



HAL
open science

Sample Optimality and All-for-all Strategies in Personalized Federated and Collaborative Learning

Mathieu Even, Laurent Massoulié, Kévin Scaman

► **To cite this version:**

Mathieu Even, Laurent Massoulié, Kévin Scaman. Sample Optimality and All-for-all Strategies in Personalized Federated and Collaborative Learning. 2022. hal-03550407

HAL Id: hal-03550407

<https://hal.science/hal-03550407v1>

Preprint submitted on 1 Feb 2022

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

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

Sample Optimality and *All-for-all* Strategies in Personalized Federated and Collaborative Learning

Mathieu Even¹, Laurent Massoulié^{1,2} and Kévin Scaman¹

¹Inria - Département d'informatique de l'ENS

²MSR-Inria Joint Centre

ABSTRACT. In personalized Federated Learning, each member of a potentially large set of agents aims to train a model minimizing its loss function averaged over its local data distribution. We study this problem under the lens of stochastic optimization. Specifically, we introduce information-theoretic lower bounds on the number of samples required from all agents to approximately minimize the generalization error of a fixed agent. We then provide strategies matching these lower bounds, in the *all-for-one* and *all-for-all* settings where respectively one or all agents desire to minimize their own local function. Our strategies are based on a *gradient filtering* approach: provided prior knowledge on some notions of distances or discrepancies between local data distributions or functions, a given agent filters and aggregates stochastic gradients received from other agents, in order to achieve an optimal bias-variance trade-off.

1. INTRODUCTION

A central task in Federated Learning [McMahan et al., 2017, Kairouz et al., 2019] is the training of a common model from local data sets held by individual agents. A typical application is when users (*e.g.* mobile phones, hospitals) want to make predictions (*e.g.* next-word prediction, treatment prescriptions), but each has access to very few data samples, hence the need for collaboration. As highlighted by many recent works (*e.g.* Hanzely et al. [2020], Mansour et al. [2020]), while training a global model yields better statistical efficiency on the combined datasets of all agents by increasing the number of samples linearly in the number of agents, this approach can suffer from a dramatically poor generalization error on local datasets. A solution to this generalization issue is the training of *personalized* models, a midway between a shared model between agents and models trained locally without any coordination.

An ideal approach would take the best of both worlds: increased statistical efficiency by using more samples, while keeping local generalization errors low. This raises the fundamental question: what is the optimal bias/variance tradeoff between personalization and coordination, and how can it be achieved?

We formulate the personalized federated learning problem as follows, studying it under the lens of stochastic optimization [Bottou et al., 2018]. Consider $N \in \mathbb{N}^*$ agents denoted by integers $1 \leq i \leq N$, each desiring to minimize its own local function $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, while sharing their stochastic gradients. Since only a limited number of samples are locally available, we focus on *stochastic gradient descent*-like algorithms, where agents each sequentially compute stochastic gradients g_i^k such that $\mathbb{E}[g_i^k] = \nabla f_i$. In order to reduce the sample complexity, *i.e.* the number of samples or stochastic gradients required to reach small generalization error, agents thus need to use stochastic gradients from other agents, that are *biased* since in general $\mathbb{E}[g_i^k] \neq \nabla f_j$. We first consider the *all-for-one* objective, where a single agent i wants to minimize its local function (local generalization error), using its own information as well as information from agents $j \neq i$. We then address the *all-for-all* objective where all agents want to minimize their local function (local generalization error) in parallel using shared information. Our algorithms are based on a *gradient filtering* approach: in both *all-for-one* and *all-for-all* objectives, upon reception of stochastic gradients $(g_j^k)_j$, agent i *filters* these gradients and aggregates them using some weights λ_j into $\sum_j \lambda_j g_j^k$, in order to achieve some bias/variance trade-off.

Contributions and outline of the paper. In this paper, we consider oracle models where at each step $k = 1, 2, \dots$, one (or all) agent(s) may draw a sample according to its (their) local distribution. We aim at computing the number of stochastic gradients sampled from all agents, required to reach a small generalization error in both *all-for-one* and *all-for-all* formulations, in terms of: biases (distances between functions or distributions), regularity and noise assumptions. The oracle models, main assumptions and problem formulations are given in Section 2. Our main contributions are then as follows:

(i) In Section 3 we prove *information theoretic* lower bounds: to reach a target generalization error $\varepsilon > 0$ for a fixed agent i , no algorithm can achieve a reduction in the number of oracle calls by a factor larger than the total number of agents ε -close –in a suitable sense– to agent i .

(ii) We next study a weighted gradient averaging algorithm for the *all-for-one* problem, matching this lower bound.

(iii) We then propose in Section 5 a parallel extension of the simple weighted gradient averaging algorithm that yields an efficient algorithm for the *all-for-all* problem. In this algorithm, agents compute stochastic gradients at *their* local estimate, and broadcast it to other agents who may use these to update their own estimates. For $x^k = (x_1^k, \dots, x_N^k)$ where x_i^k is the local estimate of agent i at iteration k , updates of the *all-for-all* algorithm write as:

$$x^{k+1} = x^k - \eta W g^k,$$

where $g^k = (g_1^k, \dots, g_N^k)$ for an unbiased stochastic gradient g_i^k of function f_i , a step size η , and a carefully chosen matrix W . Agents i thus use stochastic gradients that are doubly biased, as gradients of a “wrong function” f_j instead of f_i computed at a “wrong location” x_j^k instead of x_i^k .

Related works. *Federated Learning* is a paradigm in machine learning where training is done collaboratively among several agents, taking into account privacy constraints [McMahan et al., 2017, Konečný et al., 2016, Kairouz et al., 2019, Wang et al., 2019]. A central task is the training of a common model for all agents, for which both *centralized* approaches orchestrated by a server and *decentralized* approaches with no central coordinator [Nedich et al., 2018] have been considered. The algorithms we propose in this paper are well suited for a decentralized implementation.

As observed in Hanzely et al. [2020], training a common model for all users can lead to poor generalization on certain tasks such as e.g. next-word prediction. To improve both accuracy and fairness, *personalized* models thus need to be learnt for each agent [Li et al., 2020, Mohri et al., 2019, Yu et al., 2021]. Approaches to personalization include fine-tuning [Cheng et al., 2021, Khodak et al., 2019], transfer learning techniques [Tripuraneni et al., 2020, Wang et al., 2019]. Hanzely et al. [2020], Fallah et al. [2020] among others formulate personalization in FL as the training of local models with a regularization term that enforces collaboration between users. We refer the interested reader to Kulkarni et al. [2020] for a broader survey of Personalized Federated Learning.

While the goal of personalization is to minimize local generalization errors, the above cited works do not provide theoretical guarantees over the sample complexity to obtain small local errors, but instead control errors on a regularized problem, in terms of communication rounds or full gradients used, and not in terms of samples used. Deng et al. [2020], Mansour et al. [2020] among others provide generalization errors under a statistical learning framework that depend on VC-dimensions and on distances between each local data distribution and the mixture of all datasets. Donahue and Kleinberg [2021a,b] study the bias-variance trade-off between collaboration and personalization for mean estimation in a game-theoretic framework. Chayti et al. [2021], Grimberg et al. [2021], Beaussart et al. [2021] frame personalization as a stochastic optimization problem with biased gradients and are the works closest to ours. They consider the training of a single agent with biased gradients from another group of agents, i.e. the *all-for-one* problem and obtain performance guarantees in terms of distance between individual function f_i and the average $N^{-1} \sum_j f_j$. In contrast, we obtain more general performance bounds based on distance bounds between all pairs of functions f_i, f_j (or equivalently, pairs of local distributions). In addition, we prove matching lower bounds. Finally, we also consider the *all-for-all* problem, for which we obtain efficient algorithms, and leverage in that setting the use of weaker bias assumptions than in the *all-for-one* problem.

2. PROBLEM STATEMENT AND ASSUMPTIONS

We now detail our objectives and the necessary technical assumptions. We consider general stochastic gradient methods and formulate our problems, assumptions and algorithms accordingly. Our lower bounds apply to a more specific problem, namely generalization error minimization, or GEM, where functions f_j of all agents j are all obtained from the same loss function ℓ . These lower bounds a fortiori apply to the more general stochastic gradient setup of our algorithms.

Stochastic (sub)-gradients. Let $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, $1 \leq i \leq N$ be agent i ’s local function to be minimized (e.g. the average of a loss function over its local data distribution). At every iteration

$k = 1, 2, \dots$, agent i may access unbiased *i.i.d.* estimates $g_i^k(x)$ of $\nabla f_i(x)$ (or of some subgradient $s(x) \in \partial f_i(x)$, depending on the regularity assumptions made):

$$\mathbb{E} [g_i^k(x)] \in \partial f_i(x), \quad x \in \mathbb{R}^d, k \geq 0, 1 \leq i \leq N.$$

We consider two different objectives. In the *all-for-one* setting, using information from all agents, a fixed agent i desires to minimize its own local function f_i (typically, in collaborative GEM, f_i is the generalization error $f_i(x) = \mathbb{E}_{\mathcal{D}_i}[\ell(x, \xi_i)]$). In the *all-for-all* setting, all agents want to minimize their own local function, using shared information.

Oracle models. To specify the information shared between agents via access to stochastic gradients, we define the following oracles. The *synchronous oracle* (*resp. asynchronous oracle*) lets at every iteration all agents (*resp. only one agent*) sample a stochastic gradient. After K queries from the synchronous oracle (*resp. asynchronous oracle*), each agent will have sampled K (*resp. K/N on average*) stochastic gradients for a total of NK (*resp. K*) in the whole set of agents. In the *all-for-one* setting where a fixed agent i desires to minimize its own local function f_i , all agents j can take the same $x_j^k = x_i^k$ as argument of their stochastic gradient, resulting in a simple stochastic gradient descent with biased gradients. The difficulty in analyzing *all-for-all* algorithms lies in the fact that stochastic gradients sampled by agent j are doubly biased for agent i , being potentially computed at some x_j^k distinct from x_i^k .

Synchronous and asynchronous oracles

At iterations $k = 1, 2, \dots$:

- 1:** A set of agents S^k is chosen: $S^k = \{1, \dots, N\}$ for the *synchronous oracle*, $S^k = \{i_k\}$ where i_k is sampled uniformly at random amongst agents, for the *asynchronous oracle*. ;
- 2:** For all $j \in S^k$, agent j chooses some y_j^k as a (possibly random) function of all previous stochastic gradients and iterates, and samples $g_j^k(y_j^k)$;
- 3:** All agents can then perform an update using these new stochastic gradients and all previous ones.

For fixed target precision $\varepsilon > 0$, the objective is to find, using T_ε samples from all agents in total, queried with synchronous or asynchronous oracles, models with local generalization error ε . We prove matching lower and upper bounds on the *sample complexity* T_ε in the *all-for-one* and *all-for-all* settings.

Regularity, Bias and Noise Assumptions. Throughout the paper, we assume that each function f_i is minimized over \mathbb{R}^d , and we denote by x_i^* such a minimizer.

Bias assumptions on function discrepancies. For some non-negative weights $(b_{ij})_{1 \leq i, j \leq N} \in \mathbb{R}^{+N \times N}$, $(\tilde{b}_{ij})_{1 \leq i, j \leq N} \in \mathbb{R}^{+N \times N}$ and $m \geq 0$, either one of the two following assumptions will be made. Note that if differentiable, we have $\mathbb{E} [g_i^k(x)] = \nabla f_i(x)$.

B.1 For all $1 \leq i, j \leq N$, optimal model x_j^* for j generalizes (with generalization error b_{ij}) for agent i :

$$f_i(x_j^*) - f_i(x_i^*) \leq b_{ij}. \quad (1)$$

B.2 For all $1 \leq i, j \leq N$ and for all $x \in \mathbb{R}^d$, $k \geq 0$:

$$\|\mathbb{E} [g_i^k(x)] - \mathbb{E} [g_j^k(x)]\|^2 \leq \tilde{b}_{ij}. \quad (2)$$

Two sets of regularity and noise assumptions. The two different noise and regularity assumptions we shall consider are as follows.

N.1 For all $1 \leq i \leq N$, f_i is convex and for some $B > 0$, all $k \geq 1$ and $x \in \mathbb{R}^d$,

$$\mathbb{E} \left[\|g_i^k(x)\|^2 \right] \leq B^2. \quad (3)$$

N.2 For all $1 \leq i \leq N$, f_i μ -strongly convex, L -smooth, and we write $\kappa = \frac{L}{\mu}$. For $\sigma > 0$, all $k \geq 1$ and $x \in \mathbb{R}^d$:

$$\mathbb{E} \left[\|g_i^k(x) - \nabla f_i(x)\|^2 \right] \leq \sigma^2. \quad (4)$$

Our regularity assumptions are standard¹; our bias (and gradient dissimilarity) assumptions are less classical, and can be derived from Wassertstein-like distribution-based distances in the GEM setting (Section 6 and Appendix A). Assumption (B.2) will be used in the *all-for-one* setting, while the less restrictive (B.1) will be used in the *all-for-all* setting. It quantifies by how much agents' objectives differ. To the best of our knowledge it has not been previously used in Federated Learning. It is related to (B.2) through $f_i(x_j^*) - f_i(x_i^*) \leq \frac{1}{2\mu} \|\nabla f_i(x_j^*) - \nabla f_i(x_i^*)\|^2 = \frac{1}{2\mu} \|\nabla f_i(x_j^*)\|^2$ under μ -PL or strong convexity assumptions. More generally, we always have $f_i(x_j^*) - f_i(x_i^*) \leq \|x_i^* - x_j^*\| \|\nabla f_i(x_j^*)\|$. Our lower bounds apply to functions that verify both these bias assumptions, and *a fortiori* verify the assumptions made in our upper bounds where only one of these bias assumptions is required.

Finally, our upper bounds will use the following settings, that we refer to as **Setting 1** and **Setting 2**:

- (1) Asynchronous oracle is used, assumption (N.1) holds.
- (2) Synchronous oracle is used, assumption (N.2) holds.

A key instance of our problem is *collaborative generalization error minimization (GEM)*: let \mathcal{D}_i for $1 \leq i \leq N$ be a probability distribution on a set Ξ (agent i 's local distribution, *not* its empirical distribution), $\ell : \mathcal{X} \times \Xi \rightarrow \mathbb{R}$ a loss function, and define the following local objective function, that agent i aims at minimizing:

$$f_i(x) = \mathbb{E}_{\xi_i \sim \mathcal{D}_i} [\ell(x, \xi_i)], \quad x \in \mathcal{X}. \quad (5)$$

Function f_i is thus the generalization error on agent i 's local distribution. Stochastic gradients are in that case of the form $g_i^k(x) = \nabla_x \ell(x, \xi_i^k)$ where $\xi_i^k \sim \mathcal{D}_i$. Counting the number of stochastic gradients used in the whole set of agents to reach a precision ε for f_i thus reduces to computing the number of samples required from all agents to obtain local generalization error ε for agent i . In Appendix A we discuss how our bias assumptions follow from bounds on distances between distributions $\mathcal{D}_i, \mathcal{D}_j$ together with practical scenarios where such bounds can be obtained.

Notation: in the rest of the paper, variables t or T denote the number of stochastic gradients g_i^k sampled (or data item sampled from personal distribution) from all agents, while variables k or K denote the iterates of the algorithms or equivalently to the number of oracle calls made.

3. IT LOWER BOUNDS ON THE SAMPLE COMPLEXITY

In this section, we prove lower bounds on the sample complexity of the *all-for-one* problem. Our lower bounds apply to collaborative GEM, i.e. functions $(f_i)_{1 \leq i \leq N}$ of the form (5), for some shared loss function ℓ and user distributions $\mathcal{D}_1, \dots, \mathcal{D}_N$.

An oracle $\phi : \mathbb{R}^{N \times d} \rightarrow \mathcal{I}$ is a random function that answers some $\phi(x) \in \mathcal{I}$ where \mathcal{I} is an information set, for every query $x \in \mathbb{R}^{N \times d}$. We adapt the definitions of Agarwal et al. [2012] of sample complexity for *SGD* to our personalization problem. Formally, the *first-order* oracle we defined in Section 2 (for either synchronous or asynchronous oracles) and that we write as $\phi((\mathcal{D}_i)_{i=1, \dots, N}, \ell)$ for shared loss function ℓ and user distributions $\mathcal{D}_1, \dots, \mathcal{D}_N$, returns for $x \in \mathbb{R}^{N \times d}$:

$$\phi((\mathcal{D}_i)_i, \ell)(x) = \left(i, x_i, \xi_i, \ell(x_i, \xi_i), g_i^k(x_i) \right)_{i \in S^k},$$

where $\xi_i \sim \mathcal{D}_i$ and S^k is the set of agents returning a stochastic gradient at step k , according to the synchronous or asynchronous oracle in use. Given distributions and a loss function $((\mathcal{D}_i)_i, \ell)$, we denote by \mathbb{M} the set of all methods $\mathcal{M} = (\mathcal{M}_T)_{T \geq 0}$: for any $T \geq 0$, \mathcal{M}_T makes K oracle calls from oracle $\phi((\mathcal{D}_i)_i, \ell)$ while using T stochastic gradient samples from all agents ($T = NK$ with the synchronous oracle, $T = K$ with the asynchronous oracle), and returns $x_i^T \in \mathbb{R}^d$ for agent i . For a set \mathbb{D} of couples of distributions and loss function $((\mathcal{D}_j)_j, \ell)$ defining functions $(f_i)_{1 \leq i \leq N}$, we are interested in lower-bounding:

$$\inf_{\mathcal{M} \in \mathbb{M}} \sup_{((\mathcal{D}_j)_j, \ell) \in \mathbb{D}} \mathcal{T}_i^\varepsilon \left(\mathcal{M}, ((\mathcal{D}_j)_j, \ell) \right),$$

¹In fact our approach could be extended to other regularity assumptions, e.g. functions f_i with $\mu = 0$ as in Appendix E.4, non-convex, or satisfying a Polyak-Lojasiewicz inequality; we do not pursue this here for lack of space. We also generalize (B.2) to function dissimilarities such as those in Karimireddy et al. [2020], in Appendix D, (B.3).

where $\mathcal{T}_i^\varepsilon(\mathcal{M}, ((\mathcal{D}_j)_j, \ell))$ is the number of samples required from all agents (uniformly divided between agents) to reach generalization error $\varepsilon > 0$ for agent i and writes as:

$$\inf \left\{ T \in \mathbb{N}^* \text{ such that } \mathbb{E} \left[f_i(x_i^T) - \min_{x \in \mathcal{X}} f_i(x) \right] \leq \varepsilon \right\}.$$

For $(b_{ij}) \in \mathbb{R}^{+N \times N}$, $r > 0$, $B > 0$, let $\mathbb{D}(r, b, B)$ be the set of all $((\mathcal{D}_i)_{1 \leq i \leq N}, \ell)$ for probability distributions $(\mathcal{D}_i)_{1 \leq i \leq N}$ on a probability space $\Xi \subset \mathbb{R}^D$ such that all functions f_i parameterized by the losses and distributions in this set verify $\|x_i^*\| \leq r$, assumption (N.1), and bias assumptions (B.1) and (B.2) for b and $\tilde{b}_{ij} = (b_{ij}/r)^2$.

Similarly, for $\sigma > 0$, $\mu, L > 0$, let $\mathbb{D}_\mu^L(r, b, B)$ be the set of all $((\mathcal{D}_i)_{1 \leq i \leq N}, \ell)$ for probability distributions $(\mathcal{D}_i)_{1 \leq i \leq N}$ on a probability space $\Xi \subset \mathbb{R}^D$ such that all functions f_i parameterized by this set verify $\|x_i^*\| \leq r$, assumption (N.2), and bias assumptions (B.1) and (B.2) for b and $\tilde{b} = \mu b$ here.

Theorem 1 (IT lower bound). *Let $\varepsilon \in (0, 1/16)$, and assume that either the synchronous or asynchronous oracle is used, that (b_{ij}) verifies the triangle inequality $b_{ij} \leq b_{ik} + b_{kj}$ for all $1 \leq i, j, k \leq N$. For some constant $C > 0$ independent of the problem and any $i \in \{1, \dots, N\}$:*

$$\inf_{\mathcal{M} \in \mathbb{M}((\mathcal{D}_j)_j, \ell) \in \mathbb{D}(r, b, B)} \sup \mathcal{T}_i^\varepsilon(\mathcal{M}, ((\mathcal{D}_j)_j, \ell)) \geq \frac{Cr^2 B^2 N}{\varepsilon^2 \mathcal{N}_i^\varepsilon(\frac{b}{4})},$$

where $\mathcal{N}_i^\varepsilon(b) = \sum_j \mathbb{1}_{\{b_{ij} \leq \varepsilon\}}$ is the number of agents j verifying $b_{ij} \leq \varepsilon$. Under strong-convexity and smoothness assumptions for $\mu = L = 1/r^2$, we have, for some $C' > 0$:

$$\inf_{\mathcal{M} \in \mathbb{M}} \sup_{((\mathcal{D}_j)_j, \ell) \in \mathbb{D}_{\mu=1/r^2}^{L=1/r^2}(r, b, \sigma)} \mathcal{T}_i^\varepsilon(\mathcal{M}, ((\mathcal{D}_j)_j, \ell)) \geq \frac{C'r^2 \sigma^2 N}{\varepsilon \mathcal{N}_i^\varepsilon(\frac{b}{4})}.$$

The proof of these lower bounds (Appendix B) builds on lower bounds based on Fano's inequality [Duchi and Wainwright, 2013] for stochastic gradient descent Agarwal et al. [2012] or for information limited statistical estimation [Zhang et al., 2013, Duchi and Rogers, 2019], adapted to personalization.

Lower-bound interpretation. Theorem 1 states that, given the knowledge of (b_{ij}) and B^2 (or σ^2 , μ and L), there exist difficult instances of the problem that satisfy assumption (N.1) (or (N.2)) and both bias assumptions (B.1) and (B.2) for b , such that the number of samples from all agents (generated through the synchronous or asynchronous oracle) needed to obtain a generalization error of ε for an agent i is lower-bounded by the right hand sides of the equations in Theorem 1.

Collaboration speedup. The factor $N \times CB^2 r^2 \varepsilon^{-2}$ is reminiscent of stochastic gradient descent, and is present in Agarwal et al. [2012]: without cooperation, this is the sample complexity of SGD for a fixed agent (for T samples from all agents, only T/N are taken from a given agent i). Cooperation appears in the factor $1/\mathcal{N}_i^\varepsilon(b/4)$: the sample complexity is inversely proportional to the number of agents j that have functions (or distributions) similar to that of i . One cannot hope for better than a linear speedup proportional to agents $\varepsilon/4$ -close to i , in the sense that the functions parameterized by the loss function and distributions verify $f_i(x_j^*) - f_i(x_i^*) \leq b_{ij}/4$ and $\|\nabla f_i(x) - \nabla f_j(x)\|^2 \leq \tilde{b}_{ij}$ for all $x \in \mathbb{R}^d$, where $\tilde{b} = (b/r)^2$ in the first case, and $\tilde{b} = b/r^2 = \mu b$ in the second.

Bias assumptions and collaborative GEM. In Section 6, we relate $\|\nabla f_i - \nabla f_j\|$ to distribution-based distances $d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j)$ in the GEM setting, where $d_{\mathcal{H}}$ is a 1-Wassertstein-like distance. We prove that the distributions leading to the functions used to prove Theorem 1 verify $d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq b_{ij}$; thus, the distribution based distances also sharply characterize the sample complexity. The requirement for the lower bound to apply that (b_{ij}) verifies the triangle inequality is thus quite natural, when translating bias assumptions into distribution-based distances.

4. WEIGHTED GRADIENT AVERAGING IS OPTIMAL IN THE *All-for-one* SETTING

We now prove that a weighted gradient averaging (WGA) algorithm is optimal in the *all-for-one* setting. Fix the agent i who desires to minimize its local function f_i (no longer assumed to derive from a loss function). The iterates of the weighted gradient averaging algorithm write as, for both

synchronous and asynchronous oracles:

$$x^{k+1} = x^k - \eta \sum_{j=1}^N \lambda_j G^k(x)_j, \quad (6)$$

for some step size $\eta > 0$ and stochastic vector $\lambda \in \mathbb{R}^N$, and:

$$G^k(x)_i = \begin{cases} N \mathbb{1}_{\{i=i_k\}} g_{i_k}^k(x) & (\text{asynchronous}) \\ g_i^k(x) & (\text{synchronous oracle}). \end{cases} \quad (7)$$

In the WGA algorithm, all agents involved keep the same local estimates x_i^k . As defined in Section 2, (g_i^k) are *i.i.d.* non-biased stochastic (sub)gradients of functions f_i . The WGA algorithm thus defined is more general than previous ones, through the introduction of parameter λ . In the GEM setting, WGA is equivalent to training a model on the mixture of distributions $(\mathcal{D}_j)_j$, with weights $(\lambda_j)_j$. We have the following convergence guarantees in terms of *sample complexity* for the minimization of a function f_i with WGA iterations.

Theorem 2. *Let $\varepsilon > 0$, and set $\lambda_j = \frac{\mathbb{1}_{\{b_{ij} < \varepsilon/2\}}}{\mathcal{N}_i^\varepsilon(2b)}$, $\mathcal{N}_i^\varepsilon$ being defined in Theorem 1. Generate $(x^k)_k$ using (6).*

2.1 Under Setting 1. *Assume that for all $k \geq 0$ and i , $\|x^k - x_i^*\| \leq D$ for some $D > 0$ and bias assumption (B.2) holds for $\tilde{b} = (b/D)^2$. For $K = T_\varepsilon$ and for $\eta = \sqrt{\frac{D^2}{2KNB^2 \sum_{j=1}^N \lambda_j^2}}$, we have $\mathbb{E} \left[\frac{1}{K} \sum_{0 \leq k < K} f_i(x^k) - f_i(x_i^*) \right] \leq \varepsilon$ using a total number $T_\varepsilon(i)$ of stochastic gradients from all agents of:*

$$T_\varepsilon(i) \leq \frac{4D^2B^2}{\varepsilon^2} \frac{N}{\mathcal{N}_i^\varepsilon(2b)}.$$

2.2 Under Setting 2. *Denote $\kappa = L/\mu$, assume that one of bias assumptions (B.2) or (B.3) holds for $\tilde{b} = \mu b$ (and any $m \geq 0$). For $K = T_\varepsilon/N$ and $\eta = \min(1/(2L), \frac{1}{\mu K} \ln(\frac{F_\lambda^0 \mu^2 K}{L\sigma^2} \sum_j \lambda_j^2))$ where $F_\lambda^0 = \sum_j \lambda_j (f_j(x^0) - f_j(x_j^*))$, we have $\mathbb{E} [f_i(x^K) - f_i(x_i^*)] \leq \varepsilon$ using a total number $T_\varepsilon(i)$ of stochastic gradients from all agents of:*

$$T_\varepsilon(i) \leq \tilde{\mathcal{O}} \left(\frac{\kappa \sigma^2}{\mu \varepsilon} \frac{N}{\mathcal{N}_i^\varepsilon(2b)} \right).$$

The personalization-dependent factor of both these sample complexities (the second factor $\frac{N}{\mathcal{N}_i^\varepsilon(2b)}$) matches that of the lower bound in Theorem 1. While in Theorem 2.1 upper and lower bounds (up to constant factors) exactly match, in the strongly convex and smooth case (Theorem 2.2), the lower bound is for $\kappa = 1$, and in that particular case, lower and upper bounds match; we conjecture that in the more general case, the optimal sample complexity as a linear dependency in $\sqrt{\kappa}$. In terms of *gradient filtering*, an optimal strategy is thus to filter and keep stochastic gradients sent by agents that are ε -close, in the sense of (B.2). Theorem 4 presents a tuning of the algorithm for a fixed target precision ε ; using a *doubling trick* as we do in Appendix E.3 yields time-adaptive algorithms, that achieve the same sample complexity (up to constant factors) for any $\varepsilon > 0$.

In the GEM setting, the proposed WGA algorithm consists in training a model over the distribution $\mathcal{D}_i^\varepsilon$, a mixture of the distributions $\{\mathcal{D}_j : d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq \varepsilon/2\}$, where $d_{\mathcal{H}}$ is a distribution-based distance, leading to an increase in the number of samples available by a factor $\mathcal{N}_i^\varepsilon$ (thus reducing the variance), while keeping the bias of order ε .

5. All-for-all STRATEGIES

We now focus on the more challenging *all-for-all* setting: all agents desire to minimize their local function. We first present a naive approach, that proves to be sample-optimal, but that cannot be used in practice for large scale problems, as it requires a number computations per oracle call for each agent that increases with N . We thus introduce and analyze the *all-for-all* algorithm, the main algorithmic contribution of our paper.

Naive Approach. Each agent i keeps N shared local models x_1^k, \dots, x_N^k , where x_j^k estimates x_j^* at iteration k (the knowledge of x_j^k needs to be shared by all agents). At each iteration k , when a sample ξ_j^k is obtained at agent j , it is used by that agent to compute unbiased estimates of $\nabla f_j(x_j^k)$ for all $i \in [N]$. These are then broadcast to all agents i that verify $b_{ij} \leq \varepsilon$ (target precision). Agents then perform N weighted gradient averaging algorithms simultaneously, leading to the following guarantees. Under the assumptions of Theorem 2.1, we have $\max_{1 \leq i \leq N} \mathbb{E} [f_i(x_i^k) - f_i(x_i^*)] \leq \varepsilon$, with a number of data items sampled from personal distribution T_ε of:

$$T_\varepsilon \leq \frac{4D^2B^2}{\varepsilon^2} \max_{1 \leq i \leq N} \frac{N}{\mathcal{N}_i^\varepsilon(2b)}.$$

Under the assumptions of Theorem 2.2, we have $\max_{1 \leq i \leq N} \mathbb{E} [f_i(x_i^k) - f_i(x_i^*)] \leq \varepsilon$, with a number of data item sampled from personal distribution T_ε of:

$$T_\varepsilon \leq \tilde{\mathcal{O}} \left(\frac{\kappa\sigma^2}{\mu\varepsilon} \max_{1 \leq i \leq N} \frac{N}{\mathcal{N}_i^\varepsilon(2b)} \right).$$

Using the *all-for-one* lower bounds of Theorem 1, this proves to be optimal. Yet, the memory requirements and computation/communication costs of this approach are forbiddingly high for large N (they scale with $\mathcal{N}_i^\varepsilon$ for agent i), hence the *all-for-all* algorithm below, that relaxes the metric considered by controlling the averaged local errors $N^{-1} \sum_{1 \leq i \leq N} \mathbb{E} [f_i(x_i^k) - f_i(x_i^*)]$ instead of worst-case ones $\max_{1 \leq i \leq N} \mathbb{E} [f_i(x_i^k) - f_i(x_i^*)]$. Furthermore, relaxing the quantity controlled leads to weaker necessary bias assumptions.

Algorithm 1 *All-for-all* algorithm

- 1: Step size $\eta > 0$, matrix $W \in \mathbb{R}^{N \times N}$
 - 2: Initialization $x_1^0 = \dots = x_N^0 \in \mathbb{R}^d$ (x_i^0 at agent i)
 - 3: **for** $k = 0, 1, 2, \dots, K - 1$ **do**
 - 4: Agents $j \in S^k$ (activated agents) compute stochastic gradients $g_j^k(x_j^k)$ and broadcast it to all agents i such that $W_{ij} > 0$.
 - 5: For $i = 1, \dots, N$, update

$$x_i^{k+1} = x_i^k - \eta \sum_{j \in S^k} W_{ij} g_j^k(x_j^k)$$
 - 6: **end for** Return x_i^K for agent i
-

The All-for-all algorithm. We now present the *all-for-all* algorithm (AFA), an adaptation of the weighted gradient averaging algorithm to the *all-for-all* setting, where all agents desire to use all stochastic gradients computed. For $1 \leq i \leq N$, initialize $x_i^0 = x_0 \in \mathbb{R}^d$. At iteration k , let $x_i^k \in \mathbb{R}^d$ be agent i 's current estimate of x_i^* , and denote $x^k = (x_i^k)_{1 \leq i \leq N} \in \mathbb{R}^{N \times d}$. For a step size $\eta > 0$ and a symmetric non-negative matrix $W \in \mathbb{R}^{N \times N}$, iterates of the *all-for-all* algorithm are generated with Algorithm 1. In Theorem 3, we control the averaged local generalization error amongst all agents:

$$F^k = \frac{1}{N} \sum_{i=1}^N f_i(x_i^k) - f_i(x_i^*), \quad k \geq 0.$$

Theorem 3 (*All-for-all* algorithm). *Let $K > 0$, $\eta > 0$, and W a symmetric non-negative random matrix of the form $W = \Lambda \Lambda^\top$ for some stochastic matrix $\Lambda = (\lambda_{ij})_{1 \leq i, j \leq N}$. Let $(x_i^k)_{k \geq 0, 1 \leq i \leq N}$ be generated with Algorithm 1. Assume that bias assumption (B.1) holds for some (b_{ij}) .*

3.a Under Setting 1. *Using step size $\eta = \sqrt{\frac{2N \|x^0 - x^\Lambda\|^2}{KB^2 \sum_{i,j} \lambda_{ij}^2}}$:*

$$\mathbb{E} \left[\frac{1}{K} \sum_{k=0}^{K-1} F^k \right] \leq \sqrt{\frac{2B^2 \sum_{i=1}^N \|x_i^0 - x_i^\Lambda\|^2}{NK} \sum_{1 \leq i, j \leq N} \lambda_{ij}^2} + \frac{1}{N} \sum_{1 \leq i, j \leq N} \lambda_{ij} b_{ij},$$

3.b Under Setting 2. If $\eta = \frac{1}{\mu K} \ln\left(\frac{NF^0\mu^2K}{L\sigma^2\sum_{ij}\lambda_{ij}^2}\right) \wedge \frac{1}{2L}$:

$$\begin{aligned} \mathbb{E}[F^K] &\leq F^0 e^{-\frac{\kappa}{2L}} + \tilde{O}\left(\frac{L\sigma^2}{K\mu^2N} \sum_{1\leq i,j\leq N} \lambda_{ij}^2\right) \\ &\quad + \frac{1}{N} \sum_{1\leq i,j\leq N} \lambda_{ij} b_{ij}. \end{aligned}$$

In 3.a, x^Λ is a minimizer of the function (8). In terms of sample complexity, we have the following.

Theorem 4 (All-for-all sample complexity). Let $\varepsilon > 0$, and set $W = \Lambda\Lambda^\top$ for $\lambda_{ij} = \frac{\mathbb{1}_{\{b_{ij}\leq\varepsilon/2\}}}{N_i^\varepsilon(2b)}$.

(1) Under the same assumptions as in Theorem 3.a, we have $\mathbb{E}\left[\frac{1}{K}\sum_{k=0}^{K-1} F^k\right] \leq \varepsilon$ for a total number T_ε of data item sampled from personal distribution from all agents of:

$$T_\varepsilon \leq \frac{4D^2B^2}{\varepsilon^2} \sum_{i=1}^N \frac{1}{N_i^\varepsilon(2b)},$$

where D^2 bounds all $\|x_i^0 - x_i^\Lambda\|^2$, $1 \leq i \leq N$.

(2) Under the same assumptions as in Theorem 3.b, we have $\mathbb{E}[F^K] \leq \varepsilon$ for a total number T_ε of data item sampled from personal distribution from all agents of:

$$T_\varepsilon \leq 2 \max\left(\frac{\kappa\sigma^2}{\varepsilon\mu} \sum_{i=1}^N \frac{1}{N_i^\varepsilon(2b)}, N\kappa\right) \ln(\varepsilon^{-1}F^0).$$

Collaborative speedup. In Setting 1, denoting $T_\varepsilon(i)$ the sample complexity in the *all-for-one* setting (that matches the corresponding lower bound), we observe that in the *all-for-all* regime, we obtain $T_\varepsilon \leq \frac{1}{N} \sum_i T_\varepsilon(i)$. The speedup in comparison with a no-collaboration strategy (all agents locally performing SGD) is $\frac{1}{N} \sum_i \frac{1}{N_i^\varepsilon(2b)}$: the mean of all *all-for-one* speed-ups. Similarly in Setting 2, the speed-up in comparison with the no-collaboration setting is still $\frac{1}{N} \sum_i \frac{1}{N_i^\varepsilon(2b)}$ in the statistical regime. In Setting 2, one could obtain $\sqrt{\kappa}$ instead of κ in the optimization term, using accelerated gradient methods (with additive noise here), leading to a faster convergence to the statistical regime (first term of the max in Theorem 4.2), where we have the collaboration speedup. We present in Appendix E.3 a time-adaptive variant with varying step sizes and matrices, that share the same sample complexity, and Theorem 8 in Appendix E.4 is the case $\mu = 0$.

Assumption (B.1). Importantly, controlling the mean of local generalization errors leads to a much weaker bias assumption: instead of requiring a uniform control of gradient norms (as is done in all previous works, and in our *all-for-one* setting), our *all-for-all* algorithm only requires that for some (b_{ij}) , we have $f_i(x_j^*) - f_i(x_i^*) \leq b_{ij}$: in the GEM setting *e.g.*, even when gradients of f_i and f_j may differ a lot, if agent j 's optimal model generalizes well-enough under agent i 's distribution, they should collaborate. We believe this notion of function proximity that we leverage in this setting to be the weakest possible.

Intuition behind the algorithm. Perhaps surprisingly, matrix W is in general *not* a gossip matrix (*i.e.* such that $W\mathbf{1} = \mathbf{1}$): agent i does not aggregate a convex combination of stochastic gradients, but a combination with scalars that do not necessarily sum to 1. In the collaborative GEM setting, we thus cannot say that the *all-for-all* algorithm acts as if, in parallel, each agent i trains a model on the mixture of distributions \mathcal{D}_j with weights W_{ij} . In fact, as the analysis shows below, agent i trains a model on the mixture of distributions, with weights λ_{ij} , if Λ is a stochastic square root of matrix W ($\Lambda\Lambda^\top = W$). Thus, the *all-for-all* algorithm is exactly a paralleled version of weighted gradient averagings, where agent i uses the stochastic vector $(\lambda_{ij})_{1\leq j\leq N}$. In order to account for inter-dependencies between agents that do not directly share information, the *all-for-all gradient filtering* uses weights W_{ij} to aggregate information, instead of λ_{ij} . Propagating information using a matrix W , that induces a similarity graph G_W on $\{1, \dots, N\}$, such that $(ij) \in E_W$ if $W_{ij} > 0$, is quite natural [Vanhaesebrouck et al., 2017, Bellet et al., 2018]; yet, ours is the first analysis to give such precise generalization error bounds, through the use of a stochastic optimization framework.

Degrees of freedom offered by W . In comparison to Theorem 3, the classical personalized FL approaches that consider personalized local models of the form $x_i = \bar{x} - \delta_i$, where \bar{x} is some global quantity shared by all agents, perturbed (and personalized) by some local quantity δ_i (e.g. averaging between local and a global models), can be seen as the special instances where, for all i , we have $\lambda_{ii} = 1 - \alpha_i$ and $\lambda_{ij} = \frac{\alpha_i}{N-1}$ if $i \neq j$ for some α_i , and leads to bias terms of the form $\frac{1}{N} \sum_i \frac{\alpha_i}{N-1} \sum_{j \neq i} b_{ij}$. Full and naive collaboration (a single model trained for all users) corresponds to $\lambda_{ij} = 1/N$ for all i, j , and leads to a bias term of $\frac{1}{N^2} \sum_{i,j} b_{ij}$. The degrees of freedom offered by our matrix W (and by coefficients λ_{ij}) enable pairwise agent adaptation, and tighter generalization guarantees and bias/variance tradeoffs.

Proof sketch of convergence guarantees. Since brutally analyzing convergence of the iterates (x^k) generated with $x^{k+1} = x^k - WG^k$ seems impossible due to both gradient biases and model biases between agents, we study these iterates through the introduction of a different but related problem. This approach is in fact similar to some decentralized optimization ones, where a dual problem or a related energy function is often introduced [Scaman et al., 2019, Even et al., 2021], upon which well-studied algorithms are applied. The related problem we formulate is different from and more flexible than all the different personalized FL problems in the literature [Hanzely et al., 2020, T. Dinh et al., 2020], that consider regularization terms that enforce consensus. For $\lambda = (\lambda_{ij})_{1 \leq i, j \leq N}$ a stochastic matrix (such that for all $1 \leq i \leq N$, we have $\sum_{j=1}^N \lambda_{ij} = 1$), let f^Λ be defined as:

$$f^\Lambda(y) = \bar{f}(\Lambda y), \quad y \in \mathbb{R}^{N \times d}, \quad (8)$$

where $\bar{f} = \frac{1}{N} \sum_i f_i$. Gradient descent on f^Λ writes as

$$y^{k+1} = y^k - \eta \Lambda^\top \nabla \bar{f}(\Lambda y^k)$$

where $\nabla \bar{f}(x) = \frac{1}{N} (\nabla f_i(x_i))_{1 \leq i \leq N}$ for any $x \in \mathbb{R}^{N \times d}$. Importantly, notice that denoting $x^k = \Lambda y^k$ and since $W = \Lambda \Lambda^\top$, we have the recursion

$$x^{k+1} = x^k - \eta W \nabla \bar{f}(x^k)$$

making an analysis of the iterates (x^k) possible. In our case, we however use stochastic gradients given by our oracles. The full gradient $\nabla \bar{f}(x)$ is thus replaced by $(G^k(x_i)_i)_i$ defined at Equation (7). Defining $(y^k)_k$ with the recursion:

$$y^{k+1} = y^k - \eta ((\Lambda G^k(y^k)_i)_i),$$

initialized at $y_1^0 = x_1^0 = \dots = y_N^0 = x_N^0$ we have $x^k = \Lambda y^k$, for all $k \geq 0$, where (x^k) is generated using Algorithm 1. As a consequence, controlling in Theorem 3 the function values $\frac{1}{N} \sum_i f_i(x_i^k)$ is equivalent to controlling $f^\Lambda(y^k)$ (these two quantities are equal). The bias-variance trade-off thus writes as, where x^Λ minimizes f^Λ and $x^* = (x_i^*)_{1 \leq i \leq N}$:

$$F^k \leq \underbrace{f^\Lambda(x^\Lambda) - \bar{f}(x^*)}_{\text{Bias term}} + \underbrace{f^\Lambda(y^k) - f^\Lambda(x^\Lambda)}_{\text{Optimization and variance terms}}.$$

□

6. COLLABORATIVE GEM

In this section, we place ourselves in the GEM setting (loss function ℓ , distributions $(\mathcal{D}_i)_i$). We briefly elaborate (see Appendix A for further details) on the collaborative GEM setting we used through the paper, clearly define the distribution-based distances related to bias assumptions (B.1) and (B.3), and introduce toy problems under which the bias upperbounds may be known.

6.1. Distribution-based distances. For \mathcal{H} a set of functions from Ξ to \mathbb{R}^d and $\mathcal{D}, \mathcal{D}'$ two probability distributions on Ξ , we define:

$$d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}') = \sup_{h \in \mathcal{H}} \|\mathbb{E}[h(\xi) - h(\xi')]\|,$$

where $\xi \sim \mathcal{D}$ and $\xi' \sim \mathcal{D}'$. $d_{\mathcal{H}}$ is a pseudo-distance on the set of probability measures on Ξ . The definition of $d_{\mathcal{H}}$ is motivated by the fact that, for fixed $x \in \mathcal{X}$ and $1 \leq i, j \leq N$, we have:

$$\|\nabla f_i(x) - \nabla f_j(x)\| \leq d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j),$$

if $(\xi \in \Xi \mapsto \nabla_x \ell(x, \xi)) \in \mathcal{H}$ where $(\xi_i, \xi_j) \sim \mathcal{D}_i \times \mathcal{D}_j$. Thus, if for all $1 \leq i, j \leq N$, $d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq \hat{b}_{ij}$ for some weights \hat{b}_{ij} , bias assumption (B.1) holds for $b_{ij} = \sup_{k,l} \|x_k^* - x_l^*\| \hat{b}_{ij}$ (and for $\hat{b}_{ij}^2/(2\mu)$ under μ -PL assumption), and bias assumption (B.2) holds for $\tilde{b}_{ij} = \hat{b}_{ij}^2$.

6.2. Weakly supervised setting. In a scenario where, from a large pool of unlabelled data (distributions $(\hat{\mathcal{D}}_i)_i$), requiring a label led to a sample from \mathcal{D}_i but were costly, agents would benefit from computing distance-based distributions between unlabelled distributions $(\hat{\mathcal{D}}_i)_i$ (possible thanks to a large amount of unlabelled samples); these computed distances would then help to reduce the number of labelled samples required.

6.3. Geometric structure of the agents and infinitely-many-agents. A toy problem is to assume that agent distributions are drawn in an *i.i.d.* fashion from a continuous set of possible distributions $\Theta \subset \mathbb{R}^p$: $\theta_1, \dots, \theta_N \sim \nu$ where ν is a density over the set Θ . Making $N \rightarrow \infty$, we end up with a continuum of agents of distributions $(\mathcal{D}_\theta)_{\theta \in \Theta}$ with density ν . In that setting, under a bias assumption of the form $d_{\mathcal{H}}(\mathcal{D}_\theta, \mathcal{D}_{\theta'}) \leq \hat{b}(\theta, \theta') \leq \|\theta - \theta'\|^q$ for all θ, θ' , the collaboration speedup of Theorem 4.1 (for instance), writes as:


$$\int_{\Theta} (f_\theta(x_\theta) - f_\theta(x_\theta^*)) \nu(d\theta) \leq \varepsilon,$$

with a total number of samples drawn of

$$T_\varepsilon \leq \mathcal{O}_{\varepsilon \rightarrow 0} \left(B^2 D^2 \varepsilon^{-2-p/q} \right),$$

under regularity assumptions on ν and \hat{b} . This approach is made more rigorous in Appendix A.4.

6.4. Illustration of our theory. To test the robustness of our algorithms and our theory, we use in our experiments the mean estimation problem used in the strongly convex and smooth lower bound of Theorem 1: $\ell(x, \xi) = \frac{1}{2} \|x - \xi\|^2$, $x \in \mathbb{R}^d$ for ξ a d -dimensional Bernoulli random variable. We use the time-adaptive version of Algorithm 1, and we place ourselves in the setting of Section A.4, where agent distributions are drawn from a distribution of distributions: for $N = 100$ agents, we draw $(p_i)_{1 \leq i \leq N}$ *i.i.d.* uniformly distributed in $[0, 1]$, and p_i is the parameter of agent i 's Bernoulli variables. We draw 10^3 samples for each agent, and we compute for $1 \leq t \leq T$ the averaged local generalization errors (here the error from the mean), namely $\frac{1}{N} \sum_{i=1}^N \frac{1}{2} \|x_i^t - p_i\|^2$, where x_i^t is the output of the algorithm after t samples drawn. As expected and as illustrated in Figure 1(a), our *all-for-all* approach benefits from both *no-collaboration* (each agent locally estimating its mean p_i with only its locally available samples, in orange in the graph) and *single-model* approaches (a fully centralized minimization of $\frac{1}{N} \sum_i f_i(x)$, in green in the graph), through both a convergence to the true mean, and a non-asymptotic acceleration. In Figure 1(b), we study the effect of noise on the estimation of bias parameters b_{ij} (that here correspond to $b_{ij} = \frac{1}{2}(p_i - p_j)^2$), by taking as inputs in algorithms $b_{ij}^{\text{noisy}} = \frac{1}{2}(p_i + n_i - p_j - n_j)^2$, where n_i are *i.i.d.* uniformly distributed in $[-\text{noise}, \text{noise}]$, for different noise values (Figure 1(b)). *All-for-all* algorithm appears to be quite robust to noise: for small noise amplitudes (0.02 and 0.08, and even 0.18), performances are not too degraded. For (abusively) large noise values (0.32 and 0.5), the non-asymptotic speedup is kept,



(a) Comparison of *All-for-all* with naive benchmarks (b) Effect of noise in the estimation of bias parameters b_{ij}

FIGURE 1. **All-for-all** algorithm in practice

with degraded asymptotic performances. Still, in the range of parameters considered, these perform better than no-collaboration and single model approaches.

DISCUSSION AND CONCLUSION.

In this paper, we quantified in term of function biases (b_{ij}), stochastic gradient noise or amplitudes (σ^2 or B^2), target precision $\varepsilon > 0$ and functions regularity parameters, the benefit of collaboration between agents for shared minimization using stochastic gradient algorithms. Our lower bound (Theorem 1) states that in the *all-for-one* setting, assumption parameters being fixed, no algorithm matches the performances of weighted gradient averagings in terms of sample optimality. Another lower bound (and corresponding upper-bounds) that would be worth investigating is: without any knowledge on biases b_{ij} , what is the worst case complexity of an algorithm that would thus need to *learn who to learn with*? We extended weighted gradient averagings to the *all-for-all* setting, technically more challenging, through the introduction of a related problem, simplifying the analysis of such algorithms.

The main limitation of our work lies in the assumption that upper-bounds on the biases b_{ij} are known. We investigate in Section A scenarii in which this assumption is valid. Yet, as mentioned above, *learning who to learn with and how* is a challenging question worth tackling. Extending our work to model agnosticism together with data heterogeneity would require bias assumptions such as our distribution-based ones in Section A. Finally, fairness issues might be raised by our approaches in the *all-for-all* setting: we consider the average errors amongst agents and therefore do not ensure bounds on the supremum of all local errors. Still, incentives to collaborate and send gradients to other users are the hope of being helped back (matrix W is symmetric in Algorithm 1).

REFERENCES

- Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, and Martin J. Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, 2012.
- Heinz H. Bauschke, Jérôme Bolte, and Marc Teboulle. A descent lemma beyond lipschitz gradient continuity: First-order methods revisited and applications. *Mathematics of Operations Research*, 42(2):330–348, 2017.
- Martin Beaussart, Felix Grimberg, Mary-Anne Hartley, and Martin Jaggi. WAFFLE: Weighted Averaging for Personalized Federated Learning. *arXiv:2110.06978 [cs]*, October 2021.
- Aurélien Bellet, Rachid Guerraoui, Mahsa Taziki, and Marc Tommasi. Personalized and private peer-to-peer machine learning. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 473–481, 2018.
- Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, January 2018.
- El Mahdi Chayti, Sai Praneeth Karimireddy, Sebastian U. Stich, Nicolas Flammarion, and Martin Jaggi. Linear speedup in personalized collaborative learning, 2021.
- Gary Cheng, Karan Chadha, and John Duchi. Fine-tuning is Fine in Federated Learning. *arXiv:2108.07313 [cs, math, stat]*, August 2021.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, April 2005.
- Yuyang Deng, Mohammad Mahdi Kamani, and Mehrdad Mahdavi. Adaptive Personalized Federated Learning. *arXiv:2003.13461 [cs, stat]*, November 2020. arXiv: 2003.13461.
- Kate Donahue and Jon Kleinberg. Model-sharing games: Analyzing federated learning under voluntary participation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6): 5303–5311, May 2021a.
- Kate Donahue and Jon Kleinberg. Optimality and stability in federated learning: A game-theoretic approach. In *Advances in Neural Information Processing Systems*, 2021b.
- Radu Alexandru Dragomir, Mathieu Even, and Hadrien Hendriks. Fast stochastic bregman gradient methods: Sharp analysis and variance reduction. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 2815–2825. PMLR, 18–24 Jul 2021.

- John Duchi and Ryan Rogers. Lower bounds for locally private estimation via communication complexity. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1161–1191, Phoenix, USA, June 2019. PMLR.
- John C. Duchi and Martin J. Wainwright. Distance-based and continuum fano inequalities with applications to statistical estimation. 2013.
- Mathieu Even, Raphaël Berthier, Francis Bach, Nicolas Flammarion, Hadrien Hendrikx, Pierre Gaillard, Laurent Massoulié, and Adrien Taylor. Continuized accelerations of deterministic and stochastic gradient descents, and of gossip algorithms. In *Advances in Neural Information Processing Systems*, 2021.
- Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach. *Advances in Neural Information Processing Systems*, 33:3557–3568, 2020.
- Felix Grimberg, Mary-Anne Hartley, Sai P. Karimireddy, and Martin Jaggi. Optimal Model Averaging: Towards Personalized Collaborative Learning. *arXiv:2110.12946 [cs, stat]*, October 2021. arXiv: 2110.12946.
- Filip Hanzely, Slavomír Hanzely, Samuel Horváth, and Peter Richtárik. Lower Bounds and Optimal Algorithms for Personalized Federated Learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 2304–2315. Curran Associates, Inc., 2020.
- Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrede Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and Open Problems in Federated Learning. *arXiv:1912.04977 [cs, stat]*, December 2019. arXiv: 1912.04977.
- Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In *Proceedings of the 37th International Conference on Machine Learning*, pages 5132–5143. PMLR, November 2020. ISSN: 2640-3498.
- Mikhail Khodak, Maria-Florina F Balcan, and Ameet S Talwalkar. Adaptive Gradient-Based Meta-Learning Methods. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Jakub Konečný, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated Optimization: Distributed Machine Learning for On-Device Intelligence. *arXiv:1610.02527 [cs]*, October 2016. arXiv: 1610.02527.
- Viraj Kulkarni, Milind Kulkarni, and Aniruddha Pant. Survey of Personalization Techniques for Federated Learning. *arXiv:2003.08673 [cs, stat]*, March 2020. arXiv: 2003.08673.
- Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. Fair resource allocation in federated learning. In *International Conference on Learning Representations*, 2020.
- Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three Approaches for Personalization with Applications to Federated Learning. *arXiv:2002.10619 [cs, stat]*, July 2020. arXiv: 2002.10619.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282. PMLR, 20–22 Apr 2017.
- Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh. Agnostic Federated Learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 4615–4625. PMLR, May 2019. ISSN: 2640-3498.

- Angelia Nedich, Alex Olshevsky, and Michael G. Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, May 2018.
- Kevin Scaman, Francis Bach, Sébastien Bubeck, Yin Tat Lee, and Laurent Massoulié. Optimal convergence rates for convex distributed optimization in networks. *Journal of Machine Learning Research*, 20(159):1–31, 2019.
- Canh T. Dinh, Nguyen Tran, and Josh Nguyen. Personalized Federated Learning with Moreau Envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020.
- Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7852–7862. Curran Associates, Inc., 2020.
- Paul Vanhaesebrouck, Aurélien Bellet, and Marc Tommasi. Decentralized Collaborative Learning of Personalized Models over Networks. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 509–517. PMLR, 20–22 Apr 2017.
- Kangkang Wang, Rajiv Mathews, Chloé Kiddon, Hubert Eichner, Françoise Beaufays, and Daniel Ramage. Federated Evaluation of On-device Personalization. *arXiv:1910.10252 [cs, stat]*, October 2019. arXiv: 1910.10252.
- Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. Salvaging Federated Learning by Local Adaptation. *arXiv:2002.04758 [cs, stat]*, October 2021. arXiv: 2002.04758.
- Yuchen Zhang, John Duchi, Michael I Jordan, and Martin J Wainwright. Information-theoretic lower bounds for distributed statistical estimation with communication constraints. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.

APPENDIX A. THE CASE OF COLLABORATIVE GENERALIZATION ERROR MINIMIZATION (GEM)

In most of the related works and in this paper, the quantities equivalent to our bias assumptions are assumed to be known by the optimizer. Yet, while this assumption makes analyses possible, the knowledge of these quantities is a strong assumption. In this section, we provide settings and toy problems under which these assumptions are natural. We first begin by introducing distribution based distances between agents, in order to control gradient biases in the setting where functions are of the form (5) for a shared loss ℓ and local distributions $(\mathcal{D}_i)_i$, for collaborative GEM. We then present two settings for this problem.

A.1. Distribution-based distances. In order to quantify the bias/variance tradeoff that appears when using stochastic gradients computed by other agents, we define the following notion of distance.

Definition 1. For \mathcal{H} a set of functions from Ξ to \mathbb{R}^d and $\mathcal{D}, \mathcal{D}'$ two probability distributions on Ξ , we define:

$$d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}') = \sup_{h \in \mathcal{H}} \|\mathbb{E}[h(\xi)] - \mathbb{E}[h(\xi')]\|,$$

where $\xi \sim \mathcal{D}$ and $\xi' \sim \mathcal{D}'$. $d_{\mathcal{H}}$ is a pseudo-distance on the set of probability measures on Ξ .

Some instances of these distances encompass:

- (1) For $\mathcal{H} = \mathcal{H}_{\text{Affine}}$ the set of 1-Lipschitz affine functions, we have $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}') = \|\mathbb{E}[\xi] - \mathbb{E}[\xi']\|$.
- (2) For $\mathcal{H} = \mathcal{H}_{\text{Lipschitz}}$ the set of 1-Lipschitz functions on Ξ , we have $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}') = \mathcal{W}_1(\mathcal{D}, \mathcal{D}')$ the 1-Wasserstein distance (or EMD, earth mover distance).
- (3) For $\mathcal{H} = \mathcal{H}_{\text{Bounded}}$ the set of functions whose values lie in a space of diameter 1, we have $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}') = d_{\text{TV}}(\mathcal{D}, \mathcal{D}')$ the total variation distance.
- (4) For $\mathcal{H} = \mathcal{H}_{\text{Loc.Bounded}}$ defined as $\{h : \Xi \rightarrow \mathbb{R}^d \text{ s.t. } \forall \xi \in \Xi, \text{Diam}(h(\mathcal{B}(\xi, 1) \cap \Xi)) \leq 1\}$, we also obtain the total variation distance.

The definition of $d_{\mathcal{H}}$ is motivated by the fact that, for fixed $x \in \mathcal{X}$ and $1 \leq i, j \leq N$, we have:

$$\begin{aligned} \|\nabla f_i(x) - \nabla f_j(x)\| &= \|\mathbb{E}[\nabla_x \ell(x, \xi_i) - \nabla_x \ell(x, \xi_j)]\| \\ &\leq d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j). \end{aligned}$$

if $(\xi \in \Xi \mapsto \nabla_x \ell(x, \xi)) \in \mathcal{H}$ where $(\xi_i, \xi_j) \sim \mathcal{D}_i \times \mathcal{D}_j$.

We thus add a fourth bias assumptions to our list in Section 2, called distribution-based bias assumption: for some function set \mathcal{H} , for all $1 \leq i, j \leq N$, for all $x \in \mathbb{R}^d$, $d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq \hat{b}_{ij}$ for some non negative weights \hat{b} , and for all $x \in \mathbb{R}^d$, $(\xi \in \Xi \mapsto \nabla_x \ell(x, \xi)) \in \mathcal{H}$. This assumptions is verified in the following settings for mean estimation and regression.

- (1) For $\ell(x, \xi) = \frac{1}{2}(x - \xi)^2$ (mean estimation), distribution-based bias assumption holds for $\mathcal{H}_{\text{Affine}}$, and the distance to consider is thus $\|\mathbb{E}[\xi_i] - \mathbb{E}[\xi_j]\|$.
- (2) For quadratic loss (linear regression) $\ell(x, \xi) = \frac{1}{2}(a^\top x - b)^2$ where $\xi = (aa^\top, ba^\top)$, we have $\mathcal{H} = D\mathcal{H}_{\text{Affine}}$, where D is the diameter of the space in which are iterates lie, and the distance to consider is thus a scaled 1-Wasserstein distance between distributions
- (3) For logistic regression $\ell(x, \xi) = \log(1 + e^{-ba^\top x})$ where $\xi = ba^\top$, we have $\mathcal{H} = (1 + \frac{1}{e})\mathcal{H}_{\text{Lipschitz}}$ and the distance to consider is thus the 1-Wasserstein distance between distributions (scaled with a constant factor).
- (4) For the hinge loss $\ell(x, \xi) = \max(0, 1 - ba^\top x)$, for $\xi = ba^\top$ we have $\mathcal{H} = D\mathcal{H}_{\text{Bounded}}$ where D is the diameter of the space in which are iterates lie, and the distance to consider between distributions is thus a scaled total variation distance.

Proposition 1. If for all $1 \leq i, j \leq N$, $d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq \hat{b}_{ij}$ (i.e. distribution-based bias assumption holds), then:

- (1) Bias assumption (B.1) holds for $b_{ij} = D^* \hat{b}_{ij}$;
- (2) Bias assumption (B.2) holds for $\tilde{b}_{ij} = \hat{b}_{ij}^2$.

A.2. Lower-bounds (Theorem 1) in term of distribution based distances. Theorem 1 being formulated as IT-lower bounds for collaborative GEM, the distribution-based distances we introduced perfectly adapt to it. In fact, in the proof of the lower bounds, the distributions we build verify $d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq b_{ij}$ if $\mathcal{H} \subset \mathcal{H}_{\text{Loc.Bounded}}$, so that the following directly holds.

Theorem 5. Let $\varepsilon \in (0, 1/16)$. Assume that the function set \mathcal{H} satisfies:

$$\mathcal{H}_{\text{Affine}} \subset \mathcal{H} \subset \mathcal{H}_{\text{Loc.Bounded}},$$

and assume that either the synchronous or asynchronous oracle is used. Assume that (b_{ij}) verifies the triangle inequality $b_{ij} \leq b_{ik} + b_{kj}$ for all $1 \leq i, j, k \leq N$. Let $\mathbb{D}(r, b, B)$ be the set $\mathbb{D}(r, b, B)$, but where the bias assumptions are replaced by the distribution-based bias assumption ($d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq b_{ij}$), and similarly for $\tilde{\mathbb{D}}_{\mu=1/r^2}^{L=1/r^2}(r, b, \sigma)$. We have, for some constant $C > 0$ independent of the problem and any $i \in \{1, \dots, N\}$:

$$\inf_{\mathcal{M} \in \mathbb{M}_{((\mathcal{D}_j)_j, \ell) \in \mathbb{D}(r, b, B)}} \sup \mathcal{T}_i^\varepsilon \left(\mathcal{M}, ((\mathcal{D}_j)_j, \ell) \right) \geq \frac{CB^2 r^2 N}{\varepsilon^2 \mathcal{N}_i^\varepsilon(\frac{b}{4})},$$

where $\mathcal{N}_i^\varepsilon(b) = \sum_j \mathbb{1}_{\{b_{ij} \leq \varepsilon\}}$ is the number of agents j verifying $b_{ij} \leq \varepsilon$. Under strong-convexity and smoothness assumptions for $\mu = L = 1/r^2$, this bound becomes, for some constant $C' > 0$:

$$\inf_{\mathcal{M} \in \mathbb{M}_{((\mathcal{D}_j)_j, \ell) \in \tilde{\mathbb{D}}_{\mu=1/r^2}^{L=1/r^2}(r, b, \sigma)}} \sup \mathcal{T}_i^\varepsilon \left(\mathcal{M}, ((\mathcal{D}_j)_j, \ell) \right) \geq \frac{C' r^2 \sigma^2 N}{\varepsilon \mathcal{N}_i^\varepsilon(\frac{b}{4})}.$$

A.3. Weak supervision. We explicit in this section another paradigm in which the knowledge of some distribution-based bias assumptions is realistic, in a weakly supervised learning setting. More precisely, assume that the desired task is a classification or regression one: random variables ξ are of the form $\xi = (\zeta, \beta)$ for some $\zeta \in Z \subset \mathbb{R}^{D_1}$ and some label $\beta \in \mathbb{R}^{D_2}$. For every agent $1 \leq i \leq N$, let \mathcal{D}_i^ζ be the ζ -marginal of \mathcal{D}_i . In our weakly supervised learning setting, we assume that there exists some (eventually random) function \mathcal{L} that maps unlabelled data ζ to their labelled ones *i.e.* for some $\mathcal{L} : Z \rightarrow \Xi$, the random variable $\mathcal{L}(\zeta_i)$ is of same law as $\xi_i \sim \mathcal{D}_i$ for $\zeta_i \sim \mathcal{D}_i^\zeta$. While many samples may be accessible and drawn from \mathcal{D}_i^ζ for every agent $1 \leq i \leq N$, labelling data *i.e.* applying function \mathcal{L} in order to recover samples of law \mathcal{D}_i is assumed to be costly. For instance, applying \mathcal{L} in medical data analysis could require the help from an expert (a radiologist, ...).

Agents would benefit from computing distances between distributions \mathcal{D}_i using marginals \mathcal{D}_i^ζ : under mild assumptions on the function \mathcal{L} (Lipschitzness with high probability, mainly), distances $d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j)$ are upperbounded by distances between \mathcal{D}_i^ζ and \mathcal{D}_j^ζ . Thus, in this setting, agents would benefit from collaborating: computing distribution-based distances would not require expert advices, while reducing the total number of expert calls required by reducing the number of samples needed to be drawn from each \mathcal{D}_i through collaboration.

A.4. Infinitely-Many-Agents Limit and Geometrical Prior Knowledge. We introduce the *infinitely-many-agents* limit, by taking $N \rightarrow \infty$ with added structure on the problem below. This toy-problem illustrates a geometric structure under which the knowledge of distribution-based bias assumptions is realistic. Let \mathcal{P} be a set of probability laws and $\Theta \subset \mathbb{R}^p$ be a set parameterizing \mathcal{P} in the sense that \mathcal{P} can be written as $\{\mathcal{D}_\theta, \theta \in \Theta\}$ where $\mathcal{D}_\theta, \theta \in \Theta$ are probability laws on Ξ . For all $1 \leq i \leq N$, we assume that there exists $\theta_1, \dots, \theta_N \in \Theta$ such that $\mathcal{D}_i = \mathcal{D}_{\theta_i}$, and that the agents are drawn from a probability distribution with density ν on Θ : $(\theta_i)_i$ is a sequence of *i.i.d.* vectors sampled from Θ with density ν .

This aims at modelling a *geometric* structure on the agents distributions: agents i and j such that θ_i and θ_j are nearby in the set Θ share similar distributions. We quantify this in the following way: there exists a continuous function $b : \Theta^2 \rightarrow \mathbb{R}^+$ such that for any $(\theta, \theta') \in \Theta^2$, $d_{\mathcal{H}}(\theta, \theta') \leq b(\theta, \theta')$. In sensor networks, Θ is some domain of \mathbb{R}^2 or \mathbb{R}^3 , each agent is a sensor located at a physical position in Θ , whose goal is to predict based on noisy local observations (*e.g.* temperature, pollution, ...) that are space-continuous. More generally, if agents model a population (agent i being an individual in the population) in some geographical location $\Theta \subset \mathcal{S}_2$ (sphere in \mathbb{R}^3) (ranging from cities, countries, continents, or the whole wide world), and if we are interested in speech-recognition models, a fair assumption is to have similar local distributions for nearby agents. The distribution ν is then the geographical density of agents.

The *infinitely-many-agents* limit is obtained by taking $N \rightarrow \infty$ with our geometric structure on the agents: we obtain a continuum of agents parameterized by $\theta \in \Theta$, with density ν on Θ . For every $\theta \in \Theta$, we write:

$$f_\theta(x) = \mathbb{E}_{\xi \sim \mathcal{D}_\theta}[\ell(x, \xi)], \quad x \in \mathcal{X},$$

and we denote by x_θ^* a minimizer over \mathcal{X} of f_θ . The algorithm and model we consider are defined as follows, inspired by Oracle 2.

Infinitely-many-agents oracle and algorithm

At each time-step $t = 1, 2, \dots$:

- 1:** An agent parameterized by $\theta_t \sim \nu$ is ‘awakened’, draws $\xi_{\theta_t}^t$ is drawn from data distribution \mathcal{D}_{θ_t} and computes a stochastic gradient $g_t = \nabla_x \ell(y_t, \xi_{\theta_t}^t)$.
- 2:** Agents $\theta \in \Theta$ such that $w(\theta, \theta_t) \leq s_\varepsilon$ for some fixed s_ε receive g_t .
- 3:** Upon reception of g_t , agents may update their local estimates.

At finite time-horizon $T > 0$, only a finite number of agents may have ‘awakened’ and drawn (at most) one sample from their local distribution; yet, we aim at obtaining upper-bounds on the local generalization errors for all agents. Applying our AFO and AFA algorithms, under additional assumptions, yields the following bounds.

Proposition 2. *Assume that ν is continuous on Θ , and that for all $\theta, \theta' \in \Theta$, we have $b(\theta, \theta') \leq \|\theta - \theta'\|^q$. Assume that assumption N.1 holds and that iterates lie in a space of diameter D .*

- (1) *All-for-one: let $\varepsilon > 0$ and $\theta \in \text{int } A$. There exists a procedure such that $T_\varepsilon(\theta)$ the time needed to reach a precision ε for agent θ is upper-bounded by:*

$$T_\varepsilon(\theta) \leq \mathcal{O}_{\varepsilon \rightarrow 0} \left(B^2 D^2 \nu(\theta) \varepsilon^{-2-p/q} \right).$$

- (2) *All-for-all: for any $\varepsilon > 0$, there exists a procedure making T_ε oracle calls and answering $(x_\theta)_\theta \in \mathcal{X}^\Theta$ such that:*

$$\int_{\Theta} (f_\theta(x_\theta) - f_\theta(x_\theta^*)) \nu(d\theta) \leq \varepsilon,$$

with

$$T_\varepsilon \leq \mathcal{O}_{\varepsilon \rightarrow 0} \left(B^2 D^2 \varepsilon^{-2-p/q} \right).$$

Proof. For finite number of agents N , we have $\theta_1, \dots, \theta_N$ drawn *i.i.d.* from ν . Using weighted gradient averagings and Theorem 2.1, we have for agent i , and denoting $\theta = \theta_i$:

$$T_\varepsilon(\theta) \leq 2B^2 D^2 \varepsilon^{-2} \frac{\mathcal{N}_i^\varepsilon(2b)}{N}.$$

Here, $\mathcal{N}_i^\varepsilon(2b) \leq |\{1 \leq j \leq N : \|\theta - \theta_j\| \leq (\varepsilon/2)^{1/q}\}| = \sum_{j=1}^N \mathbf{1}_{\{\|\theta - \theta_j\| \leq (\varepsilon/2)^{1/q}\}}$. Using the strong law of large number, we have almost surely, as $N \rightarrow \infty$:

$$\frac{\mathcal{N}_i^\varepsilon(2b)}{N} \leq \int_{\mathcal{B}(\theta, (\varepsilon/2)^{1/q})} \nu(z) dz + o(1).$$

Thus, for $N \rightarrow \infty$:

$$T_\varepsilon(\theta) \leq 2B^2 D^2 \int_{\mathcal{B}(\theta, (\varepsilon/2)^{1/q})} \nu(z) dz.$$

Then, making $\varepsilon \rightarrow 0$, we have $\int_{\mathcal{B}(\theta, (\varepsilon/2)^{1/q})} \nu(z) dz \sim_{\varepsilon \rightarrow 0} \nu(\theta) \text{Vol}(\mathcal{B}(\theta, (\varepsilon/2)^{1/q})) = \mathcal{O}(\nu(\theta) \varepsilon^{-p/q})$ (ν is continuous). We thus have the result for θ such that $\nu(\theta) > 0$ in the *all-for-one* setting.

In the *all-for-all* setting, we use Theorem 4.1 in for a finite number of agents, and we similarly use the strong law of large numbers as $N \rightarrow \infty$. \square

APPENDIX B. PROOF OF LOWER-BOUNDS (THEOREM 1 AND COROLLARY 1)

B.1. General framework to prove lower bounds. The idea is that, when optimizing a function $f(x) = \mathbb{E}[\ell(x, \xi)]$ and finding a good approximation of a minimizer x^* , we learn some information on the distribution \mathcal{D} over which samples are drawn. In order to prove lower bounds, we construct a loss function ℓ , and distributions $\mathcal{D}_1^\alpha, \dots, \mathcal{D}_N^\alpha$, where α is a random parameter. We argue that minimizing (in the *all-for-all* or *all-for-one* settings) the objective function up to a certain precision

gives a good estimator (quantified) of the random seeds α . Then, using Fano inequality, we bound the efficiency of such an estimator in terms of number of oracle calls, obtaining a lower bound on the sample complexity. This approach is inspired by Agarwal et al. [2012], who prove IT-lower bounds for stochastic gradient descent. We adapt their proof technique to the personalized and multi-agent setting.

Constructing difficult loss functions. For any two functions $f, g : \mathbb{R}^d \rightarrow \mathbb{R}$, we define the discrepancy measure $\rho(f, g)$ as:

$$\rho(f, g) = \inf_{x \in \mathbb{R}^d} \left\{ f(x) + g(x) - f(x_f^*) - g(x_f^*) \right\},$$

which is a pseudo metrics. Now, for a finite set \mathcal{V} of parameters, let $\mathcal{G}(\delta) = \{g_\alpha^\delta, \alpha \in \mathcal{V}\}$ be a set of functions indexed by \mathcal{V} , that depend on δ (fixed in the set). The dependency of each $g_\alpha \in \mathcal{G}(\delta)$ is left implicit in the following subsections. We define:

$$\psi(\delta) = \inf_{f, g \in \mathcal{G}(\delta)} \rho(f, g).$$

Minimizing is Bernoulli parameters identification. The two following lemmas justify that optimizing a function $g_\alpha \in \mathcal{G}(\delta)$ to a precision of order $\psi(\delta)$ is more difficult than estimating the parameter α .

Lemma 1 (Agarwal et al. [2012]). *For any $x \in \mathbb{R}^d$, there can be at most one function g_α in $\mathcal{G}(\delta)$ such that:*

$$g_\alpha(x) - \inf_{\mathbb{R}^d} g_\alpha < \frac{\psi(\delta)}{3}.$$

Lemma 2 (Agarwal et al. [2012]). *Assume that for some fixed but unknown $\alpha \in \mathcal{V}$ there exists a method \mathcal{M}_T based on the data $\phi = \{X_1, \dots, X_T\}$ that returns x^T (function of ϕ) satisfying an error of:*

$$\mathbb{E} \left[g_\alpha(x^T) - \min_{x \in \mathbb{R}^d} g_\alpha(x) \right] < \frac{\psi(\delta)}{9},$$

where the mean is taken over the randomness of both the oracle Φ , the method \mathcal{M}_T and $\alpha \in \mathcal{V}$ if random. Then, there exists a hypothesis test $\hat{\alpha} : \phi \rightarrow \mathcal{V}$ such that:

$$\max_{\alpha \in \mathcal{V}} \mathbb{P}_\phi(\hat{\alpha} \neq \alpha) \leq \frac{1}{3}.$$

Suppose now that the parameter α in the previous Lemma is chosen uniformly at random in \mathcal{V} . Let $\hat{\alpha} : \phi \rightarrow \mathcal{V}$ be a hypothesis test estimating α . By Fano inequality [Cover and Thomas, 2005], we have:

$$\mathbb{P}(\hat{\alpha} \neq \alpha) \geq 1 - \frac{I(\phi, \alpha) + \ln(2)}{|\mathcal{V}|}, \quad (9)$$

where $I(\phi, \alpha)$ is the mutual information between ϕ and α , that we need to upper-bound. Combining Fano inequality with Lemmas 1 and 2, fixing a target error $\varepsilon = \psi(\delta)$, we obtain a lower bound on the sample complexity T_ε :

$$\frac{1}{3} \geq \mathbb{P}_\phi(\hat{\alpha} \neq \alpha) \geq 1 - \frac{I(\phi_{T_\varepsilon}, \alpha) + \ln(2)}{|\mathcal{V}|},$$

where ϕ_{T_ε} is the information contained in T_ε oracle calls. If we have an equality of the form $I(\phi_{T_\varepsilon}, \alpha) = T_\varepsilon I(\phi_1, \alpha)$, this gives:

$$T_\varepsilon \geq \frac{\frac{2}{3}|\mathcal{V}| - \ln(2)}{I(\phi_1, \alpha)}. \quad (10)$$

Playing with the different parameters $\delta, \alpha, \mathcal{V}$ gives lower bounds in our *all-for-one* settings. We refer the interested reader to Chapter 2 in Cover and Thomas [2005] for Fano inequality and mutual information.

B.2. Applying this in the *all-for-one* setting, first part of Theorem 1. We first prove the lower bounds in the case of the asynchronous oracle (one data item sampled from personal distribution at each iteration).

We define the loss $\ell : \mathbb{R}^d \times \Xi$ where $\Xi \subset \mathbb{R}^d$ and $\delta > 0$:

$$\forall (x, \xi) \in \mathcal{X} \times \Xi, \quad \ell(x, \xi) = \frac{1}{d} \sum_{i=1}^d \left[\xi f_i^+(x) + (1 - \xi) f_i^-(x) - \delta x_i \right],$$

where

$$f_i^+(x) = |x_i + 1/2| \quad f_i^-(x) = |x_i - 1/2|.$$

Now, for any $\alpha \in \{-1, 1\}^d$, $\delta \in [0, 1/2]$ being fixed, we define:

$$g_\alpha(x) = \frac{1}{d} \sum_{i=1}^d \left[\left(\frac{1}{2} + \alpha_i \delta \right) f_i^+(x) + \left(\frac{1}{2} - \alpha_i \delta \right) f_i^-(x) - \delta x_i \right]. \quad (11)$$

For simplicity, assume that $r^2 = d$ and $B^2 = 1$. Let $\delta > 0$ a free parameter. Let $\mathcal{V} = \{\alpha^1, \dots, \alpha^K\} \subset \{-1, 1\}^d$ be a subset of the hypercube such that for all $k \neq l$,

$$\frac{1}{2} \sum_{i=1}^d |\alpha_i^k - \alpha_i^l| \geq \frac{d}{4},$$

i.e. \mathcal{V} is a $d/4$ -packing of the hypercube. We assume that $-\mathbf{1}$ (d -dimensional vector with -1 at all its entries) is not in \mathcal{V} and that for all $\alpha \in \mathcal{V}$, $\frac{1}{2} \sum_{i=1}^d |\alpha_i + 1| \geq \frac{d}{4}$. We know that we can set $|\mathcal{V}| \geq (2/\sqrt{e})^{d/2}$. Without loss of generality, we prove a lower bound in the case where the agent that desires to minimize its local function is indexed by 1. For any $i = 1, \dots, N$, let \mathcal{D}_i be the probability distribution on $\{0, 1\}^d$ of the following random variable:

$$\text{Ber}\left(\frac{1}{2} + \delta_i \alpha_k\right) \epsilon_k \quad \text{where} \quad \delta_i = (\delta - b_{i1})^+,$$

where $s^+ = \max(0, s)$ for $s \in \mathbb{R}$, k is taken uniformly at random in $\{1, \dots, d\}$, (ϵ_k) is the canonical basis of \mathbb{R}^d , and $\text{Ber}(p)$ is a Bernoulli random variable, independent of k .

The function $f_1(x) = \mathbb{E}[\ell(x, \xi_1)]$ for $\xi_1 \sim \mathcal{D}_1$ verifies $f_1 = g_\alpha$ for g_α defined in (11). For our fixed $\delta > 0$ and any $\alpha, \beta \in \mathcal{V}$, g_α is minimized at $x^\alpha = -\alpha/2$ and we have:

$$g_\alpha(x^\alpha) = \frac{1}{d} \sum_{k=1}^d \frac{1 - \delta \alpha_k}{2}.$$

The discrepancies thus write as:

$$\rho(g_\alpha, g_\beta) = \frac{\delta}{d} \sum_{k=1}^d |\alpha_k - \beta_k|,$$

so that $\psi(\delta) \geq \delta/4$. Each g_α is minimized for $x^\alpha = -\alpha/2 \in \mathcal{B}_\infty(0, 1/2) \subset \mathcal{B}(0, \sqrt{d}/2) \subset \mathcal{B}(x^0, r)$ for $x^0 = 0$ (where \mathcal{B}_∞ denotes a ball for the infinity norm): we are in the case $r = \sqrt{d}$.

We now bound the quantities $f_i(x_j^*) - f_i(x_i^*)$. For any $i = 1, \dots, N$, we have $x_i^* = x^\alpha = -\alpha/2$ if $\delta_i \geq \delta/2$, and $x_i^* = \mathbf{1}/2$ otherwise. For $1 \leq i, j \leq N$, if both $\delta_i \geq \delta/2$ and $\delta_j \geq \delta/2$, then $x_i^* = x_j^* = x^\alpha$, and $f_i(x_j^*) - f_i(x_i^*) = 0$. Similarly, if both $\delta_i < \delta/2$ and $\delta_j < \delta/2$, $x_i^* = x_j^* = \mathbf{1}/2$ and $f_i(x_j^*) - f_i(x_i^*) = 0$. If $\delta_i \geq \delta/2$ and $\delta_j < \delta/2$, then

$$\begin{aligned} f_i(x_j^*) - f_i(x_i^*) &= f_i(\mathbf{1}/2) - f_i(x^\alpha) \\ &= \frac{1}{d} \sum_{k=1}^d \frac{|\alpha_k - 1|}{2} (2\delta_i - \delta) \\ &\leq 2\delta_i - \delta \\ &\leq 2(\delta_i - \delta_j). \end{aligned}$$

Since $\delta_i - \delta_j \leq b_{ij}/2$ using 1-Lipschitzness of the positive value and the triangle inequality verified by the vector b , we have $f_i(x_j^*) - f_i(x_i^*) \leq b_{ij}$. If $\delta_i < \delta/2$ and $\delta_j > \delta/2$, we similarly have $f_i(x_j^*) - f_i(x_i^*) \leq \frac{1}{d} \sum_{k=1}^d \frac{|\alpha_k - 1|}{2} (\delta - 2\delta_i) \leq b_{ij}$.

Since the functions f_i are differentiable at all points where $x_k \neq 1/2, -1/2$, bounding $\|\nabla f_i(x) - \nabla f_j(x)\|^2$ at the points where f_i and f_j are differentiable is enough to have (B.2). For $x_k < -1/2$, we have $\nabla_k f_i(x) = \nabla_k f_j(x) = -1 - \delta$ and $\nabla_k f_i(x) = \nabla_k f_j(x) = 1 - \delta$ for $x_k > 1/2$. For $-1/2 < x_k < 1/2$, we have $\nabla_k f_i(x) - \nabla_k f_j(x) = 2(\delta_i - \delta_j)\alpha_k$, so that $\|\nabla f_i(x) - \nabla f_j(x)\|^2 \leq 4(\delta_i - \delta_j)^2 \leq b_{ij}^2/d$ at all points where f_i and f_j are differentiable. The gradients scale with r , hence the dependency in r (we recall that $r^2 = d$ in our case).

The mutual information $I(\phi_T, \alpha)$ writes as:

$$\begin{aligned} I(\phi_T, \alpha) &= TI(\phi_1, \alpha) \\ &= \frac{T}{dN} \sum_{k=1}^d \sum_{i=1}^N I\left(\text{Ber}\left(\frac{1}{2} + \mathbf{1}_{b_{1i} \leq \delta} \alpha_k \delta_i\right), \alpha_k\right) \\ &\leq \frac{C_1 T}{N} \sum_{i=1}^N \mathbf{1}_{b_{1i} < \delta} \delta_i^2 \\ &= \frac{C_1 T \mathcal{N}_1^\delta(b)}{N} \delta^2, \end{aligned}$$

for some numerical constant $C > 0$ independent of the problem. Using Equation (10) then leads to, for $\varepsilon = \psi(\delta)$:

$$T_\varepsilon \geq C \frac{dN}{\varepsilon^2 \mathcal{N}_1^\varepsilon(4b)}, \quad (12)$$

where $C > 0$ is a numerical constant independent of the problem.

Equation (12) is obtained for the choice of loss function and distributions defined above, and the functions defined are in $\mathbb{D}(r, b, B)$ for $r = \sqrt{d}$, the fixed weights b , and $B^2 = 1$. Assumption $\mathcal{B}_\infty(0, 1/2) \subset \mathcal{B}(0, \sqrt{d}/2) \subset \mathcal{X}$ is in Theorem 1 the assumption that $\|x^0 - x_i^*\| \leq r$. In order to obtain a dependency in $r > 0$, one simply has to consider the loss $(x, \xi) \mapsto \ell(\frac{\sqrt{d}}{2r}x, \xi)$. In order to obtain the dependency in B^2 , we need to modify the distribution, and take $\mathcal{D}'_i = \text{Ber}(1/B^2)B^2\mathcal{D}_i$. In this case, we have a noise amplitude of order B^2 instead of order 1, and a factor $1/B^2$ appears in the mutual information.

The proof naturally extends to Oracle 1: the mutual information is just N times bigger in that case.

In terms of function distances $d_{\mathcal{H}}$ under the assumptions on \mathcal{H} of Theorem 5, we have, since the distributions are Bernoulli:

$$d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq |\delta_i - \delta_j| \leq b_{ij}.$$

We thus proved the first part of Theorem 1.

B.3. Proof of the strongly convex and smooth lower bound: second part of Theorem 1.

Let

$$\ell(x, \xi) = \frac{1}{2} \|x - \xi\|^2$$

for $x, \xi \in \mathbb{R}^d$ and, for fixed $\delta > 0$ and any $\alpha \in \{-1, 1\}^d$:

$$g_\alpha(x) = \frac{1}{2d} \sum_{k=1}^d \left(x_k^2 + 1 - 2\left(\frac{1}{2} + \alpha_k \delta\right) x_k \right), \quad x \in \mathcal{X}.$$

We keep the same notations as last subsection ($\psi(\delta), \rho$). We have:

$$\rho(g_\alpha, g_\beta) = \frac{\delta^2}{d} \sum_{k=1}^d |\alpha_k - \beta_k|.$$

Similarly to last subsection, this leads to $\psi(\delta) \geq \delta^2/4$ since \mathcal{V} is a $d/4$ -packing of the hypercube.

We keep the same distributions $\mathcal{D}_1, \dots, \mathcal{D}_N$ as last subsection, replacing δ_i by $\delta_i = (\delta - \sqrt{b_{1i}})^+$. Agent 1 is still the one that wants to minimize its local generalization error. The mutual information is thus:

$$I(\phi_T, \alpha) \leq \frac{C_1 T \mathcal{N}_1^\delta(\sqrt{b})}{N} \delta^2.$$

Setting the target precision as $\varepsilon = \delta^2/4$, we obtain:

$$T_\varepsilon \geq C' \frac{dN}{\varepsilon \mathcal{N}_1^\varepsilon(4b)}.$$

The loss function and distributions built verify our regularity assumptions for $\mu = 1/d$, $L = 1/d$, noise $\sigma^2 \leq 1$.

We verify that for all $1 \leq j, k \leq N$, we have $f_j(x_k^*) - f_j(x_j^*) \leq b_{kj}$. We first notice that $x_j^* = \frac{1}{d}(\frac{1}{2} + \delta_j \alpha_l)_{1 \leq l \leq d}$, so that:

$$\begin{aligned} f_j(x_k^*) - f_j(x_j^*) &= \frac{1}{d} \|x_j^* - x_k^*\|^2 \\ &= (\delta_i - \delta_k)^2 \\ &\leq (\sqrt{b_{1j}} - \sqrt{b_{1k}})^2 \\ &\leq |b_{1j} - b_{1k}| \\ &\leq b_{jk}, \end{aligned}$$

since the weights b verify the triangle inequality. Under the assumptions of Theorem 5 on \mathcal{H} , we have, in terms of distribution-based distances:

$$d_{\mathcal{H}}(\mathcal{D}_i, \mathcal{D}_j) \leq |\delta_i - \delta_j| \leq \sqrt{b_{ij}}.$$

The minimum of each g_α is attained at $x^\alpha = \frac{1}{2} + \delta\alpha$, we thus need to assume as in last subsection that r is of order \sqrt{d} , and a rescaling leads to the dependency in r . The dependency in σ^2 is obtained as in last subsection.

APPENDIX C. All-for-all LOWER BOUND

Theorem 1 immediately yields a Corollary providing lower bounds for the all-to-all problem:

Corollary 1. *Let $\varepsilon > 0$, b verifying the triangle inequality, $B > 0, r > 0$. Assume that there exists a method \mathcal{M} and some $T > 0$ such that, for all $((\mathcal{D}_j)_j, \ell) \in \mathbb{D}(r, b, B)$, \mathcal{M} returns (x_1, \dots, x_N) verifying:*

$$\frac{1}{N} \sum_{i=1}^N f_i(x_i) - f_i(x_i^*) \leq \varepsilon.$$

where T is the number of stochastic gradients sampled from all agents. Then we have:

$$T \geq \frac{Cr^2B^2}{N^2\varepsilon^2} \sum_{i=1}^N \frac{1}{\mathcal{N}_i^\varepsilon(\frac{b}{4N})}.$$

Replacing $\mathbb{D}(r, b, B)$ by $\mathbb{D}_{\mu=1/r^2}^{L=1/r^2}(r, b, \sigma)$, we have:

$$T \geq \frac{C'r^2\sigma^2}{N\varepsilon} \sum_{i=1}^N \frac{1}{\mathcal{N}_i^\varepsilon(\frac{b}{4N})}.$$

APPENDIX D. All-for-one UPPER-BOUNDS

We provide the following relaxation of bias assumption B.2, a generalization of classical function dissimilarities assumptions Karimireddy et al. [2020], Chayti et al. [2021] to our setting.

B.3 For all $1 \leq i \leq N$, for any stochastic vector $\lambda \in \mathbb{R}^N$ (non-negative entries that sum to 1), and for all $x \in \mathbb{R}^d$, where $g_\lambda(x) = \sum_{j=1}^N \lambda_j \nabla f_j(x)$:

$$\|\nabla f_i(x) - g_\lambda(x)\|^2 \leq m \|g_\lambda(x)\|^2 + \sum_{j=1}^N \lambda_j \tilde{b}_{ij}. \quad (13)$$

In a similar setting (N agents with heterogeneous functions), Karimireddy et al. [2020], Deng et al. [2020], Chayti et al. [2021] assume that for a fixed agent (or for any agent) i , $\|\nabla f_i(x) - \nabla \bar{f}(x)\|^2 \leq m \|\nabla \bar{f}(x)\|^2 + \zeta^2$ (or similar assumptions) where \bar{f} is the mean of all functions. (B.2) is the simplest relaxation of this to the more general setting where all agents want to minimize their local function, while (B.3) is a relaxation that keeps the less-restrictive first term for some $m \geq 0$.

We prove the following intermediate result.

Theorem 6 (WGA). *Let $(x^k)_{k \geq 0}$ be generated with (6), for some fixed stochastic vector $\lambda \in \mathbb{R}^N$, stepsize $\eta > 0$. Let $1 \leq i \leq N$ fixed. Let $b \in \mathbb{R}^{+N \times N}$. Denote $F_i^k = f_i(x^k) - f_i(x_i^*)$.*

6.a Use the asynchronous oracle. Assume that for all $k \geq 0$ and i , $\|x^k - x_i^\| \leq D$ for some $D > 0$. If f_i is convex, noise assumption (N.1) holds for some $B > 0$ and bias assumption (B.2) holds for $\tilde{b} = (b/D)^2$, for any $K > 0$ and for $\eta = \sqrt{\frac{D^2}{2KNB^2 \sum_{j=1}^N \lambda_j^2}}$:*

$$\mathbb{E} \left[\frac{1}{K} \sum_{0 \leq k < K} F_i^k \right] \leq \sqrt{\frac{2D^2B^2}{K} \sum_{j=1}^N \lambda_j^2} + \sum_{1 \leq j \leq N} \lambda_j b_{ij}.$$

6.b Use the synchronous oracle. Assume that all f_j are μ -strongly convex and L -smooth (denote $\kappa = L/\mu$), satisfy noise assumption N.2 for some $\sigma > 0$, that one of bias assumptions (B.2) or (B.3) holds for $\tilde{b} = \mu b$ (and any $m \geq 0$). For any $K > 0$ and $\eta = \min(1/(2L), \frac{1}{\mu K} \ln(\frac{F_\lambda^0 \mu^2 K}{L\sigma^2} \sum_j \lambda_j^2))$ where $F_\lambda^0 = \sum_j \lambda_j (f_j(x^0) - f_j(x_j^))$, we have:*

$$\mathbb{E} [F_i^K] \leq 2(m+1)\kappa \left(F_\lambda^0 e^{-\frac{K}{2\kappa}} + \tilde{\mathcal{O}} \left(\frac{L\sigma^2}{\mu^2 K} \sum_{1 \leq j \leq N} \lambda_j^2 \right) \right) + \sum_{1 \leq j \leq N} \lambda_j b_{ij}.$$

This leads to a more general sample complexity in Setting 2, where we can use (B.3).

Theorem 7. *Let $\varepsilon > 0$, and set $\lambda_j = \frac{\mathbb{1}_{\{b_{ij} \leq \varepsilon/2\}}}{N_i^\varepsilon(2b)}$, N_i^ε being defined in Theorem 1. Generate $(x^k)_k$ using (6). Under Setting 1. Assume that for all $k \geq 0$ and i , $\|x^k - x_i^*\| \leq D$ for some $D > 0$ and bias assumption (B.2) holds for $\tilde{b} = (b/D)^2$. For $K = T_\varepsilon$ and for $\eta = \sqrt{\frac{D^2}{2KNB^2 \sum_{j=1}^N \lambda_j^2}}$,*

we have $\mathbb{E} \left[\frac{1}{K} \sum_{0 \leq k < K} f_i(x^k) - f_i(x_i^) \right] \leq \varepsilon$ using a total number T_ε of stochastic gradients from all agents of:*

$$T_\varepsilon \leq \frac{4D^2B^2}{\varepsilon^2} \frac{N}{N_i^\varepsilon(2b)}.$$

7.2 Under Setting 2. Denote $\kappa = L/\mu$, assume that one of bias assumptions (B.2) or (B.3) holds for $\tilde{b} = \mu b$ (and any $m \geq 0$). For $K = T_\varepsilon/N$ and $\eta = \min(1/(2L), \frac{1}{\mu K} \ln(\frac{F_\lambda^0 \mu^2 K}{L\sigma^2} \sum_j \lambda_j^2))$ where $F_\lambda^0 = \sum_j \lambda_j (f_j(x^0) - f_j(x_j^))$, we have $\mathbb{E} [f_i(x^K) - f_i(x_i^*)] \leq \varepsilon$ using a total number T_ε of stochastic gradients from all agents of:*

$$T_\varepsilon \leq \tilde{\mathcal{O}} \left(\frac{(m+1)\kappa^2\sigma^2}{\mu\varepsilon} \frac{N}{N_i^\varepsilon(2b)} \right),$$

and of

$$T_\varepsilon \leq \tilde{\mathcal{O}} \left(\frac{\kappa\sigma^2}{\mu\varepsilon} \frac{N}{N_i^\varepsilon(2b)} \right),$$

if $m \leq 1/2$.

In Theorem 6.b upper-bound, the first term is the optimization term, and vanishes exponentially quickly in front of the second and third ones (if they are non null), and is not reminiscent of our personalized problem. The second is the noise term (referred to as the statistical term), while the third is a bias term, making appear as advertised in the introduction a *bias-variance* trade-off. In Theorem 6.a, only the statistical and the bias terms are present. In both cases, a right choice of λ_{ij} matches the lower bound of Theorem 1 in terms of sample complexity, for $\varepsilon > 0$ small enough such that the optimization term vanishes in front of the statistical one. In Theorem 6.a, D plays the role of r in the lower bound (and can be more restrictive). Parameters μ, L in Theorem 6.b can take any value, while they are fixed in our lower bounds (with $\kappa = 1$).

For the two proofs below, we write

$$g_\lambda^k(x) = \sum_{j=1}^N \lambda_j G^k(x)_j.$$

D.1. Proof of Theorem 6.a. Assume first that the f_j are differentiable. Observe first that under noise assumption (N.1) and with the asynchronous oracle:

$$\mathbb{E} \left[\|g_\lambda^k(x)\|^2 \right] \leq \sum_{j=1}^N N \lambda_j^2 B^2, \quad x \in \mathbb{R}^d.$$

We now proceed following classical SGD analysis, but with biased gradients here. Denote $e_i^k = \mathbb{E} \left[\|x^k - x_i^*\|^2 \right]$. We have:

$$\begin{aligned} e_i^{k+1} &= e_i^k - 2\mathbb{E} \left[\left\langle \sum_{j=1}^N \lambda_j \nabla f_j(x^k), x^k - x_i^* \right\rangle \right] + \eta^2 \mathbb{E} \left[\|g_\lambda^k(x^k)\|^2 \right] \\ &\leq e_i^k - 2\eta \mathbb{E} \left[\langle \nabla f(x^k), x^k - x_i^* \rangle \right] + \mathbf{2\eta} \mathbb{E} \left[\left\langle \nabla f(x^k) - \sum_{j=1}^N \lambda_j \nabla f_j(x^k), x^k - x_i^* \right\rangle \right] + \eta^2 \sum_{j=1}^N N \lambda_j^2 B^2, \end{aligned}$$

where the bold term is the bias term, that we bound by $D \sum_j \lambda_j \sqrt{\tilde{b}_{ij}} = \sum_j \lambda_j b_{ij}$ using bias assumption (2) and $\tilde{b}_{ij} = (b_{ij}/D)^2$. Then, since $-2\eta \mathbb{E} \left[\langle \nabla f(x^k), x^k - x_i^* \rangle \right] \leq -2\eta \mathbb{E} F_i^k$, we have, summing the above inequality for $1 \leq k \leq K$:

$$\sum_{k=0}^{K-1} F_i^k \leq \frac{e_i^0}{2\eta} + K \frac{\eta}{2} N B^2 \sum_{j=1}^N \lambda_j^2 + K \sum_{j=1}^N \lambda_j b_{ij}.$$

We obtain the desired result by dividing by K and for the choice of η as in Theorem 3.a. In the case where we have access to subdifferentials only, the proof stays identical (as in Lemma 6).

D.2. Proof of Theorem 6.b. Denoting $f^\lambda(x) = \sum_{j=1}^N \lambda_j f_j(x)$, observe that f^λ is μ -strongly convex and L -smooth, and for all $x \in \mathbb{R}^d$, under the synchronous oracle and noise assumption (N.2):

$$\begin{aligned} \mathbb{E} [g_\lambda^k(x)] &= \nabla f^\lambda(x), \\ \mathbb{E} \left[\|g_\lambda^k(x) - \nabla f^\lambda(x)\|^2 \right] &\leq \sigma^2 \sum_{j=1}^N \lambda_j^2. \end{aligned} \tag{14}$$

Applying Lemma 4 (SGD under strong convexity and smoothness assumptions), we have, where x^λ minimizes f^λ :

$$f^\lambda(x^k) - f^\lambda(x^\lambda) \leq F_\lambda^0 e^{-\frac{k}{2\kappa}} + \tilde{\mathcal{O}} \left(\frac{L\sigma^2}{\mu^2 k} \sum_{1 \leq j \leq N} \lambda_j^2 \right).$$

Then, under either bias assumption (B.2) or (B.3), we have:

$$\begin{aligned} \|\nabla f_i(x^k)\|^2 &\leq 2\|\nabla f_i(x^k) - \nabla f^\lambda(x^k)\|^2 + 2\|\nabla f^\lambda(x^k)\|^2 \\ &\leq 2(m\|\nabla f^\lambda(x^k)\|^2 + \sum_{j=1}^N \lambda_{ij} \tilde{b}_{ij}) + 2\|\nabla f^\lambda(x^k)\|^2 \\ &\leq 2 \sum_{j=1}^N \lambda_{ij} \tilde{b}_{ij} + 4L(m+1)(f^\lambda(x^k) - f^\lambda(x^\lambda)). \end{aligned}$$

We conclude using $f_i(x^k) - f_i(x_i^*) \leq \frac{1}{2\mu} \|\nabla f_i(x^k)\|^2$ and $\tilde{b}_{ij} = \mu b_{ij}$.

This result is valid under any $m \geq 0$, and gives the first part of Theorem 7.2. If $m \leq 1/2$, one can use Theorem 1 in Chayti et al. [2021] with the set of agents j such that $b_{ij} \leq \varepsilon$, agent 0 is agent i , with $\alpha = 1/\mathcal{N}_i^\varepsilon(2b)$, to obtain the sample complexity without the κ factor.

APPENDIX E. All-for-all UPPER-BOUNDS

We first begin with the following simple lemma.

Lemma 3. *If Λ is a stochastic matrix and if bias assumption (B.1) holds ($f_i(x_j^*) - f_i(x_i^*) \leq b_{ij}$ for all $1 \leq i, j \leq N$):*

$$f^\Lambda(x^\Lambda) - \bar{f}(x^*) \leq \frac{1}{N} \sum_{1 \leq i, j \leq N} \lambda_{ij} b_{ij}.$$

Proof. Writing the optimality of x^Λ gives:

$$\begin{aligned} f^\Lambda(x^\Lambda) &\leq f^\Lambda(x^*) \\ &= \frac{1}{N} \sum_i f_i \left(\sum_j \lambda_{ij} x_j^* \right) \\ &\leq \frac{1}{N} \sum_{1 \leq i, j \leq N} \lambda_{ij} f_i(x_j^*), \end{aligned}$$

where we used convexity of each f_i . Then, subtracting $\bar{f}(x^*)$ and using stochasticity of Λ :

$$f^\Lambda(x^\Lambda) - \bar{f}(x^*) \leq \frac{1}{N} \sum_{1 \leq i, j \leq N} \lambda_{ij} (f_i(x_j^*) - f_i(x_i^*)).$$

□

E.1. Proof of Theorem 3.a. Using Lemma 3 and the proof sketch ($f^\Lambda(y^k) = \bar{f}(x^k)$ and the bias variance decomposition):

$$\bar{f}(x^k) - \bar{f}(x^*) \leq f^\Lambda(y^k) - f^\Lambda(x^\Lambda) + \frac{1}{N} \sum_{1 \leq i, j \leq N} \lambda_{ij} b_{ij},$$

where (y^k) is generated using simple SGD on f^Λ :

$$y^{k+1} = y^k - \eta \nabla G_\Lambda^k(y^k),$$

where

$$G_\Lambda^k(y) = (\lambda_{i_k j} g_{i_k}^k((\Lambda y^k)_{i_k}))_{1 \leq j \leq N}.$$

f^Λ is convex, and we have under noise assumption (N.1) for the asynchronous oracle, under differentiable assumptions:

$$\begin{aligned} \mathbb{E} [G_\Lambda^k(y)] &= \nabla f^\Lambda(y), \\ \mathbb{E} [\|G_\Lambda^k(y)\|^2] &\leq \frac{B^2}{N} \sum_{1 \leq i, j \leq N} \lambda_{ij}^2. \end{aligned}$$

If we only have subdifferentials, we still have $\mathbb{E} [G_\Lambda^k(y)] \in \partial f^\Lambda(y)$. Using Lemma 6 (SGD under these assumptions), we obtain the desired result.

E.2. Proof of Theorem 3.b. Under noise assumption (N.2) and for the synchronous oracle, we still have:

$$y^{k+1} = y^k - \eta \nabla G_\Lambda^k(y^k),$$

for

$$G_\Lambda^k(y) = \frac{1}{N} \left(\sum_{i=1}^N \lambda_{ij} g_i^k((\Lambda y^k)_i) \right)_{1 \leq j \leq N},$$

that verifies:

$$\begin{aligned} \mathbb{E} [G_\Lambda^k(y)] &= \nabla f^\Lambda(y), \\ \mathbb{E} [\|G_\Lambda^k(y) - \nabla f^\Lambda(y)\|^2] &\leq \frac{\sigma^2}{N^2} \sum_{1 \leq i, j \leq N} \lambda_{ij}^2. \end{aligned}$$

The function f^Λ is however not necessarily strongly convex. However, since $\nabla^2 f^\Lambda(y) = \Lambda^\top \nabla^2 \bar{f}(\Lambda y) \Lambda$ and \bar{f} is L/N -smooth and μ/N -strongly convex, f^Λ is L/N -relatively smooth and μ/N -relatively strongly convex [Bauschke et al., 2017] with respect to $\frac{1}{2} \|y\|_W^2 = \frac{1}{2} y^\top W y$. Note also that the

spectral radius of W is 1, since Λ is stochastic. Instead of using stochastic Bregman gradient descent (e.g. Dragomir et al. [2021]), we use Lemma 4: classical SGD that naturally generalizes to relative smoothness and strong convexity assumptions, when the mirror map is quadratic.

Algorithm 2 *All-for-all algorithm: Time-adaptive variant*

- 1: Stepsizes $(\eta_k)_{k \geq 0}$, matrixes $(W^{(k)})_{k \geq 0}$
- 2: Initialization $x_1^0 = \dots = x_N^0 \in \mathbb{R}^d$
- 3: **for** $k = 0, 1, 2, \dots$ **do**
- 4: Agents $j \in S^k$ (activated agents) compute stochastic gradients $g_j^k(x_j^k)$ and broadcast it to all agents i such that $W_{ij}^{(k)} > 0$.
- 5: For $i = 1, \dots, N$, update

$$x_i^{k+1} = x_i^k - \eta_k \sum_{j \in S^k} W_{ij}^{(k)} g_j^k(x_j^k)$$

- 6: **end for**
-

E.3. Time-adaptive variant.

Proposition 3 (Time-adaptive all-for-all). *Assume that the same assumptions as in Theorem 3.a hold. For any $\varepsilon > 0$, denote $\Lambda_\varepsilon = (\frac{\mathbb{1}_{2b_{ij} \leq \varepsilon}}{N_i^\varepsilon(2b)})$, $K_\varepsilon = \frac{4D^2 B^2}{\varepsilon^2} \sum_{i=1}^N \frac{1}{N_i^\varepsilon(2b)}$ and $\eta(\varepsilon) = \frac{2ND^2}{K_\varepsilon B^2 \sum_i 1/N_i^\varepsilon(2b)}$. For any $k \geq 0$, there exists $p_k \geq 0$ such that $\sum_{q=0}^{p_k-1} T_{2^{-q}} \leq k < \sum_{q=0}^{p_k} T_{2^{-q}}$, and choose:*

$$\eta_k = \eta(2^{-p_k}), \quad W^{(k)} = \Lambda_{2^{-p_k}} \Lambda_{2^{-p_k}}^\top.$$

Then, for any $\varepsilon > 0$, the iterates generated by Algorithm 2 reach averaged precision ε for a number of data item sampled from personal distribution of:

$$T_\varepsilon \leq \frac{16D^2 B^2}{\varepsilon^2} \sum_{i=1}^N \frac{1}{N_i^\varepsilon(2b)},$$

Similarly, under strong-convexity and smoothness assumptions as in Theorem 3.b, one would obtain the corresponding sample complexity, up to a factor 2.

E.4. Setting 2 under $\mu = 0$.

Theorem 8 (All-for-all convex case). *Let $K > 0$, $\eta > 0$, and W a symmetric non-negative random matrix of the form $W = \Lambda \Lambda^\top$ for some stochastic matrix $\Lambda = (\lambda_{ij})_{1 \leq i, j \leq N}$. Let $(x_i^k)_{k \geq 0, 1 \leq i \leq N}$ be generated with Algorithm 1. Assume that bias assumption (B.1) holds for some (b_{ij}) . Under Setting 2, for $\mu = 0$ (convex case), if $\eta = \frac{N^{3/2} D}{\sigma \sqrt{\sum_{i,j} \lambda_{ij}^2}} \wedge \frac{N}{2L}$, where $\|x_i^0 - x_i^\Lambda\| \leq D$ for all i , we have:*

$$\mathbb{E}[F^K] \leq \frac{2LD^2}{K} + \sqrt{\frac{2D^2 \sigma^2}{KN} \sum_{1 \leq i, j \leq N} \lambda_{ij}^2 + \sum_{1 \leq i, j \leq N} \lambda_{ij} b_{ij}}.$$

Consequently, for $\varepsilon > 0$, and a choice of $W = \Lambda \Lambda^\top$ for $\lambda_{ij} = \frac{\mathbb{1}_{\{b_{ij} \leq \varepsilon/2\}}}{N_i^\varepsilon(2b)}$ and under the same assumptions, the all-for-all algorithm reaches $\mathbb{E}[F^K] \leq \varepsilon$ for a total number T_ε of data item sampled from personal distributions of:

$$T_\varepsilon \leq \max \left(\frac{8D^2 \sigma^2}{\varepsilon^2} \sum_{i=1}^N \frac{1}{N_i^\varepsilon(2b)}, \frac{4NLD^2}{\varepsilon} \right)$$

Proof. Combining Lemma 5 with the noise variance and regularity parameters of F^Λ . \square

APPENDIX F. STOCHASTIC OPTIMIZATION TOOLBOX

F.1. SGD under strongly convex and smooth assumptions.

Lemma 4 (SGD, s.c. and smooth). *Define $\|x\|_A^2 = x^\top Ax$ for some non-negative and symmetric matrix A . Let $f : \mathcal{X} \rightarrow \mathbb{R}$ μ -relatively strongly convex and L -relatively smooth with respect to $\frac{1}{2}\|x\|_A^2$. Let $(f_t, g_t)_{t \geq 0}$ be first order oracle calls such that for all $t \geq 0$:*

$$\forall x \in \mathcal{X}, \quad \begin{cases} \mathbb{E}[f_t(x)] = f(x), \\ \mathbb{E}[g_t(x)] = \nabla f(x), \\ \mathbb{E}[\|g_t(x) - \nabla f(x)\|^2] \leq \sigma^2, \end{cases}$$

for some $\sigma > 0$. Let L_A be the largest eigenvalue of A , and assume that $L_A \leq 1$ (our result generalizes to any L_A). Let $(x_t)_{t \geq 0}$ be generated with:

$$\forall t \geq 0, \quad x^{t+1} = x^t - \eta g_t(x^t),$$

for a fixed stepsize $\frac{1}{2L} \geq \eta > 0$, and assume that all the iterates lie in \mathcal{X} . Assume that f is minimized over \mathcal{X} at some interior point x^* . We have for any $T > 0$:

$$\mathbb{E}[f(x^T) - f(x^*)] \leq e^{-\eta\mu T} (f(x^0) - f(x^*)) + \frac{\eta L \sigma^2}{\mu}.$$

For fixed $T > 0$, setting $\eta = \min(1/(2L), \frac{1}{\mu T} \ln(\frac{f_0 \mu^2 T}{L \sigma^2}))$ gives:

$$\mathbb{E}[f(x^T) - f(x^*)] \leq e^{-\frac{\mu}{2L} T} (f(x^0) - f(x^*)) + \frac{L \sigma^2}{\mu^2 T} \ln\left(\frac{f_0 \mu^2 T}{L \sigma^2}\right).$$

Thus, for fixed target precision $\varepsilon > 0$, using stepsize $\eta_\varepsilon = \min(\frac{\mu \varepsilon}{2L \sigma^2}, \frac{1}{2L})$ and setting $T_\varepsilon = \lceil \ln(\varepsilon^{-1}(f(x^0) - f(x^*))) \frac{1}{\eta_\varepsilon \mu} \rceil$, we have:

$$f\left(\frac{1}{T_\varepsilon} \sum_{t < T_\varepsilon} x^t\right) - f(x^*) \leq \varepsilon,$$

with a number of oracle calls

$$T_\varepsilon \leq \max\left(\frac{2L \sigma^2}{\varepsilon \mu^2}, \frac{2L}{\mu}\right) \ln(\varepsilon^{-1}(f(x^0) - f(x^*))).$$

Proof. For some $t \geq 0$, denoting $f_t = \mathbb{E}[f(x^{t+1}) - f(x^*)]$, using relative smoothness, unbiasedness of the stochastic gradients and then relative strong convexity:

$$\begin{aligned} f_{t+1} - f_t &\leq -\eta \mathbb{E}[\|\nabla f(x^t)\|^2] + \frac{\eta^2 L}{2} \mathbb{E}[\|g_t\|_A^2] \\ &\leq -\eta \mathbb{E}[\|\nabla f(x^t)\|^2] + \frac{\eta^2 L L_A}{2} \mathbb{E}[\|g_t\|^2] \\ &\leq -\eta \left(1 - \frac{\eta L L_A}{2}\right) \mathbb{E}[\|\nabla f(x^t)\|^2] + \frac{\eta^2 L L_A \sigma^2}{2}. \end{aligned}$$

Using relative strong convexity of f , we have:

$$\begin{aligned} \|\nabla f(x^t)\|^2 &\geq \frac{1}{L_A} \|\nabla f(x^t)\|_A^2 \\ &\leq \frac{2\mu}{L_A} f_t, \end{aligned}$$

yielding, for $\eta < 1/(LL_A)$

$$f_{t+1} - f_t \leq -2\eta \frac{\mu}{L_A} f_t + \frac{\eta^2 L L_A \sigma^2}{2}.$$

Then, for some $T > 0$ and since $L_A \leq 1$, sum the above inequality multiplied by $(1 - \eta\mu)^{-t-1}$:

$$\begin{aligned} \sum_{0 \leq t \leq T-1} (1 - \eta\mu)^{-t-1} f_{t+1} - (1 - \eta\mu)^{-t} f_t &\leq \frac{\eta^2 L \sigma^2}{2} \sum_{0 \leq t \leq T-1} (1 - \eta\mu)^{-t-1} \\ &\leq \frac{\eta^2 L \sigma^2}{2} \frac{(1 - \eta\mu)^{-T-1}}{\eta\mu}, \end{aligned}$$

leading to the desired result. \square

F.2. SGD under smoothness and convexity assumptions.

Lemma 5 (SGD, convex and smooth). *Let $f : \mathcal{X} \rightarrow \mathbb{R}$ convex and L -smooth. Let $(f_t, g_t)_{t \geq 0}$ be first order oracle calls such that for all $t \geq 0$:*

$$\forall x \in \mathcal{X}, \quad \begin{cases} \mathbb{E}[f_t(x)] = f(x), \\ \mathbb{E}[g_t(x)] = \nabla f(x), \\ \mathbb{E}[\|g_t(x) - \nabla f(x)\|^2] \leq \sigma^2, \end{cases}$$

for some $\sigma > 0$. Let $(x_t)_{t \geq 0}$ be generated with:

$$\forall t \geq 0, \quad x^{t+1} = x^t - \eta g_t(x^t),$$

for a fixed stepsize $\frac{1}{2L} \geq \eta > 0$, and assume that all the iterates lie in \mathcal{X} . Assume that f is minimized over \mathcal{X} at some interior point x^* . Denote, for $T > 0$, $\bar{x}^T = \frac{1}{T} \sum_{0 \leq t < T} x^t$. We have for any $T > 0$:

$$\mathbb{E}[f(\bar{x}^T) - f(x^*)] \leq \frac{\|x^0 - x^*\|^2}{\eta T} + \eta \sigma^2.$$

For fixed $T > 0$, setting $\eta = \min\left(1/(2L), \frac{\|x^0 - x^*\|}{\sigma \sqrt{T}}\right)$ gives:

$$\mathbb{E}[f(x^T) - f(x^*)] \leq \frac{2\|x^0 - x^*\|\sigma}{\sqrt{T}} + \frac{2L\|x^0 - x^*\|^2}{T}.$$

Thus, for fixed target precision $\varepsilon > 0$, we have:

$$f\left(\frac{1}{T_\varepsilon} \sum_{t < T_\varepsilon} x^t\right) - f(x^*) \leq \varepsilon,$$

with a number of oracle calls

$$T_\varepsilon \leq 4\|x^0 - x^*\|^2 (4\sigma^2 \varepsilon^{-2} + L\varepsilon^{-1}).$$

Proof. Denote $e_t = \mathbb{E}[\|x^t - x^*\|^2]$. We have:

$$\begin{aligned} e_{t+1} - e_t &\leq -2\eta \langle \nabla f(x^t), x^t - x^* \rangle + \eta^2 \mathbb{E}[\|g_t^2\|] \\ &\leq -2\eta f_t + \eta^2 (\mathbb{E}[\|\nabla f(x^t)\|^2] + \sigma^2) \\ &\leq -2\eta(1 - L\eta)f_t + \eta^2 \sigma^2 \\ &\leq -\eta f_t + \eta^2 \sigma^2, \end{aligned}$$

where we used convexity and smoothness assumptions. Then, summing:

$$\sum_{t < T} f_t \leq \frac{e_0 - e_T}{\eta} + T\eta\sigma^2,$$

leading to the desired result. \square

F.3. SGD under convex assumptions.

Lemma 6 (SGD, convex). *Let $f : \mathcal{X} \rightarrow \mathbb{R}$ convex. Let $(f_t, g_t)_{t \geq 0}$ be first order oracle calls such that for all $t \geq 0$:*

$$\forall x \in \mathcal{X}, \quad \begin{cases} \mathbb{E}[f_t(x)] = f(x), \\ \mathbb{E}[g_t(x)] \in \partial f(x), \\ \mathbb{E}[\|g_t(x)\|^2] \leq B^2, \end{cases}$$

for some $B > 0$. For $p_{\mathcal{X}}$ the projection on the convex set \mathcal{X} , let $(x_t)_{t \geq 0}$ be generated with:

$$\forall t \geq 0, \quad x^{t+1} = p_{\mathcal{X}}(x^t - \eta g_t(x^t)),$$

for a fixed stepsize $\eta > 0$. Assume that f is minimized over \mathcal{X} at some interior point x^* . We have for any $T > 0$ and for $\eta = \sqrt{\frac{2D^2}{B^2T}}$ for $D^2 \geq \|x^0 - x^*\|^2$:

$$f\left(\frac{1}{T} \sum_{t < T} x^t\right) - f(x^*) \leq \sqrt{\frac{2}{T}} BD.$$

Thus, for fixed target precision $\varepsilon > 0$, setting $T_\varepsilon = \lceil 2\varepsilon^{-2} B^2 D^2 \rceil$ and using stepsize $\eta_\varepsilon = \frac{\sqrt{2}\varepsilon}{B^2}$, we have:

$$f\left(\frac{1}{T_\varepsilon} \sum_{t < T_\varepsilon} x^t\right) - f(x^*) \leq \varepsilon.$$

Proof. For any $y \in \mathcal{X}$, using properties of $p_{\mathcal{X}}$:

$$\begin{aligned} \|x^{t+1} - y\|^2 &= \|p_{\mathcal{X}}(x^t - \eta g_t(x^t)) - y\|^2 \\ &\leq \|x^t - \eta g_t(x^t) - y\|^2 \\ &= \|x^t - y\|^2 + \eta^2 \|g_t(x^t)\|^2 - 2\eta \langle g_t(x^t), x^t - y \rangle. \end{aligned}$$

Taking the mean and $y = x^*$, with $e_t = \mathbb{E}[\|x^t - x^*\|^2]$ and $f_t = \mathbb{E}[f(x^t) - f(x^*)]$:

$$\begin{aligned} e_{t+1} &\leq e_t + \eta^2 B^2 - 2\eta \langle \nabla f(x^t), x^t - y \rangle \\ &\leq e_t + \eta^2 B^2 - 2\eta f_t, \end{aligned}$$

where we used convexity of f . Thus, $f_t \leq \frac{e_t - e_{t+1}}{2\eta} + \frac{\eta B^2}{2}$, and by summing this inequality for $0 \leq t < T$:

$$\sum_{t < T} f_t \leq \frac{e_0 - e_T}{2\eta} + \frac{\eta T B^2}{2} \leq \frac{e_0}{\eta} + \frac{\eta T B^2}{2}.$$

Thus, for $\eta = \sqrt{\frac{2e_0}{B^2T}}$, we have the result using convexity of f . \square