

Data Similarity-Based One-Shot Clustering for Multi-Task Hierarchical Federated Learning

Abdulmoneam Ali Ahmed Arafa
Department of Electrical and Computer Engineering
University of North Carolina at Charlotte, NC 28223
aali28@charlotte.edu *aarafa@charlotte.edu*

Abstract—We address the problem of cluster identity estimation in a hierarchical federated learning setting in which users work toward learning different tasks. To overcome the challenge of task heterogeneity, users need to be grouped in a way such that users with the same task are in the same group, conducting training together, while sharing the weights of feature extraction layers with the other groups. Toward that end, we propose a *one-shot* clustering algorithm that can effectively identify and group users based on their *data similarity*. This enables more efficient collaboration and sharing of a common layer representation within the federated learning system. Our proposed algorithm not only enhances the clustering process, but also overcomes challenges related to privacy concerns, communication overhead, and the need for prior knowledge about learning models or loss function behaviors. We validate our proposed algorithm using various datasets such as CIFAR-10 and Fashion MNIST, and show that it outperforms the baseline in terms of accuracy and variance reduction.

I. INTRODUCTION

Multi-Task Learning (MTL) and Personalized Federated Learning (PFL) have gained significant attention in recent years due to their potential in accelerating learning and improving generalization [1], [2]. In PFL, instead of learning a single global model, users seek to learn a personalized model while still leveraging the data and learning progress of other users, even if their specific model structures or parameters might differ. Likewise, MTL facilitates the simultaneous learning of multiple tasks using different models by leveraging the shared knowledge and information within the models' common representation.

The inherent characteristics of multi-task learning (MTL) meet the requirements of Hierarchical Federated Learning (HFL). The core idea of HFL is to alleviate the communication bottleneck with the parameter server [3], [4]. From a communication perspective, MTL offers the advantage of reducing the number of weights that need to be shared between users. By integrating MTL with HFL, groups of users can be associated with local parameter servers (LPSs), enabling them to jointly train their models within these LPSs. They can then communicate with the Global Parameter Server (GPS) to access the updated weights of common layers from other LPSs. This approach is expected to significantly alleviate the communication bottleneck encountered in Federated Learning

(FL). The main challenge is how to efficiently design this multi-task hierarchical federated learning (MT-HFL) setting in a way that does not result in divergence or slow learning.

Related Work. In [5], the authors propose an alternating minimization algorithm that iterates between two phases: minimizing the loss function and identifying the cluster identities. They show that their algorithm converges, given a good initialization, with a strongly convex and smooth loss function. In [6], the authors assume the existence of a matrix that captures the relationships between tasks. This matrix, either known beforehand or estimated during the training, is incorporated into the objective function along with a carefully chosen regularization term that depends on this task relation matrix. The training process then proceeds with this augmented objective function. Reference [7] evaluates the cosine similarity between the weights of each neural network layer among different pairs of users and considers this value as a metric to cluster users. The work in [8] focuses mainly on a multi-task linear regression model, and tackles feature heterogeneity using an augmented dataset-based approach. In [9], the authors assume a pre-clustered task structure, where each task is associated with a distinct cloud parameter server. Their objective is to determine the optimal number of local iterations per cloud server and global iterations, ensuring convergence within given resource budget and usage cost constraints.

Contributions. Based on the literature, clustering approaches for MT-HFL can be broadly categorized into two types. The first type assumes that the model parameters for the same task are similar between users. It leverages this assumption by clustering users based on their model parameters, typically through iterative clustering algorithms that revisit the clustering task throughout the training process, as early-iteration weights may not be sufficiently informative for accurate cluster identification. The second type involves designing a specific auxiliary loss function that addresses both feature heterogeneity and task-relation similarity. Therefore the question we address in this work is as follows:

Can we leverage feature heterogeneity among users to get a one-shot clustering algorithm for MT-HFL?

The essence of this work is to turn the feature heterogeneity among users from a challenge into an *opportunity*. Our main objective is to efficiently cluster users among LPSs while preserving their data privacy and minimizing communication

costs, and to do so *independently* of the model or the class of loss function at the users. Specifically, enabling each LPS to conduct training for a different task, determined by its associated users, and to collaborate with the other LPSs by sharing the common representation layers through the GPS.

To address the proposed question, we employ a variant of the data valuation method proposed in [10]. Considering data valuation in clustering users has three advantages. Firstly, it provides one-shot clustering, with the decision being taken *before* training begins. Secondly, it provides an *optimum* user clustering solution that can be achieved in a distributed manner among users without intervention from the LPS or GPS. By optimum, we mean that our proposed algorithm successfully generates clusters that align with each user’s specific task preferences, *as if it had prior knowledge of their individual tasks*. Thirdly, it is important to note that our approach does *not* rely on any prior knowledge regarding the model architecture, loss function, or the need for dataset augmentation. In summary, we mainly solve the concerns about one-shot clustering algorithms detailed in [5].

Our proposed algorithm can be summarized in two main steps: the first step estimates the feature similarity among users, as elaborated in Section II-B. Subsequently, we associate similar users with a specific LPS, as explained in Section II-C. Finally, we conduct a standard federated averaging training within each LPS, then share only subsets of the feature weights for aggregation with the GPS. To evaluate the effectiveness of our proposed clustering algorithm, we compare it with a random clustering approach that disregards the information regarding data similarity between users.

Notation. We use lowercase letters for scalars, bold lowercase letters for vectors, and bold uppercase letters for matrices throughout this paper.

II. DATA SIMILARITY MT-HFL CLUSTERING

A. System Model

We consider an MT-HFL system with one GPS and a set of LPSs. The number of LPSs is equivalent to the number of tasks, T , that users in the system are aiming to learn. The set of users is defined as $\mathcal{N} = \{1, 2, \dots, N\}$. Each user has its own data set $\{(\mathbf{x}_l, y_l)\}_{l=1}^{n_i}$, with $\mathbf{x}_l \in \mathbb{R}^m$ and $y_l \in \{1, \dots, C\}$, where C is the number of classes and n_i is the number of samples available at user i .

The main goal is to cluster users into T clusters, without violating their data privacy, such that each cluster only contains users that seek to learn the same task. Users within the same cluster would then communicate with a dedicated LPS towards learning a model for their specific task. The LPSs, in turn, communicate with the GPS to learn a common layer representation to be shared among all clusters.

We assume that users who are looking to learn a specific task have most of their training samples drawn from that particular task in the training dataset. In addition, we assume that all users are honest and trustworthy.

B. Data Similarity Estimation

In order to cluster users into groups that have samples from the same task training dataset, we propose a variant of the data valuation approach in [10] amenable to the FL setting. In particular, let us consider that user i has its own raw data arranged in a matrix $\mathbf{X}_i \in \mathbb{R}^{n_i \times m}$. Our objective is to group user i with user j that has a data matrix \mathbf{X}_j drawn from the same distribution as \mathbf{X}_i (and therefore would be interested in learning the same task). The authors in [10] propose a metric called data relevance that measures the similarity between the data of a seller and the data of a buyer, assisted by a broker. However, the assumptions made in their approach do not meet the FL data privacy requirements. We therefore propose a modified version of this metric by focusing on how to evaluate it in a distributed manner, without sharing users’ data, without requiring any joint training among users, and with only at most *one* iteration with the GPS.

Specifically, each user i is required to compute an eigen decomposition of the Gram matrix of their data features. The Gram matrix is obtained by weighting the features with the number of samples of user i as follows:

$$\frac{1}{n_i} \Phi(\mathbf{X}_i)^T \Phi(\mathbf{X}_i), \quad (1)$$

where $\Phi(\cdot)$ is a feature mapping function belonging to $\mathbb{R}^{n_i \times d}$ with $d < m$. The resulting eigen decomposition produces a set of eigenvalues arranged in a vector $\boldsymbol{\lambda}_i \in \mathbb{R}^d$ with their associated eigenvectors stacked in a matrix $\mathbf{V}_i \in \mathbb{R}^{d \times d}$. Then, the users share their eigenvector matrix, \mathbf{V}_i , among themselves. Based on the matrix received, user i can now estimate how much their data, \mathbf{X}_i , vary in the direction of each column of \mathbf{V}_j from user j . This is achieved by projecting the eigenvectors of other users onto their data and evaluating the Euclidean norm as follows:

$$\hat{\lambda}_k^{(j)} = \left\| \frac{1}{n_i} \Phi(\mathbf{X}_i)^T \Phi(\mathbf{X}_i) \mathbf{v}_k^{(j)} \right\|, \quad \forall k \in [d], \quad (2)$$

where $\mathbf{v}_k^{(j)}$ is the k th eigenvector of user j .

Reaching this step, one can expect that if two users, i and j , have relatively similar features, i.e., their feature matrices lie close to each other in the feature space, then the values of their eigenvalues, both the original ($\boldsymbol{\lambda}_i$) and the projected ($\hat{\boldsymbol{\lambda}}_j$), should be relatively close as well. Therefore, each user i can now compute the relevance [10] of their data to user j based on $\boldsymbol{\lambda}_i$ and $\hat{\boldsymbol{\lambda}}_j$ by evaluating:

$$\lambda_k^{(i,j)} = \frac{\min\{\lambda_k^{(i)}, \hat{\lambda}_k^{(j)}\}}{\max\{\lambda_k^{(i)}, \hat{\lambda}_k^{(j)}\}}, \quad \forall k \in [d], \quad (3)$$

$$r(i, j) = \prod_{k=1}^d (\lambda_k^{(i,j)})^{\frac{1}{d}}, \quad (4)$$

where the $\max\{.,.\}$ in (3) is to normalize the value of $r(i, j)$.

After that, the GPS will receive from each user i its data relevance estimates with respect to every other user $j \in \mathcal{N}$, as shown in (4), and in turn estimates the average data similarity




Fig. 1: Illustration of our proposed algorithm.

between each pair of users as follows:

$$\mathbf{R}(i, j) = \frac{r(i, j) + r(j, i)}{2}, \forall i, j \in \mathcal{N}. \quad (5)$$

Note that \mathbf{R} is a symmetric $N \times N$ matrix.

C. Clustering Decision

Since \mathbf{R} captures the pairwise data similarity/relevance between every user pair, the GPS adopts the Hierarchical Agglomerative Clustering algorithm (HAC) [11] to cluster users among the LPSs.

To better illustrate our idea, we show an example of matrix \mathbf{R} in Table I below. In this example, the CIFAR-10 dataset is distributed among five users such that users 1 and 2 have samples from class labels: {plane, car, ship, truck}, while users 3, 4, and 5 have samples from class labels: {bird, cat, deer, dog, frog, horse}.

Users	User 1	User 2	User 3	User 4	User 5
User 1	1	0.97	0.31	0.31	0.32
User 2	0.97	1	0.31	0.32	0.32
User 3	0.31	0.31	1	0.97	0.98
User 4	0.31	0.32	0.97	1	0.98
User 5	0.32	0.32	0.98	0.98	1

TABLE I: Example of the data similarity matrix \mathbf{R} on the CIFAR-10 dataset with two different tasks.

In this example, the HAC algorithm initially assigns each user to a single cluster, resulting in five clusters. The algorithm proceeds to merge each close pair of users. For example, users 1 and 2 are very close, so they merge together. The same holds for users 4 and 5. Next, the algorithm merges user 3 into the cluster formed by users 4 and 5. The output of the algorithm is a binary tree known as a dendrogram. By varying the cutting height of this tree, different numbers of clusters can be obtained. The number of clusters needed in this example is 2, since there are two tasks to be learned. In general, the HAC algorithm is adjusted to produce T clusters.

D. MT-HFL Training Procedure

Once the assigned users are determined for each LPS, users can start the training process. Specifically, each LPS will individually execute a standard federated averaging algorithm (FedAvg) [12] for a specified number of local iterations. Subsequently, they will share the weights of the common layers with the GPS. During each global communication round, the GPS aggregates the shared weights from the LPSs and broadcasts them back to initiate the next local iteration. These weighted average weights, combined with the LPSs' last weights for the remaining layers, serve as the starting point for the subsequent iteration. The training process continues until the predefined number of global iterations, G , is reached. The steps outlined above are summarized in Algorithms 1 and 2 below.

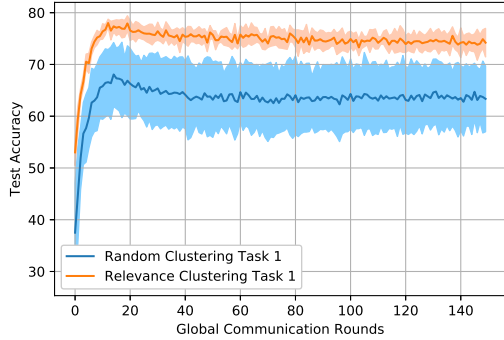
Pictorially, our proposed approach is depicted in Fig. 1

Algorithm 1 Main Algorithm

- 1: **Input:** number of tasks T ; number of global iterations G
- 2: Execute Algorithm 2: Data Similarity Clustering
- 3: **for** $r \in [G]$ **do**
- 4: **for** $i \in [T]$ **do**
- 5: Conduct FedAvg at each LPS
- 6: Share the weights of first common layers with GPS
- 7: GPS aggregates the shared weights and broadcasts them back to each LPS
- 8: **end for**
- 9: **end for**

Algorithm 2 Data Similarity Clustering

- 1: Perform an eigenvalue decomposition:
- 2: **for** $i \in \mathcal{N}$ **do**
- 3: $\lambda_i, V_i = \text{eigen}(\frac{1}{n_i} \Phi(\mathbf{X}_i)^T \Phi(\mathbf{X}_i))$
- 4: Share V_i with the other users
- 5: **end for**
- 6: Users compute the estimated the eigen values:
- 7: **for** $i \in \mathcal{N}$ **do**
- 8: **for** $j \in \mathcal{N}$ **do**
- 9: $\hat{\lambda}_k = \|\frac{1}{n_i} \Phi(\mathbf{X}_i)^T \Phi(\mathbf{X}_j) \mathbf{v}_k^{(j)}\|, \forall k \in [d]$
- 10: Compute $\lambda_k^{(i,j)} = \frac{\min\{\lambda_k^{(i)}, \hat{\lambda}_k\}}{\max\{\lambda_k^{(i)}, \hat{\lambda}_k\}}, \forall k \in [d]$
- 11: Apply (4) and share with the GPS
- 12: **end for**
- 13: **end for**
- 14: GPS uses (5) to compute the average similarity matrix \mathbf{R} :
- 15: **for** $i \in \mathcal{N}$ **do**
- 16: $\mathbf{R}(i, j) = \frac{r(i,j)+r(j,i)}{2}$ and $\mathbf{R}(j, i) \leftarrow \mathbf{R}(i, j)$
- 17: **end for**
- 18: Feed \mathbf{R} to the HAC algorithm as input.
- 19: Based on the desired number of clusters, T , and the output of the HAC algorithm, the memberships of the users to each cluster is obtained



(a) CIFAR-10 Task 1 LPS.



(b) CIFAR-10 Task 2 LPS.

Fig. 2: Performance on the CIFAR-10 dataset with two tasks/clusters.

III. EXPERIMENTS

We now present simulation results for the proposed data similarity-based clustering algorithm to demonstrate its superiority in achieving higher accuracy compared to the baseline approach. Furthermore, our approach exhibits a lower variance in the results. The subtle distinction between our proposed algorithm and the baseline is that the baseline ignores the similarity in the data between users and randomly clusters them.

Datasets and Models. We consider an image classification supervised learning task on the CIFAR-10 dataset [13] and the Fashion MNIST dataset [14]. For CIFAR-10, a convolution neural network (CNN) is adopted with two 5x5 convolution layers, two 2x2 max pooling layers, two fully connected layers with 120 and 84 units, respectively, ReLU activation, a final softmax output layer, and cross-entropy loss. In the case of Fashion MNIST, we adopt a multi-layer perceptron (MLP) architecture. The MLP comprises two fully connected layers: the first layer has dimensions (784, 32) and utilizes ReLU activation, followed by the second layer with dimensions (32, 10) and a log softmax activation. The loss function is the negative log likelihood loss. We conduct six experimental runs and calculate the average performance across these experiments.

Beating the Baseline. In Fig. 2, we demonstrate the superiority of our proposed algorithm over random clustering using CIFAR-10. We define two tasks: the first task involves learning {plane, car, ship, truck} labels, while the second task is to learn labels for {bird, cat, deer, dog, frog, horse}. For each task, we distribute the corresponding labels among five users. Additionally, to introduce some level of label diversity and to test the robustness of our proposed algorithm, we assign 10% labels from the other task to each user. The two LPSs are only sharing the weights of the two convolution layers with the GPS (common layers).

To further validate the effectiveness of our proposed clustering algorithm, we conduct simulations involving three distinct tasks and an unbalanced distribution of task labels. In Fig. 3, we show the superiority of the proposed algorithm using the Fashion MNIST dataset [14]. We divide the dataset into three tasks: the first task focuses on learning clothes-related labels

(Task 1), the second task is dedicated to learning shoe-related labels (Task 2), and the third task is to learn bag-related labels (Task 3). We have ten users: five of them have the majority of Task 1 labels, three users have the majority of Task 2 labels, and two users have Task 3 labels. Each user also has a minority of labels from other tasks.

In particular, the number of samples for each task is unbalanced. Specifically, Task 1 has a significantly larger number of samples compared to Task 2, and Task 3 has the smallest number of samples. In addition to that, only two users have the majority of Task 3 samples. This is why the performance of Task 3 is relatively lower in the random clustering baseline due to the lower possibility of choosing the two specific users related to this task out of the total of ten users. The significance of our proposed algorithm becomes evident in addressing such an unbalanced distribution of data. Unlike CIFAR-10, here we do not need a feature mapping Φ , since the raw data itself is informative and has a lower dimension ($m = 784$). For CIFAR-10, the raw data has a high dimension ($m = 3072$), and are not inherently informative. Therefore, following the approach in the literature [10], [8], [7], we utilize the convolution layers' weights of an ImageNet-pre-trained ResNet18 to evaluate the feature vector for each data point at users. *However, when users start training for learning, they start from random weights.*

Dataset/Distribution Robustness. Furthermore, our algorithm performs well when dealing with two distinct distributions. In particular, we investigated the similarity between users' labels using the CIFAR-10 and CIFAR-100 datasets in a three-user setting. User 1 has labels related to vehicles from CIFAR-10, while users 2 and 3 have CIFAR-100 data, with user 2 having labels representing vehicles, and user 3 having labels from the remaining classes of CIFAR-100.

The results, as shown in Table II, show the efficacy of our algorithm in successfully matching users with similar labels, even when such labels originate from different datasets. This showcases the algorithm's ability to handle diverse label distributions and facilitate effective collaboration among users.

Communication Improvement. According to equation (3), we observe that we need to discard the eigenvalues that are




Fig. 3: Performance on the Fashion MNIST with three tasks/clusters.

Users	User 2	User 3
User 1	0.62	0.39

TABLE II: Similarity between users having different datasets: CIFAR-10 (User 1) and CIFAR-100 (Users 2 and 3).




Fig. 4: Impact of the number of shared eigenvectors on the relevance value.

extremely small, as the multiplication will be highly drifted by such small values, even if there is only one such eigenvalue. To illustrate this, consider an example where $\hat{\lambda}_{100} = 0.000001$ and $\lambda_{100} = 1$. This observation, along with our goal to minimize the needed communication, motivates us to explore the optimal number of eigenvectors necessary to differentiate between users during the clustering step. Using the same setting as in the Fashion MNIST experiment, we select the first fifty eigenvectors corresponding to the fifty largest eigenvalues. As shown in Fig. 4, it reveals that using only five eigenvectors, we can effectively cluster users with different tasks. This implies that instead of exchanging the entire matrix \mathbf{V} with dimensions (784×784) , users can exchange a much smaller matrix with dimensions (5×784) and still accurately compute their relatedness to each other. In Fig. 4, Users 0 and 3 have samples from the same task, which is different from the task associated with User 6 (having samples from Task 2) and user 9 (having samples from Task 3). This finding also emphasizes the power of our clustering algorithm over existing ones in the literature that mainly rely on the weights of the model, which are of much higher dimensions, at least on the order of thousands.

IV. CONCLUSION

A one-shot clustering data similarity-based algorithm is proposed in an MT-HFL setting, investigating the effects of clustering users based on their data on the overall accuracy of FL. Our proposed clustering algorithm enabled us to perfectly cluster users without violating their privacy. Furthermore, we achieved this with low-dimensional information exchange and independently of the learning model or the loss function class. For future investigations, we plan on studying the robustness of the proposed algorithm when the exchanged eigenvectors are noisy, in addition to exploring the possibility of adding an extra layer of privacy to our algorithm without compromising its clustering ability.

REFERENCES

- [1] Y. Zhang and Q. Yang. A survey on multi-task learning. *IEEE Trans. Knowl. Data Eng.*, 34(12):5586–5609, December 2022.
- [2] O. Marfoq, G. Neglia, A. Bellet, L. Kameni, and R. Vidal. Federated multi-task learning under a mixture of distributions. In *Proc. NeurIPS*, December 2021.
- [3] L. Liu, J. Zhang, S. H. Song, and K. B. Letaief. Client-edge-cloud hierarchical federated learning. In *Proc. IEEE ICC*, June 2020.
- [4] A. Ali and A. Arafa. Delay sensitive hierarchical federated learning with stochastic local updates. arXiv:2302.04851.
- [5] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran. An efficient framework for clustered federated learning. *IEEE Trans. Inf. Theory*, 68(12):8076–8091, December 2022.
- [6] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. In *Proc. NeurIPS*, 2017.
- [7] R. Cai, X. Chen, S. Liu, J. Srinivasa, M. Lee, R. Kompella, and Z. Wang. Many-task federated learning: A new problem setting and a simple baseline. In *Proc. IEEE/CVF CVPRW*, June 2023.
- [8] R. Liu, J. Yang, and C. Shen. Exploiting feature heterogeneity for improved generalization in federated multi-task learning. In *Proc. ISIT*, June 2023.
- [9] X. Yi, R. Li, C. Peng, F. Wang, J. Wu, and Z. Zhao. RHFedMTL: Resource-aware hierarchical federated multi-task learning. arXiv:2306.00675.
- [10] M. M. Amiri, F. Berdoz, and R. Raskar. Fundamentals of task-agnostic data valuation. In *Proc. AAAI*, June 2023.
- [11] Kevin P. Murphy. *Probabilistic Machine Learning: An introduction*. MIT Press, 2022.
- [12] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proc. AISTATS*, April 2017.
- [13] A. Krizhevsky. Learning multiple layers of features from tiny images. 2009, <https://www.cs.toronto.edu/~kriz/cifar.html>.
- [14] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. arXiv:1708.07747.