*, 54(6) (2021) 1–25.
- [15] R. K. Halder, et al., Enhancing K-nearest neighbor algorithm: A comprehensive review and performance analysis of modifications, *Journal of Big Data*, 11(1) (2024) 113.
- [16] E. Y. Boateng and D. A. Abaye, A review of the logistic regression model with emphasis on medical research, *Journal of Data Analysis and Information Processing*, 7(4) (2019) 190.
- [17] B. Ghogh, et al., Locally linear embedding and its variants: Tutorial and survey, *arXiv preprint* (2020) arXiv:2011.10925.
- [18] D. A. Pisner and D. M. Schnyer, Support vector machine, *Machine Learning*, Academic Press (2020) 101–121.
- [19] A. Tedyyana, et al., Enhance telecommunication security through the integration of support vector machines, *Int. J. Adv. Comput. Sci. Appl.*, 15(3) (2024).
- [20] T. N. Rincy and R. Gupta, Ensemble learning techniques and its efficiency in machine learning: A survey, *Proc. IDEA Conf., IEEE* (2020).
- [21] K. O. Lye, S. Mishra and R. Molinaro, A multi-level procedure for enhancing accuracy of machine learning algorithms, *European Journal of Applied Mathematics*, 32(3) (2021) 436–469.
- [22] H. Luan and C.-C. Tsai, A review of using machine learning approaches for precision education, *Educational Technology & Society*, 24(1) (2021) 250–266.
- [23] W. Cullerne Bown, Sensitivity and specificity versus precision and recall, and related dilemmas, *Journal of Classification*, 41(2) (2024) 402–426.

-
- [24] A. Humphrey, et al., Machine-learning classification of astronomical sources: Estimating F1-score in the absence of ground truth, *Monthly Notices of the Royal Astronomical Society: Letters*, 517(1) (2022) L116–L120.
 - [25] Z. Chen, Z. Li, J. Huang, S. Liu and H. Long, An effective method for anomaly detection in industrial Internet of Things using XGBoost and LSTM, *Scientific Reports*, 14(1) (2024) 23969.
 - [26] H. N. Fakhouri, et al., A comprehensive study on the role of machine learning in 5G security: Challenges, technologies, and solutions, *Electronics*, 12(22) (2023) 4604.
 - [27] N. Gadey, S. D. Pande and A. Khamparia, Enhancing 5G and IoT network security: A multi-model deep learning approach for attack classification, *CRC Press* (2024) 1–23.
 - [28] T. Senevirathna, et al., A survey on XAI for 5G and beyond security: Technical aspects, challenges and research directions, *IEEE Communications Surveys & Tutorials*, 27(2) (2024) 941–973.
 - [29] D. D. Kulkarni, S. Rathore and R. K. Jaiswal, Intrusion detection system for IoT networks using neural networks with extended Kalman filter, *Proc. ICCCN, IEEE* (2021).
 - [30] P. Ruzafa-Alcázar, et al., Intrusion detection based on privacy-preserving federated learning for the industrial IoT, *IEEE Trans. Ind. Inform.*, 19(2) (2021) 1145–1154.
 - [31] S. A. Mahmud, et al., Privacy-preserving federated learning-based intrusion detection technique for cyber-physical systems, *Mathematics*, 12(20) (2024) 3194.
 - [32] C. Hazman, et al., A smart model integrating LSTM and XGBoost for improving IoT-enabled smart cities security, *Cluster Computing*, 28(1) (2025) 70.
 - [33] S. Shibu, et al., Anomaly detection using deep learning approach for IoT smart city applications, *Multimedia Tools and Applications*, 84(17) (2025) 17929–17949.
 - [34] M. Aly and M. H. Behiry, Enhancing anomaly detection in IoT-driven factories using logistic boosting, random forest, and SVM: A comparative machine learning approach, *Scientific Reports*, 15(1) (2025) 23694.