
Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach

Alireza Fallah
EECS Department
Massachusetts Institute of Technology
Cambridge, MA 02139
afallah@mit.edu

Aryan Mokhtari
ECE Department
University of Texas at Austin
Austin, TX 78712
mokhtari@austin.utexas.edu

Asuman Ozdaglar
EECS Department
Massachusetts Institute of Technology
Cambridge, MA 02139
asuman@mit.edu

Abstract

In Federated Learning, we aim to train models across multiple computing units (users), while users can only communicate with a common central server, without exchanging their data samples. This mechanism exploits the computational power of all users and allows users to obtain a richer model as their models are trained over a larger set of data points. However, this scheme only develops a *common* output for all the users, and, therefore, it does not adapt the model to each user. This is an important missing feature, especially given the *heterogeneity* of the underlying data distribution for various users. In this paper, we study a personalized variant of the federated learning in which our goal is to find an *initial shared model* that current or new users can easily adapt to their local dataset by performing one or a few steps of gradient descent with respect to their own data. This approach keeps all the benefits of the federated learning architecture, and, by structure, leads to a more personalized model for each user. We show this problem can be studied within the Model-Agnostic Meta-Learning (MAML) framework. Inspired by this connection, we study a *personalized* variant of the well-known Federated Averaging algorithm and evaluate its performance in terms of gradient norm for non-convex loss functions. Further, we characterize how this performance is affected by the closeness of underlying distributions of user data, measured in terms of distribution distances such as Total Variation and 1-Wasserstein metric.

1 Introduction

In Federated Learning (FL), we consider a set of n users that are all connected to a central node (server), where each user has access only to its local data [1]. In this setting, the users aim to come up with a model that is trained over all the data points in the network without exchanging their local data with other users or the central node due to privacy issues or communication limitations. More

formally, if we define $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ as the loss corresponding to user i , the goal is to solve

$$\min_{w \in \mathbb{R}^d} f(w) := \frac{1}{n} \sum_{i=1}^n f_i(w). \quad (1)$$

In particular, consider a supervised learning setting, where f_i represents expected loss over the data distribution of user i , i.e.,

$$f_i(w) := \mathbb{E}_{(x,y) \sim p_i} [l_i(w; x, y)], \quad (2)$$

where $l_i(w; x, y)$ measures the error of model w in predicting the true label $y \in \mathcal{Y}_i$ given the input $x \in \mathcal{X}_i$, and p_i is the distribution over $\mathcal{X}_i \times \mathcal{Y}_i$. The focus of this paper is on a data *heterogeneous* setting where the probability distribution p_i of users are not identical. To illustrate this formulation, consider the example of training a Natural Language Processing (NLP) model over the devices of a set of users. In this problem, p_i represents the empirical distribution of words and expressions used by user i . Hence, $f_i(w)$ can be expressed as $f_i(w) = \sum_{(x,y) \in \mathcal{S}_i} p_i(x, y) l_i(w; x, y)$, where \mathcal{S}_i is the data set corresponding to user i and $p_i(x, y)$ is the probability that user i assigns to a specific word which is proportional to the frequency of using this word by user i .

Indeed, each user can solve its local problem defined in (2) without any exchange of information with other users; however, the resulted model may not generalize well to new samples as it has been trained over a small number of samples. If users cooperate and exploit the data available at all users, then their local models could obtain stronger generalization guarantees. A conventional approach for achieving this goal is minimizing the aggregate of local functions defined in (1). However, this scheme only develops a *common* output for all the users, and therefore, it does not adapt the model to each user. In particular, in the *heterogeneous* settings where the underlying data distribution of users are not identical, the resulted global model obtained by minimizing the average loss could perform arbitrarily poorly once applied to the local dataset of each user. In other words, the solution of problem (1) is not *personalized* for each user. To highlight this point, recall the NLP example, where although the distribution over the words and expressions varies from one person to another, the solution to problem (1) provides a shared answer for all users, and, therefore, it is not fully capable of achieving a user-specific model.

In this paper, we overcome this issue by considering a modified formulation of the federated learning problem which incorporates personalization (Section 2). Building on the Model-Agnostic Meta-Learning (MAML) problem formulation introduced in [2], the goal of this new formulation is to find an initial point shared between all users which performs well *after* each user updates it with respect to its own loss function, potentially by performing a few steps of a gradient-based method. This way, while the initial model is derived in a distributed manner over all users, the final model implemented by each user differs from other ones based on her or his own data. We study a Personalized variant of the FedAvg algorithm, called Per-FedAvg, designed for solving the proposed personalized FL problem (Section 3). In particular, we elaborate on its connections with the original FedAvg algorithm [3], and also, discuss a number of considerations that one needs to take into account for implementing Per-FedAvg. We also establish the convergence properties of the proposed Per-FedAvg algorithm for minimizing non-convex loss functions (Section 4). In particular, we characterize the role of data heterogeneity and closeness of data distribution of different users, measured by distribution distances, such as Total Variation (TV) or 1-Wasserstein, on the convergence of Per-FedAvg.

Related Work. Recently we have witnessed significant progress in developing novel methods that address different challenges in FL; see [4, 5]. In particular, there have been several works on various aspects of FL, including preserving the privacy of users [6–9] and lowering communication cost [10–13]. Several work develop algorithms for the homogeneous setting, where the data points of all users are sampled from the same probability distribution [14–17]. More related to our paper, there are several works that study statistical heterogeneity of users’ data points in FL [18–23], but they do not attempt to find a *personalized solution* for each user.

The centralized version of model-agnostic meta-learning (MAML) problem was first proposed in [2] and followed by a number of papers studying its empirical characteristics [24–29] as well as its convergence properties [30, 31]. In this work, we focus on the convergence of MAML methods for the FL setting that is more challenging as nodes perform multiple local updates before sending their updates to the server, which is not considered in previous theoretical works on meta-learning.

Recently, the idea of personalization in FL and its connections with MAML has gained a lot of attention. In particular, [32] considers a formulation and algorithm similar to our paper, and elaborates

on the empirical success of this framework. Also, recently, there has been a number of other papers that have studied different combinations of MAML-type methods with FL architecture from an empirical point of view [33, 34]. However, our main focus is on developing a theoretical understating regarding this formulation, where we characterize the convergence of the Per-FedAvg, and the role of this algorithm’s parameters on its performance. Besides, in our numerical experiment section, we show how the method studied in [32] may not perform well in some cases, and propose another algorithm which addresses this issue. In addition, an independent and concurrent work [35] studies a similar formulation theoretically for the case of strongly convex functions. The results in [35] are completely different from ours, as they study the case that the functions are strongly convex and exact gradients are available, while we study nonconvex functions, and also address gradient stochasticity.

Using meta-learning and multi-task learning to achieve personalization is not limited to MAML framework. In particular, [36] proposes ARUBA, a meta-learning algorithm inspired by online convex optimization, and shows that applying it to FedAvg improves its performance. A similar idea is later used in [37] to design differentially private algorithms with application in FL. Also in [38], the authors use multi-task learning framework and propose a new method, MOCHA, to address the statistical and systems challenges, including data heterogeneity and communication efficiency. Their proposed multi-task learning scheme also leads to a set of solutions that are more user-specific. A detailed survey on the connections of FL and multi-task and meta-learning can be found in [4, 5]. Also, in [39], the authors consider a framework for training a mixture of a single global model and local models, leading to a personalized solution for each user. A similar idea has been studied in [40], where the authors propose an adaptive federated learning algorithm that learns a mixture of local and global models as the personalized model.

2 Personalized Federated Learning via Model-Agnostic Meta-Learning

As we stated in Section 1, our goal in this section is to show how the fundamental idea behind the Model-Agnostic Meta-Learning (MAML) framework in [2] can be exploited to design a personalized variant of the FL problem. To do so, let us first briefly recap the MAML formulation. Given a set of tasks drawn from an underlying distribution, in MAML, in contrast to the traditional supervised learning setting, the goal is not finding a model which performs well on all the tasks in expectation. Instead, in MAML, we assume we have a limited computational budget to update our model after a new task arrives, and in this new setting, we look for an *initialization* which performs well *after* it is updated with respect to this new task, possibly by one or a few steps of gradient descent. In particular, if we assume each user takes the initial point and updates it using one step of gradient descent with respect to its own loss function, then problem (1) changes to

$$\min_{w \in \mathbb{R}^d} F(w) := \frac{1}{n} \sum_{i=1}^n f_i(w - \alpha \nabla f_i(w)), \tag{3}$$

where $\alpha \geq 0$ is the stepsize. The strength of this formulation is that, not only it allows us to maintain the advantages of FL, but also it captures the difference between users as either existing or new users can take the solution of this new problem as an initial point and slightly update it with respect to their own data. Going back to the NLP example, this means that the users could take this resulting initialization and update it by going over their own data \mathcal{S}_i and performing just one or few steps of gradient descent to obtain a model that works well for their own dataset.

As mentioned earlier, for the considered heterogeneous model of data distribution, solving problem (1) is not the ideal choice as it returns a single model that even after a few steps of local gradient may not quickly adjust to each users local data. On the other hand, by solving (3) we find an initial model (Meta-model) which is trained in a way that after one step of local gradient leads to a good model for each individual user. This formulation can also be extended to the case that users run a few steps of gradient update, but to simplify our notation we focus on the single gradient update case. We would like to mention that the problem formulation in (3) for FL was has been proposed independently in another work [32] and studied numerically. In this work, we focus on the theoretical aspect of this problem and seek a provably convergent method for the case that the functions f_i are nonconvex.

3 Personalized FedAvg

In this section, we present the Personalized FedAvg (Per-FedAvg) method to solve (3). This algorithm is inspired by FedAvg, but it is designed to find the optimal solution of (3) instead of (1). In FedAvg, at each round, the server chooses a fraction of users with size rn ($r \in (0, 1]$) and sends its current model to these users. Each selected user i updates the received model based on its own loss function f_i and by running $\tau \geq 1$ steps of stochastic gradient descent. Then, the active users return their updated models to the server. Finally, the server updates the global model by computing the average of the models received from these selected users, and then the next round follows. Per-FedAvg follows the same principles. First, note that function F in (3) can be written as the average of *meta-functions* F_1, \dots, F_n where the meta-function F_i associated with user i is defined as

$$F_i(w) := f_i(w - \alpha \nabla f_i(w)). \quad (4)$$

To follow a similar scheme as FedAvg for solving problem (3), the first step is to compute the gradient of local functions, which in this case, the gradient ∇F_i , that is given by

$$\nabla F_i(w) = (I - \alpha \nabla^2 f_i(w)) \nabla f_i(w - \alpha \nabla f_i(w)). \quad (5)$$

Computing the gradient $\nabla f_i(w)$ at every round is often computationally costly. Hence, we take a batch of data \mathcal{D}^i with respect to distribution p_i to obtain an unbiased estimate $\tilde{\nabla} f_i(w, \mathcal{D}^i)$ given by

$$\tilde{\nabla} f_i(w, \mathcal{D}^i) := \frac{1}{|\mathcal{D}^i|} \sum_{(x,y) \in \mathcal{D}^i} \nabla l_i(w; x, y). \quad (6)$$

Similarly, the Hessian $\nabla^2 f_i(w)$ in (5) can be replaced by its unbiased estimate $\tilde{\nabla}^2 f_i(w, \mathcal{D}^i)$.

At round k of Per-FedAvg, similar to FedAvg, first the server sends the current global model w_k to a fraction of users \mathcal{A}_k chosen uniformly at random with size rn . Each user $i \in \mathcal{A}_k$ performs τ steps of stochastic gradient descent locally and with respect to F_i . In particular, these local updates generate a local sequence $\{w_{k+1,t}^i\}_{t=0}^{\tau}$ where $w_{k+1,0}^i = w_k$ and, for $\tau \geq t \geq 1$,

$$w_{k+1,t}^i = w_{k+1,t-1}^i - \beta \tilde{\nabla} F_i(w_{k+1,t-1}^i), \quad (7)$$

where β is the local learning rate (stepsize) and $\tilde{\nabla} F_i(w_{k+1,t-1}^i)$ is an estimate of $\nabla F_i(w_{k+1,t-1}^i)$ in (5). Note that the stochastic gradient $\tilde{\nabla} F_i(w_{k+1,t-1}^i)$ for all local iterates is computed using independent batches \mathcal{D}_t^i , $\mathcal{D}_t^{\prime i}$, and $\mathcal{D}_t^{\prime\prime i}$ as follows

$$\tilde{\nabla} F_i(w_{k+1,t-1}^i) := \left(I - \alpha \tilde{\nabla}^2 f_i(w_{k+1,t-1}^i, \mathcal{D}_t^{\prime\prime i}) \right) \tilde{\nabla} f_i(w_{k+1,t-1}^i - \alpha \tilde{\nabla} f_i(w_{k+1,t-1}^i, \mathcal{D}_t^i), \mathcal{D}_t^{\prime i}). \quad (8)$$

Note that $\tilde{\nabla} F_i(w_{k+1,t-1}^i)$ is a biased estimator of $\nabla F_i(w_{k+1,t-1}^i)$ due to the fact that $\tilde{\nabla} f_i(w_{k+1,t-1}^i - \alpha \tilde{\nabla} f_i(w_{k+1,t-1}^i, \mathcal{D}_t^i), \mathcal{D}_t^{\prime i})$ is a stochastic gradient that contains another stochastic gradient inside.

Once, the local updates $w_{k+1,\tau}^i$ are evaluated, users send them to the server, and the server updates its global model by averaging over the received models, i.e., $w_{k+1} = \frac{1}{rn} \sum_{i \in \mathcal{A}_k} w_{k+1,\tau}^i$.

Note that as in other MAML methods [2, 31], the update in (7) can be implemented in two stages: First, we compute $\tilde{w}_{k+1,t}^i = w_{k+1,t-1}^i - \alpha \tilde{\nabla} f_i(w_{k+1,t-1}^i, \mathcal{D}_t^i)$ and then evaluate $w_{k+1,t}^i$ by $w_{k+1,t}^i = \tilde{w}_{k+1,t}^i - \beta (I - \alpha \tilde{\nabla}^2 f_i(w_{k+1,t-1}^i, \mathcal{D}_t^{\prime\prime i})) \tilde{\nabla} f_i(\tilde{w}_{k+1,t}^i, \mathcal{D}_t^{\prime i})$. Indeed, it can be verified the outcome of the these two steps is equivalent to the update in (7). To simplify the notation, throughout the paper, we assume that the size of \mathcal{D}_t^i , $\mathcal{D}_t^{\prime i}$, and $\mathcal{D}_t^{\prime\prime i}$ is equal to D , D' , and D'' , respectively, and for any i and t . The steps of Per-FedAvg are depicted in Algorithm 1.

4 Theoretical Results

In this section, we study the convergence properties of the Personalized FedAvg (Per-FedAvg) method. We focus on nonconvex settings, and characterize the overall communication rounds between server and users to find an ϵ -approximate first-order stationary point, where its formal definition follows.

Definition 4.1. A random vector $w_\epsilon \in \mathbb{R}^d$ is called an ϵ -approximate First-Order Stationary Point (FOSP) for problem (3) if it satisfies $\mathbb{E}[\|\nabla F(w_\epsilon)\|^2] \leq \epsilon$.

Algorithm 1: The proposed Personalized FedAvg (Per-FedAvg) Algorithm

Input: Initial iterate w_0 , fraction of active users r .
for $k : 0$ to $K - 1$ **do**
 Server chooses a subset of users \mathcal{A}_k uniformly at random and with size rn ;
 Server sends w_k to all users in \mathcal{A}_k ;
for all $i \in \mathcal{A}_k$ **do**
 Set $w_{k+1,0}^i = w_k$;
for $t : 1$ to τ **do**
 Compute the stochastic gradient $\tilde{\nabla} f_i(w_{k+1,t-1}^i, \mathcal{D}_t^i)$ using dataset \mathcal{D}_t^i ;
 Set $\tilde{w}_{k+1,t}^i = w_{k+1,t-1}^i - \alpha \tilde{\nabla} f_i(w_{k+1,t-1}^i, \mathcal{D}_t^i)$;
 Set $w_{k+1,t}^i = w_{k+1,t-1}^i - \beta (I - \alpha \tilde{\nabla}^2 f_i(w_{k+1,t-1}^i, \mathcal{D}_t^i)) \tilde{\nabla} f_i(\tilde{w}_{k+1,t}^i, \mathcal{D}_t^i)$;
end for
 Agent i sends $w_{k+1,\tau}^i$ back to server;
end for
 Server updates its model by averaging over received models: $w_{k+1} = \frac{1}{rn} \sum_{i \in \mathcal{A}_k} w_{k+1,\tau}^i$;
end for

Next, we formally state the assumptions required for proving our main results.

Assumption 1. Functions f_i are bounded below, i.e., $\min_{w \in \mathbb{R}^d} f_i(w) > -\infty$.

Assumption 2. For every $i \in \{1, \dots, n\}$, f_i is twice continuously differentiable and L_i -smooth, and also, its gradient is bounded by a nonnegative constant B_i , i.e.,

$$\|\nabla f_i(w)\| \leq B_i, \quad \|\nabla f_i(w) - \nabla f_i(u)\| \leq L_i \|w - u\| \quad \forall w, u \in \mathbb{R}^d. \quad (9)$$

As we discussed in Section 3, the second-order derivative of all functions appears in the update rule of Per-FedAvg Algorithm. Hence, in the next Assumption, we impose a regularity condition on the Hessian of each f_i which is also a customary assumption in the analysis of second-order methods.

Assumption 3. For every $i \in \{1, \dots, n\}$, the Hessian of function f_i is ρ_i -Lipschitz continuous, i.e.,

$$\|\nabla^2 f_i(w) - \nabla^2 f_i(u)\| \leq \rho_i \|w - u\| \quad \forall w, u \in \mathbb{R}^d. \quad (10)$$

To simplify the analysis, in the rest of the paper, we define $B := \max_i B_i$, $L := \max_i L_i$, and $\rho := \max_i \rho_i$ which can be, respectively, considered as a bound on the norm of gradient of f_i , smoothness parameter of f_i , and Lipschitz continuity parameter of Hessian $\nabla^2 f_i$, for $i = 1, \dots, n$.

Our next assumption provides upper bounds on the variances of gradient and Hessian estimations.

Assumption 4. For any $w \in \mathbb{R}^d$, the stochastic gradient $\nabla l_i(x, y; w)$ and Hessian $\nabla^2 l_i(x, y; w)$, computed with respect to a single data point $(x, y) \in \mathcal{X}_i \times \mathcal{Y}_i$, have bounded variance, i.e.,

$$\mathbb{E}_{(x,y) \sim p_i} [\|\nabla l_i(x, y; w) - \nabla f_i(w)\|^2] \leq \sigma_G^2, \quad (11)$$

$$\mathbb{E}_{(x,y) \sim p_i} [\|\nabla^2 l_i(x, y; w) - \nabla^2 f_i(w)\|^2] \leq \sigma_H^2. \quad (12)$$

Finally, we state our last assumption which characterizes the *similarity* between the tasks of users.

Assumption 5. For any $w \in \mathbb{R}^d$, the gradient and Hessian of local functions $f_i(w)$ and the average function $f(w) = \sum_{i=1}^n f_i(w)$ satisfy the following conditions

$$\frac{1}{n} \sum_{i=1}^n \|\nabla f_i(w) - \nabla f(w)\|^2 \leq \gamma_G^2, \quad \frac{1}{n} \sum_{i=1}^n \|\nabla^2 f_i(w) - \nabla^2 f(w)\|^2 \leq \gamma_H^2. \quad (13)$$

Assumption 5 captures the diversity between the gradients and Hessians of users. Note that under Assumption 2, the conditions in Assumption 5 are automatically satisfied for $\gamma_G = 2B$ and $\gamma_H = 2L$. However, we state this assumption separately to highlight the role of similarity of functions corresponding to different users in convergence analysis of Per-FedAvg. In particular, in the following subsection, we highlight the connections between this assumption and the similarity of distributions p_i for the case of supervised learning (2) under two different distribution distances.

4.1 On the Connections of Task Similarity and Distribution Distances

Recall the definition of f_i in (2). Note that Assumption 5 captures the similarity of loss functions of different users. Hence, a fundamental question here is whether this has any connection with the closeness of distributions p_i . We study this connection by considering two different distances: Total Variation (TV) distance and 1-Wasserstein distance. Throughout this subsection, we assume all users have the same loss function $l(\cdot; \cdot)$ over the same set of inputs and labels, i.e., $f_i(w) := \mathbb{E}_{z \sim p_i} [l(z; w)]$ where $z := (x, y) \in \mathcal{Z} := \mathcal{X} \times \mathcal{Y}$. Also, let $p = \frac{1}{n} \sum_i p_i$ denote the average of all users' distributions.

• **Total Variation (TV) Distance:** For distributions q_1 and q_2 over countable set \mathcal{Z} , their TV distance is given by $\|q_1 - q_2\|_{TV} = \frac{1}{2} \sum_{z \in \mathcal{Z}} |q_1(z) - q_2(z)|$. If we assume a stronger version of Assumption 2 holds where for any $z \in \mathcal{Z}$ and $w \in \mathbb{R}^d$, we have $\|\nabla_w l(z; w)\| \leq B$ and $\|\nabla_w^2 l(z; w)\| \leq L$, then Assumption 5 holds with (check Appendix B)

$$\gamma_G^2 = 4B^2 \frac{1}{n} \sum_{i=1}^n \|p_i - p\|_{TV}^2, \quad \gamma_H^2 = 4L^2 \frac{1}{n} \sum_{i=1}^n \|p_i - p\|_{TV}^2. \quad (14a)$$

This simple derivation shows that γ_G and γ_H exactly capture the difference between the probability distributions of the users in a heterogeneous setting.

• **1-Wasserstein Distance:** The 1-Wasserstein distance between two probability distributions q_1 and q_2 over a metric space \mathcal{Z} defined as $W_1(q_1, q_2) := \inf_{q \in Q(q_1, q_2)} \int_{\mathcal{Z} \times \mathcal{Z}} d(z_1, z_2) dq(z_1, z_2)$, where $d(\cdot, \cdot)$ is a distance function over metric space \mathcal{Z} and $Q(q_1, q_2)$ denotes the set of all measures on $\mathcal{Z} \times \mathcal{Z}$ with marginals q_1 and q_2 on the first and second coordinate, respectively. Here, we assume all p_i have bounded support (note that this assumption holds in many cases as either \mathcal{Z} itself is bounded or because we normalize the data). Also, we assume that for any w , the gradient $\nabla_w l(z; w)$ and the Hessian $\nabla_w^2 l(z; w)$ are both Lipschitz with respect to parameter z and distance $d(\cdot, \cdot)$, i.e.,

$$\|\nabla_w l(z_1; w) - \nabla_w l(z_2; w)\| \leq L_{\mathcal{Z}} d(z_1, z_2), \quad \|\nabla_w^2 l(z_1; w) - \nabla_w^2 l(z_2; w)\| \leq \rho_{\mathcal{Z}} d(z_1, z_2). \quad (15)$$

Then, Assumption 5 holds with (check Appendix B)

$$\gamma_G^2 = L_{\mathcal{Z}}^2 \frac{1}{n} \sum_{i=1}^n W_1(p_i, p)^2, \quad \gamma_H^2 = \rho_{\mathcal{Z}}^2 \frac{1}{n} \sum_{i=1}^n W_1(p_i, p)^2. \quad (16)$$

This derivation does not require Assumption 2 and holds when (15) are satisfied. Finally, consider a special case where the data distributions are homogeneous, and each p_i is an empirical distribution drawn from a distribution p_u with sample size m . In this case, we have $W_1(p_i, p_u) = \mathcal{O}(1/\sqrt{m})$ [41]. Hence, since W_1 is a distance, it is easy to verify that $\gamma_G, \gamma_H = \mathcal{O}(1/\sqrt{m})^1$.

4.2 Convergence Analysis of Per-FedAvg Algorithm

In this subsection, we derive the overall complexity of Per-FedAvg for achieving an ϵ -first-order stationary point. To do so, we first prove the following intermediate result which shows that under Assumptions 2 and 3, the local meta-functions $F_i(w)$ defined in (4) and their average function $F(w) = (1/n) \sum_{i=1}^n F_i(w)$ are smooth.

Lemma 4.2. *Recall the definition of $F_i(w)$ in (4) with $\alpha \in [0, 1/L]$. If Assumptions 2 and 3 hold, then F_i is smooth with parameter $L_F := 4L + \alpha\rho B$. As a consequence, the average function $F(w) = (1/n) \sum_{i=1}^n F_i(w)$ is also smooth with parameter L_F .*

Assumption 4 provides upper bounds on the variances of gradient and Hessian estimation for functions f_i . To analyze the convergence of Per-FedAvg, however, we require upper bounds on the bias and variance of gradient estimation of F_i . We derive these bounds in the following lemma.

Lemma 4.3. *Recall the definition of the gradient estimate $\tilde{\nabla} F_i(w)$ in (8) which is computed using \mathcal{D} , \mathcal{D}' , and \mathcal{D}'' that are independent batches with size D , D' , and D'' , respectively. If Assumptions 2-4*

¹While our focus here is to elaborate on the dependence of Wasserstein distance on the number of samples, it is worth noting that one drawback of this bound is that the convergence speed of Wasserstein distance in dimension is exponentially slow.

hold, then for any $\alpha \in [0, 1/L]$ and $w \in \mathbb{R}^d$ we have

$$\begin{aligned} \left\| \mathbb{E} \left[\tilde{\nabla} F_i(w) - \nabla F_i(w) \right] \right\| &\leq \frac{2\alpha L \sigma_G}{\sqrt{D}}, \\ \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \right] &\leq \sigma_F^2 := 12 \left[B^2 + \sigma_G^2 \left[\frac{1}{D'} + \frac{(\alpha L)^2}{D} \right] \right] \left[1 + \sigma_H^2 \frac{\alpha^2}{4D''} \right] - 12B^2. \end{aligned}$$

To measure the tightness of this result, we consider two special cases. First, if the exact gradients and Hessians are available, i.e., $\sigma_G = \sigma_H = 0$, then $\sigma_F = 0$ as well which is expected as we can compute exact ∇F_i . Second, for the classic federated learning problem, i.e., $\alpha = 0$ and $F_i = f_i$, we have $\sigma_F = \mathcal{O}(1)\sigma_G^2/D'$ which is tight up to constants.

Next, we use the similarity conditions for the functions f_i in Assumption 5 to study the similarity between gradients of the functions F_i .

Lemma 4.4. *Recall the definition of $F_i(w)$ in (4) and assume that $\alpha \in [0, 1/L]$. Suppose that the conditions in Assumptions 2, 3, and 5 are satisfied. Then, for any $w \in \mathbb{R}^d$, we have*

$$\frac{1}{n} \sum_{i=1}^n \|\nabla F_i(w) - \nabla F(w)\|^2 \leq \gamma_F^2 := 3B^2\alpha^2\gamma_H^2 + 192\gamma_G^2.$$

To check the tightness of this result, we focus on two special cases as we did for Lemma 4.3. First, if ∇f_i are all equal, i.e., $\gamma_G = \gamma_H = 0$, then $\gamma_F = 0$. This is indeed expected as all ∇F_i are equal to each other in this case. Second, for the classic federated learning problem, i.e., $\alpha = 0$ and $F_i = f_i$, we have $\gamma_F = \mathcal{O}(1)\gamma_G$ that is optimal up to a constant factor given the conditions in Assumption 5.

Theorem 4.5. *Consider the objective function F defined in (3) for the case that $\alpha \in (0, 1/L]$. Suppose that the conditions in Assumptions 1-4 are satisfied, and recall the definitions of L_F , σ_F , and η_F from Lemmas 4.2-4.4. Consider running Algorithm 1 for K rounds with τ local updates in each round and with $\beta \leq 1/(10\tau L_F)$. Then, the following first-order stationary condition holds*

$$\begin{aligned} \frac{1}{\tau K} \sum_{k=0}^{K-1} \sum_{t=0}^{\tau-1} E \left[\|\nabla F(\bar{w}_{k+1,t})\|^2 \right] &\leq \frac{4(F(w_0) - F^*)}{\beta\tau K} \\ &+ \mathcal{O}(1) \left(\beta L_F (1 + \beta L_F \tau (\tau - 1)) \sigma_F^2 + \beta L_F \gamma_F^2 \left(\frac{1-r}{r(n-1)} + \beta L_F \tau (\tau - 1) \right) + \frac{\alpha^2 L^2 \sigma_G^2}{D} \right), \end{aligned}$$

where $\bar{w}_{k+1,t}$ is the average of iterates of users in \mathcal{A}_k at time t , i.e., $\bar{w}_{k+1,t} = \frac{1}{rn} \sum_{i \in \mathcal{A}_k} w_{k+1,t}^i$, and in particular, $\bar{w}_{k+1,0} = w_k$ and $\bar{w}_{k+1,\tau} = w_{k+1}$.

Note that σ_F is not a constant, and as expressed in Lemma 4.3, we can make it arbitrary small by choosing batch sizes D , D' , or D'' large enough. To see how tight our result is, we again focus on special cases. Let $\alpha = 0$, $\tau = 1$, and $r = 1$. In this case, Per-FedAvg reduces to stochastic gradient descent, where the only source of stochasticity is the batches of gradient. In this case, the second term in the right hand side reduces to $\mathcal{O}(\beta L_F \sigma_F^2)$ where, here, σ_F^2 itself is equal to σ_G^2/D . This is the classic result for stochastic gradient descent for nonconvex functions, and we recover the lower bounds [42]. Also, it is worth noting that the term $\alpha^2 L^2 \sigma_G^2/D$ appears in the upper bound due to the fact that $\tilde{\nabla} F_i(w)$ is a *biased* estimator of $\nabla F_i(w)$. This bias term will be eliminated if we assume that we have access to the exact gradients at training time (see the discussion after Lemma 4.3), which is, for instance, the case in [35], where the authors focus on the deterministic case.

Next, we characterize the choices of τ , K , and β in terms of the required accuracy ϵ to obtain the best possible complexity bound for the result in Theorem 4.5.

Corollary 4.6. *Suppose the conditions in Theorem 4.5 are satisfied. If we set the number of local updates as $\tau = \mathcal{O}(\epsilon^{-1/2})$, number of communication rounds with the server as $K = \mathcal{O}(\epsilon^{-3/2})$, and stepsize of Per-FedAvg as $\beta = \epsilon$, then we find an $\mathcal{O}(\epsilon + \frac{\alpha^2 \sigma_G^2}{D})$ -first-order stationary point of F .*

The result in Corollary 4.6 shows that to achieve an $\mathcal{O}(\epsilon + \frac{\alpha^2 \sigma_G^2}{D})$ -first-order stationary point of F the Per-FedAvg algorithm requires $K = \mathcal{O}(\epsilon^{-3/2})$ rounds of communication between users and the server. Indeed, by setting $D = \mathcal{O}(\epsilon^{-1})$ or setting the meta-step stepsize as $\alpha = \mathcal{O}(\epsilon^{1/2})$ Per-FedAvg can find an ϵ -first-order stationary point of F for any arbitrary $\epsilon > 0$.

Remark 4.7. *The result of Theorem 4.5 and Corollary 4.6 provide an upper bound on the average of $E [\|\nabla F(\bar{w}_{k+1,t})\|^2]$ for all $k \in \{0, 1, \dots, K-1\}$ and $t \in \{0, 1, \dots, \tau-1\}$. However, one concern here is that due to the structure of Algorithm 1, for any k , we only have access to $\bar{w}_{k+1,t}$ for $t = 0$. To address this issue, at any iteration k , the center can choose $t_k \in \{0, 1, \dots, \tau-1\}$ uniformly at random, and ask all the users in \mathcal{A}_k to send w_{k+1,t_k}^i back to the server, in addition to $w_{k+1,\tau}^i$. By following this scheme we can ensure that the same upper bound also holds for the expected average models at the server, i.e., $\frac{1}{K} \sum_{k=0}^{K-1} E [\|\nabla F(\bar{w}_{k+1,t_k})\|^2]$.*

Remark 4.8. *It is worth noting that it is possible to achieve the same complexity bound using a diminishing stepsize. We will further discuss this at the end of Appendix G.*

5 Numerical Experiments

In this section, we numerically study the role of personalization when the data distributions are heterogeneous. In particular, we consider the multi-class classification problem over MNIST [43] and CIFAR-10 [44] datasets and distribute the training data between n users as follows: (i) Half of the users, *each* have a images of *each* of the first five classes; (ii) The rest, *each* have $a/2$ images from only *one* of the first five classes and $2a$ images from only *one* of the other five classes (see Appendix I for an illustration). We set the parameter a as $a = 196$ and $a = 68$ for MNIST and CIFAR-10 datasets, respectively. This way, we create an example where the distribution of images over all the users are different. Similarly, we divide the test data over the nodes with the same distribution as the one for the training data. Note that for this particular example in which the user’s distributions are significantly different, our goal is not to achieve state-of-the-art accuracy. Rather, we aim to provide an example to compare the various approaches for obtaining personalization in the heterogenous setting. Indeed, by using more complex neural networks the results for all the considered algorithms would improve; however, their relative performance would stay the same.

We focus on three algorithms: The first method that we consider is the FedAvg method, and, to do a fair comparison, we take the output of the FedAvg method, and update it with one step of stochastic gradient descent with respect to the test data, and then evaluate its performance. The second and third algorithms that we consider are two different efficient approximations of Per-FedAvg. Similarly, we evaluate the performance of these methods for the case that one step of local stochastic gradient descent is performed during test time. To formally explain these two approximate versions of Per-FedAvg, note that the implementation of Per-FedAvg requires access to second-order information which is computationally costly. To address this issue, we consider two different approximations:

(i) First, we replace the gradient estimate with its first-order approximation which ignores the Hessian term, i.e., $\tilde{\nabla} F_i(w_{k+1,t-1}^i)$ in (8) is approximated by $\tilde{\nabla} f_i(w_{k+1,t-1}^i - \alpha \tilde{\nabla} f_i(w_{k+1,t-1}^i, \mathcal{D}_t^i), \mathcal{D}_t^i)$. This is the same idea deployed in First-Order MAML (FO-MAML) in [2], and it has been studied empirically for the federated learning setting in [32]. We refer to this algorithm as Per-FedAvg (FO).

(ii) Second, we use the idea of the HF-MAML, proposed in [31], in which the Hessian-vector product in the MAML update is replaced by difference of gradients using the following approximation: $\nabla^2 \phi(w)u \approx (\nabla \phi(u + \delta v) - \nabla \phi(u - \delta v))/\delta$. We refer to this algorithm as Per-FedAvg (HF).

As shown in [31], for small stepsize at test time α both FO-MAML and HF-MAML perform well, but as α becomes large, HF-MAML outperforms FO-MAML in the centralized setting. A more detailed discussion on Per-FedAvg (FO) and Per-FedAvg (HF) is provided in Appendix H. Moreover, there we discuss how our analysis can be extended to these two methods. Note that the model obtained by any of these three methods is later updated using one step of stochastic gradient descent at the test time, and hence they have the same budget at the test time.

We use a neural network with two hidden layers with sizes 80 and 60, and we use Exponential Linear Unit (ELU) activation function. We take $n = 50$ users in the network, and run all three algorithms for $K = 1000$ rounds. At each round, we assume rn agents with $r = 0.2$ are chosen to run τ local updates. The batch sizes are $D = D' = 40$ and the learning rate is $\beta = 0.001$. Part of the code is adopted from [45]. Note that the reported results for all the considered methods corresponds to the average test accuracy among all users, after running one step of local stochastic gradient descent.

The test accuracy results along with the 95% confidence intervals are reported in Table 1. For MNIST dataset, both Per-FedAvg methods achieve a marginal gain compared to FedAvg. However, the achieved gain from using Per-FedAvg (HF) compared to FedAvg is more significant for CIFAR-10

Table 1: Comparison of test accuracy of different algorithms given different parameters

Dataset	Parameters	Algorithms		
		FedAvg + update	Per-FedAvg (FO)	Per-FedAvg (HF)
MNIST	$\tau = 10, \alpha = 0.01$	75.96% \pm 0.02%	78.00% \pm 0.02%	79.85% \pm 0.02%
	$\tau = 4, \alpha = 0.01$	60.18 % \pm 0.02%	64.55% \pm 0.02%	70.94% \pm 0.03%
CIFAR-10	$\tau = 10, \alpha = 0.001$	40.49% \pm 0.07%	46.98% \pm 0.1%	50.44% \pm 0.15%
	$\tau = 4, \alpha = 0.001$	38.38% \pm 0.07%	34.04% \pm 0.08%	43.73% \pm 0.11%
	$\tau = 4, \alpha = 0.01$	35.97% \pm 0.17%	25.32% \pm 0.18%	46.32% \pm 0.12%
	$\tau = 4, \alpha = 0.01$, diff. hetero.	58.59% \pm 0.11%	37.71% \pm 0.23%	71.25% \pm 0.05%

dataset. In particular, we have three main observations here: (i) For $\alpha = 0.001$ and $\tau = 10$, Per-FedAvg (FO) and Per-FedAvg (HF) perform almost similarly, and better than FedAvg. In addition, decreasing τ leads to a decrease in the performance of all three algorithms, which is expected as the total number of iterations decreases. (ii) Next, we study the role of α . By increasing α from 0.001 to 0.01, for $\tau = 4$, the performance of Per-FedAvg (HF) improves, which could be due to the fact that model adapts better with user data at test time. However, as discussed above, for larger α , Per-FedAvg (FO) performance drops significantly. (iii) Third, we examine the effect of changing the level of data heterogeneity. To do so, we change the data distribution of half of the users that have $a/2$ images from one of the first five classes by removing these images from their dataset. As the last line of Table 1 shows, Per-FedAvg (HF) performs significantly better than FedAvg under these new distributions, while Per-FedAvg (FO) still suffers from the issue we discussed in (ii). In summary, the more accurate implementation of Per-FedAvg, i.e., Per-FedAvg (HF), outperforms FedAvg in all cases and leads to a more personalized solution.

6 Conclusion

We considered the Federated Learning (FL) problem in the heterogeneous case, and studied a *personalized* variant of the classic FL formulation in which our goal is to find a proper initialization model for the users that can be quickly adapted to the local data of each user after the training phase. We highlighted the connections of this formulation with Model-Agnostic Meta-Learning (MAML), and showed how the decentralized implementation of MAML, which we called Per-FedAvg, can be used to solve the proposed personalized FL problem. We also characterized the overall complexity of Per-FedAvg for achieving first-order optimality in nonconvex settings. Finally, we provided a set of numerical experiments to illustrate the performance of two different first-order approximations of Per-FedAvg and their comparison with the FedAvg method, and showed that the solution obtained by Per-FedAvg leads to a more personalized solution compared to the solution of FedAvg.

Broader Impact

Federated Learning (FL) provides a framework for training machine learning models efficiently and in a distributed manner. Due to these favorable properties, it has gained significant attention and has been deployed in a broad range of applications with critical societal benefits. These applications go from healthcare systems, where machine learning models can be trained while preserving patients’ privacy, to image classification and NLP models, where tech companies can improve their neural networks without requiring users to share their data with a server or other users. In our work, we study one of the challenges in FL, which is the personalization aspect. The main question that we try to answer from a theoretical point of view is whether we can have a user-oriented variant of classic FL algorithms that can adapt to each user data while enjoying the distributed architecture of FL. We show the answer is positive, and provide rigorous theoretical guarantees for algorithms that can be used in all applications mentioned above to achieve more personalized models in FL framework. Indeed, this result could have a broad impact on improving the quality of users’ models in several applications that deploy federated learning such as healthcare systems.

Acknowledgments and Disclosure of Funding

Research was sponsored by the United States Air Force Research Laboratory and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. Alireza Fallah acknowledges support from MathWorks Engineering Fellowship. The research of Aryan Mokhtari is supported by NSF Award CCF-2007668.

References

- [1] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [2] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning*, (Sydney, Australia), 06–11 Aug 2017.
- [3] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, vol. 54 of *Proceedings of Machine Learning Research*, (Fort Lauderdale, FL, USA), pp. 1273–1282, PMLR, 20–22 Apr 2017.
- [4] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, *et al.*, “Advances and open problems in federated learning,” *arXiv preprint arXiv:1912.04977*, 2019.
- [5] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, 2020.
- [6] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Privacy aware learning,” *Journal of the ACM (JACM)*, vol. 61, no. 6, p. 38, 2014.
- [7] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang, “Learning differentially private recurrent language models,” *arXiv preprint arXiv:1710.06963*, 2017.
- [8] N. Agarwal, A. T. Suresh, F. X. X. Yu, S. Kumar, and B. McMahan, “cpsgd: Communication-efficient and differentially-private distributed sgd,” in *Advances in Neural Information Processing Systems*, pp. 7564–7575, 2018.
- [9] W. Zhu, P. Kairouz, B. McMahan, H. Sun, and W. Li, “Federated heavy hitters discovery with differential privacy,” in *International Conference on Artificial Intelligence and Statistics*, pp. 3837–3847, 2020.
- [10] A. Reisizadeh, A. Mokhtari, H. Hassani, A. Jadbabaie, and R. Pedarsani, “Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization,” in *International Conference on Artificial Intelligence and Statistics*, pp. 2021–2031, 2020.
- [11] X. Dai, X. Yan, K. Zhou, K. K. Ng, J. Cheng, and Y. Fan, “Hyper-sphere quantization: Communication-efficient sgd for federated learning,” *arXiv preprint arXiv:1911.04655*, 2019.
- [12] D. Basu, D. Data, C. Karakus, and S. Diggavi, “Qsparse-local-sgd: Distributed sgd with quantization, sparsification and local computations,” in *Advances in Neural Information Processing Systems*, pp. 14668–14679, 2019.
- [13] Z. Li, D. Kovalev, X. Qian, and P. Richtárik, “Acceleration for compressed gradient descent in distributed and federated optimization,” *arXiv preprint arXiv:2002.11364*, 2020.
- [14] S. U. Stich, “Local sgd converges fast and communicates little,” *arXiv preprint arXiv:1805.09767*, 2018.

- [15] J. Wang and G. Joshi, “Cooperative sgd: A unified framework for the design and analysis of communication-efficient sgd algorithms,” *arXiv preprint arXiv:1808.07576*, 2018.
- [16] F. Zhou and G. Cong, “On the convergence properties of a k-step averaging stochastic gradient descent algorithm for nonconvex optimization,” in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pp. 3219–3227, 2018.
- [17] T. Lin, S. U. Stich, K. K. Patel, and M. Jaggi, “Don’t use large mini-batches, use local SGD,” in *8th International Conference on Learning Representations, ICLR, 2020*.
- [18] 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.
- [19] A. K. Sahu, T. Li, M. Sanjabi, M. Zaheer, A. Talwalkar, and V. Smith, “On the convergence of federated optimization in heterogeneous networks,” *arXiv preprint arXiv:1812.06127*, 2018.
- [20] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, “Scaffold: Stochastic controlled averaging for on-device federated learning,” *arXiv preprint arXiv:1910.06378*, 2019.
- [21] F. Haddadpour and M. Mahdavi, “On the convergence of local descent methods in federated learning,” *arXiv preprint arXiv:1910.14425*, 2019.
- [22] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, “On the convergence of fedavg on non-iid data,” *arXiv preprint arXiv:1907.02189*, 2019.
- [23] A. K. R. Bayoumi, K. Mishchenko, and P. Richtarik, “Tighter theory for local sgd on identical and heterogeneous data,” in *International Conference on Artificial Intelligence and Statistics*, pp. 4519–4529, 2020.
- [24] A. Antoniou, H. Edwards, and A. Storkey, “How to train your MAML,” in *International Conference on Learning Representations*, 2019.
- [25] Z. Li, F. Zhou, F. Chen, and H. Li, “Meta-SGD: Learning to learn quickly for few-shot learning,” *arXiv preprint arXiv:1707.09835*, 2017.
- [26] E. Grant, C. Finn, S. Levine, T. Darrell, and T. Griffiths, “Recasting gradient-based meta-learning as hierarchical bayes,” in *International Conference on Learning Representations*, 2018.
- [27] A. Nichol, J. Achiam, and J. Schulman, “On first-order meta-learning algorithms,” *arXiv preprint arXiv:1803.02999*, 2018.
- [28] L. Zintgraf, K. Shiarli, V. Kurin, K. Hofmann, and S. Whiteson, “Fast context adaptation via meta-learning,” in *Proceedings of the 36th International Conference on Machine Learning*, pp. 7693–7702, 2019.
- [29] H. S. Behl, A. G. Baydin, and P. H. S. Torr, “Alpha MAML: adaptive model-agnostic meta-learning,” 2019.
- [30] P. Zhou, X. Yuan, H. Xu, S. Yan, and J. Feng, “Efficient meta learning via minibatch proximal update,” in *Advances in Neural Information Processing Systems 32*, pp. 1534–1544, Curran Associates, Inc., 2019.
- [31] A. Fallah, A. Mokhtari, and A. Ozdaglar, “On the convergence theory of gradient-based model-agnostic meta-learning algorithms,” in *International Conference on Artificial Intelligence and Statistics*, pp. 1082–1092, 2020.
- [32] F. Chen, M. Luo, Z. Dong, Z. Li, and X. He, “Federated meta-learning with fast convergence and efficient communication,” *arXiv preprint arXiv:1802.07876*, 2018.
- [33] Y. Jiang, J. Konečný, K. Rush, and S. Kannan, “Improving federated learning personalization via model agnostic meta learning,” *arXiv preprint arXiv:1909.12488*, 2019.

- [34] T. Li, M. Sanjabi, and V. Smith, “Fair resource allocation in federated learning,” *arXiv preprint arXiv:1905.10497*, 2019.
- [35] S. Lin, G. Yang, and J. Zhang, “A collaborative learning framework via federated meta-learning,” *arXiv preprint arXiv:2001.03229*, 2020.
- [36] M. Khodak, M.-F. F. Balcan, and A. S. Talwalkar, “Adaptive gradient-based meta-learning methods,” in *Advances in Neural Information Processing Systems*, pp. 5915–5926, 2019.
- [37] J. Li, M. Khodak, S. Caldas, and A. Talwalkar, “Differentially private meta-learning,” *arXiv preprint arXiv:1909.05830*, 2019.
- [38] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, “Federated multi-task learning,” in *Advances in Neural Information Processing Systems*, pp. 4424–4434, 2017.
- [39] F. Hanzely and P. Richtárik, “Federated learning of a mixture of global and local models,” *arXiv preprint arXiv:2002.05516*, 2020.
- [40] Y. Deng, M. M. Kamani, and M. Mahdavi, “Adaptive personalized federated learning,” *arXiv preprint arXiv:2003.13461*, 2020.
- [41] E. del Barrio, E. Giné, and C. Matrán, “Central limit theorems for the wasserstein distance between the empirical and the true distributions,” *Annals of Probability*, pp. 1009–1071, 1999.
- [42] Y. Arjevani, Y. Carmon, J. C. Duchi, D. J. Foster, N. Srebro, and B. Woodworth, “Lower bounds for non-convex stochastic optimization,” *arXiv preprint arXiv:1912.02365*, 2019.
- [43] Y. LeCun, “The mnist database of handwritten digits,” <http://yann.lecun.com/exdb/mnist/>, 1998.
- [44] A. Krizhevsky, G. Hinton, *et al.*, “Learning multiple layers of features from tiny images,” 2009.
- [45] J. Langelaar, “Mnist neural network training and testing,” *MATLAB Central File Exchange*, 2019.
- [46] C. Villani, *Optimal transport: old and new*, vol. 338. Springer Science & Business Media, 2008.

Supplementary Material

A Intermediate Notes

Note that the gradient Lipschitz assumption, i.e., the second inequality in (9), also implies that f_i satisfies the following conditions for all $w, u \in \mathbb{R}^d$:

$$-L_i I_d \preceq \nabla^2 f_i(w) \preceq L_i I_d, \quad (17a)$$

$$|f_i(w) - f_i(u) - \nabla f_i(u)^\top (w - u)| \leq \frac{L_i}{2} \|w - u\|^2. \quad (17b)$$

B Proofs of results in Subsection 4.1

B.1 TV Distance

Note that

$$\begin{aligned} \|\nabla f_i(w) - \nabla f(w)\| &= \left\| \sum_{z \in \mathcal{Z}} \nabla_w l(z; w) (p_i(z) - p(z)) \right\| \\ &\leq \sum_{z \in \mathcal{Z}} \|\nabla_w l(z; w)\| |p_i(z) - p(z)| \\ &\leq B \sum_{z \in \mathcal{Z}} |p_i(z) - p(z)| = 2B \|p_i - p\|_{TV} \end{aligned} \quad (18)$$

where the second inequality holds due to the assumption that $\|\nabla_w l(z; w)\| \leq B$ for any w and z . Plugging (18) in $\frac{1}{n} \sum_{i=1}^n \|\nabla f_i(w) - \nabla f(w)\|^2$, gives us the desired result. The other result on Hessians can be proved similarly.

B.2 1-Wasserstein Distance

We claim that for any i and $w \in \mathbb{R}^d$, we have

$$\|\nabla f_i(w) - \nabla f(w)\| \leq L_{\mathcal{Z}} W_1(p_i, p), \quad (19)$$

which will immediately give us one of the two results. To show this, first, note that

$$\begin{aligned} \|\nabla f_i(w) - \nabla f(w)\| &= \sup_{v \in \mathbb{R}^d: \|v\| \leq 1} v^\top (\nabla f_i(w) - \nabla f(w)) \\ &= \sup_{v \in \mathbb{R}^d: \|v\| \leq 1} \left(\mathbb{E}_{z \sim p_i} [v^\top \nabla l(z; w)] - \mathbb{E}_{z \sim p} [v^\top \nabla l(z; w)] \right) \end{aligned}$$

Thus, we need to show for any $v \in \mathbb{R}^d$ with $\|v\| \leq 1$, we have

$$\mathbb{E}_{z \sim p_i} [v^\top \nabla l(z; w)] - \mathbb{E}_{z \sim p} [v^\top \nabla l(z; w)] \leq L_{\mathcal{Z}} W_1(p_i, p). \quad (20)$$

Next, note that since p_i and p both have bounded support, by Kantorovich-Rubinstein Duality [46], we have

$$W_1(p_i, p) = \sup \{ \mathbb{E}_{z \sim p_i} [g(z)] - \mathbb{E}_{z \sim p} [g(z)] \mid \text{continuous } g: \mathcal{Z} \rightarrow \mathbb{R}, \text{Lip}(g) \leq 1 \}. \quad (21)$$

Using this result, to show (20), it suffices to show $g(z) = v^\top \nabla l(z; w)$ is $L_{\mathcal{Z}}$ -Lipschitz. Note that Cauchy-Schwarz inequality implies

$$\|v^\top \nabla l(z_1; w) - v^\top \nabla l(z_2; w)\| \leq \|v\| \|\nabla l(z_1; w) - \nabla l(z_2; w)\| \leq L_{\mathcal{Z}} d(z_1, z_2) \quad (22)$$

where the last inequality is obtained using $\|v\| \leq 1$ along with (15).

Finally, note that we can similarly show the result for γ_H by considering the fact that

$$\begin{aligned} \|\nabla^2 f_i(w) - \nabla^2 f(w)\| &= \max_{\xi \in \{1, -1\}} \sup_{v \in \mathbb{R}^d: \|v\| \leq 1} \xi v^\top (\nabla^2 f_i(w) - \nabla^2 f(w)) v \\ &= \max_{\xi \in \{1, -1\}} \sup_{v \in \mathbb{R}^d: \|v\| \leq 1} \xi \left(\mathbb{E}_{z \sim p_i} [v^\top \nabla^2 l(z; w) v] - \mathbb{E}_{z \sim p} [v^\top \nabla^2 l(z; w) v] \right) \end{aligned}$$

and taking the functions $g(z) = v^\top \nabla^2 l(z; w) v$ and $g(z) = -v^\top \nabla^2 l(z; w) v$ along with using Kantorovich-Rubinstein Duality Theorem again.

C Proof of Lemma 4.2

Recall that

$$\nabla F_i(w) = (I - \alpha \nabla^2 f_i(w)) \nabla f_i(w - \alpha \nabla f_i(w)). \quad (23)$$

Given this, note that

$$\begin{aligned} & \|\nabla F_i(w_1) - \nabla F_i(w_2)\| \\ &= \left\| (I - \alpha \nabla^2 f_i(w_1)) \nabla f_i(w_1 - \alpha \nabla f_i(w_1)) - (I - \alpha \nabla^2 f_i(w_2)) \nabla f_i(w_2 - \alpha \nabla f_i(w_2)) \right\| \\ &= \left\| (I - \alpha \nabla^2 f_i(w_1)) (\nabla f_i(w_1 - \alpha \nabla f_i(w_1)) - \nabla f_i(w_2 - \alpha \nabla f_i(w_2))) \right. \\ & \quad \left. + ((I - \alpha \nabla^2 f_i(w_1)) - (I - \alpha \nabla^2 f_i(w_2))) \nabla f_i(w_2 - \alpha \nabla f_i(w_2)) \right\| \end{aligned} \quad (24)$$

$$\begin{aligned} & \leq \|I - \alpha \nabla^2 f_i(w_1)\| \|\nabla f_i(w_1 - \alpha \nabla f_i(w_1)) - \nabla f_i(w_2 - \alpha \nabla f_i(w_2))\| \\ & \quad + \alpha \|\nabla^2 f_i(w_1) - \nabla^2 f_i(w_2)\| \|\nabla f_i(w_2 - \alpha \nabla f_i(w_2))\| \end{aligned} \quad (25)$$

where (24) is obtained by adding and subtracting $(I - \alpha \nabla^2 f_i(w_1)) \nabla f_i(w_2 - \alpha \nabla f_i(w_2))$ and the last inequality follows from the triangle inequality and the definition of matrix norm. Now, we bound two terms of (25) separately.

First, note that by (17a), $\|I - \alpha \nabla^2 f_i(w_1)\| \leq 1 + \alpha L$. Using this along with smoothness of f_i , we have

$$\begin{aligned} & \|I - \alpha \nabla^2 f_i(w_1)\| \|\nabla f_i(w_1 - \alpha \nabla f_i(w_1)) - \nabla f_i(w_2 - \alpha \nabla f_i(w_2))\| \\ & \leq (1 + \alpha L)L \|w_1 - \alpha \nabla f_i(w_1) - w_2 + \alpha \nabla f_i(w_2)\| \\ & \leq (1 + \alpha L)L (\|w_1 - w_2\| + \alpha \|\nabla f_i(w_1) - \nabla f_i(w_2)\|) \\ & \leq (1 + \alpha L)L(1 + \alpha L)\|w_1 - w_2\| \\ & \leq 4L\|w_1 - w_2\|, \end{aligned} \quad (26)$$

where we used smoothness of f_i along with $\alpha \leq 1/L$.

For the second term, Using (9) in Assumption 2 along with Assumption 3 implies

$$\alpha \|\nabla^2 f_i(w_1) - \nabla^2 f_i(w_2)\| \|\nabla f_i(w_2 - \alpha \nabla f_i(w_2))\| \leq \alpha \rho B \|w_1 - w_2\|. \quad (27)$$

Putting (26) and (27) together, we obtain the desired result.

D Proof of Lemma 4.3

Recall that the expression for the stochastic gradient $\tilde{\nabla} F_i(w)$ is given by

$$\tilde{\nabla} F_i(w) = (I - \alpha \tilde{\nabla}^2 f_i(w, \mathcal{D}'')) \tilde{\nabla} f_i(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}') \quad (28)$$

which can be written as

$$\tilde{\nabla} F_i(w) = (I - \alpha \nabla^2 f_i(w) + e_1) (\nabla f_i(w - \alpha \nabla f_i(w)) + e_2). \quad (29)$$

Note that in the above expression e_1 and e_2 are given by

$$e_1 = \alpha (\nabla^2 f_i(w) - \tilde{\nabla}^2 f_i(w, \mathcal{D}')),$$

and

$$e_2 = \tilde{\nabla} f_i(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}') - \nabla f_i(w - \alpha \nabla f_i(w)).$$

Based on Assumption 4, it can be easily shown that

$$\mathbb{E}[e_1] = 0, \quad (30a)$$

$$\mathbb{E}[\|e_1\|^2] \leq \alpha^2 \frac{\sigma_H^2}{D''}. \quad (30b)$$

Next, we proceed to bound the first and second moments of e_2 . To do so, first note that e_2 can also be written as

$$\begin{aligned} e_2 &= \left(\tilde{\nabla} f_i(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}') - \nabla f_i(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D})) \right) \\ & \quad + \left(\nabla f_i(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D})) - \nabla f_i(w - \alpha \nabla f_i(w)) \right). \end{aligned} \quad (31)$$

Note that, conditioning on \mathcal{D} , the first term is zero mean and the second term is deterministic. Therefore,

$$\begin{aligned} \|\mathbb{E}[e_2]\| &= \left\| \mathbb{E} \left[\nabla f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}) \right) - \nabla f_i \left(w - \alpha \nabla f_i(w) \right) \right] \right\| \\ &\leq \mathbb{E} \left[\left\| \nabla f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}) \right) - \nabla f_i \left(w - \alpha \nabla f_i(w) \right) \right\| \right] \\ &\leq \alpha L \mathbb{E} \left[\left\| \tilde{\nabla} f_i(w, \mathcal{D}) - \nabla f_i(w) \right\| \right] \end{aligned} \quad (32)$$

$$\leq \frac{\alpha L \sigma_G}{\sqrt{D}}, \quad (33)$$

where (32) is obtained using smoothness of f_i . The last inequality is also obtained using

$$\mathbb{E} \left[\left\| \tilde{\nabla} f_i(w, \mathcal{D}) - \nabla f_i(w) \right\|^2 \right] \leq \frac{\sigma_G^2}{D}. \quad (34)$$

In addition, we have

$$\begin{aligned} \mathbb{E} [\|e_2\|^2] &= \mathbb{E} [\mathbb{E} [\|e_2\|^2 | \mathcal{D}]] \\ &= \mathbb{E} \left[\left\| \tilde{\nabla} f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}' \right) - \nabla f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}) \right) \right\|^2 \right] \\ &\quad + \mathbb{E} \left[\left\| \nabla f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}) \right) - \nabla f_i \left(w - \alpha \nabla f_i(w) \right) \right\|^2 \right] \\ &\leq \frac{\sigma_G^2}{D'} + L^2 \alpha^2 \mathbb{E} \left[\left\| \tilde{\nabla} f_i(w, \mathcal{D}) - \nabla f_i(w) \right\|^2 \right] \end{aligned} \quad (35)$$

$$\leq \sigma_G^2 \left(\frac{1}{D'} + \frac{(\alpha L)^2}{D} \right) \quad (36)$$

where (36) follows from (34), and (35) is obtained using smoothness of f_i along with the fact that

$$\mathbb{E} \left[\left\| \tilde{\nabla} f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}' \right) - \nabla f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}) \right) \right\|^2 \right] \leq \frac{\sigma_G^2}{D'}.$$

Next, note that, by comparing (29) and (5), along with the fact that e_1 and e_2 are independent, and e_1 is zero-mean (30a), we have

$$\mathbb{E} \left[\tilde{\nabla} F_i(w) - \nabla F_i(w) \right] = (I - \alpha \nabla^2 f_i(w)) \mathbb{E} [e_2]. \quad (37)$$

Hence, by taking the norm of both sides, we obtain

$$\begin{aligned} \left\| \mathbb{E} \left[\tilde{\nabla} F_i(w) - \nabla F_i(w) \right] \right\| &= \left\| (I - \alpha \nabla^2 f_i(w)) \mathbb{E} [e_2] \right\| \\ &\leq \left\| (I - \alpha \nabla^2 f_i(w)) \right\| \|\mathbb{E} [e_2]\| \end{aligned} \quad (38)$$

where the last inequality follows from the definition of matrix norm. Now, using (33) along with the fact that $\|I - \alpha \nabla^2 f_i(w)\| \leq 1 + \alpha L \leq 2$ gives us the first result in Lemma 4.3.

To show the other result, note that, by comparing (29) and (5), along with the matrix norm definition, we have

$$\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\| \leq \|I - \alpha \nabla^2 f_i(w)\| \|e_2\| + \|e_1\| \|\nabla f_i(w - \alpha \nabla f_i(w))\| + \|e_1\| \|e_2\|. \quad (39)$$

As a result, by the Cauchy-Schwarz inequality $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$ for $a, b, c \geq 0$, we have

$$\begin{aligned} &\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \\ &\leq 3 \|I - \alpha \nabla^2 f_i(w)\|^2 \|e_2\|^2 + 3 \|e_1\|^2 \|\nabla f_i(w - \alpha \nabla f_i(w))\|^2 + 3 \|e_1\|^2 \|e_2\|^2. \end{aligned} \quad (40)$$

By taking expectation, and using the fact that $\|I - \alpha \nabla^2 f_i(w)\| \leq 1 + \alpha L \leq 2$ and

$$\|\nabla f_i(w - \alpha \nabla f_i(w))\| \leq B,$$

we have

$$\mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \right] \leq 3B^2 \mathbb{E} [\|e_1\|^2] + 12 \mathbb{E} [\|e_2\|^2] + 3 \mathbb{E} [\|e_1\|^2] \mathbb{E} [\|e_2\|^2] \quad (41)$$

where we also used the fact that e_1 and e_2 are independent as \mathcal{D}'' is independent from \mathcal{D} and \mathcal{D}' . Plugging (30b) and (36) in (41), we obtain

$$\begin{aligned} & \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \right] \\ & \leq 3B^2 \alpha^2 \frac{\sigma_H^2}{D''} + 12\sigma_G^2 \left(\frac{1}{D'} + \frac{(\alpha L)^2}{D} \right) + 3\alpha^2 \sigma_G^2 \sigma_H^2 \left(\frac{1}{D' D''} + \frac{(\alpha L)^2}{D D''} \right) \end{aligned}$$

which gives us the desired result.

E Proof of Lemma 4.4

Recall that

$$\nabla F_i(w) = (I - \alpha \nabla^2 f_i(w)) \nabla f_i(w - \alpha \nabla f_i(w)). \quad (42)$$

which can be expressed as

$$\nabla F_i(w) = (I - \alpha \nabla^2 f(w) + E_i) (\nabla f(w - \alpha \nabla f(w)) + r_i) \quad (43)$$

where

$$E_i = \alpha (\nabla^2 f(w) - \nabla^2 f_i(w)), \quad (44)$$

$$r_i = \nabla f_i(w - \alpha \nabla f_i(w)) - \nabla f(w - \alpha \nabla f(w)). \quad (45)$$

First, note that, by Assumption 5, we have

$$\frac{1}{n} \sum_{i=1}^n \|E_i\|^2 = \alpha^2 \gamma_H^2. \quad (46)$$

Second, note that

$$\begin{aligned} \|r_i\| & \leq \|\nabla f_i(w - \alpha \nabla f_i(w)) - \nabla f_i(w - \alpha \nabla f(w))\| \\ & \quad + \|\nabla f_i(w - \alpha \nabla f(w)) - \nabla f(w - \alpha \nabla f(w))\| \\ & \leq \alpha L \|\nabla f_i(w) - \nabla f(w)\| + \|\nabla f_i(w - \alpha \nabla f(w)) - \nabla f(w - \alpha \nabla f(w))\| \end{aligned} \quad (47)$$

where the last inequality is obtained using (9) in Assumption 2. Now, by using $(a+b)^2 \leq 2(a^2+b^2)$, we have

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \|r_i\|^2 \\ & \leq \frac{2}{n} \sum_{i=1}^n \left((\alpha L)^2 \|\nabla f_i(w) - \nabla f(w)\|^2 + \|\nabla f_i(w - \alpha \nabla f(w)) - \nabla f(w - \alpha \nabla f(w))\|^2 \right) \\ & \leq 2(1 + (\alpha L)^2) (\gamma_G^2 + \gamma_G^2) \end{aligned} \quad (48)$$

$$\leq 8\gamma_G^2. \quad (49)$$

where the second inequality follows from Assumption 5 and the last inequality is obtained using $\alpha L \leq 1$. Next, recall that the goal is to bound the variance of $\nabla F_i(w)$ when i is drawn from a uniform distribution. We know that by subtracting a constant from a random variable, its variance does not change. Thus, variance of $\nabla F_i(w)$ is equal to variance of $\nabla F_i(w) - (I - \alpha \nabla^2 f(w)) \nabla f(w - \alpha \nabla f(w))$. Also, the variance of the latter is bounded by its second moment, and hence,

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \|\nabla F_i(w) - \nabla F(w)\|^2 \leq \frac{1}{n} \sum_{i=1}^n \left\| E_i \nabla f(w - \alpha \nabla f(w)) + (I - \alpha \nabla^2 f(w)) r_i + E_i r_i \right\|^2 \\ & \leq \frac{1}{n} \sum_{i=1}^n \left(\|E_i \nabla f(w - \alpha \nabla f(w))\|^2 + \|(I - \alpha \nabla^2 f(w)) r_i\|^2 + \|E_i r_i\|^2 \right) \end{aligned} \quad (50)$$

Therefore, using $\|\nabla f(w - \alpha \nabla f(w))\| \leq B$ along with $\|I - \alpha \nabla^2 f(w)\| \leq 1 + \alpha L$ and Cauchy-Schwarz inequality $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$ for $a, b, c \geq 0$, we obtain

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \|\nabla F_i(w) - \nabla F(w)\|^2 &\leq 3 \left(B^2 \frac{1}{n} \sum_{i=1}^n \|E_i\|^2 + (1 + \alpha L)^2 \frac{1}{n} \sum_{i=1}^n \|r_i\|^2 + \frac{1}{n} \sum_{i=1}^n \|E_i r_i\|^2 \right) \\ &\leq 3 \left(B^2 \frac{1}{n} \sum_{i=1}^n \|E_i\|^2 + 4 \frac{1}{n} \sum_{i=1}^n \|r_i\|^2 + \frac{1}{n} \sum_{i=1}^n \|E_i\|^2 \|r_i\|^2 \right) \end{aligned} \quad (51)$$

where the last inequality is obtained using $\alpha L \leq 1$ along with $\|E_i r_i\| \leq \|E_i\| \|r_i\|$ which comes from the definition of matrix norm. Finally, to complete the proof, notice that we have

$$\frac{1}{n} \sum_{i=1}^n \|E_i\|^2 \|r_i\|^2 \leq \max_i \|E_i\|^2 \left(\frac{1}{n} \sum_{i=1}^n \|r_i\|^2 \right) \quad (52)$$

$$\leq \max_i \|E_i\|^2 (8\gamma_G^2) \quad (53)$$

$$\leq 32(\alpha L)^2 \gamma_G^2 \leq 32\gamma_G^2 \quad (54)$$

where (53) follows from (49) and the last line is obtained using $\alpha L \leq 1$ along with the fact that $\|\nabla^2 f_i(w)\| \leq L$, and thus,

$$\frac{\|E_i\|}{\alpha} = \|\nabla^2 f(w) - \nabla^2 f_i(w)\| \leq 2L. \quad (55)$$

Plugging (53) in (51) along with (46) and (49), we obtain the desired result.

F An Intermediate Result

Proposition F.1. *Recall from Section 3 that at any round $k \geq 1$, and for any agent $i \in \{1, \dots, n\}$, we can define a sequence of local updates $\{w_{k,t}^i\}_{t=0}^\tau$ where $w_{k,0}^i = w_{k-1}$ and, for $\tau \geq t \geq 1$,*

$$w_{k,t}^i = w_{k,t-1}^i - \beta \tilde{\nabla} F_i(w_{k,t-1}^i). \quad (56)$$

We further define the average of these local updates at round k and time t as $w_{k,t} = 1/n \sum_{i=1}^n w_{k,t}^i$. Suppose that the conditions in Assumptions 2-4 are satisfied. Then, for any $\alpha \in [0, 1/L]$ and any $t \geq 0$, we have

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \|w_{k,t}^i - w_{k,t}\| \right] \leq 2\beta t (1 + 2\beta L_F)^{t-1} (\sigma_F + \gamma_F), \quad (57a)$$

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \|w_{k,t}^i - w_{k,t}\|^2 \right] \leq 4\beta^2 (1 + \frac{1}{\phi}) t \left(1 + \phi + 16(1 + \frac{1}{\phi}) \beta^2 L_F^2 \right)^{t-1} (2\sigma_F^2 + \gamma_F^2) \quad (57b)$$

where $\phi > 0$ is an arbitrary positive constant and L_F , σ_F , and γ_F are given in Lemmas 4.2, 4.3, and 4.4, respectively.

Before stating the proof, note that an immediate consequence of this result is the following corollary:

Corollary F.2. *Under the same assumptions as Proposition F.1, and for any $\beta \leq 1/(10\tau L_F)$, we have*

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \|w_{k,t}^i - w_{k,t}\| \right] \leq 4\beta t (\sigma_F + \gamma_F), \quad (58a)$$

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \|w_{k,t}^i - w_{k,t}\|^2 \right] \leq 35\beta^2 t \tau (2\sigma_F^2 + \gamma_F^2) \quad (58b)$$

for any $0 \leq t \leq \tau$.

Proof. Let

$$S_t := \frac{1}{n} \sum_{i=1}^n \mathbb{E} [\|w_{k,t}^i - w_{k,t}\|] \quad (59)$$

where $S_0 = 0$ since $w_{k,0}^i = w_{k-1}$ for any i . Note that

$$\begin{aligned}
S_{t+1} &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} [\|w_{k,t+1}^i - w_{k,t+1}\|] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| w_{k,t}^i - \beta \tilde{\nabla} F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \left(w_{k,t}^j - \beta \tilde{\nabla} F_j(w_{k,t}^j) \right) \right\| \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| w_{k,t}^i - \frac{1}{n} \sum_{j=1}^n w_{k,t}^j \right\| \right] + \beta \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \tilde{\nabla} F_j(w_{k,t}^j) \right\| \right]. \tag{60}
\end{aligned}$$

Note that the first term in (60) is in fact S_t and the second one can be upper bounded as follows

$$\begin{aligned}
&\frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \tilde{\nabla} F_j(w_{k,t}^j) \right\| \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \nabla F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}^j) \right\| \right] + \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \nabla F_i(w_{k,t}^i) - \tilde{\nabla} F_i(w_{k,t}^i) \right\| \right] \\
&\quad + \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n \left\| \nabla F_j(w_{k,t}^j) - \tilde{\nabla} F_j(w_{k,t}^j) \right\| \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \nabla F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}^j) \right\| \right] + 2\beta\sigma_F
\end{aligned}$$

where the last inequality is obtained using Lemma 4.3. By substituting this in (60), we obtain

$$S_{t+1} \leq S_t + 2\beta\sigma_F + \beta \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \nabla F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}^j) \right\| \right]. \tag{61}$$

If we define $\eta_i := \nabla F_i(w_{k,t}^i) - \nabla F_i(w_{k,t})$, using (61), we obtain

$$\begin{aligned}
S_{t+1} &\leq S_t + 2\beta\sigma_F + \beta \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \nabla F_i(w_{k,t}) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}) \right\| \right] \\
&\quad + \beta \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \eta_i - \frac{1}{n} \sum_{j=1}^n \eta_j \right\| \right]. \tag{62}
\end{aligned}$$

Note that, by Lemma 4.2,

$$\|\eta_i\| \leq L_F \|w_{k,t}^i - w_{k,t}\|, \tag{63}$$

and thus,

$$\frac{1}{n} \sum_{i=1}^n \|\eta_i\| \leq L_F S_t. \tag{64}$$

As a result, and by using (62), we have

$$\begin{aligned}
S_{t+1} &\leq (1 + 2\beta L_F) S_t + 2\beta\sigma_F + \beta \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \nabla F_i(w_{k,t}) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}) \right\| \right] \\
&\leq (1 + 2\beta L_F) S_t + 2\beta(\sigma_F + \gamma_F) \tag{65}
\end{aligned}$$

where the last inequality is obtained using Lemma 4.4. Using (65) recursively, we obtain

$$S_{t+1} \leq \left(\sum_{j=0}^t (1 + 2\beta L_F)^j \right) 2\beta(\sigma_F + \gamma_F) \leq 2\beta(t+1)(1 + 2\beta L_F)^t (\sigma_F + \gamma_F) \tag{66}$$

which completes the proof of (57a). To prove (57b), let

$$\Sigma_t := \frac{1}{n} \sum_{i=1}^n \mathbb{E} [\|w_{k,t}^i - w_{k,t}\|^2]. \quad (67)$$

Similarly $\Sigma_0 = 0$. Note that

$$\begin{aligned} \Sigma_{t+1} &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} [\|w_{k,t+1}^i - w_{k,t+1}\|^2] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| w_{k,t}^i - \beta \tilde{\nabla} F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \left(w_{k,t}^j - \beta \tilde{\nabla} F_j(w_{k,t}^j) \right) \right\|^2 \right] \\ &\leq \frac{1+\phi}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| w_{k,t}^i - \frac{1}{n} \sum_{j=1}^n w_{k,t}^j \right\|^2 \right] \\ &\quad + \beta^2 \frac{1+1/\phi}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \tilde{\nabla} F_j(w_{k,t}^j) \right\|^2 \right] \end{aligned} \quad (68)$$

$$\leq (1+\phi)\Sigma_t + \beta^2 \frac{1+1/\phi}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \tilde{\nabla} F_j(w_{k,t}^j) \right\|^2 \right] \quad (69)$$

where (68) is obtained using $\|a+b\|^2 \leq (1+\phi)\|a\|^2 + (1+1/\phi)\|b\|^2$ for any arbitrary positive real number ϕ . To bound the second term in (69), note that

$$\begin{aligned} \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \tilde{\nabla} F_j(w_{k,t}^j) \right\|^2 \right] &\leq 2\mathbb{E} \left[\left\| \nabla F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}^j) \right\|^2 \right] \\ &\quad + 2\mathbb{E} \left[\left\| \left(\tilde{\nabla} F_i(w_{k,t}^i) - \nabla F_i(w_{k,t}^i) \right) + \frac{1}{n} \sum_{j=1}^n \left(\nabla F_j(w_{k,t}^j) - \tilde{\nabla} F_j(w_{k,t}^j) \right) \right\|^2 \right]. \end{aligned} \quad (70)$$

Now, we bound the second term in (70). Using Cauchy-Schwarz inequality

$$\left\| \sum_{l=1}^{n+1} a_l b_l \right\|^2 \leq \left(\sum_{l=1}^{n+1} \|a_l\|^2 \right) \left(\sum_{l=1}^{n+1} \|b_l\|^2 \right) \quad (71)$$

with $a_1 = \tilde{\nabla} F_i(w_{k,t}^i) - \nabla F_i(w_{k,t}^i)$, $b_1 = 1$ and $a_l = 1/\sqrt{n}$ ($\tilde{\nabla} F_{l-1}(w_{k,t}^{l-1}) - \nabla F_{l-1}(w_{k,t}^{l-1})$), $b_l = 1/\sqrt{n}$, for $l = 2, \dots, n+1$, implies

$$\begin{aligned} &\mathbb{E} \left[\left\| \left(\tilde{\nabla} F_i(w_{k,t}^i) - \nabla F_i(w_{k,t}^i) \right) + \frac{1}{n} \sum_{j=1}^n \left(\nabla F_j(w_{k,t}^j) - \tilde{\nabla} F_j(w_{k,t}^j) \right) \right\|^2 \right] \\ &\leq 2\mathbb{E} \left[\left\| \tilde{\nabla} F_i(w_{k,t}^i) - \nabla F_i(w_{k,t}^i) \right\|^2 + \frac{1}{n} \sum_{j=1}^n \left\| \nabla F_j(w_{k,t}^j) - \tilde{\nabla} F_j(w_{k,t}^j) \right\|^2 \right] \\ &\leq 4\sigma_F^2 \end{aligned} \quad (72)$$

where the last inequality is obtained using Lemma 4.3. Plugging (72) in (70) and using (69), we obtain

$$\begin{aligned} \Sigma_{t+1} &\leq (1+\phi)\Sigma_t + 8\left(1 + \frac{1}{\phi}\right)\beta^2\sigma_F^2 \\ &\quad + 2\left(1 + \frac{1}{\phi}\right)\beta^2\frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \nabla F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}^j) \right\|^2 \right]. \end{aligned} \quad (73)$$

Now, it remains to bound the last term in (73). Recall $\eta_i = \nabla F_i(w_{k,t}^i) - \nabla F_i(w_{k,t})$. First, note that, using $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$, we have

$$\|\nabla F_i(w_{k,t}^i) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t}^j)\|^2 \leq 2\|\nabla F_i(w_{k,t}) - \frac{1}{n} \sum_{j=1}^n \nabla F_j(w_{k,t})\|^2 + 2\|\eta_i - \frac{1}{n} \sum_{j=1}^n \eta_j\|^2. \quad (74)$$

Substituting this bound in (73) and using Lemma 4.4 yields

$$\Sigma_{t+1} \leq (1 + \phi)\Sigma_t + 4(1 + \frac{1}{\phi})\beta^2(2\sigma_F^2 + \gamma_F^2) + 4(1 + \frac{1}{\phi})\beta^2 \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\|\eta_i - \frac{1}{n} \sum_{j=1}^n \eta_j\|^2 \right]. \quad (75)$$

Note that, using Cauchy-Schwarz inequality (71) with $a_1 = \eta_i, b_1 = 1$ and $a_l = 1/\sqrt{n}\eta_{l-1}, b_l = 1/\sqrt{n}$ for $l = 2, \dots, n+1$, implies

$$\begin{aligned} \|\eta_i - \frac{1}{n} \sum_{j=1}^n \eta_j\|^2 &\leq 2 \left(\|\eta_i\|^2 + \frac{1}{n} \sum_{j=1}^n \|\eta_j\|^2 \right) \\ &\leq 2L_F^2 \left(\|w_{k,t}^i - w_{k,t}\|^2 + \frac{1}{n} \sum_{j=1}^n \|w_{k,t}^j - w_{k,t}\|^2 \right) \end{aligned} \quad (76)$$

where the last inequality is obtained using Lemma 4.2 which states

$$\|\eta_i\| \leq L_F \|w_{k,t}^i - w_{k,t}\|. \quad (77)$$

Plugging (76) in (75) implies

$$\Sigma_{t+1} \leq \left(1 + \phi + 16(1 + \frac{1}{\phi})\beta^2 L_F^2 \right) \Sigma_t + 4(1 + \frac{1}{\phi})\beta^2(2\sigma_F^2 + \gamma_F^2). \quad (78)$$

As a result, similar to (66), we obtain

$$\Sigma_{t+1} \leq 4\beta^2(1 + \frac{1}{\phi})(t+1) \left(1 + \phi + 16(1 + \frac{1}{\phi})\beta^2 L_F^2 \right)^t (2\sigma_F^2 + \gamma_F^2) \quad (79)$$

which gives us the desired result (57b).

Finally, to show (58), first note that for any n , we know

$$(1 + \frac{1}{n})^n \leq e. \quad (80)$$

Using this, along with the assumption $\beta \leq 1/(10L_F\tau)$ and the fact that $e^{0.2} \leq 2$, we immediately obtain (58a). To show the other one (58b), we use (57b) with $\phi = 1/(2\tau)$:

$$\begin{aligned} \phi + 16(1 + \frac{1}{\phi})\beta^2 L_F^2 &= \frac{1}{2\tau} + 16(1 + 2\tau)\beta^2 L_F^2 \\ &\leq \frac{1}{2\tau} + 16(1 + 2\tau) \frac{1}{100\tau^2} \\ &\leq \frac{1}{\tau} \end{aligned} \quad (81)$$

where the first inequality follows from the assumption $\beta \leq 1/(10L_F\tau)$ and the last inequality is obtained using the trivial bound $1 + 2\tau \leq 3\tau$. Finally, using (81) along with (80) completes the proof. \square

G Proof of Theorem 4.5

Although we only ask a fraction of agents to compute their local updates in Algorithm 1, here, and just for the sake of analysis, we assume all agents perform local updates. This is just for our analysis and we will not use all agents' updates in computing w_{k+1} . Also, from Proposition F.1, recall that $w_{k,t} = 1/n \sum_{i=1}^n w_{k,t}^i$.

Let \mathcal{F}_{k+1}^t denote the σ -field generated by $\{w_{k+1,t}^i\}_{i=1}^n$. Note that, by Lemma 4.2, we know F is smooth with gradient Lipschitz parameter L_F , and thus, by (17b), we have

$$\begin{aligned}
& F(\bar{w}_{k+1,t+1}) \\
& \leq F(\bar{w}_{k+1,t}) + \nabla F(\bar{w}_{k+1,t})^\top (\bar{w}_{k+1,t+1} - \bar{w}_{k+1,t}) + \frac{L_F}{2} \|\bar{w}_{k+1,t+1} - \bar{w}_{k+1,t}\|^2 \\
& \leq F(\bar{w}_{k+1,t}) - \beta \nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right) + \frac{L_F}{2} \beta^2 \left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right\|^2
\end{aligned} \tag{82}$$

where the last inequality is obtained using the fact that

$$\begin{aligned}
\bar{w}_{k+1,t+1} &= \frac{1}{rn} \sum_{i \in \mathcal{A}_k} w_{k+1,t+1}^i \\
&= \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \left(w_{k+1,t}^i - \beta \tilde{\nabla} F_i(w_{k+1,t}^i) \right) \\
&= \bar{w}_{k+1,t} - \beta \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i).
\end{aligned}$$

Taking expectation from both sides of (82) yields

$$\begin{aligned}
\mathbb{E}[F(\bar{w}_{k+1,t+1})] &\leq \mathbb{E}[F(\bar{w}_{k+1,t})] - \beta \mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right) \right] \\
&\quad + \frac{L_F}{2} \beta^2 \mathbb{E} \left[\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right\|^2 \right]
\end{aligned} \tag{83}$$

Next, note that

$$\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) = X + Y + Z + \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \tag{84}$$

where

$$X = \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \left(\tilde{\nabla} F_i(w_{k+1,t}^i) - \nabla F_i(w_{k+1,t}^i) \right), \tag{85}$$

$$Y = \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \left(\nabla F_i(w_{k+1,t}^i) - \nabla F_i(w_{k+1,t}) \right), \tag{86}$$

$$Z = \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \left(\nabla F_i(w_{k+1,t}) - \nabla F_i(\bar{w}_{k+1,t}) \right). \tag{87}$$

We next bound the moments of X , Y , and Z , condition on \mathcal{F}_{k+1}^t . First, recall the Cauchy-Schwarz inequality

$$\left\| \sum_{i=1}^{rn} a_i b_i \right\|^2 \leq \left(\sum_{i=1}^{rn} \|a_i\|^2 \right) \left(\sum_{i=1}^{rn} \|b_i\|^2 \right). \tag{88}$$

- Using this inequality with $a_i = (\tilde{\nabla} F_i(w_{k+1,t}^i) - \nabla F_i(w_{k+1,t}^i)) / \sqrt{rn}$ and $b_i = 1 / \sqrt{rn}$, we obtain

$$\|X\|^2 \leq \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \left\| \tilde{\nabla} F_i(w_{k+1,t}^i) - \nabla F_i(w_{k+1,t}^i) \right\|^2, \tag{89}$$

and hence, by using Lemma 4.3 along with the tower rule, we have

$$\mathbb{E}[\|X\|^2] = \mathbb{E}[\mathbb{E}[\|X\|^2 \mid \mathcal{F}_{k+1}^t]] \leq \sigma_F^2. \tag{90}$$

- Regarding Y , note that by using Cauchy-Schwarz inequality (similar to what we did above) along with smoothness of F_i , we obtain

$$\|Y\|^2 \leq \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \|\nabla F_i(w_{k+1,t}^i) - \nabla F_i(w_{k+1,t})\|^2 \leq \frac{L_F^2}{rn} \sum_{i \in \mathcal{A}_k} \|w_{k+1,t}^i - w_{k+1,t}\|^2. \quad (91)$$

Again, taking expectation and using the fact that \mathcal{A}_k is chosen uniformly at random, implies

$$\begin{aligned} \mathbb{E}[\|Y\|^2] &= \mathbb{E}[\mathbb{E}[\|Y\|^2 \mid \mathcal{F}_{k+1}^t]] \\ &\leq L_F^2 \mathbb{E} \left[\mathbb{E} \left[\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \|w_{k+1,t}^i - w_{k+1,t}\|^2 \mid \mathcal{F}_{k+1}^t \right] \right] \\ &= L_F^2 \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \|w_{k,t}^i - w_{k,t}\|^2 \right] \\ &\leq 35\beta^2 L_F^2 \tau (\tau - 1) (2\sigma_F^2 + \gamma_F^2) \end{aligned} \quad (92)$$

where the last step follows from (58b) in Corollary F.2 along with the fact that $t \leq \tau - 1$.

- Regarding Z , first recall that if we have n numbers a_1, \dots, a_n with mean $\mu = 1/n \sum_{i=1}^n a_i$ and variance $\sigma^2 = 1/n \sum_{i=1}^n |a_i - \mu|^2$, and we take a subset of them $\{a_i\}_{i \in \mathcal{A}}$ with size $|\mathcal{A}| = rn$ by sampling without replacement, then we have

$$\mathbb{E} \left[\left| \frac{\sum_{i \in \mathcal{A}} a_i}{rn} - \mu \right|^2 \right] = \frac{\sigma^2}{rn} \left(1 - \frac{rn-1}{n-1} \right) = \frac{\sigma^2(1-r)}{r(n-1)}. \quad (93)$$

Using this, we have

$$\mathbb{E} [\|\bar{w}_{k+1,t} - w_{k+1,t}\|^2 \mid \mathcal{F}_{k+1}^t] \leq \frac{(1-r)/n \sum_{i=1}^n \|w_{k+1,t}^i - w_{k+1,t}\|^2}{r(n-1)}, \quad (94)$$

and hence, by taking expectation from both sides and using the tower rule along with (58b) in Corollary F.2, we obtain

$$\mathbb{E} [\|\bar{w}_{k+1,t} - w_{k+1,t}\|^2] \leq \frac{35(1-r)\beta^2 \tau (\tau - 1) (2\sigma_F^2 + \gamma_F^2)}{r(n-1)}. \quad (95)$$

Next, note that by using Cauchy-Schwarz inequality (88), with $a_i = (\nabla F_i(w_{k+1,t}) - \nabla F_i(\bar{w}_{k+1,t})) / \sqrt{rn}$ and $b_i = 1/\sqrt{rn}$, we have

$$\begin{aligned} \|Z\|^2 &\leq \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \|\nabla F_i(w_{k+1,t}) - \nabla F_i(\bar{w}_{k+1,t})\|^2 \\ &\leq \frac{L_F^2}{rn} \sum_{i \in \mathcal{A}_k} \|w_{k+1,t} - \bar{w}_{k+1,t}\|^2 = L_F^2 \|\bar{w}_{k+1,t} - w_{k+1,t}\|^2 \end{aligned} \quad (96)$$

where the last inequality is obtained using smoothness of F_i (Lemma 4.2). Now, taking expectation from both sides and using (95) yields

$$\mathbb{E}[\|Z\|^2] \leq \frac{35(1-r)\beta^2 L_F^2 \tau (\tau - 1) (2\sigma_F^2 + \gamma_F^2)}{r(n-1)}. \quad (97)$$

Now, getting back to (83), we first lower bound the term

$$\mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right) \right].$$

To do so, note that, by (84), we have

$$\begin{aligned}
& \mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right) \right] \\
&= \mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(X + Y + Z + \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right) \right] \\
&\geq \mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right) \right] - \|\mathbb{E} [\nabla F(\bar{w}_{k+1,t})^\top X]\| \\
&\quad - \frac{1}{4} \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] - \mathbb{E} [\|Y + Z\|^2]
\end{aligned} \tag{98}$$

where the last inequality is obtained using the fact that

$$\mathbb{E} [\nabla F(\bar{w}_{k+1,t})^\top (Y + Z)] \leq \frac{1}{4} \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] + \mathbb{E} [\|Y + Z\|^2].$$

Now, we bound terms in (98) separately. First, note that by tower rule we have

$$\begin{aligned}
& \mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right) \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right) \mid \mathcal{F}_{k+1}^t \right] \right] \\
&= \mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \mathbb{E} \left[\left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right) \mid \mathcal{F}_{k+1}^t \right] \right] \\
&= \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2]
\end{aligned} \tag{99}$$

where the last equality is obtained using the fact that \mathcal{A}_k is chosen uniformly at random, and thus,

$$\mathbb{E} \left[\left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right) \mid \mathcal{F}_{k+1}^t \right] = \frac{1}{n} \sum_{i=1}^n \nabla F_i(\bar{w}_{k+1,t}).$$

Second, note that

$$\begin{aligned}
\mathbb{E} [\nabla F(\bar{w}_{k+1,t})^\top X] &= \mathbb{E} \left[\mathbb{E} [\nabla F(\bar{w}_{k+1,t})^\top X \mid \mathcal{F}_{k+1}^t] \right] \\
&= \mathbb{E} [\nabla F(\bar{w}_{k+1,t})^\top \mathbb{E} [X \mid \mathcal{F}_{k+1}^t]].
\end{aligned}$$

As a result, we have

$$\begin{aligned}
\|\mathbb{E} [\nabla F(\bar{w}_{k+1,t})^\top X]\| &= \left\| \mathbb{E} [\nabla F(\bar{w}_{k+1,t})^\top \mathbb{E} [X \mid \mathcal{F}_{k+1}^t]] \right\| \\
&\leq \frac{1}{4} \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] + \mathbb{E} \left[\left\| \mathbb{E} [X \mid \mathcal{F}_{k+1}^t] \right\|^2 \right] \\
&\leq \frac{1}{4} \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] + \frac{4\alpha^2 L^2 \sigma_G^2}{D}
\end{aligned} \tag{100}$$

where the last inequality follows from Lemma 4.3. Third, note that by Cauchy-Schwarz inequality,

$$\begin{aligned}
\mathbb{E} [\|Y + Z\|^2] &\leq 2 (\mathbb{E} [\|Y\|^2] + \mathbb{E} [\|Z\|^2]) \\
&\leq 70\beta^2 L_F^2 \tau(\tau - 1) (2\sigma_F^2 + \gamma_F^2) \left(1 + \frac{1-r}{r(n-1)} \right) \\
&\leq 140\beta^2 L_F^2 \tau(\tau - 1) (2\sigma_F^2 + \gamma_F^2)
\end{aligned} \tag{101}$$

where second inequality is obtained using (92) and (97). Plugging (99), (100), and (101) in (98) implies

$$\begin{aligned} & \mathbb{E} \left[\nabla F(\bar{w}_{k+1,t})^\top \left(\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right) \right] \\ & \geq \frac{1}{2} \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] - 140\beta^2 L_F^2 \tau(\tau-1)(2\sigma_F^2 + \gamma_F^2) - \frac{4\alpha^2 L^2 \sigma_G^2}{D}. \end{aligned} \quad (102)$$

Next, we characterize an upper bound for the other term in (83):

$$\mathbb{E} \left[\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right\|^2 \right]$$

Note that, by (84) we have

$$\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right\|^2 \leq 2\|X + Y + Z\|^2 + 2\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right\|^2, \quad (103)$$

and thus, by (101) along with (90), we have

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right\|^2 \right] \\ & \leq 2\mathbb{E} \left[\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right\|^2 \right] + 4\sigma_F^2 + 560\beta^2 L_F^2 \tau(\tau-1)(2\sigma_F^2 + \gamma_F^2). \end{aligned} \quad (104)$$

Note that, $\mathbb{E} \left[\frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \mid \mathcal{F}_{k+1}^t \right] = \nabla F(\bar{w}_{k+1,t})$, since \mathcal{A}_k is chosen uniformly at random. Also, by Lemma 4.4, we have

$$\frac{1}{n} \mathbb{E} \left[\left\| \nabla F_i(\bar{w}_{k+1,t}) - \nabla F(\bar{w}_{k+1,t}) \right\|^2 \mid \mathcal{F}_{k+1}^t \right] \leq \gamma_F^2,$$

and thus, by (93), we have

$$\mathbb{E} \left[\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \nabla F_i(\bar{w}_{k+1,t}) \right\|^2 \right] \leq \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] + \frac{\gamma_F^2(1-r)}{r(n-1)}. \quad (105)$$

Plugging (105) in (104), we obtain

$$\begin{aligned} & \mathbb{E} \left[\left\| \frac{1}{rn} \sum_{i \in \mathcal{A}_k} \tilde{\nabla} F_i(w_{k+1,t}^i) \right\|^2 \right] \\ & \leq 2\mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] + \frac{2\gamma_F^2(1-r)}{r(n-1)} + 4\sigma_F^2 + 560\beta^2 L_F^2 \tau(\tau-1)(2\sigma_F^2 + \gamma_F^2). \end{aligned} \quad (106)$$

Substituting (106) and (102) in (83) implies

$$\begin{aligned} & \mathbb{E} [F(\bar{w}_{k+1,t+1})] \\ & \leq \mathbb{E} [F(\bar{w}_{k+1,t})] - \beta(1/2 - \beta L_F) \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] \\ & \quad + 140(1 + 2\beta L_F) \beta^3 L_F^2 \tau(\tau-1)(2\sigma_F^2 + \gamma_F^2) + \beta^2 L_F \left(2\sigma_F^2 + \frac{\gamma_F^2(1-r)}{r(n-1)} \right) + \frac{4\beta\alpha^2 L^2 \sigma_G^2}{D} \\ & \leq \mathbb{E} [F(\bar{w}_{k+1,t})] - \frac{\beta}{4} \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] + \beta\sigma_T^2. \end{aligned} \quad (107)$$

where

$$\sigma_T^2 := 280(\beta L_F)^2 \tau(\tau-1)(2\sigma_F^2 + \gamma_F^2) + \beta L_F \left(2\sigma_F^2 + \frac{\gamma_F^2(1-r)}{r(n-1)} \right) + \frac{4\alpha^2 L^2 \sigma_G^2}{D} \quad (108)$$

the last inequality is obtained using $\beta \leq 1/(10\tau L_F)$. Summing up (107) for all $t = 0, \dots, \tau-1$, we obtain

$$\mathbb{E} [F(w_{k+1})] \leq \mathbb{E} [F(w_k)] - \frac{\beta\tau}{4} \left(\frac{1}{\tau} \sum_{t=0}^{\tau-1} \mathbb{E} [\|\nabla F(\bar{w}_{k+1,t})\|^2] \right) + \beta\tau\sigma_T^2 \quad (109)$$

where we used the fact that $\bar{w}_{k+1,\tau} = w_{k+1}$. Finally, summing up (109) for $k = 0, \dots, K - 1$ implies

$$\mathbb{E}[F(w_K)] \leq F(w_0) - \frac{\beta\tau K}{4} \left(\frac{1}{\tau K} \sum_{k=0}^{K-1} \sum_{t=0}^{\tau-1} E[\|\nabla F(\bar{w}_{k+1,t})\|^2] \right) + \beta\tau K \sigma_T^2. \quad (110)$$

As a result, we have

$$\begin{aligned} \frac{1}{\tau K} \sum_{k=0}^{K-1} \sum_{t=0}^{\tau-1} E[\|\nabla F(\bar{w}_{k+1,t})\|^2] &\leq \frac{4}{\beta\tau K} (F(w_0) - \mathbb{E}[F(w_K)] + \beta\tau K \sigma_T^2) \\ &\leq \frac{4(F(w_0) - F^*)}{\beta\tau K} + 4\sigma_T^2 \end{aligned} \quad (111)$$

which gives us the desired result.

Remark G.1. As stated in Remark 4.8, we could easily extend our analysis to the case with diminishing stepsize. In particular, by using β_k as the stepsize at iteration k , the descent result (109) holds with $\beta = \beta_k$. Hence, summing up this equation for $k = 0, \dots, K - 1$, we recover the same complexity bounds using $\beta_k = \mathcal{O}(1/\sqrt{\tau k})$.

H On First-Order Approximations of Per-FedAvg

As we stated previously, the Per-FedAvg method, same as MAML, requires computing Hessian-vector product which is computationally costly in some applications. As a result, one may consider using the first-order approximation of the update rule for the Per-FedAvg algorithm. The main goal of this section is to show how our analysis can be extended to the case that we either drop the second-order term or approximate the Hessian-vector product using first-order techniques.

To do so, we show that it suffices to only extend the result in Lemma 4.3 for the first-order approximation settings and find $\tilde{\sigma}_F$ such that

$$\begin{aligned} \left\| \mathbb{E} \left[\tilde{\nabla} F_i(w) - \nabla F_i(w) \right] \right\| &\leq m_F, \\ \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \right] &\leq \tilde{\sigma}_F^2. \end{aligned}$$

One can easily check that the rest of analysis does not change, and the final result (Theorem 4.5) holds if we just replace σ_F by $\tilde{\sigma}_F$ and $\alpha^2 L^2 \sigma_G^2 / D$ by m_F^2 .

We next focus on two different approaches, developed for MAML formulation, for approximating the Hessian-vector product, and show how we can characterize m_F and $\tilde{\sigma}_F$ for both cases:

- **Ignoring the second-order term:** Authors in [2] suggested to simply ignore the second-order term in the update of MAML to reduce the computation cost of MAML, i.e., to replace $\tilde{\nabla} F_i(w)$ with

$$\tilde{\nabla} f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}' \right). \quad (112)$$

This approach is known as First-Order MAML (FO-MAML), and it has been shown that it performs relatively well in many cases [2]. In particular, [31] characterized the convergence properties of FO-MAML for the centralized MAML problem. Next, we characterize the mean and variance of this gradient approximation.

Lemma H.1. Assume that we estimate $\nabla F_i(w)$ by (112) where \mathcal{D} and \mathcal{D}' are independent batches with size D and D' , respectively. Suppose that the conditions in Assumptions 2-4 are satisfied. Then, for any $\alpha \in [0, 1/L]$ and $w \in \mathbb{R}^d$, we have

$$\begin{aligned} \left\| \mathbb{E} \left[\tilde{\nabla} F_i(w) - \nabla F_i(w) \right] \right\| &\leq m_F^{FO} := \alpha L \left(\frac{\sigma_G}{\sqrt{D}} + B \right), \\ \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \right] &\leq (\tilde{\sigma}_F^{FO})^2 := 2\sigma_G^2 \left(\frac{1}{D'} + \frac{(\alpha L)^2}{D} \right) + 2(\alpha L B)^2. \end{aligned}$$

Proof. In fact, in this case, $\tilde{\nabla} F_i(w)$ is approximating

$$G_i(w) := \nabla f_i(w - \alpha \nabla f_i(w)). \quad (113)$$

To bound m_F^{FO} , note that

$$\left\| \mathbb{E} \left[\tilde{\nabla} F_i(w) - \nabla F_i(w) \right] \right\| \leq \left\| \mathbb{E} \left[\tilde{\nabla} F_i(w) - G_i(w) \right] \right\| + \left\| \mathbb{E} \left[G_i(w) - \nabla F_i(w) \right] \right\| \quad (114)$$

$$\leq \frac{\alpha L \sigma_G}{\sqrt{D}} + \alpha L B \quad (115)$$

where the first term follows from (33) in the proof of Lemma 4.3 in Appendix D, and the second term is obtained using

$$\begin{aligned} \|G_i(w) - \nabla F_i(w)\| &= \alpha \left\| \nabla^2 f_i(w) \nabla f_i(w - \alpha \nabla f_i(w)) \right\| \\ &\leq \alpha \left\| \nabla^2 f_i(w) \right\| \cdot \left\| \nabla f_i(w - \alpha \nabla f_i(w)) \right\| \leq \alpha L B \end{aligned} \quad (116)$$

where the first inequality follows from the matrix norm definition and the last inequality is obtained using Assumption 2.

To characterize $\tilde{\sigma}_F^{FO}$, note that

$$\mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \right] \leq 2 \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - G_i(w) \right\|^2 \right] + 2 \mathbb{E} \left[\left\| G_i(w) - \nabla F_i(w) \right\|^2 \right]. \quad (117)$$

We bound these two terms separately. Note that we have already bounded the first term in Appendix D (see (36)), and we have

$$\mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - G_i(w) \right\|^2 \right] \leq \sigma_G^2 \left(\frac{1}{D'} + \frac{(\alpha L)^2}{D} \right). \quad (118)$$

Plugging (118) and (116) into (117), we obtain the desired result. \square

Note that while the first term in $\tilde{\sigma}_F^{FO}$ can be made arbitrary small by choosing D and D' large enough, this is not the case for the second term. However, the second term is also negligible if α is small enough. Yet this bound suggests that this approximation introduces a non-vanishing error term which is directly carried to the final result (Theorem 4.5).

• **Estimating Hessian-vector product using gradient differences:** In the context of MAML problem, it has been shown that the update of FO-MAML leads to an additive error that does not vanish as time progresses. To resolve this matter, [31] introduced another variant of MAML, called HF-MAML, which approximates the Hessian-vector product by gradient differences. More formally, the idea behind their method is that for any function g , the product of the Hessian $\nabla^2 g(w)$ by any vector v can be approximated by

$$\frac{\nabla g(w + \delta v) - \nabla g(w - \delta v)}{2\delta} \quad (119)$$

with an error of at most $\rho \delta \|v\|^2$, where ρ is the parameter for Lipschitz continuity of the Hessian of g . Building on this idea, in Per-FedAvg update rule, we can replace $\tilde{\nabla} F_i(w)$ by

$$\tilde{\nabla} f_i \left(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}' \right) - \alpha \tilde{d}_i(w) \quad (120)$$

where

$$\tilde{d}_i(w) := \frac{\tilde{\nabla} f_i \left(w + \delta \tilde{\nabla} f_i(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}'), \mathcal{D}'' \right) - \tilde{\nabla} f_i \left(w - \delta \tilde{\nabla} f_i(w - \alpha \tilde{\nabla} f_i(w, \mathcal{D}), \mathcal{D}'), \mathcal{D}'' \right)}{2\delta}. \quad (121)$$

For this approximation, we have the following result, which shows that we have an additional degree of freedom (δ) to control the error term that does not decreased with increasing batch sizes.

Lemma H.2. *Assume that we estimate $\nabla F_i(w)$ by (120) where \mathcal{D} , \mathcal{D}' , and \mathcal{D}'' are independent batches with size D , D' , and D'' , respectively. Suppose that the conditions in Assumptions 2-4 are satisfied. Then, for any $\alpha \in [0, 1/L]$ and $w \in \mathbb{R}^d$, we have*

$$\begin{aligned} \left\| \mathbb{E} \left[\tilde{\nabla} F_i(w) - \nabla F_i(w) \right] \right\| &\leq m_F^{HF} := \alpha \left(\frac{2L\sigma_G}{\sqrt{D}} + \frac{L\sigma_G}{\sqrt{D'}} + \rho\delta B^2 \right), \\ \mathbb{E} \left[\left\| \tilde{\nabla} F_i(w) - \nabla F_i(w) \right\|^2 \right] &\leq (\tilde{\sigma}_F^{HF})^2 := 6\sigma_G^2 \left(\frac{2(\alpha L)^2}{D} + \frac{2}{D'} + \frac{\alpha^2}{2\delta^2 D''} \right) + 2(\alpha\rho\delta)^2 B^4. \end{aligned}$$

Proof. Note that, this time $\tilde{\nabla}F_i(w)$ is approximating

$$G'_i(w) := \nabla f_i(w - \alpha \nabla f_i(w)) - \alpha d_i(w) \quad (122)$$

where

$$d_i(w) := \frac{\nabla f_i(w + \delta \nabla f_i(w - \alpha \nabla f_i(w))) - \nabla f_i(w - \delta \nabla f_i(w - \alpha \nabla f_i(w)))}{2\delta} \quad (123)$$

is the term approximating $\nabla^2 f_i(w) \nabla f_i(w - \alpha \nabla f_i(w))$. Below, we characterize $\tilde{\sigma}_F^{HF}$, and m_F^{HF} can be done similarly.

Similar to (117), we have

$$\mathbb{E} \left[\left\| \tilde{\nabla}F_i(w) - \nabla F_i(w) \right\|^2 \right] \leq 2\mathbb{E} \left[\left\| \tilde{\nabla}F_i(w) - G'_i(w) \right\|^2 \right] + 2\mathbb{E} \left[\left\| G'_i(w) - \nabla F_i(w) \right\|^2 \right]. \quad (124)$$

We again bound both terms separately. To simplify the notation, let us define

$$g_i(w) := \nabla f_i(w - \alpha \nabla f_i(w)), \quad \tilde{g}_i(w) := \tilde{\nabla}f_i(w - \alpha \tilde{\nabla}f_i(w, \mathcal{D}), \mathcal{D}'). \quad (125)$$

First, note that, using $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$ for $a, b, c \geq 0$, we have

$$\begin{aligned} & \left\| \tilde{\nabla}F_i(w) - G'_i(w) \right\|^2 \\ & \leq 3\|\tilde{g}_i(w) - g_i(w)\|^2 + \frac{3\alpha^2}{4\delta^2} \left\| \tilde{\nabla}f_i(w + \delta \tilde{g}_i(w), \mathcal{D}'') - \nabla f_i(w + \delta g_i(w)) \right\|^2 \\ & \quad + \frac{3\alpha^2}{4\delta^2} \left\| \tilde{\nabla}f_i(w - \delta \tilde{g}_i(w), \mathcal{D}'') - \nabla f_i(w - \delta g_i(w)) \right\|^2. \end{aligned} \quad (126)$$

Taking expectation from both sides, along with using (118), we have

$$\begin{aligned} & \mathbb{E} \left[\left\| \tilde{\nabla}F_i(w) - G'_i(w) \right\|^2 \right] \\ & \leq 3\sigma_G^2 \left(\frac{1}{D'} + \frac{(\alpha L)^2}{D} \right) + \frac{3\alpha^2}{4\delta^2} \left(\mathbb{E} \left[\left\| \tilde{\nabla}f_i(w + \delta \tilde{g}_i(w), \mathcal{D}'') - \nabla f_i(w + \delta g_i(w)) \right\|^2 \right] \right. \\ & \quad \left. + \mathbb{E} \left[\left\| \tilde{\nabla}f_i(w - \delta \tilde{g}_i(w), \mathcal{D}'') - \nabla f_i(w - \delta g_i(w)) \right\|^2 \right] \right) \\ & \leq 3\sigma_G^2 \left(\frac{\alpha^2}{2\delta^2 D''} + \frac{1}{D'} + \frac{(\alpha L)^2}{D} \right) + \frac{3\alpha^2}{4\delta^2} \left(\mathbb{E} \left[\left\| \nabla f_i(w + \delta \tilde{g}_i(w)) - \nabla f_i(w + \delta g_i(w)) \right\|^2 \right] \right. \\ & \quad \left. + \mathbb{E} \left[\left\| \nabla f_i(w - \delta \tilde{g}_i(w)) - \nabla f_i(w - \delta g_i(w)) \right\|^2 \right] \right) \end{aligned} \quad (127)$$

where (127) is obtained using the fact that \mathcal{D}'' is independent from \mathcal{D} and \mathcal{D}' which implies

$$\begin{aligned} & \mathbb{E} \left[\left\| \tilde{\nabla}f_i(w \pm \delta \tilde{g}_i(w), \mathcal{D}'') - \nabla f_i(w \pm \delta g_i(w)) \right\|^2 \right] \leq \frac{\sigma_G^2}{D''} \\ & \quad + \mathbb{E} \left[\left\| \nabla f_i(w \pm \delta \tilde{g}_i(w)) - \nabla f_i(w \pm \delta g_i(w)) \right\|^2 \right]. \end{aligned}$$

Next, note that Assumption 2 yields

$$\left\| \nabla f_i(w \pm \delta \tilde{g}_i(w)) - \nabla f_i(w \pm \delta g_i(w)) \right\| \leq \delta L \|\tilde{g}_i(w) - g_i(w)\|.$$

Plugging this bound into (127) and using (117) implies

$$\begin{aligned} \mathbb{E} \left[\left\| \tilde{\nabla}F_i(w) - G'_i(w) \right\|^2 \right] & \leq 3\sigma_G^2 \left(\frac{\alpha^2}{2\delta^2 D''} + \left(1 + \frac{(\alpha L)^2}{2}\right) \left(\frac{1}{D'} + \frac{(\alpha L)^2}{D} \right) \right) \\ & \leq 3\sigma_G^2 \left(\frac{2(\alpha L)^2}{D} + \frac{2}{D'} + \frac{\alpha^2}{2\delta^2 D''} \right) \end{aligned} \quad (128)$$

where the last inequality is obtained using $\alpha L \leq 1$.

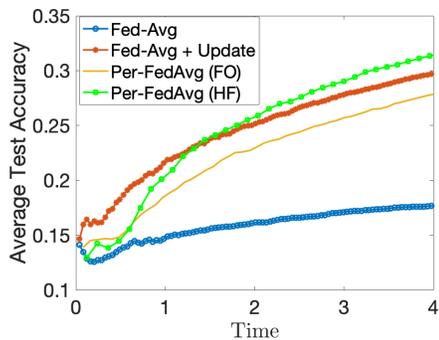
Bounding the second term in (124) is more straightforward as we have

$$\left\| G'_i(w) - \nabla F_i(w) \right\| = \alpha \left\| d_i(w) - \nabla^2 f_i(w) \nabla f_i(w - \alpha \nabla f_i(w)) \right\| \leq \alpha \rho \delta \|g_i(w)\|^2 \leq \alpha \rho \delta B^2. \quad (129)$$

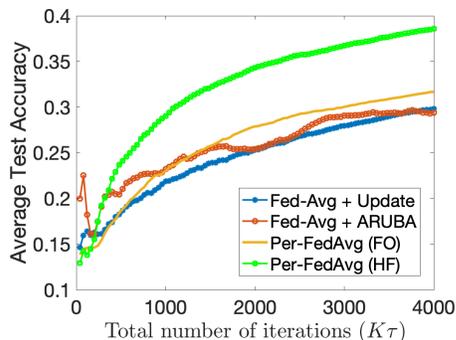
Plugging (128) and (129) into (124) gives us the desired result. \square

Table 2: Illustration of the our experiment’s setting
Image Classes

	1	2	...	5	6	7	...	10
1	a	a	...	a	0	0	...	0
2	a	a	...	a	0	0	...	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
5	a	a	...	a	0	0	...	0
6	a/2	0	...	0	2a	0	...	0
7	0	a/2	...	0	0	2a	...	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
10	0	0	...	a/2	0	0	...	2a



a: Comparison in terms of runtime



b: Comparison in terms of number of iterations

I More on Numerical Experiments

In this section, we discuss our further results on numerical experiments. We thank the anonymous reviewers for their suggestions on adding this results, and we are looking forward to further explore our method from numerical point of view in future works.

First, in Table 2. we provide an illustration of the numerical setting in Section 5.

Second, in Figure 1a, we illustrate the average test accuracy of all studied algorithms with respect to time. As this figure shows, Per-FedAvg (HF) achieves higher level of accuracy compared to the regular Fed-Avg with local updates within the same computation time.

Third, we also compare our method with ARUBA [36]. To do so, we also report the output of FedAvg+ARUBA after refinement for each user. In particular, we consider $\tau = 4$ and $K = 1000$, and also tune hyper-parameters of ARUBA for a fair comparison. The final accuracy of all algorithms is as follows: Per-FedAvg(FO): 34.04 ± 0.08 , Fed-Avg+ARUBA (with refinement): 36.74 ± 0.1 , Per-FedAvg(HF): 43.73 ± 0.11 . In Figure 1b, we have also depicted *one realization* of training path, just to provide intuition on the convergence speed of these methods.