

Privacy-Preserving Federated Intrusion Detection for Electric Vehicle Charging Stations

Arunkumar Muniswamy

School of Computer Science Engineering and Information Systems
Vellore Institute of Technology
Vellore, Tamil Nadu, India
arunkumar.m2020b@vitstudent.ac.in

Rathi R

Associate Professor Grade 1
School of Computer Science Engineering and Information Systems
Vellore Institute of Technology
Vellore, Tamil Nadu, India
rathi.r@vit.ac.in

Abstract

Electric vehicles are increasingly integrated into the charging infrastructure, making them more vulnerable to cyberattacks. In this paper, a privacy-sensitive, scalable intrusion-detection system proposes applying federated learning to a distributed multi-layer perceptron (MLP) classifier implemented on five charging-site clients. Each client preprocesses local features, selects them, and updates the model, which is then averaged on FedAvg at the server. We test the system on a multiclass EV-charging system (CICEVSE2024) and demonstrate that the federated system can effectively detect various attack types, such as large-scale network floods, achieving a best test accuracy of 92 percent, good per-class classification, and high tolerance to extreme class imbalance. The results show that federated learning enables effective collaboration among sites with conflicting charging needs, allowing them to work together without sharing raw telemetry, thereby reducing privacy risks and communication overhead. The proposed framework can also be implemented in a decentralized manner across main mobility networks and provides methods for addressing heterogeneity and imbalance in real-world IDS tasks.

Keywords

Intrusion Detection System, Federated Learning, Electric Vehicle Charging Station (EVCS), Multi-Layer Perceptron (MLP)

1. Introduction

The electric vehicle network is being built at a very high pace, transforming the transportation framework and energy consumption. The number of electric vehicles sold to consumers around the world grew 25% last year to 17.8 million (IEA, 2024). By 2025, analysts forecast that sales will hit 21.3 million units, which will provide EVs with a 24% share of the world automotive market (BloombergNEF, 2024; Wensveen et al., 2025). It is estimated that the EV market is going to increase at a compound annual growth rate of 25.32, increasing to 6.16 trillion in 2035 compared to the current level of \$561.94 billion (BloombergNEF, 2024; Wensveen et al., 2025). The infrastructure needs to be charged accordingly to the increasing growth (Dai et al., 2024).

The charging infrastructure market is also growing very fast. The EV charging infrastructure market is expected to expand at a compound annual growth rate (CAGR) of 20.3 per cent with a 2022 to 2032 projection of increasing up to 100 billion dollars (Grand View Research, 2023; Dai et al., 2024). By late 2024, there will be more than 5 million electric vehicle chargers in the world, two times more than it was two years before (IEA, 2024; Grand View Research, 2023). To a large part, this expansion is government-led: the U.S. Bipartisan Infrastructure Law is planning to install half a million publicly accessible chargers in the country over the next five years, and the EU Alternative Fuels Infrastructure Regulation is compelling high-power chargers in every large highway every 60 kilometres (U.S. Department of Energy, 2024; European Commission, 2023).

Nonetheless, this has inadvertently exposed a large attack surface to hackers by building EV charging stations as central nodes in an ecosystem of cyber-physical vulnerability that is becoming more susceptible (Wang et al., 2023; Kalla et al., 2023; Hu et al., 2024). Emerging worrying statistics on cybersecurity events caused the necessity to address these security issues as soon as possible: In 2024, the number of cyberattacks on the automotive and smart mobility ecosystem increased by 39 percent compared to last year. The entire situation is the most alarming, as the battery charging infrastructure of electric vehicles became the target of 4-6 percent of all attacks in 2023-2024 (Upstream Security, 2024; Zuo et al., 2024). Practically all of the attacks (59 out of 10) have the potential to affect millions of devices at once (Upstream Security, 2024).

The features of charging stations of electric cars predetermine a complex set of security risks: DoS attacks that cause a lack of service and, in coordinated groups, disrupt the grid, as well as man-in-the-middle attacks that intercept or modify communications, expose bill payments, allow uncontrolled control of a vehicle, or falsify bills (Kalla et al., 2023; Zuo et al., 2024). Weakness of the protocols: An insecure protocol design can allow unauthenticated control messages, which can be used to launch a larger-scale denial-of-charge and load manipulation attack (Wang et al., 2023). Recent attacks that happened in November 2024 and affected an estimated 116,000 records illustrate the ease of breaching the app and backends that enable sensitive information such as names, locations, payment details, and vehicle data to leak (Zuo et al., 2024). Firmware attacks and supply chain attacks can also create backdoors or break battery management (Hu et al., 2024). When a small portion of EVs (around 2.5) are orchestrated, it may result in major violations of voltage that may trigger blackouts across the system (Wang et al., 2023).

The conventional signature-based cybersecurity solutions are incapable of keeping up with the dynamic and changing types of threats posed by EV charging infrastructure, exposing the networks to the attacks of a new type and zero-day threats (Hu et al., 2024; Kalla et al., 2023). With the rapid development of charging system deployment to achieve high goals, speed-to-market becomes more important than security to vendors, posing a risk (Wang et al., 2023; Zuo et al., 2024). It is thus necessary to have intrusion detection systems (IDS). Recent IDS combines the practices of machine learning and deep learning to detect abnormal patterns, classify malicious actors at a live session, and monitor the traffic between vehicles, chargers, and client systems to detect unauthorized access or manipulation of the charging data and protocols (Hu et al., 2024; Zuo et al., 2024; Kalla et al., 2023).

Federated learning is a privacy-preserving algorithm that enables websites to avoid the utilization of user-sensitive data or network data to jointly train formidable anomaly detectors, thereby reducing the utilization of centralized cloud data gathering (Zuo et al., 2024; Hu et al., 2024). In contemporary applications, IDS is typically a combination of signature-based attacks on recognized threats as well as anomaly-based structures of unidentified assault, enabling it to detect ransomware, automobile identity spoofing, brute-force access initiatives, and tampering with the Battery Management System (Hu et al., 2024). It has been found that such high accuracy of detection of highly-performing solutions is active when supported by properly chosen and optimized features (Kalla et al., 2023; Zuo et al., 2024). Also, with the help of IDS - NIST guidance and other frameworks propose the application of layered detection, penetration testing, and zero-trust principles to EV charging ecosystems, regulatory and standards goals become an opportunity (NIST, 2024; Zuo et al., 2024). Succinctly, developed IDS is a viable, topical safeguard that can be applied to maintain user privacy, service stability, and grid stability as the charging networks expand (Kalla et al., 2023; Zuo et al., 2024; Hu et al., 2024).

The application of Federated Learning (FL) to intrusion detection in EV charging infrastructures is quite universal because FL can perform the collaborative learning of models in distributed charging points without revealing raw users or billing data (Hu et al., 2024; Zuo et al., 2024; Muniswamy, A., & Rathi, R., 2024). The local confidential telemetry of the charging points, and the model updates are reported without violating the privacy and breaking laws (Hu et al., 2024). FL also consumes less bandwidth, reduces the load on the cloud, and can scale to thousands of

different devices, making site-specific traffic, and improving the possibility of identifying anomalies (Hu et al., 2024; Zuo et al., 2024). Making models harder to poison centrally and combining heterogeneous and localized knowledge, FL becomes less susceptible to new vectors of attack and new models (Hu et al., 2024; Kalla et al., 2023). FL manages to compromise data privacy, operation scalability, data performance, and data resiliency (Hu et al., 2024; Zuo et al., 2024).

1.1 Objectives

1. The objective of the project is to develop a federated learning-based intrusion detection system (IDS) to monitor electric vehicle charging infrastructure to maintain the privacy of the information and to afford collaborative model training among distributed clients using local data.
2. The objective of the project is to execute a multi-layer perceptron (MLP) classifier framework on five different clients, which will conduct integrated preprocessing and feature selection and use FedAvg algorithm to aggregate models to improve detection performance.
3. To experimentally test the proposed architecture using a dataset of multiclass EV charging stations, we prove that federated learning allows identifying different types of cyberattacks with a maximum accuracy of 92 percent despite the imbalance in classes.
4. To emphasize the new value of the framework: privacy-sensitive distributed learning, efficient management of heterogeneous and imbalanced data, and scalable operation to implement cybersecurity in intelligent mobility networks of critical importance.

2. Literature Review

To eliminate the single-point-of-failure nature of the traditional FL, Husnoo et al. (2024) suggested a decentralized federated anomaly-detection system that is suitable for smart grids using peer-to-peer gossip algorithms-Random Walk and Epidemic. The Random Walk protocol is also better than the Epidemic one as it saves time that would have been used on training (around 35 per cent) than the traditional FL without a reduction in high attack-detection accuracy. The tested system on the industrial control system test sets could successfully identify various cyberattacks with data leakage. The decentralized structure can also suppress communication latency and the straggling effect, which is rather appropriate in the geographically distributed implementation in smart grids (Husnoo et al., 2024).

Xie et al. (2025) proposed a federated deep-learning-based model to identify intrusion in smart grids based on cyber-physical systems, comprising a temporal convolutional network and multi-feature fusion to identify malicious activities, such as busbar and breaker protection failure, DoS attacks, and message suppression attacks. They suggest a gradient-compression technique based on an LSTM-TCVAE model under the federated learning environment, which potentially reduces computational and communication costs and increases efficiency and security of a resource-constrained smart-grid setting (Xie et al., 2025).

Shabbir et al. (2025) examined the way that smart grids can be effectively protected through the assistance of a federated-learning enhancement scheme coupled with a fusion-based anomaly detection scheme. They also offer a customized data-poisoning attack model of smart grids and a resource-conscious mitigation model. Their scheme, which gets improved by the fusion, makes them efficient in detecting various subtle attacks in operation in distributed smart grids (Shabbir et al., 2025).

Sharma et al. (2025) proposed the smart-grid-based intrusion detection and mitigation with the help of AI in EV charging stations. Their platform identifies both cyber-physical and software-based attacks on charging stations and uses AI-based anomaly detection, and supports optimization of energy through V2G. Those test results indicated that the accuracy of the test was 96.8, the recall rate was 96.0, the F1-score was 96.4, and the detection rate of the cyberattacks was 98.9, which facilitates the provision of real-time threat information and predictive countermeasures and scalable integration to commercial-in-use EV charging infrastructures (Sharma et al., 2025).

To resolve the privacy concern in FL-based Internet-of-Vehicle networks with low computation, Manh et al. (2024) considered the IDS of data sent via homomorphic encryption to offloaded data, making it possible to process the encrypted ciphertext on the server without direct decryption. Their system has a deviation of performance of less than 0.8% over the non-encryption-based solutions, and maintains privacy during offloading at the same time as the FL utility (Duc et al., 2024).

Qian et al. (2024) suggested a federated reinforcement-learning process (FedSAC) of EV charging control on the distribution networks, posing multi-EV charge/discharge as an MDP in optimal power flow. Using such data as the ISO New England prices, the NHTS2017 behavioural travelling, and an IEEE 74-bus network, their algorithm in FedSAC improved the variability of control, reduced power variation (the standard deviation decreased by 83.28 percent in the absence of the grid reward functionality). It ensured the privacy of the driver with the help of federated learning (Qian et al., 2024).

Karunamurthy et al. (2025) proposed an ideal FL-based IoT IDS, which uses the Chimp optimization algorithm as a feature selection tool and trains the classifiers using CNN through federated learning. Applied to a single MQTT benchmark dataset and yielded a detection rate of up to 95.59 results, which is 4-10 times better than traditional ML. The strategy would identify DDoS, brute force, malformed data, and attacks based on MQTT publish floods and SlowITe (Karunamurthy et al., 2025).

Beuran et al. (2025) proposed FedMSE a semi-supervised FL architecture in the IoT intrusion detection that acquires knowledge of new threats without jeopardizing secure learning. The combination of a Shrink Autoencoder-based feature teaching and either of the Centroid one-class classifier is used to obtain local learning on each gateway. The new MSE-based aggregation gives more importance to the local models, balancing between accuracy, efficiency, and robustness (Nguyen and Beuran, 2025).

Alamer (2024) suggested PPFL-SC, a privacy-preserving FL model that is suitable to malware detection in the IoT, and this uses group-oblivious signcryption (GOSC) with regard to securing model exchange and collaborative verification on the cloud-returned results. It involves a Stackelberg game-based incentive (CMT) which incentivizes participation, reduces the number of dropouts and ensures private and accurate local gradients and aggregated models (Alamer, 2024).

Deng et al. (2024) suggested a privacy-compliant and data-oriented risk-assessment framework of smart grid that incorporates secure encryption and deep learning within the framework of FL. They create a two-step risk indicator model and a federated protocol, and homomorphic encryption, which maintain the privacy of the model parameters during training. The method involves the use of deep CNNs to process system variables and predict the risk level, which eliminates the risk-data-sharing reluctance of the operators without losing the collaborative advantage of FL (Deng et al., 2024).

2.1 Research Gaps

Although the technical advancement has been significant in the papers that have been surveyed, standardization-oriented federated learning protocols targeted at the applications of critical infrastructures continue to be demanded. The interoperability of different federated learning systems with the current infrastructures of legacy systems remains a major challenge, which needs a systematic solution.

The greater part of the research is based on simulation environments or small datasets. Difficulties in how the system might not work in the real-world environment, where there are scalable deployments, such as heterogeneity in networks, variant devices, and government laws, are to be examined. Movement of the prototypes to a production quality system should undergo rigorous testing under diverse conditions. Although the majority of works take into consideration data poisoning attacks, the resilience of federation learning protocols to realistic adversarial attacks is an interesting area of research. Continued research in the next generation attack paradigm, and the defences are essential in ensuring the integrity of critical infrastructural installations.

3. Methods

Figure 1 presents an architecture that implements a federated intrusion detection pipeline using a lightweight Multi-Layer Perceptron (LightMLP) and the Federated Averaging (FedAvg) algorithm. The central dataset CICEVSE2024 is split into stratified training and test sets; the training partition is distributed to five simulated clients, each performing local preprocessing (missing value imputation, numeric scaling, label encoding, and feature selection). Each client trains a local LightMLP initialized from the current global weights and returns model parameters to a central aggregator after a fixed number of local epochs. The server aggregates updates by weighted averaging (FedAvg) to form a new global model; this cycle repeats for 20 communication rounds. Final evaluation uses the held-out test set

and reports per-class precision, recall, F1, and multiclass ROC-AUC. Implementation details: input→64→32→output MLP with dropout=0.3, Adam optimizer (lr=0.001, weight decay=1e-4), batch size 64, 5 local epochs, and 20 rounds. This federated IDS design disseminates a central collection to five customers.

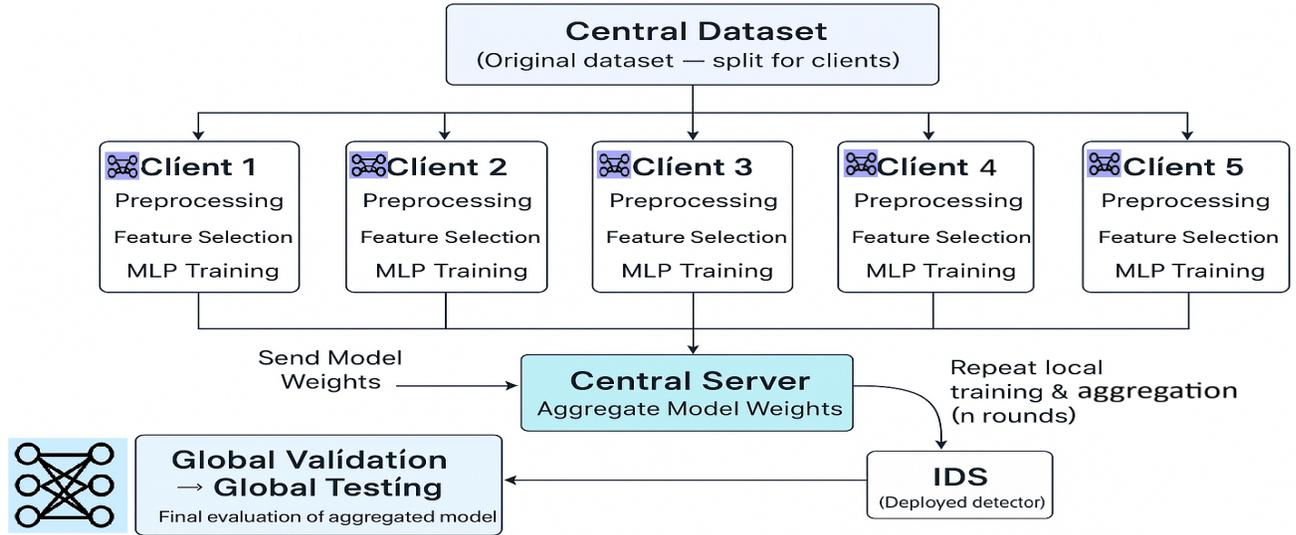


Figure 1: Federated IDS architecture

Every client preprocesses data, chooses features, and trains an MLP. Local model weights w_k^t are sent to the central server, which aggregates them using FedAvg, given in equation 1:

$$w^{t+1} = \frac{1}{K} \sum_{k=1}^K w_k^t \quad (1)$$

where $K = 5$ is the number of clients, after multiple aggregation rounds, the global model is validated and tested, and then deployed as an IDS detector for intrusion detection.

Algorithm 1 explains a federated learning (FL) protocol of a LightMLP intrusion detection model, which was already explained above.

Algorithm 1 - Federated MLP

Input: Set of the parameters — Seed (SEED), Number of Clients (N_CLIENTS), Test Ratio (TEST_RATIO), Batch Size (BATCH), Communication Rounds (ROUNDS), Local Epochs (LOCAL_EPOCHS), Learning Rate (LR)

Initialization:

```

1: set_seed(SEED)
2: create output directory
3: df ← read_csv("CICEVSE2024_NT.csv")
4: drop_na(df, "Label"); fill_na(df)
5: X ← numeric_features(df \ {"Label"}); X ← scale(X)
6: Y, CLASSES ← label_encode(df["Label"])
7: (Xtr, Xte, Ytr, Yte) ← stratified_split(X, Y, TEST_RATIO, SEED)
8: clients ← split_into_loaders(Xtr, Ytr, N_CLIENTS, BATCH)
9: model ← LightMLP(in_dim = num_cols(X), out_dim = len(CLASSES))
10: G ← model.state_dict()
// Training Loop
11: for r = 1 → ROUNDS do
12:   W_list ← []
13:   for each loader ∈ clients do
14:     cmodel ← new_model_like(model)
15:     cmodel.load_state_dict(G)
16:     w ← train_local(cmodel, loader, LOCAL_EPOCHS, LR)
17:     append(W_list, w)
  
```

```

18:   end for
19:   G ← average_weights(W_list)
20:   model.load_state_dict(G)
21:   record_accuracy(model, Xtr, Ytr, Xte, Yte)
22: end for
// Evaluation and Saving
23: (preds, probs) ← predict(model, Xte)
24: metrics ← classification_metrics(Yte, preds, probs)
25: save_plots_and_csvs(metrics, confusion_matrix(Yte, preds), OUTPUT)
26: save_model(model, OUTPUT + "global_model.pth")
end procedure

```

Algorithm 1 explains a federated learning (FL) protocol of a LightMLP intrusion detection model, which was already explained above.

4. Data Collection

The single set of data used in this paper will be CICEVSE2024, described as the multi-dimensional dataset used to measure the security of electric vehicle charging stations, which is described by Buedi et al. (2024), and the dataset can be found at the website <https://www.unb.ca/cic/datasets/evse-dataset-2024.html>. The data set consists of 1,277,367 flow entries, which are defined in terms of 42 characteristics (41 non-categorical or continuous). The only categorical variable, which is called as Label, is the target variable, the attack or benign labels or the multiclass attack. The rate of missingness is not greater than 30% per feature, and the cardinality rate of the categorical variable is low, which makes the data appropriate for both traditional machine-learning methods and deep-learning methods.

Correlation analysis using absolute Pearson correlation between each numeric attribute and an encoded representation of the Label indicates that there is a strong predictive indication in flow-timing-related and packet-statistics-related characteristics.

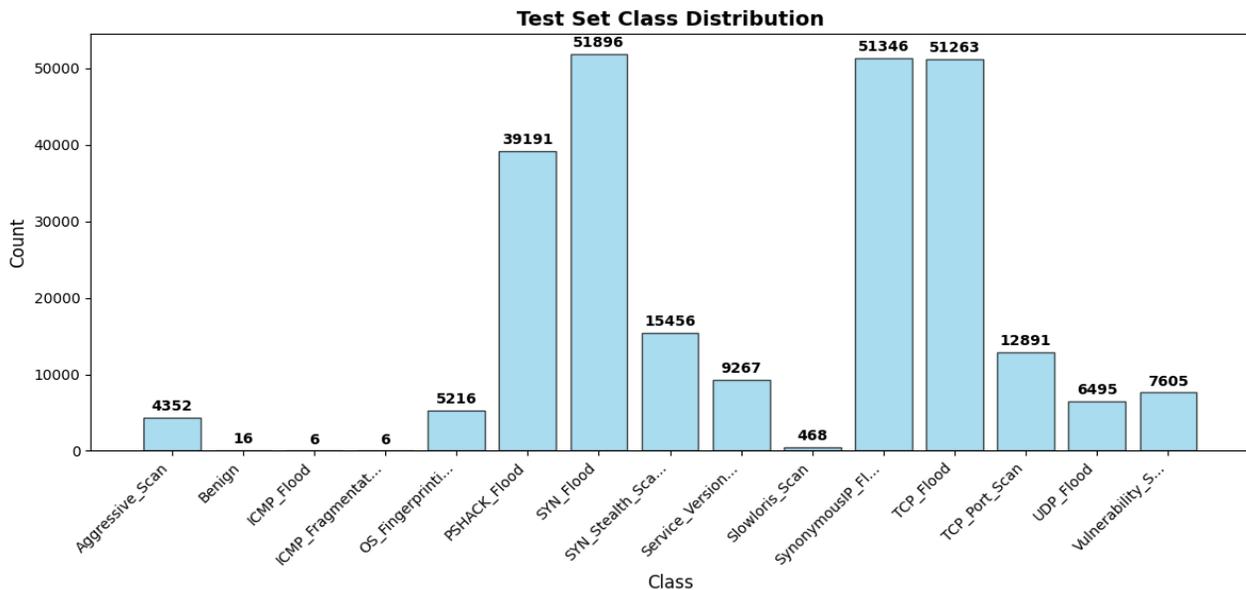


Figure 2: Class distribution in the test set (CICEVSE2024)

Attributes with the highest correlation are: dst2srclastseenms (≈ 0.581), dst2srcrstpackets (≈ 0.446), dst2srcmaxps (≈ 0.442), bidirectionalrstpackets (≈ 0.423), bidirectionalackpackets (≈ 0.386), and applicationconfidence (≈ 0.276).

These measures imply that session time, reset/ACK packet counts, packet-size distributions, and application-confidence variables are useful discriminants among the classes in question.

A bar plot was used to represent the skewed distribution of the test set, and the skewness was very high, as shown in Figure 2. The distribution is dominated by only a few classes of attacks, whose number is not large, although numerous labels are represented with minimal frequencies. The highest bars (around 50 52 -1 samples) are associated with SYNflood, SynonymousIPflood, and TCPflood, whose combination comprises a significant fraction of the data. The second level is made up of classes of mid-sized (tens of thousands down to about 9 -15), such as PSHACKflood, SYNstealthScan, and TCPportScan. Some of the classes take up the low-thousand space (e.g., AggressiveScan, UDPflood, VulnerabilityScan, OSfingerprinting). However, the most striking is the fact that some of the labels are near non-existent: Benign and some of the ICMP classes have one-digit values, and SlowlorisScan is represented by fewer than five hundred samples.

The overall performance metrics (average of performance, such as total accuracy) must be aggregated using micro-average (e.g. the most common classes), which will inevitably be skewed by the prevalent classes, which may produce the impression that is overly positive. A model that is effective in countering some of the most common SYN/TCP floods might fail catastrophically on rare but very significant incidents. In addition, the small number of test results results in a lot of noise and instability in AUC and recall estimation. To alleviate such problems, it is recommended to embrace per-class precision and recall, macro-average measures and it is also recommended to examine the confusion matrix. Also, some rebalancing schemes can be considered, including stratified validation, weighted loss, or focal loss to solve the issue of class imbalance.

5. Results and Discussion

The federated MLP achieves a global test accuracy of 91.95% (F1 \approx 91.54%). Per-class metrics (Table 1) show near-perfect detection for dominant flood classes (PSHACK_Flood, SYN_Flood, TCP_Flood, UDP_Flood) and weaker performance on minority classes (ICMP variants, OS_Fingerprinting) due to extreme class imbalance. Training and validation curves (Figure 5) show steady convergence with transient fluctuations typical of client drift in FL. The multiclass ROC micro-AUC is 0.9984, indicating strong overall separability across classes. Finally, the implementation contains logic that automatically determines the availability of a CPU or a GPU to be used in selecting the device.

We performed the experiments on the above-mentioned dataset on Python PyTorch and presented it both numerically and graphically, with explanations being provided in the subsequent subsections below.

5.1 Numerical Results

Table 1 presents the performance of the IDS model, evaluated based on Precision, Recall, F1-Score, AUC, and Support, as metrics of performance against various types of attacks. High-frequency attacks, such as PSHACKflood, SYNflood, SynonymousIPflood, TCPflood, and UDPflood, are also well detected by the model, achieving high scores (Precision, Recall, and AUC \approx 1). There are moderate performances of SYNstealthScan, ServiceVersionDetection, and TCPportScan, which reflect partial misclassifications. Other attacks, such as ICMP Flood, ICMP Fragmentation, and OS Fingerprinting, however, have low F1-scores, probably due to class imbalance and small sample sizes. In general, the model performs well on overpower attacks.

Table 1: Evaluation metrics of the different Attack types

Attack Type	Precision	Recall	F1-Score	AUC Scores	Support
Aggressive_Scan	0.269	0.332	0.297	0.9627	4,352
Benign	0.593	1	0.744	1.0000	16
ICMP_Flood	0	0	0	0.9716	6
ICMP_Fragmentation	0	0	0	0.9979	6
OS_Fingerprinting	0.528	0.13	0.209	0.9694	5,216
PSHACK_Flood	1	1	1	1.0000	39,191
SYN_Flood	1	0.997	0.999	1.0000	51,896

SYN_Stealth_Scan	0.629	0.908	0.743	0.9905	15,456
Service_Version_Detection	0.56	0.624	0.59	0.9824	9,267
Slowloris_Scan	0.954	0.31	0.468	0.9902	468
SynonymousIP_Flood	0.997	1	0.999	1.0000	51,346
TCP_Flood	1	1	1	1.0000	51,263
TCP_Port_Scan	0.848	0.738	0.789	0.9952	12,891
UDP_Flood	1	0.998	0.999	1.0000	6,495
Vulnerability_Scan	0.713	0.434	0.54	0.9858	7,605

5.2 Graphical Results

Figure 3 shows a multiclass confusion matrix, which provides a detailed description of the plot. A row is a real, true traffic/attack class, and a column is a model prediction. The darker, larger squares down the diagonal numbers: this means that the model has accurately proposed a large number of instances of that kind; pale or off-diagonal squares signify systematic errors (confusions).

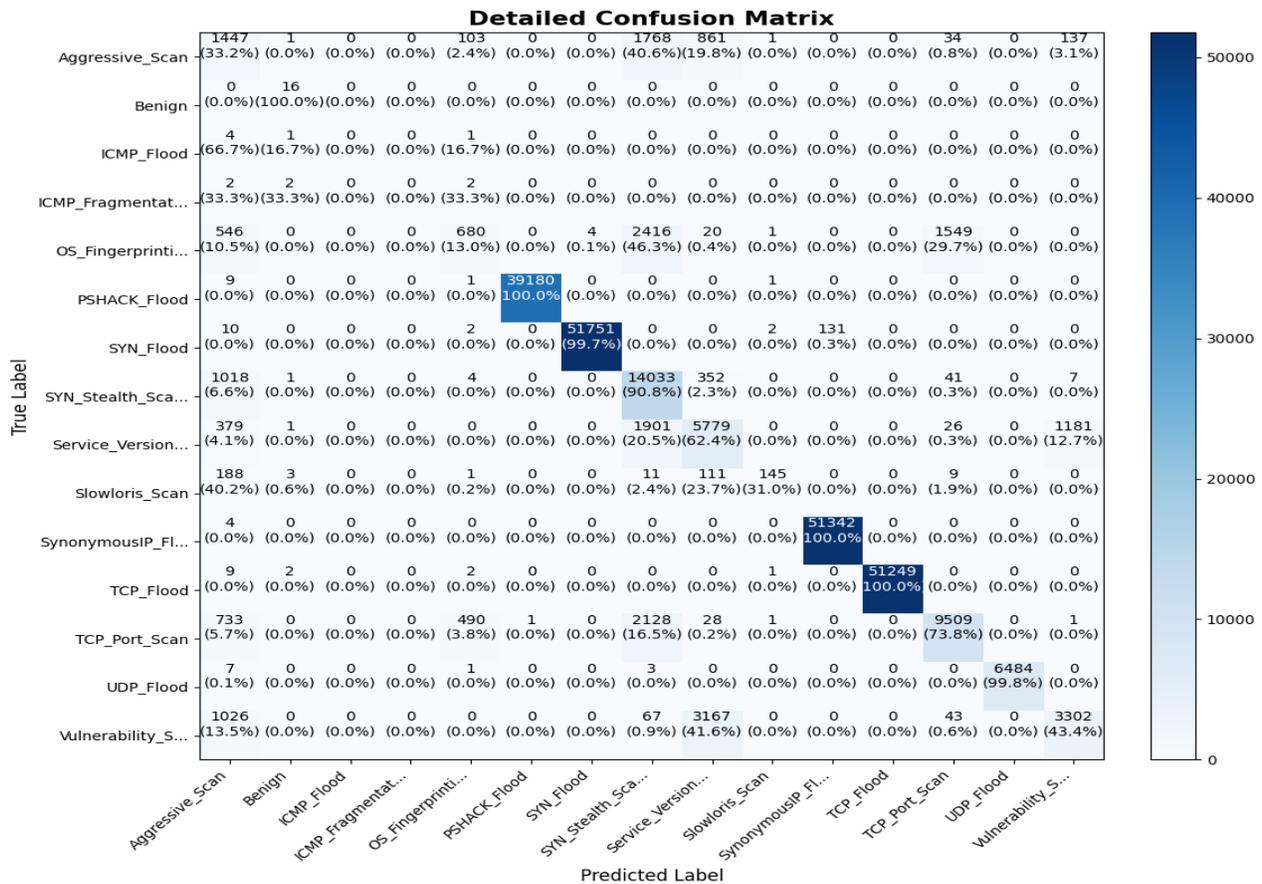


Figure 3: Multiclass confusion matrix (test set)

A few classifications of flood (PSHACK_Flood, SYN_Flood, Synonymous_Flood, TCP_Flood) have giant dark diagonal cells - those are reliably identified by this model (tens of thousands of successful predictions). In comparison to other classes, some classes are small (as seen in the example of a ratio of only 16 samples of the Benign class, which is ideally classified), indicating a drastic imbalance of the classes: a small segment can be ideally classified only because of the few examples.

The off-diagonal patterns demonstrate working inaccuracies: scans and fingerprinting attacks (Aggressive_Scan, SYNStealth_Scan, TCPPort_Scan, OSFingerprint, ServiceVersion, Vulnerability_Scan) intrude at each other, which implies the possibility of their features coinciding, or having particular packet behaviour. Slowloris scanning and some forms of ICMP are included with other predictions in a scattered manner, meaning poor recall.

Figure 4 shows one of the robust multiclass ROCs with a very high micro-average ROC (AUC = 0.9984); overall, the model can be considered to classify the classes with almost perfect accuracy. Most of the individual class curve concentrates around the upper-left part and quite a few of the courses have an AUC of 1.00 (Benign, PSHACK_Flood, SYN_Flood, Synonymous_Flood, TCP_Flood, UDP_Flood), meaning that the classifier is in perfect generic discrimination on the label set upon which validation is performed. The rest of the classes have high scores (e.g., ICMP_Fragmentation as 0.9979, TCPPort_Scan as 0.9952, SYNStealth_Scan as 0.9905), and some are also stronger but rather weak (Aggressive_Scan as 0.9627, ICMP_Flood as 0.9716, OS_Fingerprinting as 0.9694).

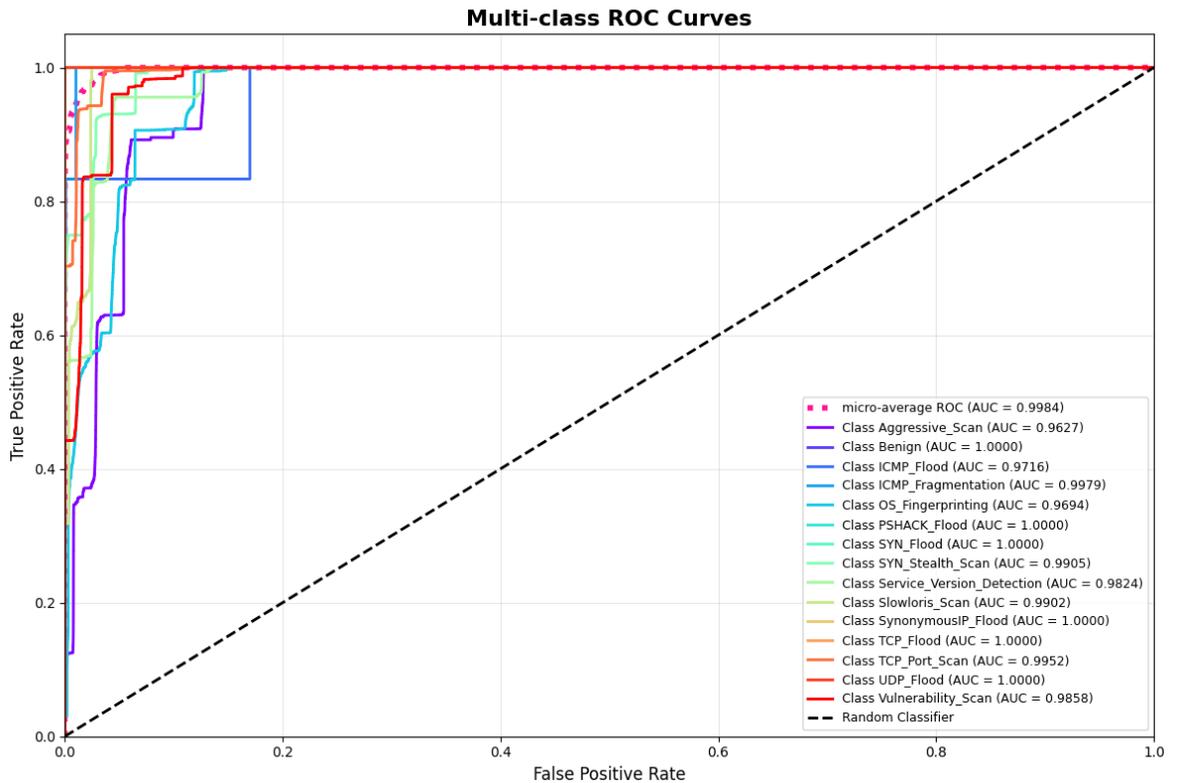


Figure 4: Multiclass ROC curves and micro-AUC (0.9984) showing global separability across classes.

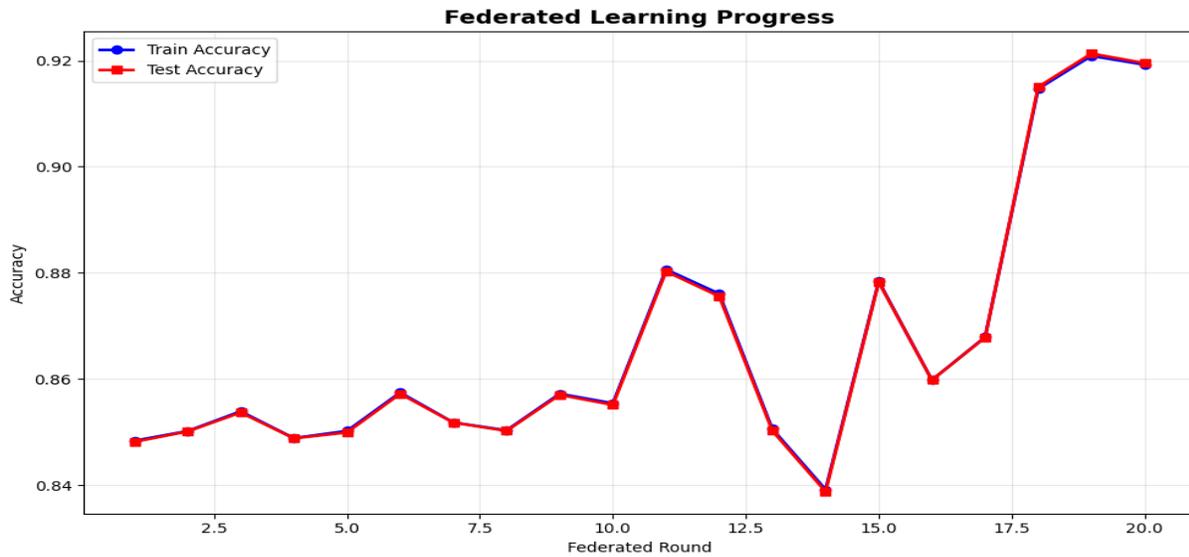


Figure 5: Training and test accuracy vs. communication rounds (20 rounds).

The sharp increases close to the y-axis are associated with extensive, accurate, favourable rates at very low false positive rates, which is an ideal feature of intrusion detection.

The results of the model shown in Figure 5 indicate that the accuracy of the model on both training and test data is better as the federation round increases to 20. Accuracies showed in the initial ten rounds have very limited variance and are quite proximate (around 0.84-0.86), meaning that the model taking place across the globe is learning steadily and generalizing favorably. Therefore, the training and testing lines barely overlap, which implies that there are few indicators of overfitting. We clearly have a jump at round 11 to some point of 0.88, and slightly reduced and then a steep fall at round 14 to 0.839. The interim training variations are typical of federated systems, and usually they signal temporary client drift (a different client pushing divergent updates), a change of the clients involved, or a local learning-rate / optimization phenomenon.

The model recovers ever since round 14, and its optimum accuracy occurs around round 18-19 (approximately 0.92). The latter advancement implies that the aggregation approach is ceding, so that the updates that the clients obtain are becoming more stable or the hyperparameters (a reduced learning rate, more local epochs, or better client sampling) allow the model to learn functional patterns. As the train and test curves are correlated across, the end high accuracy is probably an actual improvement in generalization, as opposed to overfitting.

The suggested federated IDS decreases the centralization of raw telemetry, which decreases privacy risk and cloud storage costs. Nonetheless, trade-offs in deployment need to be taken into account: (i) communication overhead - periodic model changes occupy bandwidth and cost of operation; client selection can be applied through model compression or sparsification to cut down the communication overhead; (ii) latency - model updates at the client end are local exhibiting low latency, but model updates at the global end rely on network round trip time as well as the aggregation period (design choice); (iii) scalability - FedAvg scales to a large number of clients, but stragglers and heterogeneity would necessitate client selection or as there should be a cost- benefit analysis of cloud costs against centralized training and additional complexity in orchestration of FL as well as possible benefits in privacy and regulatory compliance.

5.3 Proposed Improvements

Table 2 indicates that federated learning is almost superior to all other baseline approaches, with an accuracy of about 91.9% and an F1 score of about 91.5%. Compared to them, however, only the two oldest models (Random Forest and KNN) have about 75% performance (Buedi et. al., 2024). The large performance difference indicates that the federated pipeline can be more generalized. The overall updating weights of model parameters on a pool of clients by using FedAvg can efficiently use diverse local data across partitions and therefore avoid overfitting to a given partition and

increase coverage across classes. In addition, the technique has useful features, such as privacy (raw data are stored on the client) and resistance to local noise.

Table 2: Comparison of Evaluation metrics with the proposed model and Existing Models

Model	Accuracy	F1 Score	Precision	Recall
Decision Tree	71.57	71.60	71.63	71.57
KNN	75.11	74.43	74.40	75.11
Adaboost	50.81	50.59	53.57	50.81
MLP	54.59	52.43	52.16	54.59
Naive Bayes	63.61	59.28	56.40	63.62
Logistic Regression	57.75	48.98	44.15	57.75
Random Forest	75.43	74.98	74.83	75.43
SVM	70.61	65.63	73.70	70.61
Federated Learning	91.95	91.54	92.26	91.95

Table 3: Comparison of this work with similar studies

Study (Author & Year)	Model/Approach	Domain/Dataset	Accuracy (%)	F1-Score (%)	Remarks/Improvement Observed
Hu et al. (2024)	Federated MLP IDS	EV Charging Stations	90.3	89.7	High detection but moderate generalization; limited heterogeneity handling
Zuo et al. (2024)	Privacy preserving FLIDS	Smart EV Infrastructure	91.2	90.9	Focus on privacy preservation using FedAvg and encryption
Karunamurthy et al. (2025)	Optimal FL with Chimp Optimization	IoT Intrusion Detection	95.6	94.8	Strong feature optimization but higher computational cost
Nguyen & Beuran (2025)	FedMSE (Semi-supervised FL)	IoT Networks	94.7	94.1	Maintains quality with minimal data sharing
Shabbir et al. (2025)	Fusion-Enhanced FL	Smart Grid IDS	96.2	95.4	High adversarial resistance with energy-aware defense
Proposed Study (2025)	Federated MLP (FedAvg, 5 clients)	EV Charging Stations	91.95	91.54	Balances privacy preservation, scalability, and communication efficiency

Table 3 is a summary of the recently conducted federated intrusion-detection research as well as the proposed study. In both EV and IoT spheres, the accuracy lies between 90.3 % and 96.2 %, and the F1 scores follow the same pattern; in fact, the accuracy of Hu et al. (2024) is 91.2 % with solid detection in EV charging data but less coverage of heterogeneity, whereas FedAvg and encryption ensure the privacy of the results in Zuo et al. (2024), with 90.3 % accuracy. Karunamurthy et al. (2025) achieve better performance (95.6) by using federated learning coupled with Chimp optimization on feature selection at the cost of computational cost. FedMSE by Nguyen and Beuran (2025) indicates that semi-supervised FL can achieve the quality of models without requiring significant label sharing, and Fusion-enhanced FL, as Shabbir et al. (2025) reveal, can increase adversarial resilience and energy-aware defences. Our federated MLP (2025) achieves an impressive 91.95% accuracy and 91.54% F1 and offers the advantage of privacy, scalability and communication efficiency, and further improvements in the architecture or feature-selection should be made to address the outstanding performance gap and support the deployment of our solution in heterogeneous EV fleets.

6. Conclusion

We introduced a federated MLP with privacy guarantees on EV charging infrastructure intrusion detection and tested it on the CICEVSE2024 dataset, resulting in an overall accuracy of 91.95% (F1= 91.5) and maintaining privacy of client data. The findings illustrate that FL is viable in practical use in distributed EVCS monitoring, particularly when the attack is of a high volume of floods. To go to practice, future design should deal with adversarial security (model poisoning and backdoor attacks), economical communication (compression and sparsification), and energy-related training (minimizing on-machine compute and power consumption). On-device anomaly scoring and the hybrid architectures (local deep feature extractors) with federated aggregation to enhance minority-class detection are also to be explored in order to reduce the operational overhead and improve the detection of minority classes.

References

- Alamer, A., 'A privacy-preserving federated learning with a secure collaborative framework for malware detection models using Internet of Things resources', *Internet of Things*, 25, pp. 101015–101015, 2024. doi: <https://doi.org/10.1016/j.iot.2023.101015>
- BloombergNEF, *Electric Vehicle Outlook 2024*. BloombergNEF, 2024, [online] Available at: <https://about.bnef.com/electric-vehicle-outlook/> [Accessed 14 Oct. 2025].
- Brownlee, J., *Micro, Macro & Weighted Averages of F1 Score, Clearly Explained*, Towards Data Science, 2020. Available at: <https://towardsdatascience.com/micro-macro-weighted-averages-of-f1-score-clearly-explained-b603420b292f/> (Accessed: 15 October 2025).
- Buedi, E.D., Ghorbani, A.A., Dadkhah, S. and Ferreira, R.L., July. Enhancing ev charging station security using a multi-dimensional dataset: Cicevse2024. In *IFIP Annual Conference on Data and Applications Security and Privacy* (pp. 171-190). Cham: Springer Nature Switzerland, 2024.
- Dai, H., Wang, C., Zhang, X., Liu, T. & Lu, J., 2024. Recent advances and trends in electric vehicle charging infrastructure and energy storage. *Energy Reports*, 10, pp.1367–1383.
- Deng, S., Zhang, L. and Yue, D. (2024) 'Data-driven and privacy-preserving risk assessment method based on federated learning for smart grids', *Communications Engineering*, 3(1). doi: <https://doi.org/10.1038/s44172-024-00300-6>
- Duc, M.B., Nguyen, C.-H., Hoang, D.T. and Nguyen, D.N. (2024) 'Homomorphic Encryption-Enabled Federated Learning for Privacy-Preserving Intrusion Detection in Resource-Constrained IoV Networks', arXiv.org, arXiv:2407.18503. Available at: <https://arxiv.org/abs/2407.18503>
- European Commission, 2023. Regulation (EU) 2023/1804 on the deployment of alternative fuels infrastructure. *Official Journal of the European Union*, L234, pp.1–50.
- Evidently AI (2025) *Accuracy, Precision, and Recall in Multiclass Classification*, EvidentlyAI. Available at: <https://www.evidentlyai.com/classification-metrics/multi-class-metrics> (Accessed: 15 October 2025).
- Grand View Research, 2023. *Electric Vehicle Charging Infrastructure Market Size, Share & Trends Analysis Report*. Grand View Research, [online] Available at: <https://www.grandviewresearch.com/industry-analysis/electric-vehicle-ev-charging-infrastructure-market> [Accessed 14 Oct. 2025].
- Hu, R., Liang, B., Han, D. & Wang, F.-Y., 2024. A federated learning-based intrusion detection model in EV charging infrastructure. *IEEE Transactions on Transportation Electrification*, 10(2), pp.1650–1664.
- Husnoo, M.A., Anwar, A., Haque, M.E. and Mahmood, A.N., 'Decentralized Federated Anomaly Detection in Smart Grids: A P2P Gossip Approach', arXiv.org, arXiv:2407.15879, 2024. Available at: <https://arxiv.org/abs/2407.15879>
- IEA (International Energy Agency), 2024. *Global EV Outlook 2024*. IEA, [online] Available at: <https://www.iea.org/reports/global-ev-outlook-2024> [Accessed 14 Oct. 2025].
- Kalla, A., Chockalingam, S.P., Bansal, A. & Venugopal, K.R., 2023. Security threats and detection approaches in EV charging stations for smart grids. *IEEE Internet of Things Journal*, 10(24), pp.21508–21520.
- Karunamurthy, A., Vijayan, K., Kshirsagar, P.R. and Tan, K.T. (2025) 'An optimal federated learning-based intrusion detection for IoT environment', *Scientific Reports*, 15(1). doi: <https://doi.org/10.1038/s41598-025-93501-8>
- Muniswamy, A., & Rathi, R. (2024). A detailed review on enhancing the security in Internet of Things-Based Smart City Environment using Machine learning algorithms. *IEEE Access*. <https://doi.org/10.1109/ACCESS.2024.3450180>
- Nguyen, V.T. and Beuran, R. (2025) 'FedMSE: Semi-supervised federated learning approach for IoT network intrusion detection', *Computers & Security*, pp. 104337–104337. doi: <https://doi.org/10.1016/j.cose.2025.104337>

- NIST (National Institute of Standards and Technology), 2024. Cybersecurity Framework Profile for Electric Vehicle Extreme Fast Charging Infrastructure. NIST, [online] Available at: <https://csrc.nist.gov/publications/detail/nistir/8473/final> [Accessed 14 Oct. 2025].
- Qian, J., Wang, Y., Zhang, Y., Xu, H., Wu, Q., Li, Y., Niu, Q. and Zhang, Q. (2024) ‘Federated Reinforcement Learning for Electric Vehicles Charging Control on Distribution Networks’, IEEE Internet of Things Journal, 11(3), pp. 5511–5525. doi: <https://doi.org/10.1109/jiot.2023.3306826>
- Shabbir, A., Manzoor, H.U., Zoha, A. and Halim, Z. (2025) ‘Smart grid security through fusion-enhanced federated learning against adversarial attacks’, Engineering Applications of Artificial Intelligence, 157, p. 111169. doi: <https://doi.org/10.1016/j.engappai.2025.111169>
- Sharma, A., Rani, S. and Shabaz, M. (2025) ‘Artificial intelligence-augmented smart grid architecture for cyber intrusion detection and mitigation in electric vehicle charging infrastructure’, Scientific Reports, 15(1). doi: <https://doi.org/10.1038/s41598-025-04984-4>
- U.S. Department of Energy, 2024. National Electric Vehicle Infrastructure (NEVI) Program Update. U.S. Department of Energy, [online] Available at: <https://www.energy.gov/nevi> [Accessed 14 Oct. 2025].
- Upstream Security, 2024. Automotive Cybersecurity Report 2024. Upstream Security, [online] Available at: <https://upstream.auto/automotive-cybersecurity-report-2024/> [Accessed 14 Oct. 2025].
- Wang, Y., Yuan, Y., Sun, M., Zheng, H. & Liu, H., 2023. Cybersecurity and privacy challenges in electric vehicle charging infrastructure: Threats, vulnerabilities, and countermeasures. IEEE Transactions on Smart Grid, 14(5), pp.4865–4878.
- Wensveen, P.J., Wiedmann, T.O. & Lenzen, M., 2025. The environmental footprint of electric vehicles: A global assessment for 2024–2035. Nature Communications, 16, article 1554.
- Xie, R., Wang, B. and Xu, X. (2025) ‘A novel federated deep learning for intrusion detection in smart grid cyber-physical systems’, Engineering Applications of Artificial Intelligence, 162, p. 112404. doi: <https://doi.org/10.1016/j.engappai.2025.112404>
- Zuo, R., Wu, H., Wang, R., Wang, H. & Zhou, A., 2024. Federated intrusion detection for smart EV charging: A privacy-preserving approach. IEEE Transactions on Industrial Informatics, 20(2), pp.2301–2312.

Appendix A — Evaluation metrics and aggregation methods

In each case, c is assumed to be the positive class with all other classes being assumed negative. The equations below (2) up to (8) demonstrate various evaluation metrics. Let:

TP_c = true positives of class c , FP_c = false positives of class c , FN_c = false negatives of class c , n_c = the true number of samples of a single class (Brownlee, J. 2020, Evidently AI 2025).

Precision for class c (how many predicted c were correct):

$$Precision_c = \frac{TP_c}{FP_c + TP_c} \quad (2)$$

Recall for class c (how many true c were found):

$$Recall_c = \frac{TP_c}{FN_c + TP_c} \quad (3)$$

F1-score for class c (harmonic mean of precision & recall):

$$F1_c = 2 \cdot \frac{Precision_c \cdot Recall_c}{Precision_c + Recall_c} = \frac{2 \cdot TP_c}{2 \cdot TP_c + FP_c + FN_c} \quad (4)$$

Let K be the number of classes, $N = \sum_c n_c$, and class support $w_c = n_c/N$.

Macro average — treat all classes equally (simple mean of per-class scores):

$$Precision_{macro} = \frac{1}{K} \sum_{c=1}^K Precision_c, Recall_{macro} = \frac{1}{K} \sum_{c=1}^K Recall_c, F1_{macro} = \frac{1}{K} c = \sum_{c=1}^K F1_c \quad (5)$$

Weighted average — weight per-class score by class support w_c :

$$Precision_{weighted} = \sum_{c=1}^K w_c \cdot Precision_c, F1_{weighted} = \sum_{c=1}^K w_c \cdot F1_c \quad (6)$$

Micro average — sum raw counts, then compute metric (global):

$$TP_{micro} = \sum_c TP_c, FP_{micro} = \sum_c FP_c, FN_{micro} = \sum_c FN_c \quad (7)$$

$$Precision_{micro} = \frac{TP_{micro}}{TP_{micro} + FP_{micro}}, Recall_{micro} = \frac{TP_{micro}}{TP_{micro} + FN_{micro}} \quad (8)$$

Note: for single-label multiclass predictions,

$$Precision_{micro} = Recall_{micro} = \text{overall accuracy.}$$

Biographies

M. Arunkumar received a bachelor's degree from REC (Anna University Affiliated), Walajapet, Vellore, in 2005, and a master's degree from Vel Tech Dr. Rangarajan Dr. Sagunthala R & D Institute of Science and Technology, Avadi, Chennai, in 2012. He is currently a Research Scholar with the School of Computer Science, Engineering, and Information Systems at the Vellore Institute of Technology, Vellore, Tamil Nadu, India. He has over fourteen years of teaching experience and has published review articles and research papers in SCI journals. His research interests include artificial intelligence, security, the Internet of Things, and machine learning.

R. Rathi received her M.E. in Computer Science and Engineering from Anna University, India, in 2006 and B. Tech from the University of Madras, Tamil Nadu, India, in 2003. Currently, she serves as an Associate Professor in the School of Computer Science, Engineering, and Information Systems at Vellore Institute of Technology, Vellore, India. She received her PhD at Vellore Institute of Technology University, Vellore, India. She has over 15 years of experience in teaching and research, and has published around 20 papers. Her research interests include Rough sets, Knowledge discovery databases, Genetic algorithms, deep learning, machine learning, and federated learning.