
FLIX: A Simple and Communication-Efficient Alternative to Local Methods in Federated Learning

Elnur Gasanov
KAUST

Ahmed Khaled
Princeton University

Samuel Horváth
KAUST

Peter Richtárik
KAUST

Abstract

Federated Learning (FL) is an increasingly popular machine learning paradigm in which multiple nodes try to collaboratively learn under privacy, communication and multiple heterogeneity constraints. A persistent problem in federated learning is that it is not clear what the optimization objective should be: the standard average risk minimization of supervised learning is inadequate in handling several major constraints specific to federated learning, such as communication adaptivity and personalization control. We identify several key desiderata in frameworks for federated learning and introduce a new framework, FLIX, that takes into account the unique challenges brought by federated learning. FLIX has a standard finite-sum form, which enables practitioners to tap into the immense wealth of existing (potentially non-local) methods for distributed optimization. Through a smart initialization that does not require any communication, FLIX does not require the use of local steps but is still provably capable of performing dissimilarity regularization on par with local methods. We give several algorithms for solving the FLIX formulation efficiently under communication constraints. Finally, we corroborate our theoretical results with extensive experimentation.

1 INTRODUCTION

There is a wealth of data today that is decentralized among many clients and which can not be centralized for privacy reasons. Federated Learning (FL) aims to

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

enable machine learning in this decentralized setting while respecting data privacy. Application domains of federated learning include healthcare, learning language models for virtual keyboards, and speech recognition (Kairouz et al., 2019). The promise of federated learning is that by participating in a distributed training process, clients can learn better machine learning models than they can using only their own data. The main cost in using federated learning over local training lies in the network bandwidth used for the distributed training process. Hence, federated learning must be flexible enough to provide a benefit to users without a prohibitive communication cost. The standard formulation of FL is to cast it as an optimization problem of the form

$$\min_{x \in \mathbb{R}^d} \left[f(x) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n f_i(x) \right], \quad (\text{ERM})$$

where f_i is the loss function on client i . Thus, the goal of classical FL is for the n clients to collaboratively learn a single model, $x^* = \arg \min f$, to be deployed on all clients. Recent development shows that using a single model for all clients can be severely detrimental to individual performance on many clients (Yu et al., 2020), defeating the purpose of joining distributed training. Furthermore, (ERM) offers no clear tunable knobs that can accommodate constraints on the network bandwidth. Motivated by these concerns, the chief question of our paper is:

Can we find a formulation for federated learning that is flexible enough to accommodate the needs of federated learning, yet also solvable using standard methods?

1.1 Key properties of the FLIX framework

Our main contribution is *FLIX* (Federated Learning mIXtures), a novel and flexible formulation for federated learning: define $\alpha_i > 0$ to be the *personalization parameter for node i* , and let $x_i \stackrel{\text{def}}{=} \min_{x \in \mathbb{R}^d} f_i(x)$ be the local solution to the i -th objective—note that x_i

can be found by solely running a local optimizer, and hence computing it requires no communication at all. If f_i is nonconvex, then x_i can be merely a stationary point, i.e. a point such that $\nabla f_i(x_i) = 0$. The FLIX problem is

$$\min_{x \in \mathbb{R}^d} \tilde{f}(x) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i) x_i). \quad (\text{FLIX})$$

Once we find a solution x_* of sufficient quality for (FLIX), we deploy $T_i(x_*) = \alpha_i x_* + (1 - \alpha_i) x_i$ on node i as its final personalized model. We now enumerate some of the key properties of (FLIX):

- Efficiently solvable as a finite-sum problem.** Formulation (FLIX) preserves the standard finite-sum formulation of empirical risk minimization. Moreover, it preserves problem structure: we show (in Section 3) that when the f_i are smooth (and/or convex), \tilde{f} is also smooth (resp. convex). Hence, we can leverage standard empirical risk minimization method for federated learning under communication and/or personalization constraints. In Section 3 we instantiate several such methods and show they enjoy fast convergence guarantees under standard assumptions on the functions f_1, \dots, f_n .
- Adaptive to communication constraints.** Communication efficiency is an important concern federated learning, as often bandwidth is valuable and limited (Konečný et al., 2016; Li et al., 2019; Kairouz et al., 2019). The FLIX formulation is adaptive to communication constraints: observe that computing x_i , a precondition to solving (FLIX), requires no communication at all, and can be done purely locally on node i . If $\alpha_i = 0$, then no communication at all is needed to compute the personalized model $T_i(x_*)$. By varying α_i between 0 and 1, we can control the amount of communication needed to compute $T_i(x_*)$. We show (in Section 3) that given a communication budget of R steps, we can find parameters α_i that allows us to solve (FLIX) in no more than R communication steps.
- Adaptivity to personalization.** Our end-goal in federated learning is to generalize well *on each client*: this means that the solution deployed on node i should be tailored to its local data distribution, which may differ from the data distributions on other nodes. Hence, we want to use the data from other nodes to train the model deployed on node i only if that data is useful on its own distribution. Indeed, a naive application of empirical risk minimization can have terrible effects in federated learning situations where the pure local model (i.e. x_i) does better than average empirical risk minimizer (Yu et al., 2020). In FLIX, varying α_i enables us to amplify or reduce the

effect of other objectives on the solution deployed on node i . In situations where the data on all of the nodes is sufficiently heterogeneous, we set α_i to be small and the effect of other data on node i will be negligible. On the other hand, when the data on the different nodes is related we may set α_i to be closer to 1. We observe a benefit to varying α in this manner in practice: Figure 1 shows the effect of varying the α_i on real data (see Section 4 for the details and for other experiments).

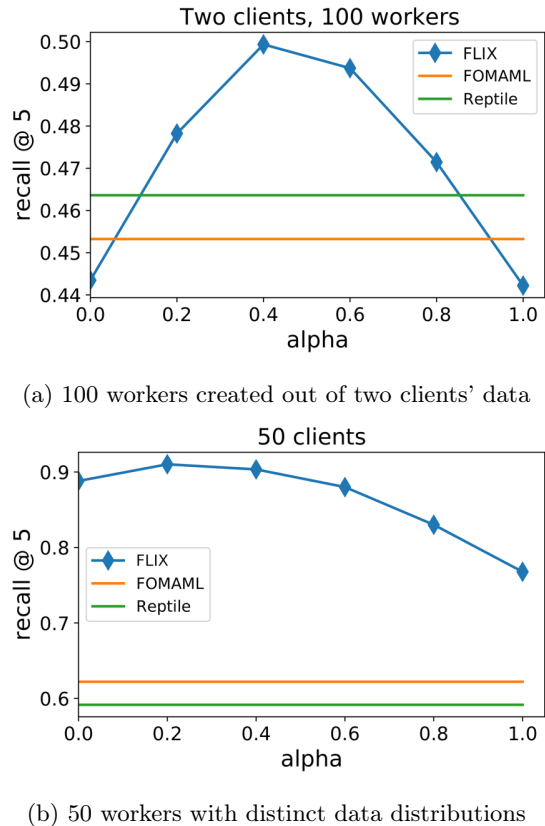


Figure 1: Test accuracy of FLIX model for different personalization parameter values, FOMAML and Reptile. $\alpha_i = \alpha$ is set to the value indicated on horizontal axis. FOMAML and Reptile are independent from the personalization parameter α . Plots correspond to different data splittings.

FLIX fills a gap that is unsatisfied by existing methods. To the best of our knowledge, there is no other method for federated learning that is efficiently solvable via standard algorithms and also adaptive to communication and personalization constraints, and indeed both constraints are important in practice (Li et al., 2020). We believe the key properties we enumerate can also serve as natural desiderata in the development of new formulations and methods for federated learning.

1.2 Related work

Personalization has garnered significant recent interest in federated learning as personalized models often perform well in practice compared to non-personalized models (Jiang et al., 2019; Yu et al., 2020). FLIX is a *model mixture* method: the personalized solution is a mixture of a global model and a local model. In recent work, Deng et al. (2020) and Mansour et al. (2020) propose model mixture methods and prove their statistical benefits, while Zec et al. (2021) introduce a similar formulation based on the mixture of experts framework. Unfortunately, we show in the supplementary material that, from the perspective of optimization and without additional data, the formulations in all three works are trivially minimized at the local minimizers x_1, \dots, x_n . An alternative to model mixing is mixing in function space, where we optimize a mixture of objectives rather than a model mixture. This mixture is often constructed to control model variance: examples of this approach can be found in (Hanzely and Richtárik, 2020; T. Dinh et al., 2020; Huang et al., 2021; Dinh et al., 2021). In FLIX, we take the model mixture approach as it allows us to use pretraining to better solve the problem while still regularizing model variance (see Section 2.1). A parallel line of work applies meta-learning methods like MAML to federated learning (Jiang et al., 2019; Fallah et al., 2020a): in Section 2.2 we motivate FLIX by taking MAML as our starting point. Chen et al. (2021) discuss the statistical limits of personalization and show that either solving empirical risk minimization or local training is optimal, depending on certain problem parameters; However, as of yet there is no single optimal adaptive algorithm (from the statistical perspective). There are several other techniques in federated learning that can be combined with our approach for better results, such as clustering (Sattler et al., 2020) or robust optimization (Reisizadeh et al., 2020).

2 THE FLIX FORMULATION

In this section we reintroduce and motivate the FLIX formulation in detail. We define the FLIX objective as

$$\tilde{f}(x; \alpha_1, \dots, \alpha_n, x_1, \dots, x_n) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i) x_i), \quad (1)$$

where $\alpha_i \in (0, 1)$ is the personalization coefficient for node i and x_i is the minimizer of f_i , for all $i = 1, 2, \dots, n$. We will use $\tilde{f}(x)$ to refer to the objective in (1) when the α_i and x_i are clear from the

context. The FLIX problem is then

$$\min_{x \in \mathbb{R}^d} \left[\tilde{f}(x) = \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i) x_i) \right]. \quad (2)$$

Let $\alpha = [\alpha_1, \dots, \alpha_n]$ be the vector of the personalization coefficients. If $x_* = x_*(\alpha)$ is a solution of (2), we call $T_i(x; \alpha_i, x_i) = \alpha_i x_* + (1 - \alpha_i) x_i$ the *deployed solution on node i* . Like with \tilde{f} , we will refer to the deployed solution on node i as $T_i(x)$ when x_i and α_i are clear from the context.

2.1 Motivation 1: from Local GD to FLIX

The most popular algorithm for solving federated learning problems is the Federated Averaging algorithm (Kairouz et al., 2019), also known as Local (Stochastic) Gradient Descent (Local GD/SGD). Local GD alternates steps of local computation on each node with steps of communication and aggregation. More concretely, the Local GD update is:

$$x_{t+1}^i = \begin{cases} x_t^i - \gamma \nabla f_i(x_t^i) & \text{if } t \bmod H \neq 0 \\ \frac{1}{n} \sum_{i=1}^n [x_t^i - \gamma \nabla f_i(x_t^i)] & \text{if } t \bmod H = 0 \end{cases},$$

where H is the number of local steps. Early papers on federated learning (such as e.g. (Konečný et al., 2016)) motivated local methods as communication-efficient ways of solving (ERM), but subsequent theoretical development reveals that local methods are, in fact, quite bad solvers for (ERM) whenever there is significant statistical heterogeneity among the clients (Woodworth et al., 2020). Moreover, Pathak and Wainwright (2020) show that for the linear least-squares problem, Local GD converges to a different point than the minimizer of (ERM). More generally, the fixed points of Local GD can be very different from the minimizer of (ERM) whenever $H > 1$ (Malinovskiy et al., 2020). Hanzely and Richtárik (2020) show that a mild variant of Local GD can be interpreted as SGD applied on the nd -dimensional regularized objective f_λ defined by

$$f_\lambda(y_1, y_2, \dots, y_n) \stackrel{\text{def}}{=} \left[\frac{1}{n} \sum_{i=1}^n f_i(y_i) + \frac{\lambda}{2n} \sum_{i=1}^n \|y_i - \bar{y}\|^2 \right],$$

where $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$ is the counterpart of x in (ERM), and where λ is a regularization parameter determined according to the number of local steps. Objective f_λ is the summation of two terms: the first asks that each node i finds a solution y_i that minimizes its local objective well, while the regularizer $\psi(y_1, \dots, y_n) = \frac{1}{2n} \sum_{i=1}^n \|y_i - \bar{y}\|^2$ forces the solutions y_1, y_2, \dots, y_n to be close to their average \bar{y} . Hence, Local GD incentivizes finding *personalized* solutions y_1, y_2, \dots, y_n that have small population variance. Hanzely and

Richtárik (2020) note that as the λ parameter varies between 0 and ∞ , the solutions found by Local GD interpolate between the pure local optimal models (i.e. $x_i = \operatorname{argmin}_x f_i(x)$) and the solution of the global problem x_* (the minimizer of (ERM)). Our starting point is then the following observation:

The solutions y_1, \dots, y_n found by Local GD are an *implicit mixture* of the local minimizers x_1, \dots, x_n and the global empirical risk minimizer x_* .

Rather than seeking an implicit mixture of the local and global optimal models, we instead propose to find an *explicit* mixture of the local optimal models and a global model: given any global model x (not necessarily the empirical risk minimizer), we choose coefficients $\alpha_1, \alpha_2, \dots, \alpha_n$ (all between 0 and 1) and then deploy on node i the mixture

$$T_i(x) = \alpha_i x + (1 - \alpha_i) x_i, \quad (3)$$

we may then choose x as *the best* such global model by explicitly solving the optimization problem

$$\min_{x \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i) x_i),$$

and this is exactly the FLIX formulation. Observe that coefficients $\alpha_1, \alpha_2, \dots, \alpha_n$ should regularize the population variance of the deployed solutions $T_1(x), T_2(x), \dots, T_n(x)$, as in local methods. We show this rigorously for equal α_i in Proposition 3. Our development thus leads us to a natural framework that captures the strength of local methods while also satisfying the desiderata specified in Section 1.1.

2.2 Motivation 2: from model-agnostic meta-learning to FLIX

We now motivate FLIX differently by starting with *personalization via fine-tuning*. The ordinary formulation of the federated learning problem (ERM) asks for a single global model to be used on all clients. If the clients are sufficiently heterogeneous, a single model may perform badly on many of them (Jiang et al., 2019). Personalizing a global model to each of the users’ custom data is often beneficial in practice; For example, Wang et al. (2019) study the benefits of personalizing language models for a virtual keyboard application used by tens of millions of users. They observe that a sizeable fraction of the users benefit from personalization. Personalization is often done in two steps:

Step I: initial model training. Find a “good” global model x_{global} .

Step II: fine-tuning. Personalize the global model x_{global} on each client to get the *personalized* local models x_i .

Methods that fit this framework are known as *finite-tuning approaches*: they include the model-agnostic meta-learning (MAML) family of methods (Finn et al., 2017). In addition to its practical popularity, recent theoretical investigations reveal that fine-tuning approaches, such as MAML, are also beneficial from a statistical perspective (Fallah et al., 2021; Chua et al., 2021). In MAML, we find x_{global} by optimizing for the loss after a single step of gradient descent, i.e. the MAML objective is

$$\text{Find } x_{\text{global}} \in \operatorname{argmin}_{x \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n f_i(x - \gamma \nabla f_i(x)), \quad (4)$$

where γ is a given stepsize. Once x_{global} is found, we may then fine-tune it by running gradient descent for a number of steps on each node i locally using its own objective f_i (Finn et al., 2017). To gain further insight into what fine-tuning is doing, we now consider the case when each f_i is a quadratic function. Because this problem is amenable to analysis, several authors have used it to study the theoretical properties of MAML (Collins et al., 2020; Charles and Konečný, 2021; Gao and Sener, 2020), and we follow in their footsteps. We assume that each f_i can be written as $f_i(x) = \frac{1}{2} x^T A_i x - b_i^T x + c_i$, where $A_i \in \mathbb{R}^{d \times d}$ is a positive definite matrix, $b_i \in \mathbb{R}^d$, and $c_i \in \mathbb{R}$. Now suppose that we have some initial global model x^0 , and we fine-tune it by running gradient descent for H steps on node i : the next proposition shows the final iterate is a matrix-weighted average of the initial solution and the optimal local solution:

Proposition 1. *Suppose that we run gradient descent for H steps on the quadratic objective $f_i = \frac{1}{2} x^T A_i x - b_i^T x + c$ starting from x^0 with stepsize $\gamma > 0$. Suppose that the stepsize satisfies $\gamma \leq \frac{1}{L_i}$, where $L_i = \lambda_{\max}(A_i)$. Then the final iterate x_i^H can be written as*

$$x_i^H = (I - J_i^H) x_i + J_i^H x^0,$$

where x_i minimizes f_i and $J_i \in \mathbb{R}^{d \times d}$ is a matrix with maximum eigenvalue smaller than 1, i.e. $\lambda_{\max}(J) < 1$.

The proof of Proposition 1 and all subsequent proofs are relegated to the supplementary. Plugging the result of Proposition 1 into Equation (4), observe that in MAML we find the initial model x^0 by solving the problem

$$\min_{x \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n f_i((I - J_i)x_i + J_i x).$$

Hence, MAML is optimizing *for a specific weighted average* of the initial model x^0 and the local solutions x_1, x_2, \dots, x_n . We thus propose to dispense with the specific matrix J_i and instead optimize an average weighted with an arbitrary constant α_i :

$$\min_{x \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i) x_i),$$

and this is exactly the FLIX formulation. Observe that by choosing α_i we may accomplish a similar effect to multiplying by J_i^H for large H , so this is a valid approximation at least for large H . This gives FLIX a new interpretation as an approximate generalized MAML, where we optimize the global model for performance after potentially many gradient descent steps rather than just a single step. By exploiting information about the objective's optimum x_i (available at zero communication cost in federated learning), FLIX takes into account higher-order information that ordinary MAML can not.

3 THEORY AND ALGORITHMS

In this section we aim to develop algorithms to solve (FLIX) in a communication-efficient manner. Before discussing concrete algorithms, we study a few algorithm-independent properties of (FLIX) that will come in handy for understanding the formulation and proving convergence bounds. The following proposition shows that the formulation preserves smoothness and convexity. This is in contrast to other meta-learning methods such as MAML, where the objective may generally be nonsmooth (Fallah et al., 2020b).

Proposition 2. *Suppose that each objective f_i is L_i -smooth. That is, for any $x, y \in \mathbb{R}^d$ we have $\|\nabla f_i(x) - \nabla f_i(y)\| \leq L_i \|x - y\|$. Then the FLIX objective \tilde{f} defined in (1) is L_α -smooth for $L_\alpha \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i$. If each f_i is convex, then \tilde{f} is also convex. If each f_i is μ_i -strongly convex, then \tilde{f} is μ_α strongly convex for $\mu_\alpha \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 \mu_i$.*

Our next result offers some insight into the variance-regularizing effect of the α_i : in particular, when all the α_i are equal, increasing α in (FLIX) directly decreases the variance of the deployed local models from their mean. As discussed in Section 2.1, this is a key property of local descent methods that the FLIX formulation captures.

Proposition 3. *Suppose that $\alpha_1 = \alpha_2 = \dots = \alpha_n = \beta$ in the FLIX formulation (FLIX). Let $T_1(x), T_2(x), \dots, T_n(x)$ be the deployed models defined in (3). If y_1, \dots, y_n are vectors in \mathbb{R}^d and \bar{y} is their mean, we define $V(y_1, \dots, y_n)$ as the population vari-*

ance

$$V(y_1, \dots, y_n) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \|y_i - \bar{y}\|^2.$$

Then,

$$V(T_1(x), \dots, T_n(x)) = (1 - \beta)^2 V(x_1, \dots, x_n).$$

One-shot learning is a learning paradigm where we may use only a single round of communication to solve the federated learning problem (Guha et al., 2019; Salehkaileybar et al., 2019). When the personalization parameters are small enough, we can provably solve the FLIX problem with a single round of communication by computing a certain weighted average of the local solutions x_1, x_2, \dots, x_n . We first consider if each x_i is a minimizer of f_i , then we can get an approximate minimizer:

Theorem 1. *Suppose that each objective f_i is L_i -smooth, that each x_i minimizes f_i , and let $\hat{L} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n L_i$. Given the pure local models x_1, x_2, \dots, x_n , define the weighted average*

$$x^{\text{avg}} \stackrel{\text{def}}{=} \sum_{i=1}^n w_i x_i, \quad w_i \stackrel{\text{def}}{=} \frac{\alpha_i^2 L_i}{n L_\alpha}, \quad L_\alpha \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i. \quad (5)$$

We further define the constants

$$D \stackrel{\text{def}}{=} \max_{i,j=1,\dots,n,i \neq j} \|x_i - x_j\| \quad \text{and} \quad V \stackrel{\text{def}}{=} \sum_{i=1}^n w_i \|x_i - x^{\text{avg}}\|^2 \quad (6)$$

Fix any $\epsilon > 0$. Assume that either $\max_{i=1,\dots,n} \alpha_i \leq \sqrt{2\epsilon}/\sqrt{\hat{L}D}$, or $\alpha_i = \beta$ for all i and $\beta \leq \sqrt{2\epsilon}/\sqrt{\hat{L}D}$. Then x^{avg} is an ϵ -approximate minimizer of (FLIX).

We can relax this to the requirement that each x_i is a stationary point, in which case we also get another stationary point using weighted averaging, see Theorem 3 in the supplementary for details. We point out that this implies FLIX applies equally well to objectives with minimizers (e.g. when the f_i are strongly convex) as well as when the objective is non-convex and the goal is attaining a stationary point rather than minimization. For larger α_i than that required by Theorem 1 or Theorem 3, we need more communication rounds than one. In the next subsection, we describe how distributed gradient descent can be used to solve the problem.

3.1 Distributed gradient descent

The simplest approach to solving (FLIX) is via distributed gradient descent: given the local models

x_1, x_2, \dots, x_n (precomputed before starting the process) and an initial global model x^0 , we run the update

$$x^{k+1} = x^k - \frac{\gamma}{n} \sum_{i=1}^n \alpha_i \nabla f_i (\alpha_i x^k + (1 - \alpha_i) x_i). \quad (\text{DGD})$$

Each iteration of (DGD) requires that every node sends its local gradient $\nabla \tilde{f}_i(x^k)$ to the server. The server then averages the received gradients, computes the new iterate x^{k+1} , and broadcasts it to the nodes. The next theorem shows that under smoothness and strong convexity, distributed gradient descent converges linearly to the solution of (FLIX).

Theorem 2. *Suppose that each f_i in (FLIX) is L_i -smooth and μ_i -strongly convex. Define x^{avg}, L_α , and \hat{L} by (5) and V, D by (6). Suppose that we run (DGD) for K iterations starting from $x^0 = x^{\text{avg}}$. Then the following hold:*

i) *If the α_i are allowed to be arbitrary, then for $\alpha_{\max} \stackrel{\text{def}}{=} \max_{i=1, \dots, n} \alpha_i$ we have*

$$\tilde{f}(x^k) - \min_{x \in \mathbb{R}^d} \tilde{f}(x) \leq \left(1 - \frac{\mu_\alpha}{L_\alpha}\right)^K \frac{\alpha_{\max}^2 \hat{L} D}{2}.$$

ii) *If $\alpha_i = \beta$ for all i , then*

$$\tilde{f}(x^k) - \min_{x \in \mathbb{R}^d} \tilde{f}(x) \leq \left(1 - \frac{\hat{\mu}}{\hat{L}}\right)^K \frac{\beta^2 \hat{L} V}{2}, \quad (7)$$

where $\hat{\mu} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \mu_i$.

There are four ways of making the right hand side in (2) small (i.e. decrease the communication complexity):

- **Communicate more.** Increase the number of communications K .
- **Homogeneous data.** The variance $V = \sum_{i=1}^n w_i \|x_i - x^{\text{avg}}\|^2$ can be seen as a measure of data heterogeneity. More homogeneous data sets will have smaller V , which leads to better performance.
- **Train simpler models.** Focusing attention on models with smaller \hat{L} (adjust model design), or larger $\hat{\mu}$ (e.g., add more regularization).
- **Put more weight on local models.** If we prefer local models to the global model, then α_i is small, and hence fewer communications are needed to achieve any given accuracy.

Armed with Theorem 2, we make good on our promise in Section 1.1 and show that FLIX can be solved using any communication budget. Looking at (2) we see that for any fixed $\epsilon > 0$ we have

$$\tilde{f}(x^k) - \min_{x \in \mathbb{R}^d} \tilde{f}(x) \leq \epsilon$$

as long as

$$\beta \leq Aq^k,$$

where $A = \sqrt{2\epsilon/(\hat{L}V)}$ and $q = 1/\sqrt{1 - \frac{\hat{\mu}}{\hat{L}}}$.

Putting this together leads to the following observations:

- If $\beta = 0$, the problem can be solved with 0 communications (i.e., each device i independently computes the pure local model x_i).
- If $0 < \beta \leq A$, the problem can be solved with 1 communication (i.e., compute x^{avg}). This follows from Theorem 1, and also from the more general result Theorem 2 by setting $K = 0$.
- If $A < \beta \leq Aq$, the problem can be solved with 2 communications (1 communication to compute $x^0 = x^{\text{avg}}$, followed by one iteration of distributed GD).
- If $Aq^{k-1} < \beta \leq Aq^k$, the problem can be solved with $k+1$ communications (1 communication to compute $x^0 = x^{\text{avg}}$, followed by K iterations of distributed gradient descent).
- If $\beta = 1$, we need 1 communication to compute $x^0 = x^{\text{avg}}$, followed by $k \geq \frac{\hat{L}}{\hat{\mu}} \log \frac{\hat{L}V}{2\epsilon}$ iterations of distributed gradient descent. This recovers the standard communication complexity of gradient descent needed to find the optimal solution of the average risk minimization problem (ERM).

In the supplementary, we develop other algorithms for solving (FLIX) such as distributed gradient descent with compression (Alistarh et al., 2017) and DIANA (Mishchenko et al., 2019). We point out that Theorem 8 and Theorem 9 analyze the compressed gradient descent and DIANA algorithms applied in the non-convex case, showing explicit dependence of the convergence rates on α , thus this convergence theory applies to convex as well as non-convex objectives. Because (FLIX) has a standard finite-sum form, many more algorithms can be used to solve it, e.g. accelerated minibatch SGD (Cotter et al., 2011) or SARAH (Nguyen et al., 2017). We leave the task of finding the optimal algorithm for (FLIX) to future work.

4 EXPERIMENTS

Logistic regression with l_2 regularizer. For our first experiment, we consider a setup where each device runs regularized logistic regression:

$$f_i(x) := \frac{1}{k_i} \sum_{j=1}^{k_i} [\log(1 + \exp(-a_{i,j}^\top x))] + \frac{\lambda}{2} \|x\|^2, \quad (8)$$

where $a_{i,j} \in \mathbb{R}^d$ are given for all $j = 1 \dots k_i$, k_i is a number of data points associated with device i and λ is a regularization parameter. The objective given by (8) is typically used in classification problems. It is not difficult to show that f_i is $1/4k_i \left\| \sum_{i=1}^{k_i} a_i a_i^\top \right\| + \lambda$ -smooth and λ -strongly convex.

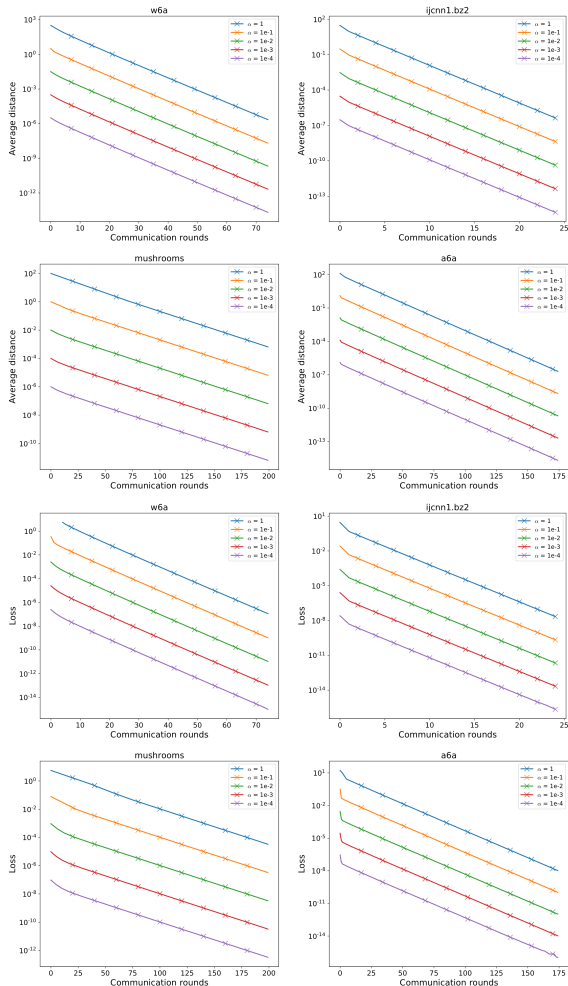


Figure 2: Squared averaged distance $1/n \sum_{i=1}^n \|x_i - x_i^*\|^2$ and loss $f(x) - f^*$ vs. # of communication rounds of Gradient Descent for logistic regression with l_2 regularizer. $\alpha_i = \alpha$ is set to the value indicated in the legend.

We use four datasets from LIBSVM (Chang and Lin, 2011) for (8): w6a, mushrooms, ijcnn1.bz2, a6a. We set all α_i 's to be equal. We divide data equally between all machines while preserving the order of data points such that i -th machine owns data with indices $[(i-1)r/n] + 1$ up to $[ir/n]$, where r is the total number of data points. We set the regularization parameter λ to 0.1. To find pure local models for each machine we run gradient descent with step size $1/L_i$, where L_i is f_i 's smoothness parameter until the norm of the

gradient is below 10^{-6} . We set the condition number to be $\kappa = \frac{1}{n} \sum_{i=1}^n L_i / \lambda$.

In this experiment, we investigate the convergence of gradient descent and look into the dependence between convergence and value of α . As expected, Figure 2 confirms that smaller values of α lead to better convergence as we rely more on local solutions and thus start closer to the optimal solution. Also note that the speed of the convergence appears to be constant among different values of α which is also predicted by our theory as the same α on each machine does not affect the conditioning of the global problem, see Proposition 2. We also use the DIANA algorithm (Mishchenko et al., 2019) with the random sparsification (also known as Rand- k compression), where set k coordinates to zero at random before communicating gradients to the server. Figure 3 shows the effect of varying k on the convergence of the method in terms of communication rounds. We consider 7 values of k linearly spaced between 1 and d . Similar to Figure 2, we observe that smaller values of α lead to better convergence. The rate at which the algorithm converges linearly is controlled by the compression constant $\omega + 1 = d/k$ or the effective conditioning $\kappa(\omega + 1)/n$. This is in line with the theory for DIANA (Mishchenko et al., 2019).

Table 1: Test accuracy comparisons of different methods after 3k rounds of training on EMNIST and Shakespeare datasets. Numbers next to FLIX model denote the value of $\alpha_i = \alpha$.

MODEL	EMNIST	SHAKESPEARE
FLIX, 0.1	0.862	0.5620
FLIX, 0.3	0.8691	0.5291
FLIX, 0.5	0.8634	0.5718
FLIX, 0.7	0.8652	0.5146
FLIX, 0.9	0.8397	0.5642
FedAvg	0.8663	0.5629

For completeness, we include convex experiments with unregularized logistic loss, extra experiments for all the combinations of 5 datasets and 3 algorithms – gradient descent, compressed gradient descent, and DIANA, as well as a detailed experimental description in the supplementary material (see Section 11).

Generalization experiment 1: Fitting Sine Functions with two-layer neural networks. Following Finn et al. (2017) and Zhou et al. (2019), we show the generalization advantages of the proposed method on the following regression problem. We define i -th client's function $f_i(x) = a_i \sin(x + b_i)$, where amplitude a_i and phase b_i lie in the intervals $[0.1, 0.5]$ and $[0, 2\pi]$, respectively. For each client, we fix a_i and b_i




Figure 3: Loss $f(x) - f^*$ vs. # of communication rounds of DIANA for logistic regression problem l_2 regularizer, k is a sparsification parameter of Random- k compressor.

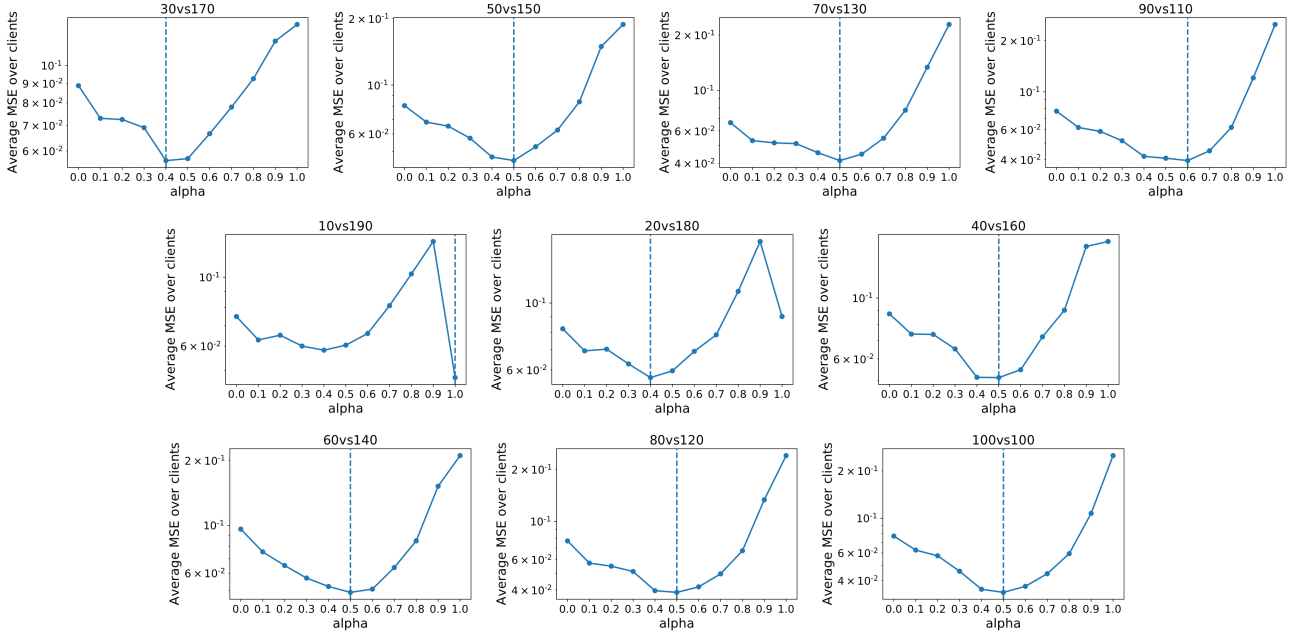


Figure 4: Average MSE vs. personalization parameter α .

and sample 50 points uniformly at random from the interval $[-5.0, 5.0]$. We measure regression fit in terms of mean squared error (MSE) loss. To train a model each client adopts a neural net with 2 hidden layers of size 40 with tanh activation. Further technical details are deferred to the appendix. For the experiment, we first sample 2 pairs $\{a_i, b_i\}$ and each of 200 clients is assigned one pair, we investigate different proportion–(30, 170), (50, 150), (70, 130), (90, 110). We then train our FLIX formulation with $\alpha_i = \alpha = 0.1, 0.2, \dots, 1$. For testing for each client generates a new dataset of size 2000. Figure 4 shows average MSE over clients against different values of α for different proportions. As this figure indicates, optimal α for which test average MSE is minimal can dramatically outperform the edge cases of either global model for all tasks or personalized model trained only on the local dataset. This underlines the main benefit of FLIX formulation,

which aims to find an optimal trade-off between local personalization and a single global model.

Generalization experiment 2: Comparison to FOMAML and Reptile with mutli-class logistic regression. Inspired by Reddi et al. (2021), we conduct a similar experiment to compare generalization capabilities, i.e., test accuracy, of FLIX and its two baselines FOMAML (Finn et al., 2017), and Reptile (Nichol et al., 2018) on the following tasks. For the first experiment (see Figure 1a), we take 500 train data points of two clients (with client ids ‘00000267’ and ‘00000459’) from the Stack Overflow dataset (TensorFlow Developers, 2021) and divide them among 100 workers so that there are 50 workers with 10 train data points from the first client and another 50 with 10 data points from the second client. For the second experiment (see Figure 1b), a worker gets 90 train data

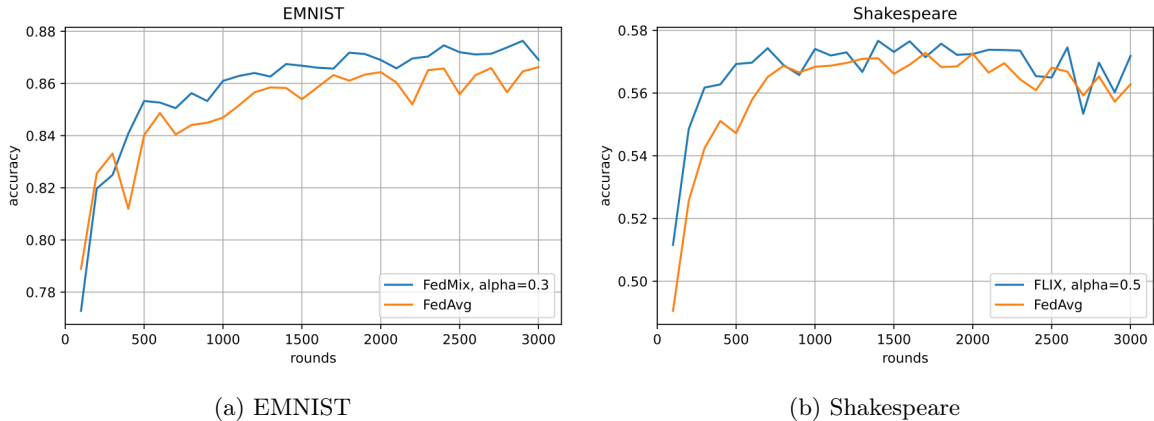


Figure 5: Generalization progress of FLIX vs. FedAvg on high-scale datasets.

points from a distinct client. For both experiments, each objective component f_i is a cross-entropy loss for multi-class logistic regression. Further technical details and the hyperparameters tuning for a fair comparison can be found in the supplementary. In the test phase, for each client, we used a hold-out testing dataset of size 300 (the same dataset has been used for workers related to the same client in the first experiment). It can be observed from Figure 1a, that for wide range of α_i , $\alpha_i \in \{0.2, 0.4, 0.6, 0.8\}$ FLIX exhibits a better generalization than its classical meta-learning competitors (FOMAML and Reptile), and it can lead to improvement of up to 11% in recall@5. Figure 1b shows that in the more real-world scenario FLIX outperforms FOMAML and Reptile while showing its best test accuracy in non-edge α . More experimental details and figures are included in the supplementary (Section 11).

Generalization experiment 3: Comparison to FedAvg with CNNs and LSTMs on large-scale datasets. For this section, we implement all the methods using the Python FedJAX library (Ro et al., 2020). We compare FLIX’s solution versus pure global obtained by FedAvg on two large-scale FL datasets: the EMNIST (Caldas et al., 2018) and Shakespeare (McMahan et al., 2017) datasets, solving a character prediction task on the first dataset and a next-character prediction task on the second dataset. For the first task, we adopt the EMNIST CNN model as described in (Reddi et al., 2021). To choose the hyperparameters for FLIX, we run grid search on a five-dimensional grid: pure local model (PLM) batch sizes [10, 20, 50, 100], pure local model learning rates [10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 1], FLIX batch sizes [20, 50, 100, 200], and FLIX learning rates [10^{-10} , 10^{-9} , \dots , 1]. We run SGD for 100 local epochs to compute PLMs for each client; with PLMs computed, we run Adam to compute FLIX global model. At each round of training in FLIX, we apply client subsampling and randomly choose ten clients to participate in the

training. For the final training, we choose the five hyperparameters showing the highest accuracy on a hold-out validation (different from the test) dataset (cut from train EMNIST dataset). For training FedAvg, we use the same hyperparameters and setting used by (Reddi et al., 2021). For the second task, we use the next-character prediction model described in (McMahan et al., 2017). We choose hyperparameters as follows: for the pure local model learning rates, we search over the grid [10^{-1} , $10^{-0.5}$, 10^0 , $10^{0.5}$, 10^1], for FLIX learning rates, we search over the grid [10^{-2} , $10^{-1.5}$, \dots , 10^2], for the batch sizes for PLM and FLIX we search over the grid [1, 4, 10, 20]. To compute PLMs, we run SGD for 25 epochs on each client locally. We then use SGD to compute the FLIX global model. At each round, only ten clients were chosen uniformly at random to participate in the training. We report test accuracy of the final models in Table 1 and show the generalization progress for the best alphas in Figure 5. As Table 1 shows the generalization capability of FLIX is superior to the purely global model obtained by FedAvg on both datasets.

5 ACKNOWLEDGMENTS

We would like to thank David Pugh for his support in configuring experiments on the KAUST HPC cluster.

References

Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. QSGD: communication-efficient SGD via gradient quantization and encoding. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing*

- Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 1709–1720, 2017. (Cited on page 6)
- Sebastian Caldas, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. LEAF: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097*, 2018. (Cited on page 9)
- Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, 2011. (Cited on page 7)
- Zachary Charles and Jakub Konečný. Convergence and accuracy trade-offs in federated learning and meta-learning. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 2575–2583. PMLR, 13–15 Apr 2021. (Cited on page 4)
- Shuxiao Chen, Qinqing Zheng, Qi Long, and Weijie J. Su. A theorem of the alternative for personalized federated learning. *arXiv preprint arXiv:2103.01901*, 2021. (Cited on page 3)
- Kurtland Chua, Qi Lei, and Jason D. Lee. How fine-tuning allows for effective meta-learning. *arXiv preprint arXiv:2105.02221*, 2021. (Cited on page 4)
- Liam Collins, Aryan Mokhtari, and Sanjay Shakkottai. How does the task landscape affect MAML performance? *arXiv preprint arXiv:2010.14672*, 2020. (Cited on page 4)
- Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. In John Shawe-Taylor, Richard S. Zemel, Peter L. Bartlett, Fernando C. N. Pereira, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*, pages 1647–1655, 2011. (Cited on page 6)
- Yuyang Deng, Mohammad Mahdi Kamani, and Mehrdad Mahdavi. Adaptive personalized federated learning. *arXiv preprint arXiv:2003.13461*, 2020. (Cited on pages 3 and 38)
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2019. (Cited on page 39)
- Canh T. Dinh, Tung T. Vu, Nguyen H. Tran, Minh N. Dao, and Hongyu Zhang. FedU: A unified framework for federated multi-task learning with laplacian regularization. *arXiv preprint arXiv:2102.07148*, 2021. (Cited on page 3)
- Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 3557–3568. Curran Associates, Inc., 2020a. URL <https://proceedings.neurips.cc/paper/2020/file/24389bfe4fe2eba8bf9aa9203a44cdad-Paper.pdf>. (Cited on page 3)
- Alireza Fallah, Aryan Mokhtari, and Asuman E. Ozdaglar. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. In Silvia Chiappa and Roberto Calandra, editors, *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pages 1082–1092. PMLR, 2020b. (Cited on page 5)
- Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Generalization of model-agnostic meta-learning algorithms: Recurring and unseen tasks. *arXiv preprint arXiv:2102.03832*, 2021. (Cited on page 4)
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR, 2017. (Cited on pages 4, 7, and 8)
- Katelyn Gao and Ozan Sener. Modeling and optimization trade-off in meta-learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. (Cited on page 4)
- Eduard Gorbunov, Filip Hanzely, and Peter Richtárik. A unified theory of sgd: Variance reduction, sampling, quantization and coordinate descent. *arXiv preprint arXiv:1905.11261*, 2019. (Cited on pages 23, 24, 25, 28, and 29)
- Neel Guha, Ameet Talwalkar, and Virginia Smith. One-shot federated learning. *arXiv preprint arXiv:1902.11175*, 2019. (Cited on page 5)
- Filip Hanzely and Peter Richtárik. Federated learning of a mixture of global and local models. *arXiv:2002.05516*, 2020. (Cited on page 3)

- Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. Personalized cross-silo federated learning on non-iid data. *arXiv preprint arXiv:2007.03797*, 2021. (Cited on page 3)
- Yihan Jiang, Jakub Konečný, Keith Rush, and Sreeram Kannan. Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*, 2019. (Cited on pages 3 and 4)
- Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, 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 preprint arXiv:1912.04977*, 2019. (Cited on pages 1, 2, and 3)
- Ahmed Khaled and Peter Richtárik. Better theory for SGD in the nonconvex world. *arXiv preprint arXiv:2002.03329*, 2020. (Cited on pages 33 and 34)
- Ahmed Khaled, Othmane Sebbouh, Nicolas Loizou, Robert M. Gower, and Peter Richtárik. Unified analysis of stochastic gradient methods for composite convex and smooth optimization. *arXiv preprint arXiv:2006.11573*, 2020. (Cited on pages 24, 31, and 33)
- Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. In *NIPS Private Multi-Party Machine Learning Workshop*, 2016. (Cited on pages 2 and 3)
- Dingwei Li, Qinglong Chang, Lixue Pang, Yanfang Zhang, Xudong Sun, Jikun Ding, and Liang Zhang. More industry-friendly: Federated learning with high efficient design. *arXiv preprint arXiv:2012.08809*, 2020. (Cited on page 2)
- Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *arXiv preprint arXiv:1908.07873*, 2019. (Cited on page 2)
- Zhize Li and Peter Richtárik. A unified analysis of stochastic gradient methods for nonconvex federated optimization. *arXiv preprint arXiv:2006.07013*, 2020. (Cited on page 35)
- Grigory Malinovskiy, Dmitry Kovalev, Elnur Gasanov, Laurent Condat, and Peter Richtárik. From local SGD to local fixed-point methods for federated learning. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 6692–6701. PMLR, 2020. (Cited on page 3)
- Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*, 2020. (Cited on pages 3 and 38)
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera 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. (Cited on page 9)
- Konstantin Mishchenko, Eduard Gorbunov, Martin Takác, and Peter Richtárik. Distributed learning with compressed gradient differences. *arXiv preprint arXiv:1901.09269*, 2019. (Cited on pages 6, 7, and 23)
- Yurii Nesterov. *Lectures on Convex Optimization*. Springer International Publishing, 2018. doi: 10.1007/978-3-319-91578-4. (Cited on pages 15 and 22)
- Lam M. Nguyen, Jie Liu, Katya Scheinberg, and Martin Takác. SARAH: A novel method for machine learning problems using stochastic recursive gradient. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 2613–2621. PMLR, 2017. (Cited on page 6)
- Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. (Cited on page 8)
- Reese Pathak and Martin J. Wainwright. FedSplit: an algorithmic framework for fast federated optimization. *arXiv preprint arXiv:2005.05238*, 2020. (Cited on page 3)
- Sashank J. Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Hugh Brendan McMahan. Adaptive federated optimization. In *International Conference*

- on Learning Representations (ICLR)*, 2021. (Cited on pages 8, 9, and 39)
- Amirhossein Reisizadeh, Farzan Farnia, Ramtin Pedarsani, and Ali Jadbabaie. Robust federated learning: the case of affine distribution shifts. *arXiv preprint arXiv:2006.08907*, 2020. (Cited on page 3)
- Jae Hun Ro, Ananda Theertha Suresh, and Ke Wu. Fed-JAX: Federated learning simulation with JAX, 2020. URL <http://github.com/google/fedjax>. (Cited on page 9)
- Saber Salehkaleybar, Arsalan Sharifnassab, and S. Jamaloddin Golestani. One-shot federated learning: Theoretical limits and algorithms to achieve them. *arXiv preprint arXiv:1905.04634*, 2019. (Cited on page 5)
- Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–13, 2020. doi: 10.1109/TNNLS.2020.3015958. (Cited on page 3)
- Canh T. Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with moreau envelopes. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21394–21405. Curran Associates, Inc., 2020. (Cited on page 3)
- TensorFlow Developers. TensorFlow Federated Stack Overflow dataset, 2021. URL https://www.tensorflow.org/federated/api_docs/python/tff/simulation/datasets/stackoverflow/load_data. (Cited on page 8)
- Kangkang Wang, Rajiv Mathews, Chloé Kiddon, Hubert Eichner, Françoise Beaufays, and Daniel Ramage. Federated evaluation of on-device personalization. *arXiv preprint arXiv:1910.10252*, 2019. (Cited on page 4)
- Blake Woodworth, Kumar Kshitij Patel, and Nathan Srebro. Minibatch vs local SGD for heterogeneous distributed learning. *arXiv preprint arXiv:2006.04735*, 2020. (Cited on page 3)
- Zheng Xu, Karan Singhal, Zachary Charles, Ziyu Liu, Advait Gadhikar, and Shanshan Wu. Federated optimization. *GitHub repository*, 2021. (Cited on page 39)
- Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. Salvaging federated learning by local adaptation. *arXiv preprint arXiv:2002.04758*, 2020. (Cited on pages 1, 2, and 3)
- Edvin Listo Zec, Olof Mogren, John Martinsson, Leon René Sütfeld, and Daniel Gillblad. Specialized federated learning using a mixture of experts. *arXiv preprint arXiv:2010.02056*, 2021. (Cited on pages 3 and 38)
- Pan Zhou, Xiaotong Yuan, Huan Xu, Shuicheng Yan, and Jiashi Feng. Efficient meta learning via mini-batch proximal update. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 1534–1544. Curran Associates, Inc., 2019. (Cited on page 7)

APPENDIX

Contents

1	INTRODUCTION	1
1.1	Key properties of the FLIX framework	1
1.2	Related work	3
2	THE FLIX FORMULATION	3
2.1	Motivation 1: from Local GD to FLIX	3
2.2	Motivation 2: from model-agnostic meta-learning to FLIX	4
3	THEORY AND ALGORITHMS	5
3.1	Distributed gradient descent	5
4	EXPERIMENTS	6
5	ACKNOWLEDGMENTS	9
6	BASIC FACTS AND NOTATION	15
7	PROOFS FOR SECTIONS 2 AND 3	16
7.1	Proof of Proposition 1	16
7.2	Proofs for algorithm-independent results	16
7.2.1	Proof of Proposition 2	16
7.2.2	A proposition for bounding the gradient norm	17
7.2.3	Proof of Proposition 3	18
7.3	Results on one shot averaging	18
7.4	Distributed Gradient Descent	22
8	OTHER ALGORITHMS	23
8.1	Strongly convex objectives	23
8.1.1	DCGD	23
8.1.2	DIANA	26
8.2	Convex objectives	30
8.2.1	DCGD	30
8.2.2	DIANA	33
8.3	Nonconvex objectives	33
8.3.1	DCGD	34
8.3.2	DIANA	35

9 DISCUSSION OF OTHER MODEL MIXTURE METHODS	38
10 EXPERIMENTAL DETAILS	39
11 EXTRA EXPERIMENTS	41

6 BASIC FACTS AND NOTATION

We will use the triangle inequality for norms:

$$\left\| \sum_{i=1}^n a_i \right\| \leq \sum_{i=1}^n \|a_i\|. \quad (9)$$

We say that a function g is L_g -smooth if for any $x, y \in \mathbb{R}^d$ we have

$$\|g(x) - g(y)\| \leq L_g \|x - y\|. \quad (10)$$

Note that (10) implies

$$g(x) \leq g(y) + \langle \nabla g(y), x - y \rangle + \frac{L_g}{2} \|x - y\|^2. \quad (11)$$

We say that a function g is μ_g -strongly convex for $\mu_g > 0$ if for all $x, y \in \mathbb{R}^d$ we have

$$g(x) \geq g(y) + \langle \nabla g(y), x - y \rangle + \frac{\mu_g}{2} \|x - y\|^2. \quad (12)$$

If (12) holds with $\mu_g = 0$, we say that g is convex.

We say $\mathcal{C} \in \mathbb{B}^d(\omega)$ is a compression operator if \mathcal{C} is unbiased (i.e., $\mathbb{E}[\mathcal{C}(x)] = x$ for all $x \in \mathbb{R}^d$) and if the second moment is bounded as

$$\mathbb{E}\|\mathcal{C}(x) - x\|^2 \leq \omega \|x\|^2 \quad \forall x \in \mathbb{R}^d. \quad (13)$$

Note that if $\mathcal{C} \in \mathbb{B}^d(\omega)$, then

$$\mathbb{E}\|\mathcal{C}(x)\|^2 \leq (1 + \omega)\|x\|^2, \quad \forall x \in \mathbb{R}^d. \quad (14)$$

Let f be a convex function. Then the Bregman divergence associated with f for points $x, y \in \mathbb{R}^d$ is defined in the following way:

$$D_f(x, y) = f(x) - f(y) - \langle \nabla f(y), x - y \rangle.$$

[Nesterov \(2018\)](#) shows that if function f is convex and L -smooth, then for all x and y

$$\|\nabla f(x) - \nabla f(y)\|^2 \leq 2LD_f(x, y). \quad (15)$$

Let $a, b \in \mathbb{R}^n$ be arbitrary vectors. Then, it holds that

$$\|a + b\|^2 \leq 2(\|a\|^2 + \|b\|^2). \quad (16)$$

7 PROOFS FOR SECTIONS 2 AND 3

This section collects the proofs of all propositions and theorems mentioned in the paper.

7.1 Proof of Proposition 1

In Section 2.2 we provided the motivation for FLIX through the lens of fine-tuning. We used the following proposition on fine-tuning quadratics:

Proposition 1. *Suppose that we run gradient descent for H steps on the quadratic objective $f_i = \frac{1}{2}x^T A_i x - b_i^T x + c$ starting from x^0 with stepsize $\gamma > 0$. Suppose that the stepsize satisfies $\gamma \leq \frac{1}{L_i}$, where $L_i = \lambda_{\max}(A_i)$, and suppose that A_i is positive definite. Then the final iterate x_i^H can be written as*

$$x_i^H = (I - J_i^H) x_i + J_i^H x^0, \quad (17)$$

where x_i minimizes f_i and $J_i \in \mathbb{R}^{d \times d}$ is a matrix with maximum eigenvalue smaller than 1, i.e. $\lambda_{\max}(J) < 1$.

Proof. The gradient descent update is

$$\begin{aligned} x_i^{t+1} &= x_i^t - \gamma \nabla f(x_i^t) \\ &= x_i^t - \gamma(A_i x_i^t - b_i) \\ &= (I - \gamma A_i) x_i^t + \gamma b_i. \end{aligned} \quad (18)$$

Note that because x_i minimizes f_i , we have $\nabla f_i(x_i) = 0$ by first-order optimality. Hence,

$$\nabla f_i(x_i) = 0 \iff A_i x_i = b_i.$$

Using this in (18),

$$x_i^{t+1} = (I - \gamma A_i) x_i^t + \gamma A_i x_i$$

Subtracting x_i from both sides,

$$\begin{aligned} x_i^{t+1} - x_i &= (I - \gamma A_i) x_i^t + (\gamma A_i - I) x_i \\ &= (I - \gamma A_i) (x_i^t - x_i). \end{aligned}$$

Iterating the above equality for H steps we get

$$x_i^H - x_i = (I - \gamma A_i)^H (x^0 - x_i).$$

Rearranging the terms we get (17) with $J_i \stackrel{\text{def}}{=} (I - \gamma A_i)$. Observe that when $\gamma \leq \frac{1}{\lambda_{\max}(A_i)}$ and $\lambda_{\min}(A_i) > 0$ we have that $\lambda_{\max}(I - \gamma A_i) < 1$. \square

7.2 Proofs for algorithm-independent results

7.2.1 Proof of Proposition 2

Proposition 2. *Suppose that each objective f_i is L_i -smooth. That is, for any $x, y \in \mathbb{R}^d$ we have*

$$\|\nabla f_i(x) - \nabla f_i(y)\| \leq L_i \|x - y\|. \quad (19)$$

Then the FLIX objective \tilde{f} defined in (1) is L_α -smooth for $L_\alpha \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i$. If each f_i is convex, then \tilde{f} is also convex. If each f_i is μ_i -strongly convex, then \tilde{f} is μ_α strongly convex for $\mu_\alpha \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 \mu_i$.

Proof. We separate the proofs in two cases: when each f_i is smooth, and when each f_i is (strongly) convex.

(i) Suppose that each f_i is L_i -smooth, and let $x, y \in \mathbb{R}^d$. Then by direct computation we have

$$\begin{aligned} \left\| \nabla \tilde{f}(x) - \nabla \tilde{f}(y) \right\| &= \left\| \frac{1}{n} \sum_{i=1}^n \alpha_i [\nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) - \nabla f_i(\alpha_i y + (1 - \alpha_i)y_i)] \right\| \\ &\stackrel{(9)}{\leq} \frac{1}{n} \sum_{i=1}^n \|\alpha_i\| \|\nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) - \nabla f_i(\alpha_i y + (1 - \alpha_i)y_i)\| \\ &\stackrel{(19)}{\leq} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i \|x - y\| = L_\alpha \|x - y\|, \end{aligned}$$

hence \tilde{f} is L_α -smooth.

(ii) Suppose that each f_i is μ_i -strongly convex for $\mu_i \geq 0$ (where $\mu_i = 0$ corresponds to just convexity). Then for $x, h \in \mathbb{R}^d$ we have

$$\begin{aligned} \tilde{f}(x+h) &= \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i(x+h) + (1 - \alpha_i)x_i) \\ &\stackrel{(12)}{\geq} \frac{1}{n} \sum_{i=1}^n \left[f_i(\alpha_i x + (1 - \alpha_i)x_i) + \langle \nabla f_i(\alpha_i x + (1 - \alpha_i)x_i), \alpha_i h \rangle + \frac{\mu_i}{2} \|\alpha_i h\|^2 \right] \\ &= \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i)x_i) + \left\langle \frac{1}{n} \sum_{i=1}^n \alpha_i \nabla f_i(\alpha_i x + (1 - \alpha_i)x_i), h \right\rangle \\ &\quad + \frac{1}{n} \sum_{i=1}^n \frac{\alpha_i^2 \mu_i}{2} \|h\|^2. \\ &= \tilde{f}(x) + \langle \nabla \tilde{f}(x), h \rangle + \frac{\mu_\alpha}{2} \|h\|^2, \end{aligned}$$

hence \tilde{f} is μ_α -strongly convex if all the μ_i are positive (resp. convex if they are equal to 0). □

7.2.2 A proposition for bounding the gradient norm

Proposition 2 provides us with a simple way to bound the gradient of \tilde{f} in terms of the distance to the optimum x^α :

Proposition 4. Define x^α as the solution to the FLIX problem (FLIX):

$$x^\alpha \stackrel{\text{def}}{=} \operatorname{argmin}_{x \in \mathbb{R}^d} \left[\tilde{f}(x) = \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i)x_i) \right].$$

Define L_α as in Proposition 2. Then for any $x \in \mathbb{R}^d$ we have

$$\left\| \nabla \tilde{f}(x) \right\| \leq L_\alpha \|x - x^\alpha\|, \tag{20}$$

and

$$\tilde{f}(x) - \tilde{f}(x^\alpha) \leq \frac{L_\alpha}{2} \|x - x^\alpha\|^2. \tag{21}$$

Proof. By Proposition 2, we have for any $x, y \in \mathbb{R}^d$ that

$$\left\| \nabla \tilde{f}(x) - \nabla \tilde{f}(y) \right\| \leq L_\alpha \|x - y\|.$$

Putting $y = x^\alpha$ and using that $\nabla \tilde{f}(x^\alpha) = 0$ we get (20). The L_α -smoothness of \tilde{f} implies

$$\tilde{f}(x) \leq \tilde{f}(y) + \langle \nabla \tilde{f}(y), x - y \rangle + \frac{L_\alpha}{2} \|x - y\|^2.$$

Putting $y = x^\alpha$ recovers (21). □

7.2.3 Proof of Proposition 3

Proposition 3. *Suppose that $\alpha_1 = \alpha_2 = \dots = \alpha_n = \beta$ in the FLIX formulation (FLIX). Let $T_1(x), T_2(x), \dots, T_n(x)$ be the deployed models defined in (3). If y_1, \dots, y_n are vectors in \mathbb{R}^d and \bar{y} is their mean, we define $V(y_1, \dots, y_n)$ as the population variance $V(y_1, \dots, y_n) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \|y_i - \bar{y}\|^2$. Then,*

$$V(T_1(x), T_2(x), \dots, T_n(x)) = (1 - \beta)^2 V(x_1, x_2, \dots, x_n).$$

Proof. By direct computation observe

$$\begin{aligned} V(T_1(x), T_2(x), \dots, T_n(x)) &= \frac{1}{n} \sum_{i=1}^n \left\| T_i(x) - \frac{1}{n} \sum_{j=1}^n T_j(x) \right\|^2 \\ &= \frac{1}{n} \sum_{i=1}^n \left\| \beta x + (1 - \beta)x_i - \frac{1}{n} \sum_{j=1}^n (\beta x + (1 - \beta)x_j) \right\|^2 \\ &= \frac{1}{n} \sum_{i=1}^n \left\| (1 - \beta) \left(x_i - \frac{1}{n} \sum_{j=1}^n x_j \right) \right\|^2 \\ &= (1 - \beta)^2 V(x_1, x_2, \dots, x_n). \end{aligned}$$

□

7.3 Results on one shot averaging

Before proving Theorems 1 and 3, we will need the following lemma which shows that the gradient of the FLIX and the functional suboptimality can be bound using a weighted average of the iterate norms:

Lemma 1. *Suppose that each f_i is L_i -smooth and that each x_i is a stationary point of f_i . Then for any $x \in \mathbb{R}^d$ we have,*

$$\left\| \nabla \tilde{f}(x) \right\| \leq \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i \|x - x_i\|, \quad (22)$$

and,

$$\left\| \nabla \tilde{f}(x) \right\|^2 \leq \frac{1}{n} \sum_{i=1}^n \alpha_i^4 L_i^2 \|x - x_i\|^2, \quad (23)$$

and,

$$\tilde{f}(x) \leq \frac{1}{n} \sum_{i=1}^n f_i(x_i) + \frac{1}{2n} \sum_{i=1}^n \alpha_i^2 L_i \|x - x_i\|^2. \quad (24)$$

Proof. For the first inequality, using the fact that $\nabla f_i(x_i) = 0$ for all i , and applying the triangle inequality (9) and the L_i -smoothness of f_i , we get

$$\begin{aligned} \left\| \nabla \tilde{f}(x) \right\| &= \left\| \frac{1}{n} \sum_{i=1}^n \alpha_i \nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) \right\| \\ &= \left\| \frac{1}{n} \sum_{i=1}^n \alpha_i [\nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) - \nabla f_i(x_i)] \right\| \\ &\stackrel{(9)}{\leq} \frac{1}{n} \sum_{i=1}^n |\alpha_i| \left\| \nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) - \nabla f_i(x_i) \right\| \\ &\stackrel{(10)}{\leq} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i \|x - x_i\|. \end{aligned}$$

For the second inequality, we use that $\nabla f_i(x_i) = 0$, the convexity of the squared norm, and smoothness:

$$\begin{aligned}
 \|\nabla \tilde{f}(x)\|^2 &= \left\| \frac{1}{n} \sum_{i=1}^n \alpha_i \nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) \right\|^2 \\
 &= \left\| \frac{1}{n} \sum_{i=1}^n \alpha_i [\nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) - \nabla f_i(x_i)] \right\|^2 \\
 &\leq \frac{1}{n} \sum_{i=1}^n \alpha_i^2 \|\nabla f_i(\alpha_i x + (1 - \alpha_i)x_i) - \nabla f_i(x_i)\|^2 \\
 &\stackrel{(10)}{\leq} \frac{1}{n} \sum_{i=1}^n \alpha_i^4 L_i^2 \|x - x_i\|^2.
 \end{aligned}$$

Using L_i -smoothness of f_i and that $\nabla f_i(x_i) = 0$, we get

$$\begin{aligned}
 f_i(\alpha_i x + (1 - \alpha_i)x_i) &= f_i(x_i + \alpha_i(x - x_i)) \\
 &\stackrel{(11)}{\leq} f_i(x_i) + \langle \nabla f_i(x_i), \alpha_i(x - x_i) \rangle + \frac{\alpha_i^2 L_i}{2} \|x - x_i\|^2 \\
 &= f_i(x_i) + \frac{\alpha_i^2 L_i}{2} \|x - x_i\|^2.
 \end{aligned}$$

Averaging the above inequality yields (24). \square

As communicated earlier in the paper, we can solve (FLIX) by taking a weighted average of the pure local models x_i if the α_i are small enough. This is a consequence of preprocessing step where we compute x_i , which is communication-free, and also a consequence of the new formulation, where for small α the FLIX objective is less responsive to argument change (see Proposition 2).

Theorem 1. *Suppose that each objective f_i is L_i -smooth, that x_i minimizes f_i , and let $\hat{L} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n L_i$. Given the pure local models x_1, x_2, \dots, x_n , define the weighted average*

$$x^{\text{avg}} \stackrel{\text{def}}{=} \sum_{i=1}^n w_i x_i, \quad w_i \stackrel{\text{def}}{=} \frac{\alpha_i^2 L_i}{n L_\alpha}, \quad L_\alpha \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i. \quad (25)$$

We further define the constants

$$D \stackrel{\text{def}}{=} \max_{i,j=1,\dots,n,i \neq j} \|x_i - x_j\|^2, \quad \text{and} \quad V \stackrel{\text{def}}{=} \sum_{i=1}^n w_i \|x_i - x^{\text{avg}}\|^2 \quad (26)$$

Fix any $\epsilon > 0$. Assume that either $\max_{i=1,\dots,n} \alpha_i \leq \sqrt{2\epsilon}/\sqrt{\hat{L}D}$, or $\alpha_i = \beta$ for all i and $\beta \leq \sqrt{2\epsilon}/\sqrt{\hat{L}D}$. Then x^{avg} is an ϵ -approximate minimizer of (FLIX). That is,

$$\tilde{f}(x^{\text{avg}}; \alpha_1, \dots, \alpha_n, x_1, \dots, x_n) - \min_{x \in \mathbb{R}^d} \tilde{f}(x; \alpha_1, \dots, \alpha_n, x_1, \dots, x_n) \leq \epsilon.$$

Proof. Define

$$\mathcal{E}_\alpha(x) \stackrel{\text{def}}{=} \frac{1}{2n} \sum_{i=1}^n \alpha_i^2 L_i \|x - x_i\|^2. \quad (27)$$

Note that since the vectors $\{x_i\}$ are known, the expression $\mathcal{E}_\alpha(x)$ can be minimized in x , leading to a weighted average of the pure local models:

$$x^{\text{avg}} \stackrel{\text{def}}{=} \sum_{i=1}^n w_i x_i, \quad w_i \stackrel{\text{def}}{=} \frac{\alpha_i^2 L_i}{n L_\alpha}, \quad L_\alpha \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i. \quad (28)$$

Note that x^{avg} , can be computed using a single communication round. By plugging (28) into (27), we can evaluate the error of the average model:

$$\mathcal{E}(x^{avg}) = \frac{L_\alpha}{2} \sum_{i=1}^n w_i \|x^{avg} - x_i\|^2 = \frac{L_\alpha V_\alpha}{2}, \quad (29)$$

where

$$V_\alpha \stackrel{\text{def}}{=} \sum_{i=1}^n w_i \|x^{avg} - x_i\|^2 = \mathbb{E} \|x_i - \mathbb{E}x_i\|^2 \quad (30)$$

can be interpreted as the *variance of the local optimal models* $\{x_i\}$. We now give two ways how $\mathcal{E}(x^{avg})$ can be bounded:

- Assume the customization parameters $\{\alpha_i\}$ are allowed to be arbitrary. Since x^{avg} is in the convex hull of the set $\{x_1, \dots, x_n\}$, we have $\|x^{avg} - x_i\| \leq \max_{i,j} \|x_i - x_j\|$. Therefore,

$$V_\alpha \leq D \stackrel{\text{def}}{=} \max_{i,j} \|x_i - x_j\|^2. \quad (31)$$

Moreover,

$$L_\alpha = \frac{1}{n} \sum_{i=1}^n \alpha_i^2 L_i \leq \alpha_{\max}^2 \bar{L}, \quad (32)$$

where $\alpha_{\max} \stackrel{\text{def}}{=} \max_i \alpha_i$ and $\bar{L} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n L_i$. By plugging (31) and (32) into (29), we get

$$\mathcal{E}(x^{avg}) \leq \frac{\alpha_{\max}^2 \bar{L} D}{2}. \quad (33)$$

- Assume the customization parameters $\{\alpha_i\}$ are all equal: $\alpha_i = \beta$ for all i . Then $w_i = \frac{L_i}{\sum_j L_j}$, and hence x^{avg} and $V_\alpha = V$ are independent of β . Since $L_\alpha = \beta^2 \bar{L}$, by plugging these expressions into (29), we get

$$\mathcal{E}(x^{avg}) \leq \frac{\beta^2 \bar{L} V}{2}. \quad (34)$$

Note that $V \leq D$.

Equations (33) and (34) yield the following observation: fix any $\epsilon > 0$ and assume one of the following conditions holds: (1) The maximum personalization parameter satisfies

$$\alpha_{\max} \leq \sqrt{\frac{2\epsilon}{\bar{L}D}}. \quad (35)$$

Or, (2) All the personalization parameters are equal, i.e., $\alpha_i = \beta$ for all i , and

$$\beta \leq \sqrt{\frac{2\epsilon}{\bar{L}V}}. \quad (36)$$

Then

$$\mathcal{E}(x^{avg}) \leq \epsilon. \quad (37)$$

Then combining (24) from Lemma 1 with (37) we get that the weighted average of the local optimal models x^{avg} satisfies

$$\tilde{f}(x^{avg}) \leq \frac{1}{n} \sum_{i=1}^n f_i(x_i) + \epsilon \leq \frac{1}{n} \sum_{i=1}^n f_i(x^\alpha) + \epsilon = \tilde{f}(x^\alpha) + \epsilon,$$

where $x^\alpha = \operatorname{argmin}_{x \in \mathbb{R}^d} \tilde{f}(x)$ and where we used that x_i minimizes f_i . This shows the second part of this Theorem's claim. \square

The next theorem extends the previous result to stationary points instead of minimizers.

Theorem 3. *Suppose that each objective f_i is L_i -smooth and that x_i is a stationary point of f_i , and let $\hat{L} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n L_i$. Given the pure local models x_1, x_2, \dots, x_n , define the weighted average*

$$x^{\text{avg}} \stackrel{\text{def}}{=} \sum_{i=1}^n w_i x_i, \quad w_i \stackrel{\text{def}}{=} \frac{\alpha_i^4 L_i^2}{n \hat{L}_\alpha}, \quad \hat{L}_\alpha^2 \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \alpha_i^4 L_i^2. \quad (38)$$

Fix any $\epsilon > 0$ and define D and V as in Theorem 1. Assume that either $\max_{i=1, \dots, n} \alpha_i \leq \sqrt{2\epsilon}/\sqrt{\hat{L}D}$, or $\alpha_i = \beta$ for all i and $\beta \leq \sqrt{2\epsilon}/\sqrt{\hat{L}D}$. Then x^{avg} is a stationary point of (FLIX). That is,

$$\|\nabla \tilde{f}(x^{\text{avg}})\| \leq \epsilon.$$

Proof. We will proceed very similarly to Theorem 1. Let

$$\zeta_\alpha(x) = \frac{1}{n} \sum_{i=1}^n \alpha_i^4 L_i^2 \|x - x_i\|^2.$$

It is easy to show that (38) minimizes $\zeta_\alpha(x)$, giving

$$\zeta_\alpha(x^{\text{avg}}) = \frac{\hat{L}_\alpha^2 V_\alpha}{2}, \quad (39)$$

where

$$V_\alpha \stackrel{\text{def}}{=} \sum_{i=1}^n w_i \|x^{\text{avg}} - x_i\|^2.$$

- Assume the customization parameters $\{\alpha_i\}$ are allowed to be arbitrary. Since x^{avg} is in the convex hull of the set $\{x_1, \dots, x_n\}$, we have $\|x^{\text{avg}} - x_i\| \leq \max_{i,j} \|x_i - x_j\|$. Therefore,

$$V_\alpha \leq D \stackrel{\text{def}}{=} \max_{i,j} \|x_i - x_j\|^2. \quad (40)$$

Moreover,

$$\hat{L}_\alpha^2 = \frac{1}{n} \sum_{i=1}^n \alpha_i^4 L_i^2 \leq \alpha_{\max}^4 \hat{L}^2, \quad (41)$$

where $\alpha_{\max} \stackrel{\text{def}}{=} \max_i \alpha_i$ and $\hat{L}^2 \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n L_i^2$. By plugging (40) and (41) into (39), we get

$$\zeta_\alpha(x^{\text{avg}}) \leq \alpha_{\max}^4 \hat{L}^2 D. \quad (42)$$

- Assume the customization parameters $\{\alpha_i\}$ are all equal: $\alpha_i = \beta$ for all i . Then $w_i = \frac{L_i^2}{\sum_j L_j^2}$, and hence x^{avg} and $V_\alpha = V$ are independent of β . Since $L_\alpha^2 = \beta^4 \hat{L}^2$, by plugging these expressions into (39), we get

$$\zeta_\alpha(x^{\text{avg}}) \leq \beta^4 \hat{L}^2 V. \quad (43)$$

Combining (42) and (43) we get that if (1) $\alpha_{\max} \leq \left(\frac{\epsilon^2}{\hat{L}^2 D}\right)^{1/4}$ or (2) $\alpha_i = \beta$ and $\beta \leq \left(\frac{\epsilon^2}{\hat{L}^2 V}\right)^{1/4}$ then

$$\zeta_\alpha(x^{\text{avg}}) \leq \epsilon^2.$$

It remains to use (23) from Lemma 1 and then take square roots. \square

Algorithm 1 Distributed gradient descent for FLIX.

Require: Number of communication rounds K , stepsize γ , initial point x^0

- 1: **for** $k = 0, 1, \dots, K - 1$ **do**
- 2: **for** $i = 1, \dots, n$ in parallel **do**
- 3: Compute $\nabla \tilde{f}_i(x^k) = \alpha_i \nabla f_i(\alpha_i x^k + (1 - \alpha_i) x_i)$ and communicate it to server.
- 4: Average and broadcast the new iterate

$$x^{k+1} = x^k - \frac{\gamma}{n} \sum_{i=1}^n \nabla \tilde{f}_i(x^k).$$

7.4 Distributed Gradient Descent

Gradient descent is a simple but yet informative way to solve many optimization problems. Here we provide the proof for the convergence of DGD.

Theorem 2. *Suppose that each f_i in (FLIX) is L_i -smooth and μ_i -strongly convex. Define x^{avg} , L_α , and \hat{L} by (5) and V, D by (6). Suppose that we run (DGD) for K iterations starting from $x^0 = x^{\text{avg}}$. Then the following hold:*

i) *If the α_i are allowed to be arbitrary, then for $\alpha_{\max} \stackrel{\text{def}}{=} \max_{i=1, \dots, n} \alpha_i$ we have*

$$\tilde{f}(x^k) - \min_{x \in \mathbb{R}^d} \tilde{f}(x) \leq \left(1 - \frac{\mu_\alpha}{L_\alpha}\right)^k \frac{\alpha_{\max}^2 \hat{L} D}{2}.$$

ii) *If $\alpha_i = \beta$ for all i , then*

$$\tilde{f}(x^k) - \min_{x \in \mathbb{R}^d} \tilde{f}(x) \leq \left(1 - \frac{\hat{\mu}}{\hat{L}}\right)^k \frac{\beta^2 \hat{L} V}{2},$$

where $\hat{\mu} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \mu_i$.

Proof. Recall the standard result that gradient descent for an L -smooth and μ -strongly convex objective g satisfies for any initial point x^0

$$g(x^k) - g_* \leq (1 - \gamma\mu)^k (g(x_0) - g_*), \quad (44)$$

where $g_* = \min_{x \in \mathbb{R}^d} g(x)$. For a proof, see (Nesterov, 2018). Note that by Proposition 2 we have that \tilde{f} is L_α -smooth and μ_α -strongly convex. Specializing (44) to this case yields

$$\tilde{f}(x^k) - \min_{x \in \mathbb{R}^d} \tilde{f}(x) \leq \left(1 - \frac{\mu_\alpha}{L_\alpha}\right)^k \left(\tilde{f}(x_0) - \min_{x \in \mathbb{R}^d} \tilde{f}(x)\right) \quad (45)$$

Note that our initialization is the same x^{avg} from Theorem 1. Plugging the bounds of that theorem into (45) yields the theorem's claims. \square

Table 2: Convergence results for Distributed Compressed Gradient Descent and DIANA in different settings. All the constants are independent of α .

ALGORITHM	ASSUMPTION	CONVERGENCE GUARANTEE	COROLLARY
DCGD	smoothness, str cvx	$\mathbb{E}\ x^k - x^*\ ^2 \leq (1 - \gamma_0\bar{\mu})^k C_1 + C_2$ $\mathbb{E}[\tilde{f}(x^k) - \tilde{f}^*] \leq ((1 - \gamma_0\bar{\mu})^k C_1 + C_2)\alpha^2$	1
DIANA	smoothness, str cvx	$\mathbb{E}\ x^k - x^*\ ^2 \leq (1 - \rho)^k C$ $\mathbb{E}[\tilde{f}(x^k) - \tilde{f}^*] \leq (1 - \rho)^k C\alpha^2$	2
DCGD	smoothness, cvx	$\mathbb{E}[\tilde{f}(\bar{x}^k) - \tilde{f}(x^*)] \leq \frac{1}{k}C_1\alpha^2 + C_2\alpha$	3
DIANA	smoothness, cvx	$\mathbb{E}[\tilde{f}(\bar{x}^k) - \tilde{f}(x^*)] \leq \frac{1}{k}(C_1\alpha^2 + C_2\alpha)$	4
DCGD	smoothness	$\min_{0 \leq t \leq k-1} \mathbb{E}\ \nabla \tilde{f}(x^t)\ ^2 \leq \frac{(1+C_1\gamma_0^2)^k C_2}{\gamma_0^k} \alpha^2$	5
DIANA	smoothness	$\mathbb{E}\ \nabla \tilde{f}(\hat{x})\ ^2 \leq \frac{C}{k} \alpha^2$	6

8 OTHER ALGORITHMS

In this section, we include additional algorithms to solve (FLIX) formulation, namely, DCGD (Gorbunov et al., 2019) and DIANA (Mishchenko et al., 2019). In Table 2, we display convergence guarantees for these algorithms. One can see that similarly to our results in Section 3 both of the algorithms requires less iterations in terms of convergence in both local deploy iterates $T_i(x^k)$'s and functional value $\tilde{f}(x^k)$. In addition, for the standard (ERM) problem, i.e., $\alpha = 1$, we recover the best known convergence guarantees. Below, we provide a derivation of these claims.

8.1 Strongly convex objectives

8.1.1 DCGD

We firstly introduce the DCGD algorithm followed by general convergence results that we later exploit to obtain a convergence guarantee for DCGD applied to (FLIX).

Algorithm 2 Distributed Compressed Gradient Descent with different noise levels ω_i

Require: $x^0 \in \mathbb{R}^d$, learning rate γ

- 1: **for** $k = 0, 1, 2, \dots$ **do**
 - 2: Broadcast x^k to all workers
 - 3: **for** $i = 1, \dots, n$ in parallel **do**
 - 4: Evaluate $\nabla f_i(x^k)$
 - 5: $g_i^k = \mathcal{C}_i(\nabla f_i(x^k))$
 - 6: $g^k = \frac{1}{n} \sum_{i=1}^n g_i^k$
 - 7: $x^{k+1} = x^k - \gamma g^k$
-

Lemma 2. Suppose each f_i is L_i -smooth and convex, and f is L -smooth. Let $\mathcal{C}_i : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be randomized compression operators satisfying $\mathcal{C}_i \in \mathbb{B}^d(\omega_i)$. Let $g_k = \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla f_i(x^k))$. Then

$$\mathbb{E}\|g^k - \nabla f(x^*)\|^2 \leq 2 \left(L + \frac{2 \max\{L_i \omega_i\}}{n} \right) D_f(x^k, x^*) + \sigma_{DCGD}, \quad (46)$$

where $\sigma_{DCGD} = \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^*)\|^2$.

Proof. By applying bias-variance decomposition to $\mathbb{E}\|g^k - \nabla f(x^*)\|^2$, we get

$$\mathbb{E}\|g^k - \nabla f(x^*)\|^2 = \|\nabla f(x^k) - \nabla f(x^*)\|^2 + \mathbb{E}\|g^k - \nabla f(x^k)\|^2.$$

Since all functions are convex holds, function f as a linear combination of convex functions is also convex. That is why the first term enjoys the 'classic' bound for *convex* and smooth functions expressed in Bregman divergence between x^k and x^* (see inequality 15):

$$\|\nabla f(x^k) - \nabla f(x^*)\|^2 \leq 2LD_f(x^k, x^*). \quad (47)$$

We start looking into the second term by expanding the quadratic:

$$\begin{aligned} & \mathbb{E}\|g^k - \nabla f(x^k)\|^2 \\ &= \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla f_i(x^k)) - \frac{1}{n} \sum_{i=1}^n \nabla f_i(x^k) \right\|^2 \\ &= \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^n (\mathcal{C}_i(\nabla f_i(x^k)) - \nabla f_i(x^k)) \right\|^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}\|\mathcal{C}_i(\nabla f_i(x^k)) - \nabla f_i(x^k)\|^2 + \\ & \quad \frac{1}{n^2} \sum_{i \neq j} \mathbb{E}\langle \mathcal{C}_i(\nabla f_i(x^k)) - \nabla f_i(x^k), \mathcal{C}_j(\nabla f_j(x^k)) - \nabla f_j(x^k) \rangle. \end{aligned}$$

Since $\mathcal{C}_i \in \mathbb{B}(\omega_i) \forall i$ is drawn independently, the expected value of each scalar product in the second sum becomes the scalar product of expected values, each of which is zero due to unbiasedness of a compression operator (see Section 6):

$$\begin{aligned} & \mathbb{E}\langle \mathcal{C}_i(\nabla f_i(x^k)) - \nabla f_i(x^k), \mathcal{C}_j(\nabla f_j(x^k)) - \nabla f_j(x^k) \rangle \\ &= \underbrace{\langle \mathbb{E}\mathcal{C}_i(\nabla f_i(x^k)) - \nabla f_i(x^k), \mathbb{E}\mathcal{C}_j(\nabla f_j(x^k)) - \nabla f_j(x^k) \rangle}_{=0} = 0. \end{aligned}$$

Then, we apply 13 (the second property of compressed operators) to get the final upper-bound on $\mathbb{E}\|g^k - \nabla f(x^k)\|^2$.

$$\mathbb{E}\|g^k - \nabla f(x^k)\|^2 = \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}\|\mathcal{C}_i(\nabla f_i(x^k)) - \nabla f_i(x^k)\|^2 \leq \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k)\|^2.$$

Curiously but not surprisingly, the result resembles the Law of large numbers (indeed, if each squared gradient norm and noise level are bounded above by values R and ω_0 respectively, then the right side is bounded by $\frac{\omega_0 R}{n}$, which converges to zero as n converges to infinity), what partially verifies the correctness of the proof. The error $\mathbb{E}\|g^k - \nabla f(x^k)\|^2$ is now estimated by squared gradient norms at iterate x^k , which is, in general, a random point generated by CGD. Since we already have some dependence on Bregman divergence in 47, which perfectly fits assumptions for unified theories from papers [Gorbunov et al. \(2019\)](#) and [Khaled et al. \(2020\)](#) required in later theorems, we express the term in the right side of the last inequality in Bregman divergence, too. By subtracting and adding the same vector $\nabla f_i(x^*)$ inside the norm operator we get

$$\begin{aligned} & \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k)\|^2 = \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - \nabla f_i(x^*) + \nabla f_i(x^*)\|^2 \\ & \leq \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - \nabla f_i(x^*)\|^2 + \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^*)\|^2, \end{aligned}$$

where in the second line we used (16). Applying again the bound from 15 to the first term and using the linearity of Bregman divergence we have

$$\begin{aligned} \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - \nabla f_i(x^*)\|^2 &\leq \frac{2}{n^2} \sum_{i=1}^n \omega_i 2L_i D_{f_i}(x^k, x^*) \\ &\leq \frac{4 \max\{L_i \omega_i\}}{n^2} \sum_{i=1}^n D_{f_i}(x^k, x^*) = \frac{4 \max\{L_i \omega_i\}}{n} D_f(x^k, x^*). \end{aligned}$$

Incorporation of the results above gives the statement. \square

Proposition 5. *Suppose each f_i is L_i -smooth and μ_i -strong. Assume $\alpha_i \equiv \alpha \in \mathbb{R}$ for all i . For any $x^0 \in \mathbb{R}^d$, it holds that*

$$\|x^0 - x^*\|^2 \leq \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2}. \quad (48)$$

Proof.

$$\begin{aligned} \|x^0 - x^*\|^2 &\stackrel{(12)}{\leq} \frac{2}{\mu_\alpha} (\tilde{f}(x^0) - \tilde{f}(x^*)) \stackrel{\text{Proposition 2}}{=} \frac{2}{\bar{\mu} \alpha^2} (\tilde{f}(x^0) - \tilde{f}(x^*)) \stackrel{(24)}{\leq} \frac{2}{\bar{\mu} \alpha^2} \mathcal{E}_\alpha(x^0) \\ &\stackrel{(27)}{=} \frac{1}{\bar{\mu} \alpha^2} \frac{1}{n} \sum_{i=1}^n L_i \alpha^2 \|x^0 - x_i\|^2 = \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2}. \end{aligned}$$

\square

Proposition 6. *Suppose each f_i is L_i -smooth and μ_i -strong. Assume $\alpha_i \equiv \alpha \in \mathbb{R}$ for all i . Then,*

$$\overline{\|x_i - x^*\|^2} \leq \frac{\max_i L_i}{\bar{\mu}} \max_{i,j} \|x_i - x_j\|^2. \quad (49)$$

Proof. According to Proposition 5, for any x_l^* where $l \in \{1, \dots, n\}$ it follows that

$$\|x_l - x^*\|^2 \leq \frac{1}{\bar{\mu}} \overline{L_i \|x_l - x_i\|^2} \leq \frac{1}{\bar{\mu}} \max_i L_i \cdot \overline{\|x_l - x_i\|^2} \leq \frac{1}{\bar{\mu}} \max_i L_i \cdot \max_{i,j} \|x_j - x_i\|^2.$$

Taking the average of both sides finishes the proof. \square

Theorem 4. *Assume all conditions of Lemma 2 hold, and each function f_i is μ_i -strongly convex. Then, if $\gamma \leq \frac{1}{L_\alpha + \frac{2 \max\{L_i \alpha_i^2 \omega_i\}}{n}}$, then*

$$\mathbb{E} \|x^k - x^*\|^2 \leq (1 - \gamma \mu_\alpha)^k \|x^0 - x^*\|^2 + \frac{2\gamma}{\mu_\alpha n^2} \sum_{i=1}^n \omega_i \|\nabla [f_i(T_i(x^*))]\|^2. \quad (50)$$

Proof. The lemma is a direct corollary of Theorem 4.1 from Gorbunov et al. (2019) and Proposition 2 with constants $A = L_\alpha + \frac{2 \max\{L_i \alpha_i^2 \omega_i\}}{n}$, $D_1 = \sigma_{DCGD} = \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla [f_i(T_i(x^*))]\|^2$, $B = 0, \sigma_k^2 \equiv 0, \rho = 1, C = 0, D_2 = 0$. \square

Corollary 1. *Assume all conditions of Theorem 4 hold, and $\alpha_i \equiv \alpha \forall i$. If $\gamma = \frac{1}{\alpha^2 (\bar{L} + \frac{2 \max_i L_i}{n})} = \gamma_0 \frac{1}{\alpha^2}$, then*

$$\mathbb{E} \|x^k - x^*\|^2 \leq (1 - \gamma_0 \bar{\mu})^k \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2} + \frac{2\gamma_0 \omega (\max_i L_i)^2 \max_{i,j} \|x_i - x_j\|^2}{\bar{\mu}^2 n} \quad (51)$$

and

$$\mathbb{E} [\tilde{f}(x^k) - \tilde{f}^*] \leq ((1 - \gamma_0 \bar{\mu})^k C_1 + C_2) \alpha^2, \quad (52)$$

where $C_1 = \frac{\bar{L}_i}{2\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2}$ and $C_2 = \frac{\bar{L}_i \gamma_0 \omega (\max_i L_i)^2 \max_{i,j} \|x_i - x_j\|^2}{\bar{\mu}^2 n}$.

Proof. Notice that

$$\begin{aligned}
 & \sum_{i=1}^n \omega_i \|\nabla[f_i(\alpha_i x^* + (1 - \alpha_i)x_i)]\|^2 = \sum_{i=1}^n \omega_i \|\alpha_i \nabla f_i(\alpha_i x^* + (1 - \alpha_i)x_i)\|^2 \\
 & = \omega \alpha^2 \sum_{i=1}^n \|\nabla f_i(\alpha x^* + (1 - \alpha)x_i)\|^2 \stackrel{\nabla f_i(x_i)=0}{=} \omega \alpha^2 \sum_{i=1}^n \|\nabla f_i(\alpha x^* + (1 - \alpha)x_i) - \nabla f_i(x_i)\|^2 \\
 & \stackrel{(10)}{\leq} \omega \alpha^2 \sum_{i=1}^n L_i \|\alpha(x^* - x_i)\|^2 \leq \omega \alpha^4 \max_i L_i n \|x^* - x_i\|^2 \stackrel{(49)}{\leq} \frac{\omega \alpha^4 n (\max_i L_i)^2}{\bar{\mu}} \max_{i,j} \|x_i - x_j\|^2.
 \end{aligned}$$

This brings us to the following bound on the neighbourhood:

$$\begin{aligned}
 & \frac{2\gamma}{\mu_\alpha n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^*))]\|^2 = \frac{2\gamma_0}{\alpha^4 \bar{\mu} n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^*))]\|^2 \\
 & \leq \frac{2\gamma_0 \omega (\max_i L_i)^2 \max_{i,j} \|x_i - x_j\|^2}{\bar{\mu}^2 n}.
 \end{aligned} \tag{53}$$

We further investigate the rate of convergence

$$1 - \gamma \mu_\alpha \stackrel{\text{Proposition 2}}{=} 1 - \gamma_0 \frac{1}{\alpha^2} \bar{\mu} \alpha^2 = 1 - \gamma_0 \bar{\mu}. \tag{54}$$

As can be seen, the rate of convergence does not depend on α .

Combining last two results we obtain the following dependence on α for CGD convergence:

$$\begin{aligned}
 & \mathbb{E}\|x^k - x^*\|^2 \\
 & \stackrel{(50)}{\leq} (1 - \gamma \mu_\alpha)^k \|x^0 - x^*\|^2 + \frac{2\gamma}{\mu_\alpha n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^*))]\|^2 \\
 & \stackrel{(54)}{=} (1 - \gamma_0 \bar{\mu})^k \|x^0 - x^*\|^2 + \frac{2\gamma}{\mu_\alpha n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^*))]\|^2 \\
 & \stackrel{(48)}{\leq} (1 - \gamma_0 \bar{\mu})^k \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2} + \frac{2\gamma}{\mu_\alpha n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^*))]\|^2 \\
 & \stackrel{(53)}{\leq} (1 - \gamma_0 \bar{\mu})^k \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2} + \frac{2\gamma_0 \omega (\max_i L_i)^2 \max_{i,j} \|x_i - x_j\|^2}{\bar{\mu}^2 n} \\
 & = (1 - \gamma_0 \bar{\mu})^k \tilde{C}_1 + \tilde{C}_2,
 \end{aligned} \tag{55}$$

where constants $\tilde{C}_1 = \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2}$ and $\tilde{C}_2 = \frac{2\gamma_0 \omega (\max_i L_i)^2 \max_{i,j} \|x_i - x_j\|^2}{\bar{\mu}^2 n}$ do not depend on α .

In terms of suboptimality convergence for CGD we have

$$\mathbb{E}[\tilde{f}(x^k) - \tilde{f}^*] \stackrel{(11)}{\leq} \frac{L_\alpha}{2} \mathbb{E}\|x^k - x^*\|^2 \stackrel{L_\alpha \stackrel{\text{def}}{=} \bar{L}}{=} \frac{\bar{L}}{2} \alpha^2 \mathbb{E}\|x^k - x^*\|^2 \stackrel{(55)}{\leq} ((1 - \gamma_0 \bar{\mu})^k C_1 + C_2) \alpha^2, \tag{56}$$

where $C_1 = \frac{\tilde{C}_1 \bar{L}}{2}$ and $C_2 = \frac{\tilde{C}_2 \bar{L}}{2}$. \square

8.1.2 DIANA

We follow the same procedure as for the previous subsection. Below, we introduce DIANA algorithm, followed by general theorem, which is then applied to (FLIX).

Algorithm 3 DIANA with different noise levels ω_i

Require: $x^0, h_1^0, \dots, h_n^0 \in \mathbb{R}^d$, $h^0 = \frac{1}{n} \sum_{i=1}^n h_i^0$

- 1: **for** $k = 0, 1, 2, \dots$ **do**
 - 2: Broadcast x^k to all workers
 - 3: **for** $i = 1, \dots, n$ in parallel **do**
 - 4: $\Delta_i^k = \nabla f_i(x^k) - h_i^k$
 - 5: Sample $\hat{\Delta}_i^k \sim \mathcal{C}_i(\Delta_i^k)$
 - 6: $h_i^{k+1} = h_i^k + \beta_i \hat{\Delta}_i^k$
 - 7: $\hat{g}_i^k = h_i^k + \hat{\Delta}_i^k$
 - 8: $g^k = \frac{1}{n} \sum_{i=1}^n \hat{g}_i^k = h^k + \frac{1}{n} \sum_{i=1}^n \hat{\Delta}_i^k$
 - 9: $x^{k+1} = x^k - \gamma g^k$
 - 10: $h^k = \frac{1}{n} \sum_{i=1}^n h_i^{k+1} = h^k + \frac{1}{n} \sum_{i=1}^n \beta_i \hat{\Delta}_i^k$
-

Lemma 3. Suppose each f_i is L_i -smooth and convex, and f is L -smooth. Let $\mathcal{C}_i : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be randomized compression operators satisfying $\mathcal{C}_i \in \mathbb{B}^d(\omega_i)$. Let $g_k = \frac{1}{n} \sum_{i=1}^n h_i^k + \mathcal{C}_i(\nabla f_i(x^k) - h_i^k)$. Then

$$\mathbb{E} \|g^k - \nabla f(x^*)\|^2 \leq 2 \left(L + \frac{2 \max\{L_i \omega_i\}}{n} \right) D_f(x^k, x^*) + \frac{2}{n} \sigma_k^2, \quad (57)$$

where $\sigma_k^2 = \frac{1}{n} \sum_{i=1}^n \omega_i \|h_i^k - \nabla f_i(x^*)\|^2$.

Proof. We start with bias-variance decomposition:

$$\mathbb{E} \|g^k - \nabla f(x^*)\|^2 = \|\nabla f(x^k) - \nabla f(x^*)\|^2 + \mathbb{E} \|g^k - \nabla f(x^k)\|^2, \quad (58)$$

where the first term in RHS bounded due to convexity and smoothness of function f : $\|\nabla f(x^k) - \nabla f(x^*)\|^2 \stackrel{(15)}{\leq} 2LD_f(x^k, x^*)$. Function f is, indeed, convex as a linear combination of convex functions f_i .

When expanding the second term in RHS of equation 58, we encounter scalar products, each of which is zero in expectation due to independence of \mathcal{C}_i for all i and unbiasedness of a compression operator (see the definition in Section 6).

$$\begin{aligned} & \mathbb{E} \|g^k - \nabla f(x^k)\|^2 \\ &= \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n (h_i^k + \mathcal{C}_i(\nabla f_i(x^k) - h_i^k)) - \frac{1}{n} \sum_{i=1}^n \nabla f_i(x^k) \right\|^2 \\ &= \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla f_i(x^k) - h_i^k) - (\nabla f_i(x^k) - h_i^k) \right\|^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \|\mathcal{C}_i(\nabla f_i(x^k) - h_i^k) - (\nabla f_i(x^k) - h_i^k)\|^2 \\ & \quad + \underbrace{\frac{1}{n^2} \sum_{i \neq j} \mathbb{E} \langle \mathcal{C}_i(\nabla f_i(x^k) - h_i^k) - (\nabla f_i(x^k) - h_i^k), \mathcal{C}_j(\nabla f_j(x^k) - h_j^k) - (\nabla f_j(x^k) - h_j^k) \rangle}_{=0}. \end{aligned}$$

Then, we apply bounded variance property of compression operators:

$$\begin{aligned} \mathbb{E} \|g^k - \nabla f(x^k)\|^2 &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \|\mathcal{C}_i(\nabla f_i(x^k) - h_i^k) - (\nabla f_i(x^k) - h_i^k)\|^2 \\ &\leq \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - h_i^k\|^2. \end{aligned}$$

We subtract and add $\nabla f_i(x^*)$ inside each norm operator to split $\nabla f_i(x^k)$ and h_i^k from each other using (16):

$$\begin{aligned} \mathbb{E}\|g^k - \nabla f(x^k)\|^2 &\leq \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - h_i^k\|^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - \nabla f_i(x^*) - (h_i^k - \nabla f_i(x^*))\|^2 \\ &\leq \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - \nabla f_i(x^*)\|^2 + \frac{2}{n^2} \sum_{i=1}^n \omega_i \|h_i^k - \nabla f_i(x^*)\|^2 \end{aligned}$$

While leaving the second term in the last line unchanged (it is basically $2\sigma_k/n$), we apply (15) to each norm of gradient differences in the first term and use max-function over $L_i\omega_i$ to take out Bregman divergences:

$$\begin{aligned} \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla f_i(x^k) - \nabla f_i(x^*)\|^2 &\leq \frac{2}{n^2} \sum_{i=1}^n \omega_i 2L_i D_{f_i}(x^k, x^*) \\ &\leq \frac{4 \max\{\omega_i L_i\}}{n^2} \sum_{i=1}^n D_{f_i}(x^k, x^*) = \frac{4 \max\{\omega_i L_i\}}{n} D_f(x^k, x^*), \end{aligned}$$

where in the last part we used linearity property of Bregman divergence. It remains to incorporate all results to get the statement. \square

Lemma 4. *Assuming all conditions of Lemma 3 hold, let $h_i^{k+1} = h_i^k + \beta_i \mathcal{C}_i(\nabla f_i(x^k) - h_i^k)$ and $\beta_i \leq \frac{1}{\omega_i+1}$, where $i \in \{1, \dots, n\}$. Then,*

$$\sigma_{k+1}^2 \leq (1 - \min \beta_i) \sigma_k^2 + 2 \max\{\beta_i \omega_i L_i\} D_f(x^k, x^*). \quad (59)$$

Proof. Gorbunov et al. (2019) state for $\beta_i \leq \frac{1}{\omega_i+1}$

$$\mathbb{E}\|h_i^{k+1} - \nabla f_i(x^*)\|^2 \leq (1 - \beta_i) \|h_i^k - \nabla f_i(x^*)\|^2 + 2\beta_i L_i D_{f_i}(x^k, x^*). \quad (60)$$

From this it follows that

$$\begin{aligned} \sigma_{k+1}^2 &= \frac{1}{n} \sum_{i=1}^n \omega_i \|h_i^{k+1} - \nabla f_i(x^*)\|^2 \\ &\leq \frac{1}{n} \sum_{i=1}^n (1 - \beta_i) \omega_i \|h_i^k - \nabla f_i(x^*)\|^2 + \frac{1}{n} \sum_{i=1}^n 2\beta_i \omega_i L_i D_{f_i}(x^k, x^*) \\ &\leq (1 - \min \beta_i) \sigma_k^2 + \frac{1}{n} \sum_{i=1}^n 2\beta_i \omega_i L_i D_{f_i}(x^k, x^*) \\ &\leq (1 - \min \beta_i) \sigma_k^2 + 2 \max\{\beta_i \omega_i L_i\} \frac{1}{n} \sum_{i=1}^n D_{f_i}(x^k, x^*) \\ &= (1 - \min \beta_i) \sigma_k^2 + 2 \max\{\beta_i \omega_i L_i\} D_f(x^k, x^*). \end{aligned}$$

\square

Theorem 5. *Assume all conditions of Lemmas 3, 4 hold, and each function f_i is μ_i -strongly convex. If $\gamma \leq \frac{1}{L_\alpha + \frac{2 \max\{L_i \alpha_i^2 \omega_i\}}{n} + \frac{4 \max\{\beta_i \omega_i L_i \alpha_i^2\}}{n \min \beta_i}}$, then*

$$\mathbb{E}[\mathcal{D}^k] \leq \max \left\{ (1 - \gamma \mu_\alpha)^k, \left(1 - \frac{1}{2} \min \beta_i\right)^k \right\} \mathcal{D}^0, \quad (61)$$

where $\mathcal{D}^k = \|x^k - x^*\|^2 + \frac{4}{n \min \beta_i} \gamma^2 \sigma_k^2$ and L_α is defined as in Proposition 2.

Proof. The lemma is a direct corollary of Theorem 4.1 from Gorbunov et al. (2019), 3 and 4 and strong convexity with constants: $A = L_\alpha + \frac{2 \max\{L_i \alpha_i^2 \omega_i\}}{n}$, $B = \frac{2}{n}$, $D_1 = 0$, $\rho = \min \beta_i$, $C = \max\{\beta_i \omega_i L_i \alpha_i^2\}$, $D_2 = 0$, $M = \frac{4}{n \min \beta_i} = \frac{2B}{\rho} \geq \frac{4(1 + \max \omega_i)}{n}$. \square

Corollary 2. Assume all conditions of Theorem 5 hold and $\alpha_i \equiv \alpha \forall i$. If $\gamma = \frac{1}{\alpha^2} \cdot \frac{1}{L + \frac{2 \max\{L_i \omega_i\}}{n} + \frac{4 \max\{\beta_i \omega_i L_i\}}{n \min \beta_i}} =: \gamma_0 \frac{1}{\alpha^2}$, then

$$\mathbb{E} \|x^k - x^*\|^2 \leq \max \left\{ (1 - \gamma_0 \bar{\mu})^k, \left(1 - \frac{1}{2} \min \beta_i\right)^k \right\} \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2} \quad (62)$$

and

$$\mathbb{E} [\tilde{f}(x^k)] - \tilde{f}^* \leq \max \left\{ (1 - \gamma_0 \bar{\mu})^k, \left(1 - \frac{1}{2} \min \beta_i\right)^k \right\} C \alpha^2, \quad (63)$$

where $C = \frac{\bar{L}_i}{2} \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2}$.

Proof. Let us investigate the quantity \mathcal{D}^0 . We assume memory h_0^i in DIANA algorithm is equal to $\nabla[f_i(T_i(x^0))]$ for $i \in \{1, \dots, n\}$. Plugging in definitions and initializations we obtain expanded definition of \mathcal{D}^0 :

$$\begin{aligned} \mathcal{D}^0 &\stackrel{\text{def}}{=} \|x^0 - x^*\|^2 + \frac{4}{n \min \beta_i} \gamma^2 \sigma_0^2 \stackrel{\gamma = \gamma_0 \frac{1}{\alpha^2}}{=} \|x^0 - x^*\|^2 + \frac{4}{n \min \beta_i} \gamma_0^2 \frac{1}{\alpha^4} \sigma_0^2 \\ \sigma_0^2 &\stackrel{\text{def}}{=} \|x^0 - x^*\|^2 + \frac{4}{n \min \beta_i} \gamma_0^2 \frac{1}{\alpha^4} \frac{1}{n} \sum_{i=1}^n \omega_i \|h_0^i - \nabla[f_i(T_i(x^*))]\|^2 \\ &\stackrel{h_0^i \text{ initialisation}}{=} \|x^0 - x^*\|^2 + \frac{4}{n \min \beta_i} \gamma_0^2 \frac{1}{\alpha^4} \frac{1}{n} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^0))] - \nabla[f_i(T_i(x^*))]\|^2 \end{aligned}$$

We further upper bound the sum in the last term using smoothness properties of functions f_i :

$$\begin{aligned} &\frac{1}{n} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^0))] - \nabla[f_i(T_i(x^*))]\|^2 \\ &= \frac{1}{n} \sum_{i=1}^n \omega_i \|\nabla[f_i(\alpha x^0 + (1 - \alpha)x_i^*)] - \nabla[f_i(\alpha x^* + (1 - \alpha)x_i^*)]\|^2 \\ &= \frac{1}{n} \sum_{i=1}^n \omega_i \|\alpha(\nabla f_i(\alpha x^0 + (1 - \alpha)x_i^*) - \nabla f_i(\alpha x^* + (1 - \alpha)x_i^*))\|^2 \\ &= \alpha^2 \frac{1}{n} \sum_{i=1}^n \omega_i \|\nabla f_i(\alpha x^0 + (1 - \alpha)x_i^*) - \nabla f_i(\alpha x^* + (1 - \alpha)x_i^*)\|^2 \\ &\leq \alpha^2 \frac{1}{n} \sum_{i=1}^n \omega_i L_i^2 \|\alpha(x^0 - x^*)\|^2 = \alpha^4 \overline{\omega_i L_i^2} \|x^0 - x^*\|^2. \end{aligned}$$

Plugging this result back into the definition of \mathcal{D}^0 and applying further bounds we get

$$\begin{aligned}
 \mathcal{D}^0 &\leq \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \|x^0 - x^*\|^2 \\
 &\stackrel{(12)}{\leq} \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{2}{\mu_\alpha} (\tilde{f}(x^0) - \tilde{f}(x^*)) \\
 &\stackrel{\text{Proposition 2}}{=} \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{2}{\bar{\mu} \alpha^2} (\tilde{f}(x^0) - \tilde{f}(x^*)) \\
 &\stackrel{(24)}{\leq} \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{2}{\bar{\mu} \alpha^2} \mathcal{E}_\alpha(x^0) \\
 &\stackrel{(27)}{=} \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{1}{\bar{\mu} \alpha^2} \frac{1}{n} \sum_{i=1}^n L_i \alpha^2 \|x^0 - x_i\|^2 \\
 &= \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2},
 \end{aligned}$$

what shows that \mathcal{D}^0 can be upper bounded by value *independent* of α . Since $M\gamma^2\sigma_k^2 \geq 0$ $\mathcal{D}^k = \|x^k - x^*\|^2 + M\gamma^2\sigma_k^2 \geq \|x^k - x^*\|^2$, and that is why

$$\begin{aligned}
 \mathbb{E}\|x^k - x^*\|^2 &\leq \mathbb{E}\mathcal{D}^k \leq \max \left\{ (1 - \gamma\mu_\alpha)^k, \left(1 - \frac{1}{2} \min \beta_i\right)^k \right\} \mathcal{D}^0 \\
 &= \max \left\{ \left(1 - \left(\gamma_0 \frac{1}{\alpha^2}\right) (\bar{\mu} \alpha^2)\right)^k, \left(1 - \frac{1}{2} \min \beta_i\right)^k \right\} \mathcal{D}^0 \\
 &= \max \left\{ (1 - \gamma_0 \bar{\mu})^k, \left(1 - \frac{1}{2} \min \beta_i\right)^k \right\} \mathcal{D}^0 \\
 &\leq \max \left\{ (1 - \gamma_0 \bar{\mu})^k, \left(1 - \frac{1}{2} \min \beta_i\right)^k \right\} \left(1 + \frac{4}{n \min \beta_i} \gamma_0^2 \overline{\omega_i L_i^2}\right) \frac{1}{\bar{\mu}} \overline{L_i \|x^0 - x_i\|^2},
 \end{aligned}$$

what means that the convergence of iterates is completely independent from α . To derive the result for the convergence of functional values we use L_α -smoothness of function \tilde{f} (see Proposition 2): $\mathbb{E}[\tilde{f}(x^k)] - \tilde{f}^* \leq \frac{L_\alpha}{2} \mathbb{E}\|x^k - x^*\|^2 = \alpha^2 \frac{\bar{L}_i}{2} \mathbb{E}\|x^k - x^*\|^2$. □

8.2 Convex objectives

8.2.1 DCGD

As for the previous section, we first introduce the algorithm, followed by general theorem, which is then applied to (FLIX).

Theorem 6. *Assume all conditions of Lemma 2 hold. Let $0 < \gamma \leq \frac{1}{4\left(L_\alpha + \frac{2 \max\{L_i \omega_i \alpha_i^2\}}{n}\right)}$. Then,*

$$\mathbb{E}[\tilde{f}(\bar{x}^k) - \tilde{f}(x^*)] \leq \frac{2(\tilde{f}(x^0) - \tilde{f}(x^*))}{k} + \frac{\|x^0 - x^*\|^2}{\gamma k} + 2\gamma\sigma_{DCGD}, \quad (64)$$

where $\bar{x}^k = (1/k) \sum_{j=1}^k x^j$.

Proof. The theorem is a corollary of Corollary 4.1 from [Khaled et al. \(2020\)](#) with constants

$$\begin{aligned} A &= L_\alpha + \frac{2 \max\{L_i \omega_i \alpha_i^2\}}{n}, \quad B = 0, \quad \sigma_k^2 \equiv 0, \\ \rho &= 1, \quad C = 0, \quad D_2 = 0, \\ D_1 &= \sigma_{DCGD} = \frac{2}{n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x^*))]\|^2 \\ &= \frac{2}{n^2} \sum_{i=1}^n \omega_i \alpha_i^2 \|\nabla f_i(T_i(x^*))\|^2. \end{aligned}$$

□

Corollary 3. *Assume all conditions of Lemma 2 hold. Let $\alpha_i = \alpha$, $\omega_i = \omega$ for all i , $\gamma = \frac{1}{4(\bar{L}_i + \frac{2\omega \max_i L_i}{n})\alpha^2}$ and $\sup_{\alpha \in [0,1]} \inf_{x^* \in \mathcal{X}} \|x^0 - x^*\| = R$. Then,*

$$\begin{aligned} &\mathbb{E}[\tilde{f}(\bar{x}^k) - \tilde{f}(x^*)] \\ &\leq \alpha^2 \cdot \frac{2}{k} \left(\overline{L_i} \|x^0 - x_i\|^2 + 2 \left(\bar{L}_i + \frac{2 \max\{L_i \omega_i\}}{n} \right) R^2 \right) \\ &\quad + \alpha \cdot \frac{\omega \max_i L_i}{n \bar{L}_i + 2\omega \max_i L_i} \left(\tilde{f}(x^*(0)) - \frac{1}{n} \sum_{i=1}^n f_i(x_i) \right) \\ &= \frac{1}{k} C_1 \alpha^2 + C_2 \alpha. \end{aligned}$$

Proof. Let us investigate $\tilde{f}(x^0) - \tilde{f}(x^*)$. Let $\alpha_i \equiv \alpha \in \mathbb{R}$ for all i . Since $\min \frac{1}{n} \sum_{i=1}^n f_i(T_i(x)) \geq \frac{1}{n} \sum_{i=1}^n \min f_i(T_i(x)) = \frac{1}{n} \sum_{i=1}^n \min f_i(x)$ and each f_i is L_i -smooth, we get

$$\begin{aligned} \tilde{f}(x^0) - \tilde{f}(x^*) &\leq \tilde{f}(x^0) - \frac{1}{n} \sum_{i=1}^n f_i(x_i) \\ &= \frac{1}{n} \sum_{i=1}^n f_i(\alpha x^0 + (1-\alpha)x_i) - f_i(x_i) \\ &\stackrel{(10)}{\leq} \frac{1}{n} \sum_{i=1}^n L_i \|\alpha(x^0 - x_i)\|^2 \\ &= \alpha^2 \overline{L_i} \|x^0 - x_i\|^2. \end{aligned}$$

This observation shows that the dependence between $\tilde{f}(x^0) - \tilde{f}(x^*)$ and α is quadratic. $\tilde{f}(x^0) - \tilde{f}(x^*)$ diminishes to zero as α goes to zero what proves our intuition: the closer α is to zero, the less steps the algorithm needs to make till convergence.

We verify that the stepsize in the corollary fits the restriction on stepsizes in Theorem 6.

$$\begin{aligned} \frac{1}{\gamma} &= 4 \left(L_\alpha + \frac{2 \max\{L_i \omega_i \alpha_i^2\}}{n} \right) \\ \alpha_i \equiv \alpha, \text{ Proposition 2} &\stackrel{\equiv}{=} 4 \left(\bar{L}_i + \frac{2 \max\{L_i \omega_i\}}{n} \right) \cdot \alpha^2. \end{aligned} \tag{65}$$

Further assuming that $\inf_{x^* \in \mathcal{X}} \|x^0 - x^*\|$ is bounded above for all $\alpha \in [0, 1]$ by value $R \in \mathbb{R}$, we obtain the following result:

$$\frac{1}{\gamma} \|x^0 - x^*\|^2 \leq 4 \left(\bar{L}_i + \frac{2 \max\{L_i \omega_i\}}{n} \right) R^2 \alpha^2.$$

The neighborhood governing term is $\gamma\sigma_{DCGD}$. Let us assume $\alpha_i \equiv \alpha$ and $\omega_i \equiv \omega$ for all i .

$$\begin{aligned}\sigma_{DCGD} &= \frac{2}{n^2} \sum_{i=1}^n \omega_i \alpha_i^2 \|\nabla f_i(T_i(x^*))\|^2 \\ T_i(x) &\stackrel{\text{def}}{=} \frac{2}{n^2} \sum_{i=1}^n \omega_i \alpha_i^2 \|\nabla f_i(\alpha_i x^* + (1 - \alpha_i)x_i)\|^2 \\ \alpha_i &\stackrel{\text{def}}{=} \alpha, \omega_i \stackrel{\text{def}}{=} \omega \quad \frac{2\omega\alpha^2}{n^2} \sum_{i=1}^n \|\nabla f_i(\alpha_i x^* + (1 - \alpha_i)x_i)\|^2\end{aligned}$$

Since x_i^* minimizes $f_i(x)$, $\nabla f_i(x_i^*) = 0$. Thus,

$$\begin{aligned}&\frac{2\omega\alpha^2}{n^2} \sum_{i=1}^n \|\nabla f_i(\alpha_i x^* + (1 - \alpha_i)x_i)\|^2 \\ &= \frac{2\omega\alpha^2}{n^2} \sum_{i=1}^n \|\nabla f_i(\alpha_i x^* + (1 - \alpha_i)x_i) - \nabla f_i(x_i)\|^2\end{aligned}$$

Inequality (15) further bounds each norm of gradient differences.

$$\begin{aligned}&\frac{2\omega\alpha^2}{n^2} \sum_{i=1}^n \|\nabla f_i(\alpha_i x^* + (1 - \alpha_i)x_i) - \nabla f_i(x_i)\|^2 \\ &\leq \frac{4\omega\alpha^2}{n^2} \sum_{i=1}^n L_i (f_i(\alpha_i x^* + (1 - \alpha_i)x_i) - f_i(x_i)) \\ &\leq \frac{4\omega\alpha^2 \max_i L_i}{n^2} \sum_{i=1}^n (f_i(\alpha_i x^* + (1 - \alpha_i)x_i) - f_i(x_i)) \\ &= \frac{4\omega\alpha^2 \max_i L_i}{n} \left(\tilde{f}(x^*) - \frac{1}{n} \sum_{i=1}^n f_i(x_i) \right)\end{aligned}$$

In the rightmost term point x^* depends on α because it is an optimal point of a problem dependent on α . But the observation below explicitly reveals the dependence. Since each f_i is convex,

$$\begin{aligned}\tilde{f}(x) &= \frac{1}{n} \sum_{i=1}^n f_i(\alpha x + (1 - \alpha)x_i) \\ &\leq \alpha \frac{1}{n} \sum_{i=1}^n f_i(x) + (1 - \alpha) \frac{1}{n} \sum_{i=1}^n f_i(x_i).\end{aligned}$$

Taking the minimum of both sides of inequality above we get to

$$\tilde{f}(x^*) \leq \alpha \tilde{f}(x^*(0)) + (1 - \alpha) \frac{1}{n} \sum_{i=1}^n f_i(x_i).$$

We established that

$$\sigma_{DCGD} \leq \frac{4\omega \max_i L_i}{n} \left(\tilde{f}(x^*(0)) - \frac{1}{n} \sum_{i=1}^n f_i(x_i) \right) \alpha^3. \quad (66)$$

Together with equality 65 it brings us to the following bound on the neighbourhood:

$$\gamma\sigma_{DCGD} \leq \frac{\omega \max_i L_i}{n\bar{L}_i + 2\omega \max_i L_i} \left(\tilde{f}(x^*(0)) - \frac{1}{n} \sum_{i=1}^n f_i(x_i) \right) \alpha = C\alpha,$$

where C does not depend on α .

□

8.2.2 DIANA

As for the previous section, we first introduce the algorithm, followed by general theorem, which is then applied to (FLIX).

Theorem 7. *Assume all conditions of Lemmas 3 and 4 hold. Let $0 < \gamma \leq \frac{1}{4\left(L_\alpha + \frac{2 \max_i \{L_i \alpha_i^2 \omega_i\}}{n} + \frac{4 \max_i \{\beta_i \omega_i L_i \alpha_i^2\}}{n \min_i \beta_i}\right)}$.*

Then,

$$\mathbb{E}[\tilde{f}(\bar{x}^k) - \tilde{f}(x^*)] \leq \left(2(\tilde{f}(x^0) - \tilde{f}(x^*)) + \frac{4\gamma}{n \min \beta_i} \sigma_0^2 + \frac{1}{\gamma} \|x^0 - x^*\|^2\right) \frac{1}{k}. \quad (67)$$

Proof. The theorem is a corollary of Corollary 4.1 from Khaled et al. (2020) with constants

$$A = L_\alpha + \frac{2 \max\{L_i \omega_i \alpha_i^2\}}{n}, B = \frac{2}{n}, D_1 = 0, \\ \rho = \min_i \beta_i, C = \max_i \{\beta_i \omega_i L_i \alpha_i^2\}, D_2 = 0.$$

□

Corollary 4. *Assume all conditions of Lemmas 3 and 4 hold. Let $\alpha_i = \alpha$, $\omega_i = \omega$, $h_i^0 = 0 \in \mathbb{R}^d$ for all i , $\gamma = \frac{1}{4\left(\bar{L}_i + \frac{6\omega \max_i L_i}{n}\right)\alpha^2}$, $\sup_{\alpha \in [0,1]} \inf_{x^* \in \mathcal{X}} \|x^0 - x^*\| = R$. Then,*

$$\mathbb{E}[f(\bar{x}^k) - f(x^*)] \leq \frac{1}{k}(C_1 \alpha^2 + C_2 \alpha),$$

where $C_1 = 2\left(\bar{L}_i \|x^0 - x_i\|^2 + 2\left(\bar{L}_i + \frac{6\omega \max_i L_i}{n}\right) R^2\right)$ and $C_2 = \frac{\omega(\omega+1) \max_i L_i}{(n\bar{L}_i + 6\omega \max_i L_i)} \left(\tilde{f}(x^*(0)) - \frac{1}{n} \sum_{i=1}^n f_i(x_i)\right)$.

Proof. Let us investigate stepsize γ . The statement of the corollary requires the stepsize equals its maximum value $\frac{1}{4\left(L_\alpha + \frac{2 \max_i \{L_i \alpha_i^2 \omega_i\}}{n} + \frac{4 \max_i \{\beta_i \omega_i L_i \alpha_i^2\}}{n \min_i \beta_i}\right)}$. Since $\alpha_i = \alpha$, $\omega_i = \omega$, $\beta_i = \frac{1}{1+\omega}$ for all i ,

$$\gamma = \frac{1}{4\left(L_\alpha + \frac{2 \max_i \{L_i \alpha_i^2 \omega_i\}}{n} + \frac{4 \max_i \{\beta_i \omega_i L_i \alpha_i^2\}}{n \min_i \beta_i}\right)} = \frac{1}{4\left(\bar{L}_i + \frac{6\omega \max_i L_i}{n}\right)\alpha^2}. \quad (68)$$

From definitions it follows that $\sigma_0^2 = \frac{n}{2} \sigma_{DCGD}$ if $h_i^0 = 0 \in \mathbb{R}^d$. Thus,

$$\sigma_0^2 = \frac{n}{2} \sigma_{DCGD} \stackrel{(66)}{\leq} 2\omega \max_i L_i \left(\tilde{f}(x^*(0)) - \frac{1}{n} \sum_{i=1}^n f_i(x_i)\right) \alpha^3.$$

That is why

$$\frac{4\gamma}{n \min \beta_i} \sigma_0^2 \stackrel{\beta_i = \frac{1}{\omega+1}}{=} \frac{4\gamma(\omega+1)}{n} \sigma_0^2 \stackrel{(68)}{=} \frac{\omega+1}{n\left(\bar{L}_i + \frac{6\omega \max_i L_i}{n}\right)\alpha^2} \sigma_0^2 \\ = \frac{\omega+1}{(n\bar{L}_i + 6\omega \max_i L_i)\alpha^2} \sigma_0^2 \leq \frac{2\omega(\omega+1) \max_i L_i}{(n\bar{L}_i + 6\omega \max_i L_i)} \left(\tilde{f}(x^*(0)) - \frac{1}{n} \sum_{i=1}^n f_i(x_i)\right) \alpha.$$

Remaining terms are analyzed as for DCGD in convex case. □

8.3 Nonconvex objectives

In this section we analyze algorithms in the nonconvex case of functions. Throughout this section we assume that x_i^* is a stationary point of function f_i . For the sake of convenience we transfer Lemma 1 from Khaled and Richtárik (2020) to here.

Lemma 5 (Lemma 1 from Khaled and Richtárik (2020)). *Suppose $f_i^* = \min f_i(x)$ exists and f_i is L_i -smooth. Then, for any $x \in \mathbb{R}^d$ it holds that*

$$\|\nabla f_i(x)\|^2 \leq 2L_i(f_i(x) - f_i^*).$$

Throughout this section we will use notation $f^* \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n f_i^*$.

8.3.1 DCGD

Theorem 8. Suppose $f_i^* = \min f_i(x)$ exists and f_i is L_i -smooth for each i . Then, if $\gamma \leq \frac{1}{L_\alpha}$,

$$\min_{0 \leq t \leq k-1} \mathbb{E} \|\nabla \tilde{f}(x^t)\|^2 \leq \frac{2 \left(1 + \frac{2L_\alpha \gamma^2 \max\{L_i \omega_i \alpha_i^2\}}{n}\right)^k}{\gamma k} (\tilde{f}(x^0) - f^*). \quad (69)$$

Proof. The proof is a direct application of Theorem 2 in (Khaled and Richtárik, 2020) to our setting. First, let us analyze $\mathbb{E} \|g(x)\|^2$.

$$\begin{aligned} \mathbb{E} \|g(x)\|^2 &= \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla[f_i(T_i(x))]) \right\|^2 \\ &= \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla[f_i(T_i(x))]) - \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla[f_i(T_i(x))]) \right] \right\|^2 + \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla[f_i(T_i(x))]) \right\|^2 \\ &= \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla[f_i(T_i(x))]) - \frac{1}{n} \sum_{i=1}^n \nabla[f_i(T_i(x))] \right\|^2 + \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \nabla f_i(T_i(x)) \right\|^2 \\ &= \mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\nabla[f_i(T_i(x))]) - \frac{1}{n} \sum_{i=1}^n \nabla[f_i(T_i(x))] \right\|^2 + \|\nabla \tilde{f}(x)\|^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \|\mathcal{C}_i(\nabla[f_i(T_i(x))]) - \nabla[f_i(T_i(x))]\|^2 + \|\nabla \tilde{f}(x)\|^2 \\ &\leq \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x))]\|^2 + \|\nabla \tilde{f}(x)\|^2. \end{aligned}$$

To analyze the first term we refer ourselves to Lemma 5:

$$\begin{aligned} \|\nabla[f_i(T_i(x))]\|^2 &\stackrel{T_i \text{ def}}{=} \|\nabla[f_i(\alpha_i x + (1 - \alpha_i)x_i)]\|^2 = \alpha_i^2 \|\nabla f_i(\alpha_i x + (1 - \alpha_i)x_i)\|^2 \\ &\stackrel{\text{Lemma 5}}{\leq} \alpha_i^2 2L_i (f_i(\alpha_i x + (1 - \alpha_i)x_i) - f_i^*) \end{aligned}$$

Plugging this in the previous inequality we get

$$\begin{aligned} \mathbb{E} \|g(x)\|^2 &\leq \frac{1}{n^2} \sum_{i=1}^n \omega_i \|\nabla[f_i(T_i(x))]\|^2 + \|\nabla \tilde{f}(x)\|^2 \\ &\leq \frac{1}{n^2} \sum_{i=1}^n \omega_i \alpha_i^2 2L_i (f_i(\alpha_i x + (1 - \alpha_i)x_i) - f_i^*) + \|\nabla \tilde{f}(x)\|^2 \\ &\leq \frac{2 \max\{L_i \omega_i \alpha_i^2\}}{n} \left(\frac{1}{n} \sum_{i=1}^n f_i(\alpha_i x + (1 - \alpha_i)x_i) - \frac{1}{n} \sum_{i=1}^n f_i^* \right) + \|\nabla \tilde{f}(x)\|^2 \\ &= \frac{2 \max\{L_i \omega_i \alpha_i^2\}}{n} (\tilde{f}(x) - f^*) + \|\nabla \tilde{f}(x)\|^2, \end{aligned}$$

what means that Assumption 2 from Khaled and Richtárik (2020) holds with $A = \frac{2 \max\{L_i \omega_i \alpha_i^2\}}{n}$, $B = 1$, $C = 0$. Then, as Theorem 2 from the same paper states, for stepsize $\gamma \leq \frac{1}{L_\alpha}$ it holds that

$$\min_{0 \leq t \leq k-1} \mathbb{E} \|\nabla \tilde{f}(x^t)\|^2 \leq \frac{2 \left(1 + \frac{2L_\alpha \gamma^2 \max\{L_i \omega_i \alpha_i^2\}}{n}\right)^k}{\gamma k} (\tilde{f}(x^0) - f^*).$$

□

Corollary 5. *Suppose assumptions of Theorem 8 hold, $\alpha_i \equiv \alpha$, $\omega_i \equiv \omega$. Let $\gamma = \gamma_0 \frac{1}{\alpha^2} \leq \frac{1}{L_i \alpha^2}$ and $\Delta_0 = \sup_{\alpha} \tilde{f}(x^0) - f^*$.*

$$\min_{0 \leq t \leq k-1} \mathbb{E} \|\nabla \tilde{f}(x^t)\|^2 \leq \alpha^2 \frac{2 \left(1 + \frac{2\bar{L}\omega\gamma_0^2 \max_i L_i}{n}\right)^k}{\gamma_0 k} \Delta_0. \quad (70)$$

Proof. We note that when $\alpha_i \equiv \alpha$, $L_\alpha = \bar{L}\alpha^2$, which follows from the definition of L_α . Plugging equations $L_\alpha = \bar{L}\alpha^2$, $\gamma = \gamma_0 \frac{1}{\alpha^2}$, $\alpha_i \equiv \alpha$ and $\omega_i \equiv \omega$ into the main result of Theorem 8 we get

$$\min_{0 \leq t \leq k-1} \mathbb{E} \|\nabla \tilde{f}(x^t)\|^2 \leq \alpha^2 \frac{2 \left(1 + \frac{2\bar{L}\omega\gamma_0^2 \max_i L_i}{n}\right)^k}{\gamma_0 k} (\tilde{f}(x^0) - f^*).$$

Noting that $\tilde{f}(x^0) - f^* \leq \Delta_0$ we finish the proof. \square

8.3.2 DIANA

Theorem 9. *Suppose $f_i^* = \min f_i(x)$ exists and f_i is L_i -smooth for each i . Let $\beta_k \equiv \beta$ and $\omega_k \equiv \omega$ for all $k \in \{1, \dots, n\}$. Suppose $\beta \in \frac{1}{1+\omega} \left[1, \min \left\{1 + \frac{1}{2(1+\omega)}, \sqrt{\frac{3}{2} + \omega + \frac{1}{2\omega}}\right\}\right]$ and $\gamma \leq \frac{1}{L_\alpha} \cdot \min \left\{\frac{1}{2\sqrt{\eta_0}}, \frac{2}{\sqrt{1+8\eta_0+1}}\right\}$, $\eta_0 = \frac{(1+\omega)\omega(1+2(1+\omega))}{2}$, and $h_i^0 = \nabla[f_i(T_i(x^0))]$. Let \hat{x} be a point chosen uniformly at random among iterates x^0, x^1, \dots, x^{k-1} generated by DIANA. Then,*

$$\mathbb{E} \|\nabla \tilde{f}(\hat{x})\|^2 \leq \frac{2}{\gamma k} (\tilde{f}(x^0) - f^*). \quad (71)$$

Proof. The proof is following the proof of Theorem 1 in Li and Richtárik (2020) with small modifications: we rewrite the main proposition in the different way and give slightly different bound on the stepsize. First, Lemma 9 in Li and Richtárik (2020) says that

$$\mathbb{E} \|g^k\|^2 \leq \|\nabla f(x^k)\|^2 + \frac{1+\omega}{n} \sigma_k^2, \quad (72)$$

$$\mathbb{E} \sigma_{k+1}^2 \leq (1-\rho) \sigma_k^2 + \frac{\omega(1+r)L_\alpha^2 \gamma^2}{1+\omega} \|\nabla f(x)\|^2, \quad (73)$$

where

$$\sigma_k^2 = \frac{\omega}{(1+\omega)n} \sum_{i=1}^n \|\nabla[f_i(T_i(x^k))] - h_i^k\|^2 \quad (74)$$

$$\rho = \theta - \frac{\omega(1+r)}{n} L_\alpha^2 \gamma^2 \quad (75)$$

$$\theta = \min\{1 - \beta^2 \omega, 2\beta - \frac{1-\beta}{r} - \beta^2 - \beta^2 \omega\} \quad (76)$$

and r is an arbitrary positive number. We choose $r = 2(1+\omega) > 0$. For the sake of readability, we define function ζ in the following way

$$2\beta - \frac{1-\beta}{2(1+\omega)} - \beta^2 - \beta^2 \omega = -(\omega+1)\beta^2 + \left(2 + \frac{1}{2(1+\omega)}\right)\beta - \frac{1}{2(1+\omega)} =: \zeta(\beta)$$

That is why we can write that $\theta \in \min\{1 - \beta^2 \omega, \zeta(\beta)\}$.

Second, let us analyze the lower bound for θ . Note that for β lying in the range $\frac{1}{1+\omega} \left[1, \min \left\{1 + \frac{1}{2(1+\omega)}, \sqrt{\frac{3}{2} + \omega + \frac{1}{2\omega}}\right\}\right]$ we have

$$\begin{aligned} 1 - \beta^2 \omega &\geq 1 - \frac{\omega}{(1+\omega)^2} \left(\frac{3}{2} + \omega + \frac{1}{2\omega}\right) = 1 - \frac{\omega}{(1+\omega)^2} \frac{(1+2\omega)(1+\omega)}{2\omega} \\ &= 1 - \frac{1+2\omega}{2(1+\omega)} = \frac{1}{2(1+\omega)}, \end{aligned}$$

and

$$\begin{aligned}\zeta\left(\frac{1}{1+\omega}\right) &= -\frac{1}{1+\omega} + \left(2 + \frac{1}{2(1+\omega)}\right) \frac{1}{1+\omega} - \frac{1}{2(1+\omega)} \\ &= \frac{1}{2(1+\omega)} + \frac{1}{2(1+\omega)^2} \geq \frac{1}{2(1+\omega)},\end{aligned}$$

and

$$\begin{aligned}\zeta\left(\frac{1}{1+\omega} + \frac{1}{2(1+\omega)^2}\right) &= -\frac{1}{1+\omega} - \frac{1}{(1+\omega)^2} - \frac{1}{4(1+\omega)^3} \\ &\quad + \left(2 + \frac{1}{2(1+\omega)}\right) \left(\frac{1}{1+\omega} + \frac{1}{2(1+\omega)^2}\right) - \frac{1}{2(1+\omega)} \\ &= \frac{1}{2(1+\omega)} + \frac{1}{2(1+\omega)^2} \geq \frac{1}{2(1+\omega)}.\end{aligned}$$

Since ζ is a quadratic function with the ends of the parabola pointed downwards, for all β in the range we have $\zeta(\beta) \geq \frac{1}{2(1+\omega)}$, which altogether means that $\theta \geq \frac{1}{2(1+\omega)} > 0$.

If we enforce

$$\frac{\omega(1+r)}{n} L_\alpha^2 \gamma^2 \leq \frac{1}{4(1+\omega)}$$

or

$$\gamma \leq \frac{1}{2L_\alpha} \sqrt{\frac{n}{(1+\omega)\omega(1+r)}} = \frac{1}{2L_\alpha} \sqrt{\frac{n}{(1+\omega)\omega(1+2(1+\omega))}},$$

then $\rho = \theta - \frac{\omega(1+r)}{n} L_\alpha^2 \gamma^2 \geq \frac{1}{4(1+\omega)}$.

Now let us switch to the proof of the convergence. We first note that due to smoothness of f_i s and inequality 72 we get

$$\begin{aligned}\mathbb{E}f(x^{k+1}) &\leq f(x^k) + \mathbb{E}\langle \nabla f(x^k), x^{k+1} - x^k \rangle + \frac{L_\alpha}{2} \mathbb{E}\|x^{k+1} - x^k\|^2 \\ &= f(x^k) - \gamma \|\nabla f(x^k)\|^2 + \frac{L_\alpha \gamma^2}{2} \mathbb{E}\|g^k\|^2 \\ &\leq f(x^k) - \left(\gamma - \frac{L_\alpha \gamma^2}{2}\right) \|\nabla f(x^k)\|^2 + \frac{L_\alpha \gamma^2}{2} \frac{1+\omega}{n} \sigma_k^2.\end{aligned}$$

Let us fix $\xi > 0$. Then according to 73 we have

$$\begin{aligned}\mathbb{E}[f(x^{k+1}) - f^* + \xi \sigma_{k+1}^2] \\ \leq f(x^k) - f^* - \left(\gamma - \frac{L_\alpha \gamma^2}{2} - \xi \frac{\omega(1+r)L_\alpha^2 \gamma^2}{1+\omega}\right) \|\nabla f(x^k)\|^2 + \left(\xi(1-\rho) + \frac{L_\alpha \gamma^2}{2} \frac{1+\omega}{n}\right) \sigma_k^2.\end{aligned}$$

Let us notate $\Delta^k = f(x^k) - f^* + \xi \sigma_k^2$ and set $\xi = \frac{L_\alpha \gamma^2}{2\rho} \frac{1+\omega}{n}$. Then from previous inequality it follows that

$$\mathbb{E}\Delta^{k+1} \leq \Delta^k - \left(\gamma - \frac{L_\alpha \gamma^2}{2} - \xi \frac{\omega(1+r)L_\alpha^2 \gamma^2}{1+\omega}\right) \|\nabla f(x^k)\|^2. \quad (77)$$

Let us define γ' the coefficient in front of $\|\nabla f(x^k)\|^2$ in the last inequality. When is γ' larger than $\frac{\gamma}{2}$?

$$\begin{aligned}\gamma - \frac{L_\alpha \gamma^2}{2} - \xi \frac{\omega(1+r)L_\alpha^2 \gamma^2}{1+\omega} &= \gamma - \frac{L_\alpha \gamma^2}{2} - \frac{L_\alpha \gamma^2}{2\rho} \frac{1+\omega}{n} \frac{\omega(1+r)L_\alpha^2 \gamma^2}{1+\omega} \\ &= \gamma - \frac{L_\alpha \gamma^2}{2} - \frac{L_\alpha \gamma^2 \omega(1+r)L_\alpha^2 \gamma^2}{2\rho n} = \gamma - \frac{\gamma}{2} \left(L_\alpha \gamma + \frac{L_\alpha^3 \gamma^3 \omega(1+r)}{\rho n}\right) \geq \frac{\gamma}{2},\end{aligned}$$

if

$$L_\alpha \gamma + \frac{L_\alpha^3 \gamma^3 \omega(1+r)}{\rho n} \leq 1.$$

Let us define $y := L_\alpha \gamma$ and $a := \frac{\omega(1+r)}{n}$. Then according to (75) $\rho = \theta - ay^2$. Plugging this back to the last inequality we get

$$\begin{aligned} y + \frac{ay^3}{\theta - ay^2} &\leq 1 \\ \Leftrightarrow \frac{y\theta - ay^3 + ay^3}{\theta - ay^2} &\leq 1 \\ \Leftrightarrow \frac{y\theta}{\theta - ay^2} &\leq 1 \\ \stackrel{\theta - ay^2 = \rho > 0}{\Leftrightarrow} y\theta &\leq \theta - ay^2 \\ \stackrel{\theta > 0}{\Leftrightarrow} \frac{a}{\theta} y^2 + y - 1 &\leq 0. \end{aligned}$$

The last inequality holds for all y lying between zero and the positive root of the quadratic equation. The positive root is $\frac{\sqrt{1 + \frac{4a}{\theta}} - 1}{\frac{2a}{\theta}} = \frac{\frac{4a}{\theta}}{2a(\sqrt{1 + \frac{4a}{\theta}} + 1)} = \frac{2}{\sqrt{1 + \frac{4a}{\theta}} + 1}$. That means that the last inequality holds if

$$y \leq \frac{2}{\sqrt{1 + \frac{4a}{\theta}} + 1}$$

or the tightest bound for γ achieved on smallest θ gives

$$\gamma \leq \frac{1}{L_\alpha} \frac{2}{\sqrt{1 + \frac{4\omega(1+2(1+\omega))}{n\theta_{\min}} + 1}} = \frac{1}{L_\alpha} \frac{2}{\sqrt{1 + \frac{8(1+\omega)\omega(1+2(1+\omega))}{n} + 1}}.$$

Going back to 77 we finally get

$$\begin{aligned} \mathbb{E}\Delta^{k+1} &\leq \Delta^k - \frac{\gamma}{2} \|\nabla f(x^k)\|^2 \\ \mathbb{E}\|\nabla f(x^k)\|^2 &\leq \frac{2}{\gamma} (\mathbb{E}\Delta^k - \mathbb{E}\Delta^{k+1}). \end{aligned}$$

That is why

$$\mathbb{E}\|\nabla f(\hat{x})\|^2 = \frac{1}{k} \sum_{t=0}^{k-1} \mathbb{E}\|\nabla f(x^t)\|^2 \leq \frac{2}{\gamma k} (\Delta^0 - \mathbb{E}\Delta^k) \leq \frac{2}{\gamma k} \Delta^0.$$

The proof is written for general function $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$. To apply the theorem for our setting, we replace f by \tilde{f} . \square

Corollary 6. *Suppose all conditions of Theorem 9 hold and $\alpha_i \equiv \alpha$ for all i from $\{1, \dots, n\}$. Let $\gamma = \frac{1}{\alpha^2} \gamma_0$, where $\gamma_0 = \frac{1}{L} \min \left\{ \frac{1}{2\sqrt{\eta_0}}, \frac{2}{\sqrt{1+8\eta_0+1}} \right\}$. Let $\Delta_0 = \sup_x \tilde{f}(x^0) - f^*$. Then,*

$$\mathbb{E}\|\nabla \tilde{f}(\hat{x})\|^2 \leq \alpha^2 \frac{2}{\gamma_0 k} \Delta_0. \quad (78)$$

Proof. According to the definition of $L_\alpha = \frac{1}{n} \sum_{i=1}^n L_i \alpha_i^2$, in the case of equal α_i s we get $L_\alpha = \bar{L} \alpha^2$. It remains to notice that the stepsize in the corollary condition still satisfies the condition on the stepsize γ in Theorem 9 and $\Delta_0 \geq \tilde{f}(x) - f^*$. \square

9 DISCUSSION OF OTHER MODEL MIXTURE METHODS

Deng et al. (2020) propose that each client solves the local problem

$$\min_{v \in \mathbb{R}^d} f_i(\alpha_i v + (1 - \alpha_i) w^*),$$

where w^* is the minimizer of (ERM). This does not result in any personalization since the “personalized” solution on each node is a reparameterization of each local solution x_i . Furthermore, the convergence theory that Deng et al. (2020) develop does not recover the linear convergence of gradient descent at $\alpha_i = 1$.

Mansour et al. (2020) introduce a similar method, MAPPER, where they propose to solve

$$\min_{z, \alpha_i, y_i} \frac{1}{n} \sum_{i=1}^n f_i(\alpha_i y_i + (1 - \alpha_i) z).$$

Again, this objective is trivially minimized by setting $\alpha_i = 1$, $y_i = \min_{x \in \mathbb{R}^d} f_i(x)$, and $z = 0$ (i.e. with no personalization at all).

Zec et al. (2021) also introduce a similar formulation based on the mixture of experts framework, where they propose to first learn the minimizer x_* of (ERM), learn the optimal local models x_1, x_2, \dots, x_n , and then learn a mixture of both the global and local models (i.e. the α_i) on each client. Unfortunately, this is also ill-defined, as $\alpha_i = 1$ will always perform best on the local training set, and hence if there is no additional data the optimization process cannot improve over the local minimizers x_1, \dots, x_n .

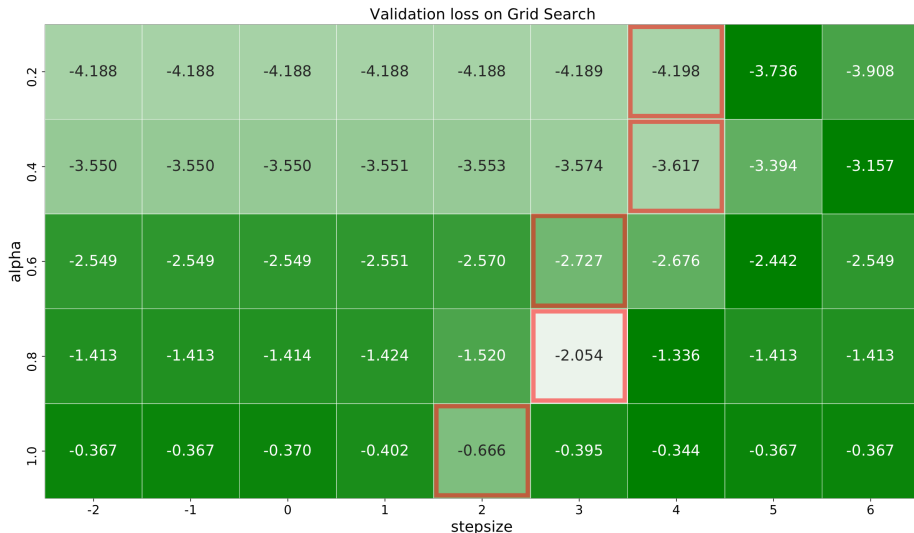


Figure 6: Grid-Search of stepsizes for FLIX for the first Stack Overflow experiment. Values in cells are logarithms of validation losses. Horizontal axis corresponds to logarithm of scaled step-sizes. Red frames correspond to minimum value in a row. Colors are scaled row-wise.

10 EXPERIMENTAL DETAILS

Generalization experiment 1: Fitting Sine Functions. We train each local model i until the gradient norm $\|\nabla f_i(x)\|$ is below 10^{-2} . Stopping criteria for global models is the same but with respect to FLIX formulation: $\|\frac{1}{n} \sum_{i=1}^n \nabla f_i(T_i(x))\| < 10^{-2}$. Gradient descent with line search at each iteration was used as a local and a meta optimizer.

Generalization experiment 2: Comparison to FOMAML and Reptile. We preprocess the raw data in the same way as in Reddi et al. (2021) for Stackoverflow logistic regression task, i.e., each feature vector is a bag-of-words representation of a user’s sentence, each label vector is a binary vector showing if a sentence relates to a particular question tag or not. Word vocabulary for the feature dataset is restricted to the 10000 most common words. We restrict the task to the 500 most used tags. For preprocessing, we used Tensorflow computational procedures from Xu et al. (2021).

To select 50 clients for the Figure 1, we map first 5000 clients from train dataset to vector space with BERT Devlin et al. (2019) and run k-means with 10 clusters. The first 50 clients from the first cluster have been selected for the experiment.

The hold-out validation dataset is of size 100, for the first experiment, and of size 110, for the second one. To compute pure local models we run gradient descent until the norm of the gradient is less than $[10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}]$, respectively. We train pure local models until the gradient norm is less than 10^{-4} , for the first experiments, and 10^{-5} , for the second, as this tolerance level was observed to have the lowest generalization error on the validation dataset in our experiments.

Then, we use a grid-search to find optimal stepsizes for gradient descent used for training FLIX. Our grid for step-sizes are $[5 \cdot 10^{-2}, 5 \cdot 10^{-1}, \dots, 5 \cdot 10^6]$. The results are presented in Figures 6 and 7.

As table shows, the smaller α is, the higher step-size the task needs to achieve the best generalization, which is in line with our convergence results, see Theorem 2 and Table 2.

To train FOMAML and Reptile, we set the number of inner steps to five and grid search outer and inner loop step-sizes. The explored outer step-sizes are the same as for FLIX. Inner step-sizes iterate over the set $[10^{-3}, 10^{-2}, \dots, 10]$.




Figure 7: Grid-Search of stepsizes for FLIX for the second Stack Overflow experiment. Values in cells are logarithms of validation losses. Horizontal axis corresponds to logarithm of scaled step-sizes. Red frames correspond to minimum value in a row. Colors are scaled row-wise. Grey cells correspond to nan values.

Table 3: Combinations of the parameters that achieve the highest accuracy. See Section 4 for the description of the grids.

PARAMETER	EMNIST	SHAKESPEARE
Personalization parameter α	0.3	0.5
PLM batch size	10	4
PLM learning rate	0.001	0.1
FLIX batch size	20	1
FLIX learning rate	0.01	$10^{0.5}$

After grid search, we run gradient descent for each value of alpha of FLIX, FOMAML, and Reptile for 10 000 and 50000 iterations in the first and seconds experiments accordingly and report obtained test accuracy.

Generalization experiment 3: Comparison to FedAvg with CNNs and LSTMs on large-scale datasets. In Table 3 we indicate the best parameters for the tasks discussed in Section 4.

11 EXTRA EXPERIMENTS

Logistic regression with l_2 regularizer. On figures 8, 9, 10, and 11 we present the plots for regularized logistic regression loss for all datasets mentioned in the previous section and for all 3 algorithms: GD, CGD and DIANA. Number of machines n is set to 50 for DIANA and GD, and is set to 8 for CGD. Plots showing the dependence between loss and communication cost show that although in terms of communication rounds DIANA and CGD lose to classical GD, they are better with respect to communication cost, which is of more practical importance. We run all the algorithms with their best theoretical step-sizes.

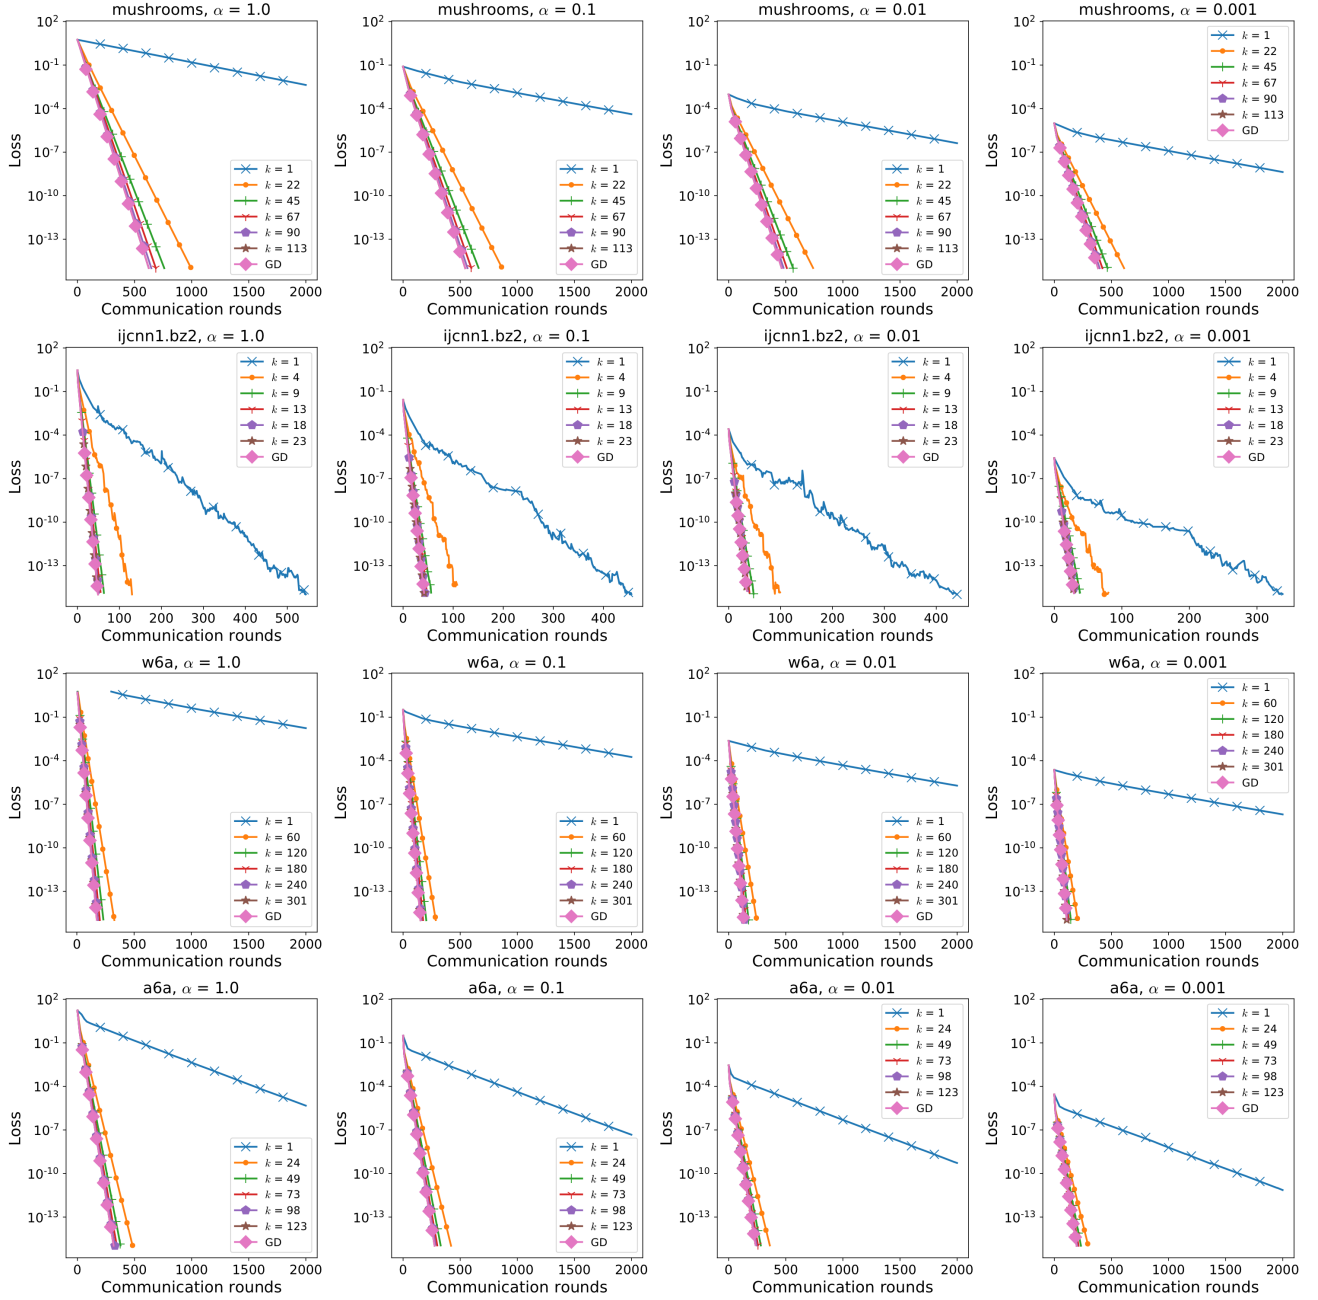


Figure 8: Loss $f(x) - f^*$ vs. # of communication rounds of DIANA and GD for logistic regression problem l_2 regularizer for four datasets, k is a sparsification parameter of Random- k compressor.

Unregularized logistic regression. Setting $\lambda = 0$ in problem 8, we obtain unregularized logistic regression task, which is well-known to be a convex problem. Similarly to regularized case, we present four plots (12, 13, 14,

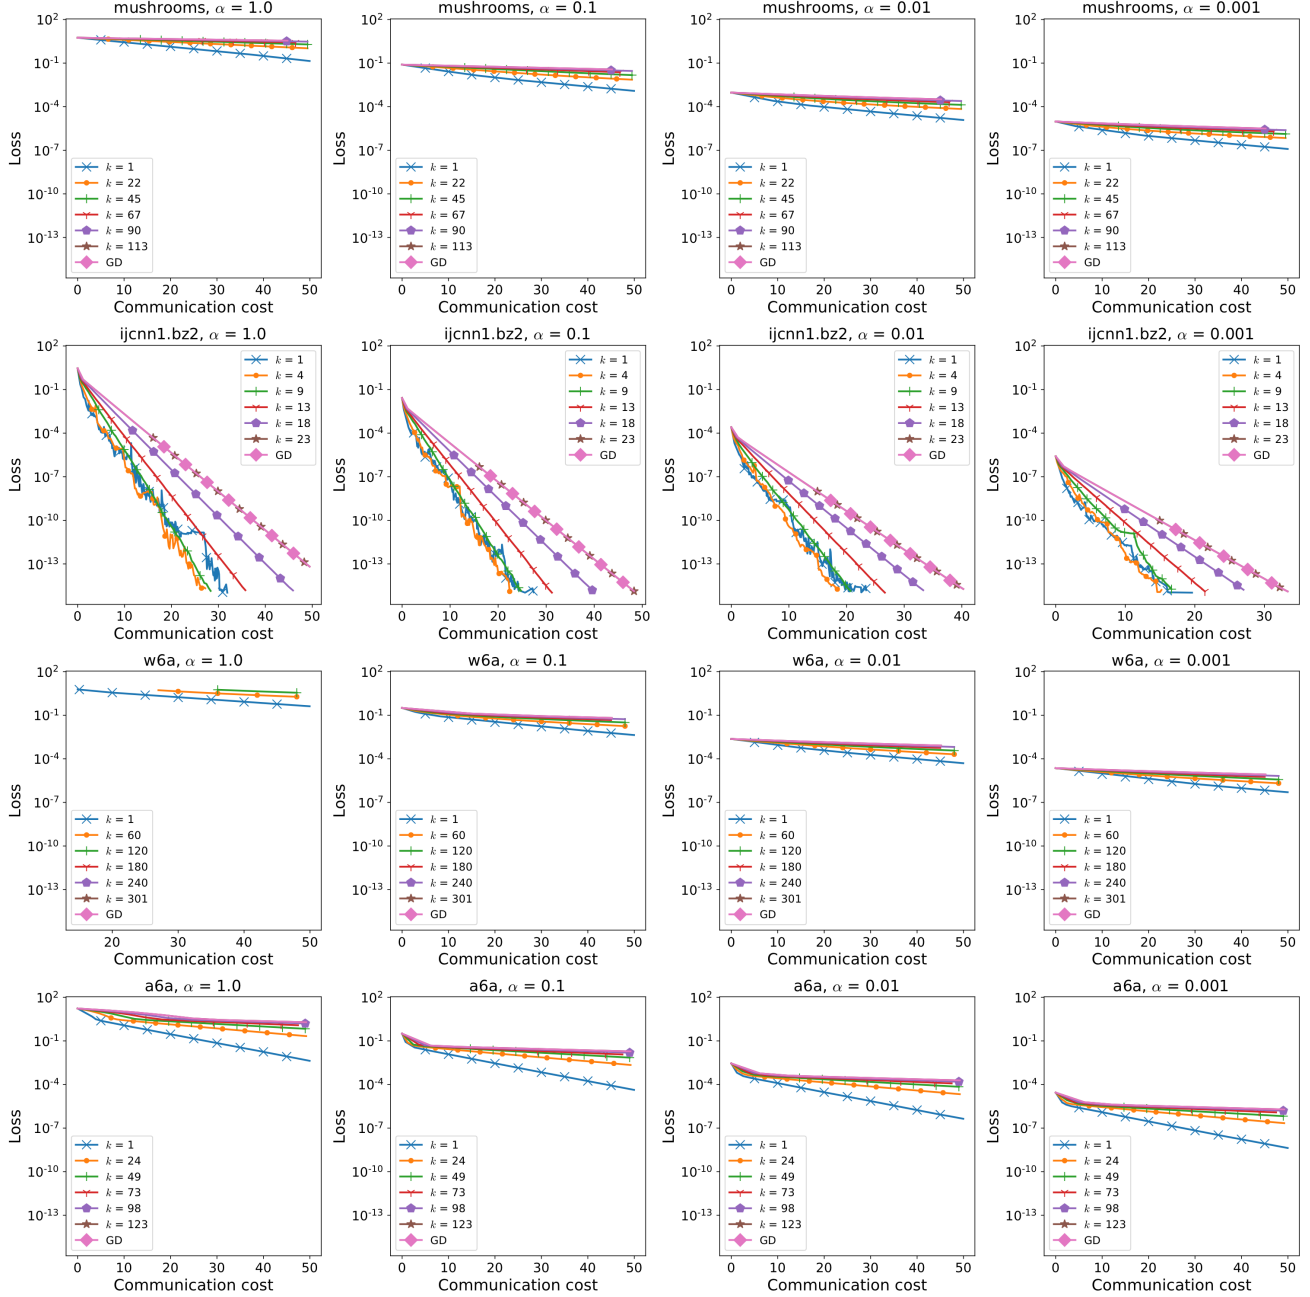


Figure 9: Loss $f(x) - f^*$ vs. communication cost in unit of thousands of float numbers of DIANA and GD for logistic regression problem l_2 regularizer for four datasets, k is a sparsification parameter of Random- k compressor.

and 15) exhibiting convergence of three algorithms in terms of communication rounds and overall communication cost for four LIBSVM datasets. Number of machines n is set to 100 for DIANA and GD, and to 8 for CGD. We run all the algorithms with their best theoretical step-sizes.

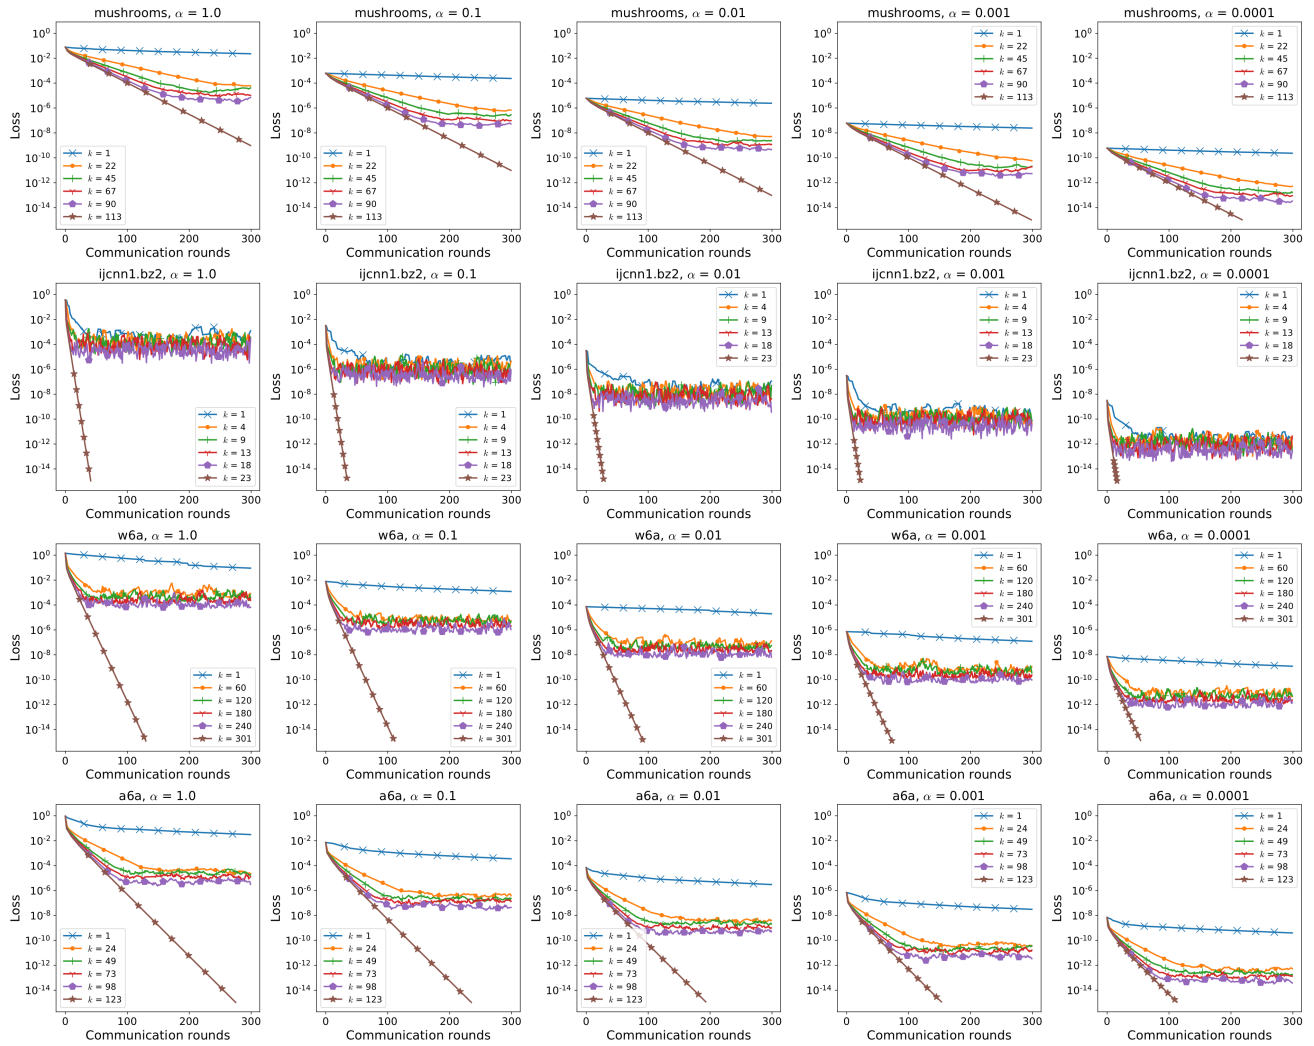


Figure 10: Loss $f(x) - f^*$ vs. # of communication rounds of CGD for logistic regression problem l_2 regularizer for four datasets, k is a sparsification parameter of Random- k compressor.

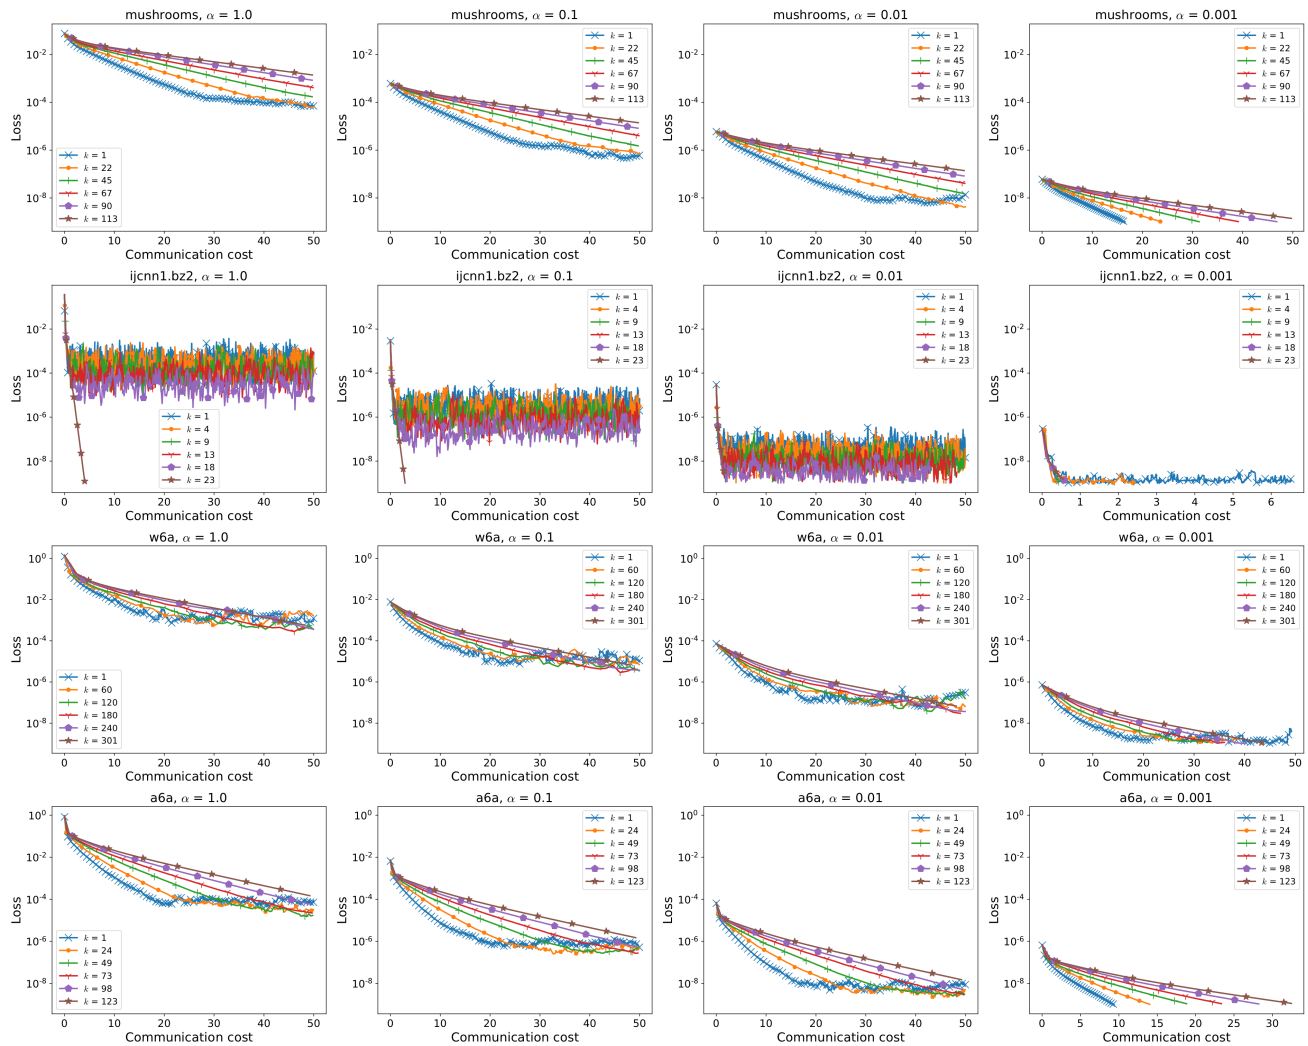


Figure 11: Loss $f(x) - f^*$ vs. communication cost in unit of thousands of float numbers of CGD for logistic regression problem l_2 regularizer for four datasets, k is a sparsification parameter of Random- k compressor.

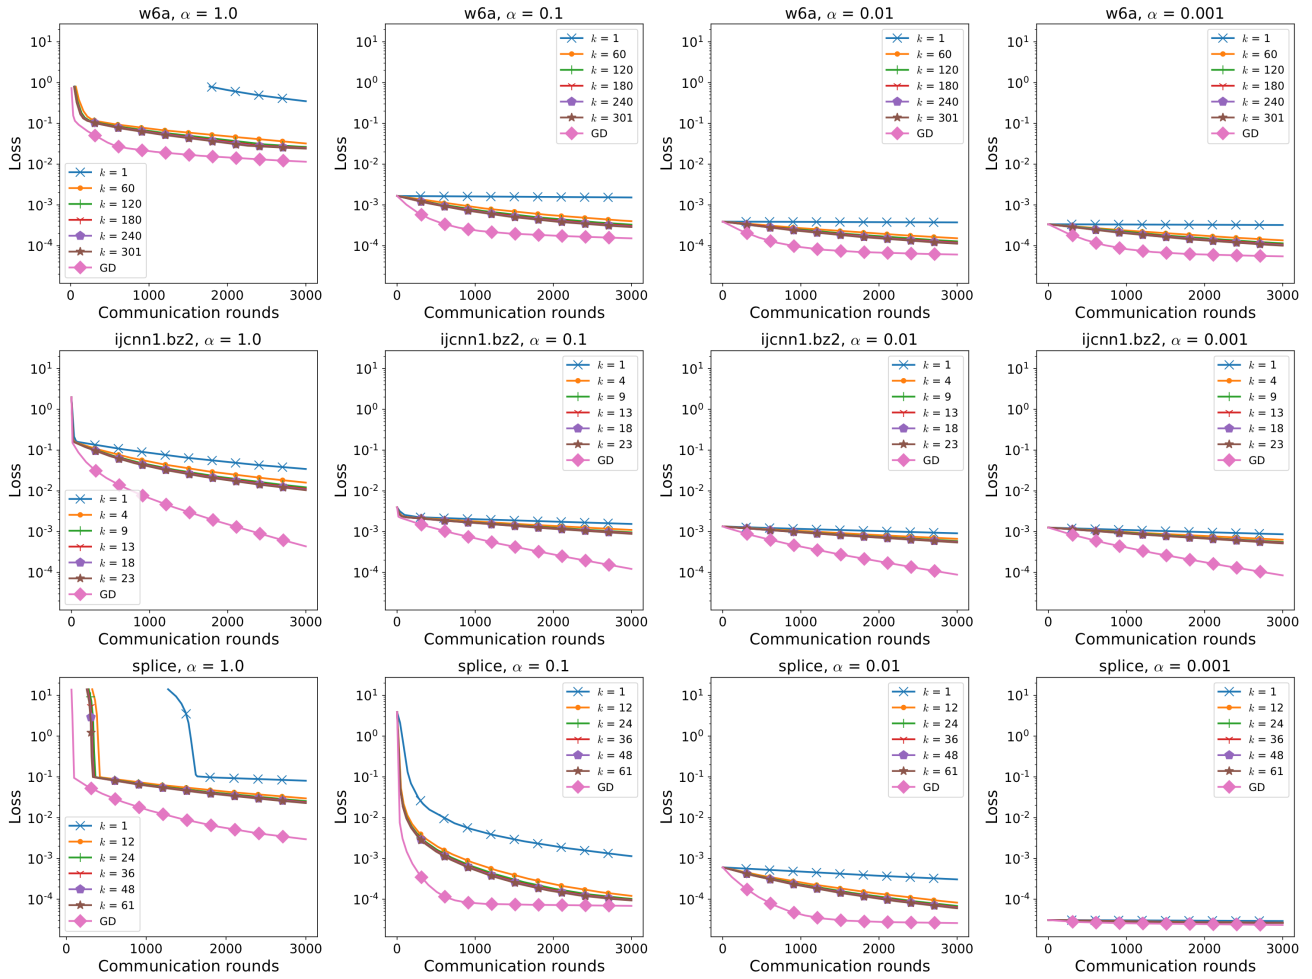


Figure 12: Loss $f(x) - f^*$ vs. # of communication rounds of DIANA and GD for unregularized logistic regression problem for four datasets, k is a sparsification parameter of Random- k compressor.

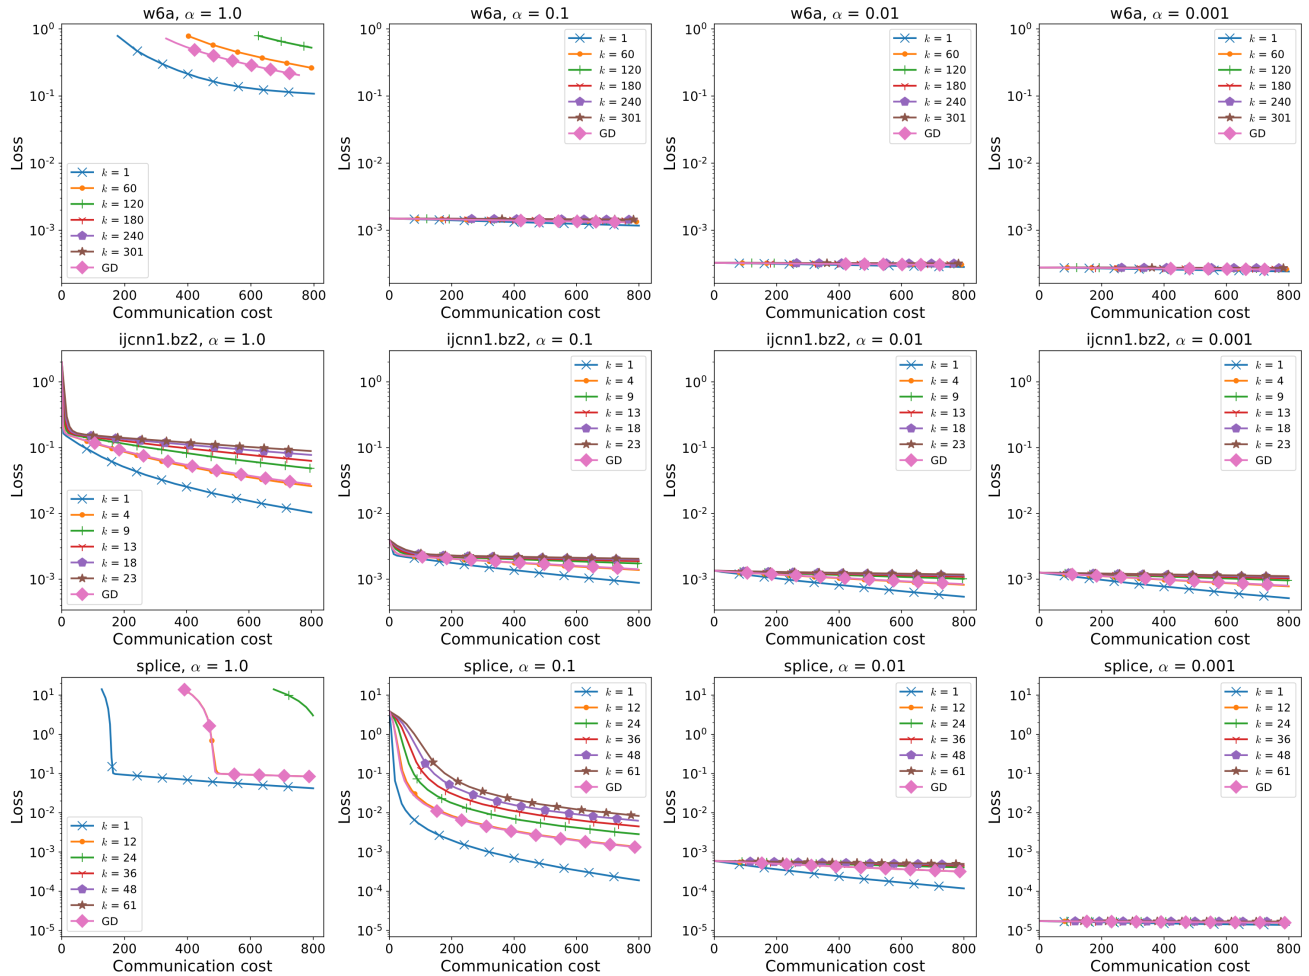


Figure 13: Loss $f(x) - f^*$ vs. communication cost in unit of thousands of float numbers of DIANA and GD for unregularized logistic regression problem for four datasets, k is a sparsification parameter of Random- k compressor.

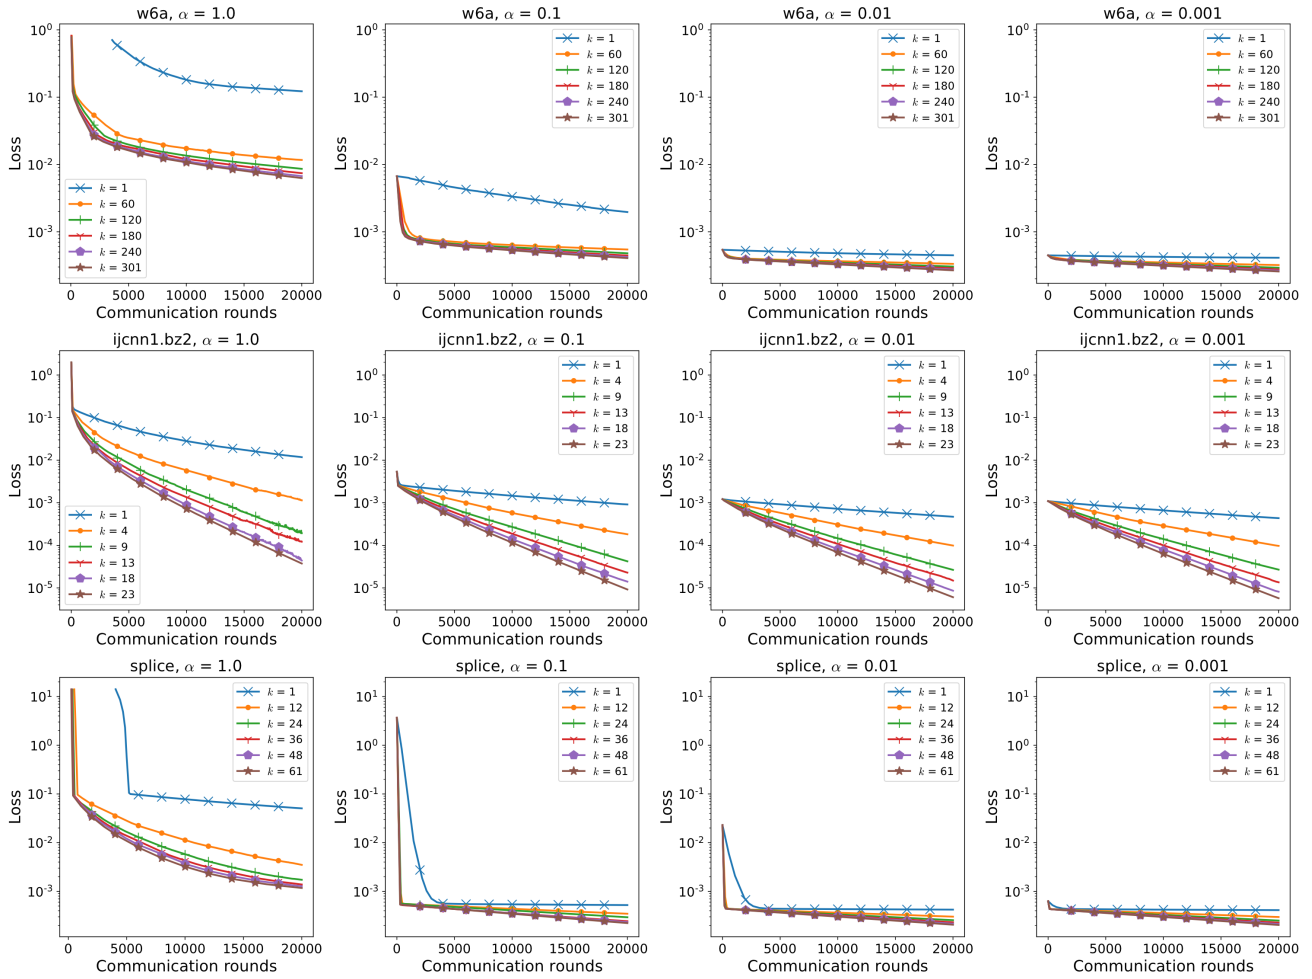


Figure 14: Loss $f(x) - f^*$ vs. # of communication rounds of CGD for unregularized logistic regression problem for four datasets, k is a sparsification parameter of Random- k compressor.

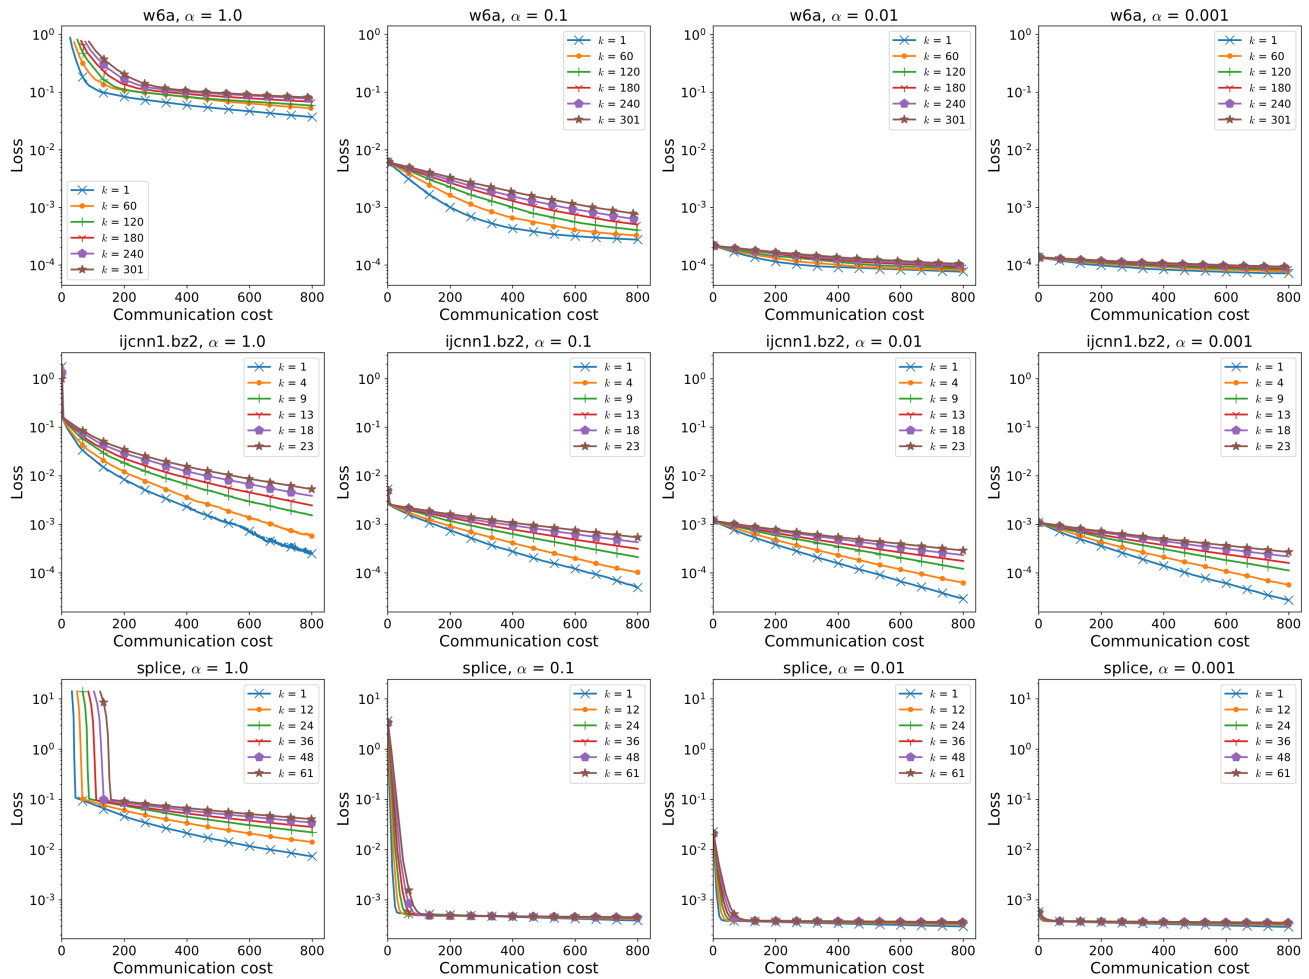


Figure 15: Loss $f(x) - f^*$ vs. communication cost in unit of thousands of float numbers of CGD for unregularized logistic regression problem for four datasets, k is a sparsification parameter of Random- k compressor.