PRINCE: A Proactive Client Selection in Federated Learning for Connected and Autonomous Vehicles

Amanda Lopes[1], John Sousa[1], Lucas Bastos[2], Lucas Pacheco[1], Iago Medeiros[1], Denis Rósario[1], and Eduardo Cerqueira[1]

[1]Federal University of Para (UFPA)

[2]Federal University of South and Southeast of Para (UNIFESPA)

{amanda.lopes, john.sousa, lucas.pacheco}@itec.ufpa.br,

bastos.lucas@unifesspa.edu.br, {iagomedeiros, denis, cerqueira}@ufpa.br

Abstract. *Federated Learning (FL) enables cooperative training among Connected and Autonomous Vehicles (CAVs) while preserving data privacy. However, the volatility of vehicular environments, characterized by frequent link interruptions and high mobility, poses a significant obstacle to system robustness, often leading to client failures (e.g., connection, resource, aborts) that degrade global model performance. In this paper, we introduce PRINCE (Proactive Reliability-driven INtelligent Client sElection), a framework that integrates stochastic mobility modeling directly into the FL decision-making loop. In its operation, PRINCE synergizes Shannon Entropy to quantify the informational value of local data with a probabilistic mobility model to proactively filter unstable nodes before selection. Evaluation results demonstrate that PRINCE achieves a final accuracy of 83.90% and a training success rate of 61.32%. Crucially, our approach outperforms state-of-the-art reactive baselines, delivering gains of up to 9.22% in accuracy and a* $3.5\times$ *improvement in resource efficiency.*

1. Introduction

The implementation of 5G networks enables the transition of CAVs from theoretical concepts to operational systems, primarily by enabling the generation and manipulation of substantial volumes of data [Elbir et al. 2022]. For example, through multimodal sensors, CAVs can perceive their surroundings and make intelligent decisions by integrating locally obtained data [Chellapandi et al. 2023]. In this sense, CAVs consider a set of sensors to feed Machine Learning (ML) models for various predictions, such as collision risk assessment and maneuver prediction of nearby vehicles [Deva Hema and Rajeeth Jaison 2024]. These models allow the vehicle's comprehension of its surroundings, enabling CAVs to precisely interpret the environment and make informed decisions during driving [Pan et al. 2024]. This capability allows vehicles to perceive the environment and use independent decision-making capabilities while continuously exchanging information with other vehicles and the infrastructure [Lobato et al. 2022].

There is a considerable privacy concern with CAV, as vehicle data have the capacity to disclose sensitive information about CAV, drivers, and passengers. In addition, the amount of data required for ML training might consume network bandwidth. In this context, Federated Learning (FL) emerges as a promising paradigm for taking advantage of the computational capabilities of vehicles without compromising vehicle privacy or network bandwidth. In this sense, FL acts as a privacy protection layer for CAV applications by deploying model data packets instead of raw data packets, which also reduces the utilization of communication resources [Chellapandi et al. 2023]. Specifically, during each communication round, the server selects a particular subset of CAV to receive the global model, participate in training using their local data, and exchange model parameters without exposing raw sensing data [Lobato et al. 2024]. Subsequently, a designated policy merges the shared local models into an improved global model. This methodology not only strengthens privacy but also enhances training efficiency in a distributed, interconnected environment.

Despite the advantages of FL, the highly dynamic nature of the vehicular environment introduces critical variables that can severely degrade the performance of the global model [Gutierrez et al. 2024]. Client failures, caused by network instability, insufficient computational resources, or abrupt topology failures, represent major obstacles to system scalability and convergence [Zhang et al. 2021]. In this sense, clients might not provide their local model updates, interfering with FL's ability to learn effectively, and thus a given CAV client with high-quality data can become a model detractor if it is unable to complete the training and transmission cycle [Vardhan et al. 2025]. In addition, clients might experience different failure rates due to their heterogeneous composition. Most traditional FL approaches operate reactively to these failures, wasting communication resources and processing time by selecting clients that subsequently fail [Maroua 2024].

Recent works have sought to optimize the client selection process by focusing on data heterogeneity (*i.e.*, non-independent and identically distributed: Non-IID), where [Sousa et al. 2025] introduced an entropy-based method to identify diverse data and used the Minimal Repair Models (MRM) to replace failed clients reactively. Although effective in describing data representativeness, such methods operate predominantly reactively or focus only on instantaneous diversity, neglecting the historical behavioral stability of the node within the network [Chen and Vikalo 2024]. Hence, to our knowledge, there is a significant gap in the literature on client selection mechanisms that consider the need for information-rich data (*i.e.*, high entropy) alongside robustness guarantees against the volatility inherent in vehicular mobility [Smestad and Li 2023].

This paper introduces a Proactive Client Selection mechanism that integrates stochastic mobility modeling into the FL decision-making process, called PRINCE. The main contribution of this paper is the introduction of a hybrid scoring algorithm that penalizes predicted failure risks and rewards entropy, applying adaptive computing policies to maximize convergence in hostile scenarios. Unlike traditional approaches that react to failures, PRINCE addresses the problem through three complementary steps to anticipate network failures. The first step consists of filtering based on the Received Signal Strength Indicator (RSSI), which serves as a line of defense against unsuitable nodes by removing candidates that do not meet the minimum link stability requirement. The second step focuses on reliability, where PRINCE considers a Markov chain model to analyze interaction history and calculate transition probabilities between states (*i.e.*, success; network connection failure; resource failure; abort decision), allowing the system to predict and avoid unstable vehicles before selection. The third step focuses on information quality, in which PRINCE uses Shannon entropy to assess the relevance of local data. The simulation results show that PRINCE outperforms traditional approaches, demonstrating relevant performance. Furthermore, PRINCE has proven capable of maintaining training continuity while ensuring model convergence, even in unfavorable dynamic situations.

The remainder of this paper is organized as follows: Section 2 lists the works relevant to this research. Section 3 details the proposed PRINCE approach and its mathematical foundations. Section 4 presents and analyzes the simulation results in hostile scenarios. Finally, Section 5 concludes the topic of the paper.

2. Related Work

Previous studies have extensively examined the challenges of FL, focusing primarily on data heterogeneity and client instability. [Zhu et al. 2021] analyzed the impact of Non-IID data on local devices, noting that this characteristic significantly degrades the performance of the model compared to centralized learning. The authors reviewed strategies for addressing attribute skew and label skew, but emphasized that existing algorithmic solutions still struggle to close the performance gap in dynamic scenarios involving frequent client failures.

[de Souza et al. 2024] proposed ACSP-FL, which uses distributed accuracy to measure global model convergence. ACSP-FL focuses on selecting clients that require more training (lower accuracy) and employs processing time and transmission parameters to qualify available resources. Although effective in terms of communication efficiency, this approach reacts to client metrics at the time of selection without necessarily predicting future availability or ensuring statistical diversity of data in unstable vehicular environments. Specifically, ACSP-FL measures processing time and transmission parameters reactively, assuming that current resource availability persists throughout training, whereas our Markov-based approach predicts temporal state transitions to proactively avoid clients that are likely to fail mid-round due to mobility-induced disconnections.

[Huang et al. 2022] investigated the problem of joint selection considering effective participation and fairness. They introduced E3CS, a stochastic selection strategy that seeks to balance training efficiency with the reduction of bias caused by failures. The experimental results showed that E3CS accelerates convergence compared to traditional methods. However, E3CS focuses on mitigating the effects of failures during aggregation rather than avoiding the selection of failure-prone nodes. We do not include E3CS in our experimental comparison because its Multi-Armed Bandit formulation requires a priori knowledge of arm rewards, which is not directly applicable to time-varying vehicular scenarios, and it addresses fairness post-selection through aggregation weighting rather than reliability pre-selection through mobility prediction.

[Sousa et al. 2025] proposed the ECS-HDSR mechanism that integrates entropy-based selection with a Minimal Repair Model (MRM). ECS-HDSR prioritizes data diversity by selecting clients with high entropy and addresses failures by replacing disconnected clients with similar nodes identified via Hausdorff Distance. Despite improving robustness and accuracy in scenarios with high dropout rates, it operates under a reactive paradigm: it repairs the system by replacing the client after the failure occurs, which can introduce latency and computational costs during the search for substitutes.

Table 1 summarizes the analyzed approaches. We observe that existing strategies operate predominantly under a reactive paradigm, addressing data heterogeneity but measuring link stability reactively (after selection) rather than proactively predicting temporal stability before selection. While ECS-HDSR considers link quality during reactive repair, it does not integrate mobility prediction into the initial selection decision, leading to wasted resources on clients likely to fail. In contrast, this paper introduces a proactive selection strategy that integrates Markov Chain modeling with Shannon Entropy to anticipate instability. By predicting reliability rather than reacting to failures, our method selects, from the outset, a set of clients that maximizes both informational quality and the probability of training completion.

Table 1. Comparison of Client Selection Approaches in Vehicular FL

Method	Selection Paradigm	Mobility-Aware	Data Quality	Link Stability
ACSP-FL [de Souza et al. 2024]	Reactive (Accuracy)	✗	✗	✗
E3CS [Huang et al. 2022]	Stochastic (Fairness)	✗	✗	✗
ECS-HDSR [Sousa et al. 2025]	Reactive (Repair)	✗	✓ (Entropy)	✓
Ours	**Proactive (Predictive)**	**✓ (Markov)**	**✓ (Entropy)**	**✓ (RSSI)**

3. Proactive Client Selection in Federated Learning for Connected and Autonomous Vehicles

This section presents the PRINCE (Proactive Reliability-driven INtelligent Client sElection) mechanism for optimizing the selection and aggregation of clients for CAV applications. In its operation, the system employs a hybrid approach composed of network-aware pre-filtering, stochastic reliability modeling, and mobility prediction. The following sections detail system model and PRINCE operations.

3.1. Scenario overview

Figure 1 shows an overview of the vehicular network scenario to address the challenges of high mobility and connection failures in vehicular FL. In this scenario, we consider an FL scenario involving n vehicles $V = v_1, ..., v_n$, where each vehicle acts as a client connected to a central base station b. The server operates in a distributed manner at the edge of the network, with each base station managing its clients within its coverage area. The learning round is divided into three main steps: pre-filtering RSSI, Markov chain calculation, and adaptive aggregation.

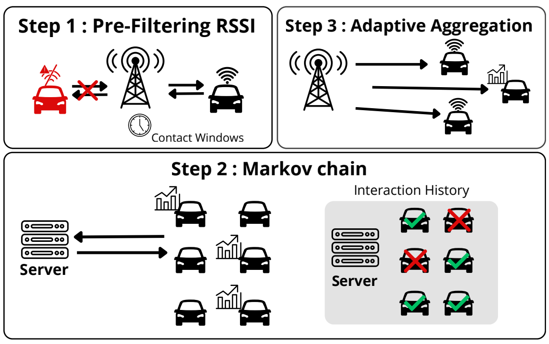

Figure 1. PRINCE Architecture

PRINCE manages the physical movement of vehicles, where speed determines how long vehicles remain in the station's coverage area. The GTSRB dataset [Guo et al. 2019] is used for training, and training is performed over multiple rounds to ensure that the global model achieves optimal stability. PRINCE mobility considers different driving profiles, such as taxis, buses, deliveries, and commuters, which influence the probabilities of success or failure calculated from the Markov chain.

During the pre-filtering step, the central server applies an RSSI-based mechanism to eliminate clients that cannot reliably communicate with the base station. This prevents wasted computation on clients that are likely to experience connection failures. Specifically, for each client v_i, we compute the RSSI using a path loss model, as shown in Eq. 1, where P_{tx} represents the transmission power, $n = 2$ is the path loss exponent for urban environments, $PL_0 = 46.3$ dB is the reference path loss at 1 meter for 3.5 GHz is the reference path loss, d_v is the distance to the base station, and $X_\sigma \sim \mathcal{N}(0, 8^2)$ models log-normal shadow fading.

$$\text{RSSI}_v = P_{tx} - [PL_0 + 10 \cdot n \cdot \log_{10}(d_v)] + X_\sigma \tag{1}$$

We define the contact window τ_v as the estimated time duration during which vehicle v remains within communication range of its serving base station. This metric is fundamental to resource planning, as it determines whether a vehicle can complete a training round before departing the coverage area. Only clients with $\text{RSSI}_v > \text{RSSI}_{th}$ are considered for subsequent selection steps. Although this exclusion imposes a trade-off in terms of participation fairness, it is strictly necessary to prioritize effective contributions, preventing the selection of clients whose physical layer limitations lead to transmission failures. This filtering eliminates vehicles with a high probability of failure, ensuring that the server interacts only with clients that have satisfactory connectivity to complete the training rounds. RSSI and contact window τ_v are intrinsically coupled through vehicular mobility: vehicles with longer predicted contact windows typically experience stronger signal strength due to favorable positioning relative to the base station. The contact window(τ_v) acts as a resource guarantee, validating whether the estimated dwell time is greater than the time required to complete the training round. A bigger contact window (τ_v) generally leads to a more favorable RSSI value estimation for that client.

After the training step, the central server considers the Markov chain Importance-Weighted Aggregation to prioritize updates from valuable clients. In this sense, the global model update $w^{(t+1)}$ is calculated by weighting the local updates $w_v^{(t)}$ based on an importance score I_v, as shown in Eq. 2, where $\mathcal{S}_t$ represents the set of successful clients. The importance score I_v integrates the context of data quality and reliability, ensuring that clients with higher entropy and better link quality contribute more significantly to the global model.

$$\mathbf{w}^{(t+1)} = \sum_{v \in \mathcal{S}_t} \frac{I_v}{\sum_{k \in \mathcal{S}_t} I_k} \cdot \mathbf{w}_v^{(t)} \tag{2}$$

This is essential to ensure the generalization of learning in a highly specific data scenario, such as the vehicular ecosystem. Each vehicle has traffic data limited to its context, and the global model allows a vehicle to benefit from the knowledge acquired in another vehicle without the need to share raw or sensitive data. In the Non-IID context, prioritizing clients with high entropy ensures that rare information is given greater weight, preventing the overall model from becoming biased.

Finally, in the adaptive aggregation stage, the server ends the learning cycle by consolidating only the updates from clients in $\mathcal{S}_t$ (successful completions). Unlike traditional FedAvg which weights updates uniformly by dataset size, our adaptive aggregation uses the importance score I_v (Equation 2) to weight contributions differentially. By prioritizing vehicles with greater connection stability and high entropy, adaptive aggregation ensures that the global model evolves, with reliable and informative updates receiving proportionally greater influence while mitigating the impact generated by abrupt connection losses.

3.2. Proactive Selection Operations

The proposed mechanism relies on determining the specific stability and utility of each client to ensure efficient learning. In this sense, this work considers the client reliability model to capture the temporal dynamics of node behavior, which predicts the training success of each vehicle during the rounds. This model takes into account basic reliability, RSSI, contact windows, and Markov chain.

Given the high mobility of vehicles, the stability of a client is not deterministic. To this end, we model the interaction history as a Discrete-Time Markov Chain (DTMC). We define the state space $\mathcal{Z} = \{S, C, R, A\}$, representing Success, Connection Failure, Resource Failure, and Abort, respectively. The Success state (S) represents vehicles with good connectivity and resource availability to start and finish training rounds, during which the server receives the weights within the time limit. Connection Failure (C) occurs when the vehicle connection lacks sufficient signal strength to establish an initial primary connection with the server, or when the Contact Window (τ_v) expires. Resource Failure (R) includes scenarios where the vehicle's available hardware resources are insufficient to sustain the training task, due to critical battery levels, overheating, or high computational load; therefore, the client rejects connection. Finally, Abort Failure (A) means that the process was voluntarily terminated by the system or by the client before completion.

The transition dynamics are governed by the Matrix $\mathbf{P}$, where each element p_{ij} represents the probability of transitioning from state i to j, as shown in Eq. 3. Therefore, for example, p_{ss} is the probability of transition from a Success state to remaining at the same Success state. p_{sc} is the probability of transition from Success to Connection Failure. So, the Matrix $\mathbf{P}$ is structured in a 4x4 dimension, where all the possible outcomes are calculated.

$$\mathbf{P} = \begin{bmatrix} p_{SS} & p_{SC} & p_{SR} & p_{SA} \\ p_{CS} & p_{CC} & p_{CR} & p_{CA} \\ p_{RS} & p_{RC} & p_{RR} & p_{RA} \\ p_{AS} & p_{AC} & p_{AR} & p_{AA} \end{bmatrix} \tag{3}$$

The base Failure Risk (ρ_v) is calculated by adding the transition probabilities to any failure state (C, R, and A), as shown in Eq. 4. This formulation allows the system to empirically estimate the likelihood that a client completes a round based on historical windows.

$$\rho_v = \sum_{f \in C,R,A} p_{s_v,f} \tag{4}$$

Afterwards, we consider mobility prediction to anticipate client departures. Specifically, we consider a Kalman Filter to estimate the position $\mathbf{p}_v$ and speed $\mathbf{v}_v$ of the vehicle. We employ a constant-velocity motion model with state vector $\mathbf{x}_v = [\mathbf{p}_v^T, \mathbf{v}_v^T]^T$, where position and velocity are updated recursively using standard Kalman equations with process noise accounting for acceleration variability and measurement noise reflecting GPS uncertainty. It is important to mention that other mobility predictors could be used, such as ARIMA or LSTM. We choose the Kalman [Zhang et al. 2023] for its real-time prediction performance, which is essential for identifying which vehicles are most likely to leave the coverage area before completing their training rounds. The predicted Contact Window τ_v is calculated geometrically by solving Eq. 5, which determines when the vehicle's trajectory intersects the coverage boundary.

$$|\mathbf{p}_v + \mathbf{v}_v \cdot t - \mathbf{b}| = r_b \tag{5}$$

It is important to combine reliability and mobility into a coherent selection metric. In addition to RSSI, the inclusion of the contact window (τ_v) serves as a resource-conservation mechanism, defining the time interval during which the vehicle remains within range of a base station's signal. In this sense, the basic risk of failure is modulated by the normalized contact time $\tilde{\tau}_v$ to produce a mobility-sensitive risk estimate $\rho_v^{mod} = \rho_v \cdot (1 - \tilde{\tau}_v)$, where $\tilde{\tau}_v = \min(\tau_v/\tau_{max}, 1)$ normalizes the predicted contact window to the range [0,1]. This ensures that customers with longer contact windows are less penalized by historical instability, as their $(1 - \tilde{\tau}_v)$ factor approaches zero, reducing the overall risk penalty.

To mitigate the bias introduced by non-IID data, we consider Shannon Entropy to quantify the class diversity of the local dataset, where $p_c^{(v)}$ represents the proportion of samples belonging to class c in vehicle v's local dataset, and K is the total number of classes. A vehicle with a higher entropy H_v contributes more to the generalization of the global model, as represented by Eq. 6.

$$H_v = -\sum_{c=1}^{K} p_c^{(v)} \log_2 p_c^{(v)} \tag{6}$$

Finally, the selection decision is based on a Multi-factor Utility Function (U_v) that balances diversity, risk, mobility, and historical accuracy (A_v), as shown in Eq. 7. All components are normalized to [0,1]: $\tilde{H}_v = H_v/\log_2(K)$, $\tilde{\tau}_v = \min(\tau_v/\tau_{round}, 1)$, and $\tilde{A}_v = 1 - A_v$.

$$U_v = \alpha \cdot \tilde{H}_v + \gamma \cdot \tilde{\tau}_v - \beta \cdot \rho_v^{mod} - \delta \cdot \tilde{A}_v \tag{7}$$

In this equation, the hyperparameters α, β, γ, and δ dynamically tune the selection strategy based on the vehicular context. Specifically, $\alpha \cdot \tilde{H}_v$ rewards high Shannon entropy to maximize informational gain from diverse distributions, while $\gamma \cdot \tilde{\tau}_v$ incorporates mobility constraints by favoring extended contact windows to prevent transmission aborts. Conversely, $\beta \cdot \rho_v^{mod}$ penalizes historical instability based on the Markov risk estimate. Finally, the term $-\delta \cdot \tilde{A}_v$ implements a distributed hard-example mining mechanism adapted from [de Souza et al. 2024], by targeting clients with lower local accuracy, which typically possess underrepresented data distributions in vehicular Non-IID scenarios, effectively focusing training on edge cases where the global model currently underperforms.

Hyperparameters values α, β, γ, and δ were empirically tuned via ablation studies over $[0.1, 0.5]$ to maximize convergence speed and validation accuracy. The resulting configuration prioritizes proactive failure avoidance, balancing informational quality

(Entropy/Hard-mining) with physical reliability (Markov/Mobility constraints) to ensure robust aggregation.

Algorithm 1 Proactive Client Selection with Multi-factor Utility

---

Require: Candidate set $\mathcal{V}$, number of clients n, hyperparameters $\alpha, \beta, \gamma, \delta$, thresholds $\mathrm{RSSI}_{\min}, \tau_R$, epochs $E_{\max}, E_{\min}$
Ensure: Set of selected clients $\mathcal{V}^*$ with assigned local epochs

// τ_R: Resource failure risk threshold (e.g., 0.3) for epoch adaptation
// Step 1: RSSI Pre-filtering and Utility Computation
$\mathcal{V}_{eligible} \leftarrow \{v \in \mathcal{V} : \mathrm{RSSI}_v \geq \mathrm{RSSI}_{\min}\}$
// Step 2: Utility Computation
for each vehicle $v \in \mathcal{V}_{eligible}$ **do**
    Estimate transition matrix $\mathbf{P}_v$
    Calculate base failure risk ρ_v
    Predict contact time τ_v
    Calculate modulated risk ρ_v^{mod}
    Calculate local entropy H_v
    Retrieve historical accuracy $\tilde{A}_v$
    Calculate utility score U_v
end for
// Step 3: Selection
$\mathcal{V}^* \leftarrow$ top-n vehicles sorted by U_v in descending order
// Step 4: Adaptive Epoch Assignment
for each vehicle $v \in \mathcal{V}^*$ **do**
    if $p_{s_v R} > \tau_R$ **then** ▷ High resource failure risk detected
        $E_v \leftarrow E_{\min}$ ▷ Reduce epochs for at-risk clients
    else
        $E_v \leftarrow E_{\max}$
    end if
end for
return $\mathcal{V}^*$ with local epochs $\{E_v\}_{v \in \mathcal{V}^*}$

The algorithm 1 summarizes PRINCE operations to optimize the set of participating vehicles while ensuring robustness against mobility-induced failures. The process begins with a pre-filtering step, where the candidate set $\mathcal{V}$ is screened based on the RSSI, immediately discarding vehicles below $RSSI_{min}$ to prevent resource waste on unstable links. Subsequently, for each eligible client, the algorithm computes a multi-factor utility score U_v by integrating the base failure risk ρ_v (estimated via Markov Chain modeling) and the predicted contact window τ_v with local data quality metrics, specifically Shannon Entropy H_v and historical accuracy $\tilde{A}_v$. The transition matrix $\mathbf{P}_v$ is estimated from interaction history via maximum likelihood, with uniform prior for cold-start vehicles and exponential smoothing ($\lambda = 0.9$) to balance recent observations with historical context. Finally, to further mitigate dropout risks, the algorithm applies an Adaptive Epoch Assignment strategy: selected clients identified as having tight contact windows or higher risk profiles are assigned a reduced workload E_{min} to ensure successful transmission before moving out of range, while stable clients utilize the maximum capacity E_{max}.

4. Evaluation

This section presents the evaluation of the PRINCE mechanism in a realistic federated vehicle learning scenario. We compare its performance against state-of-the-art client selection approaches, focusing on model accuracy, loss, convergence behavior, success rates, failure rates, and distribution of training outcomes. We first describe the experimental setup, including the dataset, simulation environment, and baseline methods. We then analyze the obtained results and discuss the observed tradeoffs.

4.1. Experimental Setup

In this paper, the simulation environment is built upon a custom Python-based framework[1] utilizing the PyTorch library for local model training and aggregation orchestration. We chose a controlled simulation to isolate the performance of the client selection mechanism from lower-level network issues, such as packet collisions or handover delays, that fall outside the scope of this study. To strictly enforce reproducibility, all random seeds for data partitioning and weight initialization were fixed. All experiments were repeated 33 times with different random seeds, and the reported metrics represent mean values across these runs. We considered the German Traffic Sign Recognition Benchmark (GTSRB) dataset [Guo et al. 2019], which contains 43 classes of traffic signs and more than 50,000 labeled images. This dataset represents a realistic vehicle perception task in which autonomous vehicles must reliably recognize traffic signs to support safe navigation. The learning model is a convolutional neural network with batch normalization, composed of four convolutional blocks followed by fully connected layers.

We simulate a vehicular FL environment as modeled in Section 3.1. Table 2 summarizes the main simulation parameters. The simulation environment consists of a $1500\,\mathrm{m} \times 1500\,\mathrm{m}$ urban area covered by four base stations. In this environment, vehicles move according to realistic mobility patterns and are categorized into distinct profiles (taxi, bus, delivery, and commuter), inspired by a similar setup used by [Chen et al. 2025]. In this sense, each vehicle is characterized by different speeds and movement behaviors. Specifically, the network comprises 58 vehicles, from which the server selects 11 clients ($\approx 20\%$) per communication round. The training process spans 200 rounds, ensuring the global model achieves optimal stability. The temporal dynamics are governed by a strict 30-second round duration, which defines the window for local training and weight transmission. Since vehicles move continuously, their positional displacement between consecutive rounds is determined by the constant-velocity motion model applied over this 30-second interval, with exact distances varying based on the assigned mobility profile.

We modeled three types of failures in our simulation environment: Connection, Resource, and Abort failures. Unlike static probability models, these failures arise dynamically from physical simulation conditions. Specifically, connection failures occur when the RSSI falls below the threshold ($RSSI_{th}$) or when handover instability occurs. $RSSI_{th} = -85$ dbM value represents the transition between low-quality and unusable signal strength [Mangipudi et al. 2025]. Resource failures result from insufficient battery, overheating, or high computing capacity load, while Aborts occur when the predicted contact time is insufficient to complete a training round. Consequently, the observed fail rate is directly dependent upon the estimated failure risk (ρ_v), demonstrating strong alignment between predicted and actual failure outcomes, serving as an empirical validation of the transition probabilities defined in the Markov Chain model presented in Equation 4.

[1] https://github.com/Twinte/PRINCE---Proactive

Table 2. Simulation Parameters

Parameter	Value
Number of clients	58 vehicles
Clients per round	11 ($\approx$20%)
Number of rounds	200
Local epochs	5
Batch size	32
Learning rate	0.01
Non-IID distribution	Dirichlet $\alpha = 0.3$
Grid size	$1500\,\text{m} \times 1500\,\text{m}$
Base stations	4
BS coverage radius	$550\,\text{m}$
RSSI threshold	-85 dBm
Round duration	30 seconds

To model realistic non-IID data across vehicles, we adopt a Dirichlet distribution with a concentration parameter $\alpha = 0.3$. This configuration produces heterogeneous data partitions in which each vehicle observes a skewed subset of traffic sign classes, reflecting the fact that vehicles traveling on different routes encounter different traffic conditions and signage. We compare the PRINCE mechanism with four baseline client selection strategies: i) **Random Selection** randomly selects k clients per round without considering data characteristics or reliability; ii) **Entropy-only** selects clients with the highest label entropy, prioritizing data diversity while ignoring reliability. iii) **ACSP-FL** [de Souza et al. 2024] selects clients with below-average local accuracy, highlighting clients that require further training; and **ECS-HDSR** [Sousa et al. 2025] considers an entropy-based approach with a reserve pool to substitute failing clients using Hausdorff distance.

All algorithms are compared using the following metrics: accuracy, loss, success rate, convergence speed, distribution of training outcomes by success, connection failures, resource failures, aborts, and fail rate. Accuracy and loss are used as primary performance indicators, where the former measures the accuracy of classifications in the dataset (GTSRB) and the latter quantifies model error. Convergence speed reflects the process speed, while the success and failure rates measure communication efficiency.

4.2. Results

Figure 2 shows the accuracy and loss over the rounds for the evaluated mechanisms. In this sense, Figure 2(a) presents the evolution of accuracy over communication rounds for the evaluated mechanisms. By analyzing the results, we observe that the PRINCE mechanism (*i.e.*, blue line) achieves faster and more stable convergence compared to the baseline mechanisms. While methods like ACSP-FL and Entropy-only exhibit significant fluctuations due to the inclusion of unreliable clients, the PRINCE maintains a steady upward trajectory, reaching a final accuracy plateau of approximately. 84%. This stability is a direct result of the Markov-based prediction, which prevents the selection of nodes likely to disconnect during the weight transmission step.

Figure 2(b) highlights the convergence of training loss for the analyzed mechanism. The PRINCE mechanism consistently maintains the lowest loss throughout the optimization process, achieving a final value of 0.628. In contrast, reactive baselines such as Random Selection and ACSP-FL exhibit spikes in the loss function, indicative of model divergence due to failed aggregations or non-IID data bias. Our approach effectively mitigates this by balancing the Shannon Entropy (data quality) with the RSSI pre-filtering (link quality), ensuring that only high-quality updates are aggregated.

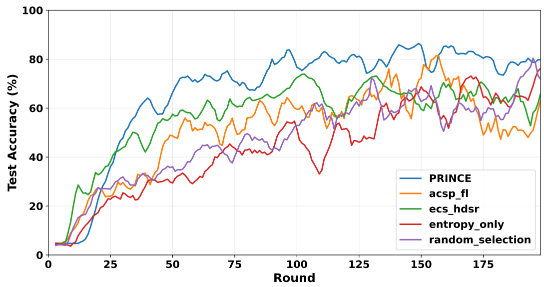

(a) Accuracy over rounds

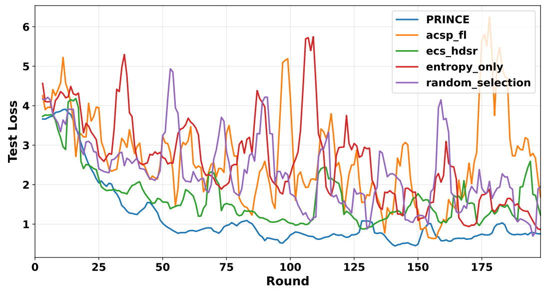

(b) Loss over the rounds

Figure 2. Accuracy and Loss over the rounds for the evaluated mechanisms

Table 3 summarizes the overall performance of all the mechanisms evaluated. The PRINCE mechanism achieves the highest final accuracy (83.90%) and success rate (61.32%), demonstrating that its reliability-aware selection and learning-impact optimization can be effectively combined.

Table 3. Overall Performance Comparison

Method	Accuracy	Loss	Success Rate
PRINCE	**83.90%**	**0.628**	**61.32%**
Entropy-only	74.68%	0.910	27.05%
Random	72.84%	0.948	25.18%
ECS-HDSR	71.48%	0.987	35.73%
ACSP-FL	68.45%	1.246	17.32%

Figure 3 shows the convergence and success rate for the evaluated mechanisms. For instance, Figure 3(a) reports the number of communication rounds required to reach accuracy thresholds of 60% and 70%. PRINCE converges to a 60% accuracy in only 35 rounds, which is $1.4\times$ faster than ECS-HDSR, $1.6\times$ faster than ACSP-FL and $3\times$ faster than Random selection. The rapid convergence of the PRINCE is particularly relevant for time-sensitive vehicle applications, where models must be deployed quickly.

Figure 3(b) shows the success rate per round. Although baselines exhibit high volatility, often dropping below 20% success in congested or low-coverage scenarios, the PRINCE mechanism maintains robust performance, consistently staying above the 60% mark. This improvement is attributed to the Markov chain reliability model, which learns temporal patterns from historical client behavior. Clients who have experienced recent failures exhibit higher transition probabilities toward failure states, and modulated risk ρ_v^{mod} captures this tendency, allowing the system to preemptively avoid them.

Figure 4 illustrates the distribution of the training results, categorizing them into Success, Connection Failure, Resource Failure, and Abort. Visual analysis reveals that connection failures (red bars) are the dominant failure mode for all baseline mechanisms, accounting for more than 79% of interruptions. However, PRINCE drastically reduces this component. By applying the RSSI pre-filtering mechanism, we reduced connection failures by approximately 2.2 to $2.5\times$ compared to baseline approaches (698 vs. 1,511 to 1,738), validating the hypothesis that signal-aware selection is crucial for vehicular FL.

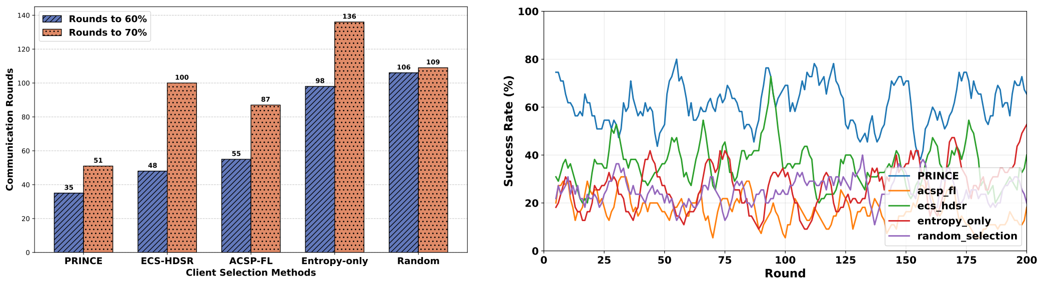

(a) Convergence Speed (Rounds to Reach Threshold)

(b) Success Rate Comparison over rounds

Figure 3. Convergence and Success Rate for the evaluated mechanisms

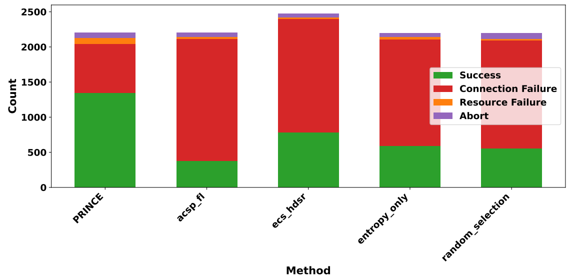

Figure 4. Distribution of training outcomes

Table 4 provides the numerical breakdown of these outcomes. Although the PRINCE mechanism exhibits slightly higher resource failures (81) than some baselines, this is expected behavior: since the RSSI pre-filtering successfully eliminates connection-prone clients, the remaining failures become proportionally more visible. Nevertheless, the total failure rate of 38.68% is substantially lower than the 82.68% observed in ACSP-FL, which confirms the resource efficiency of PRINCE. From a practical perspective, these results demonstrate that combining reliability-aware filtering with need-based client prioritization is essential for vehicular FL systems operating in high mobility. The PRINCE effectively minimizes the waste of computational resources on clients destined to fail, maximizing the utility of every communication round.

Table 4. Training Outcome Distribution

Method	Success	Conn. Fail	Res. Fail	Abort	Fail Rate
PRINCE	1349	698	81	72	**38.68%**
ECS-HDSR	786	1609	27	52	64.27%
Entropy-only	595	1511	38	56	72.95%
Random	554	1536	26	84	74.82%
ACSP-FL	381	1738	22	59	82.68%

5. Conclusion

This paper presents a proactive client selection mechanism for vehicular FL that addresses the critical challenge of client failures in dynamic mobile environments. The proposed approach integrates three key components: RSSI-based pre-filtering to eliminate connection-prone clients, Markov chain reliability modeling to predict failure probabilities from historical behavior, and a multi-factor utility function that combines data diversity, reliability, connectivity, and learning potential. By targeting reliable clients with lower local accuracy, the method implements a distributed hard-example mining strategy that forces the optimization process to focus on edge cases where the global model currently underperforms, accelerating convergence without compromising training stability.

Evaluation results on the GTSRB dataset demonstrated that the PRINCE mechanism outperforms the latest baselines across all metrics. The PRINCE achieved 83.90% accuracy with a 61.32% success rate, representing improvements of 9.22 percentage points in accuracy and 3.5× in resource efficiency compared to baseline methods. The PRINCE also converged to 60% accuracy in only 35 rounds, which is up to 3× faster than competing approaches. Furthermore, RSSI pre-filtering reduced connection failures by 2.2 to 2.5×, confirming that network instability is the dominant failure mode in vehicular federated learning and can be effectively mitigated through signal-aware client selection.

Acknowledgment

This research was partially sponsored by CNPq grant 404186/2021-1, CAPES, the Brazilian Ministry of Science, Technology, and Innovations, with resources from Law nº 8,248, of October 23, 1991, within the scope of PPI-SOFTEX, coordinated by Softex, and published under Arquitetura Cognitiva (Phase 3), DOU 01245.003479/2024-10. This research is part of the INCT of Intelligent Communications Networks and the Internet of Things (ICoNIoT) funded by CNPq (proc. 405940/2022-0) and Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) Finance Code 88887.954253/2024-00This work was funded by the National Institute of Science and Technology in Artificial Intelligence Applied to Smart and Sustainable Cities in the Brazilian Amazon (INCT IAmazônia - inctiamazonia.org.br) (Grant nº 409001/2024-4).

References

Chellapandi, V. P., Yuan, L., Brinton, C. G., Żak, S. H., and Wang, Z. (2023). Federated learning for connected and automated vehicles: A survey of existing approaches and challenges. *IEEE Transactions on Intelligent Vehicles*, 9(1):119–137.

Chen, D., Deng, T., Jia, J., Feng, S., and Yuan, D. (2025). Mobility-aware decentralized federated learning with joint optimization of local iteration and leader selection for vehicular networks. *Computer Networks*, 263:111232.

Chen, H. and Vikalo, H. (2024). Heterogeneity-guided client sampling: Towards fast and efficient non-iid federated learning. *Advances in Neural Information Processing Systems*, 37:65525–65561.

de Souza, A. M., Maciel, F., da Costa, J. B., Bittencourt, L. F., Cerqueira, E., Loureiro, A. A., and Villas, L. A. (2024). Adaptive client selection with personalization for communication efficient Federated Learning. *Ad Hoc Networks*, 157:103462.

Deva Hema, D. and Rajeeth Jaison, T. (2024). Efficient collision risk prediction model for autonomous vehicle using novel optimized lstm based deep learning framework. *International Journal of Intelligent Transportation Systems Research*, 22(2):352–362.

Elbir, A. M., Soner, B., Çöleri, S., Gündüz, D., and Bennis, M. (2022). Federated learning in vehicular networks. In *IEEE International Mediterranean Conference on Communications and Networking (MeditCom)*, pages 72–77.

Guo, X., Zhao, C., and Wang, Y. (2019). Traffic sign recognition based on joint convolutional neural network model. the 2nd International Conference on Big Data Technologies (ICBDT '19), page 200–203, New York, NY, USA. ACM.

Gutierrez, D. M. J., Solans, D., Heikkilä, M. A., Vitaletti, A., Kourtellis, N., Anagnostopoulos, A., and Chatzigiannakis, I. (2024). Non-iid data in federated learning: A survey with taxonomy, metrics, methods, frameworks and future directions. *ArXiv*, abs/2411.12377.

Huang, T., Lin, W., Shen, L., Li, K., and Zomaya, A. Y. (2022). Stochastic client selection for federated learning with volatile clients. *IEEE Internet of Things Journal*, 9(20):20055–20070.

Lobato, W., Costa, J. B. D. D., Souza, A. M. d., Rosário, D., Sommer, C., and Villas, L. A. (2022). Flexe: Investigating federated learning in connected autonomous vehicle simulations. In *2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall)*, pages 1–5.

Lobato, W., Da Costa, J. B., Gonzalez, L. F., Cerqueira, E., Rosário, D., Sommer, C., and Villas, L. A. (2024). Entropy and mobility-based model assignment for multi-model vehicular federated learning. In *2nd International Conference on Federated Learning Technologies and Applications (FLTA)*, pages 8–15. IEEE.

Mangipudi, S. et al. (2025). Cohere - congestion-aware offloading and handover via empirical rat evaluation for multi-rat networks. *arXiv preprint arXiv:2511.00439*.

Maroua, D. (2024). A state-of-the-art on federated learning for vehicular communications. *Vehicular Communications*, 45:100709.

Pan, H., Luo, M., Wang, J., Huang, T., and Sun, W. (2024). A Safe Motion Planning and Reliable Control Framework for Autonomous Vehicles. *IEEE Transactions on Intelligent Vehicles*, 9(4):4780–4793.

Smestad, C. and Li, J. (2023). A systematic literature review on client selection in federated learning. In *Proceedings of the 27th International Conference on Evaluation and Assessment in Software Engineering*, page 2–11. ACM.

Sousa, J., Ribeiro, E., Bustincio, R., Bastos, L., Morais, R., Cerqueira, E., and Rosário, D. (2025). Enhancing robustness in federated learning using minimal repair and dynamic adaptation in a scenario with client failures. *Annals of Telecommunications*, 80(9-10):885–899.

Vardhan, H., Yu, X., Rosing, T., and Mazumdar, A. (2025). Client Selection in Federated Learning with Data Heterogeneity and Network Latencies.

Zhang, B., Yu, W., Jia, Y., Huang, J., Yang, D., and Zhong, Z. (2023). Predicting vehicle trajectory via combination of model-based and data-driven methods using kalman filter. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 238(8):2437–2450.

Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., and Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems*, 216:106775.

Zhu, H., Xu, J., Liu, S., and Jin, Y. (2021). Federated learning on non-IID data: A survey. *Neurocomputing*, 465:371–390.