Evaluation of Client Selection Mechanisms in Vehicular Federated Learning Environments with Client Failures

John Sousa[1], Eduardo Ribeiro[1], Lucas Bastos[1] Denis Rosário[1], Allan M. de Sousa[2], and Eduardo Cerqueira[1]

[1] Federal Univeristy of Pará (UFPA), Belém, Brazil

[2]University of Campinas (UNICAMP), Campinas, Brazil

{john.sousa, lucas.bastos, eduardo.ribeiro}@itec.ufpa.br,

allanms@unicamp.br, {denis, cerqueira}@ufpa.br

Abstract. *Federated Learning (FL) emerges as a promising solution to enable collaborative model training for autonomous vehicles while preserving privacy and communication overhead issues. An efficient selection of clients to participate in the training process is still challenging, especially in scenarios with statistical heterogeneity of data distribution and client failure events. Client failure is an uncontrollable event in the training process that reduces accuracy, convergence, and speed. Therefore, investigating the performance of client selection mechanisms in this scenario is crucial. This paper presents a reliability and robustness analysis of entropy-based client selection mechanisms in FL environments with client failure. The results demonstrated that entropy-based selection outperformed the other methods regarding training loss, accuracy, and AUC, particularly in high client dropout scenarios. These findings show the importance of considering entropy data for client selection when addressing the challenges posed by client failure in FL scenarios.*

1. Introduction

Data privacy preservation emerges as a paramount concern in smart cities, especially in sensitive domains, which become notably complex to address in the context of Connected and Autonomous Vehicles (CAVs). CAVs are equipped with onboard sensors, including cameras, RADAR, LiDAR, and proximity and temperature sensors, to collect multi-modal data, such as navigation, perception, obstacle detection, and vehicle control. CAVs rely on vehicular network technology to enable data sharing with neighbors and edge servers, providing data processing for a cooperative understanding of the environment among vehicles and infrastructure entities [Zhang et al. 2023b]. Vision-related tasks, such as steering wheel angle prediction [Zhang et al. 2021b], traffic sign recognition [Stergiou et al. 2022], semantic segmentation [Fantauzzo et al. 2022], object detection [Jallepalli et al. 2021], and driver monitoring [Yuan et al. 2023] typically use images captured by the camera as the data source. In this context, Deep Learning (DL) plays a pivotal role with its ability to extract meaningful patterns and insights from large datasets. By leveraging these datasets, services such as route optimization, predictive maintenance, real-time decision-making, and personalized in-vehicle experiences can be enhanced [Pervej et al. 2023].

Traditional Machine Learning (ML) approaches are based on cloud-centric architecture where data is stored and processed centrally on the cloud server

[Zhang et al. 2021a]. However, the widespread data sharing between CAVs and servers poses significant privacy risks and demands substantial network bandwidth. In response to these challenges, there is an urgent need for a privacy-preserving distributed ML solution for CAV environments. In addition, it is believed that the future of ML and cloud computing schemes will be distributed at the network edges [Zhang et al. 2023c]. In recent years, ML-driven Federated Learning (FL) has been gaining much attention in this area due to its decentralized nature, allowing data training locally on devices, enabling multiple clients (*i.e.*, CAVs) to collaboratively train a shared model without sharing individual information [Lian et al. 2022].

FL relies on a robust and always connected client selection mechanism deployed at the edge server to choose a group of clients with valuable samples for the model training at each communication round [Xiong et al. 2023, de Souza et al. 2024]. These selected clients receive the global model, conduct training based on their local data, and then share their model parameters instead of transmitting their raw sensing data, as described in [Sousa et al. 2023]. Afterward, a given aggregation policy aggregates the shared local models at the cloud or edge servers to produce an accurate global model. Finally, the updated global model is distributed to the clients. In this way, FL allows ongoing learning by adapting the ML model without sharing raw data, provides privacy preservation by keeping the collected data stored on the CAVS, and avoids the potential communication overhead that the heavy data traffic of CAV information can cause. Hence, integrating FL in CAV systems opens up various possibilities for enhancing vehicular intelligence while addressing privacy, security, and communication challenges [Zhang et al. 2023b].

The data distribution highly impacts the client selection mechanism since data are not independent or have different statistical distributions, *i.e.*, non-IID data scenarios. This statistical heterogeneity results in lower classification accuracy, introducing non-representativeness issues and potentially decreasing model accuracy and fairness among the participating entities. In this way, it is essential to develop a client selection mechanism that can handle non-IID data in dynamic and mobile environments without compromising classification accuracy in FL over CAV scenarios [Nguyen et al. 2022]. In addition, clients might fail to provide their local model updates, interfering with the FL's ability to learn effectively [Zhang et al. 2021a]. These failures result from different reasons, such as, insufficient computing resources, client abort, network failure, etc, where clients might experience different rates of failure owing to their heterogeneous composition [Huang et al. 2022]. In this way, only a subset of clients can complete local training and transmit the model updates in each round, reducing the accuracy, convergence, and training speed [Zhang et al. 2021a].

When clients fail to contribute their local model updates, the overall training data available for the global model update is reduced, and it obtains a biased update that deviates from the desired global model [Sun et al. 2023]. This reduction in training data leads to slower convergence of the global model and decreased model accuracy. In this context, it is essential to design a robust and reliable client selection mechanism for FL over CAV systems, which can be based on random clustering, entropy, and other approaches [Shanmugarasa et al. 2023]. Entropy-based client selection is a promising approach for an FL over CAV scenario with client failure since entropy enables identifying the most relevant client with more diverse data for learning models, capturing the heterogeneity of

FL over CAV scenario. In this sense, clients with high entropy ensure that the learned models represent the entire network and capture the scenario variations to improve the accuracy of round training more robustly. However, it is essential to understand the impact of arbitrary client failure and how it affects the performance of an entropy-based client selection mechanism, which are the central questions of this paper.

In this paper, we assess the reliability and robustness of an entropy-based client selection mechanism where clients have a higher probability to fail due to various kinds of reasons and indifferent levels of frequency [Sousa et al. 2023]. In its operations, the entropy-based mechanism measures the randomness or unpredictability of a system based on the entropy of client data, where it prioritizes clients with more diverse and representative data. In this way, it only selects a subset of clients with the better-suited data for model training, *i.e.*, clients are selected based on their data distribution heterogeneity. We assessed the reliability and accuracy of the client selection mechanism with different client failure rate in a non-IID scenario. Simulation results demonstrate that, even in the face of client failures, our entropy-based selection strategy consistently outperforms existing client selection mechanisms. These findings motivate the detailed analysis of client failures effects in FL environments presented in this work.

The remainder of this paper is structured as follows. Section 2 presents an overview of well-known client selection works and their main drawbacks. Section 3 presents the system model, and the entropy-based client selection mechanism. Section 4 discusses the simulation results. Finally, Section 5 concludes the paper and presents some future work directions.

2. Related Work

Previous studies have examined the challenges of FL in vehicular networks, mainly focusing on non-IID data and biased device data distributions. For example, [Zhu et al. 2021] noted that non-IID data on local devices significantly impacts model performance, contrasting with centralized learning. This research assesses the effects of non-IID data on parametric and non-parametric ML models across different FL settings. It reviews prior studies, proposes specific strategies, and evaluates the advantages and disadvantages of these methods. Moreover, client failures complicate training with heterogeneous data, intensifying the non-IID issue, with existing algorithmic solutions unable to fully bridge the gap in local and global loss minimization. However, this research must extensively address resilience and how these algorithms fare under dynamic conditions such as client failures.

In our previous work, [Sousa et al. 2023], we proposed a novel entropy-based client selection, which ranks the entropy of label data of the users in the area and selects the 25% highest entropy values. Despite the promising results, a more detailed investigation regarding the resilience of the selection method in the face of the challenges of the network, such as client failures, is needed.

[Shanmugarasa et al. 2023] highlighted issues stemming from security, privacy concerns, and the intricacies of FL processes, particularly the increased computational burden on clients. These challenges may impact specific clients or affect the entire network, with privacy management being a universal concern. The study concludes that collaborative efforts between servers, platforms, and clients are imperative to effectively address client-side challenges in the FL ecosystem. While advocating for collaborative solutions, the work only extensively explores the intersection of these challenges with a scenario with client failures. A more nuanced comparison with a work emphasizing failures analysis could delve into how client-related challenges impact the overall robustness and performance of FL algorithms, especially in adverse conditions.

[Wang and Xu 2023] explored the challenges posed by client failures in FL, emphasizing a key distinction from client sampling, noting that failure introduces uncontrollable client participation. While previous studies often focused on actively managing client participation through sampling, Wang et al. bring attention to passive partial participation. In this scenario, clients fail involuntarily due to external events, an aspect less explored in existing literature. This perspective adds valuable insights into the impact of unplanned client failure on the performance and robustness of FL algorithms.

[Souza et al. 2023] addressed the communication challenges and scalability issues by dynamically adapting the number of participating devices and training rounds through a client selection strategy that selects the clients whose accuracy falls below the average. Using a containerized environment, DEEV showcases significant reductions of up to 60% in communication overhead and an impressive 90% in computation overhead compared to existing approaches. Its robust performance in scenarios with non-IID data underscores its potential for enhancing FL model efficiency. However, the work considers only an environment where every client is available and is stationary, a scenario different from the scenarios usually present in vehicular networks.

[Sun et al. 2023] studied the convergence performance of the classic Federated Learning Average (FedAVG) aggregation algorithm in scenarios involving arbitrary client failures. The theoretical analysis indicated that client failures lead to biased updates in each training iteration. When employing the commonly used strategy of a decaying learning rate, the model trained by FedAvg may, in the worst-case scenario, exhibit oscillations around a stationary point of the global loss function. A cross-device federated learning system simulation was carried out to validate these findings, incorporating various client failure patterns.

[Huang et al. 2022] investigated the vital topic of client selection in a fluctuating environment. They acknowledged that choosing particular clients for each synchronous round in FL training significantly affects both training efficiency and the ultimate performance of the model. In the context of heterogeneous clients experiencing varying degrees of training failure, their research defined the client selection problem by simultaneously considering effective participation and fairness. Seeking to balance training efficiency with reduced bias, the authors introduced E3CS, a stochastic client selection strategy. The experimental results using a public dataset showed that E3CS leads to quicker convergence towards a predetermined model accuracy while retaining the same level of final model accuracy compared to leading-edge selection methods.

Upon reviewing the challenges faced by previous FL environments' client selection mechanisms, it becomes clear that understanding and investigating the impact of client failures is crucial for the robustness and reliability of FL in connected and autonomous vehicles (CAVs). Client failures disrupt the learning process and skew model performance, degrading performance and affecting the viability of client selection methods; thus, an analysis of these FL systems under special conditions becomes important.

3. System Model

This section outlines a common CAV-based FL scenario involving client failures, where clients may fail for various reasons and at different frequencies. We then detail an entropy-based client selection method and its operational mode. This method assesses the data entropy of each CAV and selects the top-ranked clients for the next training round based on their local data. Subsequently, we present the system model and the operations of the entropy-based mechanism.

3.1. FL over CAVs Environment

We envisage a scenario involving a set of n CAVs navigating an urban environment. Each CAV, denoted by an index i within the range $[1, n]$ and represented as $C = \{c_1, c_2, c_3, ..., c_n\}$. Every CAV c_i moves in a specific direction and maintains a speed s_i within the range of the minimum speed (s_{min}) and maximum speed (s_{max}). Each CAV c_i is equipped with onboard sensors and collects data crucial for ML applications, such as recognition or image classification. In this way, each CAV c_i has local dataset $D_i \in \{D_1, \ldots, D_n\}$ distributed in a non-IID manner, which contains a set of *features* $x_{k,i}$ with $k \in \{1, \ldots, \|D_i\|\}$ associated with a *label* $y_{k,i}$.

Furthermore, each CAV c_i is equipped with a Vehicle-to-Infrastructure (V2I) communication interface, such as, Dedicated Short Range Communication (DSRC) or 5G, which is used to communicate with the edge server ES through the core network. The edge server ES that plays a pivotal role in distributing ML parameters for the initial or updated global model ω to all CAVs during each communication round μ. Moreover, the edge server ES assumes responsibility for collecting and analyzing entropy data, and also for model aggregation.

We considered the typical FL architecture, where the process starts with the initialization of a global model M_g on a central server. At each communication round μ, a subset of k CAVs denoted as $V = \{v_1, v_2, v_3, ..., v_k\}$ is selected to receive the global model M_g and perform the training based on its Dataset D_i. Each selected client v_i is able to train a model architecture A to obtain the local model W_i based on the local dataset D_i. In this way, each client v_i trains the local model W_i to minimize a loss function l for better convergence with a minimum accuracy value across users. Specifically, the local loss $l(W_i, D_i)$ is defined as the average loss based on the prediction error, across all predictions for the dataset D_i using the weights W_i, which is computed based on Eq. 1.

$$l(W_i, D_i) = \frac{1}{\|D_i\|} \sum_{k=1}^{\|D_i\|} f(W_i, x_{k,i}, y_{k,i}) \tag{1}$$

In the aggregation phase, the model updates, *i.e.*, learned parameters or gradients W_i, are sent periodically to the edge server ES, which applies a given aggregation policy, such as, FedAVG. Specifically, FedAVG computes an average of the shared local models W_i at edge server ES to produce an accurate global model M_g, which is transmitted back to the participating CAVs. In addition, the edge server ES defines the number of k selected clients based on a client selection mechanism, such as, Entropy-Based.

3.2. Client Failure Model

Client failure in FL over CAVs refers to clients' cessation of active participation in the collaborative model training process [Wang and Xu 2023]. This phenomenon could result from different factors, such as vehicle mobility due to intermittent connectivity issues during transitions between Roadside Units(RSU); connectivity problems caused by temporary or permanent disconnection due to network disruptions; intentional withdrawal, where clients opt out voluntarily due to privacy concerns or limited resources; and resource limitations, as seen in devices with constrained battery life choosing to drop out strategically.

To better illustrate the concept, we consider a standard FL algorithm with clients working together to train the same global model. In a scenario of client failure, only a random subset of selected clients will be trained in each round due to a failure event. This failure affects the training process of a FL system, causing reduced accuracy, increased bias, and compromised fairness. They impede model convergence, and client mobility leads to inconsistent data contributions. Reliable and robust FL algorithms must adjust to sporadic client participation caused by failures. Strategies that accommodate intermittent client presence and optimize model aggregation under varying network conditions enhance stability in dynamic FL environments.

Figure 1 depicts a representation of a typical FL over CAV environment with client failures. In this scenario, a given client v_i has a probability $P(v_i)$ to be selected for participation in the FL process, which depends on the metric used to select the client. For instance, the probability $P(v_i)$ can be a random number in a random selection, can be proportional to the entropy that represents the data randomness or unpredictability in an entropy-based mechanism [Sousa et al. 2023] or can be clients with accuracy lower than the average of all participating clients [Souza et al. 2023]. Furthermore, let $Q(v_i)$ represents the probability of a client v_i being disconnected from the training process, which can be influenced by different factors, such as network stability, device power, or communication issues.

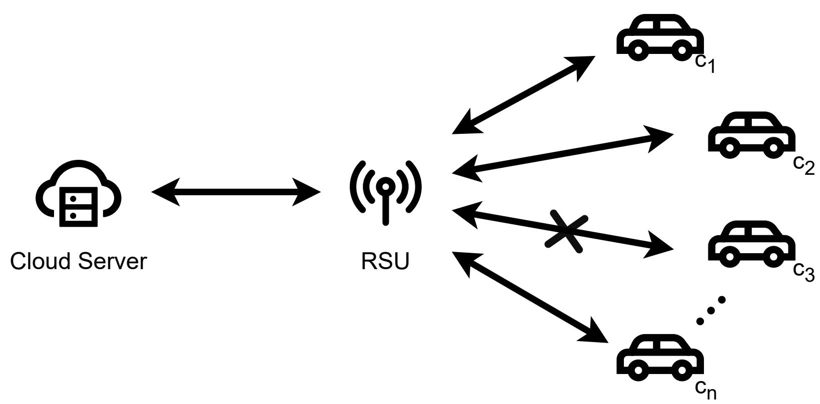

Figure 1. Representation of client failures in a FL over CAV environment

The overall probability of a client v_i being selected and then dropping out during training is given by the product of the selection and failure probabilities:

$$P_{failure}(v_i) = P(v_i) \times Q(v_i) \quad (2)$$

If we want to consider multiple clients potentially failing independently, we can define an overall failure probability for the entire set of clients:

$$P_{totalfailure}(V) = \Pi_{i=1}^{k} P_{failure}(v_i) \quad (3)$$

3.3. Entropy-Based Client Selection Mechanism

FL for CAV applications presents unique challenges due to the ever-changing nature of CAV mobility and the critical objectives of enhanced privacy and reduced server load. The issues of client failure and varying data diversity can significantly impact the FL process. To tackle these challenges effectively, it is worth considering the adoption of entropy as a promising criterion for client selection. Specifically, entropy plays a vital role in information theory, serving as a crucial measure for quantifying a system's level of randomness or disorder. The use of entropy in evaluating FL algorithms holds the potential to identify the most relevant and diverse data, contributing to the development of models that effectively encapsulate the heterogeneity in the context of FL over CAVs.

In the context of CAVs, data diversity is crucial due to the inherent variability in driving behavior and network conditions. Entropy directly measures this uncertainty, providing a more comprehensive view than metrics like the Gini index, which focuses on specific aspects like class imbalance. Unlike clustering techniques that require predefined clusters, entropy facilitates efficient client selection based on data randomness. This computational efficiency aligns well with the resource limitations of CAV environments and the need to minimize server load.

In this context, an entropy-based client selection mechanism gives preference to select clients based on the entropy of their data, using it as an indicator of data diversity and representativeness. By selecting clients with high entropy, FL algorithms can ensure that the learned models represent the entire network and capture the variations in driving behavior, traffic patterns, and network connectivity. This approach also has the potential to enhance the model's robustness in the face of unpredictability in vehicular networks. Hence, the entropy-based client selection mechanism has shown promise in significantly reducing the data variability contributions and managing the challenges associated with uncertain client availability.

Figure 2 depicts the entropy-based client selection workflow, encompassing entropy calculation, local model training and testing, as well as global model aggregation and update. The communication round involves five steps: 1) the edge server ES sends the current global model M_g to all CAVs C; 2) Each CAV c_i sends its calculated data entropy $H(d_n)$ to the edge server ES; 3) the edge server ES selects a set of clients C from the set of CAVs U that meets a specified threshold θ based on entropy ranking, described as $H(d_n) \geq \theta$. These clients will be selected to perform local model training; 4) the trained local models W_i are sent to the edge server ES for aggregation; 5) the edge server ES generates an updated global model M_g based on the aggregated local models, which is then sent back to all participants. We consider the Shannon Entropy to calculate the data entropy H(x), while P(x) denotes the probability of observing a particular value x in the dataset, and log is the natural logarithm, which is described in Eq. 4.

$$H(X) = -\sum_{x} P(x) \log P(x) \tag{4}$$

Clients whose datasets have a high level of entropy are selected because they contain diverse and informative data that can improve the performance of the FL model, such as described in Eq. 5. k_m refers to the class of the data point d_{ni}, which represents an individual data point in d_n.

$$H(d_n) = -\sum_{j=1}^{m} P(k_m) \log P(k_m) \tag{5}$$

By focusing on clients with higher entropy values, the client selection mechanism benefits from including varied data. Consequently, this approach enhances the model's accuracy and applicability across different FL over CAV scenarios, with client failures for various reasons and in different frequency levels.

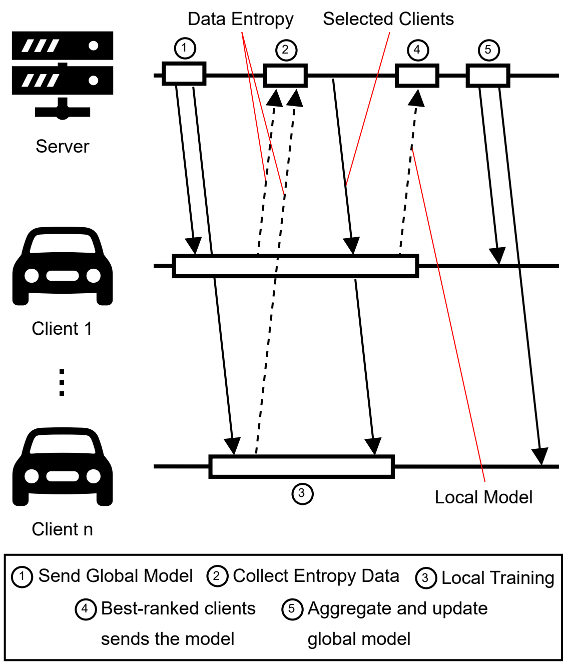

Figure 2. Entropy-Based Client Selection Mechanism

4. Evaluation

This section provides an overview of the evaluation environment and presents the results obtained. The discussion of the results revolves around analyzing the following metrics: training loss, accuracy, and Area Under the Curve (AUC).

4.1. Simulation Environment

We conducted a comprehensive simulation study using the PFLib, which is a flexible framework presented by [Zhang et al. 2023a] and available on GitHub [1]. The framework runs on a server with the following specs: i9-13900K(32), 128 GB RAM, and Dual RTX 4090 on a Ubuntu Server operating system. We consider a widely used public dataset, FMNIST, to train and test model validations. The CNN model used in the experiment has two convolutional layers with filter sizes of 5x5. A 2x2 max-pooling operation succeeds each convolutional layer. Furthermore, it is essential to consider that the data employed in this experiment follows a non-IID arrangement, resembling a realistic data distribution scenario, and is modeled using a Dirichlet distribution. This non-iid configuration was generated by a tool in PFLib, which defined the rate of the Dirichlet distribution at 0.1.

We consider a grid scenario with one km² composed of 58 clients as proposed by [Pannu et al. 2021] and use the Luxembourg SUMO Traffic (LuST) environment. Since those vehicular environments generate heterogeneous data due to the diverse behaviors of vehicles, varying speeds, and different routes, this diversity and complexity of data make the entropy-based client selection method particularly suitable. We consider a built-in feature within the framework to simulate client failures, as introduced in Section 3.2. This feature randomly selects a client to avoid sending updates and receiving models during a particular round. This capability allows us to explore the consequences of client failure on the reliability and robustness of client selection mechanisms in FL over CAV scenarios. It is worth noting that the failure rate is adjustable, providing the flexibility to control the extent of simulated failure events. We evaluate the impact of various failure rates in the scenario, considering scenarios with no failure, 16%, 33%, and 50% client failure rates. Table 1 summarizes the main simulation parameters used in our evaluation.

Table 1. Simulation parameters for experiment

Parameters	Value
Total Participant Clients	58 vehicles
Number of Rounds	100 Rounds
Learning Rate	0.001
Client Failure Rate	16%, 33%, 50%
Number of Epochs per Round	1
Network Model	CNN
Batch Size	10

We compared the three client selection methods: i) Normal selection is a baseline method that does not consider the quality or diversity of clients' data. It simply selects a random subset of clients to participate in each round of training; ii) DEEV selects clients that have lower accuracy than the average accuracy of all participating clients [Souza et al. 2023]; iii) Entropy-based client selection leverages data entropy to choose clients that contain diverse and informative data, such as introduced in Section 3.3.

The client selection mechanisms' performance was assessed using standard FL classification metrics: accuracy and loss. Accuracy is calculated by dividing correct predictions by total examples and quantifying correct predictions' proportion. All classes

[1] https://github.com/TsingZ0/PFLlib

have the same hit penalties to prevent misleadingly positive assessments, especially in imbalanced class proportions. The loss metric compares target and predicted values, evaluating the model's training data representation.

The AUC score, derived from the Receiver Operating Characteristic (ROC) curve, is significant in FL classification, measuring the model's ability to differentiate between positive and negative instances. It is valuable in imbalanced class distributions, offering nuanced evaluation and insight into discriminatory ability across threshold values.

4.2. Results

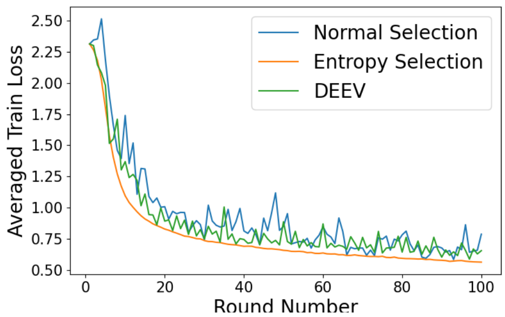

(a) Without client failures

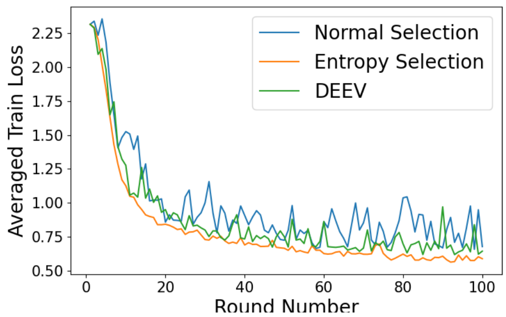

(b) 16% of client failures

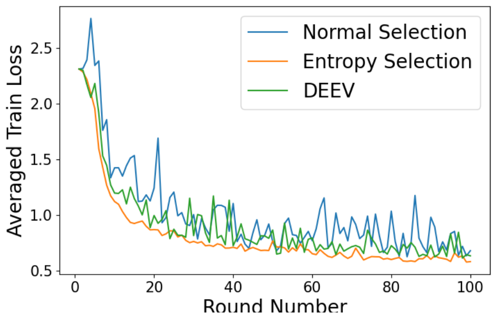

(c) 33% of client failures

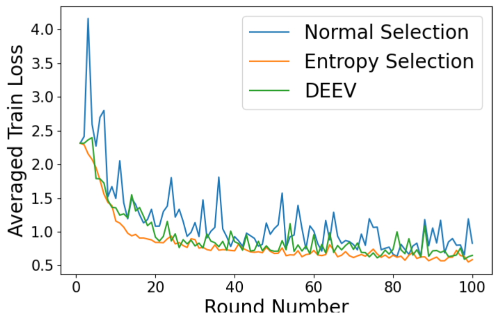

(d) 50% of client failures

Figure 3. Train Loss for different client selection mechanisms

Figure 3 shows the training loss for different client selection mechanisms under different client failure rates. By analyzing Figure 3 a), we can conclude that the entropy-based mechanism exhibits faster convergence than the other client selection mechanisms. As the client failure rate increases, all methods experience a decline in performance. However, the entropy-based mechanism maintains a slight advantage over the others. On the other hand, the DEEV strategy deteriorates to the extent that it exhibits slightly inferior performance compared to random selection, as shown in Figure 3 d).

The entropy-based client selection mechanism harnesses information entropy as its guiding principle, prioritizing clients that contribute diverse and informative data, ultimately creating a more representative model. This way, the mechanism demonstrates reduced instability in train loss metrics, exhibits faster convergence and maintains higher accuracy levels than random selection and DEEV mechanisms. This adaptability of the entropy-based mechanism to varying data distributions and the dynamic nature of FL over CAV scenarios contributes to its effectiveness in mitigating the impact of client failure with different levels of failure frequency, making it a valuable strategy for ensuring stability and top-notch performance in FL over CAV. Hence, the superiority of the entropy-based client selection mechanism shines through when faced with challenges related to client failure, as observed in [Wang and Xu 2023].

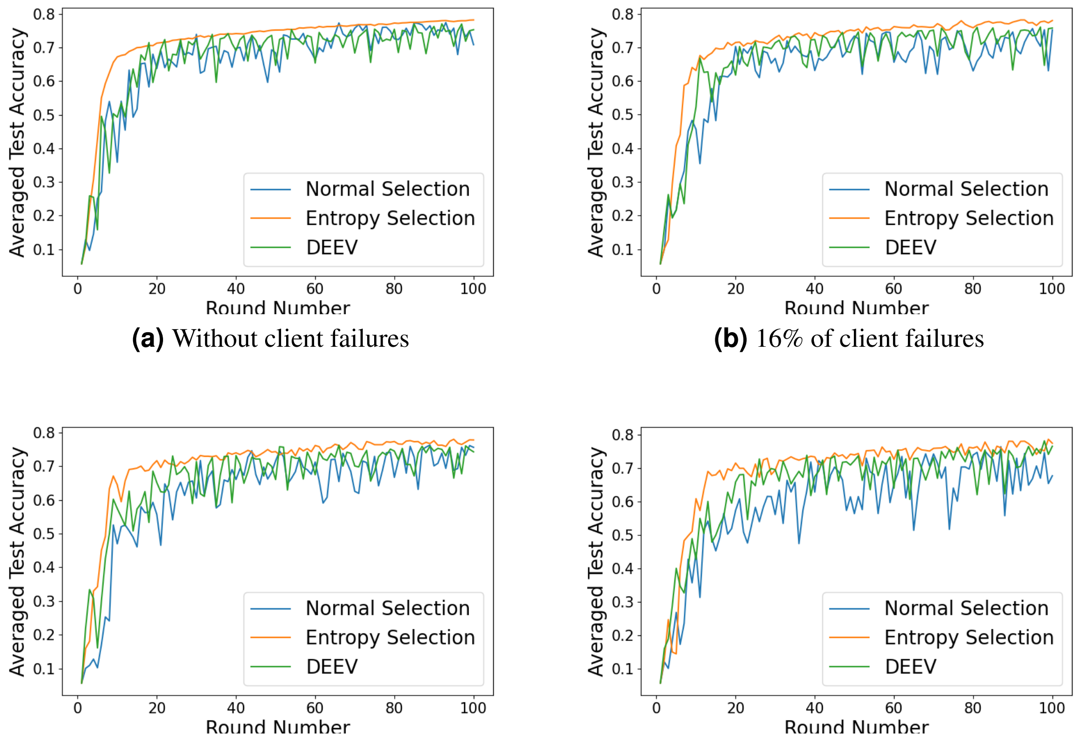

(a) Without client failures
(b) 16% of client failures
(c) 33% of client failures
(d) 50% of client failures

Figure 4. Accuracy for different client selection mechanisms

Figure 4 shows the accuracy results for different client selection mechanisms under different client failure rates. By analyzing the accuracy results, we notice a similar trend where the accuracy of all tested mechanisms deteriorates as the client failure rate increases. The entropy-based method consistently outperforms the other two mechanisms, even with high failure rates. In contrast, the DEEV mechanism exhibits a decline in performance to the extent that it falls below the performance of random selection, reaching its lowest point at a 0.5 client failure rate.

The entropy-based selection method consistently exhibited a performance advantage, even in high failure rates. This remarkable reliability and robustness can be attributed to its core principle of selecting clients with diverse datasets. Prioritizing clients based on entropy ensures the selection of clients that offer a broad spectrum of data characteristics, maintaining a robust and representative model. This mechanism proves to be effective even in challenging scenarios with high failure rates. In such situations, where clients' participation may significantly decrease, potentially losing data diversity and model accuracy, the entropy-based method shines through by preserving the model's performance and adaptability.

Figure 5 shows the AUC results for different client selection mechanisms under different client failure rates. When examining the AUC Score results, we observe a similar trend in metric degradation as observed with Accuracy and Train Loss. The entropy-based strategy consistently outperforms the other two methods, while the DEEV method experiences more significant performance degradation as client failure rates become increasingly severe. The entropy-based method's performance stability is further enhanced by its capability to mitigate the challenges associated with data skewness, a prevalent issue in non-IID data environments like those encountered in FL over CAV. In situations with high client failure rates, where data skewness is likely to be exacerbated, the entropy-based selection method ensures a well-balanced and comprehensive representation of the data. In turn, it diminishes the risks of overfitting specific client data patterns and promotes more efficient training, even when confronted with limited data resources.

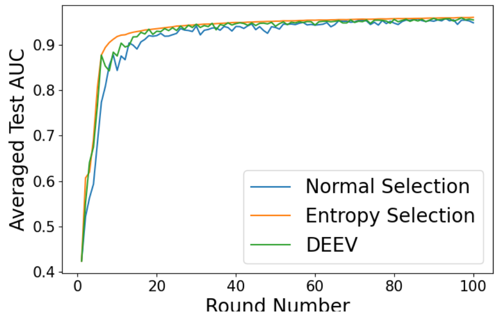

(a) Without client failures

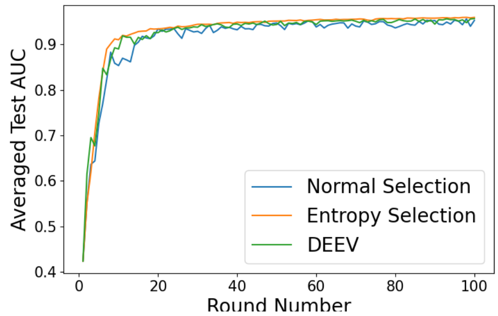

(b) 16% of client failures

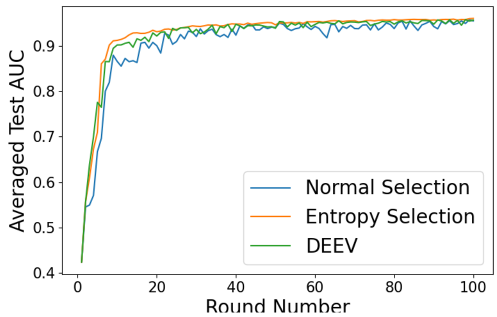

(c) 33% of client failures

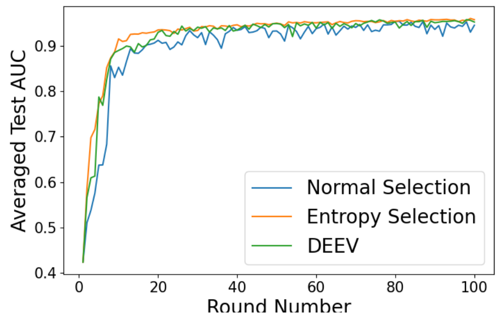

(d) 50% of client failures

Figure 5. AUC Score for different client selection mechanisms

5. Conclusion

This paper assessed the robustness and reliability of an entropy-based client selection mechanism in scenarios where vehicle dropouts can occur due to various failure events. The mechanism considers entropy to identify the most relevant and diverse data, contributing to developing models that effectively encapsulate the heterogeneity in the context of FL in CAV systems. The entropy-based client selection mechanism gives preference to select clients based on data diversity and representativeness, creating a more representative model. Simulation results presented the significance of incorporating entropy-based client selection when addressing the challenges presented by client failure events. This is because the entropy-based client selection mechanism ensures the selection of clients that offer a broad spectrum of data characteristics, resulting in the maintenance of a robust and representative model. Entropy selection demonstrated faster convergence, reduced instability, and high accuracy compared to the other mechanisms.

Our future research aims to explore adaptive client selection mechanisms capable of dynamically responding to fluctuations in network conditions and vehicular mobility patterns. Such adaptability would enhance the overall robustness of FL over CAV scenarios, ensuring their effectiveness even in dynamic and challenging environments. Additionally, we intend to investigate the integration of privacy-preserving mechanisms tailored for vehicular settings. This exploration will evaluate how such mechanisms influence the resilience of client selection strategies, contributing further to the development of secure and reliable FL frameworks in dynamic and uncertain scenarios.

Acknowledgment

This project was supported by the Ministry of Science, Technology, and Innovation of Brazil, with resources granted by the Federal Law 8.248 of October 23, 1991, under the PPI-Softex. The project was coordinated by Softex and published as Intelligent agents for mobile platforms based on Cognitive Architecture technology [01245.013778/2020-21].

References

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.

Fantauzzo, L., Fanì, E., Caldarola, D., Tavera, A., Cermelli, F., Ciccone, M., and Caputo, B. (2022). Feddrive: Generalizing federated learning to semantic segmentation in autonomous driving. In *prooceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11504–11511. IEEE.

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.

Jallepalli, D., Ravikumar, N. C., Badarinath, P. V., Uchil, S., and Suresh, M. A. (2021). Federated learning for object detection in autonomous vehicles. In *prooceedings of the IEEE Seventh International Conference on Big Data Computing Service and Applications (BigDataService)*, pages 107–114.

Lian, X., Yuan, B., Zhu, X., Wang, Y., et al. (2022). Persia: An Open, Hybrid System Scaling Deep Learning-based Recommenders up to 100 Trillion Parameters. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 3288–3298.

Nguyen, A., Do, T., Tran, M., Nguyen, B. X., Duong, C., Phan, T., Tjiputra, E., and Tran, Q. D. (2022). Deep federated learning for autonomous driving. In *proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, pages 1824–1830. IEEE.

Pannu, G. S., Ucar, S., Higuchi, T., Altintas, O., and Dressler, F. (2021). Dwell time estimation at intersections for improved vehicular micro cloud operations. *Ad Hoc Networks*, 122:102606.

Pervej, M. F., Jin, R., and Dai, H. (2023). Resource Constrained Vehicular Edge Federated Learning With Highly Mobile Connected Vehicles. *IEEE Journal on Selected Areas in Communications*, 41(6):1825–1844.

Shanmugarasa, Y., young Paik, H., Kanhere, S. S., and Zhu, L. (2023). *A systematic review of federated learning from clients' perspective: challenges and solutions*, volume 56. Springer Netherlands.

Sousa, J. L. R., Lobato, W., Rosário, D., Cerqueira, E., and Villas, L. A. (2023). Entropy-based client selection mechanism for vehicular federated environments. In *Proceedings of the 22nd Workshop on Performance of Computer and Communication Systems (WPERFORMANCE)*, pages 37–48. SBC.

Souza, A., Bittencourt, L., Cerqueira, E., Loureiro, A., and Villas, L. (2023). Dispositivos, eu escolho vocês: Seleção de clientes adaptativa para comunicação eficiente em aprendizado federado. In *Anais do XLI Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC)*, pages 1–14, Porto Alegre, RS, Brasil. SBC.

Stergiou, K. D., Psannis, K. E., Vitsas, V., and Ishibashi, Y. (2022). A federated learning approach for enhancing autonomous vehicles image recognition. In *2022 4th International Conference on Computer Communication and the Internet (ICCCI)*, pages 87–90. IEEE.

Sun, Y., Member, G. S., Mao, Y., and Zhang, J. (2023). MimiC:Combating Client Dropouts in Federated Learning by Mimicking Central Updates. *arXiv:2306.12212v3*, pages 1–17.

Wang, H. and Xu, J. (2023). Combating Client Dropout in Federated Learning via Friend Model Substitution. *arXiv preprint arXiv:2205.13222*.

Xiong, Y., Wang, R., Cheng, M., Yu, F., and Hsieh, C.-J. (2023). Feddm: Iterative distribution matching for communication-efficient federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16323–16332.

Yuan, L., Su, L., and Wang, Z. (2023). Federated transfer-ordered-personalized learning for driver monitoring application. *IEEE Internet of Things Journal*.

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

Zhang, H., Bosch, J., and Olsson, H. H. (2021b). End-to-end federated learning for autonomous driving vehicles. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.

Zhang, J., Hua, Y., Wang, H., Song, T., Xue, Z., Ma, R., and Guan, H. (2023a). Fedala: Adaptive local aggregation for personalized federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 11237–11244.

Zhang, X., Chang, Z., Hu, T., Chen, W., Zhang, X., and Min, G. (2023b). Vehicle Selection and Resource Allocation for Federated Learning-Assisted Vehicular Network. *IEEE Transactions on Mobile Computing*, PP:1–12.

Zhang, X., Liu, J., Hu, T., Chang, Z., Zhang, Y., and Min, G. (2023c). Federated learning-assisted vehicular edge computing: Architecture and research directions. *IEEE Vehicular Technology Magazine*, pages 2–11.

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