

Approximate Gradient Coding for Privacy-Flexible Federated Learning with Non-IID Data

Okko Makkonen*, Sampo Niemelä*, Camilla Hollanti*, Serge Kas Hanna†

*Department of Mathematics and Systems Analysis, School of Science, Aalto University, Finland

†Côte d’Azur University, CNRS, I3S Laboratory, France

Abstract—This work focuses on the challenges of non-IID data and stragglers/dropouts in federated learning. We introduce and explore a privacy-flexible paradigm that models parts of the clients’ local data as non-private, offering a more versatile and business-oriented perspective on privacy. Within this framework, we propose a data-driven strategy for mitigating the effects of label heterogeneity and client straggling on federated learning. Our solution combines both offline data sharing and approximate gradient coding techniques. Through theoretical results and numerical simulations on the MNIST dataset, we demonstrate that our approach enables achieving a deliberate trade-off between privacy and utility, leading to improved model convergence and accuracy while using an adaptable portion of non-private data.

I. INTRODUCTION

In the ever-evolving landscape of machine learning (ML), one paradigm has emerged as a promising solution to balance the demand for data privacy and the need for robust distributed model training: federated learning (FL) [1], [2]. This innovative approach allows multiple devices or entities, often called *clients*, to collaboratively train a shared model under the orchestration of a central server while keeping their data decentralized. The result is a model that benefits from the collective intelligence of diverse data sources without compromising individual data privacy. Traditional distributed learning [3], [4], on the other hand, involves a central server that owns all the data and has complete control over how it distributes it to edge devices to parallelize the learning process, making it more suitable for scenarios where the data is naturally centralized.

While various assumptions can be made when designing algorithms for distributed learning and analyzing their performance, the most prominent split is between assuming independent, identically distributed (IID) data and non-IID data. The IID data assumption refers to the case where the data samples across different clients are independently drawn from the same underlying distribution, i.e., share similar statistical properties, which simplifies the training process and algorithm design. While this assumption may be valid in traditional distributed learning settings where the data is centralized, it almost never holds in practical FL settings since the clients independently collect their own training data, which typically vary in both size and distribution.

In the context of FL, the term non-IID data refers to a scenario where the data available across clients is not statistically similar. Previous studies have shown that dealing with non-IID data is one of the most significant hurdles encountered by FL. More specifically, training on non-IID data introduces several drawbacks that can hinder the effectiveness of collaborative model training, such as slower convergence and

lower model accuracy [5]–[8]. Various approaches have been proposed in the literature for dealing with non-IID data in FL, including algorithmic methods, e.g. [5], [7]–[12], data-driven solutions based on data augmentation [13]–[17], and hybrid FL schemes [6], [18] with partially centralized data.

In addition to the statistical challenge of dealing with non-IID data, FL also suffers from the straggler problem. The straggler or dropout problem refers to the issue where some participating clients or devices temporarily drop out or fail to complete their participation in a FL round. This can happen because clients may join or leave the FL system dynamically or due to other practical considerations such as poor network connections and resource constraints. Previous studies have demonstrated that the presence of stragglers aggravates the problem of learning on non-IID data, resulting in diminished model quality and slower convergence [19], [20]. Coding for straggler mitigation is a popular solution that has been extensively studied in the traditional distributed learning setting, e.g. [21]–[26], and also more recently in FL [27], [28]. The key idea in these works is to first introduce training data redundancy in the distributed system and then apply coding techniques that exploit this redundancy to either approximate or recover the central model update, even when some clients drop out.

While previous works in FL treat all the clients’ local data as equally private, our work introduces and investigates a *privacy-flexible* FL paradigm that allows for a deliberate trade-off between privacy and utility, offering more adaptable privacy measures. Our motivation for studying this setting stems from practical and commercial situations where striking the right balance between privacy and utility is essential for integrating FL in business models. More precisely, in the proposed paradigm, we model a portion of each client’s data as being non-private. We argue that parts of the local data in FL can be treated as non-private for many reasons, including: (i) Participants often have diverse datasets of varying sensitivity, and some data may not contain anything private or sensitive, allowing for differentiation in privacy treatment. (ii) Privacy can be *selective*, i.e., some participants may be incentivized to sacrifice their privacy by revealing or sharing part of their raw data for enhancing the model’s performance and/or for potential financial rewards.

In this work, we introduce a novel data-driven approach that leverages the concept of privacy flexibility to improve the utility of FL while addressing the challenges posed by non-IID data and stragglers. As data can be non-IID in different ways, we specifically address *label heterogeneity*, which refers to the unequal representation of certain classes across clients. Our proposed scheme consists of two key components. First, we introduce an offline *data sharing* mechanism, executed just once before the training phase. In this process, participants share some of their non-private data with each other to reduce the statistical imbalances resulting from label heterogeneity. The data sharing

The work was supported in part by the French government through the France 2030 investment plan managed by the National Research Agency, as part of the Initiative of Excellence Université Côte d’Azur under reference number ANR-15-IDEX-01.

also creates redundancy in the training datasets as some of the data becomes available to multiple clients. The second component of our scheme capitalizes on this redundancy via *approximate gradient coding*. This coding method is designed to optimize the training process by providing an unbiased estimate of the central model update rule of gradient descent in the presence of stragglers, similar to the approach in [26]. Moreover, we theoretically demonstrate that our proposed scheme reduces the variance of the obtained estimate, suggesting faster convergence, and characterize this reduction in terms of system parameters. Additionally, we present simulation results on real data using the MNIST dataset [29], which substantiate our theoretical findings. These simulations reveal that our scheme enables achieving a convergence behavior and model accuracy that closely mirrors the IID case, contingent on the amount of shared non-private data.

II. PRELIMINARIES

A. Notation

We use bold letters for vectors and sans-serif letters for random variables, e.g., \mathbf{x} , \mathbf{X} , and \mathbf{X} represent a vector, a random variable, and a random vector, respectively. Uppercase italic letters are used for sets. Let $[n] \triangleq \{1, 2, \dots, n\}$ be the set of integers from 1 to n (inclusive), and let $\Delta^n \triangleq \{(\delta_0, \delta_1, \dots, \delta_n) \in \mathbb{R}^{n+1} \mid \sum_{i=0}^n \delta_i = 1 \text{ and } \delta_i \geq 0\}$ be the standard n -simplex. We represent the Euclidean norm of a vector \mathbf{x} by $\|\mathbf{x}\|$, and the scalar product between two vectors \mathbf{x} and \mathbf{y} by $\langle \mathbf{x}, \mathbf{y} \rangle$.

B. System Model

Consider a FL setting with $N > 1$ clients who want to use their local data to collectively train a machine learning model with the help of a central server. Suppose that the clients want to run a supervised classification task, i.e., each client has a set of training examples (\mathbf{x}, y) with features \mathbf{x} and labels $y \in [L]$, where L represents the total number of classes. Let \mathcal{D}_i be the local dataset of client $i \in [N]$, and let $\mathcal{D} = \bigcup_{i=1}^N \mathcal{D}_i$ denote the global training dataset of size M given by the union of the local datasets across all the clients. The goal of FL is to find an optimal global model β^* that minimizes a given loss function \mathcal{L} , i.e., $\beta^* = \arg \min_{\beta} \mathcal{L}(\mathcal{D}, \beta)$. The global loss can be often expressed as the sum of the individual losses evaluated at each example (\mathbf{x}, y) , i.e.,

$$\mathcal{L}(\mathcal{D}, \beta) = \sum_{(\mathbf{x}, y) \in \mathcal{D}} \mathcal{L}(\mathbf{x}, y, \beta).$$

A common approach for solving such optimization problems is using gradient-based methods. Namely, the central server applies or approximates the following gradient descent update rule at each iteration $t \in \{0, 1, \dots\}$ of the algorithm

$$\beta^{(t+1)} = \beta^{(t)} - \frac{\eta}{M} \sum_{(\mathbf{x}, y) \in \mathcal{D}} \nabla \mathcal{L}(\mathbf{x}, y, \beta^{(t)}),$$

where $\nabla \mathcal{L}$ is the gradient of the loss function and η represents the learning rate. In distributed gradient descent, the central server first sends the current model $\beta^{(t)}$ to all clients. Each client i then locally computes the individual gradients of the loss function over the examples in its local dataset \mathcal{D}_i and sends a linear combination of these gradients to the central server. The central server then aggregates these local computations and either recovers the full gradient $\nabla \mathcal{L}(\mathcal{D}, \beta^{(t)})$ or obtains

an estimate of it. The previous steps are repeated iteratively until convergence. If some of the clients are unresponsive (straggler/dropout) during certain iterations, the central server will not be able to recover the full gradient. In this case, the master's aggregate gives an estimate of the full gradient. In this work, we consider a straggling model where each client is unresponsive with probability p in any given iteration, where this probability is independent across clients and iterations.

Suppose that the global training dataset \mathcal{D} consists of a total of M training examples such that there are K_ℓ examples corresponding to class $\ell \in [L]$. Let us assume that $K_\ell = K$ for all $\ell \in [L]$, such that $M = KL$. Let X_i^ℓ denote the proportion of instances of label $\ell \in [L]$ that are owned by client $i \in [N]$. We focus on a label-heterogeneous setting where the initial proportions $\mathbf{X}^\ell = (X_1^\ell, \dots, X_N^\ell) \in \Delta^{N-1}$ are typically far from $\Theta^* \triangleq (\frac{1}{N}, \frac{1}{N}, \dots, \frac{1}{N}) \in \Delta^{N-1}$, where Θ^* corresponds to the perfectly label-homogeneous setting. We adopt the squared Euclidean distances between the label proportions and Θ^* as a measure for label heterogeneity. Some examples for modeling label-heterogeneity are the following:

- 1) The single-class label-heterogeneous setting, where each client has data belonging only to a single class [6]. For instance, for every $\ell \in [L]$, the initial label proportions could be $\mathbf{X}^\ell = (X_1^\ell, X_2^\ell, \dots, X_N^\ell) = \mathbf{e}^\ell$, where \mathbf{e}^ℓ is the ℓ^{th} standard basis vector of \mathbb{R}^N , with $N = L$. In this case, \mathbf{X}^ℓ is deterministic, with $\|\mathbf{X}^\ell - \Theta^*\|^2 = \frac{N-1}{N}$ for all $\ell \in [L]$.
- 2) The initial label distributions are random and follow a Dirichlet distribution with a small concentration parameter [30]. Namely, for every $\ell \in [L]$, $\mathbf{X}^\ell = (X_1^\ell, X_2^\ell, \dots, X_N^\ell) \sim \text{Dir}_N(\alpha)$, where $\text{Dir}_N(\alpha)$ is the Dirichlet distribution with N categories and $\alpha > 0$ is the concentration parameter. Smaller values of α (close to zero) correspond to strongly heterogeneous settings; while $\alpha = \infty$ corresponds to a perfectly homogeneous one.

Furthermore, we propose and study a privacy-flexible setting where a proportion $c \in [0, 1]$ of each client's data is considered non-private. This proportion is assumed to be evenly distributed across all classes, i.e., each client $i \in [N]$ has $C_i^\ell \triangleq cX_i^\ell K$ non-private training examples from class $\ell \in [L]$.¹ We assume that clients collect their local data independently, which implies that disclosing the non-private data of one client does not leak information about the private data of another client.

III. PROPOSED SCHEME

To mitigate the effects of label heterogeneity and client straggling, we propose using an offline data sharing scheme that generates redundancy across clients by replicating non-private data. We will focus on the proportions of some fixed label $\ell \in [L]$ and drop the superscript ℓ . Recall that $\mathbf{X} = (X_1, \dots, X_N)$ denotes the initial label proportions prior to any data sharing, and let $\mathbf{Y} = (Y_1, \dots, Y_N)$ denote the corresponding final label proportions after data sharing. The replication-based data sharing scheme is described by $\mathbf{S} = (S_1, \dots, S_N) \in \mathbb{Z}^N$, where $S_i \geq 0$ denotes the number of non-private training examples that client $i \in [N]$ receives from other clients.² Then,

$$\mathbf{Y} = \frac{\mathbf{X}K + \mathbf{S}}{K + B}, \quad (1)$$

¹To simplify notation, we assume that $cX_i^\ell K$ are integers for all i and ℓ .

²Similar to [28], client-to-client offline communication can be assumed to be routed through the central server to ensure network connectivity.

where $B = \sum_{i=1}^N S_i$ denotes the total amount of shared data.

The goal of data sharing is twofold: (i) ‘‘Break’’ label heterogeneity by generating final label proportions \mathbf{Y} that are as close as possible to Θ^* in order to reduce the effects of ‘‘non-IIDness’’. (ii) Create data redundancy across clients which we will use to ensure resilience against straggling clients.

As previously mentioned, we use the squared Euclidean distances between the label proportions and Θ^* as a measure for label heterogeneity. Prior to any data sharing the squared distance is $\|\mathbf{X} - \Theta^*\|^2$, and after data sharing, the squared distance becomes $\|\mathbf{Y} - \Theta^*\|^2$, where \mathbf{Y} is given by (1). We aim to devise a data sharing scheme, i.e., determine \mathbf{S} , that minimizes $\|\mathbf{Y} - \Theta^*\|^2$ for a given realization \mathbf{X} of \mathbf{X} . If the label counts of each client’s local data \mathcal{D}_i (private and non-private data) are assumed to be public, one can obtain an optimal deterministic data sharing scheme that minimizes $\|\mathbf{Y} - \Theta^*\|^2$ for any realization \mathbf{X} of \mathbf{X} by solving a constrained optimization problem. However, knowing the label counts of each client’s data can potentially leak some information about the private data of the clients. Therefore, we propose a randomized data sharing scheme in Section III-A that does not require the knowledge of the label counts. We provide theoretical guarantees on the performance of the randomized data sharing scheme by evaluating $\mathbb{E}[\|\mathbf{Y} - \Theta^*\|^2 \mid \mathbf{X} = \mathbf{X}]$ in terms of the initial squared distance $\|\mathbf{X} - \Theta^*\|^2$, where the expectation is taken over the randomness of the data sharing process. Furthermore, in addition to minimizing the effects of label heterogeneity, an equally important consequence of the proposed data sharing scheme is generating redundancy across clients. In Section III-B, we present an approximate gradient coding scheme that leverages this redundancy to provide resilience against straggling clients.

A. Randomized Data Sharing

To reduce the effects of label heterogeneity without the knowledge of the label counts of each client’s local data, we propose a randomized offline data sharing scheme where each client shares its non-private data with $d \in \{0, 1, \dots, N-1\}$ other clients chosen uniformly at random. Let $\tilde{C}_i = \sum_{i' \neq i} C_{i'} = c(1 - X_i)K$ denote the total number of non-private examples in \mathcal{D} that are not owned by client $i \in [N]$. The randomized data sharing scheme places each of the \tilde{C}_i non-private examples independently with probability $\frac{d}{N-1}$ at client $i \in [N]$. Thus, client $i \in [N]$ receives S_i new examples, where $S_i \sim \text{Binomial}(\tilde{C}_i, \frac{d}{N-1})$ is a random variable that follows a binomial distribution with parameters $\tilde{C}_i = c(1 - X_i)K$ and $\frac{d}{N-1}$. Hence, the new label proportions after the randomized data sharing follow from (1), with $B = dcK$.

The next theorem expresses the expected value of $\|\mathbf{Y} - \Theta^*\|^2$ in terms of the replication factor $d \in \{0, 1, \dots, N-1\}$, privacy parameter $c \in [0, 1]$, number of clients N , and the initial distance $\|\mathbf{X} - \Theta^*\|^2$ for any realization $\mathbf{X} \in \Delta^{N-1}$ of \mathbf{X} .

Theorem 1. *For any realization $\mathbf{X} \in \Delta^{N-1}$ of \mathbf{X} , the randomized data sharing generates label proportions satisfying*

$$\begin{aligned} \mathbb{E}[\|\mathbf{Y} - \Theta^*\|^2 \mid \mathbf{X} = \mathbf{X}] &= \frac{dc(N-1-d)}{(1+dc)^2(N-1)K} \\ &+ \frac{(N-1-dc)^2}{(1+dc)^2(N-1)^2} \|\mathbf{X} - \Theta^*\|^2, \end{aligned}$$

where the expectation is over the randomness of the data sharing scheme. Furthermore, for $K \gg 1$, we have

$$\mathbb{E}[\|\mathbf{Y} - \Theta^*\|^2 \mid \mathbf{X} = \mathbf{X}] \approx \frac{(N-1-dc)^2}{(1+dc)^2(N-1)^2} \|\mathbf{X} - \Theta^*\|^2.$$

The proof of Theorem 1 is available in the extended version of our paper [31]. The significance of this theorem is that for any given set of input label proportions \mathbf{X} , the expected value of the output distance $\|\mathbf{Y} - \Theta^*\|^2$ is reduced by at least a factor of $(1+dc)^2$ with respect to the initial distance $\|\mathbf{X} - \Theta^*\|^2$ for large enough L . This shows that the randomized data sharing scheme can effectively reduce heterogeneity without knowing the label counts of each client’s local data.

B. Approximate Gradient Coding

As previously mentioned, at each iteration $t \in \{0, 1, \dots\}$, the central server sends the current model $\beta^{(t)}$ to all N clients. Each client participates in any given iteration independently with probability $1-p$, i.e., the client is a straggler with probability p . Let $\mathcal{S}_i \subseteq [M]$, $i \in [N]$, be the set of indices of the training examples $(\mathbf{x}, y) \in \mathcal{D}_i$ that are owned by client i . The clients participating in a given iteration t send a linear combination of the partial gradients computed over their local data, given by

$$f_i(\beta^{(t)}) = \sum_{j \in \mathcal{S}_i} W_j \mathbf{g}_j^{(t)}, \quad (2)$$

where $\mathbf{g}_j^{(t)} \triangleq \nabla \mathcal{L}(\mathbf{x}_j, y_j, \beta^{(t)})$ denotes the partial gradient evaluated at example (\mathbf{x}_j, y_j) , and W_j is the weighting factor of example (\mathbf{x}_j, y_j) . The weighting factor of (\mathbf{x}_j, y_j) is $W_j = \frac{1}{(1-p)d_j}$, where $d_j \in \mathbb{Z}^+$ is the total number of times the example (\mathbf{x}_j, y_j) is replicated across all N clients. Note that $d_j = 1$ for all private data, and $d_j = d+1$ for all non-private data, where d is the replication factor of the randomized data sharing scheme. The central server then aggregates the local computations of the clients to obtain an estimate of the full gradient given by

$$\hat{\mathbf{g}}^{(t)} = \sum_{i=1}^N l_i^{(t)} f_i(\beta^{(t)}), \quad (3)$$

where $l_i^{(t)} = 1$ if client i is participating in iteration t , and $l_i^{(t)} = 0$ otherwise. One can easily show that $\hat{\mathbf{g}}^{(t)}$ is an unbiased estimator of the full gradient $\mathbf{g}^{(t)} = \sum_{j=1}^M \mathbf{g}_j^{(t)}$, i.e., $\mathbb{E}[\hat{\mathbf{g}}^{(t)}] = \mathbf{g}^{(t)}$ [31, Appendix B]. Note that the expectations considered in this section are over the randomness of the straggling process in iteration t , conditioned on the model $\beta^{(t)}$ given by the central server to the clients. The variance of $\hat{\mathbf{g}}^{(t)}$ depends on the data sharing scheme being used and has a direct effect on the rate of convergence of the algorithm. More specifically, it has been shown in [32] that for an unbiased estimator, the value of $\mathbb{E}[\|\hat{\mathbf{g}}^{(t)}\|^2]$ is inversely proportional to the rate of convergence under certain assumptions on the loss function. An important trait of our proposed scheme is that it reduces the variance of the estimator and thus speeds up the convergence of the algorithm. We theoretically demonstrate this phenomenon in Theorem 2 and illustrate it practically through numerical simulations in Section IV. For the single-class label-heterogeneous setting described in Section II, Theorem 2 gives a lower bound on the expected difference between the variance of $\hat{\mathbf{g}}^{(t)}$ before and after data sharing. The bound is expressed in terms of the replication factor $d \in \{0, 1, \dots, N-1\}$, privacy

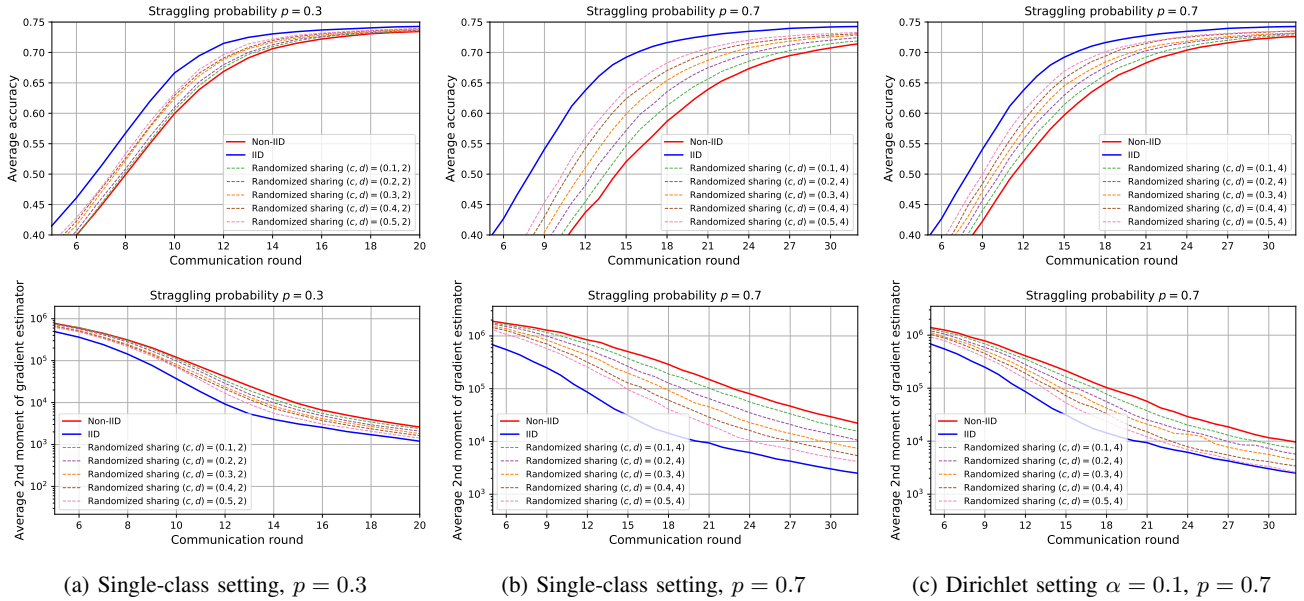


Fig. 1: Average testing accuracy and second moment of the gradient estimator (defined in (3)) as a function of the communication round (iteration). Figure 1c corresponds to the label-heterogeneous setting simulated according to the Dirichlet distribution with $\alpha = 0.1$, while 1a and 1b correspond to the single-class heterogeneous setting.

parameter $c \in [0,1]$, straggling probability $p \in [0,1)$, and other system parameters defined in Section II. The proof of this theorem is available in the extended version of our paper [31].

Theorem 2. Consider the single-class label-heterogeneous setting described in Section II and let $\hat{\mathbf{g}}_{\mathbf{X}}^{(t)}$ be the estimate of the gradient obtained in iteration t if no data is shared between clients. Let $\hat{\mathbf{g}}_{\mathbf{Y}}^{(t)}$ be the estimate of the gradient obtained in iteration t if the offline randomized data sharing scheme is applied with replication factor d and privacy parameter c . Under Assumption 1 in [31], it holds that

$$\mathbb{E} \left[\|\hat{\mathbf{g}}_{\mathbf{X}}^{(t)}\|^2 - \|\hat{\mathbf{g}}_{\mathbf{Y}}^{(t)}\|^2 \right] \geq \frac{p}{1-p} \frac{d-1}{d+1} \sum_{(j_1, j_2) \in \mathcal{J}_{\text{same}}^{\text{non-priv}}} \langle \mathbf{g}_{j_1}^{(t)}, \mathbf{g}_{j_2}^{(t)} \rangle,$$

where the expectation is over the randomness of the straggling process, and $\mathcal{J}_{\text{same}}^{\text{non-priv}}$ is the set of all index pairs $(j_1, j_2) \in [M]^2$ such that $(\mathbf{x}_{j_1}, y_{j_1})$ and $(\mathbf{x}_{j_2}, y_{j_2})$ are two examples belonging to the same class where at least one of them is non-private, with $|\mathcal{J}_{\text{same}}^{\text{non-priv}}| = MKc(2-c)$.

The result in Theorem 2 aligns with several intuitive aspects regarding the impact of the system parameters on the variance reduction achieved by the proposed scheme. For instance, the term $\frac{p}{1-p}$ indicates that the reduction in variance is more pronounced for higher values of p , which is intuitive since one would expect the benefits of data sharing to become more significant as the number of stragglers increases. Furthermore, the effects of the replication factor d and privacy parameter c are reflected by the term $\frac{d-1}{d+1}$ and the number of terms of the summation, respectively, where the latter is proportional to $c(2-c)$. By analyzing these terms, one could show that for $d > 1$ and a fixed offline communication cost characterized by the product dc , higher values of c lead to a greater reduction in variance. This observation is also intuitive, as relaxing the privacy constraint increases the diversity of redundancy generated by data sharing. Note that the theoretical analysis above relies on Assumption 1 in [31, Appendix B], which is

based on the hypothesis that gradients computed on examples from the same class tend to align well in the feature space, resulting in larger positive scalar product magnitudes. On the other hand, gradients from different classes may lack alignment or correlation, leading to lower scalar product magnitudes, reflecting the differences between the classes. This assumption is intuitive in practical scenarios where data in different classes are typically dissimilar, provided that training parameters are appropriately chosen to ensure convergence. The simulation results we provide in the next section on real data are consistent with this hypothesis, validating the variance reduction exhibited by the proposed scheme.

IV. SIMULATION RESULTS

A. Setup

We simulate a federated learning setup with $N = 10$ clients. The goal is to train a multinomial logistic regression model on the MNIST dataset [29], which is a supervised image classification task consisting of $L = 10$ different classes. We use a global training dataset \mathcal{D} of size $M = 300$, consisting of $K_\ell = K = 30$ images from each class $\ell \in [L]$ drawn uniformly at random from the 60,000 training images in MNIST. To model IID/non-IID settings, we consider multiple ways for partitioning the $M = 300$ training examples over the $N = 10$ clients: (i) IID setting, where the data is randomly shuffled, and then partitioned into 10 clients each receiving 30 examples. (ii) Two non-IID settings that follow from the single-class label-heterogeneous setting and the random Dirichlet distribution with $\alpha = 0.1$, as described in Section II.

We apply the randomized data sharing and approximate gradient coding to the two non-IID settings, and compare the performance of the trained model to the IID and non-IID cases with no data sharing. For each scenario, we run a total of 1000 simulations for 50 communication rounds (iterations) of the federated learning process and compute the average test accuracy and the second moment of the gradient estimator

(defined in (3)) as a function of the communication round. In each round, the identities of the stragglers/dropouts are determined according to the straggling model parameterized by p , as explained in Section II. We train the model using the SGD (stochastic gradient descent) optimizer with a learning rate of $\eta = 0.1$ and decay $\gamma = 0.97$.

B. Results

The average testing accuracy over 1000 simulations is plotted in Figure 1 for different parameters and compared with the baseline non-IID (no data sharing) and IID settings. These simulations show that there is a clear increase in the convergence rate as the parameter c is increased from 0 to 0.5. We also observe a slight difference in the final model accuracy achieved. We expect this difference to be more significant for more challenging datasets and more complex models. Using our proposed scheme and by manipulating the parameters c and d , we may interpolate between the non-IID and the IID settings. Furthermore, the plots of the average second moment of the gradient estimator validate our intuition and theoretical analysis in Section III-B. Namely, the results show that the increase in convergence rate is correlated with the decrease in the variance of the gradient estimator, which depends on the parameters of our scheme and the straggling behavior. The source code and additional simulation results can be found in [31].

V. CONCLUSION

In conclusion, our theoretical and numerical results demonstrate that under the privacy-flexible FL paradigm that we introduce, combining data sharing with gradient coding enables a deliberate trade-off between privacy (characterized by the parameter c) and utility (characterized by model convergence and accuracy). This trade-off is achieved at the expense of a one-time offline communication cost and additional local computation, which is a common price to pay when mitigating the effects of stragglers and non-IID data in distributed settings. In future work, we intend to analyze this trade-off for different models and datasets, and also other algorithms such as federated averaging, where multiple local model updates are performed in a single communication round.

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] P. Kairouz, H. B. McMahan, B. Avent *et al.*, "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [3] D. Jakovetic, "Distributed optimization: Algorithms and convergence rates," *PhD Thesis, Carnegie Mellon University, Pittsburgh PA, USA*, 2013.
- [4] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, V. Josifovski, J. Long, E. J. Shekita, and B.-Y. Su, "Scaling distributed machine learning with the parameter server," in *11th USENIX Symposium on operating systems design and implementation (OSDI 14)*, 2014, pp. 583–598.
- [5] T.-M. H. Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," *arXiv preprint arXiv:1909.06335*, 2019.
- [6] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-iid data," *arXiv preprint arXiv:1806.00582*, 2018.
- [7] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in *International conference on machine learning*. PMLR, 2020, pp. 5132–5143.
- [8] Y. Shi, J. Liang, W. Zhang, V. Y. Tan, and S. Bai, "Towards understanding and mitigating dimensional collapse in heterogeneous federated learning," *arXiv preprint arXiv:2210.00226*, 2022.
- [9] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, "Tackling the objective inconsistency problem in heterogeneous federated optimization," *Advances in neural information processing systems*, vol. 33, pp. 7611–7623, 2020.
- [10] A. E. Durmus, Z. Yue, M. Ramon, M. Matthew, W. Paul, and S. Venkatesh, "Federated learning based on dynamic regularization," in *International Conference on Learning Representations*, 2021.
- [11] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, "Federated multi-task learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [12] P. P. Liang, T. Liu, L. Ziyin, N. B. Allen, R. P. Auerbach, D. Brent, R. Salakhutdinov, and L.-P. Morency, "Think locally, act globally: Federated learning with local and global representations," *arXiv preprint arXiv:2001.01523*, 2020.
- [13] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.
- [14] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-L. Kim, "Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data," *arXiv preprint arXiv:1811.11479*, 2018.
- [15] J. Goetz and A. Tewari, "Federated learning via synthetic data," *arXiv preprint arXiv:2008.04489*, 2020.
- [16] M. Rasouli, T. Sun, and R. Rajagopal, "Fedgan: Federated generative adversarial networks for distributed data," *arXiv preprint arXiv:2006.07228*, 2020.
- [17] T. Yoon, S. Shin, S. J. Hwang, and E. Yang, "Fedmix: Approximation of mixup under mean augmented federated learning," *arXiv preprint arXiv:2107.00233*, 2021.
- [18] N. Yoshida, T. Nishio, M. Morikura, K. Yamamoto, and R. Yonetani, "Hybrid-fl for wireless networks: Cooperative learning mechanism using non-IID data," in *ICC 2020-2020 IEEE International Conference On Communications (ICC)*. IEEE, 2020, pp. 1–7.
- [19] Z. Charles and J. Konečný, "On the outsized importance of learning rates in local update methods," *arXiv preprint arXiv:2007.00878*, 2020.
- [20] A. Mitra, R. Jaafar, G. J. Pappas, and H. Hassani, "Linear convergence in federated learning: Tackling client heterogeneity and sparse gradients," *Advances in Neural Information Processing Systems*, vol. 34, pp. 14 606–14 619, 2021.
- [21] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *International Conference on Machine Learning*, 2017, pp. 3368–3376.
- [22] M. Ye and E. Abbe, "Communication-computation efficient gradient coding," in *International Conference on Machine Learning*. PMLR, 2018, pp. 5610–5619.
- [23] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1514–1529, 2018.
- [24] E. Ozfatura, S. Ulukus, and D. Gündüz, "Straggler-aware distributed learning: Communication–computation latency trade-off," *Entropy*, vol. 22, no. 5, 2020.
- [25] H. Wang, Z. Charles, and D. Papailiopoulos, "Erasurhead: Distributed gradient descent without delays using approximate gradient coding," *arXiv preprint arXiv:1901.09671*, 2019.
- [26] R. Bitar, M. Wootters, and S. El Rouayheb, "Stochastic gradient coding for straggler mitigation in distributed learning," *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 277–291, 2020.
- [27] S. Prakash, S. Dhakal, M. R. Akdeniz, Y. Yona, S. Talwar, S. Avestimehr, and N. Himayat, "Coded computing for low-latency federated learning over wireless edge networks," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 1, pp. 233–250, 2020.
- [28] R. Schlegel, S. Kumar, E. Rosnes, and A. G. i Amat, "CodedPaddedFL and CodedSecAgg: straggler mitigation and secure aggregation in federated learning," *IEEE Transactions on Communications*, 2023.
- [29] L. Deng, "The MNIST database of handwritten digit images for machine learning research," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.
- [30] M. Yurochkin, M. Agarwal, S. Ghosh, K. Greenewald, N. Hoang, and Y. Khazaeni, "Bayesian nonparametric federated learning of neural networks," in *International conference on machine learning*. PMLR, 2019, pp. 7252–7261.
- [31] O. Makkonen, S. Niemelä, C. Hollanti, and S. Kas Hanna, "Approximate gradient coding for privacy-flexible federated learning with non-iid data (github repository)," 2024, <https://github.com/okkomakkonen/label-heterogeneity>.
- [32] A. Rakhlin, O. Shamir, and K. Sridharan, "Making gradient descent optimal for strongly convex stochastic optimization," *arXiv preprint arXiv:1109.5647*, 2011.