



HAL
open science

Clustered Sampling: Low-Variance and Improved Representativity for Clients Selection in Federated Learning

Yann Fraboni, Richard Vidal, Laetitia Kamani, Marco Lorenzi

► **To cite this version:**

Yann Fraboni, Richard Vidal, Laetitia Kamani, Marco Lorenzi. Clustered Sampling: Low-Variance and Improved Representativity for Clients Selection in Federated Learning. ICML 2021 - 38th International Conference on Machine Learning, Jul 2021, online, United States. hal-03232421

HAL Id: hal-03232421

<https://hal.science/hal-03232421>

Submitted on 21 May 2021

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.

Clustered Sampling: Low-Variance and Improved Representativity for Clients Selection in Federated Learning

Yann Fraboni^{1,2} Richard Vidal² Laetitia Kameni² Marco Lorenzi¹

Abstract

This work addresses the problem of optimizing communications between server and clients in federated learning (FL). Current sampling approaches in FL are either biased, or non optimal in terms of server-clients communications and training stability. To overcome this issue, we introduce *clustered sampling* for clients selection. We prove that clustered sampling leads to better clients representativity and to reduced variance of the clients stochastic aggregation weights in FL. Compatibly with our theory, we provide two different clustering approaches enabling clients aggregation based on 1) sample size, and 2) models similarity. Through a series of experiments in non-iid and unbalanced scenarios, we demonstrate that model aggregation through clustered sampling consistently leads to better training convergence and variability when compared to standard sampling approaches. Our approach does not require any additional operation on the clients side, and can be seamlessly integrated in standard FL implementations. Finally, clustered sampling is compatible with existing methods and technologies for privacy enhancement, and for communication reduction through model compression.

1. Introduction

Federated learning (FL) is a training paradigm enabling different clients to jointly learn a global model without sharing their respective data. Communication can be a primary bottleneck for FL since wireless and other end-user internet connections operate at variable communication rates while being potentially unreliable. Moreover, the capacity of the aggregating server may impose constraints on the number of clients the server can communicate with at the

same time. These considerations led to significant interest in reducing the number and bandwidth of communications at every step of the FL process.

One of the most popular communication reduction strategies consists in limiting the frequency of communications at the expense of increased computation on the clients side. This is usually achieved by asking the clients to perform multiple iterations of local gradient descent before communicating their updates. In this setting, FedAvg (McMahan et al., 2017) is the first and most widely used FL algorithm, for which convergence bounds were given in (Wang et al., 2020; Li et al., 2020; Karimireddy et al., 2020; Yu et al., 2019; Khaled et al., 2020; Woodworth et al., 2020; Lin et al., 2020; Stich, 2019).

To further reduce the number of communications, the server can select a subset of clients participating at every iteration. This strategy, called client sampling, enables reducing communications to the minimum. FedAvg first proposed selecting m clients uniformly without replacement while replacing the contribution of the non-sampled clients with the current global model. However this scheme is known for being biased, since the resulting model is, in expectation, different from the deterministic aggregation of every client. To overcome this issue, (Li et al., 2018) proposes an *unbiased* sampling scheme where the new global model is created as the average of the sampled clients work. The sampling is based on a multinomial distribution (MD) whose clients probabilities corresponds to their relative sample size. While other clients sampling schemes have been proposed, most of them require additional server-clients communications and are not proven to be unbiased (Li et al., 2020; Chen et al., 2020; Nishio & Yonetani, 2019).

To the best of our knowledge, FedAvg and MD sampling are the only schemes keeping to a minimum server-clients communications. In particular, MD sampling has been proven to lead to FL optimum and shown experimentally to outperform FedAvg sampling (Li et al., 2020). In spite of its unbiasedness, MD sampling may still lead to large variance in the clients selection procedure. While unbiasedness guarantees proper clients representativity in expectation, representativity is not necessarily achieved when considering a single FL iteration. Since at each MD sampling

¹Université Côte d’Azur, Inria, Epione Research Group, France

²Accenture Labs, Sophia Antipolis, France. Correspondence to: Yann Fraboni <yann.fraboni@inria.fr>.

instance we select clients with replacement, this determines a variance in the amount of times a client is selected. This sampling variance is a leading cause of the large variability in the convergence of FL, especially in non-iid applications. Indeed, at each iteration, sampled clients improve the global model based on their data distribution, to the detriment of the data specificity of non-sampled clients.

While the literature mainly focused on the study of the behavior of FL sampling strategies in expectation, to our knowledge this study provides the first theoretical investigation of the variability properties of FL sampling. In what follows, we show that this statistical aspect is crucial to determine convergence stability and quality of FL. The contribution of this work is the introduction of *clustered sampling*, a new unbiased client sampling scheme improving MD sampling by guaranteeing smaller client selection variability, while keeping to a minimum server-clients communications. By increasing every client representativity in model aggregations, clustered sampling ensures that clients with unique distributions are more likely of being sampled, leading to smoother and faster FL convergence.

We first derive, in Section 2, the theory behind current FL sampling schemes our work is built on. We then formally introduce clustered sampling in Section 3 and prove its theoretical correctness by extending the work done in (Wang et al., 2020). We finally show the theoretical benefits of clustered sampling over MD sampling. In Section 4, we propose an implementation of clustered sampling aggregating clients based on their sample size, showing that this approach leads to reduced variance of the clients’ aggregation weights. In Section 5, we extend our sampling theory to aggregation schemes based on the similarity between clients updates, showing that this approach further reduces the variance of clients aggregation weights while improving the representation of the clients during each FL aggregation step, as compared to MD sampling. This result leads to an overall improvement of the convergence of FL. Finally, in Section 6, we experimentally demonstrate our work on a broad range of balanced and unbalanced heterogeneous dataset. The code used for this work is available here¹.

2. Related Work

Before introducing in Section 3 the core idea of clustered sampling, we first recapitulate in Section 2 the current theory behind parameter aggregation and sampling schemes for FL.

2.1. Aggregating clients local updates

In FL, we consider a set I of clients respectively owning datasets \mathcal{D}_i composed of n_i samples. FL aims at optimizing

¹https://github.com/Accenture/LLabs-Federated-Learning/tree/clustered_sampling

the average of each clients local loss function weighted by their importance p_i

$$\mathcal{L}(\theta) = \sum_{i \in I} p_i \mathcal{L}_i(\theta), \quad (1)$$

where θ represents the model parameters and $\sum_{i=1}^n p_i = 1$. While any combination of $\{p_i\}$ is possible, a common choice consists in defining $p_i = n_i/M$, where $M = \sum_{i \in I} n_i$ is the total number of sample across datasets. In this work, we adopt the same definition of the importance weights, although the theory derived below does not depend on any specific choice of the parameters $\{p_i\}$.

In this setting, to estimate a global model across clients, FedAvg (McMahan et al., 2017) is an iterative training strategy based on the aggregation of local model parameters θ_i^t . At each iteration step t , the server sends the current global model parameters θ^t to the clients. Each client updates the model by minimizing the local cost function $\mathcal{L}(\theta_i^{t+1}, \mathcal{D}_i)$ through a fixed amount of SGD initialized with θ^t . Subsequently each client returns the updated local parameters θ_i^{t+1} to the server. The global model parameters θ^{t+1} at the iteration step $t + 1$ are then estimated as a weighted average, i.e.

$$\theta^{t+1} = \sum_{i \in I} \frac{n_i}{M} \theta_i^{t+1}. \quad (2)$$

2.2. Clients’ sampling

Clients sampling is a central operation of FL. FedAvg (McMahan et al., 2017) proposes to uniformly sample a subset of participating clients S_t at every iteration while the other clients updates are replaced by the current global model, i.e.

$$\theta^{t+1} = \sum_{i \in S_t} \frac{n_i}{M} \theta_i^{t+1} + \sum_{i \notin S_t} \frac{n_i}{M} \theta^t. \quad (3)$$

The sampling scheme introduced by FedAvg is generally slow due to the attrition introduced by non-participating clients. To solve this problem, (Li et al., 2018) proposes instead to sample S_t , the subset of clients at iteration t , from a Multinomial Distribution (MD) where each client is sampled according to its relative data ratio $p_i = \frac{n_i}{M}$. The new global model is obtained as the average of each selected client, i.e.

$$\theta^{t+1} = \sum_{i \in S_t} \frac{1}{m} \theta_i^{t+1}. \quad (4)$$

By design, MD sampling is such that the aggregation of clients model updates is identical in expectation to the one obtained when considering all the clients, i.e. $\mathbb{E}_{S_t} [\theta^{t+1}] = \sum_{i \in I} p_i \theta_i^{t+1}$. Sampling schemes following this property are called *unbiased*. Notably, the sampling scheme employed by FedAvg does not satisfy this property, and it is thus prone to clients-drift (Karimireddy et al., 2020).

2.3. FL convergence with MD client sampling

Theoretical guarantees regarding the convergence of FedAvg were given in (Wang et al., 2020). The proof relies on assumptions classically used in Stochastic Gradient Descent (SGD) analysis (Bottou et al., 2018) (Assumptions 1 and 2 below), or commonly used in the federated optimization literature (Li et al., 2018; 2020; Haddadpour et al., 2019; Karimireddy et al., 2020; Wang et al., 2019) to capture the dissimilarities of local objectives (Assumption 3 below).

Assumption 1 (Smoothness). *The clients local objective function is Lipschitz smooth, that is, $\|\nabla\mathcal{L}_i(x) - \nabla\mathcal{L}_i(y)\| \leq L\|x - y\|$, $\forall i \in \{1, \dots, n\}$.*

Assumption 2 (Unbiased Gradient and Bounded Variance). *For each client i local model, the stochastic gradient $g_i(x|\xi)$ of model x evaluated on batch ξ is an unbiased estimator of the local gradient: $\mathbb{E}_\xi[g_i(x|\xi)] = \nabla\mathcal{L}_i(x)$, and has bounded variance $\mathbb{E}_\xi[\|g_i(x|\xi) - \nabla\mathcal{L}_i(x)\|^2] \leq \sigma^2$, $\forall i \in \{1, \dots, n\}$ with $\sigma^2 \geq 0$.*

Assumption 3 (Bounded Dissimilarity). *For any set of weights $\{w_i \geq 0\}_{i=1}^n$ such that $\sum_{i=1}^n w_i = 1$, there exists constants $\beta^2 \geq 1$ and $\kappa^2 \geq 0$ such that $\sum_{i=1}^n w_i \|\nabla\mathcal{L}_i(x)\|^2 \leq \beta^2 \|\sum_{i=1}^n w_i \nabla\mathcal{L}_i(x)\|^2 + \kappa^2$. If all the local loss functions are identical, then we have $\beta^2 = 1$ and $\kappa^2 = 0$.*

The following theorem was proven in (Wang et al., 2020) and provides theoretical guarantees for MD client sampling.

Theorem 1. *Under Assumptions 1 to 3, and local learning rate $\eta = \sqrt{m/NT}$, FL with FedAvg when sampling m clients with MD converges to a stationary point of $\mathcal{L}(\theta)$:*

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\|\nabla\mathcal{L}(\theta^t)\|^2 \right] \leq \mathcal{O}\left(\frac{1}{\sqrt{mNT}}\right) + \mathcal{O}\left(\frac{mN}{T}\right). \quad (5)$$

The proof of Theorem 1 can be found in (Wang et al., 2020) and shows that considering a subset of workers with MD client sampling is enough to ensure convergence of the global model to a local minimum of the federated loss function, equation (1). Following the conclusions of that work, to avoid optimizing a surrogate loss function instead of the federated one in equation (1), the server asks from every client to compute the same amount of SGD steps N .

2.4. Sampling schemes comparison

Other client sampling schemes have been proposed. For example, with (Li et al., 2020), the server sends the global model to every client before creating the new global model out of the first m updated models the server receives; with (Chen et al., 2020), the server waits for every client to send the norm of their work before selecting the m clients with the most relevant updates; with (Nishio & Yonetani, 2019),

clients transmit information about their available computation resources before the server selects m of them in function of their availability. Contrarily to FedAvg and MD sampling, these sampling schemes require additional communications and sometimes even computation from all the clients. On the contrary, FedAvg and MD sampling have the appealing property of maintaining to a minimum the amount of clients-server communications at each iteration. The global model is sent only to the sampled clients, and the amount of local work is set for those clients to only N SGD updates. To the best of our knowledge, FedAvg and MD sampling schemes are the only approaches minimizing the effective work and communications asked to the clients. Moreover, while FedAvg sampling has no theoretical guarantees regarding its convergence, neither regarding the unbiasedness of its global model, MD sampling is shown to converge to a stationary point of the global loss function (1) (Theorem 1). Based on these considerations and given that (Li et al., 2020) shows experimentally that MD sampling outperforms FedAvg, in the rest of this work we consider MD sampling as reference sampling technique.

3. Clustered sampling

In Section 3.1, we first introduce clustered sampling and prove the convergence of FL under this scheme. In Section 3.2, we show the statistical improvements brought by clustered sampling as compared to MD client sampling, in terms of reduced sampling variance, and better clients representativity across the entire FL process.

3.1. Definition of clustered sampling

Let us consider n clients participating to FL. With MD sampling, m clients are sampled from a multinomial distribution supported on $\{1, \dots, n\}$ where a client is selected in function of its data ratio p_i .

Assumption 4 (Unbiased Sampling). *A client sampling scheme is said unbiased if the expected value of the client aggregation is equal to the global deterministic aggregation obtained when considering all the clients, i.e.*

$$\mathbb{E}_{S_t} [\theta^t] = \mathbb{E}_{S_t} \left[\sum_{j \in S_t} w_j(S_t) \theta_j^t \right] := \sum_{i=1}^n p_i \theta_i^t, \quad (6)$$

where $w_j(S_t)$ is the aggregation weight of client j for subset of clients S_t .

In (Li et al., 2020), the notion of unbiased sampling is introduced by means of Assumption 4. MD sampling follows this assumption, and thus provides at every iteration an unbiased global model. However, MD sampling enables a client to be sampled from 0 to m times with non-null probability at each iteration, giving aggregation weights for every client

ranging from 0 to 1. As a result, MD sampling provides appropriate representation for every client in expectation, with however potentially large variance in the amount of times a client is selected. As a consequence, the representativity of a client at any given realization of a FL iteration may not be guaranteed. In the following we introduce clustered sampling, and show that this strategy leads to decreasing clients aggregation weight variance and better clients representativity.

We denote by W_0 the multinomial distribution with support on $\{1, \dots, n\}$ used to sample one client according to its data ratio p_i . MD sampling can be seen as sampling m times with W_0 . With clustered sampling, we propose to generalize MD sampling by sampling m clients according to m independent distributions $\{W_k(t)\}_{k=1}^m$ each of them privileging a different subset of clients based on opportune selection criteria (Section 4 and 5). With clustered sampling, the m clients can be sampled with different distributions and, at two different iterations, the set of distributions can differ. MD sampling is a special case of clustered sampling when $\forall t, \forall k \in \{1, \dots, m\}, W_k(t) = W_0$.

In the rest of this work, we denote by $r_{k,i}^t$ the probability for client i to be sampled in distribution $W_k(t)$. By construction, we have:

$$\forall k \in \{1, \dots, m\}, \sum_{i=1}^n r_{k,i}^t = 1 \text{ with } r_{k,i}^t \geq 0. \quad (7)$$

We also require clustered sampling to be unbiased. Extending Assumption 4 to m independent sampling distributions $\{W_k(t)\}_{k=1}^m$, we obtain the property:

$$\forall i \in \{1, \dots, n\}, \sum_{k=1}^m r_{k,i}^t = m p_i. \quad (8)$$

Proposition 1. *Equations (7) and (8) are sufficient conditions for clustered sampling to satisfy Assumption 4.*

Proof. Satisfying equation (7) ensures the m distributions used with clustered sampling are feasible. When sampling one client from one of the m distributions $W_k(t)$, we get:

$$\mathbb{E}_{W_k(t)} \left[\sum_{j \in W_k(t)} w_j(W_k(t)) \theta_j^t \right] = \sum_{i=1}^n r_{k,i}^t \theta_i^t. \quad (9)$$

By linearity of the expected value, the expected new global model is the average between the weighted models obtained according to each distribution $\{W_k\}_{k=1}^m$ derived in equation (9), i.e.

$$\mathbb{E}_{S_t} [\theta^t] = \sum_{k=1}^m \frac{1}{m} \sum_{i=1}^n r_{k,i}^t \theta_i^t = \sum_{i=1}^n p_i \theta_i^t, \quad (10)$$

where the second equality comes from equation (8). \square

In Theorem 2, we prove that FedAvg with clustered sampling satisfying Assumptions 1 to 3 and Proposition 1 has the same convergence bound to a FL local optimum as with FedAvg and MD sampling. The proof of Theorem 2 can be found in Appendix A.

Theorem 2. *Under Assumption 1 to 3, and local learning rate $\eta = \sqrt{m/NT}$, let's consider FL with FedAvg when sampling m clients with clustered sampling scheme satisfying Proposition 1. The same asymptotic behavior of MD sampling holds :*

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} [\|\nabla \mathcal{L}(\theta^t)\|^2] \leq \mathcal{O}\left(\frac{1}{\sqrt{mNT}}\right) + \mathcal{O}\left(\frac{mN}{T}\right) \quad (11)$$

In the proof of Theorem 2 (Appendix A), we show that this convergence bound holds for any clustered sampling scheme satisfying Proposition 1. Moreover, the convergence bound of MD sampling is a bound itself for the convergence of a general clustered sampling scheme satisfying Proposition 1. Therefore, clustered sampling enjoys better convergence guarantees than MD sampling.

3.2. Improvements provided by clustered sampling

We introduced clustered sampling in Section 3.1 and showed that under the condition of Proposition 1 it provides the same convergence bound of MD sampling. In this section, we investigate the statistical benefits of clustered sampling with respect to MD sampling.

We define by $\omega_i(S)$ the aggregation weight of client i with subset of sampled clients S , and by S_{MD} and $S_C(t)$ the subset of clients sampled at iteration t with respectively MD and clustered sampling.

We consider a clustered sampling scheme following Proposition 1. Hence, for both MD and clustered sampling, the expected aggregation equals the deterministic aggregation when considering all the clients leading to:

$$\mathbb{E}_{S_{MD}(t)} [\omega_i(S_{MD})] = \mathbb{E}_{S_C(t)} [\omega_i(S_C(t))] = p_i. \quad (12)$$

With clustered sampling, we first show that every client has a smaller aggregation weight variance. A client's aggregation weight can be written as $\omega_i(S) = \frac{1}{m} \sum_{k=1}^m \{l_k = i\}$, where l_k is the index of the k^{th} sampled client. With MD sampling, m clients are iid sampled according to $\mathcal{B}(p_i)$, a Bernoulli distribution with probability p_i , giving the following variance:

$$\text{Var}_{S_{MD}} [\omega_i(S_{MD})] = \frac{1}{m^2} m \text{Var} [\mathcal{B}(p_i)] \quad (13)$$

$$= \frac{1}{m^2} m p_i (1 - p_i). \quad (14)$$

Clustered sampling instead selects independently m clients according to the distributions $\{W_k(t)\}_{k=1}^m$. Therefore, each

client is sampled according to $\mathcal{B}(r_{k,i}^t)$ giving:

$$\text{Var}_{S_C(t)} [\omega_i(S_C(t))] = \frac{1}{m^2} \sum_{k=1}^m \text{Var} [\mathcal{B}(r_{k,i}^t)] \quad (15)$$

$$= \frac{1}{m^2} \sum_{k=1}^m r_{k,i}^t (1 - r_{k,i}^t). \quad (16)$$

By the Cauchy-Schwartz inequality, one can prove that

$$\text{Var}_{S_{MD}} [\omega_i(S_{MD})] \geq \text{Var}_{S_C(t)} [\omega_i(S_C(t))], \quad (17)$$

with equality if and only if all the m distributions are equal to the one of MD sampling, i.e. $\forall k, W_k(t) = W_0$. Exact derivation is given in Appendix B. Therefore, with clustered sampling, every client has a smaller aggregation weight variance. Another interesting statistical measure of representativity is the probability for a client to be sampled, i.e. $\mathbf{P}(\{i \in S\})$. In particular, increasing the probability for every client to be sampled is mandatory to allow a proper representation of each client's data specificity in the global model, especially in heterogeneous setting, such as with non-iid and unbalanced clients. For this statistical measure, clustered sampling also provides better guarantees than MD sampling. With MD sampling, clients are iid sampled giving

$$p(i \in S_{MD}) = 1 - p(\{i \notin S_{MD}\}) \quad (18)$$

$$= 1 - p(\{i \notin W_0\})^m \quad (19)$$

$$= 1 - (1 - p_i)^m. \quad (20)$$

Similarly, with clustered sampling we get:

$$p(i \in S_C(t)) = 1 - \prod_{k=1}^m p(\{i \notin W_k(t)\}) \quad (21)$$

$$= 1 - \prod_{k=1}^m (1 - r_{k,i}^t). \quad (22)$$

Since we assume here that clustered sampling follows Proposition 1, from equation (8), and from the inequality of arithmetic and geometric means, we get:

$$p(\{i \in S_{MD}(t)\}) \leq p(\{i \in S_C(t)\}), \quad (23)$$

with equality if and only if all the m distributions are equal to the one of MD sampling, i.e. $\forall k, W_k(t) = W_0$ (derivation in Appendix B). Therefore, with clustered sampling, every client has an higher probability of being sampled and thus is better represented throughout the FL process.

In conclusion, clustered sampling reduces clients aggregation weights variance and increases their representativity. These results are important for FL applications with heterogeneous federated dataset. Increasing a client representativity ensures that clients with unique distributions are more likely of being sampled, and can potentially lead to smoother and faster FL convergence.

Algorithm 1 Clustered sampling based on sample size

- 1: **Input:** $\{n_i\}_{i=1}^n$ clients number of samples
 - 2: Order clients by descending importance of n_i .
 - 3: $k \leftarrow 1$ distribution index.
 - 4: $q \leftarrow 0$ sum of samples.
 - 5: $M \leftarrow \sum_{i=1}^n n_i$ total number of samples.
 - 6: **for** each client $i = 1$ **to** n **do**
 - 7: $q \leftarrow q + mn_i$
 - 8: $q = a_i M + b_i$ with a_i and b_i non negative integers
 - 9: **if** $a_i > k$ **then**
 - 10: $r'_{k,i} \leftarrow M - b_{i-1}$
 - 11: $\forall l \geq k + 1$ s.t. $(a_i - 1) - l \geq 0$, $r'_{k,i} \leftarrow M$
 - 12: **end if**
 - 13: $r'_{a_i,i} \leftarrow b_i$
 - 14: $k \leftarrow a_i$
 - 15: **end for**
 - 16: **Output:** Sampling probabilities $r_{k,i} = r'_{k,i}/M$.
-

4. Clustered sampling based on sample size

We introduced and showed the convergence of unbiased clustered sampling in Section 3. Clustered sampling schemes compatible with Proposition 1 are numerous, including MD sampling. In this section, we first provide an unbiased clustered sampling scheme based on the number of samples n_i owned by each client. The proposed scheme, illustrated in Algorithm 1, is compatible with Proposition 1. In particular, we have the following theorem:

Theorem 3. *Algorithm 1 outputs m distributions for a clustered sampling satisfying Proposition 1. The complexity of the algorithm is $\mathcal{O}(n \log(n))$.*

Proof. Algorithm 1 identifies the m distributions W_k by defining m sets q_k of cardinality M in which each client i is represented with probability $r_{k,i}$. The sets are constructed as follows. We define by $n'_i = mn_i$ the total number of samples to be allocated for each client. We thus have mM samples to allocate over the m sets q_k . For each client, the integer division $n'_i = Ma_i + b_i$, identifies a_i sets for which the client must be represented with probability 1. The remaining b_i samples are allocated to the remaining $m - \sum_i a_i$ sets. This is possible by observing that $Mm = \sum_i n'_i = M(\sum_i a_i) + \sum_i b_i$, and therefore $M(m - \sum_i a_i) = \sum_i b_i$. By construction, Proposition 1 is satisfied: $|q_k| = M$ implies equation (7), while equation (8) is met since, for each client, the total number of samples distributed across the sets q_k is $n'_i = mn_i$, and thus each client is represented with proportion mp_i across all the distributions. Algorithm 1 provides the practical implementation of this scheme.

The complexity of Algorithm 1 is derived in Appendix C, where we also provide a schematic illustration of the alloca-

Algorithm 2 Clustered sampling based on model similarity

- 1: **Input:** $\{n_i\}_{i=1}^n$ clients number of samples, $\{G_i\}_{i=1}^n$ clients representative gradient, m number of sampled clients, *clustering method* (e.g. Ward method), s similarity function (e.g. Arccos)
- 2: Estimated hierarchical clustering P with *clustering method* from similarity matrix ρ with $\rho_{i,j} = s(G_i, G_j)$.
- 3: Cut P to determine $K \geq m$ groups $\{B_k\}_{k=1}^K$. We define q_k as the total number of samples of the corresponding clients: $q_k = \sum_{i \in B_k} mn_i \leq M$.
- 4: Order the groups $\{B_k\}_{k=1}^K$ by decreasing q_k .
- 5: Define clients number of samples in the m distributions $\{W_k\}_{k=1}^m$ based on the ranking of $q_k : \forall k \leq m, \forall i \in B_k, r'_{k,i} \leftarrow mn_i$.
- 6: Create a set with the clients of the remaining groups $S = \{i, u_i = mn_i\}, \forall i \in B_{m+1} \cup \dots \cup B_K$
- 7: $k \leftarrow 1$ Start considering the first distribution W_k
- 8: **repeat**
- 9: Select first client i in S with u_i samples to allocate
- 10: Determine a_i and b_i the quotient and remainder of the euclidean division of $q_k + u_i$ by M
- 11: **if** $a_i = 0$ **then**
- 12: $r'_{k,i} \leftarrow b_i$ and i removed from S
- 13: **else**
- 14: $r'_{k,i} \leftarrow M - q_k$
- 15: $u_i \leftarrow u_i - r'_{k,i}$
- 16: Remove i from S if $u_i = 0$
- 17: $k \leftarrow k + 1$
- 18: **end if**
- 19: **until** $S = \emptyset$, Proposition 1 is satisfied
- 20: **Output:** Sampling probabilities $r_{k,i} = r'_{k,i}/M$.

tion procedure. \square

We note that with Algorithm 1 a client i can be sampled up to $\lfloor mp_i \rfloor + 2$ times. This is an improvement from MD sampling where clients can be instead sampled up to m times. Since clustering is performed according to the clients sample size n_i , unless n_i changes during the learning process, Algorithm 1 needs to be run only once at the beginning of the learning process, i.e. $\{W_k(t)\}_{k=1}^m = \{W_k\}_{k=1}^m$.

5. Clustered sampling based on similarity

We have shown, in Section 3, that unbiased clustered sampling is a generalization of MD sampling providing smaller aggregation weight variance for every client, and we proposed in Section 4 an algorithm to practically fulfill Proposition 1 to obtain m distributions grouping clients based on their number of samples n_i .

In this section we extend the approach of Section 3 to define a novel clustered sampling scheme where sampling distribu-

tions are defined based on the similarity across clients. In what follows we define clients similarity based on the measure of *representative gradient*. The representative gradient is the difference between a client's updated model and the global model. Comparing clients' representative gradients at a given iteration is shown to be an effective approach for detecting similarity between FL participants (Sattler et al., 2019).

Algorithm 2 adopts this concept to define a clustered sampling scheme compatible with Proposition 1. We have:

Theorem 4. *If for every client $p_i \leq 1/m$, Algorithm 2 outputs m distributions for a clustered sampling satisfying Proposition 1. The complexity of the algorithm is in $\mathcal{O}(n^2 d + X)$, where d is the number of parameters in the model, and X is the complexity of the clustering method.*

Proof. Algorithm 2 is similar to Algorithm 1, with the additional constraint that the number of clusters K can differ from the number m of distributions. If $K = m$, the clients are already allocated in q_k sets, and the same reasoning of Algorithm 1 can be applied. If $K > m$, we consider again the partitioning problem over m sets q_k of cardinality M . We define again by $n'_i = mn_i$ the total number of samples to be allocated for each client, and we have mM samples to allocate over the m sets q_k . Differently from Algorithm 1, we initialize the allocation with the clustering. In particular, we assign to each set q_k the n'_i samples of the clients included in cluster k . By construction, each of these sets q_k has cardinality $|q_k| \leq M$. We consider the m largest sets, and distribute in these sets the remaining samples of the $K - m$ clusters until $|q_k| = M$, for each k . By construction, this allocation is possible since we have mM total number of samples to be distributed across m sets of cardinality M . As for Algorithm 1, Proposition 1 is satisfied: $|q_k| = M$ implies equation (7), while equation (8) is met since, for each client, the total number of samples distributed across sets q_k is $n'_i = mn_i$. Appendix C completes the proof on the complexity of the algorithm. \square

As for Algorithm 1, Appendix C provides a schematic for a better illustration of the algorithm. Being a clustered sampling scheme, the variance of the clients aggregation weights of Algorithm 2 is bounded (equation (16)). Moreover, since the distributions are obtained from the similarity tree resulting from the representative gradients, this scheme explicitly promotes the sampling of clients based on their similarity. Finally, with Algorithm 2, the sampling from the distributions $\{W_k\}_{k=1}^m$ can be performed even when no representative gradient is available for the clients, for example if clients have not been sampled during FL yet. In this case we simply consider a constant 0 representative gradient for those clients, and thus group them together to promote their representativity in the same distribution.

We recall that Algorithm 2 does not require to share gradients across clients, but only the difference between local and global models (a.k.a. representative gradients). Thus, the communication cost is the same of standard FL while the privacy properties of FL privacy remain identical.

We emphasize that any valid hierarchical clustering algorithm can be used in Algorithm 2. Without loss of generality, in the rest of this work we consider the Ward hierarchical clustering method (Ward, 1963), which allows to obtain a similarity tree by minimizing at every node the variance of its depending clients. This method has complexity $\mathcal{O}(n^2 \log(n))$. We finally observe that the time complexity of Algorithm 2 is not necessarily an issue, even in presence of an important amount of clients. After aggregation of the new global model, the server can sample the clients, and transmit it to them. While waiting for their local work to be completed, the server can therefore estimate the new partitioning. In this way, Algorithm 2 is equivalent to MD sampling for what concerns the process of receiving the updated models, and transmitting the new global model to the clients.

As a final observation, while Algorithm 2 is originally designed for sampling scenarios where $p_i \leq 1/m$, with few modifications it can be also used for federated datasets composed of clients with larger sample size, i.e. when $I = \{i : p_i \geq 1/m\} \neq \emptyset$, or equivalently $I = \{i : mn_i \geq M\} \neq \emptyset$. In this case, we can simply allocate those clients in specific distributions, where they are sampled with probability 1. In total, we obtain $\lfloor m \frac{n_i}{M} \rfloor$ distributions of this kind. The remaining samples $mn_i - \lfloor m \frac{n_i}{M} \rfloor M < M$ will be then redistributed according to Algorithm 2.

6. Experiments

We first show on a standard classification problem on MNIST (LeCun et al., 1998), the advantages of clustered sampling obtained with Algorithm 1 and 2 with respect to MD sampling. We consider a fully connected network with one hidden layer of 50 nodes. We create a federated dataset composed of 100 clients where each one has 500 training and 100 testing samples composed by one digit only. Each digit is owned by 10 clients, every client has the same number of samples, and the server samples 10 clients at every iteration.

We note that an ideal clustering method for this FL problem consists in creating 10 clusters each containing the 10 clients with same classes. At each FL round, we should sample a client from each cluster in order to obtain a fair representation of all the digits in the model aggregations. We call ‘target’ sampling this ideal FL scenario. In practice, the server cannot adopt ‘target’ sampling as this requires to know the clients data distributions in advance. As we shall

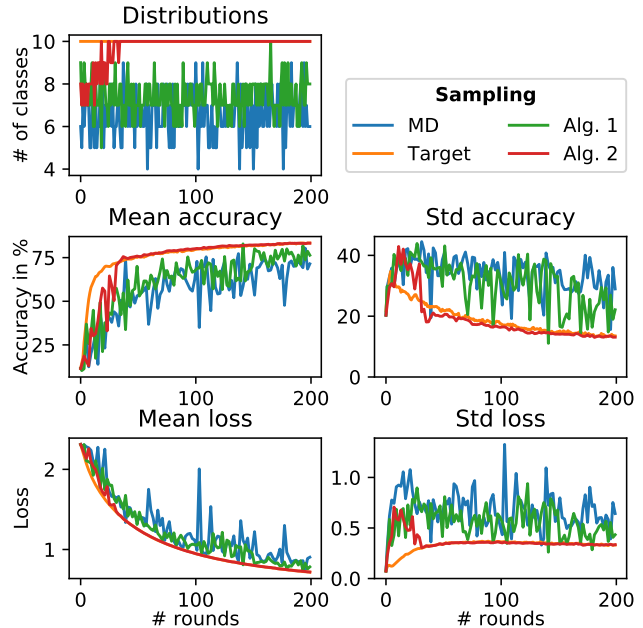


Figure 1. Comparison of MD sampling with clustered sampling of Algorithm 1 and 2 using cosine angle for the similarity measure. $n = 10$ clients from which $m = 10$ are sampled to perform $N = 50$ SGD with learning rate $lr = 0.01$ and batch size $B = 50$.

see in the rest of this section, the controlled nature of this example allows to clearly appreciate the practical benefits of clustered sampling.

We first show in Figure 1 that the FL processes obtained with Algorithm 1 and 2 both outperform MD sampling in terms of training global loss, testing accuracy, and representativity of the sampled classes. Moreover, we note that Algorithm 2 converges to the same ideal performances of ‘target’ sampling.

We also note that, with MD sampling, between 6 and 8 clients with different digits are generally chosen at each iteration round (Figure 1, top left panel). This is a practical demonstration of the sub-optimal representation of the clients heterogeneity. From a statistical perspective, with MD sampling the probability of sampling 10 different clients is $p = \frac{100!}{90!100!0} \sim 63\%$. Thus, for 37% of FL iterations, the new global model results from aggregation of less than 10 distinct clients. On the contrary, clustered sampling guarantees by construction that the aggregation will be always performed on 10 different clients. Indeed, since the dataset is balanced and the number of sampled clients $m = 10$ is a divider of the number of clients $n = 100$, every client can be allocated to one distribution only, and can be thus sampled up to once. Moreover, clustered sampling ensures that all the clients have identical aggregation weight variance. This improved data representation translates in less convergence variability at every iteration. Figure 1 il-

lustrates this result by showing noticeable improvements in terms of convergence with lower variance and better performance for training loss and testing accuracy. Moreover, with Algorithm 2, although at the early training steps some classes are not represented, the clustering strategy allows to quickly partition the 100 clients in 10 clusters and converge to the ideal distribution of ‘target’ (Figure 1, top left). As a consequence of this improved representativity of clients and classes, Algorithm 2 is associated with smoother and faster convergence processes for training loss and testing throughout iterations.

To demonstrate the benefits of clustered sampling beyond the controlled setting of MNIST, we conduct additional experiments on CIFAR10 (Krizhevsky, 2009) to investigate clustered sampling on more complex data distributions and models. CIFAR10 is composed of 32x32 images with three RGB channels of 10 different classes with 60000 samples. We use the same classifier of (McMahan et al., 2017) composed of 3 convolutional layers and 2 fully connected ones, including dropout after every convolutional layer.

To measure the influence of non-iid data distributions on the effectiveness of clustered sampling, we partition CIFAR10 using a Dirichlet distribution, $Dir(\alpha)$, giving to each client the respective partitioning across classes. The parameter α monitors the heterogeneity of the created dataset: $\alpha = 0$ assigns one class only to every client, while $\alpha \rightarrow +\infty$ gives a uniform partitioning of classes to each client. (Harry Hsu et al., 2019) provides graphical illustration of datasets obtained with such a process, and we provide in Appendix D similar illustrations for the parameters $\alpha \in \{0.001, 0.01, 0.1, 10\}$ considered in this work. To create an unbalanced federated dataset, we consider 100 clients where 10, 30, 30, 20 and 10 clients have respectively 100, 250, 500, 750, and 1000 training samples, and testing samples amounting to a fifth of their training size. All the clients consider a batch size of 50. For every CIFAR10 dataset partition, we report in this work experiments with learning rate in $\{0.001, 0.005, 0.01, 0.05, 0.1\}$ minimizing FedAvg with MD sampling training loss at the end of the learning process.

In Figure 2, we show how heterogeneity determines the improvements of clustered sampling over MD sampling for any of the four datasets. We note that the more heterogeneous a dataset is, i.e. the smaller α , the larger is the improvement provided by clustered sampling. Theorem 4 shows that clustered sampling has an identical bound as MD sampling. This is retrieved for $\alpha = 10$. and $\alpha = 0.1$ where the final performances for the two samplings are close with faster convergence for clustered sampling. However, with $\alpha = 0.01$ and $\alpha = 0.001$, clustered sampling provides faster and better convergence. Overall, with clustered sampling, the evolution of the training loss and testing accuracy

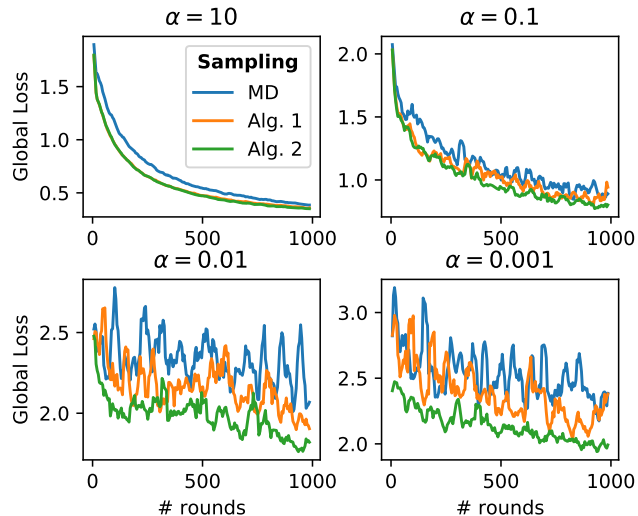


Figure 2. We investigate the improvement provided by clustered sampling on federated unbalanced datasets partitioned from CIFAR10 using a Dirichlet distribution with parameter $\alpha \in \{0.001, 0.01, 0.1, 10\}$. We use $N = 100$, $m = 10$, and respective learning rate for each dataset $lr = \{0.05, 0.05, 0.05, 0.1\}$.

are smoother processes than with MD sampling.

For sake of clarity, we note that the training losses displayed in Figure 2 is computed as the rolling mean over 50 iterations, while we provide in Appendix D the original training loss evolution. Furthermore, Appendix D reports a larger panel of experiments providing additional verification of the improvements brought by clustered sampling. In Figure 1 and 2, Algorithm 2 is computed with Arccos similarity. We show in Appendix D that with L2 and L1 we get similar improvements. We also show that increasing the amount of local work N enables clients to update models fitting better their data distribution. As a result, measuring clients similarity is easier, enabling better clustering, and leading to better performances. We also show that for any amount of sampled clients clustered sampling improves MD sampling. Finally, it is worth noticing that in none of the experimental settings considered for this paper clustered sampling underperformed with respect to MD sampling, providing further experimental evidence for our theoretical results.

7. Discussion and conclusion

In this work, we introduced clustered sampling, a novel client selection scheme in FL generalizing MD sampling, the current scheme from the state-of-the-art. We proved the correctness of clustered sampling and proposed two clustering methods implementing aggregation based on the clients number of samples, in Algorithm 1, or model similarity, in Algorithm 2. Both algorithms provide smaller weight variance for the clients aggregation process leading to better

client representativity. Consistently, clustered sampling is experimentally shown to have faster and smoother convergence in heterogeneous dataset.

The generality of clustered sampling paves the way to further investigation of clients clustering methods based on different criteria than clients sample size or model similarity. To the best of our knowledge, this work is also the first one introducing model similarity detection when sampling clients, as opposed to current approaches considering all clients at every iteration.

Finally, clustered sampling is an unbiased sampling scheme simple to implement, while not requiring to modify neither server nor clients behavior during FL training. This aspects makes clustered sampling readily compatible with existing methods and technologies for privacy enhancement and communication reduction.

Acknowledgments and Disclosure of Funding

This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002, and by the ANR JCJC project Fed-BioMed 19-CE45-0006-01. The project was also supported by Accenture. The authors are grateful to the OPAL infrastructure from Université Côte d’Azur for providing resources and support.

References

- Bottou, L., Curtis, F. E., and Nocedal, J. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018. ISSN 00361445. doi: 10.1137/16M1080173.
- Chen, W., Horvath, S., and Richtarik, P. Optimal Client Sampling for Federated Learning. *Workshop in NeurIPS: Privacy Preserving Machine Learning*, 2020.
- Haddadpour, F., Kamani, M. M., Mahdavi, M., and Cadambe, V. R. Local SGD with periodic averaging: Tighter analysis and adaptive synchronization. *Advances in Neural Information Processing Systems*, 32(2), 2019. ISSN 10495258.
- Harry Hsu, T. M., Qi, H., and Brown, M. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv*, 2019.
- Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., and Suresh, A. T. SCAFFOLD: Stochastic controlled averaging for federated learning. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 5132–5143. PMLR, 13–18 Jul 2020.
- Khaled, A., Mishchenko, K., and Richtarik, P. Tighter theory for local sgd on identical and heterogeneous data. In Chiappa, S. and Calandra, R. (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 4519–4529. PMLR, 26–28 Aug 2020.
- Krizhevsky, A. Learning multiple layers of features from tiny images. 2009.
- LeCun, Y., Bottou, L., Bengio, Y., and Ha, P. LeNet. *Proceedings of the IEEE*, (November):1–46, 1998. ISSN 00189219. doi: 10.1109/5.726791.
- Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., and Smith, V. Federated Optimization in Heterogeneous Networks. *Proceedings of the 1st Adaptive & Multitask Learning Workshop, Long Beach, California, 2019*, pp. 1–28, 2018.
- Li, X., Huang, K., Yang, W., Wang, S., and Zhang, Z. On the convergence of fedavg on non-iid data. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.
- Lin, T., Stich, S. U., Patel, K. K., and Jaggi, M. Don’t use large mini-batches, use local sgd. In *International Conference on Learning Representations*, 2020.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Singh, A. and Zhu, J. (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR.
- Nishio, T. and Yonetani, R. Client selection for federated learning with heterogeneous resources in mobile edge. In *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, pp. 1–7, 2019. doi: 10.1109/ICC.2019.8761315.
- Sattler, F., Müller, K.-R., and Samek, W. Clustered Federated Learning: Model-Agnostic Distributed Multi-Task Optimization under Privacy Constraints. pp. 1–16, 2019. URL <http://arxiv.org/abs/1910.01991>.
- Stich, S. U. Local SGD converges fast and communicates little. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=S1g2JnRcFX>.

Wang, J., Sahu, A. K., Yang, Z., Joshi, G., and Kar, S. Matcha: Speeding up decentralized sgd via matching decomposition sampling. In *2019 Sixth Indian Control Conference (ICC)*, pp. 299–300, 2019. doi: 10.1109/ICC47138.2019.9123209.

Wang, J., Liu, Q., Liang, H., Joshi, G., and Poor, H. V. Tackling the objective inconsistency problem in heterogeneous federated optimization. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

Ward, J. H. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, 1963. ISSN 01621459.

Woodworth, B., Patel, K. K., Stich, S., Dai, Z., Bullins, B., McMahan, B., Shamir, O., and Srebro, N. Is local SGD better than minibatch SGD? In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 10334–10343. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/woodworth20a.html>.

Yu, H., Yang, S., and Zhu, S. Parallel Restarted SGD with Faster Convergence and Less Communication: Demystifying Why Model Averaging Works for Deep Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):5693–5700, 2019. doi: 10.1609/aaai.v33i01.33015693.

A. Proof of Theorem 2

Theoretical guarantees regarding the convergence of FedAvg were given in (Wang et al., 2020). The proof relies on Assumptions 1 to 3. The full proof is provided in (Wang et al., 2020) for MD sampling where the MD sampling is shown to satisfy Lemma 5. In Section A.1, we reproduce the proof provided in (Wang et al., 2020) for Lemma 5 and, in Section A.2, we show that clustered sampling satisfying Proposition 1 also satisfies Lemma 5. As a result, FedAvg when sampling clients with MD or clustered sampling has identical asymptotic behavior.

Lemma 5. *Suppose we are given $z_1, z_2, \dots, z_n, x \in \mathbb{R}^d$. Let l_1, l_2, \dots, l_m be the index of the sampled clients and S be the set of sampled clients. We have*

$$\mathbb{E}_S \left[\frac{1}{m} \sum_{j=1}^m z_{l_j} \right] = \sum_{i=1}^n p_i z_i, \quad (24)$$

and

$$\begin{aligned} \mathbb{E}_S \left[\left\| \frac{1}{m} \sum_{j=1}^m z_{l_j} \right\|^2 \right] &\leq 3 \sum_{i=1}^n p_i \|z_i - \nabla \mathcal{L}_i(x)\|^2 \\ &\quad + 3 \|\nabla \mathcal{L}(x)\|^2 + \frac{3}{m} (\beta^2 \|\nabla \mathcal{L}(x)\|^2 + \kappa^2). \end{aligned} \quad (25)$$

A.1. Proof of Lemma 5 for Theorem 1 adapted from (Wang et al., 2020)

Proof. Clients are selected with MD sampling. We denote by l_1, l_2, \dots, l_m the m indices of the sampled clients which are iid sampled from a multinomial distribution supported on $\{1, \dots, n\}$ satisfying $\mathbb{P}(l_x = i) = p_i$ and $\sum_{i=1}^n p_i = 1$.

By definition, MD sampling satisfies equation (24).

Regarding equation (25), we have:

$$\begin{aligned} \frac{1}{m} \sum_{j=1}^m z_{l_j} &= \left(\frac{1}{m} \sum_{j=1}^m z_{l_j} - \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_{l_j}(x) \right) \\ &\quad + \left(\frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_{l_j}(x) - \nabla \mathcal{L}(x) \right) + \nabla \mathcal{L}(x). \end{aligned} \quad (26)$$

Using the Jensen inequality on the $\|\cdot\|^2$ operator, we get:

$$\begin{aligned} \mathbb{E} \left[\left\| \frac{1}{m} \sum_{j=1}^m z_{l_j} \right\|^2 \right] &\leq 3 \mathbb{E} \left[\left\| \frac{1}{m} \sum_{j=1}^m (z_{l_j} - \nabla \mathcal{L}_{l_j}(x)) \right\|^2 \right] \\ &\quad + 3 \mathbb{E} \left[\left\| \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_{l_j}(x) - \nabla \mathcal{L}(x) \right\|^2 \right] \\ &\quad + 3 \|\nabla \mathcal{L}(x)\|^2 \end{aligned} \quad (27)$$

Using the Jensen inequality, we get the following upper bound for the first term:

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{m} \sum_{j=1}^m (z_{l_j} - \nabla \mathcal{L}_{l_j}(x)) \right\|^2 \right] \\ & \leq \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m \|z_{l_j} - \nabla \mathcal{L}_{l_j}(x)\|^2 \right] \end{aligned} \quad (28)$$

$$= \sum_{i=1}^n p_i \|z_i - \nabla \mathcal{L}_i(x)\|^2, \quad (29)$$

where the equality follows from equation (24).

By definition, MD sampling is unbiased, i.e. $\mathbb{E}[\nabla \mathcal{L}_{l_j}(x)] = \nabla \mathcal{L}(x)$. Therefore, we get the following upper bound for the second term:

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_{l_j}(x) - \nabla \mathcal{L}(x) \right\|^2 \right] \\ & = \mathbb{E} \left[\frac{1}{m^2} \sum_{j=1}^m \|\nabla \mathcal{L}_{l_j}(x) - \nabla \mathcal{L}(x)\|^2 \right] \end{aligned} \quad (30)$$

$$= \frac{1}{m} \sum_{i=1}^n p_i \|\nabla \mathcal{L}_i(x) - \nabla \mathcal{L}(x)\|^2 \quad (31)$$

$$= \frac{1}{m} \sum_{i=1}^n p_i \|\nabla \mathcal{L}_i(x)\|^2 - \frac{1}{m} \|\nabla \mathcal{L}(x)\|^2 \quad (32)$$

$$\leq \frac{1}{m} [(\beta^2 - 1) \|\nabla \mathcal{L}(x)\|^2 + \kappa^2] \quad (33)$$

$$\leq \frac{1}{m} [\beta^2 \|\nabla \mathcal{L}(x)\|^2 + \kappa^2], \quad (34)$$

where the first inequality comes from using Assumption 3.

Finally, substituting equation (29) and (34) in equation (27) completes the proof. \square

A.2. Proof of Lemma 5 for Theorem 2

Proof. Clients are selected with clustered sampling. The m clients indices l_1, l_2, \dots, l_m are still independently sampled but no longer identically. Each index l_k is sampled from a distribution W_k . Each client can be sampled with probability $\mathbb{P}(l_k = i) = r_{k,i}$.

Clustered sampling follows Proposition 1 and therefore satisfies equation (24).

Equation (27) holds for any sampling schemes. Therefore, we also use it to prove equation (25) for clustered sampling. Using the same steps as for the proof of Lemma 5 for MD

sampling, we bound the first term of equation (27) as:

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{m} \sum_{j=1}^m (z_{l_j} - \nabla \mathcal{L}_{l_j}(x)) \right\|^2 \right] \\ & \leq \sum_{i=1}^n p_i \|z_i - \nabla \mathcal{L}_i(x)\|^2. \end{aligned} \quad (35)$$

Before bounding the second term, we define $\nabla \mathcal{L}_{W_k}(x)$ as the expected gradient of the distribution W_k with respects to the parameters x , i.e.

$$\nabla \mathcal{L}_{W_k}(x) := \mathbb{E}_{l_k \sim W_k} [\nabla \mathcal{L}_{l_k}(x)] = \sum_{i=1}^n r_{k,i} \nabla \mathcal{L}_i(x) \quad (36)$$

Using this definition, we bound the second term as

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{m} \sum_{k=1}^m \nabla \mathcal{L}_{l_k}(x) - \nabla \mathcal{L}(x) \right\|^2 \right] \\ & = \mathbb{E} \left[\left\| \frac{1}{m} \sum_{k=1}^m (\nabla \mathcal{L}_{l_k}(x) - \nabla \mathcal{L}_{W_k}(x)) \right\|^2 \right] \end{aligned} \quad (37)$$

$$= \frac{1}{m^2} \sum_{k=1}^m \mathbb{E} \left[\|\nabla \mathcal{L}_{l_k}(x) - \nabla \mathcal{L}_{W_k}(x)\|^2 \right] \quad (38)$$

$$= \frac{1}{m^2} \sum_{k=1}^m \sum_{i=1}^n r_{k,i} \|\nabla \mathcal{L}_i(x) - \nabla \mathcal{L}_{W_k}(x)\|^2 \quad (39)$$

$$= \frac{1}{m^2} \left[\sum_{i=1}^n m p_i \|\nabla \mathcal{L}_i(x)\|^2 - \sum_{k=1}^m \|\nabla \mathcal{L}_{W_k}(x)\|^2 \right] \quad (40)$$

$$\leq \frac{1}{m} [\beta^2 \|\nabla \mathcal{L}(x)\|^2 + \kappa^2], \quad (41)$$

where the last inequality comes from using Assumption 3 and equation (38) and (40) are obtained with equation (36).

Finally, substituting equation (35) and (41) in equation (27) completes the proof. \square

Equation (32) and (40) allow us to theoretically identify the convergence improvement of clustered sampling over MD sampling.

We define by $B_{MD} = \frac{1}{m} \sum_{i=1}^n p_i \|\nabla \mathcal{L}_i(x)\|^2 - \frac{1}{m} \|\nabla \mathcal{L}(x)\|^2$, equation (32), and $B_{Cl} = \frac{1}{m} \sum_{i=1}^n p_i \|\nabla \mathcal{L}_i(x)\|^2 - \frac{1}{m^2} \sum_{k=1}^m \|\nabla \mathcal{L}_{W_k}(x)\|^2$,

equation (40). Using the Jensen inequality, we get

$$-\sum_{k=1}^m \frac{1}{m^2} \|\nabla \mathcal{L}_{W_k}(x)\|^2 \leq -\frac{1}{m} \left\| \sum_{k=1}^m \frac{1}{m} \nabla \mathcal{L}_{W_k}(x) \right\|^2 \quad (42)$$

$$= -\frac{1}{m} \|\nabla \mathcal{L}(x)\|^2 \quad (43)$$

with equality if and only if $\forall k, l, \nabla \mathcal{L}_{W_k}(x) = \nabla \mathcal{L}_{W_l}(x)$. Thus, $B_{Cl} \leq B_{MD}$ with equality if and only if all the clients have the same data distribution or the considered clustered sampling is MD sampling.

B. MD and clustered sampling comparison

B.1. Client aggregation weight variance

As in Section 3, we denote by S_{MD} and $S_C(t)$ the random variables associated respectively to MD and clustered sampling. Also in Section 3, we have shown that

$$\text{Var}_{S_{MD}}[\omega_i(S_{MD})] = \frac{1}{m^2} mp_i(1-p_i), \quad (44)$$

and

$$\text{Var}_{S_C(t)}[\omega_i(S_C(t))] = \frac{1}{m^2} \sum_{k=1}^m r_{k,i}^t (1-r_{k,i}^t). \quad (45)$$

Hence, we get:

$$\text{Var}_{S_{MD}}[\omega_i(S_{MD})] - \text{Var}_{S_C(t)}[\omega_i(S_C(t))] \quad (46)$$

$$= \frac{1}{m^2} [mp_i(1-p_i) - \sum_{k=1}^m r_{k,i}^t (1-r_{k,i}^t)] \quad (47)$$

We consider an unbiased clustered sampling. Therefore, the sum of probability for client i over the m clusters satisfies $\sum_{k=1}^m r_{k,i}^t = mp_i$ giving:

$$\text{Var}_{S_{MD}}[\omega_i(S_{MD})] - \text{Var}_{S_C(t)}[\omega_i(S_C(t))] \quad (48)$$

$$= \frac{1}{m^2} [\sum_{k=1}^m r_{k,i}^t{}^2 - mp_i^2] \quad (49)$$

Using the Cauchy-Schwartz inequality, we get: $\sum_{k=1}^m r_{k,i}^t{}^2 \times \sum_{k=1}^m 1^2 \geq \left(\sum_{k=1}^m r_{k,i}^t \times 1\right)^2 = (mp_i)^2$ due to the unbiased aspect of the considered clustered sampling. As such, we get:

$$\text{Var}_{S_{MD}}[\omega_i(S_{MD})] - \text{Var}_{S_C(t)}[\omega_i(S_C(t))] \geq 0, \quad (50)$$

with equality if and only if $r_{k,i}^t = p_i$.

B.2. Probability for a client to be sampled at least once

In Section 3, we have shown that

$$p(\{i \in S_{MD}\}) = 1 - (1-p_i)^m \quad (51)$$

and

$$p(\{i \in S_C(t)\}) = 1 - \prod_{k=1}^m (1-r_{k,i}^t). \quad (52)$$

Hence, we get:

$$p(\{i \in S_{MD}\}) - p(\{i \in S_C(t)\}) \quad (53)$$

$$= \prod_{k=1}^m (1-r_{k,i}^t) - (1-p_i)^m \quad (54)$$

We consider an unbiased clustered sampling. Therefore, when using the inequality of arithmetic and geometric means, we get:

$$\prod_{k=1}^m (1-r_{k,i}^t) \leq \left(\frac{\sum_{k=1}^m (1-r_{k,i}^t)}{m}\right)^m = (1-p_i)^m, \quad (55)$$

with equality if and only if $r_{k,i}^t = p_i$. Finally, we get:

$$p(\{i \in S_{MD}\}) \geq p(\{i \in S_C(t)\}) \quad (56)$$

C. Explaining Algorithm 1 and 2

Algorithms 1 and 2 can be written in term of data ratio p_i instead of samples number n_i . While in both cases the algorithms would be correct, it turns out to be simpler to work with quantities of samples $n_i = p_i M$ instead which are integers. Therefore, without loss of generality, we denote by $r'_{k,i}$ the number of samples allocated by client i to distribution k . We retrieve the sampling probability of client i in distribution W_k with $r_{k,i} = \frac{r'_{k,i}}{M}$.

Also, without loss of generality, we prove Algorithms 1 and 2 at iteration t and therefore we use in the proofs $r_{k,i}$ and W_k instead of $r'_{k,i}$ and W_k^t .

C.1. Algorithm 1

We illustrate in Figure 3 the clients allocation scheme of Algorithm 1 introduced in Section 4, by considering how a client i is associated to the m distributions. Theorem 3 states that Algorithm 1 provides a sampling scheme satisfying Proposition 1 with complexity $\mathcal{O}(n \log(n))$ which we prove in Section 4 and in the following proof.

Proof. In term of complexity, the while loop for the client allocation, as illustrated in Figure 3, either change client or distribution at every step and is thus done in complexity $\mathcal{O}(n+m)$. Sampling client is relevant if $m < n$. Therefore the allocation complexity is equivalent to $\mathcal{O}(n+m) = \mathcal{O}(n)$. Also, sorting n elements is done in complexity $\mathcal{O}(n \log(n))$. Therefore, Algorithm 1 overall complexity is $\mathcal{O}(n \log(n))$. \square

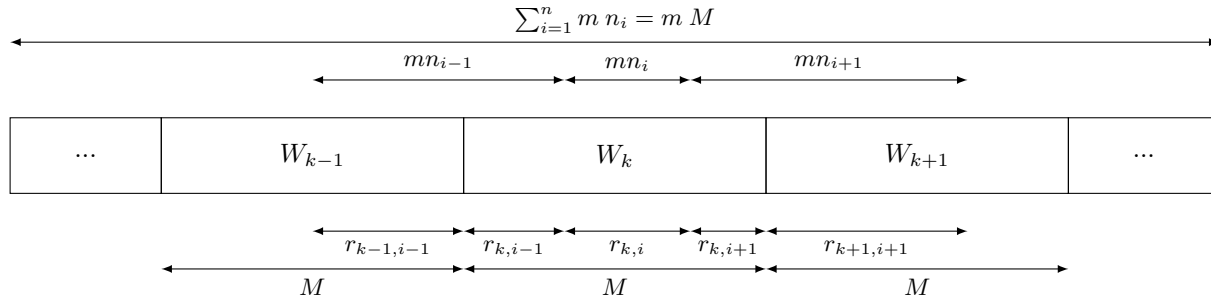


Figure 3. Illustration of the clients allocation scheme of Algorithm 1. Clients are considered in decreasing importance of their number of samples and always allocate client samples to distributions that already received samples but do not yet have M of them. As a result, after allocating a client, all distributions except at most one have 0 or M samples. Client i is only sampled in W_k because every distribution with index inferior to k are filled with clients of index inferior to i , and because there is enough room in W_k to receive all the samples that need to be allocated for client i .

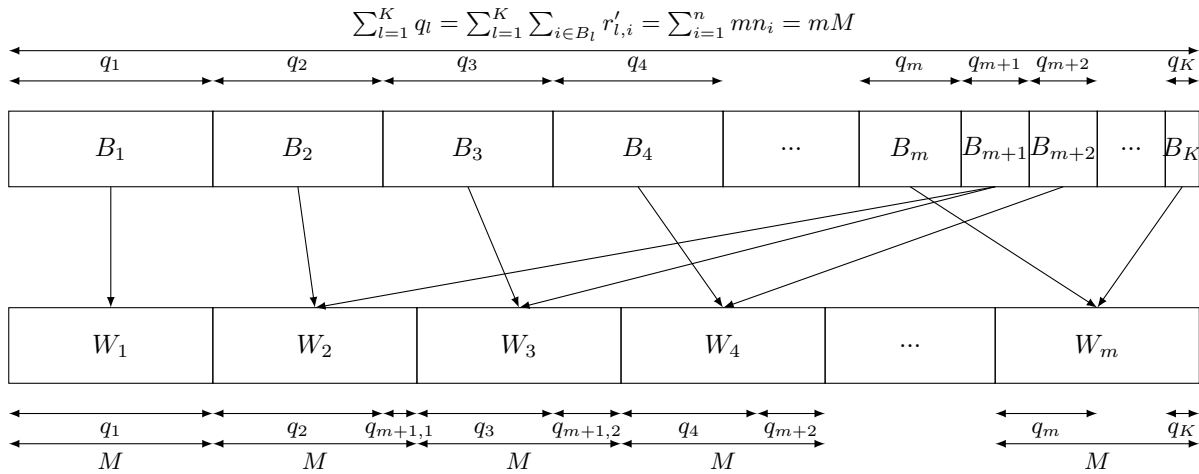


Figure 4. Illustration of the clients allocation scheme of Algorithm 2. After the tree is split in K groups of clients, the groups are ordered and we consider without loss of generality that their number of samples are inversely proportional to their index. With Algorithm 2, the first m groups, i.e. B_1 to B_m , are each associated to one distribution, i.e. W_1 to W_m . The remaining groups are considered one after the other and split among the remaining slots in the groups. Each distribution has M samples from clients participating to the FL process.

C.2. Algorithm 2

We illustrate in Figure 4, the clients allocation scheme of Algorithm 2 introduced in Section 5 by considering how a client i is associated to the m distributions. Theorem 4 states that Algorithm 2 provides a sampling scheme satisfying Proposition 1 and takes time complexity $\mathcal{O}(n^2d + X)$. We prove these statements in Section 5 and the following proof.

Proof. With identical reasoning as for Algorithm 1, clients are allocated in complexity $\mathcal{O}(n)$. Computing the similarity between two clients requires d elementary operations, where d is the number of parameters in the model, and has thus complexity $\mathcal{O}(d)$. Computing the similarity matrix requires computing $\frac{n(n-1)}{2}$ client similarities and thus has total complexity $\mathcal{O}(n^2d)$. Computing the similarity tree depends on the *clustering method* which we consider has complexity $\mathcal{O}(X)$. Transforming the tree as discussed in Section 5 requires going through its $n - 1$ nodes and thus has time complexity $\mathcal{O}(n)$. Cutting the tree requires considering at most every nodes and has thus complexity $\mathcal{O}(n)$. Lastly, the tree is cut in at most n branches and sorting them takes therefore complexity $\mathcal{O}(n \log(n))$. Finally, combining all these time complexities gives for Algorithm 2 a time complexity of $\mathcal{O}(n^2d + X)$. \square

In practice, the m distributions are computed at every iteration, while the server is required to compute the similarity between sampled clients and all the other clients. Therefore the similarity matrix can be estimated in complexity $\mathcal{O}(nmd)$, and Algorithm 2 has complexity $\mathcal{O}(nmd + X)$.

D. Additional experiments

We describe in Section 6 the different datasets used for the experiments and how we use the Dirichlet distribution to partition CIFAR10 in realistic heterogeneous federated datasets. In all the experiments, we consider a batch size of 50. For every CIFAR10 dataset partition, the learning rate is selected in $\{0.001, 0.005, 0.01, 0.05, 0.1\}$ to minimize FedAvg with MD sampling training loss.

D.1. CIFAR10 partitioning illustration

In Figure 5, we show the influence of α on the resulting federated dataset heterogeneity. $\alpha = 10$ provides almost an iid dataset and identical class percentages, column (a), and same number of samples per class, column (b). With $\alpha = 0.001$, we get a very heterogeneous dataset with almost only one class per client translating into some classes much more represented than others due to the unbalanced nature aspect of the created federated dataset, cf Section 6.

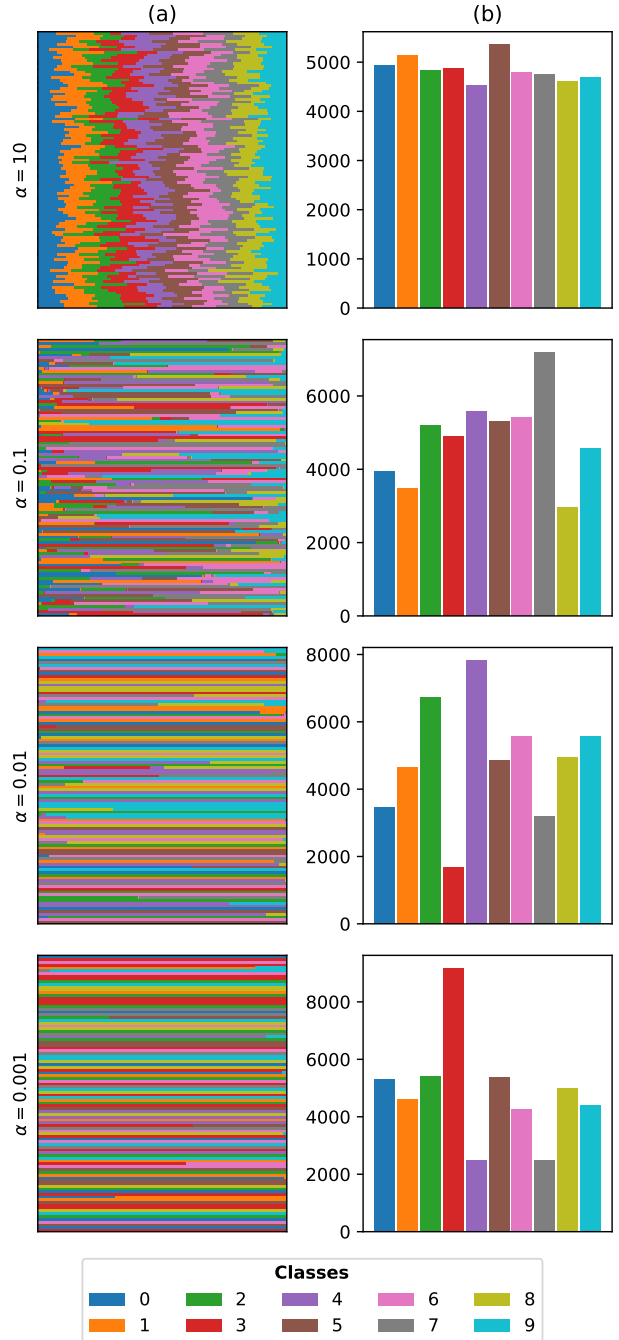


Figure 5. Effect of α on the resulting clients partitioning when using a Dirichlet distribution. Plots in column (a) represent the percentage of each class owned by the clients. Plots in column (b) give for every class its total number of samples across clients. We consider in this work $\alpha \in \{0.001, 0.01, 0.1, 10\}$.

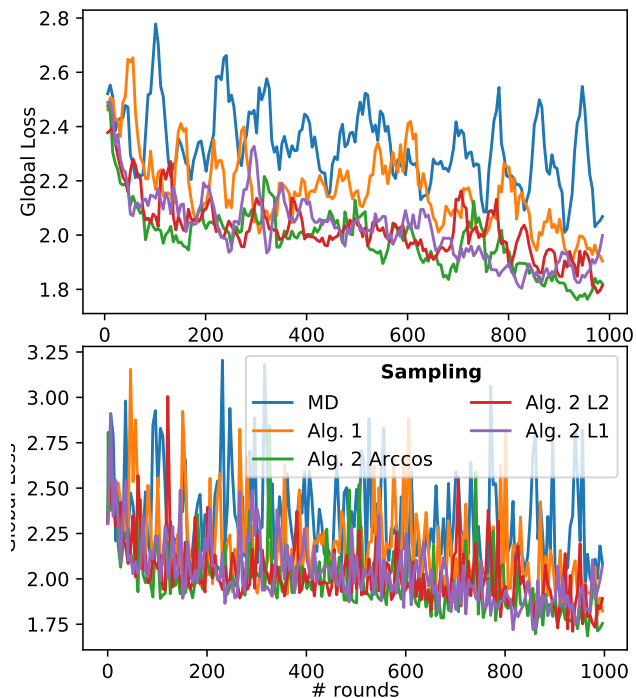


Figure 6. Effect of the similarity measure chosen for Algorithm 2 on the training loss convergence. We consider the evolution of the global loss, equation (1), in function of the server iteration t . For clarity concerns, we plot the global loss obtained with rolling mean over 50 server iterations (top) and the raw global loss (bottom). We consider CIFAR partitioned with $\text{Dir}(\alpha = 0.01)$, learning rate $lr = 0.05$, $N = 100$ SGD, and $m = 10$ sampled clients.

D.2. Influence of the similarity measure

Figure 6 shows the effect similarity measures (Arccos, L2, and L1) have on training global loss convergence. We retrieve that Algorithm 1 outperforms MD sampling by reducing clients aggregation weight variance. We remind that the hierarchical tree is obtained using Ward’s method in this work. We notice that the tree similarity measures gives similar performances when using Algorithm 2 with Ward hierarchical clustering method.. This justifies the use of Arccos similarity for the other experiments.

D.3. More details on Figure 2

For sake of clarity, we note that the training loss displayed in Figures 2 is computed as the rolling mean over 50 iterations. In Figure 7, we provide the raw training global loss with the testing accuracy at every server iteration.

D.4. Influence of m the number of sampled clients, and N the number of SGD run

We also investigates the influence the number of sampled clients m and the number of SGD run N have on the FL

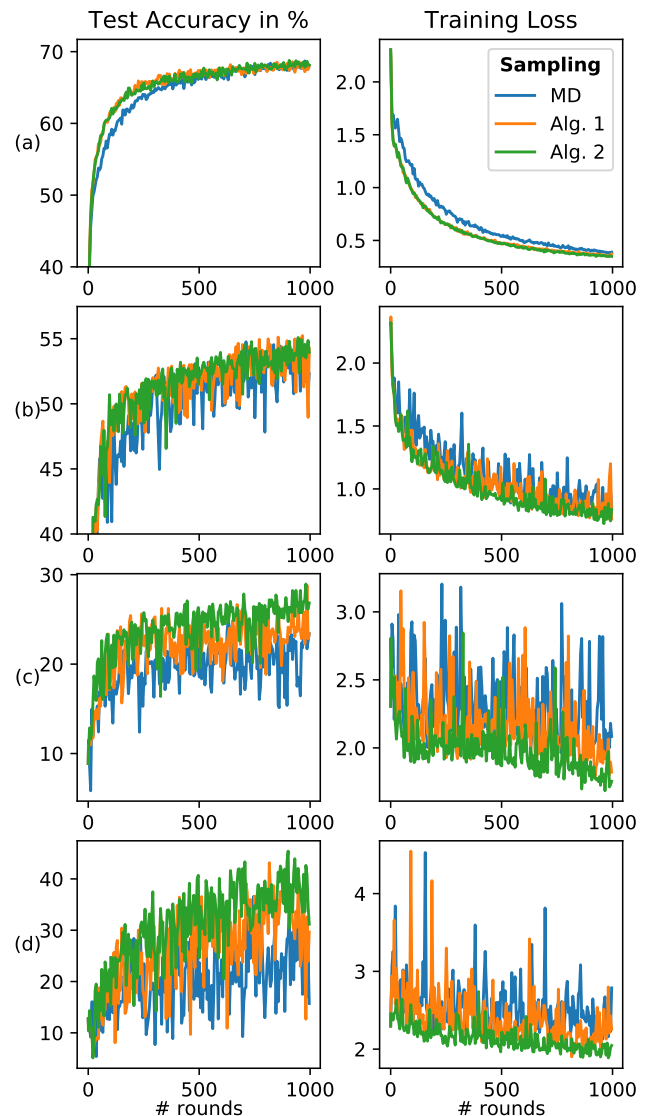


Figure 7. We investigate the improvement provided by clustered sampling on federated unbalanced datasets partitioned from CIFAR10 using a Dirichlet distribution with parameter $\alpha \in \{0.001, 0.01, 0.1, 10\}$ for respective row (a), (b), (c), (d). We use $N = 100$, $m = 10$, and respective learning rate for each dataset $lr = \{0.05, 0.05, 0.05, 0.1\}$.

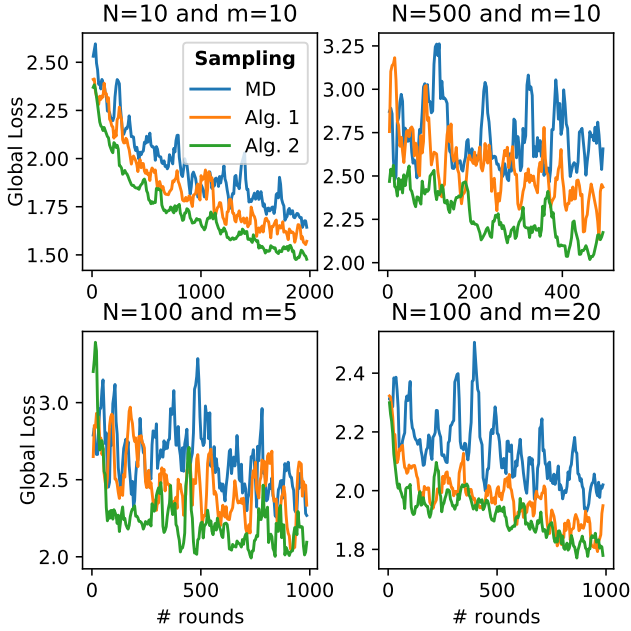


Figure 8. We consider the federated dataset partitioned from CIFAR10 using a Dirichlet distribution with parameter $\alpha = 0.01$. We investigate the influence of N , the number of SGD run by each client, and m , the number of sampled clients, on the training loss convergence. For each plot, experiments in first row use respectively $lr = \{0.1, 0.05\}$ and for second row $lr = \{0.05, 0.05\}$.

convergence speed and smoothness in Figure 8. We notice that the more important the amount of local work N is, and the faster clustered sampling convergence speed is. With more local work, clients better fit their data. In non-iid dataset this translates in more forgetting on the classes and samples which are not part of the sampled clients. Regarding the amount of sampled clients m , we notice that with a smaller amount of sampled clients the improvement of clustered sampling over MD sampling is more important. We associate this result to the better data representativity of clustered sampling. For the same reason, when we increase the number of sampled clients, we see faster convergence for both MD and clustered sampling. The performance of clustered sampling is closer but still better than the one of MD sampling.

For sake of clarity, we note that the training loss displayed in Figures 8 is computed as the rolling mean over 50 iterations. In Figure 7, we provide the raw training global loss with the testing accuracy at every server iteration.

D.5. Local regularization

With FedProx (Li et al., 2018), every client’s local loss function is equipped with a regularization term forcing the updated model to stay close to the current global model, i.e.

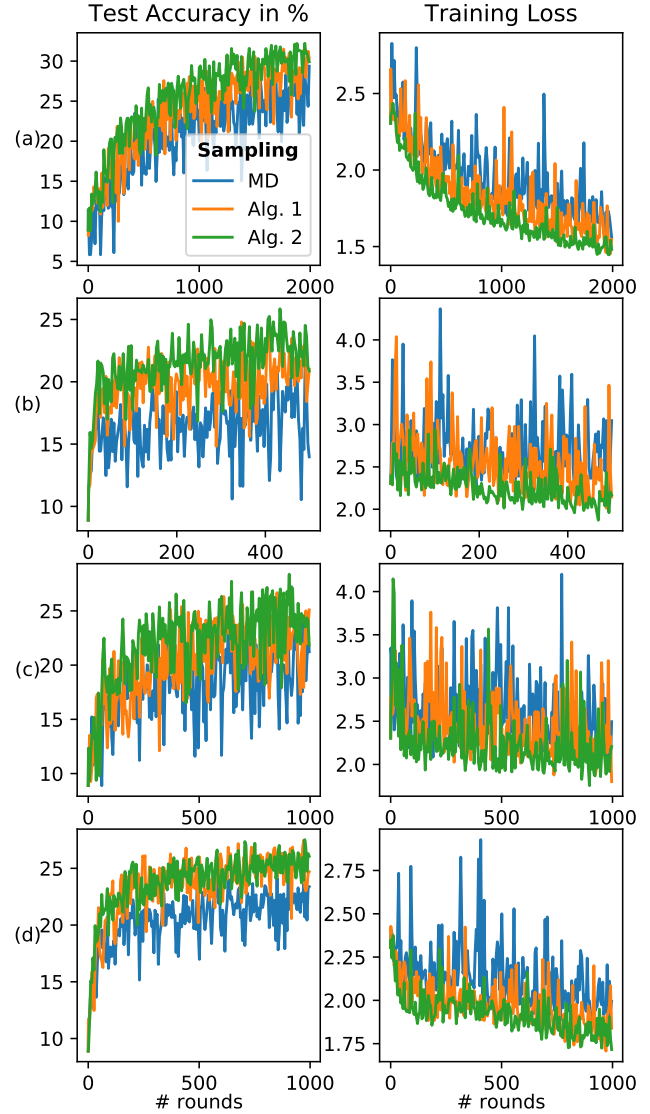


Figure 9. We consider the federated dataset partitioned from CIFAR10 using a Dirichlet distribution with parameter $\alpha = 0.01$. We investigate the influence of N the number of SGD run by each client in the first two rows with $N = 10$ and $N = 500$ for $m = 10$ and the influence of sampled clients with $m = 5$ and $m = 20$ for $N = 100$ in the last two rows. For each dataset, we use respective learning rate $lr = \{0.1, 0.05, 0.05, 0.05\}$.

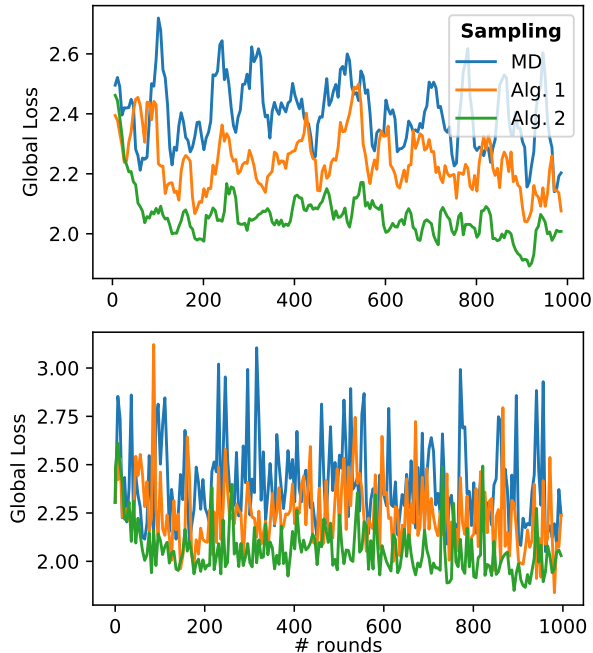


Figure 10. Training loss convergence for FL with FedProx local loss function regularization ($\mu = 0.1$). We consider CIFAR10 partitioned with $\text{Dir}(\alpha = 0.01)$, learning rate $lr = 0.05$, $m = 10$ sampled clients, and $N = 100$ SGD.

$$\mathcal{L}'_i(\theta_i^{t+1}) = \mathcal{L}_i(\theta_i^{t+1}) + \frac{\mu}{2} \|\theta_i^{t+1} - \theta^t\|^2 \quad (57)$$

where θ_{t+1} is the updated local model of client i and θ^t is the current global model. μ is the hyperparameter monitoring the regularization and is common for all the clients. This framework enables smoother federated learning processes.

We try a range of regularization term $\mu \in \{0.001, 0.01, 0.1, 1.\}$ and keep $\mu = 0.1$ maximizing the performances of FedAvg with regularization and MD sampling. We notice in Figure 10 that Algorithm 1 and 2 still outperform MD sampling.