

Measuring Data Similarity for Efficient Federated Learning: A Feasibility Study

Fernanda Famá[†], Charalampos Kalalas^{*}, Sandra Lagen^{*}, Paolo Dini^{*}

^{*}Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA)

[†]Dept. of Signal Theory and Communications, Universitat Politècnica de Catalunya, Spain

Email: {fernanda.fama, ckalalas, slagen, pdini}@cttc.es

Abstract—In multiple federated learning schemes, a random subset of clients sends in each round their model updates to the server for aggregation. Although this client selection strategy aims to reduce communication overhead, it remains energy and computationally inefficient, especially when considering resource-constrained devices as clients. This is because conventional random client selection overlooks the content of exchanged information and falls short of providing a mechanism to reduce the transmission of semantically redundant data. To overcome this challenge, we propose clustering the clients with the aid of similarity metrics, where a single client from each of the formed clusters is selected in each round to participate in the federated training. To evaluate our approach, we perform an extensive feasibility study considering the use of nine statistical metrics in the clustering process. Simulation results reveal that, when considering a scenario with high data heterogeneity of clients, similarity-based clustering can reduce the number of required rounds compared to the baseline random client selection. In addition, energy consumption can be notably reduced from 23.93% to 41.61%, for those similarity metrics with an equivalent number of clients per round as the baseline random scheme.

Index Terms—federated learning, similarity metrics, clustering, client selection

I. INTRODUCTION

The edge computing paradigm has become a breakthrough solution that allows certain processing resources to be allocated close to the data sources rather than relying only on cloud computing centers. However, edge-computing devices have limited computational and energy resources, which makes training complex machine learning (ML) models challenging. Recently, federated learning (FL) has emerged as an alternative for performing ML tasks in a distributed and privacy-preserving manner by having multiple clients collaborate without sharing their data [1]. This is because local learning clients share their parameters with a central server, which aggregates them to create a proper global model after a series of rounds.

Despite the privacy and computational sharing benefits of FL across devices, the use of multiple clients can create a communication bottleneck [2]. This may be further exacerbated by the energy consumption spent training FL models on such resource-constrained devices. Therefore, several studies have proposed techniques to select, at every round, a subset of clients to participate in training the global model. However, de-

termining how many and which clients to choose is not straightforward, therefore a random selection method is usually adopted. In turn, this may inadvertently lead to the selection of similar and redundant information for aggregation, which requires a higher commitment of resources to obtain a global model suitable for all clients [3]. Random selection may thus have a negative impact on learning efficiency, fairness, convergence and, eventually, energy consumption. Apart from this, data on client devices may not be evenly distributed in real-world FL tasks, and non-iid solutions need to be considered. The distribution divergence of non-iid data introduces significant challenges in FL [4] such as lower accuracy, delays in model communication, and slower model convergence.

Motivated by the challenges mentioned above, we hereby advocate for using similarity metrics [5] as a means of grouping FL clients with similar information. To achieve this, we consider scenarios with a skewed label distribution, and we perform cluster formulation with the aid of the following similarity metrics: cosine function; mean squared error; Euclidean, Manhattan, and Chebyshev distances; maximum mean discrepancy; Kullback-Leibler and Jensen-Shannon divergences; and Wasserstein distance. The incorporation of such metrics allows us to leverage correlations in the local information to form *semantically informative* clusters. Selecting clients from these clusters minimizes the transmission of redundant information by harnessing the underlying heterogeneity of local training datasets. Consequently, it accelerates FL training efficiently. Our twofold contribution can be summarized as follows:

- We explore the feasibility of applying different statistical similarity metrics on local data to perform informative client clustering. The incorporation of similarity metrics in client selection allows FL training to be performed in a way that promotes dissimilarity of selected clients. Our goal resides in not only reducing redundancy in the FL training phase, but also in quantifying the potential energy-efficiency gains.
- We perform an elaborate comparison between similarity-based clustering and random client selection with the aid of multiple FL performance indicators, such as accuracy, number of rounds, and energy consumption, for different degrees of skewness in label distributions. Key observations stemming from the trade-off between the number of rounds and energy consumption in FL training are made.

The present work has been funded by the European Union's Horizon 2020 Marie Skłodowska Curie Innovative Training Network Greenedge (GA. No. 953775).

This work is organized as follows. Section II presents relevant work on non-iid data distribution and client selection strategies in FL. System model considerations are provided in Section III. Section IV introduces the statistical similarity metrics (IV-A), the client selection strategy (IV-B), and the energy consumption model (IV-C) used in this work. Performance results pertaining to the evaluation of similarity-based clustering against random selection are provided in Section V. Section VI is reserved for conclusions and discussion of the path forward.

II. BACKGROUND AND RELATED WORK

The disadvantages of having all clients participating in the FL training have led several studies to investigate the development and improvement of techniques for selecting a portion of clients to join in the FL training. However, the design of an optimal selection strategy is not a straightforward task, mainly due to the detrimental effects of non-iid data distribution on the behavior of FL training [4, 6–8]. For example, an experimental study is presented in [7] that investigates six non-iid partitioning strategies for different FL algorithms. Non-iid distributions are shown to have a significant impact on the accuracy and stability of FL algorithms, and the authors propose the implementation of client subgroups in each FL round to deal with the instability caused by random sample selection. In [8], on the other hand, the authors use the distributions of the labels to deal with heterogeneity in a non-iid medical scenario and focus on computing the marginal distribution for each client as an optimization strategy.

In an effort to optimize the communication between the clients and the server, different client selection strategies have been proposed [3, 9–14]. In [13], client selection relies on the sample size and various similarity methods. As such, clients with more samples are weighted the most, which, in turn, does not prevent the transmission of redundant information, rendering this approach inefficient. It is also mentioned that measuring client similarity enables better clustering and leads to better performance. The considered similarity refers to the representative gradient, i.e., the difference between the updated model of a client and the global model, and is not extended to cluster similar clients. In [14], hierarchical clustering is implemented to group FL clients according to the similarity of their model weight updates compared to the global joint model. The authors apply only distance metrics to compute the similarity between clusters.

A contextual client selection framework is proposed in [11] for a vehicle-to-everything scenario. The solution comprises information sharing, traffic topology prediction, and client clustering at both data and network levels. The clients are grouped at the data level using their gradient as a similarity criterion. In [10], the authors used a probabilistic deep learning model to create a personalized feature extractor integrated into each client. To estimate client similarity, the server sends the model of all other clients to each client to measure the discrepancy between the two distributions. Based on the resulting discrepancy, the server then clusters the clients. The

gradient is also used in [9] as an insight to estimate the skewness of the client, using Hoeffding’s inequality. A dueling bandit is then used to select the clients with the lower skewness based on this estimate. In [3], the authors select a subset of active clients in each FL round using the stochastic greedy algorithm, which requires knowing the gradients from these clients. For this reason, a round of local model updates is performed before sharing the updates with the server.

Our study differs substantially from previous approaches. First, rather than focusing on developing a strategy for client selection that needs to be integrated into the training phase, we propose a solution focusing on the client side that can be more easily integrated into a federated system. This is because our solution is centered on providing a client selection strategy before the FL procedure starts. Second, although some works have used similarity metrics to improve FL client selection, they do not provide a rigorous comparative assessment of the metrics themselves. As such, this work aims to compare a broad set of similarity metrics to gain in-depth insights into the FL performance under different degrees of skewness in label distributions. Lastly, it is essential to obtain more efficient and sustainable FL systems (i.e., without requesting more computational, storage, or transmit power) and, at the same time, maintain or reduce energy consumption. Thus, it may not be enough to reduce the number of FL rounds in training, especially when we do not quantify the impact on energy consumption. This is a fundamental aspect that none of the previous studies have addressed. To fill these gaps, in this work we investigate the use of similarity metrics in the optimization of FL training, using the clients’ data distribution to find a trade-off between the number of rounds and the energy consumption during this process. Relevant aspects which, to the best of our knowledge, have not been investigated in a similar study before.

III. SYSTEM MODEL

In this work, we have adopted the FedAvg algorithm for distributed training [1], where a central server aggregates a weighted average of the participating clients’ model parameters. The weights are determined by the size of each client’s training dataset. We denote the number of clients as N and the local training dataset of client i as $D^i = \{(\mathbf{x}_i, \mathbf{y}_i)\}$. Focusing on multi-label classification as the learning task, \mathbf{x}_i stands for the data instances available at client i and \mathbf{y}_i represents the associated labels. The local model of client i in communication round t is expressed as w_i^t , and the global model as w_t . Aggregation is performed after each communication round, as shown in Fig. 1. Consequently, the participating clients optimize their local model before transmitting their model parameters to the server afterwards [7].

Since the data collected by an underlying system are usually non-iid [4], we henceforth focus on the use of similarity metrics to capture complex relationships among local datasets. These metrics are essential to measure the similarity or dissimilarity between data instances. Although their use is primarily attributed to transfer learning tasks, such information measures




Fig. 1. Overview of the considered problem setup in FL. Cluster formulation takes place with the aid of similarity metrics. At each communication round, client selection from clusters aims to harness the underlying heterogeneity of local training datasets.

could be applied to a wide range of problems, such as feature extraction, information reuse, and multi-task learning, to name a few. In our case, these statistical measures are used to form clusters, where client selection in each round of the federated training is performed in a way that promotes dissimilarity of selected clients, as illustrated in Fig. 1. This approach aims at a more efficient version of FedAvg scheme by avoiding the transmission of semantically redundant information to the server. As a result, given a certain accuracy threshold, the number of required rounds and the energy consumption to train the global model are expected to be reduced. Consequently, the pressure on the underlying connectivity can be alleviated by saving network resources such as bandwidth and transmission power. Moreover, non-selected clients at every round t may not perform local training, thus saving computing power.

We assume skewed label distribution that results in heterogeneity of the local models. The label distribution at each client is considered to be known. For clustering, we calculate the number of samples per label for each client, $n_{i,k}$, and we divide by the total number of samples n_i , as

$$p_{i,k} = \frac{n_{i,k}}{n_i}, \quad (1)$$

where $p_{i,k}$ represents the distribution of label k at client i . We perform this calculation for all clients and we obtain $\mathbf{P}^{N \times K}$, as

$$\mathbf{P} = \begin{bmatrix} p_{1,0} & p_{1,1} & \cdots & p_{1,K-1} \\ p_{2,0} & p_{2,1} & \cdots & p_{2,K-1} \\ \vdots & \vdots & \ddots & \vdots \\ p_{N,0} & p_{N,1} & \cdots & p_{N,K-1} \end{bmatrix}, \quad (2)$$

where K denotes the total number of labels. Each row vector \mathbf{p}_i , $i \in [1, N]$, in \mathbf{P} is equivalent to the probability mass function of the labels for client i .

IV. CLIENT SELECTION USING SIMILARITY METRICS

A. Statistical Similarity Metrics

This section presents the similarity metrics considered in our approach and relevant studies that have employed them in the past for other purposes.

a) *Cosine function*: It estimates the similarity between two vectors by measuring the angle between them. The “closer” the two vectors are, the smaller the angle between them. This widely used metric is defined as

$$\text{cos}_{i,j} = \frac{\langle \mathbf{p}_i, \mathbf{p}_j \rangle}{\|\mathbf{p}_i\| \|\mathbf{p}_j\|}, \quad \text{for } i \neq j, \quad (3)$$

where $\mathbf{p}_i, \mathbf{p}_j$ represent two different rows in \mathbf{P} , and $\|\cdot\|$ denotes the ℓ^2 -norm (i.e., magnitude) of the vector. In [15], a cosine similarity function is used for original data recovery in the context of text transmission.

b) *Mean Squared Error (MSE)*: This metric measures the average squared difference between each particular element in two different rows in \mathbf{P} . It is expressed as

$$\text{MSE}_{i,j} = \frac{1}{K} \sum_{k=1}^K (p_{i,k} - p_{j,k})^2, \quad \text{for } i \neq j, \quad (4)$$

where values closer to zero represent higher similarity. In [16], the metric is used for image quality assessment purposes.

c) *Euclidean distance*: It is a widely used metric that represents the shortest distance between two elements. Considering two different rows in \mathbf{P} , it is expressed as

$$D_{E_{i,j}} = \left(\sum_{k=1}^K (p_{i,k} - p_{j,k})^2 \right)^{\frac{1}{2}}, \quad \text{for } i \neq j. \quad (5)$$

This metric was adopted in [14] to compute the similarity between clusters in a hierarchical clustering algorithm.

d) *Manhattan distance*: This metric computes the absolute distance between two elements in \mathbf{P} , defined as

$$D_{M_{i,j}} = \sum_{k=1}^K |p_{i,k} - p_{j,k}|, \quad \text{for } i \neq j. \quad (6)$$

The metric was adopted in [14] for the same purpose as the Euclidean distance.

e) *Chebyshev distance*: This metric calculates the maximum of the absolute distance between two elements in \mathbf{P} . It is defined as

$$D_{C_{i,j}} = \sum_{k=1}^K \max_k |p_{i,k} - p_{j,k}|, \quad \text{for } i \neq j. \quad (7)$$

In [17], the authors adopted this metric to detect whether images are similar or dissimilar in the context of dimensionality reduction.

f) *Maximum Mean Discrepancy (MMD)*: It quantifies the distribution difference by computing the distance between the mean values of the instances in a reproducing kernel Hilbert space (RKHS) [18]. The calculation of MMD implies finding the RKHS function that maximizes the difference in

the expectations $E(\cdot)$ between two rows in \mathbf{P} . In our case, we consider a linear kernel (i.e., $k(x, y) = \langle x, y \rangle$) and MMD can be calculated as

$$MMD^2(\mathbf{p}_i, \mathbf{p}_j) = E_{\mathbf{p}_i}[k(x, x)] - 2E_{\mathbf{p}_i, \mathbf{p}_j}[k(x, y)] + E_{\mathbf{p}_j}[k(y, y)], \quad \text{for } i \neq j, \quad (8)$$

where $x \sim \mathbf{p}_i$, $y \sim \mathbf{p}_j$. In [19], MMD is incorporated in a neural network model as a regularization technique to reduce the distribution difference between source and target domains in transfer learning.

g) *Kullback-Leibler (KL) divergence*: It measures the statistical distance to minimize the divergence between two probability distributions, as

$$D_{KL_{i,j}} = \sum_{k=1}^K p_{i,k} \log \frac{p_{i,k}}{p_{j,k}}, \quad \text{for } i \neq j. \quad (9)$$

This metric was adopted in [20] to measure the dissimilarity between domains as a mechanism to deal with the domain adaptation problem in transfer learning.

h) *Jensen-Shannon divergence*: This metric is a symmetric version of the KL divergence, and aims to minimize the difference between two distributions [21]. For $\mathbf{p}_i, \mathbf{p}_j \in \mathcal{P}$ with $i \neq j$, it is defined as

$$D_{JSD}(\mathbf{p}_i \parallel \mathbf{p}_j) = \frac{1}{2} (D_{KL}(\mathbf{p}_i \parallel \mathbf{q}) + D_{KL}(\mathbf{p}_j \parallel \mathbf{q})), \quad (10)$$

where \mathbf{q} is a mixed distribution defined as $\mathbf{q} = \frac{1}{2}(\mathbf{p}_i + \mathbf{p}_j)$. This metric was adopted in [21] to measure the similarity of features in an activity recognition problem.

i) *Wasserstein distance*: This metric measures the minimal effort required to reconfigure \mathbf{p}_i in order to recover a distribution \mathbf{p}_j , with $i \neq j$. In particular, we consider the 1-Wasserstein distance, expressed as

$$W_1(\mathbf{p}_i, \mathbf{p}_j) = \inf_{\gamma \in \Gamma(\mathbf{p}_i, \mathbf{p}_j)} \int_{\mathcal{R} \times \mathcal{R}} |x - y| d\gamma(x, y), \quad (11)$$

where $\Gamma(\mathbf{p}_i, \mathbf{p}_j)$ denotes the set of probability measures γ on $\mathcal{R} \times \mathcal{R}$. Elements $\gamma \in \Gamma(\mathbf{p}_i, \mathbf{p}_j)$ are called couplings of \mathbf{p}_i and \mathbf{p}_j , i.e., joint distributions on $\mathcal{R} \times \mathcal{R}$ with prescribed marginals \mathbf{p}_i and \mathbf{p}_j . Intuitively, Eq. (11) implies that given a $\gamma \in \Gamma(\mathbf{p}_i, \mathbf{p}_j)$ and a pair of samples (x, y) , the value of $\gamma(x, y)$ reveals the proportion of \mathbf{p}_i 's mass at x that has to be transferred to y , in order to reconfigure \mathbf{p}_i into \mathbf{p}_j . In [22], the authors adopted the Wasserstein distance to align the feature distribution domains and capture the dissimilarity between the outputs of task-specific classifiers.

B. Client selection

In this section, we integrate the similarity metrics into our client selection strategy. As shown in Algorithm 1, based on the selected similarity metric, the pairwise calculations between rows in \mathbf{P} are performed. Next, we leverage k-medoids¹ as the clustering scheme to group the N clients

¹K-medoids is publicly available in the Python library *scikit-learn-extra*: <https://scikit-learn-extra.readthedocs.io/en/stable/modules/cluster.html#k-medoids>

into clusters. For all possible cluster numbers $c \in [2, N - 1]$, we calculate the silhouette value [23] for client i in order to determine the input value for k-medoids. The silhouette value is computed using

$$s_c(i) = \frac{d_{\text{inter}}(i) - d_{\text{intra}}(i)}{\max\{d_{\text{intra}}(i), d_{\text{inter}}(i)\}}, \quad (12)$$

where $d_{\text{intra}}(i)$ denotes the mean intra-cluster distance, i.e., between the client i and all other clients in the same cluster and $d_{\text{inter}}(i)$ denotes the smallest mean inter-cluster distance, i.e., between the client i and all other clients in any other cluster. The higher the silhouette value, the higher the probability of client i being clustered in the correct group.

Algorithm 1 Applying similarity metrics to cluster the clients in a FedAvg algorithm.

Input: FedAvg inputs, metric $m \in \mathcal{M} = \{\cos, \text{mse}, \text{mmd}, D_{kl}, D_{jsd}, W_1, D_E, D_M, D_C, \text{random}\}$, fraction of clients ϵ

Output: The final model w_t

```

1: Compute  $p_{i,k} = \frac{n_{i,k}}{n_i}, \forall i \in N, k \in K$ 
2: Construct  $\mathbf{P}$  in Eq. (2)
3: if  $m \neq \text{random}$  then
4:   Calculate pairwise similarities of  $\mathbf{P}$  rows using Eqs. (3)-(11)
5:   for cluster number  $c = 2, \dots, N - 1$  do
6:     Calculate  $s_c$  in Eq. (12)
7:   end for
8:   Select  $c$  with  $\max_c s_c$  and apply k-medoids for  $c$  clusters
9: end if
Server executes:
10: initialize  $w_0$ 
11: for each round  $t = 1, \dots, T_{\text{train}}$  do
12:   if  $m \neq \text{random}$  then
13:      $S_t \leftarrow$  (random set of  $n$  clients in  $c$  clusters)
14:   else
15:      $n \leftarrow \max(\epsilon \cdot N, 1)$  clients
16:      $S_t \leftarrow$  (random set of  $n$  clients)
17:   end if
18: end for

```

Once clustering using k-medoids is completed, the mapping of clients per cluster is available. Based on this information, the number of clients S_t that will participate in the training rounds is determined. On the other hand, when random selection is adopted, a predetermined number of clients are selected to join in each round (lines 15-16). After the client selection process, the conventional FedAvg algorithm follows.

C. Energy consumption

Besides efficient clustering, we also aim to derive energy consumption insights related to the FL process. We focus on the computational energy, being the highest component in consumption, as concluded in our previous work [24]. To this end, we evaluate the computational energy by considering a predefined set of clients S_t . In particular, the energy consumption for client i is defined as

$$e_i = P_{\text{hw},i} \cdot T_{\text{train},i}, \quad (13)$$

where $P_{\text{hw},i}$ denotes the sum of the hardware power consumption of the GPU, RAM, and CPU for client i during training time $T_{\text{train},i}$.




Fig. 2. Client clustering ($\beta = 0.05$) with the aid of (a) 1-Wasserstein and (b) Chebyshev metrics. The principal component analysis (PCA) method was used as a dimensionality reduction technique to allow the mapping of features into two principal components.

V. PERFORMANCE EVALUATION

A. Experimental setup

We implement and evaluate our approach by performing experiments on the well-known MNIST dataset [25] of 28x28 images of handwritten digits, which comprises 60000 samples as a training set and 10000 samples as a test set. For all experiments, classification is performed with the aid of a convolutional neural network architecture comprising two 5x5 convolutional layers, followed by a 2x2 max pooling layer and two fully connected layers with ReLU activation. Experiments are performed on a server equipped with a 16-core Intel Xenon CPU and 2 NVIDIA GeForce RTX 3090 GPUs. We consider an FL system with $N = 100$ clients. To generate a set of non-identical label distributions with different levels of skewness, we apply the Dirichlet distribution with a varying concentration parameter $\beta > 0$ [7]. In this way, we can allocate a diverse proportion of samples of label k to each client i . The smaller the β , the more unbalanced the label distribution.

B. Performance results

For each similarity metric listed in Section IV-A, we evaluate the number of required rounds in the FL process and the standard deviation of the achieved accuracy for 3 consecutive rounds, considering a predetermined accuracy threshold of 97%. The computational energy during FL training is also assessed with the aid of CodeCarbon² library [26] to estimate hardware power consumption according to Eq. (13). In addition, the baseline scheme of random selection with a varying number of participating clients $n \in [2, 5, 10, 15, 20, 25]$ is evaluated. For all experiments, we consider 5 different seeds and calculate the average value of the obtained results.

²CodeCarbon library is available at: <https://github.com/mlco2/codecarbon>

TABLE I
SIMILARITY-BASED CLUSTERING VS RANDOM SELECTION ($\beta = 0.05$,
ACCURACY=97%) IN FL TRAINING¹

Metric	Clients per round	Number of rounds	Energy consumption (Wh)	Acc (std)
1-Wasserstein	2	45.8	155.388	0.003542
JS-divergence	8.667	39.333	469.727	0.004051
KL-divergence	8.667	39.333	471.221	0.004051
Euclidean	10	45.8	329.746	0.002496
Chebyshev	10.2	45.8	328.94	0.001439
Manhattan	10.4	45.8	400.298	0.003918
MSE	10.4	45.8	427.103	0.005676
MMD	10.4	45.8	428.517	0.005676
Cosine	11	45.8	424.304	0.003945
Random Selection	2	215.2	204.448	0.003261
	5	113.40	377.386	0.003738
	10	87.8	563.378	0.006298
	15	63.80	621.135	0.001717
	20	57.80	840.196	0.003119
	25	49.4	726.045	0.002804

¹ We ran the algorithm five times by changing the random seed and reporting the averaged classification accuracy.

We first assess the FL performance with similarity-based clustering and with random client selection in a highly heterogeneous scenario with $\beta=0.05$. The experimental results in Table I reveal that the number of rounds required to reach convergence in the case of similarity-based clustering is lower compared to the random selection for all metrics. It is worth noting that the number of clients per round is not *a priori* defined, but rather it is determined for each metric using Algorithm 1 (line 8). A careful inspection of Table I reveals that energy consumption can be remarkably reduced from 23.93% to 41.61% compared to the baseline random selection, for those similarity metrics achieving equivalent number of clients per round ($n = 10$).




Fig. 3. Clients of a single cluster ($\beta = 0.05$, $n = 10$) formed with (a) Euclidean distance metric (b) random selection. For random selection, we select the first out of 10 clients in each FL training round.

Notably, for the 1-Wasserstein distance, the number of required rounds for convergence is ≥ 4.5 times lower than in random selection for the same number of participating clients ($n=2$). As illustrated in Fig. 2(a), a clear distinction between the two Wasserstein-based clusters can be achieved. In addition, the associated energy consumption in the training process remains at significantly lower levels compared to the random selection. On the other hand, clustering based on the Chebyshev similarity metric results in rather overlapping clusters³, as depicted in Fig. 2(b). The well-separated clusters under the 1-Wasserstein metric can be attributed to the fact that the Wasserstein space is able to subtly capture the geometry of the domain of the distributions in \mathcal{P} [27].

To demonstrate the efficacy of similarity-based clustering, Fig. 3(a) depicts the resulting client agglomeration within a single cluster when Euclidean distance is applied for $n=10$. It can be observed that similarity-based clustering allows grouping together those clients having samples predominantly of label 2, including the client having only 36 samples (client #0). On the contrary, for random client selection, no relevant criterion for the label distribution is applied, resulting in an unstructured association of clients per cluster, as shown in Fig. 3(b) for $n=10$.

While the performance gains are noticeable for highly heterogeneous scenarios (i.e., low β), we henceforth explore the feasibility of similarity-based clustering in more homogeneous setups with higher β . Table II depicts the comparative performance outcomes for similarity-based clustering and random selection for $\beta=0.1$. In this case, it is noted that only Chebyshev and Manhattan metrics achieve superior

³Note that similar behavior is observable also when using the other similarity metrics under study.

TABLE II
SIMILARITY-BASED CLUSTERING VS RANDOM SELECTION ($\beta = 0.1$, ACCURACY=97%) IN FL TRAINING.

Metric	Clients per round	Number of rounds	Energy Consumption (Wh)	Acc (std)
1-Wasserstein	2	127.2	97.725	0.004085
Euclidean	9.8	34.8	139.428	0.002823
Chebyshev	10	33.6	117.513	0.003242
MSE	10	36.6	157.028	0.004805
MMD	10	36.6	157.606	0.004805
Manhattan	10.2	31.4	122.535	0.002134
Cosine	10.2	31.4	140.241	0.001658
JS-divergence	12.2	60.8	149.1338	0.003624
KL-divergence	12.2	60.8	164.8476	0.003624
Random Selection	2	106	83.027	0.002556
	5	58.2	116.507	0.005039
	10	34	133.574	0.002512
	15	31.8	184.572	0.002197

performance in terms of both number of rounds and energy consumption compared to random selection for $n=10$. However, we remark here that even though the similarity-based clustering approach does not present such high-performance gains in more homogeneous data distributions, it does not require *a priori* information on the number of clients to be selected at every round, as in the random approach. This feature represents an added value and poses elevated merit for fast FL deployments in realistic applications. Finally, the performance gains of similarity-based clustering compared to random selection further vanish when $\beta=2$, as shown in Table III. It is noted, though, that such homogeneous settings are not common in real-world FL scenarios.

TABLE III
SIMILARITY-BASED CLUSTERING VS RANDOM SELECTION ($\beta = 2$,
ACCURACY=97%) IN FL TRAINING.

Metric	Clients per round	Number of rounds	Energy Consumption (Wh)	Acc (std)
1-Wasserstein	2	14	11.712	0.003469
JS-divergence	4	15.8	28.468	0.002598
KL-divergence	4	15.8	28.707	0.002598
MSE	20	8.6	76.105	0.001238
MMD	20	8.6	77.53	0.000872
Chebyshev	22	8.6	88.689	0.001262
Manhattan	24	10.4	101.065	0.001792
Euclidean	24.8	8.6	97.8324	0.001689
Cosine	26.6	9	88.689	0.001262
Random Selection	2	12.6	11.591	0.003823
	5	8.4	19.098	0.003137
	10	8.4	38.108	0.001083
	15	7.8	53.406	0.002364
	20	8	72.672	0.001717
	25	8	92.24	0.001712

VI. CONCLUSION AND FUTURE WORK

In this paper, we proposed a similarity-based client selection approach to reduce the transmission of semantically redundant data. To this end, we incorporated nine statistical similarity metrics in the FL client clustering process. Experimental results reveal that the more heterogeneous the local client data are, the more effective our method is, capitalizing on the lower number of required rounds and reduced energy consumption in FL training compared to random selection. In the path forward, we will direct our efforts towards integrating into our approach the quality of communication links between the participating clients and the server.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [2] O. Shahid, S. Pouriyeh, R. M. Parizi, Q. Z. Sheng, G. Srivastava, and L. Zhao, "Communication efficiency in federated learning: Achievements and challenges," *arXiv preprint arXiv:2107.10996*, 2021.
- [3] R. Balakrishnan, T. Li, T. Zhou, N. Himayat, V. Smith, and J. Bilmes, "Diverse client selection for federated learning via submodular maximization," in *Int. Conf. on Learning Representations*, 2022.
- [4] X. Ma, J. Zhu, Z. Lin, S. Chen, and Y. Qin, "A state-of-the-art survey on solving non-IID data in Federated Learning," *Future Generation Computer Systems*, vol. 135, pp. 244–258, 2022.
- [5] A. Visús, J. García, and F. Fernández, "A taxonomy of similarity metrics for markov decision processes," *arXiv preprint arXiv:2103.04706*, 2021.
- [6] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and communication-efficient federated learning from non-iid data," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 9, pp. 3400–3413, 2019.
- [7] Q. Li, Y. Diao, Q. Chen, and B. He, "Federated learning on non-iid data silos: An experimental study," in *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2022, pp. 965–978.
- [8] J. Luo and S. Wu, "Fedslid: Federated Learning with Shared Label Distribution for Medical Image Classification," *Proceedings - International Symposium on Biomedical Imaging*, 2022.
- [9] Z. Wang, Y. Zhu, D. Wang, and Z. Han, "FedACS: Federated Skewness Analytics in Heterogeneous Decentralized Data Environments," *2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQOS)*, 2021.
- [10] H. Lu, A. Thelen, O. Fink, C. Hu, and S. Laflamme, "Federated learning with uncertainty-based client clustering for fleet-wide fault diagnosis," *arXiv preprint arXiv:2304.13275*, 2023.
- [11] R. Song, L. Lyu, W. Jiang, A. Festag, and A. Knoll, "V2X-Boosted Federated Learning for Cooperative Intelligent Transportation Systems with Contextual Client Selection," *arXiv preprint arXiv:2305.11654*, 2023.
- [12] Z. Jiang, Y. Xu, H. Xu, Z. Wang, and C. Qian, "Adaptive control of client selection and gradient compression for efficient federated learning," *arXiv preprint arXiv:2212.09483*, 2022.
- [13] Y. Fraboni, R. Vidal, L. Kameni, and M. Lorenzi, "Clustered Sampling: Low-Variance and Improved Representativity for Clients Selection in Federated Learning," vol. 139, 2021.
- [14] C. Briggs, Z. Fan, and P. Andras, "Federated learning with hierarchical clustering of local updates to improve training on non-IID data," in *2020 Int. Joint Conf. on Neural Networks (IJCNN)*. IEEE, 2020, pp. 1–9.
- [15] M. Chen, Y. Wang, and H. V. Poor, "Performance optimization for wireless semantic communications over energy harvesting networks," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 8647–8651.
- [16] U. Sara, M. Akter, and M. S. Uddin, "Image quality assessment through FSIM, SSIM, MSE and PSNR—a comparative study," *Journal of Computer and Communications*, vol. 7, no. 3, pp. 8–18, 2019.
- [17] Ş. Öztürk, "Comparison of Pairwise Similarity Distance Methods for Effective Hashing," *IOP Conference Series: Materials Science and Engineering*, vol. 1099, 2021.
- [18] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2020.
- [19] M. Ghifary, W. B. Kleijn, and M. Zhang, "Domain adaptive neural networks for object recognition," in *PRICAI 2014: Trends in Artificial Intelligence: 13th Pacific Rim Int. Conf. on Artificial Intelligence, Gold Coast, QLD, Australia, December 1-5, 2014. Proceedings 13*. Springer, 2014, pp. 898–904.
- [20] F. Zhuang, X. Cheng, P. Luo, S. J. Pan, and Q. He, "Supervised representation learning: Transfer learning with deep autoencoders," in *Twenty-fourth international joint conference on artificial intelligence*, 2015.
- [21] W.-H. Chen, P.-C. Cho, and Y.-L. Jiang, "Activity recognition using transfer learning," *Sens. Mater.*, vol. 29, no. 7, pp. 897–904, 2017.
- [22] C.-Y. Lee, T. Batra, M. H. Baig, and D. Ulbricht, "Sliced wasserstein discrepancy for unsupervised domain adaptation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 10 285–10 295.
- [23] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, vol. 20, p. 53–65, 1987.
- [24] E. Guerra, F. Wilhelmi, M. Miozzo, and P. Dini, "The cost of training machine learning models over distributed data sources," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 1111–1126, 2023.
- [25] Y. LeCun and C. Cortes, "Mnist handwritten digit database," *AT&T Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, vol. 7, 2010.
- [26] V. Schmidt, K. Goyal, A. Joshi, B. Feld, L. Conell, N. Laskaris, D. Blank, J. Wilson, S. Friedler, and S. Luccioni, "Codecarbon: estimate and track carbon emissions from machine learning computing (2021)," Available: <https://github.com/mlco2/codecarbon>, 2021.
- [27] V. M. Panaretos and Y. Zemel, "Statistical Aspects of Wasserstein Distances," *Annual Review of Statistics and Its Application*, vol. 6, no. 1, pp. 405–431, 2019. [Online]. Available: <https://doi.org/10.1146%2Fannurev-statistics-030718-104938>