



HAL
open science

Comparative Evaluation of Clustered Federated Learning Methods

Michael Ben Ali, Omar El-Rifai, Imen Megdiche, André Peninou, Olivier Teste

► **To cite this version:**

Michael Ben Ali, Omar El-Rifai, Imen Megdiche, André Peninou, Olivier Teste. Comparative Evaluation of Clustered Federated Learning Methods. The 2nd IEEE International Conference on Federated Learning Technologies and Applications (FLTA24), Sep 2024, Valencia (Espagne), Spain. hal-04741739v2

HAL Id: hal-04741739

<https://hal.science/hal-04741739v2>

Submitted on 31 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparative Evaluation of Clustered Federated Learning Methods

Michael Ben Ali
IRIT, CNRS
Université Toulouse, UT3
Toulouse, France
Michael-Eddy.Ben-Ali@irit.fr

Omar El-Rifai
IRIT, CNRS
Université Toulouse, UT3
Toulouse, France
Omar.El-Rifai@irit.fr

Imen Megdiche
IRIT, CNRS
INU Champollion, ISIS Castres
Castres, France
imen.megdiche@irit.fr

André Peninou
IRIT, CNRS
Université Toulouse, UT2J
Toulouse, France
andre.peninou@irit.fr

Olivier Teste
IRIT, CNRS
Université Toulouse, UT2J
Toulouse, France
olivier.teste@irit.fr

Abstract—Over recent years, Federated Learning (FL) has proven to be one of the most promising methods of distributed learning which preserves data privacy. As the method evolved and was confronted to various real-world scenarios, new challenges have emerged. One such challenge is the presence of highly heterogeneous (often referred as non-IID) data distributions among participants of the FL protocol. A popular solution to this hurdle is Clustered Federated Learning (CFL), which aims to partition clients into groups where the distribution are homogeneous.

In the literature, state-of-the-art CFL algorithms are often tested using a few cases of data heterogeneities, without systematically justifying the choices. Further, the taxonomy used for differentiating the different heterogeneity scenarios is not always straightforward.

In this paper, we explore the performance of two state-of-the-art CFL algorithms with respect to a proposed taxonomy of data heterogeneities in federated learning (FL). We work with three image classification datasets and analyze the resulting clusters against the heterogeneity classes using extrinsic clustering metrics. Our objective is to provide a clearer understanding of the relationship between CFL performances and data heterogeneity scenarios.

Index Terms—Federated Learning, Clustering, Data Heterogeneity

I. INTRODUCTION

Federated Learning (FL) is a distributed machine learning paradigm designed to prioritize data privacy. Originally designed [17] to train a global machine learning model across multiple clients without sharing their data, only locally trained model weights and the number of samples are shared to a server that iteratively aggregates and updates the clients' models.

One of the main challenges for the convergence of FL algorithms is the presence of non-IID (non-independent and

non-identically distributed) data distribution across clients. The difficulty [24] of this scenario is illustrated with Figure 1, where a comparison between the direction of parameters of a hypothetical centralized model trained on all clients' data and the parameters of the federated model is proposed in IID and non-IID scenarios. In the case of non-IID data, the gap between the global federated model's weights ω and the centralized model's weights $\bar{\omega}$ is far larger than in the case of homogeneous distributions.


The heterogeneity of client distributions has been the subject of a multitude of studies [11], [24]. In recent research, a popular solution proposed to mitigate the problem of heterogeneous client data distribution is Clustered Federated Learning (CFL) [2], [3], [6], [7], [9]. The idea behind CFL is to relax the notion of a single global model and propose multiple personalized models for particular distributions. CFL proposes a paradigm that will group clients into clusters where ideally, client data distributions are supposed to be homogeneous and where the standard FL protocol will be applied within each cluster. It is particularly useful when we already have insight of clients' data with similar patterns.

Since a well recognized taxonomy of data heterogeneity of client datasets for FL was proposed [11] and even if CFL has been the subject of multiple research papers, the rigorous exploration of CFL algorithms regarding each cases of heterogeneous data distribution, aiming to be comprehensive and as exhaustive as possible in covering heterogeneity cases, to our knowledge, has not been addressed yet in the literature.

Our contribution in this paper is to rigorously study the effect of each cases of data heterogeneity on state-of-the-art CFL solutions, striving to be as exclusive as possible in each type without overlapping with other types. We propose to explore each heterogeneity types independently using three images classification datasets, MNIST, Fashion-MNIST and KMNIST.

The paper is organized as follows. Section II provides a

This work was carried out as part of the Data Mobility action of the Vilagil project, a project co-financed by "Toulouse Métropole" and by "France 2030" under "Territoires d'innovation" program operated by "La Banque des Territoires".



(a) All clients i , $1 \leq i \leq N$, train their local model's parameters ω_i , on homogeneous data. In this scenario, the federated model parameters ω are close to the hypothetical non-federated model with parameters $\bar{\omega}$ trained on the same data but in a centralized manner

(b) All clients i , $1 \leq i \leq N$, train their local model's parameters ω_i , on heterogeneous data. In this scenario, the federated model parameters ω drift from the non-federated model parameters $\bar{\omega}$ trained on the same data but in a centralized manner

Fig. 1: IID scenario vs Non-IID as illustrated in [24]

description of the non-iid heterogeneous data problem in FL and introduces a data heterogeneities taxonomy. In Section III, we explain in details the functionality of CFL and present two methods to classify state-of-the-art CFL solutions. In Section IV, we study potential shortcomings of state-of-the-art CFL with respect to the taxonomy described in Section II which highlights our contribution in this paper. Finally, we conclude with some remarks and future research directions.

II. NON-IID DATA DISTRIBUTION IN FEDERATED LEARNING

The original state-of-the-art federated learning algorithm and aggregation method named FedAvg [17] was thought to be capable of convergence in IID and as well in non-IID settings. Unfortunately, it was proven in subsequent studies [24] that the non-IID scenario was problematic the more the distribution of clients data shifted from one another (as illustrated in Figure 1). To have a better grasp of the different cases of data heterogeneity explained in that paper, we consider the following setup.

We suppose we have N clients with $i \in I = \{1, \dots, N\}$. For each client, we consider its associate local dataset D_i noted as a samples set of the form

$$D_i = \left\{ (x_j^{(i)}, y_j^{(i)}) \mid 1 \leq j \leq n_i \right\} \quad (1)$$

where $x = (x_j^{(i)})_{1 \leq j \leq n_i}$ represents the features and $y = (y_j^{(i)})_{1 \leq j \leq n_i}$ represents the target and with the number of local samples being $|D_i| = n_i$. In case of a FL supervised learning task with features x and target y , a statistical model operates with two levels of sampling. To access a data point, we first sample a client, denoted as i , from the statistical distribution P , which represents the available clients. Then, we draw an example (x, y) from the local data distribution of that client, denoted as $\mathcal{P}_i(x, y)$. When we talk about non-IID data in federated learning, we're mainly focused on the heterogeneity between the statistical distributions \mathcal{P}_i and \mathcal{P}_j for different clients i and j .

Each clients datasets D_i is represented by its local joint distribution $\mathcal{P}_i(x, y)$ which we can factor as follows using Bayes' rule:

$$\mathcal{P}_i(x, y) = \mathcal{P}_i(y)\mathcal{P}_i(x|y) = \mathcal{P}_i(x)\mathcal{P}_i(y|x) \quad (2)$$

where $\mathcal{P}_i(x)$ the marginal distribution of the features (respectively $\mathcal{P}_i(y)$ the target) and $\mathcal{P}_i(x|y)$ the distribution of features conditioned by the target (respectively $\mathcal{P}_i(y|x)$ the distribution of features conditioned by the target) for client $i \in I$.

A. A taxonomy of data heterogeneities in Federated Learning

Using the notations above, we consider 5 types of data distributions heterogeneity situations [17]. We should note that those situations are in principle not mutually exclusive and two clients may encounter together any numbers of those situations.

For client $i \in I$ and client $j \in I$ with $i \neq j$, we can consider D_i and D_j as having heterogeneous distributions if one or multiple of the following conditions apply:

- 1) **Concept shift on features:** (same label, different features) $\mathcal{P}_i(x|y) \neq \mathcal{P}_j(x|y)$ for clients i and j even if $\mathcal{P}_i(y) = \mathcal{P}_j(y)$. This can be artificially attained in MNIST if we rotate the images in certain devices. (see Figure 2.a.)
- 2) **Concept shift on labels:** (same features, different label) $\mathcal{P}_i(y|x) \neq \mathcal{P}_j(y|x)$ even if $\mathcal{P}_i(x) = \mathcal{P}_j(x)$. Although more likely in dataset reflecting divergence of labels concept or labels annotation error, we can illustrate such a scenario in the MNIST dataset if for instance the digits "1" in a user's dataset resembles digit "7" in another user's or if two digits labels are swapped either by error or for malicious purposes. (see Figure 2.b.).
- 3) **Feature distribution skew:** $\mathcal{P}_i(x) \neq \mathcal{P}_j(x)$ even if $\mathcal{P}_i(y|x) = \mathcal{P}_j(y|x)$. In the MNIST dataset, we can artificially create this heterogeneity using different type of writing with image erosion and dilatation simulating the type of pen used. For example, we can have a dilated image simulate bold writing with a marker pen for client i and eroded image with finer writing simulating a pencil writing for j like in Figure 2.c.
- 4) **Label distribution skew:** $\mathcal{P}_i(y) \neq \mathcal{P}_j(y)$ even if $\mathcal{P}_i(x|y) = \mathcal{P}_j(x|y)$. We can create this scenario by unevenly distributing the digits to the clients such that one client has a majority "1"s, another a majority of "4"s and so on as illustrated in Figure 2.d. It corresponds to the natural cases of imbalanced datasets across clients.
- 5) **Quantity Skew:** the volume of data differs significantly among clients. Meaning that $|D_i| \ll |D_j|$ or $|D_j| \ll |D_i|$. (see Figure 2.e.)

For a clearer understanding, we illustrate in Figure 2 each type of non-iid data distribution corresponding to the taxonomy with a corresponding example from the MNIST dataset

(The numbered items from the categories correspond to the letter items in the Figure in alphabetical order).




Fig. 2: Illustration of non-IID categories for two clients i and j with samples from the MNIST dataset.

B. Related works : Datasets and heterogeneity types in CFL literature

Even though the taxonomy defined in subsection II-A is well recognized across scientific papers, it is often ambiguously used. Most CFL papers do not clearly define which heterogeneity cases they are addressing.

To rectify this, we use the taxonomy as a guideline in Table I, where we provide an overview of the different types of datasets, clearly identifying which types of heterogeneities, as defined by the taxonomy, are explored in CFL studies. This table reveals that none of the reviewed studies have examined the global impact of all data heterogeneity types on CFL algorithms, often focusing on seemingly random or specific cases instead. Among all the papers classified in Table I, only [1]–[3], [7], [10], [18] address two heterogeneity types out of five, while all others consider only one particular case. Moreover, it is important to note that none of these papers tested the impact of quantity skew scenarios. These observations form the foundation of our research.

We also observe that most papers focus on image classification datasets like MNIST, EMNIST, FEMNIST, Fashion-MNIST, FedCelebA and CIFAR-10, with only a few addressing other type of data like text and time series. We believe the main reason behind this is the difficulty in clearly identifying the taxonomy cases for tasks other than image classification, where the cases can be more easily defined (as illustrated by the examples derived from MNIST in subsection II-A).

III. CLUSTERED FEDERATED LEARNING

As stated in previous sections, Clustered Federated Learning is one of the most recently researched method to mitigate the negative effect of data heterogeneity across clients in FL scenario.

In this Section, we define a generic cluster-based federated optimization formulation, and then, present the two type of solution approaches we found in the literature and classify them this way: namely, “Server-side Clustering” in subsection III-A and “Client-side Clustering” in subsection III-B.

Let clients $i \in \{1, \dots, N\}$ having datasets D_i each containing $|D_i| = n_i$ data samples which are supposed to span together K different data distributions $\mathcal{P}_1, \dots, \mathcal{P}_K$ such that

Dataset	Suggested Category of non-IID situation	Papers
MNIST	Concept shift on labels	[2], [3], [7], [20]
	Concept shift on features	[3], [5], [9], [10], [18], [21]
	Label distribution skew	[2], [6], [10], [13], [22]
EMNIST	Concept shift on features	[1], [18], [19]
	Label distribution skew	[1]
FEMNIST	Feature distribution skew	[6], [9], [15], [18]
	Concept shift on labels	[7]
	Label distribution skew	[21]
Fashion MNIST	Concept shift on labels	[7], [16]
	Label distribution skew	[10], [13]
	Concept shift on features	[10]
FedCelebA	Feature distribution skew	[15]
CIFAR-10	Concept shift on labels	[16], [18], [20]
	Concept shift on features	[9], [10], [19]
	Label distribution skew	[10], [18], [21], [22]
Shakespeare	Feature distribution skew	[5]
Sentiment140	Feature distribution skew	[6]
ml-100K	Concept shift on labels	[14]
ml-1M	Concept shift on labels	[14]
Ag-News	Feature distribution skew	[20]
Time series data	Feature distribution skew	[12]

TABLE I: Summary of datasets and tasks tested in the CFL literature

$1 \leq K \leq N$. That is, the total number of distribution is less than or equal to the number of clients. This implies that the clients can be partitioned into K disjoint clusters C_1, \dots, C_K where the clients’ datasets distributions are homogeneous intra-cluster. For the sake of simplicity in notations, each cluster will be considered as a disjoint partition of $I = \{1, \dots, N\}$ with the notation $i \in C_k$ meaning that the client of index i is inside cluster C_k (i.e.: $I = \bigcup_{k=1}^K C_k$). We would then have K different clusters and for $k \in \{1, \dots, K\}$ a corresponding objective function F_k to optimize the federated model on cluster C_k of the form:

$$\min_{\omega \in \mathbb{R}^d} F_k(\omega) := \sum_{i \in C_k} \frac{n_i}{\sum_{j \in C_k} n_j} f_i(\omega) \quad (3)$$

where $\sum_{j \in C_k} n_j$ correspond to the total number of samples of all clients inside cluster C_k and

$$f_i(\omega) = \mathbb{E}_{(x,y) \sim D_i} [L_i(x, y, \omega)] \quad \forall i \in C_k \quad (4)$$

that is the expected values of the expected value of the local loss function L_i calculated with feature and target (x, y) following the distribution of dataset D_i with model of parameters $\omega \in \mathbb{R}^d$.

Based on this formulation, we describe two of the main CFL solution approaches which we found in the literature.

A. Server-Side Clustering

First, and arguably the most intuitive approach as originally proposed in [8] is often referred to as gradient-based (or model-based) clustering. The studies following this approach solve a clustering problem (c.f Equation (3)) by only using

models weights. This strategy has the advantage of requiring minimum set-up over the baseline FL approach. Specifically, after a number of rounds of global FedAVG aggregations, the central server calculates clusters with k-means applied to the models weights in a "one-shot" manner and then, cluster-specific federated models are optimized separately for each cluster using the FedAVG algorithm.

Practically, in the clustering algorithm, we first define a cluster-associated representative point $\mu_k \in \mathbb{R}^d$ essential for the formation of each cluster C_k for all $k \in \{1, \dots, K\}$. The central server will determine each client $i \in \{1, \dots, N\}$ cluster membership by minimizing the distance between its model parameters $\omega_i \in \mathbb{R}^d$ and the cluster-associated representative point. The formulation can be expressed as k-means-like optimization as follows:

$$\min_{\mu_k, k \in \{1, \dots, K\}} \sum_{k=1}^K \sum_{i=1}^N \mathbf{1}_{i \in C_k} \text{dist}(\omega_i, \mu_k) \quad (5)$$

This approach is the baseline of many other CFL algorithm [5], [6], [10], [20] with **dist** being a generic distance function such as the originally proposed Euclidean distance between models parameters and $\mathbf{1}_{i \in C_k} = 1$ if client i is in cluster C_k , 0 otherwise. Since the cluster assignment is determined by the central server, throughout the remainder of this paper, this type of clustering method will be refer to as "Server-side" clustering.

B. Client-Side Clustering

An alternative to Server-side clustering is to assign the computational burden of determining cluster membership to client nodes. In this case, clients are presented with several possible models, initialized by the central server. The clients then evaluate the loss of their training data on these different cluster models and choose to adhere to the cluster whose model minimizes their loss. This specific approach is named IFCA (Iterative Federated Clustering Algorithm) by its author [9]. Once an initial cluster assignment has been chosen and at each communication round, each client iteratively re-evaluates and determines its optimal cluster. The clusters' models will be updated by the weighted average of corresponding members models parameters after clients' local training.

Formally, let $\mu_k \in \mathbb{R}^d$ be the cluster representative point of cluster C_k for all $k \in \{1, \dots, K\}$ essential for the formation of each cluster. These representative points can be initialized by the central server. Each clients $i \in \{1, \dots, N\}$ will evaluate their training data on the K different models, each initialized with parameters μ_k and update their models parameters ω_i with the cluster representative point which minimizes their local loss function:

$$\omega_i = \arg \min_{\mu_k, k \in \{1, \dots, K\}} \mathbb{E}_{(x,y) \sim D_i} [L_i(x, y, \mu_k)], \quad \forall i \in \{1, \dots, N\} \quad (6)$$

that is the representative point μ_k that minimize the expected values of the local loss function L_i calculated with feature

and target (x, y) following the distribution of dataset D_i with model of parameters μ_k .

This approach is the baseline of many other CFL algorithm [12], [15], [19] and since the cluster assignment in this case is determined by clients, throughout the remainder of this paper, this type of clustering method will be denominated as "Client-side" clustering.

IV. COMPARATIVE ANALYSIS OF STATE-OF-THE-ART CFL SOLUTION WITH RESPECT TO THE NON-IID TAXONOMY

In this section, we compare Client-side and Server-side CFL in non-iid scenarios. For this, we design experiments with heterogeneous distributions as defined in Section II-A. For simplicity, we restrict these experiments to one class of heterogeneity per client. In other words, each clients' dataset is assumed to be homogeneous in itself. Furthermore, given that the objective of CFL is to address the challenges raised by heterogeneous data distributions, we consider clusters to be well-defined if the data distributions within each cluster is homogeneous. As such, we select the initial number of clusters to match the number of data distributions across all clients. For Client-side CFL, as the final number of clusters is dictated by the algorithm, some clusters may end up with no members by the end of the process.

We use the datasets MNIST, fashion-MNIST, and KMNIST [4] for all our experiments. All three datasets contain labeled images from 0 to 9. We assign data to clients to match the desired heterogeneity distribution as described in the following Sections. This straightforward architecture was shown to perform efficiently on the MNIST dataset using CFL [9].


For Server-side CFL we use the algorithm in [8], and for Client-side CFL we use the algorithm described in [9]. We choose these algorithms as they served as baselines for later improvements in the literature. We test the CFL approaches against simple federated learning on the entire dataset, and against central learning on homogeneous subsets which we call the "oracle". These results give us respectively a lower and upper bound on our results.

The results of our experiments are summarized in Tables II, III, IV, V and VI. For each experiment, the tables show the dataset used, the models' accuracy, and the clustering-quality metrics obtained. The metrics reported are the Adjusted Rand Index (ARI), the Adjusted Mutual Information (AMI), as well as the homogeneity (*hom*), completeness (*cmplt*) and v-measure (*vm*) of the clusters. To illustrate what these metrics represent, we plot two edge cases in Figure 3. In the first case (Figure 3a), the assignment follows exactly the clients' underlying data distribution with all metrics equal to one. And in the second case (Figure 3b), the distribution of clients in clusters poorly reflects data heterogeneities with metrics close to zero.

All of these metrics are extrinsic measures calculated against the clustering ground truth (which is no other than the data heterogeneity classes). It is important to note that for CFL, the "optimal" clustering strategy might diverge from the strategy which yields the most efficient models. That is because

Server-side CFL is agnostic to the models performances and relies on the hypotheses that data homogeneity is desirable in our scenario. This might not be true for scenarios such as “quantiy-skew” but analyzing these results gives us some insights on the impact of applying these algorithms without proper knowledge of our data’s heterogeneity.

The following sections discuss the results for each type of heterogeneity scenario. In all our experiments, we used 48 client with 100 samples by label per client. The model used for the classification tasks is a fully connected neural network with ReLU activations [9], with a single hidden layer of size 200 and we set the learning rate to 0.01. For the CFL algorithms we train the models for 20 federated rounds with each round running 10 epochs while the oracle centralized model is parameterized to 50 local epochs. The associated github used for the experiments can be found at: <https://github.com/leahcimalli/Comparative-Evaluation-of-Clustered-Federated-Learning-Methods>



(a) Clients to clusters assignment which reflects the clients underlying data heterogeneity classes. In this case, all metrics equal one

(b) Clients to clusters assignment which does not reflect the underlying clients’ data heterogeneity classes. The resulting clustering metrics approach zero

Fig. 3: Illustration of two clients to clusters assignment edge cases

A. Concept shift on features

To introduce concept shift on features in our datasets, we create four heterogeneity classes by rotating the images 0, 90, 180, and 270 degrees. The distribution of images is such that each heterogeneity class is equally represented amongst clients. Additionally, any given client only has images of one type of rotations as explained above. As such, we set the number of clusters $k = 4$ in the CFL algorithms.

For Server-side CFL, the results show a significant improvement over standard FL, with performances comparable to the personalized oracle models across all experiments (cf Table II). Additionally, the clusters obtained are optimal in all Server-side experiments in the sense that they separate clients by heterogeneity perfectly. These results indicate that data heterogeneity is reflected in the models’ weights which allows k-means to find the optimal clusters.

On the other hand, with Client-side CFL, the algorithm fails to optimally cluster clients by heterogeneity, which results in performances less stellar than with server-side. Several

exp_type	dataset	accuracy	ARI	AMI	hom	cmplt	vm
FL	fashion-mnist	73.81 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	fashion-mnist	83.98 ± 0.69	n/a	n/a	n/a	n/a	n/a
client	fashion-mnist	77.31 ± 2.41	0.48	0.66	0.5	1	0.67
server	fashion-mnist	83.66 ± 1.70	1	1	1	1	1
FL	kmnist	72.83 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	kmnist	87.24 ± 1.31	n/a	n/a	n/a	n/a	n/a
client	kmnist	78.44 ± 5.55	0.7	0.85	0.75	1	0.86
server	kmnist	86.01 ± 2.00	1	1	1	1	1
FL	mnist	80.95 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	mnist	92.33 ± 0.55	n/a	n/a	n/a	n/a	n/a
client	mnist	86.26 ± 4.65	0.7	0.85	0.75	1	0.86
server	mnist	91.78 ± 1.62	1	1	1	1	1

TABLE II: Summary of results for all concept-shift-on-features experiments with average accuracy across clusters

factors can explain this. First, the initial assignment of clients to clusters is random in our experiments. This assignment can have a strong impact of the entire training process. Second, clients can be reassigned to different clusters at any point during the learning phase which leaves less room for improvement of models if the reassignment is done late in the experiment. And finally the features in this scenario very different across heterogeneity classes. This means that sub-optimal clustering will result in slightly less improvements as we see here.

From these results, It appears that personalized learning is crucial for improving performance when data features allow us to group clients with confidence such that is the case in this experiment.

B. Concept shift on labels

In the case of concept shift on labels, we introduce heterogeneity by swapping some labels in the datasets. Specifically, we create 6 heterogeneity classes: swap 1 with 7, 2 with 7, 4 with 7, 3 with 8, 5 with 6, or 7 with 9. As with previous experiments, the modified datasets are distributed across clients ensuring that heterogeneity classes are equally represented. As the resulting datasets imply diverging learning objectives, it is expected that personalized models for each heterogeneity type would also improve the models performances. This is confirmed in the results observed in Table III.

As in the previous scenario, Server-side CFL systematically reaches performances close to the oracle model’s across all datasets. Also predictably, these good performances from the Server-side CFL algorithm result from accurate clustering which is show by the perfect clustering metrics across all experiments.

For Client-side CFL, although the clustering results are not as good, the models obtained still significantly improve on the standard FL baseline. Again This confirms that personalization can still achieve good results even if the resulting clusters are sub-optimal.

exp_type	dataset	accuracy	ARI	AMI	hom	cmplt	vm
FL	fashion-mnist	68.97 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	fashion-mnist	84.48 ± 1.07	n/a	n/a	n/a	n/a	n/a
client	fashion-mnist	76.85 ± 9.33	0.81	0.92	0.87	1	0.93
server	fashion-mnist	84.17 ± 2.00	1	1	1	1	1
FL	kmnist	72.21 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	kmnist	87.98 ± 2.50	n/a	n/a	n/a	n/a	n/a
client	kmnist	76.53 ± 7.40	0.54	0.79	0.69	1	0.82
server	kmnist	86.44 ± 1.94	1	1	1	1	1
FL	mnist	75.17 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	mnist	92.17 ± 1.10	n/a	n/a	n/a	n/a	n/a
client	mnist	84.10 ± 8.89	0.81	0.92	0.87	1	0.93
server	mnist	91.57 ± 1.55	1	1	1	1	1

TABLE III: Summary of results for all concept-shift-on-labels experiments with average accuracy across clusters

C. Features distribution skew

In the case of feature distribution skew, we introduce two types of image transformations: “erosion” which makes the lines thinner, and “dilation” which makes them bolder. We also keep unaltered images as a third heterogeneity class. Consequently we set the number of clusters to $k = 3$.

In this scenario, while mostly improving performance over standard FL, the CFL algorithms have a harder time separating clients into appropriate clusters (c.f Table IV).

exp_type	dataset	accuracy	ARI	AMI	hom	cmplt	vm
FL	fashion-mnist	80.81 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	fashion-mnist	83.83 ± 2.66	n/a	n/a	n/a	n/a	n/a
client	fashion-mnist	78.69 ± 2.45	0.56	0.73	0.58	1	0.73
server	fashion-mnist	82.19 ± 2.78	1	1	1	1	1
FL	kmnist	84.49 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	kmnist	83.65 ± 10.59	n/a	n/a	n/a	n/a	n/a
client	kmnist	75.68 ± 8.90	0	0	0	1	0
server	kmnist	81.61 ± 7.49	0.54	0.64	0.6	0.71	0.65
FL	mnist	87.85 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	mnist	88.86 ± 6.86	n/a	n/a	n/a	n/a	n/a
client	mnist	80.34 ± 9.99	0	0	0	1	0
server	mnist	86.40 ± 8.18	1	1	1	1	1

TABLE IV: Summary of results for all features-distribution-skew experiments with average accuracy across clusters

For Client-side CFL, the algorithm lumps all clients into the same cluster in 2/3 cases, resulting in poor performances. These clustering results might indicate that the loss values in the different clusters are not discriminative enough. Indeed, since the heterogeneities classes are not very different, the loss values calculated at each round may be strongly influenced by the initial random distribution of clients to clusters. Also because of the nature of the Client-side algorithm, it is possible that the cluster model with all clients have had fewer rounds to learn with the entire population of client compared to standard FL model, which results in poorer performances.

With Server-side CFL, while we achieve model performances consistently better than FL, the cluster structure in the ‘k-mnist’ experiment, does not adequately separate clients




Fig. 4: Example clients to clusters distribution resulting from Server-side CFL using the k-mnist dataset for features distribution skew

by heterogeneity class. To illustrate, we see in Figure 4 that all clients with “erosion” or no transformation are assigned to cluster b , while cluster c has most clients with dilation. As opposed to the experiment with image rotations, these results suggest that the features transformations used here are not well reflected in the model weights.

D. Labels distribution skew

In the case of label distribution skew, we down-sampled some labels on each client dataset. Specifically, we create four heterogeneity classes by down-sampling 8 labels out of 10 to 0.1% in each class. This results in selecting $k = 4$ clusters for the CFL algorithms.

exp_type	dataset	accuracy	ARI	AMI	hom	cmplt	vm
FL	fashion-mnist	84.65 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	fashion-mnist	88.04 ± 2.88	n/a	n/a	n/a	n/a	n/a
client	fashion-mnist	84.55 ± 2.88	1	1	1	1	1
server	fashion-mnist	87.70 ± 3.88	1	1	1	1	1
FL	kmnist	85.64 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	kmnist	86.66 ± 3.67	n/a	n/a	n/a	n/a	n/a
client	kmnist	82.39 ± 3.43	1	1	1	1	1
server	kmnist	85.89 ± 4.43	1	1	1	1	1
FL	mnist	91.74 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	mnist	92.61 ± 1.64	n/a	n/a	n/a	n/a	n/a
client	mnist	88.57 ± 3.30	1	1	1	1	1
server	mnist	92.04 ± 3.24	1	1	1	1	1

TABLE V: Summary of results for all labels distribution skew experiments with average accuracy across clusters

As the results in Table V show, both CFL algorithms consistently succeed in separating clients by heterogeneity class while the performances remain sub-optimal. The server-side algorithm reaches models’ accuracies slightly better than standard FL while Client-side CFL has results similar to standard FL. The under performance of Client-side CFL algorithm can partly be explained by the low improvement margin available over the FL algorithm (as indicated by the results of the “oracle” model) and the iterative nature of the algorithm.

E. Quantity skew

In the case of quantity skew, we have clients with varying volumes of data, ranging from 100 percent to 20 percent of the total data distribution mentioned at the start of this section. We use this to introduce $k = 4$ heterogeneity classes, which we select as the number of clusters for CFL. Note that each client will have a balanced quantity of labels to avoid label distribution skew.

exp_type	dataset	accuracy	ARI	AMI	hom	cmplt	vm
FL	fashion-mnist	89.05 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	fashion-mnist	85.66 ± 2.04	n/a	n/a	n/a	n/a	n/a
client	fashion-mnist	82.54 ± 4.83	0	0	0	1	0
server	fashion-mnist	85.21 ± 4.76	-0.01	-0.03	0.06	0.11	0.07
FL	kmnist	93.44 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	kmnist	87.35 ± 2.48	n/a	n/a	n/a	n/a	n/a
client	kmnist	86.10 ± 4.39	0	0	0	1	0
server	kmnist	89.71 ± 4.17	0	-0.01	0.07	0.21	0.1
FL	mnist	95.72 ± 0.00	n/a	n/a	n/a	n/a	n/a
oracle	mnist	92.07 ± 0.56	n/a	n/a	n/a	n/a	n/a
client	mnist	90.33 ± 3.76	0	0	0	1	0
server	mnist	92.13 ± 3.50	0	-0.01	0.07	0.21	0.1

TABLE VI: Summary of results for all quantity skew experiments with average accuracy across clusters

It is expected that clustered federated learning with quantity skew will fail at improving performances over standard FL. Understandably, if the objective of clustering by heterogeneity class is respected, some clusters with very low volumes of data will be constituted. These low data volume clusters are expected to perform poorly. And indeed, as we see in Table VI, the results validate this intuition in that CFL algorithms do not improve over a simple FL model (We note that this is also the case for the oracle algorithm given that we separate clients by heterogeneity class as well).

It is important to note that in the standard FL model, FedAVG aggregation takes into account the number of local samples of each client. As such, clients with larger data will have a greater impact on the overall global model and potentially mitigate the discrepancy of dataset sizes among clients.

In our Server-side experiments, the clusters have no discernible logic and create models which slightly underperform. Alternatively, with Client-side CFL, the algorithm clusters all clients together which results in a set-up and performances similar to that of the baseline FL method but with less training rounds.

F. Impact of the number of clusters

All the experiments presented assume that the number of clusters is well defined. That is, the number of clusters is always set to the number of data heterogeneity classes in the dataset. In real-life, knowing the number of heterogeneity classes in advance might not always be possible. As such, even an educated guess of an expert might miss the mark on the exact number of data heterogeneities present. To evaluate the

impact of this variable, we tested the impact of the number of clusters on the accuracy for the Server-side “concept-shift-on-labels” problem with the mnist dataset.




Fig. 5: Impact of the number of clusters on the CFL results

As we see in Figure 5, when the number of clusters is very low compared to the number of heterogeneity classes (here being six like in Section IV-B), the results approach those observed with the baseline FL algorithm. As we increase the number of clusters, the accuracy increases as well. We note that the accuracy seems to stabilize if we increase the number of clusters beyond the nominal value. This means that in some circumstances (and assuming sufficient data volumes per clusters), models personalization can go beyond the number of heterogeneity classes without penalizing model performances.

V. CONCLUSION

In this study, we explored the performance of Server-side and Client-side Clustered Federated Learning (CFL) algorithms under systematized data heterogeneity conditions. Our findings show improved performances for several use-cases with more consistent results from server-side CFL.

In most cases, the accuracy improvements follow a clustering stage which successfully separates clients by heterogeneity class. This is the case in scenarios where the features or the labels are significantly different across clients such as in concept-shift-on-features and concept-shift-on-labels scenarios. On the other hand, when the differences across the heterogeneity classes are more subtle (e.g features distribution skew), even server-side CFL has a harder time separating by heterogeneity class but personalization still improves results. In the case of labels distribution skew, it seems that the heterogeneity does not have a strong impact on the model performances with standard FL achieving quite high accuracies. As such CFL method only slightly improve these results. Finally, when the data heterogeneity leads to clusters with incomplete data for learning (e.g quantity skew) it seems undesirable to use CFL at all.

This might suggest that an informed understanding of the learning environment might be key to making adequate decisions on when and how to use CFL. Fortunately, even in the quantity-skew scenario, the performances of the models were not strongly penalized by clustering. Further, in Client-side clustering, as the initial assignment of clients to clusters has a strong impact on the entire training process, it is important to make an informed decision at this initial stage. We confirmed

this intuition by starting the CFL experiments with a semi-random assignment (essentially simulating an expert opinion on the heterogeneity class of the client) and obtained improved results which approached those of server-side CFL. Further investigation in this direction is needed.

We also investigated the impact of the number of clusters on the performance of our models. We observed that, as expected, the best results are obtained when the number of clusters is close or equal to the number of data heterogeneity classes. It also appears safer to overestimate the number of heterogeneity classes (up to a certain limit) which might be explained by unaccounted heterogeneities in the data.

Further research is needed to see if these results scale beyond simple datasets, and with the presence of a combination of data heterogeneities by client.

REFERENCES

- [1] A. Augello, G. Falzone, and G. Lo Re, "DCFL: Dynamic clustered federated learning under differential privacy settings," in *2023 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, 2023.
- [2] C. Briggs, Z. Fan, and P. Andras, "Federated learning with hierarchical clustering of local updates to improve training on non-iid data," in *2020 International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2020.
- [3] F. E. Castellon, A. Mayoue, J.-H. Sublemontier, and C. Gouy-Pailler, "Federated learning with incremental clustering for heterogeneous data," in *2022 International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2022.
- [4] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto and D. Ha "Deep Learning for Classical Japanese Literature", *arxiv preprint arxiv:1812.01718*, 2018
- [5] D. K. Dennis, T. Li, and V. Smith, "Heterogeneity for the win: One-shot federated clustering", *International Conference on Machine Learning*, PMLR, 2021.
- [6] M. Duan, et al., "Fedgroup: Efficient federated learning via decomposed similarity-based clustering," in *2021 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCLOUD/SocialCom/SustainCom)*, IEEE, 2021.
- [7] M. Duan, et al., "Flexible clustered federated learning for client-level data distribution shift," *IEEE Transactions on Parallel and Distributed Systems*, 2021.
- [8] A. Ghosh, J. Hong, D. Yin, and K. Ramchandran, "Robust Federated Learning in heterogeneous environments," *arXiv preprint arXiv:1906.06629*, 2019.
- [9] A. Ghosh, J. Chung, D. Yin, K. Ramchandran, and A. Chandra, "An efficient framework for clustered federated learning," *Advances in Neural Information Processing Systems* 2020.
- [10] B. Gong, T. xing, Z. Liu, W. xi, and X. Chen, "Adaptive client clustering for efficient federated learning over non-iid and imbalanced data," *IEEE Transactions on Big Data*, 2022.
- [11] P. Kairouz, et al., "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, 2021
- [12] Y. Kim, et al., "Dynamic clustering in federated learning," in *ICC 2021-IEEE International Conference on Communications*, IEEE, 2021.
- [13] C. Li, G. Li, and P. K. Varshney, "Federated learning with soft clustering," *IEEE Internet of Things Journal*, 2021.
- [14] T. Liang et al., "Efficient one-off clustering for personalized federated learning," *Knowledge-Based Systems*, 2023.
- [15] G. Long, et al., "Multi-center federated learning: clients clustering for better personalization," *World Wide Web*, 2023.
- [16] Y. Luo, x. Liu, and J. xiu, "Energy-efficient clustering to address data heterogeneity in federated learning," in *ICC 2021-IEEE International Conference on Communications*, IEEE, 2021.
- [17] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial Intelligence and Statistics*, PMLR, 2017.
- [18] M. Mehta, and C. Shao, "A greedy agglomerative framework for clustered federated learning," *IEEE Transactions on Industrial Informatics*, 2023.
- [19] Y. Ruan, and C. Joe-Wong, "Fedsoft: Soft clustered federated learning with proximal local updating," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.
- [20] F. Sattler, K.-R. Müller, and W. Samek, "Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints," *IEEE Transactions on Neural Networks and Learning Systems*, 2019.
- [21] J. Wolfrath, N. Sreekumar, D. Kumar, Y. Wang, and A. Chandra, "Haccs: heterogeneity-aware clustered client selection for accelerated federated learning," in *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, IEEE, 2022.
- [22] Y. Xiao, J. Shu, x. Jia, and H. Huang, "Clustered federated multi-task learning with non-iid data," in *2021 IEEE 27th International Conference on Parallel and Distributed Systems (ICPADS)*, IEEE, 2021.
- [23] M. Ye, X. Fang, B. Du, P. C. Yuen, and D. Tao, "Heterogeneous federated learning: State-of-the-art and research challenges", *ACM Computing Surveys*, 2023.
- [24] Y. Zhao, et al., "Federated learning with non-iid data," *arxiv preprint arxiv:1806.00582*, 2018.