

Coalitional Game-Based Clustered Federated Learning for Mitigating Data Heterogeneity in Wireless Networks

Yu-Wei Chan, Jhih-Yu Tsai[†], Feng-Tsun Chien[‡], Dongyeon Kim[¶]

Bachelor Degree Program of Artificial Intelligence, National Taichung University of Science and Technology, Taiwan

Email: ywchan1103@nutc.edu.tw

[†]Institute of Artificial Intelligence Innovation, National Yang Ming Chiao Tung University, Taiwan

Email: arnold40905.ii12@nycu.edu.tw

[‡]Institute of Electronics, National Yang Ming Chiao Tung University, Taiwan

Email: ftchien@nycu.edu.tw

[¶]Department of Computer Science and Artificial Intelligence, Dongguk University, South Korea

Email: dongyeon@dgu.ac.kr

Abstract—Federated Learning (FL) has emerged as an effective paradigm for privacy-preserving distributed machine learning across mobile and edge devices. However, the presence of non-independent and identically distributed (non-IID) data in wireless networks often results in severe gradient divergence, slow convergence, and degraded model performance. To address these challenges, this paper proposes FedCoalition, a coalition game-based clustered federated learning framework that dynamically organizes clients into collaborative coalitions according to the similarity of their model updates. By integrating coalition formation with similarity-aware clustering, the proposed framework effectively reduces gradient divergence and enhances collaboration among clients with similar data characteristics. Furthermore, each coalition designates a Coalition Head (CH) to perform intra-coalition model aggregation and validate the aggregated model using an auxiliary validation dataset, thus improving model reliability while reducing communication overhead. Extensive experiments conducted on the MNIST and CIFAR-10 datasets under non-IID data distributions demonstrate that FedCoalition consistently achieves higher classification accuracy, faster convergence, and better robustness than existing clustered federated learning methods.

Index Terms—Federated Learning, Gradient Divergence, Coalition Game.

I. INTRODUCTION

With the rapid proliferation of Internet of Things (IoT) technologies, a vast number of mobile and edge devices continuously generate massive volumes of data through wireless networks, enabling a wide range of machine learning (ML) applications, such as image recognition, anomaly detection, and intelligent decision-making [1]. Conventional centralized learning requires transmitting raw data collected from distributed devices to a central server for model training. Such an approach incurs

substantial communication overhead, raises serious privacy concerns, and imposes considerable computational burdens on resource-constrained edge devices.

Federated Learning (FL) has emerged as an effective privacy-preserving distributed learning paradigm that enables clients to collaboratively train a global model while retaining their data locally and sharing only model updates with a central server [2]. By eliminating the need to transmit raw data, FL effectively reduces communication overhead and mitigates privacy risks, making it particularly suitable for intelligent mobile systems [1] and edge computing environments [3].

Despite these advantages, FL remains highly susceptible to data heterogeneity arising from non-independent and identically distributed (non-IID) data across participating clients. Such statistical heterogeneity often results in inconsistent local model updates, leading to severe gradient divergence, unstable optimization, slow convergence, and degraded global model performance [4]. Therefore, alleviating the adverse effects of gradient divergence under non-IID data distributions has become one of the most fundamental challenges in federated learning.

Clustered Federated Learning (CFL) mitigates the adverse effects of non-IID data by grouping clients with similar data characteristics and learning cluster-specific models. However, most existing CFL methods rely on heuristic clustering strategies, lack model validation before global aggregation, and fail to capture dynamic client interactions in wireless environments.

To overcome these limitations, we propose *FedCoalition*, a coalition game-based clustered federated learning framework built upon Hierarchical Federated Learning

(HFL) [5]. Clients are dynamically organized into coalitions according to the similarity of their model updates, while each coalition elects a Coalition Head (CH) [6] to perform intra-coalition aggregation and model selection using an auxiliary validation dataset. Furthermore, a similarity-aware coalition formation mechanism reduces gradient divergence and improves aggregation quality. Experimental results on the MNIST and CIFAR-10 datasets demonstrate that the proposed framework achieves faster convergence and higher accuracy than existing clustered federated learning methods.

The rest of this paper is organized as follows. Sec. II presents the system model and problem formulation. Sec. II presents the proposed methods. The experimental results are shown in Sec. III. Finally, Sec. IV gives the concluding remarks.

II. SYSTEM MODEL

Consider a wireless network consisting of a central server, referred to as the Model Owner (MO), and a client set $\mathcal{C} = c_1, c_2, \dots, c_N$ with N clients. The MO maintains the global model parameter w , initiates the federated learning process, and coordinates collaborative model training. Each client c_k possesses a private local dataset $\mathcal{D}_k = \{(x_k^i, y_k^i) \in \mathcal{X} \times \mathcal{Y} \mid 1 \leq i \leq n_k\}$, where x_k^i and y_k^i denote the input sample and its corresponding label, respectively. Due to privacy constraints, local datasets are never shared among clients. We assume that the local data are non-IID, and the dataset size of client c_k is $|\mathcal{D}_k| = n_k$. The objective of FL is to learn the optimal global model parameter w^* by minimizing the global loss function $\mathcal{L}$, as defined in Eq. (1).

$$\min_{w \in \mathbb{R}^d} \mathcal{L}(w) = \min_{w \in \mathbb{R}^d} \left\{ \sum_{k=1}^N p_k \mathcal{L}_k(w) \right\}, \quad (1)$$

where d is the number of model parameters of the neural network, nonnegative p_k denotes the contribution of client c_k to the global loss function, and $\mathcal{L}_k$ is the local loss function of client c_k . In FedAvg [2], $p_k = \frac{|\mathcal{D}_k|}{|\mathcal{D}|}$, where $|\mathcal{D}| = \sum_{k=1}^N |\mathcal{D}_k|$ is the total number of data samples across the client set $\mathcal{C}$.

At the beginning of each communication round r , the MO broadcasts the global model parameter w_r to all selected clients. Each client c_k then performs local training in parallel using its private dataset and the received global model via Stochastic Gradient Descent (SGD) [7].

$$w_{r+1}^k = w_r^k - \eta \nabla \mathcal{L}_k(w_r^k), \quad (2)$$

where η denotes the local learning rate of client c_k . Based on Eq. (2), the model update of client c_k at communication round r is defined as $\Delta w_r^k = w_{r+1}^k - w_r^k$. After completing local training, the MO aggregates the




Fig. 1. Illustration of the gradient divergence issue.

updated local models by weighted averaging to obtain the new global model.

A. Gradient Divergence

A fundamental challenge in federated learning is understanding how data heterogeneity affects model convergence. To address this issue, we theoretically analyze the impact of data heterogeneity using gradient divergence as a quantitative metric. The required notations are introduced as follows.

Let $w_{r,e}^{k,j}$ denote the local model weight of client c_k in coalition s_j at the e th local epoch of round r , where $e \in [0, E]$ and $r \in [1, T]$ for a fixed local number of epochs E and total number of rounds T . Let the client index set of coalition s_j be defined as $I_j = \{i \mid c_i \text{ is a member of } s_j, \forall c_i \in \mathcal{C}\}$. We denote the local loss function of client c_k in coalition s_j by $\mathcal{L}_{k,j}(\cdot)$. The loss function $\mathcal{L}_{(j)}$ of coalition¹

$$\mathcal{L}_{(j)} = \sum_{k \in I_j} p_k^j \mathcal{L}_{k,j} \quad (3)$$

where the superscript j is the coalition index, $\sum_{k \in I_j} p_k^j = 1$, and $p_k^j \geq 0$ for all $k \in I_j$.

Let $v_{r,e}^j$ denote the centralized model parameter of coalition s_j at the e th local epoch of communication round r [8]. This model corresponds to the ideal scenario in which all clients within the coalition jointly train on their combined data, and serves as a reference for quantifying the gap between decentralized and centralized training. Based on this reference, the gradient divergence is defined as follows [8].

Definition 1: The gradient divergence is defined as the difference between the gradient of the loss function of client c_k and that of its coalition s_j , which is given by

$$\delta_{k,j} = \max_w \|\nabla \mathcal{L}_{k,j}(w) - \nabla \mathcal{L}_{(j)}(w)\|. \quad (4)$$

¹ s_j is defined as follows: In this paper, the terms coalitions and clusters are considered equivalent concepts.




Fig. 2. The system architecture of FedCoalition.

The gradient divergence measures the maximum deviation between a client's local gradient and the average gradient of its coalition. A larger value indicates greater statistical heterogeneity within the coalition.

Fig. 1 illustrates the impact of data heterogeneity on federated learning. Under IID data, the local model updates $w_{r,e}^k$ closely follow the centralized trajectory $v_{r,e}$, whereas non-IID data lead to divergent optimization trajectories and slower convergence.

To alleviate this problem, we employ Clustered Federated Learning (CFL), which groups clients with similar data distributions to better align local updates, thereby improving training stability and convergence.

Assume that for any client c_k in s_j , $\mathcal{L}_{k,j}$ is λ -strongly convex, L -smooth, and M -Lipschitz continuous [8, 9]. The objective for mitigating data heterogeneity is defined as follows.

$$\min \delta = \sum_{k \in I_j} \min \delta_{k,j} \quad (5)$$

for each fixed coalition s_j . Eq. (5) shows that the gradient divergence within coalition s_j should be minimized, meaning that local updates are aligned.

B. The Proposed Architecture

The proposed CFL framework is illustrated in Fig. 2. To mitigate the adverse effects of non-IID data, we develop a hierarchical CFL framework with one-shot clustering [6] for distributed model aggregation. In the initial communication round, the model owner (MO) partitions clients into coalitions based on the similarity of their model updates. Within each coalition, a client is selected as the coalition head (CH) to perform intra-coalition model aggregation. Following [10], each CH is equipped with an auxiliary validation dataset for model evaluation. The overall training procedure consists of the following steps:

- **Step 1: Global model distribution**

At communication round r , the MO broadcasts the global model w_r to all CHs. Each CH evaluates w_r on its auxiliary validation dataset and decides

whether to distribute it to its coalition members. If w_r performs worse than the previously accepted model, it is discarded, and coalition members continue local training using the previous model received via D2D communication. This procedure is referred to as model selection [10].

- **Step 2: Local model training**

Upon receiving either w_r or a notification from the CH, the selected clients perform local training on their private datasets using SGD and then transmit their updated model parameters to the corresponding CH for intra-coalition aggregation.

- **Step 3: Intra-cluster aggregation**

Within each coalition, the CH aggregates local updates from selected clients, as described in Algorithm 1, following a FedAvg-like scheme [2]. The aggregated model is then sent to the MO. Due to similar updates within each coalition, FedCoalition achieves faster convergence.

Algorithm 1: Intra-cluster Aggregation

Input: Client set $\mathcal{C}^{(j)}$ of coalition s_j ,
 $I_j \leftarrow \{i | c_i \text{ is the member of } s_j, \forall c_i \in \mathcal{C}' \cap \mathcal{C}^{(j)}\}$,
 $\hat{w}^{(j)} \leftarrow \{\hat{w}^i | i \in I_j, \forall c_i \in \mathcal{C}' \cap \mathcal{C}^{(j)}\}$
Output: updated model parameter $\tilde{w}^{(j)}$
1: $D \leftarrow \sum_{k \in I_j} |\mathcal{D}_k|$
2: $\tilde{w}^{(j)} \leftarrow \sum_{k \in I_j} \frac{|\mathcal{D}_k|}{D} \hat{w}^k$
3: **return** $\tilde{w}^{(j)}$

C. The Proposed Clustering Methods

In this subsection, we present the clustering strategy for addressing client heterogeneity in wireless networks. We first introduce the metrics used to construct the value function [11] in the proposed coalition formation algorithm. The cosine similarity between the model updates of clients c_i and c_j at communication round r is defined as follows:

$$\text{sim}(w_r^i, w_r^j) = \frac{\langle \Delta w_r^i, \Delta w_r^j \rangle}{\|\Delta w_r^i\| \|\Delta w_r^j\|}. \quad (6)$$

The Gaussian similarity [12] between the model parameter updates of two clients c_i and c_j at round r is defined as

$$\text{gsim}(w_r^i, w_r^j) = \exp\left(\frac{-\|\Delta w_r^i - \Delta w_r^j\|^2}{2\sigma^2}\right), \quad (7)$$

where σ controls how quickly similarity decreases with the Euclidean distance between model updates. Inspired by [13], the scaled similarity between clients c_i and c_j at round r is defined as:

$$\text{scsim}(w_r^i, w_r^j) = (1 - \zeta_1) \text{gsim}(w_r^i, w_r^j) + \zeta_1 (1 + \text{sim}(w_r^i, w_r^j)), \quad (8)$$

where ζ_1 represents a weighting factor that balances the contributions of the two similarity metrics. Note that $scsim(w^i, w^j) \in [0, 1 + \zeta_1]$. We denote the normalized scaled similarity between two clients c_i and c_j at round r by $nscsim(w_r^i, w_r^j)$, which is defined as $scsim(w_r^i, w_r^j)/(1 + \zeta_1)$.

Let s be a coalition with $|s|$ clients, and let I_s denote the index set of clients in s . Then, the average normalized scaled similarity of coalition s is defined as follows:

$$avgnscsim(s) = \frac{2}{|s|(|s| - 1)} \sum_{i,j \in I_s, j \neq i} nscsim(w_r^i, w_r^j) \quad (9)$$

Finally, to satisfy individual rationality [13], we define the coalition value of coalition s as

$$v(s) = \begin{cases} \hat{v}(s) & \text{if } |s| \geq 2 \text{ and } \hat{v}(s) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

where $\hat{v}(s) = avgnscsim(s) - \frac{\zeta_2}{|s|} \sum_{i \in I_s} \mathcal{L}'(w_r^i)$, $\mathcal{L}'$ denote the validation loss of the model parameters, and ζ_2 is a scaling factor. Eqn. (10) indicates that a coalition with a higher average scaled similarity and better model quality among its clients will yield a larger coalition value.

To enable clients to decide whether to join or leave a coalition in order to maximize the value function, we introduce the concept of preference relations, which capture a client's relative evaluation of different coalitions. For any two coalitions s and s' and a client c_k , the preference relation is defined as follows:

$$s \succ_k s' \text{ iff } v(s \cup \{c_k\}) + v(s' \setminus \{c_k\}) > v(s) + v(s'), \quad (11)$$

which means that c_k prefers joining coalition s over remaining in coalition s' . In the following, we define the switch and swap operations [11] based on the above preference relation as follows:

Definition 3: Let $\mathcal{S}$ be a partition and $s, s' \in \mathcal{S}$. A switch operation occurs if client c_k leaves its current coalition s' and joins coalition s , i.e., $s \succ_k s'$. In this case, client c_k switches from s' to s .

Definition 4: Let $\mathcal{S}$ be a partition and $s, s' \in \mathcal{S}$. A swap operation occurs if client c_i leaves its current coalition s' and joins s , while client c_j leaves its current coalition s and joins s' for some $i \neq j$, i.e.,

$$v(s \cup \{c_i\} \setminus \{c_j\}) + v(s' \cup \{c_j\} \setminus \{c_i\}) > v(s) + v(s'). \quad (12)$$

The coalition formation process is based on iterative switch and swap operations. Clients dynamically adjust their coalition memberships by repeatedly performing these operations until a Nash-stable partition [11] is achieved.

TABLE I
PERFORMANCE COMPARISON WITH EXISTING METHODS UNDER NON-IID DATA DISTRIBUTIONS.

Metric	FedAvg [2]	FedCFL [8]	IFCA [15]	Ours
MNIST Acc. (%)	91.8	95.0	93.1	97.3
CIFAR-10 Acc. (%)	59.6	63.2	66.1	70.5
MNIST Loss	0.28	0.23	0.17	0.09
CIFAR-10 Loss	1.17	1.06	0.99	0.90

III. EXPERIMENTAL RESULTS

This section presents the experimental evaluation of the proposed FedCoalition framework and compares its performance with the baseline and existing approaches. We consider $N = 50$ clients with a sampling rate of 0.2. Experiments are conducted on the MNIST and CIFAR-10 datasets using a two-layer convolutional neural network (CNN) and MobileNetV3-Large [14], respectively. The total number of communication rounds is $T = 100$. In each round, every selected client performs $E = 2$ local training epochs with a learning rate of $\eta = 0.01$. The client data are partitioned according to a Dirichlet distribution with concentration parameter $\alpha = 0.6$.

We compare our proposed CFL framework with existing works:

- *FedCFL* [8]: the MO clusters clients by the cosine similarities of model weight updates using the K-means algorithm.
- *IFCA* [15]: each client trains on all k models with its local dataset and then selects the model with the lowest local loss, sending it back to the MO.

Table I compares the testing accuracy of different methods under non-IID data distributions on the MNIST and CIFAR-10 datasets. On the MNIST dataset, the proposed FedCoalition framework achieves the highest accuracy of 97.3%, outperforming FedCFL, IFCA, and FedAvg by 2.3%, 4.2%, and 5.5%, respectively. On the other hand, on the CIFAR-10 dataset, FedCoalition attains an accuracy of 70.5%, exceeding the corresponding baselines by 7.3%, 4.4%, and 10.9%, respectively. These results demonstrate the effectiveness of the proposed framework under heterogeneous data distributions.

Fig. 3 compares the testing accuracy and loss of FedCoalition with the baseline methods. FedCoalition consistently achieves higher accuracy and lower loss throughout most communication rounds, demonstrating faster and more stable convergence under heterogeneous data distributions.

IV. CONCLUSION

In this paper, we proposed a coalition formation game-based framework for federated learning to address data heterogeneity in wireless networks. By grouping clients with similar model updates into coalitions and enabling




Fig. 3. Performance comparison of FedCoalition with existing methods.

coalition heads (CHs) to perform model selection using auxiliary validation data, the proposed framework improves model aggregation quality and accelerates training convergence. Experimental results demonstrate that FedCoalition consistently achieves faster convergence and higher accuracy than existing clustered federated learning (CFL) methods.

REFERENCES

- [1] Omar El-Rifai, Michael Ben Ali, Imen Megdiche, André Peninou, and Olivier Teste, “A survey on cluster-based federated learning,” *arXiv preprint arXiv:2501.17512*, 2025.
- [2] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [3] Wenqi Fan, Ning Wei, Rongyan Xi, Ahmad Bazzi, Yue Xiu, Chadi Assi, Jing Dong, and Jing Jin, “Heterogeneous mixture-of-experts for energy-efficient multimodal isac in highly mobile networks,” *arXiv preprint arXiv:2604.06697*, 2026.
- [4] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Cavin, and Vikas Chandra, “Federated learning with non-iid data,” *arXiv preprint arXiv:1806.00582*, 2018.
- [5] Bo Xu, Wenchao Xia, Wanli Wen, Pei Liu, Haitao Zhao, and Hongbo Zhu, “Adaptive hierarchical federated learning over wireless networks,” *IEEE Transactions on Vehicular Technology*, vol. 71, no. 2, pp. 2070–2083, 2021.
- [6] Wenji He, Haipeng Yao, Tianle Mai, Fu Wang, and Mohsen Guizani, “Three-stage stackelberg game enabled clustered federated learning in heterogeneous uav swarms,” *IEEE Transactions on Vehicular Technology*, vol. 72, no. 7, pp. 9366–9380, 2023.
- [7] Léon Bottou, “Large-scale machine learning with stochastic gradient descent,” in *Proceedings of COMPSTAT’2010: 19th International Conference on Computational Statistics Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers*. Springer, 2010, pp. 177–186.
- [8] Moming Duan, Duo Liu, Xinyuan Ji, Yu Wu, Liang Liang, Xianzhang Chen, Yajuan Tan, and Ao Ren, “Flexible clustered federated learning for client-level data distribution shift,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 11, pp. 2661–2674, 2021.
- [9] Jianchun Liu, Hongli Xu, Lun Wang, Yang Xu, Chen Qian, Jinyang Huang, and He Huang, “Adaptive asynchronous federated learning in resource-constrained edge computing,” *IEEE Transactions on Mobile Computing*, vol. 22, no. 2, pp. 674–690, 2021.
- [10] Pu Tian, Weixian Liao, Wei Yu, and Erik Blasch, “Wsc: A weight-similarity-based client clustering approach for non-iid federated learning,” *IEEE Internet of Things Journal*, vol. 9, no. 20, pp. 20243–20256, 2022.
- [11] Liantao Wu, Peng Sun, Honglong Chen, Yong Zuo, Yong Zhou, and Yang Yang, “Noma-enabled multiuser offloading in multicell edge computing networks: A coalition game based approach,” *IEEE Transactions on Network Science and Engineering*, vol. 11, no. 2, pp. 2170–2181, 2024.
- [12] Andrew Ng, Michael Jordan, and Yair Weiss, “On spectral clustering: Analysis and an algorithm,” *Advances in neural information processing systems*, vol. 14, 2001.
- [13] Jianfeng Lu, Hangjian Zhang, Pan Zhou, Xiong Wang, Chen Wang, and Dapeng Oliver Wu, “Fed-law: Value-aware federated learning with individual fairness and coalition stability,” *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 9, no. 1, pp. 1049–1062, 2025.
- [14] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al., “Searching for mobilenetv3,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 1314–1324.
- [15] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran, “An efficient framework for clustered federated learning,” *Advances in neural information processing systems*, vol. 33, pp. 19586–19597, 2020.