Proceedings of the 4th International Conference on Mechatronics and Smart Systems
DOI: 10.54254/2755-2721/2026.34419

Fairness-Aware Client Scheduling Strategy for Federated Learning in Energy-Constrained Vehicular Networks

Yilu Zhang

College of Photonic and Electronic Engineering, Fujian Normal University, Fuzhou, China
zhluluann@gmail.com

Abstract. The Internet of Vehicles (IoV) generates large volumes of distributed, privacy-sensitive data across mobile clients operating under dynamic energy and communication constraints. Federated Learning (FL) offers a privacy-preserving framework for collaborative model training; however, standard aggregation schemes such as Federated Averaging (FedAvg) tend to neglect client-level performance heterogeneity, leading to severe accuracy disparities under non-independent and identically distributed (non-IID) data distributions. This paper proposes a Fairness-Aware Client Scheduling (FACS) strategy that prioritises under-performing vehicles during both the selection and aggregation phases, guided by a composite weight derived from per-client accuracy, residual battery level, and communication latency. Experiments on the MNIST handwritten-digit benchmark with 20 simulated vehicles across 100 communication rounds demonstrate that the proposed method breaks the local-optimum trap of standard FedAvg under non-IID conditions, raising global test accuracy from 89.76% to 96.96%, while reducing the communication rounds required to reach 90% accuracy from failure (never achieved) to 26 rounds. Under the IID setting, the proposed strategy introduces no measurable performance degradation relative to the unmodified baseline.

Keywords: Federated learning, Internet of Vehicles, fairness-aware scheduling, client selection

1. Introduction

With the rapid development of interconnected vehicles and intelligent transportation systems, mobile nodes operating in dynamic environments have generated unprecedented amounts of distributed and heterogeneous data. Applications such as cooperative perception, traffic flow prediction, hazard identification and autonomous driving increasingly rely on machine learning (ML) models trained on large-scale vehicle data sets [1, 2]. However, the sensitivity of data captured by on-board sensors, such as position trajectory, driving behavior and biometric signals, has caused serious privacy and regulatory concerns under the regulatory framework of the general data protection regulation (GDPR). Centralized storage of such data on remote servers will not only generate high communication costs, but also expose users to significant security risks [3].

As a promising example, federated learning (FL) came into being. It supports distributed model training without transmitting the original data to the central server [2]. Under the fl framework, each

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.34419

participating vehicle trains the local model on its private data, and only uploads the model parameters or gradients to the coordination server for aggregation. The groundbreaking FedAvg algorithm has proved that a competitive global model can be built through decentralized data, and the number of communication rounds can be significantly reduced [3]. Subsequent research will further expand FL to the field of vehicle and Internet of vehicles (IOV), in which vehicles communicate through the roadside unit (RSU) or vehicle-to-everything (V2X) links [4]. Chellapandi et al. investigated the application of FL in Connected and Automated Vehicles (CAVs), and pointed out that data heterogeneity, communication instability and resource constraints are the three major challenges currently facing [5]. Du et al. also stressed that the inherent high mobility and intermittent connection of vehicle networks make the traditional FL assumption no longer applicable, and the scheduling and aggregation mechanism needs to be specially designed [4].

The federal learning of the Internet of vehicles is facing core pain points: first, the non-IID heterogeneity of vehicle data is easy to cause client drift, weakening FedAvg convergence speed and model accuracy [6]. FedProx ensures the convergence of heterogeneous scenarios through local punishment. Existing studies have confirmed that the non-IID distribution is the main reason for the low accuracy and slow convergence of federal learning, and it is necessary to optimize the aggregation strategy and customer selection mechanism [7-10]; Second, the traditional scheme ignores the performance fairness of a single client and only optimizes the global accuracy. Zhang, LeCun, Shi and others proved from the perspective of game theory, fair scheduling, reputation mechanism and other aspects that the client selection and aggregation weight should take into account individual performance to achieve fair modeling [11-14]; Third, the energy of the Internet of vehicles is limited, and continuous training is easy to cause vehicle power failure and exit, and aggravate data deviation [15]. The existing literature carries out scheduling research from the perspective of time delay optimization, energy consumption perception, communication cost and data diversity [15-18], but there is no method to take into account selection fairness, energy constraint and communication delay under the unified scheduling framework.

The existing review reveals three unresolved research problems: first, standard aggregation algorithms like FedAvg struggle to avoid local optimum convergence under vehicle non-IID data, and random client selection often deprives low-performance vehicles of training opportunities; second, current fair selection schemes fail to simultaneously consider residual energy and communication delay, leading to data distribution deviations when energy-limited clients exit; third, a unified research framework is lacking for the coupling between client selection fairness and aggregation weights, with most studies focusing on a single dimension. These gaps have inspired this paper's core research questions:

In the energy-limited vehicle network with non-IID data distributions, a fairness-aware client scheduling strategy considering each client's accuracy, remaining battery energy, and communication delay is proposed to improve model convergence, final accuracy, and performance fairness of each client in federated learning without reducing IID data performance.

To address this question, this paper studies the following two aspects: (I) preventing the global model from converging to partial local optima under non-IID conditions by giving priority to low-accuracy vehicles in the client selection process; and (II) further amplifying fairness gains without introducing instability by incorporating each client's accuracy, data volume, and battery power into the aggregation weight.

This paper proposes a fair aware federated learning client scheduling (FACS) strategy in the Internet of vehicles environment. The main contributions are as follows. Firstly, a three factor vehicle selection score is designed to assign a higher scheduling priority to the current vehicles with

DOI: 10.54254/2755-2721/2026.34419

low test accuracy, less remaining battery energy and short communication delay, so as to ensure that the vehicles with insufficient representation get more frequent and timely training opportunities. Secondly, a composite physical perception aggregation weight is introduced to measure the contribution of each vehicle to the global model update according to the amount of local data, inversion accuracy and residual energy, so as to amplify the influence of historical ignored clients in the aggregation process. Thirdly, the proposed framework is compatible with FedAvg and FedProx local optimization methods, and can be evaluated across multiple baselines. The simulation results show that under the non-IID condition of FedAvg, FACS improves the global test accuracy from 89.76% to 96.96%, reduces the number of rounds required to reach 90% accuracy from failure to 26, and there is no measurable performance degradation under IID condition.

2. Method

2.1. Dataset preparation

MNIST data set was adopted as the experimental benchmark, including 70000 gray-scale images of handwritten digits [12], divided into 60000 training samples and 10000 test samples. Each image is stored as a 28 × 28 pixel single channel (grayscale) array, and the pixel value is in the range of [0, 255], representing one of the ten digital categories between 0 and 9. The classification task is a balanced 10 class recognition problem, in which each image carries a single ground truth label corresponding to the depicted number.

Before training, all pixel intensities are normalized to [0, 1] by dividing by 255.0, which stabilizes gradient based optimization and reduces sensitivity to pixel scale changes. Then, heterogeneous label subsets are allocated to each client through label skew partition, and 60000 training images are distributed on 20 simulated vehicles in a non-IID manner to simulate the real scene of different vehicles encountering different traffic environments. 10000 test images are reserved as a globally reserved evaluation set; After each round of communication, the aggregated global model is evaluated on the local test partition of each vehicle, and the accuracy of each client is recorded as the main performance index.

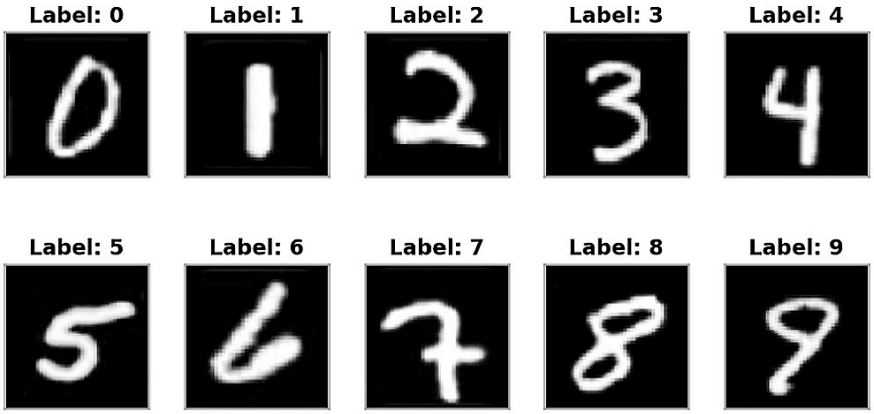


Figure 1. Representative samples from the MNIST dataset. each panel shows one image per digit class (0–9). all images are 28 × 28 pixels in grayscale [12]

All experiments are built and implemented based on python. The local model uses a two-layer convolutional neural network structure, including two convolution modules and a full connection

DOI: 10.54254/2755-2721/2026.34419

layer, and finally outputs a very similar softmax probability distribution.

2.2. Fairness-aware client scheduling strategy for federated learning

Federated learning is a distributed machine learning mode [18]. In this mode, the central server will coordinate multiple clients to train a shared global model, and there is no need to exchange original data in the process [3]. In each round of communication t, the server will send the current global model parameter $w^{(t)}$ to some selected vehicles. Each vehicle will execute the local random gradient descent (SGD) algorithm on its own private data, and then send the updated model back to the server for summary. This principle of focusing on data locality is particularly suitable for the application of Internet of vehicles (IOV) where vehicles will generate sensor data involving privacy.

The global optimisation objective minimised by standard FedAvg is:

$$F(w) = \sum_k \frac{n_k}{N} \cdot F_k(w) \tag{1}$$

where $F_k\left(w\right) = \frac{1}{n_k} \sum_{i=1}^{n_k} \mathscr{L}\left(w, x_i^k, y_i^k\right)$is the local empirical loss for vehicle k, n_k is its local data volume, $N = \sum n_k$, and $\mathscr{L}$ is the cross-entropy loss over the 10 digit classes:

$$\mathscr{L}\left(w, x_i, y_i\right) = -\sum_{c=0}^{9} 1[y_i = c] \cdot \log \hat{p}_c(w) \tag{2}$$

where $\hat{p}_c(w)$ represents the softmax probability assigned to class c, and $1[\cdot]$ is the indicating function. FedAvg minimizes equation (1) by local updates weighted by the average amount of data. This method has significant effect under the condition of independent identically distributed (IID), but under the condition of non independent identically distributed (non-IID), it will systematically bias the global model to most types of clients [6].

Fairness-aware client selection. In order to correct the systematic deviation from under represented customers, the FACS strategy is proposed. Each candidate vehicle will be assigned a selection priority score, which is inversely proportional to the current accuracy of the vehicle for each customer:

$$P_k^{(t)} = \frac{\frac{1}{Acc_k^{(t-1)}}}{\sum_{j \in S^t} \frac{1}{Acc_j^{(t-1)}}} \tag{3}$$

$Acc_k^{(t-1)}$ represents the test accuracy of vehicle K recorded at the end of the $t-1$ round, while S^t refers to the collection of vehicles that meet the RSU coverage requirements (located within a 500 meter coverage radius), battery power requirements (residual energy not less than 20% of capacity) and delay requirements (communication delay not more than 1 second). Vehicles with lower accuracy are given higher selection probability, which can ensure that the global model can gradually correct the data distribution that is not fully represented.

Composite aggregation weight. After completing the local training, each selected vehicle will upload its updated parameters. The server will use a comprehensive weight to aggregate these updates, which takes into account the amount of data and the lack of accuracy of each client:

DOI: 10.54254/2755-2721/2026.34419

$$\alpha_k^{(t)} = \frac{\frac{n_k}{N^t} \cdot P_k^{(t)}}{\sum_{j \in S^t} \frac{n_j}{N^t} \cdot P_j^{(t)}} \tag{4}$$

$$w^{(t+1)} = \sum_{k \in S^t} \alpha_k^{(t)} \cdot w_k^{(t)} \tag{5}$$

The normalized comprehensive weight $\alpha_k^{(t)}$ can ensure that the value of $\sum \alpha_k^{(t)}$ is 1, and also enhance the contribution of low-precision vehicles in data aggregation. When all clients reach the same accuracy level, equation (4) will be simplified to the standard FedAvg weight $\frac{n_k}{N^t}$, so as to ensure backward compatibility under IID conditions.

FedProx local optimisation. In the FedProx variant algorithm, each smart car will optimize an enhanced local objective function to limit the difference between the local model and the global model.

$$h_k\left(w_k; w_{global}\right) = F_k\left(w_k\right) + \frac{\mu}{2} \cdot \left\| w_k - w_{global} \right\|^2 \tag{6}$$

The near end coefficient is $\mu = 0.01$. The secondary penalty term can effectively alleviate the local model offset problem caused by large differences in the distribution of client data, and can provide the corresponding theoretical convergence guarantee in the non-IID conditions [6].

System process. In each iteration T, the complete FACS process performs the following operations in turn. First, the vehicle completes its position update on the simulated road with a length of 1000 meters, and then selects the vehicles within the coverage range of 500 meters RSU and the remaining battery power is higher than 20%, forming a qualified candidate pool for this round. The system selects the current round of training participants from the candidate pool according to equation (3), the selected vehicle receives the latest global model, and completes a round of local model iterative training based on SGD or FedProx algorithm. After the training, the model parameters will be uploaded to the server. The server will aggregate and update the model parameters through equation (4) and equation (5), and then broadcast the new global model to each vehicle node. Finally, the system evaluates the accuracy of each client model, and updates the node priority score of the next iteration synchronously.

3. Results and discussion

3.1. Results

In the process of 100 rounds of communication, this paper evaluated six experimental configurations, including three baseline groups (IID-FedAvg, non-IID-FedAvg, non-IID-FedProx) using standard random scheduling, and three suggestion groups (IID-FedAvg-FACS, non-IID-FedAvg-FACS, non-IID-FedProx-FACS) using FACS. The changes in the training process are visually displayed in Figure 1, Figure 2, Figure 3 and Figure 4.

Table 1. Summary of key metrics at communication round 100

Configuration	Setting	Global Acc.	Rounds to 90%
FedAvg (Baseline)	IID	0.9776	4
FedAvg (Baseline)	Non-IID	0.8976	Failed
FedProx (Baseline)	Non-IID	0.9537	20

DOI: 10.54254/2755-2721/2026.34419

Table 1. (continued)

FedAvg-FACS (Proposed)	IID	0.9776	4
FedAvg-FACS (Proposed)	Non-IID	0.9696	26
FedProx-FACS (Proposed)	Non-IID	0.9682	22

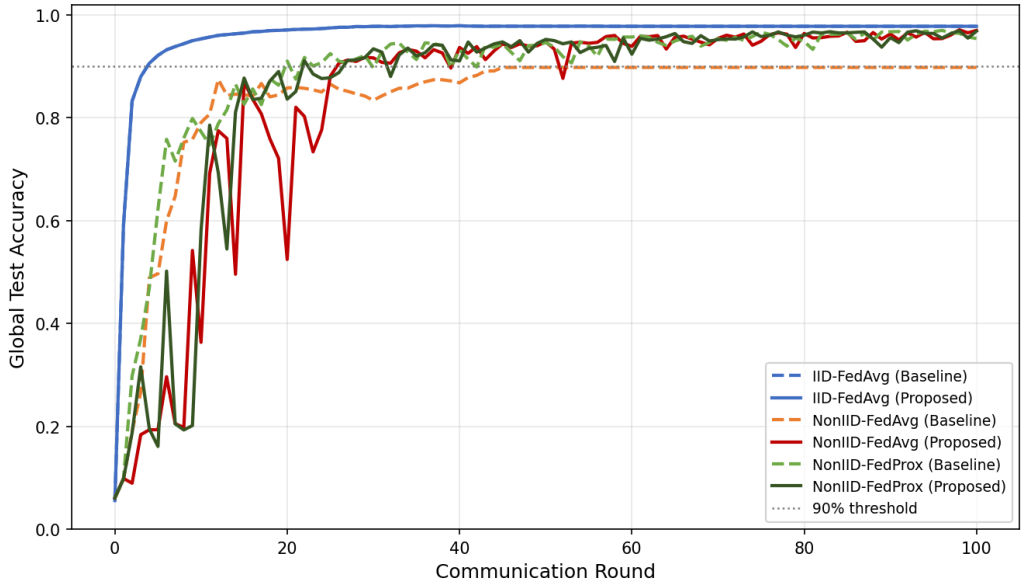


Figure 2. Global test accuracy over 100 communication rounds for all six experimental configurations. dashed curves indicate baseline (random scheduling); solid curves indicate proposed (FACS). the horizontal dotted line marks the 90% accuracy threshold (picture credit: original)

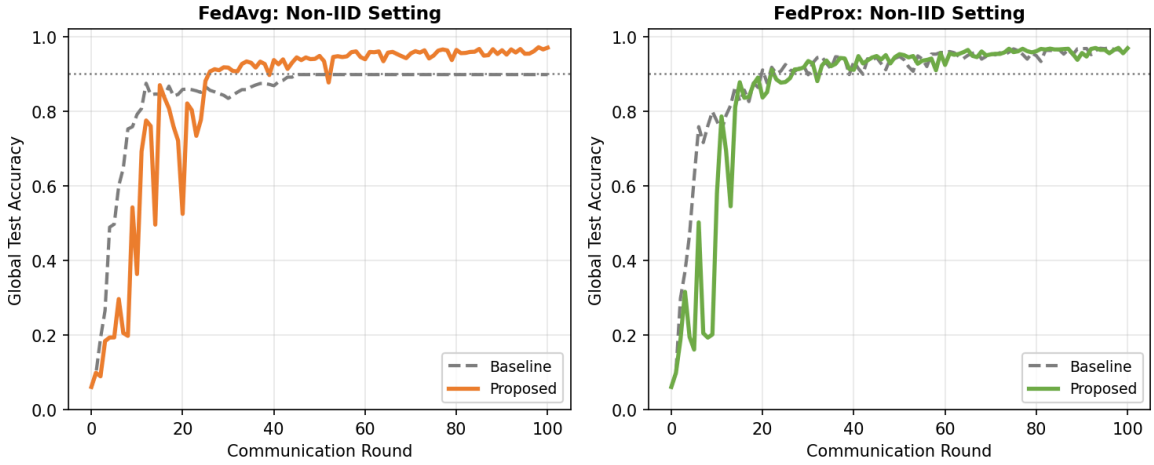


Figure 3. Side-by-side comparison of baseline vs. proposed accuracy curves under non-iid data distribution. left: fedavg setting. right: fedprox setting. the proposed strategy consistently reaches or surpasses the baseline trajectory (picture credit: original)

DOI: 10.54254/2755-2721/2026.34419

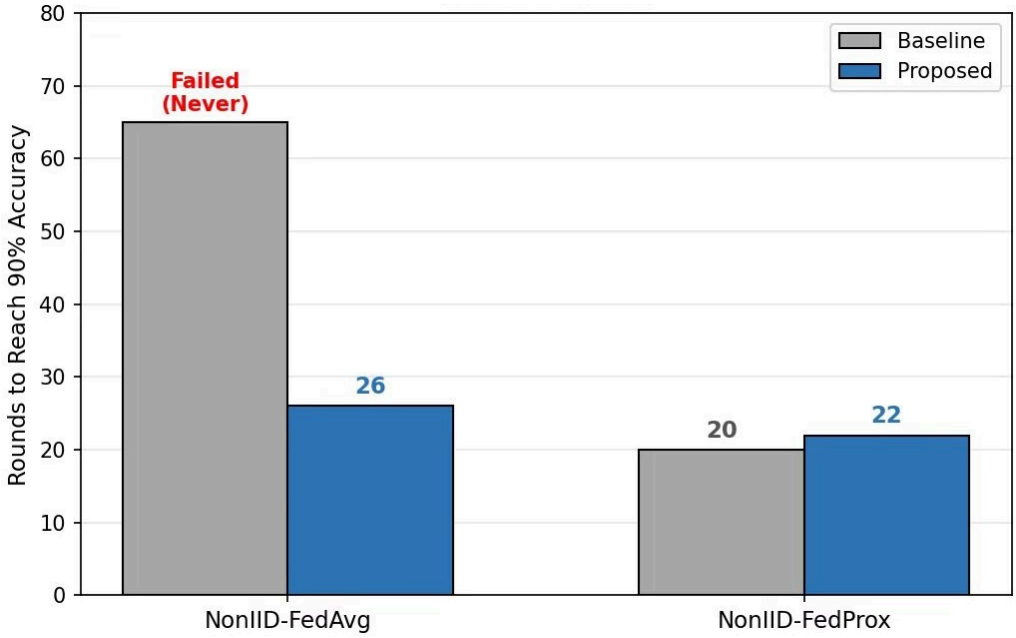


Figure 4. Communication efficiency: number of rounds required to first exceed 90% global test accuracy. lower is better. "failed" indicates the non-iid-fedavg-baseline never reached the 90% threshold within 100 rounds (picture credit: original)

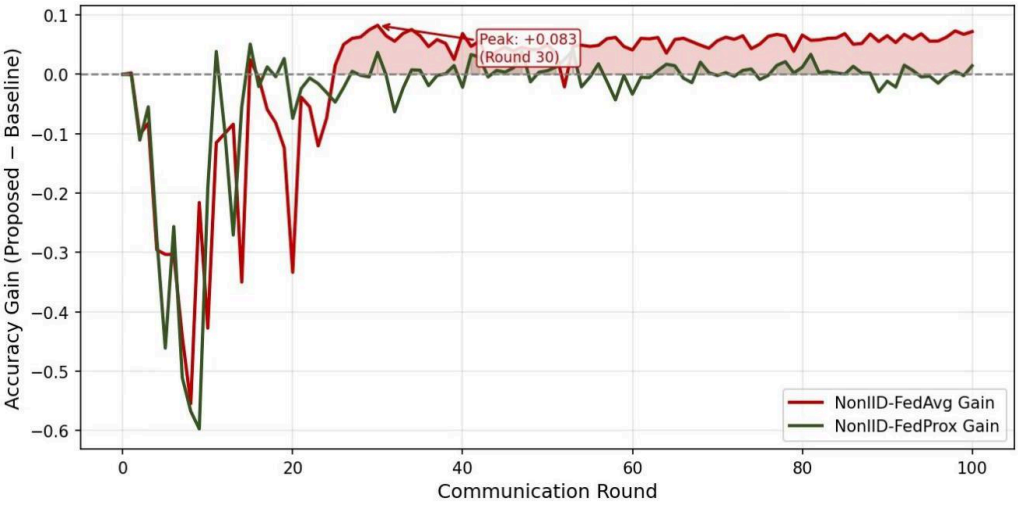


Figure 5. Per-round accuracy gain of proposed over baseline for non-iid configurations. positive values indicate that the proposed FACS strategy outperforms the corresponding baseline. the shaded region marks rounds at which the gain is positive (picture credit: original)

As shown in Figure 2 and Table 1, the final global accuracy of the IID-FedAvg-FACS configuration reached 97.76% in the 100th round, which was consistent with the performance of the IID-FedAvg-baseline configuration, and both exceeded 90% accuracy for the first time in the fourth round. This result confirms that the proposed FACS strategy will not bring measurable performance degradation when the data distribution is balanced.

Under the non independent identically distributed (non-IID) setting, the difference between different strategies is particularly obvious. The final accuracy rate of non-IID-FedAvg-baseline

DOI: 10.54254/2755-2721/2026.34419

configuration is only 89.76%, which has never exceeded the threshold of 90%, because it systematically excludes clients with poor performance, thus falling into local optimization. In contrast, the non-IID-FedAvg-FACS configuration finally achieved an accuracy of 96.96%, exceeding 90% in the 26th round, as shown in Figure 3, an increase of+7.20%. In addition, the accuracy of non-IID-FedProx-FACS configuration reached 96.82%, while the corresponding baseline configuration was only 95.37%, an increase of+1.45%. Meanwhile, the non-IID-FedProx-FACS converged to 90% in the 22nd round, while the baseline configuration reached this level in the 20th round.

Figure 5 further shows the change of accuracy improvement over time. Under the non-IID-FedAvg setting, due to the intentional introduction of clients with poor performance in the FACS strategy, its contribution will temporarily interfere with the global model at the initial stage, resulting in low accuracy. However, after the 15th round, the accuracy improvement remained positive and reached a peak of about +8% in the 18th round, which reflected the key stage of the strategy successfully guiding the global model to focus on the previously ignored data distribution. In the rest of the training, the accuracy rate was steadily improved, which confirmed the continuous advantage of the strategy on the baseline configuration.

3.2. Discussion

The results show that under the non-IID condition, the main failure mode of the standard FedAvg algorithm is the systematic bias to most types of customers, which accumulates with the training rounds, and finally makes the global model fall into the local optimum of 89-90 accuracy. The proposed FACS strategy solves this problem by reversing the relationship between customer accuracy and selection probability, so that customers with poor performance can get more frequent training opportunities, and its contribution is amplified in model aggregation. This mechanism is similar to the importance weighted sampling in statistical learning, and can dynamically correct the distribution mismatch caused by the Division of non independent identically distributed data.

At the beginning of the training, i.e. rounds 1 to 15 in Figure 4, the accuracy temporarily declined, which is the expected effect of actively introducing poor clients. The data distribution held by these customers is not learned by the current global model, so their local updates will conflict with the direction of most gradients. However, this short-term fluctuation is necessary: the FACS strategy promotes the global model to contact those data distributions that are not representative at the initial stage of training, thus avoiding the irreversible convergence of the model to biased local minima, which is the problem faced by the baseline method. This finding coincides with the theory of Li et al. [6]. Even if it needs to pay the price of short-term instability, it is also necessary to allow different customers to make contributions, which is crucial to converge to a fairer global optimal solution under the condition of non-IID.

Compared with the baseline, the non-IID-PedProx-FACS showed the narrowest performance gap, with corresponding accuracy rates of 96.82% and 95.37%, respectively, indicating that the near term in FedProx provides a positive effect on customer drift, so there is relatively limited space to further improve performance only through scheduling optimization. Nevertheless, there is still a stable gain of 1.45%, indicating that FACS provides benefits even if local optimization objective is strengthened. Future research can explore the joint adjustment of adaptive near end coefficient and fair scheduling weight to further improve the model convergence under non-IID conditions.

Some limitations of this study should be acknowledged. Firstly, MNIST dataset is used in all experiments, which is a relatively simple visualization benchmark; In more complex sensing tasks, such as lidar obstacle detection or multimodal scene understanding, the observed accuracy level may

DOI: 10.54254/2755-2721/2026.34419

overestimate the performance. Secondly, the energy dynamics model uses uniform charge and discharge parameters, and the physical battery model constraint based on actual vehicle hardware calibration is more realistic. Third, this paper evaluates a single trial for each configuration. Repeating the trial for multiple random seeds and reporting variance estimates will enhance the statistical reliability of the reported results. Addressing these limitations is a clear direction for future investigations, which may include personalized FL methods for customers with persistent accuracy defects [5].

4. Conclusion

This paper proposes a FACS strategy for Federated Learning in energy-constrained Internet of Vehicles environments. The strategy allocates vehicle selection priority according to the inverse ratio of model accuracy, uses composite weighting to enhance the aggregate contribution of underperforming customers, and solves the local optimization problem caused by non-IID data in standard FedAvg. In the MNIST experiment of 20 simulated vehicles, the global test accuracy of FACS under non-IID conditions is 7.20% higher than that of FedAvg baseline, and the target accuracy of 26 wheels is 90%, and the performance under IID conditions is not degraded. This highlights the potential of fair driven scheduling as a lightweight plug-in to enhance FL aggregation algorithm, which is suitable for distributed learning of vehicle edge computing, automatic driving and network edge privacy protection.

However, this study only verifies FACS in static simulation scenarios and simple data sets, without considering the lack of accuracy, time-varying communication links and complex real IOV environment. In the future, the strategy will be optimized to adapt to the dynamic vehicle Edge network and large-scale real traffic data, and promote its engineering application.

References

[1] Lu, Z., Qu, G., & Liu, Z. (2019). A survey on recent advances in vehicular network security, trust, and privacy. IEEE Transactions on Intelligent Transportation Systems, 20(2), 760–776.
[2] Patil, R., Yadav, R., & Mishra, K. (2025). Federated learning in vehicular networks for resource management, privacy, and security: Challenges and future directions. IEEE Access, 13, 190766-190785.
[3] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 1273–1282.
[4] Du, Z., Wu, C., Yoshinaga, T., Yau, K.-L. A., Ji, Y., & Li, J. (2020). Federated learning for vehicular internet of things: Recent advances and open issues. IEEE Open Journal of the Computer Society, 1, 45–61.
[5] Chellapandi, V. P., Yuan, L., Brinton, C. G., Żak, S. H., & Wang, Z. (2024). Federated learning for connected and automated vehicles: A survey of existing approaches and challenges. IEEE Transactions on Intelligent Vehicles, 9(1), 119–137.
[6] Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. Proceedings of the 3rd MLSys Conference, 429–450.
[7] Zhu, H., Xu, J., Liu, S., & Jin, Y. (2021). Federated learning on non-IID data: A survey. Neurocomputing, 465, 371–390.
[8] Yuan, X.-T., & Li, P. (2022). On convergence of FedProx: Local dissimilarity invariant bounds, non-smoothness and beyond. Advances in Neural Information Processing Systems. https: //arxiv.org/abs/2206.05187
[9] Jimenez G., D. M., Solans, D., Heikkila, M., Vitaletti, A., Kourtellis, N., Anagnostopoulos, A., & Chatzigiannakis, I. (2024). Non-IID data in federated learning: A survey with taxonomy, metrics, methods, frameworks and future directions. arXiv. https: //arxiv.org/abs/2411.12377
[10] Zhou, Y., Li, J., Zuo, X., & Liu, Y. (2025). Influence evaluation-based fair federated learning on non-IID data in Internet of Vehicles. IEEE Internet of Things Journal, 12(19), 40906-40918.

DOI: 10.54254/2755-2721/2026.34419

[11] Zhang, X., et al. (2023). Energy efficient federated learning in internet of vehicles: A game theoretic scheme. Transactions on Emerging Telecommunications Technologies.
[12] LeCun, Y., Cortes, C., & Burges, C. J. (1998). The MNIST database of handwritten digits. AT& T Labs. http: //yann.lecun.com/exdb/mnist/
[13] Shi, Y., Liu, Z., Shi, Z., & Yu, H. (2023). Fairness-aware client selection for federated learning. Proceedings of the 2023 IEEE International Conference on Multimedia and Expo (ICME), 324–329.
[14] Khajehali, N., Yan, J., & Chow, Y.-W. (2025). Fairness evaluation of client selection methods in federated learning: A comparative study. Proceedings of the 2025 6th International Conference on Computers and Artificial Intelligence Technology (CAIT), 27–32.
[15] Arouj, A., & Abdelmoniem, A. M. (2022). Towards energy-aware federated learning on battery-powered clients. Proceedings of the 1st ACM Workshop on Data Privacy and Federated Learning Technologies for Mobile Edge Network, 7–12.
[16] Fu, L., Zhang, H., Gao, G., Zhang, M., & Liu, X. (2023). Client selection in federated learning: Principles, challenges, and opportunities. IEEE Internet of Things Journal, 10(24), 21811–21819.
[17] Cha, N., & Chang, L. (2024). Multi-objective distributed client selection in federated learning-assisted internet of vehicles. Sensors, 24(13), 4180.
[18] Li, L., Fan, Y., Tse, M., & Lin, K.-Y. (2020). A review of applications in federated learning. Computers & Industrial Engineering, 149, 106854.