ON THE CONVERGENCE OF FEDPROX WITH EXTRAPOLATION AND INEXACT PROX

Anonymous authors
Paper under double-blind review

ABSTRACT

Enhancing the FedProx federated learning algorithm Li et al. (2020) with server-side extrapolation, Li et al. (2024a) recently introduced the FedExProx method. Their theoretical analysis, however, relies on the assumption that each client computes a certain proximal operator exactly, which is impractical since this is virtually never possible to do in real settings. In this paper, we investigate the behavior of FedExProx without this exactness assumption in the smooth and globally strongly convex setting. We establish a general convergence result, showing that inexactness leads to convergence to a neighborhood of the solution. Additionally, we demonstrate that, with careful control, the adverse effects can be effectively controlled and minimized. By linking inexactness to biased compression Beznosikov et al. (2023), we refine our analysis, highlighting robustness of extrapolation to inexact proximal updates. We also examine the local iteration complexity required by each client to achieved the required level of inexactness using various local optimizers. Our theoretical insights are validated through comprehensive numerical experiments.

1 INTRODUCTION

Distributed optimization is becoming increasingly essential in modern machine learning, especially as models grow more complex. Federated learning (FL), a decentralized approach where multiple clients collaboratively train a shared model while keeping their data locally to preserve privacy, is a key example of this trend (Konečný et al., 2016; McMahan et al., 2017). Often, a central server coordinates the process by aggregating the locally trained models from each client to update the global model without accessing the raw data. The federated average algorithm (FedAvg), introduced by McMahan et al. (2017) and Mangasarian & Solodov (1993), is one of the most popular strategies for tackling federated learning problems. The algorithm comprises three essential components: client sampling, data sampling, and local training. During its execution, the server first samples a subset of clients to participate in the training process for a given round. Each selected client then performs local training using stochastic gradient descent (SGD), with or without random reshuffling, to enhance communication efficiency, as documented by Bubeck et al. (2015); Gower et al. (2019); Moulines & Bach (2011); Sadiev et al. (2024). FedAvg has proven to be highly successful in practice, nevertheless it suffers from client drift when data is heterogeneous (Karimireddy et al., 2020).

Various techniques have been proposed to address the challenges of data heterogeneity, with FedProx, introduced by Li et al. (2020), being one notable example. Rather than having each client perform local SGD rounds, FedProx requires each client to compute a proximal operator locally. Computing the proximal operator can be regarded as an optimization problem that each client can solve locally. Proximal algorithms are advantageous when the proximal operators can be evaluated relatively easily (Parikh et al., 2014). Algorithms based on proximal operators, such as the proximal point method (PPM) (Rockafellar, 1976; Parikh et al., 2014) and its extension to the stochastic setting (SPPM) (Bertsekas, 2011; Asi & Duchi, 2019; Khaled & Jin, 2022; Richtárik & Takác, 2020; Patrascu & Necoara, 2018), offer greater stability against inaccurately specified step sizes, unlike gradient-based methods. PPM was introduced by Martinet (1972) and expanded by Rockafellar (1976). Its extension into the stochastic setting are often used in federated optimization. The stability mentioned is particularly useful when problem-specific parameters, such as the smoothness constant of the objective function, are unknown which renders determining the step size for SGD becomes challenging. Indeed, an excessively large step size in SGD leads to divergence, while a small step size ensures convergence but significantly slows down the training process.

Another approach to mitigating the slowdown caused by heterogeneity is the use of a server step size. Specifically, in FedAvg, a local step size is employed by each client to minimize their individual objectives, while a server step size is used to aggregate the 'pseudo-gradients' obtained from each client (Karimireddy et al., 2020; Reddi et al., 2021). The local step size is set relatively small to mitigate client drift, while the server step size is set larger to avoid slowdowns. However, the small step sizes result in a slowdown during the initial phase of training, which cannot be fully compensated by the large server step size (Jhunjhunwala et al., 2023). Building on the extrapolation technique employed in parallel projection methods to solve the convex feasibility problem (Censor et al., 2001; Combettes, 1997; Necoara et al., 2019), Jhunjhunwala et al. (2023) introduced FedExP as an extension of FedAvg, incorporating adaptive extrapolation as the server step size. Extrapolation involves moving further along the line connecting the most recent iterate, x_k, and the average of the projections of x_k onto different convex sets, $\mathcal{X}_i$, in the parallel projection method, which accelerates the algorithm. Extrapolation is also known as over-relaxation (Rechardson, 1911) in fixed point theory. It is a common technique to effectively accelerate the convergence of fixed point methods including gradient based algorithms and proximal splitting algorithms (Condat et al., 2023; Iutzeler & Hendrickx, 2019). Recently, Li et al. (2024a) shows that the combination of extrapolation with FedProx also results in better complexity bounds. The analysis of the resulting algorithm FedExProx reveals the relationship between the extrapolation parameter and the step size of gradient-based methods with respect to the Moreau envelope associated with the original objective function.[1] However, it relies on the assumption that each proximal operator is solved accurately, which makes it impractical and less advantageous compared to gradient-based algorithms.

1.1 Contributions

Our paper makes the following contributions, please refer to Appendix B for notation details.

- We provide the first analysis of FedExProx under *inexact proximal evaluations* in the convex regime. Unlike prior work Li et al. (2024a), our analysis removes the assumptions of exact proximal solutions and the interpolation condition. By introducing a principled notion of approximation, we establish a general convergence guarantee to a neighborhood of the solution, with a linear rate $\mathcal{O}\left(\frac{L_\gamma(1+\gamma L_{\max})}{\mu}\right)$ that matches the exact case in Li et al. (2024a).
- We show how a refined characterization of inexactness eliminates the neighborhood altogether. When the approximation error is properly bounded and the extrapolation parameter is chosen appropriately, our guarantee ensures exact convergence rather than only neighborhood convergence.
- We extend the analysis by linking inexactness to the framework of biased compression (Beznosikov et al., 2023). This yields an improved rate $\mathcal{O}\left(\frac{L_\gamma(1+\gamma L_{\max})}{\mu-4\varepsilon_2 L_{\max}}\right)$,[2] which ensures convergence to the exact solution under broader admissible bounds on inexactness. Importantly, the optimal extrapolation $1/\gamma L_\gamma$ coincides with the exact-proximal case, showing that extrapolation accelerates convergence once sufficient accuracy is attained, even with inexact proximal evaluations.
- We further analyze how the aforementioned approximations can be obtained by each client. As examples, we provide the local iteration complexity when the client employs gradient descent (GD) or Nesterov's accelerated gradient descent (AGD). Our method proves suitable for a wide range of local solvers such as Adam for the subproblems, offering substantial flexibility in practice. See Table 1 and Table 2 for a detailed comparison of complexities of all relevant quantities.
- Finally, we validate our findings through both numerical and deep learning experiments. The results show that the proposed technique of relative approximation effectively eliminates bias, and that the algorithm consistently outperforms FedProx in all cases, achieving a $2\times$ speedup with zero additional runtime overhead. In practice, the method is simple to implement, works robustly with coarse approximations, and is orthogonal to other acceleration strategies, making it immediately applicable in diverse scenarios.

[1] A tighter convergence guarantee in some cases is obtained by Anyszka et al. (2024).

[2] ε_2 is the relative approximation parameter defined in Definition 4.1. We use $\mathcal{O}(\cdot)$ to hide constant factors and $\tilde{\mathcal{O}}(\cdot)$ when logarithmic factors are omitted.

Table 1: Comparison of FedExProx (Li et al., 2024a) and our proposed inexact versions of the algorithms using different approximations. In the convergence column, we present the rate at which each algorithm converges to either the solution or a neighborhood in the globally μ-convex and interpolated setting to avoids unnecessary complications. Here, L_γ represents the smoothness constant of M^γ as defined before Theorem 3.2. The neighborhood column indicates the size of the neighborhood, the optimal extrapolation column suggests the best α for each algorithm. Constant factors are omitted, K is the total number of iterations, γ is the local step size for the proximal operator, $S(\varepsilon_2)$ defined in Theorem 4.2 is a factor of slow down due to inexactness in $(0, 1]$. For relative approximation, we present the original theory in the third row and then place the sharper analysis in the following row for comparison.

Algorithm	**Convergence**	**Neighborhood**	**Optimal Extrapolation**	**Bound on Inexactness**
FedExProx	$\exp\left(-\frac{K\mu}{L_\gamma(1+\gamma L_{\max})}\right)$	0	$\frac{1}{\gamma L_\gamma}$	NA
(NEW) FedExProx with ε_1 approximation	$\exp\left(-\frac{K\mu}{L_\gamma(1+\gamma L_{\max})}\right)$	$\varepsilon_1\left(\frac{\frac{1}{\gamma}+L_{\max}}{\mu}\right)^2$ [a]	$\frac{1}{4\gamma L_\gamma}$	NA [b]
(NEW) FedExProx with ε_2 relative approximation by biased SGD	$\exp\left(-\frac{K\mu S(\varepsilon_2)}{L_\gamma(1+\gamma L_{\max})}\right)$ [c]	0	$< \frac{1}{\gamma L_\gamma}$	$< \frac{\mu^2}{4L_{\max}^2}$
(NEW) FedExProx with ε_2 relative approximation by biased compression	$\exp\left(-\frac{K(\mu-4\varepsilon_2 L_{\max})}{L_\gamma(1+\gamma L_{\max})}\right)$	0	$\frac{1}{\gamma L_\gamma}$ [d]	$< \frac{\mu}{4L_{\max}}$

[a] Note that when $\varepsilon_1 = 0$, i.e., when the proximal operators exact, the neighborhood diminishes, and we recover the result of FedExProx by Li et al. (2024a).
[b] The convergence guarantee extends to unbounded inexactness, but as inexactness grows, the convergence neighborhood expands proportionally.
[c] Refer to Theorem 4.2 for the definition of $S(\varepsilon_2)$ and the optimal extrapolation parameter. The theory shows that the algorithm remains convergent under inexactness, with performance depending on its magnitude.
[d] Surprisingly, our sharper analysis reveals that the optimal extrapolation parameter in this case remains the same as in the exact setting, highlighting the effectiveness of extrapolation even when the proximal operators are evaluated inexactly.

Table 2: Local iteration complexities for each client using example solvers, such as GD or AGD (Nesterov, 2004). We use the i-th client as an example, where f_i is L_i-smooth and convex.

Algorithm	ε_1 **absolute approximation**	ε_2 **relative approximation**
Gradient descent	$\mathcal{O}\left((1+\gamma L_i)\log\left(\frac{\left\|x_k-\mathrm{prox}_{\gamma f_i}(x_k)\right\|^2}{\varepsilon_1}\right)\right)$ [a]	$\mathcal{O}\left((1+\gamma L_i)\log\left(\frac{1}{\varepsilon_2}\right)\right)$
Accelerate gradient descent	$\mathcal{O}\left(\sqrt{1+\gamma L_i}\log\left(\frac{\left\|x_k-\mathrm{prox}_{\gamma f_i}(x_k)\right\|^2}{\varepsilon_1}\right)\right)$	$\mathcal{O}\left(\sqrt{1+\gamma L_i}\log\left(\frac{1}{\varepsilon_2}\right)\right)$

[a] We can easily provide an upper bound of $\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2$ for determining the number of local computations needed.

1.2 Related work

Arguably, stochastic gradient descent (SGD) (Robbins & Monro, 1951; Ghadimi & Lan, 2013; Gower et al., 2019; Gorbunov et al., 2020) remains one of the foundational algorithm in the field of machine learning. One can simply formulate it as $x_{k+1} = x_k - \eta \cdot g(x_k)$, where $\eta > 0$ is a scalar step size, $g(x_k)$ is a possibly stochastic estimator of the true gradient $\nabla f(x_k)$. Various extensions have been proposed since its introduction, including compressed gradient descent (CGD) (Khirirat et al., 2018) SGD with matrix step size (Li et al., 2024b) and variance reduction (Gower et al., 2020; Johnson & Zhang, 2013; Gorbunov et al., 2021; Tyurin & Richtárik, 2024; Li et al., 2023). Gower et al. (2019) presented a framework for analyzing SGD with unbiased gradient estimator in the convex case based

on expected smoothness. However, in practice, sometimes the gradient estimator could be biased, examples include SGD with sparsified or delayed update (Alistarh et al., 2018; Recht et al., 2011). Beznosikov et al. (2023) examined biased updates in the context of compressed gradient descent. Demidovich et al. (2024) provides a framework for analyzing SGD with biased gradient estimators in the non-convex setting.

Proximal point method (PPM) was originally introduced as a method to solve variational inequalities (Martinet, 1972; Rockafellar, 1976). The transition to the stochastic case, driven by the need to efficiently address large-scale optimization problems, leads to the development of SPPM. Due to its stability and advantage over the gradient based methods, it has been extensively studied, as documented by (Patrascu & Necoara, 2018; Bianchi, 2016; Bertsekas, 2011; Gruntkowska et al., 2025). For proximal algorithms to be practical, it is commonly assumed that the proximal operator can be solved efficiently, such as in cases where a closed-form solution is available. However, in large-scale machine learning models, it is rarely possible to find such a solution in closed form. To address this issue, most proximal algorithms assume that only an approximate solution is obtained, achieving a certain level of accuracy (Khaled & Jin, 2022; Sadiev et al., 2022; Karagulyan et al., 2024). Various notions of inexactness are employed, depending on the assumptions made, the properties of the objective, and the availability of algorithms capable of efficiently finding such approximations.

Moreau envelope was first introduced to handle non-smooth functions by Moreau (1965). It is also known as the Moreau-Yosida regularization. The use of the Moreau envelope as an analytical tool to analyze proximal algorithms is not novel. Ryu & Boyd (2014) noted that running a proximal algorithm on the objective is equivalent to applying gradient methods to its Moreau envelope. Davis & Drusvyatskiy (2019) analyzed stochastic proximal point method (SPPM) for weakly convex and Lipschitz functions based on this finding. Recently, Li et al. (2024a) provided an analysis of FedProx with server-side step size in the convex case, based on the reformulation of the problem using the Moreau envelope. The role of the Moreau envelope extends beyond analyzing proximal algorithms; it has also been applied in the contexts of personalized federated learning (T Dinh et al., 2020) and meta-learning (Mishchenko et al., 2023). The mathematical properties of the Moreau envelope are relatively well understood, as documented by Jourani et al. (2014); Planiden & Wang (2019; 2016).

Projection methods initially emerged as an effective tool for solving systems of linear equations or inequalities (Kaczmarz, 1937) and were later generalized to solve the convex feasibility problem (Combettes, 1997). The parallel version of this approach involves averaging the projections of the current iterates onto all existing convex sets $\mathcal{X}_i$ to obtain the next iterate, a process that is empirically known to be accelerated by extrapolation. Numerous heuristic rules have been proposed to adaptively set the extrapolation parameter, such as those by Bauschke et al. (2006) and Pierra (1984). Only recently, the mechanism behind constant extrapolation was uncovered by Necoara et al. (2019), who developed the corresponding theoretical framework. Additionally, Li et al. (2024a) provides explanations for the effectiveness of adaptive rules, revealing the connection between the extrapolation parameter and the step size of SGD when using the Moreau envelope as the global objective.

2 Mathematical background

In this work, we are interested in the distributed optimization problem which is formulated in the following finite-sum form

$$\min_{x\in\mathbb{R}^d}\left\{f(x) := \frac{1}{n}\sum_{i=1}^{n} f_i(x)\right\}, \tag{1}$$

where $x \in \mathbb{R}^d$ is the model, n is the number of devices, $f : \mathbb{R}^d \mapsto \mathbb{R}$ is global objective, each $f_i : \mathbb{R}^d \mapsto \mathbb{R}$ is the empirical risk of model x associated with the i-th client. Each $f_i(x)$ often has the form

$$f_i(x) := \mathbb{E}_{\xi\sim\mathcal{D}_i}\left[l\left(x,\xi\right)\right], \tag{2}$$

where the loss function $l\left(x,\xi\right)$ represents the loss of model x on data point ξ over the training data $\mathcal{D}_i$ owned by client $i \in [n] := \{1, 2, \ldots, n\}$. We first give the definitions for the proximal operator and Moreau envelope, which we will be using in our analysis.

Definition 2.1 (Proximal operator)**.** The proximal operator of an extended real-valued function $\phi : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$ with step size $\gamma > 0$ and center $x \in \mathbb{R}^d$ is defined as

$$\operatorname{prox}_{\gamma\phi}(x) := \arg\min_{z\in\mathbb{R}^d} \left\{ \phi(z) + \frac{1}{2\gamma}\|z - x\|^2 \right\}.$$

It is well-known that for any proper, closed, and convex function ϕ, the proximal operator with any $\gamma > 0$ returns a singleton.

Definition 2.2 (Moreau envelope)**.** The Moreau envelope of an extended real-valued function $\phi : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$ with step size $\gamma > 0$ and center $x \in \mathbb{R}^d$ is defined as

$$M_\phi^\gamma(x) := \min_{z\in\mathbb{R}^d} \left\{ \phi(z) + \frac{1}{2\gamma}\|z - x\|^2 \right\}.$$

By the definition of Moreau envelope, it is easy to see that

$$M_\phi^\gamma(x) = \phi\left(\operatorname{prox}_{\gamma\phi}(x)\right) + \frac{1}{2\gamma}\left\|x - \operatorname{prox}_{\gamma\phi}(x)\right\|^2. \tag{3}$$

Not only are their function values related, but for any proper, closed, and convex function ϕ, the Moreau envelope is differentiable, specifically, we have:

$$\nabla M_\phi^\gamma(x) = \frac{1}{\gamma}\left(x - \operatorname{prox}_{\gamma\phi}(x)\right). \tag{4}$$

The above identity indicates that ϕ and M_ϕ^γ are intrinsically related, which plays a key role in our analysis. We also need the following assumptions on f and f_i to carry out our analysis.

Assumption 2.3. The function $f_i : \mathbb{R}^d \mapsto \mathbb{R}$ in (1) is differentiable and bounded from below for all $i \in [n]$.

Assumption 2.4. The function $f_i : \mathbb{R}^d \mapsto \mathbb{R}$ is convex for all $i \in [n]$. This means that for each f_i,

$$0 \le f_i(x) - f_i(y) - \langle \nabla f_i(y), x - y\rangle, \quad \forall x, y \in \mathbb{R}^d. \tag{5}$$

Assumption 2.5. The function $f_i : \mathbb{R}^d \mapsto \mathbb{R}$ is L_i-smooth, $L_i > 0$ for all $i \in [n]$. This means that for each f_i,

$$f_i(x) - f_i(y) - \langle \nabla f_i(y), x - y\rangle \le \frac{L_i}{2}\|x - y\|^2, \quad \forall x, y \in \mathbb{R}^d. \tag{6}$$

We will use $L_{\max}$ to denote $\max_{i\in[n]} L_i$.

Assumption 2.6. The global objective function f is μ-strongly convex, $\mu > 0$. That is

$$f(x) - f(y) - \langle \nabla f(y), x - y\rangle \ge \frac{\mu}{2}\|x - y\|^2, \quad \forall x, y \in \mathbb{R}^d.$$

These are all standard assumptions commonly used in convex optimization. For comparison with our baseline FedExProx, we optionally adopt the following interpolation condition. Unlike FedExProx, which requires interpolation for its analysis, our framework *does not depend on it*. We only optionally invoke the following condition to simplify formulas and enable a fair comparison:

Assumption 2.7. There exists $x_\star \in \mathbb{R}^d$ such that $\nabla f_i(x_\star) = 0$ for all $i \in [n]$.

This situation commonly arises in modern deep learning, where the number of parameters d far exceeds the number of training samples (Arora et al., 2019; Montanari & Zhong, 2022). The assumption also has a natural interpretation in convex feasibility problems solved by parallel projection methods (7), where the intersection of all convex sets $\mathcal{X}_i$ is assumed to be non-empty, which is precisely the interpolation assumption of each f_i being the indicator function of $\mathcal{X}_i$.

$$x_{k+1} = \frac{1}{n}\sum_{i=1}^{n} \Pi_{\mathcal{X}_i}(x_k). \tag{7}$$

For such projection methods, extrapolation is known to improve convergence both theoretically and practically (Necoara et al., 2019). Since $\operatorname{prox}_{\gamma f_i}(x_k)$ can be interpreted as a projection onto a level set of f_i, it is natural to expect that extrapolation would be effective under the same assumption. Our theory shows that this beneficial effect of extrapolation is not confined to the exact interpolated setting. We first present our algorithm as Algorithm 1. In the following sections, we provide the analysis of this algorithm under different definitions of inexactness, respectively in Section 3 and Section 4. Details on how these inexactness levels can be achieved by each client are provided in Section 5. Finally, numerical and deep learning experiments validating our results are presented in Section 6.

Algorithm 1 Inexact FedExProx

1: **Parameters:** extrapolation parameter $\alpha_k = \alpha > 0$, step size for the proximal operator $\gamma > 0$, starting point $x_0 \in \mathbb{R}^d$, number of clients n, total number of iterations K, proximal solution accuracy $\varepsilon \geq 0$.
2: **for** $k = 0, 1, 2 \ldots K-1$ **do**
3: The server broadcasts the current iterate x_k to each client
4: Each client computes an ε approximation of the solution $\tilde{x}_{i,k+1} \simeq \operatorname{prox}_{\gamma f_i}(x_k)$, and sends it back to the server
5: The server computes

$$x_{k+1} = x_k + \alpha_k \left(\frac{1}{n} \sum_{i=1}^{n} \tilde{x}_{i,k+1} - x_k \right). \tag{8}$$

6: **end for**

3 Absolute approximation in distance

As previously noted, we assume that each proximal operator is solved only approximately, and thus we need to quantify the level of inexactness. For the i-th client, the exact computation corresponds to the local subproblem

$$\min_{z \in \mathbb{R}^d} A_{k,i}^{\gamma}(z) := f_i(z) + \frac{1}{2\gamma} \|z - x_k\|^2,$$

where x_k is the current iterate and $\gamma > 0$ is the proximal step size. Because each f_i is convex, the objective $A_{k,i}^{\gamma}(z)$ is $\frac{1}{\gamma}$-strongly convex and admits a unique minimizer, namely $\operatorname{prox}_{\gamma f_i}(x_k)$. When this minimizer is computed only approximately, a natural measure of inexactness is the squared distance between the returned solution and the true minimizer, which motivates the following definition.

Definition 3.1 (Absolute approximation)**.** Given a proper, closed and convex function $\phi : \mathbb{R}^d \mapsto \mathbb{R}$, and a step size $\gamma > 0$, we say that a point $y \in \mathbb{R}^d$ is an ε_1-approximation of $\operatorname{prox}_{\gamma\phi}(x)$, if for some $\varepsilon_1 \geq 0$,

$$\left\| y - \operatorname{prox}_{\gamma\phi}(x) \right\|^2 \leq \varepsilon_1. \tag{9}$$

To analyze Algorithm 1, we transform the update rule given in (8) in the following way,

$$\begin{aligned} x_{k+1} &= x_k + \alpha_k \left(\frac{1}{n} \sum_{i=1}^{n} \left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k) \right) + \frac{1}{n} \sum_{i=1}^{n} \operatorname{prox}_{\gamma f_i}(x_k) - x_k \right) \\ &\overset{(4)}{=} x_k - \alpha_k \cdot g(x_k), \end{aligned} \tag{10}$$

where

$$g(x_k) := \frac{1}{n} \sum_{i=1}^{n} \gamma \nabla M_{f_i}^{\gamma}(x_k) + \frac{1}{n} \sum_{i=1}^{n} \left(\operatorname{prox}_{\gamma f_i}(x_k) - \tilde{x}_{i,k+1} \right) = \text{Gradient} + \text{Bias}. \tag{11}$$

This reformulation shows that Algorithm 1 can be viewed as SGD on the objective $\gamma M^{\gamma}(x) := \frac{1}{n} \sum_{i=1}^{n} \gamma M_{f_i}^{\gamma}(x)$, but with a biased gradient estimator. Unlike the unbiased case, biased SGD is less well understood, which makes extending the original FedExProx theory that corresponds to unbiased SGD to the inexact setting nontrivial. Nevertheless, we establish the following convergence guarantee, stated as a simplified version of Theorem F.1 in Appendix F, specialized to the interpolation regime.

Theorem 3.2. *Assume Assumption 2.3, 2.4, 2.5, 2.6 and additionally, Assumption 2.7 hold. If each client computes a ε_1-absolute approximation $\tilde{x}_{i,k+1}$ of $\operatorname{prox}_{\gamma f_i}(x_k)$ at every iteration, such that $\left\| \tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k) \right\|^2 \leq \varepsilon_1$. We have the following convergence guarantee for Algorithm 1: For extrapolation parameter $\alpha_k = \alpha$ satisfying $0 < \alpha \leq \frac{1}{4} \cdot \frac{1}{\gamma L_\gamma}$, where γ is the step size of the proximal operator, L_γ is the smoothness constant of M^{γ}. The last iterate x_K satisfy*

$$\mathcal{E}_K \leq \left(1 - \frac{\alpha\gamma\mu}{8(1 + \gamma L_{\max})} \right)^K \mathcal{E}_0 + \frac{4\varepsilon_1 (1 + \gamma L_{\max})}{\mu} \cdot \left(2\alpha L_\gamma + \frac{1}{\gamma} \right),$$

where $\mathcal{E}_k = \gamma M^\gamma(x_k) - \gamma M^\gamma_{\inf}$. *Specifically, when choosing* $\alpha = \frac{1}{4} \cdot \frac{1}{\gamma L_\gamma}$, *we have*

$$\Delta_K \leq \left(1 - \frac{\mu}{32 L_\gamma (1+\gamma L_{\max})}\right)^K \frac{L_\gamma (1+\gamma L_{\max})}{\mu} \cdot \Delta_0 + 12\varepsilon_1 \cdot \left(\frac{1/\gamma + L_{\max}}{\mu}\right)^2,$$

where $\Delta_K = \|x_K - x_\star\|^2$, $x_\star$ *is a minimizer of* f.

By incorporating the optional interpolation condition, Fact C.7 ensures that M^γ and f share the same minimizer, which we denote by $x_\star$. In this case, the algorithm enjoys convergence to a neighborhood around $x_\star$, whose size depends on ε_1 and the choice of γ. Smaller values of γ make the local subproblems easier to solve, though at the cost of a larger neighborhood. This can be understood intuitively: A smaller γ means less progress is made per iteration, leading to a larger accumulated error as the total number of iterations increases. The convergence guarantee remains valid for arbitrarily large ε_1, highlighting the generality of the theory. As ε_1 increases, the neighborhood expands proportionally, reflecting the natural tradeoff between approximation accuracy and convergence precision. When $\varepsilon_1 = 0$, the neighborhood diminishes, and we obtain an iteration complexity of $\tilde{\mathcal{O}}\left(\frac{L_\gamma(1+\gamma L_{\max})}{\mu}\right)$[3], which recovers the result of Li et al. (2024a) up to a constant factor. The optimal constant extrapolation parameter is $\alpha_\star = 1/4\gamma L_\gamma$, which recovers the result of Li et al. (2024a) up to a constant factor.

4 Relative approximation in distance

Theorem 3.2 offers a general theoretical framework for understanding the behavior of Algorithm 1. However, a key challenge with Algorithm 1 which utilizes inexact proximal solutions that satisfy Definition 3.1, is that, unless the proximal operators are solved exactly, convergence will always be limited to a neighborhood of the solution. The underlying reason is that, as the algorithm progresses, the gradient term in the gradient estimator $g(x_k)$ diminishes, whereas the bias term remains unchanged. Building on this observation, we propose employing a different type of approximation, specifically an approximation in relative distance, as defined below.

Definition 4.1 (Relative approximation)**.** Given a convex function $\phi : \mathbb{R}^d \mapsto \mathbb{R}$ and a stepsize $\gamma > 0$, we say that a point $y \in \mathbb{R}^d$ is a ε_2-relative approximation of $\operatorname{prox}_{\gamma\phi}(x)$, if for some $\varepsilon_2 \in [0, 1)$,

$$\left\|y - \operatorname{prox}_{\gamma\phi}(x)\right\|^2 \leq \varepsilon_2 \cdot \left\|x - \operatorname{prox}_{\gamma\phi}(x)\right\|^2. \tag{12}$$

The same concept of approximations have been studied and applied in prior research, as exemplified by Solodov & Svaiter (1999). We impose the requirement that the coefficient ε_2 be less than 1 to ensure that the next iterate is no worse than the current one. As we can observe, if the approximation satisfies Definition 4.1, both the gradient and the bias diminish as the algorithm progresses, ensuring convergence to the exact solution. Builting upon the theory of biased SGD, we can obtain the following theorem.

Theorem 4.2. *Assume all the assumptions mentioned in Theorem 3.2 also hold here. If each client only computes a* ε_2*-relative approximation* $\tilde{x}_{i,k+1}$ *in distance with* $\varepsilon_2 < \mu^2/4L_{\max}^2$, *such that* $\left\|\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k)\right\|^2 \leq \varepsilon_2 \cdot \left\|x_k - \operatorname{prox}_{\gamma f_i}(x_k)\right\|^2$ *and we are running Algorithm 1 with* $\alpha_k = \alpha$ *satisfying*

$$0 < \alpha \leq \frac{1}{\gamma L_\gamma} \cdot \frac{\mu - 2\sqrt{\varepsilon_2} L_{\max}}{\mu + 4\sqrt{\varepsilon_2} L_{\max} + 4\varepsilon_2 L_{\max}}.$$

Then the iterates generated by Algorithm 1 satisfies

$$\mathcal{E}_K \leq \left(1 - \alpha \cdot \frac{\gamma\left(\mu - 2\sqrt{\varepsilon_2} L_{\max}\right)}{4(1+\gamma L_{\max})}\right)^K \mathcal{E}_0.$$

Specifically, if we choose the largest α *possible, we have*

$$\Delta_K \leq \left(1 - \frac{\mu}{4L_\gamma(1+\gamma L_{\max})} \cdot S(\varepsilon_2)\right)^K \cdot \frac{L\gamma(1+\gamma L_{\max})}{\mu}\Delta_0,$$

[3]We leave out the log factor in $\tilde{\mathcal{O}}(\cdot)$ notation.

where $S(\varepsilon_2) := \frac{(\mu-2\sqrt{\varepsilon_2}L_{\max})\left(1-2\sqrt{\varepsilon_2}\frac{L_{\max}}{\mu}\right)}{\mu+4\sqrt{\varepsilon_2}L_{\max}+4\varepsilon_2 L_{\max}}$ *satisfies* $0 < S(\varepsilon_2) \leq 1$ *is the factor of slowing down due to inexact proximal operator evaluation.*

Observe that when $\varepsilon_2 = 0$, meaning the proximal operators are solved exactly, the optimal extrapolation is $\alpha = \frac{1}{\gamma L_\gamma}$ and the iteration complexity is $\tilde{\mathcal{O}}\left(\frac{L_\gamma(1+\gamma L_{\max})}{\mu}\right)$. This coincides with the result established in Li et al. (2024a). For inexact solutions, larger values of ε_2 reduce both α and $S(\varepsilon_2)$, illustrating the natural tradeoff between approximation accuracy and convergence speed. Here, ε_2 is required to satisfy $\varepsilon_2 = c \cdot \frac{\mu^2}{4L_{\max}^2}$ with $c < 1$, ensuring that the approximation remains within a controlled range.

Definition 4.1 is closely related to the notion of compression in SGD. In particular, since $x_k - \mathrm{prox}_{\gamma f_i}(x_k) = \gamma \nabla M_{f_i}^{\gamma}(x_k)$, the term $\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)$ can be viewed as a compressed estimator of the gradient, i.e., $\mathcal{C}(\gamma \nabla M_{f_i}^{\gamma}(x_k))$. This indicates that Algorithm 1 with approximation satisfying Definition 4.1 can be viewed as compressed gradient descent with biased compressor. We generalize the framework of Beznosikov et al. (2023) to derive the following convergence guarantee, presented as a simplified form of Theorem F.5 in Appendix F and specialized to the interpolation regime for clarity of presentation.

Theorem 4.3. *Assume all assumptions of Theorem 3.2 hold. Let the approximation* $\tilde{x}_{i,k+1}$ *all satisfies Definition 4.1 with* $\varepsilon_2 < \mu/4L_{\max}$*, that is* $\left\|\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 \leq \varepsilon_2 \cdot \left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2$*. If we are running Algorithm 1 with* $\alpha_k = \alpha \in (0, \frac{1}{\gamma L_\gamma}]$*, we have the iterates produced by it satisfying*

$$\mathcal{E}_K \leq \left(1 - \left(1 - \frac{4\varepsilon_2 L_{\max}}{\mu}\right) \cdot \frac{\gamma\mu}{4\left(1+\gamma L_{\max}\right)} \cdot \alpha\right)^K \mathcal{E}_0.$$

specifically, if we take the largest extrapolation ($\alpha = \frac{1}{\gamma L_\gamma} > 1$) *possible, we have*

$$\Delta_K \leq \left(1 - \left(1 - \frac{4\varepsilon_2 L_{\max}}{\mu}\right) \cdot \frac{\mu}{4L_\gamma\left(1+\gamma L_{\max}\right)}\right)^K \cdot \frac{L_\gamma\left(1+\gamma L_{\max}\right)}{\mu}\Delta_0.$$

The convergence guarantee obtained in this way is sharper, indeed, Theorem 4.3 suggests that as long as $\varepsilon_2 < \mu/4L$, we are able to pick $\alpha = 1/\gamma L_\gamma$[4]which is the optimal extrapolation for the exact case given in Li et al. (2024a). Notably, this implies that extrapolation remains effective, regardless of inexact proximal operator evaluations. In line with Theorem 4.2, our result shows that the method remains convergent under approximation, and notably, when $\varepsilon_2 = 0$, it recovers the guarantee of Li et al. (2024a).

5 Achieving a sufficient level of Inexactness

To comprehend the overall complexity of Algorithm 1 in theory, it is essential to examine whether the inexactness in evaluating the proximal operators can be effectively achieved. Since each $\mathrm{prox}_{\gamma f_i}(x_k)$ is computed locally by the corresponding client, the client has access to all the necessary data points for the computation. As an illustrative example, each client may perform GD.

Theorem 5.1 (Local computation via GD)**.** *Assume Assumption 2.3, 2.4 and 2.5 hold. The iteration complexity for the* i*-th client to provide an approximation using GD in the* k*-th iteration with local step size* $\eta_i = \gamma/(1+\gamma L_i)$*, satisfying Definition 3.1 is* $\mathcal{O}\left((1+\gamma L_i)\log\left(\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2/\varepsilon_1\right)\right)$*, and for Definition 4.1, it is* $\mathcal{O}\left((1+\gamma L_i)\log\left(1/\varepsilon_2\right)\right)$.

There are no constraints on ε_1, and note that $\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 \leq \left\|\gamma \nabla f(x_k)\right\|^2$ by (43), it is straightforward to adjust GD to optimize the approximation. For ε_2, we require $\varepsilon_2 < \mu/4L_{\max}$. In practice, this condition is easy to satisfy, since ε_2 can simply be chosen small enough. The complexity bounds also indicate that as the local step size γ increases, it becomes more challenging to compute the approximation. Alternatively, other algorithms can be employed to find such an approximation. For instance, by leveraging the structure in (2), SGD can be used as a local solver for the proximal

[4]It is shown in Li et al. (2024a) that $1/\gamma L_\gamma > 1$, which justifies why α is called the extrapolation parameter.

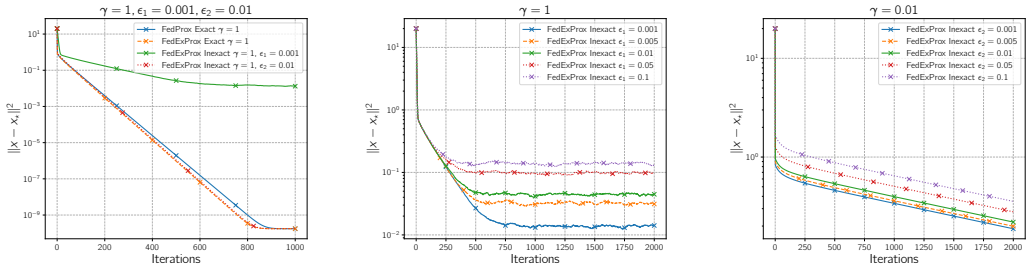

(a) (b) (c)

Figure 1: Comparison of FedProx, FedExProx with exact proximal evaluations, FedExProx absolute and relative inexactness. Figure (a) presents a comparison of the four algorithms discussed above. Figure (b) illustrates the impact of different values of ε_1 on FedExProx with absolute approximation. Figure (c) demonstrates how varying values of ε_2 affect FedExProx with relative approximation.

operator when computational resources are limited. Optimizers such as Adam are also suitable choices in this setting. Another example is to use accelerated gradient descent (AGD) (Nesterov, 2004), which yields improved iteration complexity for each client.

Theorem 5.2 (Local computation via AGD)**.** *Assume all assumptions mentioned in Theorem 5.1 hold. The iteration complexities for the i-th client to provide an approximation in the k-the iteration using AGD with local step size $\eta_i = \frac{\gamma}{1+\gamma L_i}$ and momentum parameter $\alpha_i = \frac{\sqrt{1+\gamma L_i}-1}{\sqrt{1+\gamma L_i}+1}$, satisfying Definition 3.1, Definition 4.1 are, respectively,*

$$\mathcal{O}\left(\sqrt{1+\gamma L_i}\log\left(\frac{(1+\gamma L_i)\cdot\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2}{\varepsilon_1}\right)\right); \quad \mathcal{O}\left(\sqrt{1+\gamma L_i}\log\left(\frac{1+\gamma L_i}{\varepsilon_2}\right)\right)$$

6 Experiments

We refer the readers to Appendix J for the details of the settings and the corresponding experiments. Finally, we provide empirical evidence to support our theoretical findings. We refer the readers to Appendix J for the details of deep learning experiments. See Figure 1 for an numerical experiment in a simple illustrative setting. In Figure 1 (a), we compare the performance of FedProx, exact FedExProx, ε_1-inexactFedExProx and ε_2-inexact FedExProx. Interestingly, FedExProx with relative approximations delivers strong performance when ε_2 is appropriately selected, and in some cases, it even outperforms FedProx with exact updates. This demonstrates the effectiveness of server extrapolation despite inexact proximal evaluations.

In Figure 1 (b), we present a comparison of FedExProx with absolute approximations under different levels of inexactness ε_1. In all cases, the algorithm converges to a neighborhood of the solution, with larger inexactness resulting in a larger neighborhood.

In Figure 1 (c), we compare FedExProx with relative approximations under varying levels of inexactness ε_2. In all cases, the algorithm converges to the exact solution, validating the effectiveness of relative approximation in eliminating bias.

7 Reproducibility Statement

All assumptions are stated clearly in the main text, and each theorem is presented together with its corresponding assumptions. Complete proofs of the theoretical results are provided in the appendix. Experimental details are described in Appendix J, where we also include a link to an anonymous code repository to facilitate reproducibility.

References

Dan Alistarh, Torsten Hoefler, Mikael Johansson, Nikola Konstantinov, Sarit Khirirat, and Cédric Renggli. The convergence of sparsified gradient methods. *Advances in Neural Information*

Processing Systems, 31, 2018.

Mihai Anitescu. Degenerate nonlinear programming with a quadratic growth condition. *SIAM Journal on Optimization*, 10(4):1116–1135, 2000.

Wojciech Anyszka, Kaja Gruntkowska, Alexander Tyurin, and Peter Richtárik. Tighter performance theory of fedexprox. *arXiv preprint arXiv:2410.15368*, 2024.

Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pp. 322–332. PMLR, 2019.

Hilal Asi and John C Duchi. Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019.

Heinz H Bauschke, Patrick L Combettes, and Serge G Kruk. Extrapolation algorithm for affine-convex feasibility problems. *Numerical Algorithms*, 41:239–274, 2006.

Amir Beck. *First-order methods in optimization*. SIAM, 2017.

Dimitri P Bertsekas. Incremental proximal methods for large scale convex optimization. *Mathematical Programming*, 129(2):163–195, 2011.

Aleksandr Beznosikov, Samuel Horváth, Peter Richtárik, and Mher Safaryan. On biased compression for distributed learning. *Journal of Machine Learning Research*, 24(276):1–50, 2023.

Pascal Bianchi. Ergodic convergence of a stochastic proximal point algorithm. *SIAM Journal on Optimization*, 26(4):2235–2260, 2016.

Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

Y Censor, T Elfving, and GT Herman. Averaging strings of sequential iterations for convex feasibility problems. In *Studies in Computational Mathematics*, volume 8, pp. 101–113. Elsevier, 2001.

Patrick L Combettes. Convex set theoretic image recovery by extrapolated iterations of parallel subgradient projections. *IEEE Transactions on Image Processing*, 6(4):493–506, 1997.

Laurent Condat, Daichi Kitahara, Andrés Contreras, and Akira Hirabayashi. Proximal splitting algorithms for convex optimization: A tour of recent advances, with new twists. *SIAM Review*, 65 (2):375–435, 2023.

Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019.

Yury Demidovich, Grigory Malinovsky, Igor Sokolov, and Peter Richtárik. A guide through the zoo of biased SGD. *Advances in Neural Information Processing Systems*, 36, 2024.

Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.

Pinghua Gong and Jieping Ye. Linear convergence of variance-reduced stochastic gradient without strong convexity. *arXiv preprint arXiv:1406.1102*, 2014.

Eduard Gorbunov, Filip Hanzely, and Peter Richtárik. A unified theory of SGD: Variance reduction, sampling, quantization and coordinate descent. In *International Conference on Artificial Intelligence and Statistics*, pp. 680–690. PMLR, 2020.

Eduard Gorbunov, Konstantin P Burlachenko, Zhize Li, and Peter Richtárik. MARINA: Faster nonconvex distributed learning with compression. In *International Conference on Machine Learning*, pp. 3788–3798. PMLR, 2021.

Robert M Gower, Mark Schmidt, Francis Bach, and Peter Richtárik. Variance-reduced methods for machine learning. *Proceedings of the IEEE*, 108(11):1968–1983, 2020.

Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. SGD: General analysis and improved rates. In *International Conference on Machine Learning*, pp. 5200–5209. PMLR, 2019.

Kaja Gruntkowska, Hanmin Li, Aadi Rane, and Peter Richtárik. The ball-proximal (=" broximal") point method: a new algorithm, convergence theory, and applications. *arXiv preprint arXiv:2502.02002*, 2025.

Franck Iutzeler and Julien M Hendrickx. A generic online acceleration scheme for optimization algorithms via relaxation and inertia. *Optimization Methods and Software*, 34(2):383–405, 2019.

Divyansh Jhunjhunwala, Shiqiang Wang, and Gauri Joshi. FedExP: Speeding up federated averaging via extrapolation. In *International Conference on Learning Representations*, 2023.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26, 2013.

Abderrahim Jourani, Lionel Thibault, and Dariusz Zagrodny. Differential properties of the moreau envelope. *Journal of Functional Analysis*, 266(3):1185–1237, 2014.

Stefan Kaczmarz. Approximate solution of systems of linear equations. *International Journal of Control*, 57(6):1269–1271, 1937.

Avetik Karagulyan, Egor Shulgin, Abdurakhmon Sadiev, and Peter Richtárik. Spam: Stochastic proximal point method with momentum variance reduction for non-convex cross-device federated learning. *arXiv preprint arXiv:2405.20127*, 2024.

Sai Praneeth Karimireddy, Quentin Rebjock, Sebastian Stich, and Martin Jaggi. Error feedback fixes signsgd and other gradient compression schemes. In *International Conference on Machine Learning*, pp. 3252–3261. PMLR, 2019.

Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pp. 5132–5143. PMLR, 2020.

Ahmed Khaled and Chi Jin. Faster federated optimization under second-order similarity. In *The Eleventh International Conference on Learning Representations*, 2022.

Sarit Khirirat, Hamid Reza Feyzmahdavian, and Mikael Johansson. Distributed learning with compressed gradients. *arXiv preprint arXiv:1806.06573*, 2018.

Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 8, 2016.

Hanmin Li, Avetik Karagulyan, and Peter Richtárik. Variance reduced distributed non-convex optimization using matrix stepsizes. *arXiv preprint arXiv:2310.04614*, 2023.

Hanmin Li, Kirill Acharya, and Peter Richtárik. The power of extrapolation in federated learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024a.

Hanmin Li, Avetik Karagulyan, and Peter Richtárik. Det-CGD: Compressed gradient descent with matrix stepsizes for non-convex optimization. In *International Conference on Learning Representations*, 2024b.

Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020.

Ji Liu and Stephen J Wright. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.

Olvi L Mangasarian and Mikhail V Solodov. Backpropagation convergence via deterministic nonmonotone perturbed minimization. *Advances in Neural Information Processing Systems*, 6, 1993.

Bernard Martinet. *Algorithmes pour la résolution de problèmes d'optimisation et de minimax*. PhD thesis, Université Joseph-Fourier-Grenoble I, 1972.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pp. 1273–1282. PMLR, 2017.

Konstantin Mishchenko, Slavomir Hanzely, and Peter Richtárik. Convergence of first-order algorithms for meta-learning with Moreau envelopes. *arXiv preprint arXiv:2301.06806*, 2023.

Andrea Montanari and Yiqiao Zhong. The interpolation phase transition in neural networks: Memorization and generalization under lazy training. *The Annals of Statistics*, 50(5):2816–2847, 2022.

Jean-Jacques Moreau. Proximité et dualité dans un espace Hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299, 1965.

Eric Moulines and Francis Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. *Advances in Neural Information Processing Systems*, 24, 2011.

Ion Necoara, Peter Richtárik, and Andrei Patrascu. Randomized projection methods for convex feasibility: Conditioning and convergence rates. *SIAM Journal on Optimization*, 29(4):2814–2852, 2019.

Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*. Kluwer Academic Publishers, 2004.

Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1 (3):127–239, 2014.

Andrei Patrascu and Ion Necoara. Nonasymptotic convergence of stochastic proximal point methods for constrained convex optimization. *Journal of Machine Learning Research*, 18(198):1–42, 2018.

Guy Pierra. Decomposition through formalization in a product space. *Mathematical Programming*, 28:96–115, 1984.

Chayne Planiden and Xianfu Wang. Strongly convex functions, Moreau envelopes, and the generic nature of convex functions with strong minimizers. *SIAM Journal on Optimization*, 26(2):1341–1364, 2016.

Chayne Planiden and Xianfu Wang. Proximal mappings and Moreau envelopes of single-variable convex piecewise cubic functions and multivariable gauge functions. *Nonsmooth Optimization and Its Applications*, pp. 89–130, 2019.

Boris T Polyak. Gradient methods for solving equations and inequalities. *USSR Computational Mathematics and Mathematical Physics*, 4(6):17–32, 1964.

LF Rechardson. The approximate arithmetical solution by finite difference of physical problems involving differential equations, with an application to the stresses in a masonary dam. *R. Soc. London Phil. Trans. A*, 210:307–357, 1911.

Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *Advances in Neural Information Processing Systems*, 24, 2011.

Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečnỳ, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. *International Conference on Learning Representations*, 2021.

Peter Richtárik and Martin Takác. Stochastic reformulations of linear systems: algorithms and convergence theory. *SIAM Journal on Matrix Analysis and Applications*, 41(2):487–524, 2020.

Peter Richtárik, Igor Sokolov, and Ilyas Fatkhullin. Ef21: A new, simpler, theoretically better, and practically faster error feedback. *Advances in Neural Information Processing Systems*, 34: 4384–4396, 2021.

Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 400–407, 1951.

R Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.

Ernest K Ryu and Stephen Boyd. Stochastic proximal iteration: a non-asymptotic improvement upon stochastic gradient descent. *Author website, early draft*, 2014.

Abdurakhmon Sadiev, Dmitry Kovalev, and Peter Richtárik. Communication acceleration of local gradient methods via an accelerated primal-dual algorithm with an inexact prox. *Advances in Neural Information Processing Systems*, 35:21777–21791, 2022.

Abdurakhmon Sadiev, Grigory Malinovsky, Eduard Gorbunov, Igor Sokolov, Ahmed Khaled, Konstantin Burlachenko, and Peter Richtárik. Don't compress gradients in random reshuffling: compress gradient differences. *Advances in Neural Information Processing Systems*, 37:84523–84605, 2024.

Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Interspeech*, volume 2014, pp. 1058–1062. Singapore, 2014.

Mikhail V Solodov and Benar F Svaiter. A hybrid projection-proximal point algorithm. *Journal of convex analysis*, 6(1):59–70, 1999.

Canh T Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with Moreau envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020.

Alexander Tyurin and Peter Richtárik. DASHA: Distributed nonconvex optimization with communication compression and optimal oracle complexity. In *International Conference on Learning Representations*, 2024.

CONTENTS

1 Introduction **1**
1.1 Contributions . . . 2
1.2 Related work . . . 3

2 Mathematical background **4**

3 Absolute approximation in distance **6**

4 Relative approximation in distance **7**

5 Achieving a sufficient level of Inexactness **8**

6 Experiments **9**

7 Reproducibility statement **9**

A LLM Usage **15**

B Notations **15**

C Facts and lemmas **15**

D Theory of biased SGD **17**

E Theory of biased compression **18**

F Full Statements of Theorems **19**

G Further discussion **20**

H Analysis of inexact FedExProx in the client sampling setting **20**
H.1 Relative approximation in distance . . . 20
H.2 Absolute approximation in distance . . . 21

I Proof of theorems and lemmas **21**
I.1 Proof of Lemma C.9 . . . 21
I.2 Proof of Theorem 3.2 . . . 22
I.3 Proof of Theorem 4.2 . . . 24
I.4 Proof of Theorem 4.3 . . . 27
I.5 Proof of Theorem 5.1 . . . 30
I.6 Proof of Theorem 5.2 . . . 30
I.7 Proof of Theorem F.1 . . . 31
I.8 Proof of Remark F.3 . . . 32

I.9 Proof of Theorem F.5 32
I.10 Proof of Theorem H.1 34

J Experiments **37**
J.1 Numerical experiments 37
J.2 Deep learning experiments 40

A LLM USAGE

Large Language Models (LLMs) were used exclusively for editing purposes, including grammar, spelling, and word choice. No content generation, analysis, or substantive contributions were made by LLMs.

B NOTATIONS

Throughout the paper, we use the notation $\|\cdot\|$ to denote the standard Euclidean norm defined on $\mathbb{R}^d$ and $\langle\cdot,\cdot\rangle$ to denote the standard Euclidean inner product. Given a differentiable function $f:\mathbb{R}^d\mapsto\mathbb{R}$, its gradient is denoted as $\nabla f(x)$. We use the notation $D_f(x,y)$ to denote the Bregman divergence associated with a function $f:\mathbb{R}^d\mapsto\mathbb{R}$ between x and y. The notation $\inf f$ is used to denote the minimum of a function $f:\mathbb{R}^d\mapsto\mathbb{R}$. We use $\mathrm{prox}_{\gamma\phi}(x)$ to denote the proximity operator of function $\phi:\mathbb{R}^d\mapsto\mathbb{R}$ with $\gamma>0$ at $x\in\mathbb{R}^d$, and $M_\phi^\gamma(x)$ to denote the corresponding Moreau Envelope. We denote the average of the Moreau envelope of each local objective f_i by the notation $M^\gamma:\mathbb{R}^d\mapsto\mathbb{R}$. Specifically, we define $M^\gamma(x)=\frac{1}{n}\sum_{i=1}^n M_f^\gamma(x)$. Note that $M^\gamma(x)$ has an implicit dependence on γ, its smoothness constant is denoted by L_γ, convexity constant as μ_γ. We say an extended real-valued function $f:\mathbb{R}^d\mapsto\mathbb{R}\cup\{+\infty\}$ is proper if there exists $x\in\mathbb{R}^d$ such that $f(x)<+\infty$. We say an extended real-valued function $f:\mathbb{R}^d\mapsto\mathbb{R}\cup\{+\infty\}$ is closed if its epigraph is a closed set. We use the notation $\mathcal{E}_k=\gamma M^\gamma(x_k)-\gamma M_{\inf}^\gamma$ to denote the function value suboptimality of γM^γ at x_k, and $\Delta_k=\|x_k-x_\star\|^2$ to denote the squared distance. The notation $\mathcal{O}(\cdot)$ is used to describe complexity while omitting constant factors, whereas $\tilde{\mathcal{O}}(\cdot)$ is used when logarithmic factors are also omitted. For approximation $y\in\mathbb{R}^d$ of $\mathrm{prox}_{\gamma f}(x)$, we use ε_1 as the accuracy of absolute approximation such that $\left\|y-\mathrm{prox}_{\gamma f}(x)\right\|^2\le\varepsilon_1$, and we use ε_2 as the accuracy of relative approximation such that $\left\|y-\mathrm{prox}_{\gamma f}(x)\right\|^2\le\varepsilon_2\cdot\left\|x-\mathrm{prox}_{\gamma f}(x)\right\|^2$.

C FACTS AND LEMMAS

Fact C.1 (Young's inequality)**.** For any two vectors $x,y\in\mathbb{R}^d$, the following inequality holds,

$$\|x+y\|^2\le 2\|x\|^2+2\|y\|^2. \tag{13}$$

Fact C.2 (Property of convex smooth functions)**.** Let $\phi:\mathbb{R}^d\mapsto\mathbb{R}$ be differentiable. The following statements are equivalent:

1. ϕ is convex and L-smooth.
2. $0\le 2D_\phi(x,y)\le L\|x-y\|^2$ for all $x,y\in\mathbb{R}^d$.
3. $\frac{1}{L}\|\nabla\phi(x)-\nabla\phi(y)\|^2\le 2D_\phi(x,y)$ for all $x,y\in\mathbb{R}^d$.

The notation $D_\phi(x,y)$ denotes the Bregman divergence associate with ϕ at $x,y\in R^d$, defined as

$$D_\phi(x,y)=\phi(x)-\phi(y)-\langle\nabla\phi(y),x-y\rangle.$$

The following two facts establish that the convexity and smoothness of a function $\phi:\mathbb{R}^d\mapsto\mathbb{R}$ ensure the convexity and smoothness of its Moreau envelope.

Fact C.3 (Convexity of Moreau envelope)**.** (Beck, 2017, Theorem 6.55) Let $\phi : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$ be a proper and convex function. Then M_ϕ^γ is a convex function.

Fact C.4 (Smoothness of Moreau envelope)**.** (Li et al., 2024a, Lemma 4) Let $\phi : \mathbb{R}^d \mapsto \mathbb{R}$ be a convex and L-smooth function. Then M_ϕ^γ is $\frac{L}{1+\gamma L}$-smooth.

The following fact illustrates the relationship between the minimizer of a function ϕ and its Moreau envelope M_ϕ^γ.

Fact C.5 (Minimizer equivalence)**.** (Li et al., 2024a, Lemma 5) Let $\phi : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$ be a proper, closed and convex function. Then for any $\gamma > 0$, ϕ and M_ϕ^γ has the same set of minimizers.

In our case, we assume each f_i from (1) is convex and L_i-smooth. Therefore by Fact C.3 and Fact C.4, we know that each $M_{f_i}^\gamma$ is also convex and $\frac{L_i}{1+\gamma L_i}$-smooth. This means that $M_\gamma = \frac{1}{n}\sum_{i=1}^n M_{f_i}^\gamma$ is also convex and smooth. We denote its smoothness constant as L_γ, and the following fact provides a range for this constant.

Fact C.6 (Global convexity and smoothness)**.** (Li et al., 2024a, Lemma 7) Let each f_i be proper, closed convex and L_i-smooth. Then M^γ is convex and L_γ-smooth with

$$\frac{1}{n^2}\sum_{i=1}^n \frac{L_i}{1+\gamma L_i} \le L_\gamma \le \frac{1}{n}\sum_{i=1}^n \frac{L_i}{1+\gamma L_i}.$$

The following fact establishes that the minimizer of f and M^γ are the same.

Fact C.7 (Global minimizer equivalence)**.** (Li et al., 2024a, Lemma 8) If we let every $f_i : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$ be proper, closed and convex, then $f(x) = \frac{1}{n}\sum_{i=1}^n f_i(x)$ has the same set of minimizers and minimum as

$$M^\gamma(x) = \frac{1}{n}\sum_{i=1}^n M_{f_i}^\gamma(x),$$

if we are in the interpolation regime and $0 < \gamma < \infty$.

The above fact demonstrates that running SGD on the objective M^γ will lead us to the correct destination, as the minimizers of M^γ and f are identical in our setting. In problem (1), if we assume that f is strongly convex, then we have M^γ satisfies the following star strong convexity inequality.

Fact C.8 (Star strong convexity)**.** (Li et al., 2024a, Lemma 11) Assume Assumption 2.3, 2.4, 2.5 and 2.6 hold, then the convex function $M^\gamma(x)$ satisfies the following inequality,

$$M^\gamma(x) - M_{\inf}^\gamma \ge \frac{\mu}{1+\gamma L_{\max}} \cdot \frac{1}{2}\|x - x_\star\|^2,$$

for any $x \in \mathbb{R}^d$ and a minimizer $x_\star$ of $M^\gamma(x)$.

The above fact implies that the strong convexity of f translates to the star strong convexity of M^γ. Star strong convexity is also known as quadratic growth (QG) condition (Anitescu, 2000). In the case of a convex function, it is also known as optimal strong convexity (Liu & Wright, 2015) and semi-strong convexity (Gong & Ye, 2014). It is known that for a convex function satisfying quadratic growth condition, it also satisfies the Polyak-Lojasiewicz inequality (Polyak, 1964) which is described by the following lemma. Notice that since Algorithm 1 can be viewed as running SGD with objective γM^γ and a fixed step size $\alpha_k = \alpha$, we describe the inequality based on γM^γ in the following lemma.

Lemma C.9 (PL-inequality)**.** *Let Assumption 2.3, 2.4, 2.5 and 2.6 hold, then $\gamma M^\gamma(x)$ satisfies the following Polyak-Lojasiewicz inequality,*

$$\|\gamma \nabla M^\gamma(x)\|^2 \ge 2 \cdot \frac{\gamma\mu}{4(1+\gamma L_{\max})}\left(\gamma M^\gamma(x) - \gamma M_{\inf}^\gamma\right), \tag{14}$$

where $x \in \mathbb{R}^d$ is an arbitrary vector and $x_\star$ is a minimizer of $M^\gamma(x)$.

D THEORY OF BIASED SGD

For completeness, we provide the theory of biased SGD we used to analyze our algorithm in this paper. It is adapted from Demidovich et al. (2024), which offers a comprehensive study of various assumptions employed in the analysis of SGD with biased gradient updates. In addition, the authors introduced a new set of assumptions, referred to as the Biased ABC assumption, which are less restrictive than all previous assumptions. The authors provided convergence guarantees for SGD with biased gradient updates in the non-convex and convex setting. Specifically, they considered the case of minimizing a function $f : \mathbb{R}^d \mapsto \mathbb{R}$,

$$\min_{x \in \mathbb{R}^d} f(x),$$

with

$$x_{k+1} = x_k - \eta g(x_k), \tag{biased SGD}$$

where $\eta > 0$ is the stepsize, $g(x_k)$ is a possibly stochastic and biased gradient estimator. They introduced the biased ABC assumption,

Assumption D.1 (Biased-ABC)**.** (Demidovich et al., 2024, Assumption 9) There exists constants $A, B, C, b, c \geq 0$ such that the gradient estimator $g(x)$ for every $x \in \mathbb{R}^d$ satisfies

$$\begin{aligned} \langle \nabla f(x), \mathbb{E}[g(x)] \rangle &\geq b \left\| \nabla f(x) \right\|^2 - c \\ \mathbb{E}\left[\left\| g(x) \right\|^2 \right] &\leq 2A \left(f(x) - f_{\inf} \right) + B \left\| \nabla f(x) \right\|^2 + C. \end{aligned}$$

A convergence guarantee was provided for biased SGD under Assumption D.1 given that f is $\widehat{L}$-smooth and $\widehat{\mu}$-PL, that is, there exists $\widehat{\mu} > 0$, such that

$$\left\| \nabla f(x) \right\|^2 \geq 2\widehat{\mu} \left(f(x) - f_{\inf} \right),$$

for all $x \in \mathbb{R}^d$.

Theorem D.2 (Theory of biased SGD)**.** *(Demidovich et al., 2024, Theorem 4) Let f be $\widehat{L}$-smooth and $\widehat{\mu}$-PL and Assumption D.1 hold. If we choose a step size η satisfying*

$$0 < \eta < \min \left\{ \frac{\widehat{\mu} b}{\widehat{L} \left(A + \widehat{\mu} B \right)}, \frac{1}{\widehat{\mu} b} \right\}. \tag{15}$$

Then we have

$$\mathbb{E}[f(x_k) - f_{\inf}] \leq \left(1 - \eta \widehat{\mu} b \right)^k \left(f(x_0) - f_{\inf} \right) + \frac{LC\eta}{2\widehat{\mu} b} + \frac{c}{\widehat{\mu} b}.$$

Under the special case of

$$\frac{\widehat{\mu} b}{\widehat{L} \left(A + \widehat{\mu} B \right)} < \frac{1}{\widehat{\mu} b},$$

The range of the step size can be simplified to

$$0 < \eta \leq \frac{\widehat{\mu} b}{\widehat{L} \left(A + \widehat{\mu} B \right)},$$

and if we take the largest possible step size, we have

$$\mathbb{E}[f(x_k) - f_{\inf}] \leq \left(1 - \frac{\widehat{\mu}^2 b^2}{\widehat{L} \left(A + \widehat{\mu} B \right)} \right)^k \left(f(x_0) - f_{\inf} \right) + \frac{LC}{2\widehat{L} \left(A + \widehat{\mu} B \right)} + \frac{c}{\widehat{\mu} b}.$$

The constants C, c determine whether the algorithm is converging to the exact solution or just a neighborhood. For $g(x) = \nabla f(x)$, clearly we have $A = 0, B = 1, b = 1, C = 0, c = 0$, and there is no neighborhood. This is expected because the algorithm reduces to standard GD The iteration complexity is give by $\tilde{\mathcal{O}}\left(\frac{\widehat{L}}{\widehat{\mu}} \right)$, which is also expected for GD.

E THEORY OF BIASED COMPRESSION

In this section, we present the theory of SGD with biased compression. The theory is adapted from Beznosikov et al. (2023). The authors introduced theory for analyzing compressed gradient descent (CGD) with biased compressor, both in the single node case and in the distributed case when the objective function is assumed to be strongly convex. Here, we are only concerned with the single node case because distributed compressed gradient descent (DCGD) with biased compressor may fail to converge. To address this issue, error feedback mechanism (Seide et al., 2014; Karimireddy et al., 2019; Richtárik et al., 2021) is needed. In the single node case, the authors considered solving

$$\min_{x\in\mathbb{R}^d} f(x),$$

where $f:\mathbb{R}^d\mapsto\mathbb{R}$ is $\widehat{L}$-smooth and $\widehat{\mu}$-strongly convex, with the following compressed gradient descent algorithm

$$x_{k+1}=x_k-\eta\mathcal{C}\left(\nabla f(x_k)\right), \tag{CGD}$$

where $\mathcal{C}:\mathbb{R}^d\mapsto\mathbb{R}$ are potentially biased compression operators, $\eta>0$ is a step size. The author proved that if certain conditions on $\mathcal{C}$ is satisfied, a corresponding convergence guarantee can then be established. Three classes of compressor/mapping were introduced.

Definition E.1 (Class $\mathbb{B}^1$)**.** We say a mapping $\mathcal{C}\in\mathbb{B}^1(\alpha,\beta)$ for some $\alpha,\beta>0$ if

$$\alpha\left\|x\right\|^2\leq\mathbb{E}\left[\left\|\mathcal{C}(x)\right\|^2\right]\leq\beta\left\langle\mathbb{E}[\mathcal{C}(x)],x\right\rangle,\qquad\forall x\in\mathbb{R}^d.$$

Definition E.2 (Class $\mathbb{B}^2$)**.** We say a mapping $\mathcal{C}\in\mathbb{B}^2(\xi,\beta)$ for some $\xi,\beta>0$ if

$$\max\left\{\xi\left\|x\right\|^2,\frac{1}{\beta}\mathbb{E}\left[\left\|\mathcal{C}(x)\right\|^2\right]\right\}\leq\left\langle\mathbb{E}[\mathcal{C}(x)],x\right\rangle,\qquad\forall x\in\mathbb{R}^d.$$

Definition E.3 (Class $\mathbb{B}^3$)**.** We say a mapping $\mathcal{C}\in\mathbb{B}^3(\delta)$ for some $\delta>0$, if

$$\mathbb{E}\left[\left\|\mathcal{C}(x)-x\right\|^2\right]\leq\left(1-\frac{1}{\delta}\right)\left\|x\right\|^2.$$

The authors proved the following theorem about the convergence of the algorithm, the notation $\mathcal{F}_k$ is used to denote $\mathbb{E}[f(x_k)]-f_{\inf}$, with $\mathcal{F}_0=f(x_0)-f_{\inf}$,

Theorem E.4. *Let $\mathcal{C}\in\mathbb{B}^1(\alpha,\beta)$. Then we have $\mathcal{F}_k\leq\left(1-\alpha/\beta\eta\widehat{\mu}\left(2-\eta\beta\widehat{L}\right)\right)\mathcal{F}_{k-1}$, as long as $0\leq\eta\leq\frac{2}{\beta\widehat{L}}$. If we choose $\eta=\frac{1}{\beta\widehat{L}}$, we have*

$$\mathcal{F}_k\leq\left(1-\frac{\alpha}{\beta^2}\cdot\frac{\widehat{\mu}}{\widehat{L}}\right)^K\mathcal{F}_0. \tag{16}$$

Let $\mathcal{C}\in\mathbb{B}^2(\xi,\beta)$. Then we have $\mathcal{F}_k\leq\left(1-\xi\eta\left(2-\eta\beta\right)\widehat{L}\right)\mathcal{F}_{k-1}$, as long as $0\leq\eta\leq\frac{2}{\beta\widehat{L}}$. If we choose $\eta=\frac{1}{\beta\widehat{L}}$, we have

$$\mathcal{F}_k\leq\left(1-\frac{\xi}{\beta}\cdot\frac{\widehat{\mu}}{\widehat{L}}\right)^k\mathcal{F}_0. \tag{17}$$

Let $\mathcal{C}\in\mathbb{B}^3(\delta)$. Then we have $\mathcal{F}_k\leq\left(1-\frac{1}{\delta}\eta\widehat{\mu}\right)\mathcal{F}_{k-1}$, as long as $0\leq\eta\leq\frac{1}{\widehat{L}}$. If we choose $\eta=\frac{1}{\widehat{L}}$, we have

$$\mathcal{F}_k\leq\left(1-\frac{1}{\delta}\cdot\frac{\widehat{\mu}}{\widehat{L}}\right)^k\mathcal{F}_0. \tag{18}$$

Notice that when $\mathcal{C}(x)=x$, that is, when no compression happens, we have $\alpha=\beta=\xi=\delta=1$. In this case, the iteration complexity of CGD is given by $\tilde{\mathcal{O}}\left(\frac{\widehat{L}}{\widehat{\mu}}\right)$ and we recover the result of GD. It is worth noting that Theorem E.4 remains valid if the condition of f being $\widehat{\mu}$-strongly convex is replaced with f being $\widehat{\mu}$-PL.

F FULL STATEMENTS OF THEOREMS

Theorem F.1. *Assume Assumption 2.3, 2.4, 2.5, and 2.6 hold for our objective. Suppose each client computes an ε_1-absolute approximation $\tilde{x}_{i,k+1}$ of $\operatorname{prox}_{\gamma f_i}(x_k)$ at every iteration, i.e.,*

$$\left\|\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k)\right\|^2 \leq \varepsilon_1.$$

Then for Algorithm 1, with extrapolation parameter $\alpha_k = \alpha$ satisfying $\alpha = \frac{1}{4\gamma L_\gamma}$, where γ is the proximal step size and μ_γ, L_γ is the convexity and smoothness constant of M^γ, the last iterate x_K satisfies

$$\Delta_K^f \leq \left(1 - \frac{\mu_\gamma}{32L_\gamma}\right)^K \frac{2L_\gamma}{\mu_\gamma}\Delta_0^f + \frac{12\varepsilon_1}{\gamma\mu_\gamma^2} + 2\left\|x_\star^f - x_\star^M\right\|^2,$$

where $\Delta_K^f = \left\|x_K - x_\star^f\right\|^2$, $x_\star^f$ is a minimizer of f, and $x_\star^M$ is a minimizer of M^γ.

As expected, in the general case without the interpolation condition, the algorithm converges only to an enlarged neighborhood of the solution. In the following remarks, we show that this neighborhood is bounded and controllable, which allows us to set it aside temporarily and simplify the discussion.

Remark F.2 (Controlling the neighborhood size)**.** The additional term $2\|x_\star^f - x_\star^M\|^2$ arises because the algorithm minimizes M^γ rather than f, and M^γ can be viewed as a personalized federated learning objective:

$$\frac{1}{n}\sum_{i=1}^n F_i(x) = \frac{1}{n}\sum_{i=1}^n M_{f_i}^\gamma(x) := \frac{1}{n}\sum_{i=1}^n \min_{z\in\mathbb{R}^d}\left\{f_i(z) + \tfrac{1}{2\gamma}\|z - x\|^2\right\}.$$

The size of this neighborhood is controlled by the proximal parameter γ (the local step size). As $\gamma \to 0$, we have $M_{f_i}^\gamma \to f_i$, implying $\|x_\star^f - x_\star^M\| = 0$. In contrast, as $\gamma \to \infty$, any $x \in \mathbb{R}^d$ becomes a minimizer of the envelope, and $\|x_\star^f - x_\star^M\| \to \infty$. Thus, while γ can be chosen arbitrarily, meaningful convergence guarantees require it to be sufficiently small.

Remark F.3 (Upper bound on the neighborhood)**.** When the proximal parameter γ is chosen sufficiently small, specifically $\gamma L_i \leq 1$ for all $i \in [n]$, the distance between the minimizers of f and M^γ admits the bound

$$\|x_\star^M - x_\star^f\|^2 \leq \frac{\gamma(1+\gamma L_{\max})}{\mu(1-\gamma L_{\max})^2}\,\sigma_\star^2,$$

where the heterogeneity measure is defined as

$$\sigma_\star^2 := \frac{1}{n}\sum_{i=1}^n \|\nabla f_i(x_\star^f)\|^2.$$

Remark F.4. Notice that choosing a smaller γ makes the local subproblems easier to approximate, since each step corresponds to less local progress. This increases the total number of global iterations required, but reduces the computational burden on each client per iteration.

Theorem F.5. *Assume Assumption 2.3, 2.4, 2.5, and 2.6 hold for our objective. Let the approximation $\tilde{x}_{i,k+1}$ all satisfies Definition 4.1 with $\varepsilon_2 < \mu_\gamma(\gamma + \frac{1}{L_{\max}})/4$, that is*

$$\left\|\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k)\right\|^2 \leq \varepsilon_2 \cdot \left\|x_k - \operatorname{prox}_{\gamma f_i}(x_k)\right\|^2.$$

If we are running Algorithm 1 with $\alpha_k = \alpha = 1/\gamma L_\gamma$, we have the corresponding iterates satisfying

$$\Delta_K^f \leq \left(1 - \frac{\mu_\gamma}{L_\gamma}\left(1 - \frac{4\varepsilon_2 L_{\max}}{\mu_\gamma(1+\gamma L_{\max})}\right)\right)^K \cdot \frac{2L_\gamma}{\mu_\gamma}\Delta_0^f + 2\left\|x_\star^f - x_\star^M\right\|^2,$$

Remark F.6 (Sanity check)**.** In the absence of inexactness, the rate naturally reduces to the standard exponential convergence rate, as expected. As in Theorem F.1, the iterates converge to an additional neighborhood, whose control we have already discussed above.

G FURTHER DISCUSSION

The motivation behind FedExProx stems from the parallel projection method Combettes (1997) of solving the convex feasibility problem. Initially, it was observed that extrapolation can accelerate the parallel projection method (in this convex interpolation setting). Given the similarity between projections and proximal operators (the latter can be viewed as a projection to a level set of the function), FedExProx was developed. One of our contributions is to demonstrate that the acceleration effect extends to a broader setting, going beyond exactness and convex feasibility, holding even without interpolation, exact evaluations, convexity, or smoothness. Empirically, we illustrate the effectiveness of extrapolation with inexact proximal updates through deep learning experiments, showing a $2\times$-$4\times$ speedup with no additional overhead.

In the heterogeneous setting, the solution of the envelope objective $x_\star^M$ naturally corresponds to the personalized federated learning framework. Instead of enforcing consensus on a single global minimizer across all clients, the envelope objective aggregates the client-specific objectives through their Moreau envelopes, which balance proximity to the shared iterate with fidelity to each client's local loss. As a result, $x_\star^M$ can be interpreted as a personalized solution that adapts to client heterogeneity while still maintaining coordination through the shared global model. This perspective aligns with recent work in personalized federated learning, where the goal is to strike a tradeoff between global generalization and local adaptation.

H ANALYSIS OF INEXACT FEDEXPROX IN THE CLIENT SAMPLING SETTING

In this section, we will discuss the case where we do client sampling in algorithm 1, we first formulate the algorithm as below. For the sake of simplicity, we use τ-nice sampling as an example.

Algorithm 2 Inexact FedExProx with τ-nice sampling

1: **Parameters:** extrapolation parameter $\alpha_k = \alpha > 0$, step size for the proximal operator $\gamma > 0$, starting point $x_0 \in \mathbb{R}^d$, number of clients n, size of minibatch τ, total number of iterations K, proximal solution accuracy $\varepsilon_2 \geq 0$.
2: **for** $k = 0, 1, 2 \ldots K-1$ **do**
3: The server broadcasts the current iterate x_k to a selected set of client S_k of size τ
4: Each selected client computes a ε approximation of the solution $\tilde{x}_{i,k+1} \simeq \operatorname{prox}_{\gamma f_i}(x_k)$, and sends it back to the server
5: The server computes

$$x_{k+1} = x_k + \alpha_k \left(\frac{1}{\tau} \sum_{i \in S_k} \tilde{x}_{i,k+1} - x_k \right). \tag{19}$$

6: **end for**

H.1 RELATIVE APPROXIMATION IN DISTANCE

The biased compression theory: Similar to Theorem E.4, we initially apply the theory from Beznosikov et al. (2023), as it provides improved results in the full-batch scenario. We first define the compressing mapping $\mathcal{C}_\tau$ in this case,

$$\mathcal{C}_\tau\left(\gamma \nabla M^\gamma(x_k)\right) = \frac{1}{\tau} \sum_{i \in S_k} \left(\gamma \nabla M_{f_i}^\gamma(x_k) - \left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k) \right) \right). \tag{20}$$

One can verify for every x_k and ε_2-approximation $\tilde{x}_{i,k+1}$ of $\operatorname{prox}_{\gamma f_i}(x_k)$, we have

$$\mathcal{C}_\tau \in \mathbb{B}^3 \left(\delta = \frac{\mu}{\mu - 4\varepsilon_2 L_{\max} - \frac{n-\tau}{\tau(n-1)} \left[4\left(2+\varepsilon_2\right) L_{\max} - 2\mu \right]} \right)$$

In the case of $\tau = n$, we have $\mathcal{C}_n \in \mathbb{B}^3\left(\frac{\mu}{\mu - 4\varepsilon_2 L_{\max}}\right)$, which recovers the result of (41). When $\tau = 1, \varepsilon_2 = 0$, however, this is problematic, as $\mathcal{C}_1 \in \mathbb{B}^3\left(\delta = \frac{\mu}{3\mu - 8L_{\max}}\right)$. Notice that we require

$\delta > 0$, so we require $3\mu > 8L_{\max}$ which only holds in a very restrictive setting. This is due to the stochasticity contained in (20), which arises from client sampling.

Theory of biased SGD: Built upon the theory of biased SGD, one can obtain a convergence guarantee.

Theorem H.1. *Assume 2.3, 2.4, 2.5, 2.6 and 2.7 hold. Let the approximation $\tilde{x}_{i,k+1}$ all satisfies Definition 4.1 with $\varepsilon_2 < \frac{\mu^2}{4L_{\max}^2}$, that is*

$$\left\|\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k)\right\|^2 \leq \varepsilon_2 \cdot \left\|x_k - \operatorname{prox}_{\gamma f_i}(x_k)\right\|^2,$$

holds for all client i at iteration k. If we are running Algorithm 2 with minibatch size τ and extrapolation parameter $\alpha_k = \alpha > 0$ satisfying

$$\alpha \leq \frac{1}{\gamma L_\gamma} \cdot \frac{\mu - 2\sqrt{\varepsilon_2}L_{\max}}{\mu + 4\varepsilon_2 L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} + \frac{n-\tau}{\tau(n-1)} \cdot \left(4L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} - \mu\right)}$$

Then the iterates generated by Algorithm 2 satisfies

$$\mathbb{E}[\mathcal{E}_K] \leq \left(1 - \alpha \cdot \frac{\gamma\left(\mu - 2\sqrt{\varepsilon_2}L_{\max}\right)}{4\left(1 + \gamma L_{\max}\right)}\right)^K \mathcal{E}_0. \tag{21}$$

Specifically, if we choose the largest α possible, we have

$$\mathbb{E}[\Delta_K] \leq \left(1 - \frac{\mu}{4L_\gamma\left(1 + \gamma L_{\max}\right)} \cdot S\left(\varepsilon_2, \tau\right)\right)^K \cdot \frac{L\gamma\left(1 + \gamma L_{\max}\right)}{\mu}\Delta_0,$$

where $S(\varepsilon_2, \tau)$ is defined as

$$S(\varepsilon_2, \tau) := \frac{\left(\mu - 2\sqrt{\varepsilon_2}L_{\max}\right)\left(1 - 2\sqrt{\varepsilon_2}\frac{L_{\max}}{\mu}\right)}{\mu + 4\varepsilon_2 L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} + \frac{n-\tau}{\tau(n-1)} \cdot \left(4L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} - \mu\right)},$$

satisfying

$$0 < S(\varepsilon_2, \tau) \leq 1.$$

Notice that we have $S(\varepsilon_2, \tau = n) = S(\varepsilon_2)$, which appears in Theorem 4.2. For the special case when $\varepsilon_2 = 0$, every proximal operator is solved exactly. The range of α becomes,

$$0 < \alpha \leq \frac{1}{\gamma L_\gamma} \cdot \frac{\mu}{\frac{n-\tau}{\tau(n-1)} \cdot 4L_{\max} + \frac{n(\tau-1)}{\tau(n-1)}\mu}.$$

According to Li et al. (2024a),

$$0 < \alpha \leq \frac{1}{\gamma L_\gamma} \cdot \frac{L_\gamma\left(1 + \gamma L_{\max}\right)}{\frac{n-\tau}{\tau(n-1)}L_{\max} + \frac{n(\tau-1)}{\tau(n-1)} \cdot L_\gamma\left(1 + \gamma L_{\max}\right)}.$$

This matches our expectation. When client sampling is used together with biased compressors, it does not necessarily guarantee any benefits.

H.2 Absolute approximation in distance

Similarly to Theorem H.1, by applying the theory of biased SGD (Demidovich et al., 2024), we can derive a convergence guarantee for the minibatch case. For brevity, we do not include the details here.

I Proof of theorems and lemmas

I.1 Proof of Lemma C.9

Using Fact C.8, we have

$$M^\gamma(x) - M_{\inf}^\gamma \geq \frac{\mu}{1 + \gamma L_{\max}} \cdot \frac{1}{2}\left\|x - x_\star\right\|^2, \tag{22}$$

where $x \in \mathbb{R}^d$ is any vector, $x_\star$ is a minimizer of M^γ, by Fact C.5, it is also a minimizer of f. Since we assume each function f_i is convex, by Fact C.3, we know that $M_{f_i}^\gamma$ is also convex. As a result, the average of $M_{f_i}^\gamma$, M^γ is also a convex function. Utilizing the convexity of M^γ, we have,

$$M_{\inf}^\gamma \geq M^\gamma(x) + \langle \nabla M^\gamma(x), x_\star - x \rangle .$$

Rearranging terms we get,

$$\langle \nabla M^\gamma(x), x - x_\star \rangle \geq M^\gamma(x) - M_{\inf}^\gamma. \tag{23}$$

As a result, we have

$$\langle \nabla M^\gamma(x), x - x_\star \rangle \overset{(22)+(23)}{\geq} \frac{\mu}{1+\gamma L_{\max}} \cdot \frac{1}{2} \|x - x_\star\|^2 .$$

Using Cauchy-Schwarz inequality, we have

$$\|\nabla M^\gamma(x)\| \|x - x_\star\| \geq \langle \nabla M^\gamma(x), x - x_\star \rangle \geq \frac{\mu}{1+\gamma L_{\max}} \cdot \frac{1}{2} \|x - x_\star\|^2 .$$

When $\|x - x_\star\| > 0$, the above inequality leads to

$$\|\nabla M^\gamma(x)\| \geq \frac{\mu}{2(1+\gamma L_{\max})} \cdot \|x - x_\star\| , \tag{24}$$

which also holds when $\|x - x_\star\| = 0$. Now using (23) and (24), we obtain

$$\begin{aligned} M^\gamma(x) - M_{\inf}^\gamma &\overset{(23)}{\leq} \langle \nabla M^\gamma(x), x - x_\star \rangle \\ &\leq \|\nabla M^\gamma(x)\| \|x - x_\star\| \overset{(24)}{\leq} \frac{2(1+\gamma L_{\max})}{\mu} \|\nabla M^\gamma(x)\|^2 . \end{aligned}$$

A simple rearranging of terms result in

$$\|\gamma \nabla M^\gamma(x)\|^2 \geq 2 \cdot \frac{\gamma\mu}{4(1+\gamma L_{\max})} (\gamma M^\gamma(x) - \gamma M_{\inf}^\gamma) .$$

Up till here we have already proved the statement in the lemma, but we want to look at the strongly constant μ of f a little bit. In order to provide an upper bound of μ, we notice that due to Fact C.4, each $M_{f_i}^\gamma$ is $\frac{L_i}{1+\gamma L_i}$-smooth and therefore M^γ is smooth. We use the notation L_γ to denote its smoothness constant. Applying the smoothness of $M^\gamma(x)$, we have

$$M^\gamma(x) \leq M^\gamma(x_\star) + \langle \nabla M^\gamma(x_\star), x - x_\star \rangle + \frac{L_\gamma}{2} \|x - x^\star\|^2 .$$

Utilizing the fact that $\nabla M^\gamma(x_\star) = 0$, we have

$$M^\gamma(x) - M_{\inf}^\gamma \leq \frac{L_\gamma}{2} \|x - x_\star\|^2 \tag{25}$$

Combining (25) and (22), we can deduce that

$$\frac{\mu}{1+\gamma L_{\max}} \cdot \frac{1}{2} \|x - x_\star\|^2 \leq M^\gamma(x) - M_{\inf}^\gamma \leq \frac{L_\gamma}{2} \|x - x_\star\|^2 .$$

which results in the estimate that

$$\mu \leq L_\gamma (1 + \gamma L_{\max}) . \tag{26}$$

I.2 Proof of Theorem 3.2

Let us first recall that after reformulation, Algorithm 1 can be written as

$$x_{k+1} = x_k - \alpha \cdot g(x_k),$$

where $g(x_k)$ is defined as

$$g(x_k) := \frac{1}{n} \sum_{i=1}^n \gamma \nabla M_{f_i}^\gamma(x_k) - \frac{1}{n} \sum_{i=1}^n \left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k) \right) .$$

We view this as running full batch biased SGD with stepsize α and global objective $\gamma M^{\gamma}(x)$. We first examine if Assumption D.1 (Biased-ABC) holds for arbitrary x_k. Since we are in the full batch case, it is easy to see that

$$\mathbb{E}[g(x_k)] = g(x_k).$$

Since our objective now is $\gamma M^{\gamma}(x)$, we have that

$$\begin{aligned}\langle \gamma \nabla M^{\gamma}(x_k), g(x_k)\rangle &= \left\langle \gamma \nabla M^{\gamma}(x_k), \gamma \nabla M^{\gamma}(x_k) - \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\rangle \\ &= \|\gamma \nabla M^{\gamma}(x_k)\|^2 - \underbrace{\left\langle \gamma \nabla M^{\gamma}(x_k), \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\rangle}_{:=P_1}.\end{aligned}$$

Now let us focus on P_1, we have the following upper bound,

$$P_1 \quad \leq \quad \frac{1}{2}\|\gamma \nabla M^{\gamma}(x_k)\|^2 + \frac{1}{2}\left\|\frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2 \overset{(9)}{\leq} \frac{1}{2}\|\gamma \nabla M^{\gamma}(x_k)\|^2 + \frac{\varepsilon_1}{2}.$$

As a result, we have

$$\langle \gamma \nabla M^{\gamma}(x_k), g(x_k)\rangle \geq \frac{1}{2}\|\gamma \nabla M^{\gamma}(x_k)\| - \frac{\varepsilon_1}{2},$$

which holds for arbitrary x_k. This suggests that $b = \frac{1}{2}, c = \frac{\varepsilon_1}{2}$. On the other hand,

$$\begin{aligned}\mathbb{E}\left[\|g(x_k)\|^2\right] \quad &= \quad \left\|\gamma \nabla M^{\gamma}(x_k) + \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2 \\ &\overset{(13)}{\leq} \quad 2\|\gamma \nabla M^{\gamma}(x_k)\|^2 + 2\left\|\frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2 \\ &\overset{(9)}{\leq} \quad 2\|\gamma \nabla M^{\gamma}(x_k)\|^2 + 2\varepsilon_1.\end{aligned}$$

Thus, we can choose $A = 0, B = 2, C = 2\varepsilon_1$. Since we have assumed Assumption 2.4 (Individual convexity) and Assumption 2.5 (Smoothness), it is easy to see that M^{γ} is smooth, and we denote its smoothness constant as L_{γ}. It is therefore straightforward to see that our global objective γM^{γ} is γL_{γ}-smooth. We also assume f is μ-strongly convex, which by Fact C.8 indicates that M^{γ} is $\frac{\mu}{1+\gamma L_{\max}}$ star strongly convex. We immediately obtain using Lemma C.9 that γM^{γ} is $\frac{\gamma\mu}{4(1+\gamma L_{\max})}$-PL. Now, we have validated all the assumptions for using Theorem D.2. Applying Theorem D.2, we obtain that when the extrapolation parameter satisfies

$$0 < \alpha < \frac{1}{4}\cdot \min\left\{\frac{1}{\gamma L_{\gamma}}, \frac{2(1+\gamma L_{\max})}{\gamma\mu}\right\},$$

the last iterate x_K of Algorithm 1 with each proximal operator solved inexactly according to Definition 2.1 satisfies

$$\mathcal{E}_K \leq \left(1 - \frac{\alpha\gamma\mu}{8(1+\gamma L_{\max})}\right)^K \mathcal{E}_0 + \frac{8\varepsilon_1 \alpha L_{\gamma}(1+\gamma L_{\max})}{\mu} + \frac{4\varepsilon_1(1+\gamma L_{\max})}{\gamma\mu},$$

where $\mathcal{E}_k = \gamma M^{\gamma}(x_k) - M^{\gamma}_{\inf}$. Let us now prove that

$$\frac{1}{\gamma L_{\gamma}} < \frac{2(1+\gamma L_{\max})}{\gamma\mu}.$$

This is equivalent to prove

$$\mu < 2L_{\gamma}(1+\gamma L_{\max}),$$

which is always true since (26) holds. As a result, we can simplify the range of the extrapolation parameter to

$$0 < \alpha \leq \frac{1}{4\gamma L_\gamma}.$$

If we pick the largest possible α, we have

$$\mathcal{E}_K \leq \left(1 - \frac{\mu}{32 L_\gamma \left(1 + \gamma L_{\max}\right)}\right)^K \mathcal{E}_0 + \frac{6\varepsilon_1 \left(1 + \gamma L_{\max}\right)}{\gamma \mu}.$$

This result is not directly comparable to that of Li et al. (2024a). However, using smoothness of γL_γ, if we denote $\Delta_k = \|x_k - x_\star\|^2$ where $x_\star$ is a minimizer of both M^γ and f since we assume we are in the interpolation regime (Assumption 2.7), we have

$$\mathcal{E}_0 \leq \frac{\gamma L_\gamma}{2} \Delta_0.$$

Using star strong convexity, we have

$$\mathcal{E}_K \geq \frac{\gamma \mu}{2\left(1 + \gamma L_{\max}\right)} \Delta_K.$$

As a result, we can transform the above convergence guarantee into

$$\Delta_K \leq \left(1 - \frac{\mu}{32 L_\gamma \left(1 + \gamma L_{\max}\right)}\right)^K \frac{L_\gamma \left(1 + \gamma L_{\max}\right)}{\mu} \cdot \Delta_0 + 12\varepsilon_1 \cdot \left(\frac{1/\gamma + L_{\max}}{\mu}\right)^2.$$

This completes the proof.

I.3 Proof of Theorem 4.2

Since we based our analysis on the theory of biased SGD, we first verify the validity of Assumption D.1.

Finding b and c: Let us start with finding a lower bound on $\langle \gamma \nabla M^\gamma\left(x_k\right), \mathbb{E}[g(x_k)]\rangle$. We have

$$\begin{aligned}
\langle \gamma \nabla M^\gamma\left(x_k\right), \mathbb{E}[g(x_k)]\rangle &= \left\langle \gamma \nabla M^\gamma\left(x_k\right), \gamma \nabla M^\gamma\left(x_k\right) - \frac{1}{n}\sum_{i=1}^n \left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}\left(x_k\right)\right)\right\rangle \\
&= \left\|\gamma \nabla M^\gamma\left(x_k\right)\right\|^2 - \left\langle \gamma \nabla M^\gamma\left(x_k\right), \frac{1}{n}\sum_{i=1}^n \left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}\left(x_k\right)\right)\right\rangle \\
&\geq \left\|\gamma \nabla M^\gamma\left(x_k\right)\right\|^2 - \left\|\gamma \nabla M^\gamma\left(x_k\right)\right\| \cdot \left\|\frac{1}{n}\sum_{i=1}^n \left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}\left(x_k\right)\right)\right\|,
\end{aligned}$$

where the last inequality is obtained using Cauchy-Schwarz inequality. We then utilize the convexity of $\|\cdot\|$ and obtain,

$$\begin{aligned}
\langle \gamma \nabla M^\gamma\left(x_k\right), \mathbb{E}[g(x_k)]\rangle &\geq \left\|\gamma \nabla M^\gamma\left(x_k\right)\right\|^2 - \left\|\gamma \nabla M^\gamma\left(x_k\right)\right\| \cdot \frac{1}{n}\sum_{i=1}^n \left\|\left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}\left(x_k\right)\right)\right\| \\
&\overset{(12)}{\geq} \left\|\gamma \nabla M^\gamma\left(x_k\right)\right\|^2 - \sqrt{\varepsilon_2}\left\|\gamma \nabla M^\gamma\left(x_k\right)\right\| \cdot \frac{1}{n}\sum_{i=1}^n \left\|x_k - \operatorname{prox}_{\gamma f_i}\left(x_k\right)\right\| \\
&= \left\|\gamma \nabla M^\gamma\left(x_k\right)\right\|^2 - \sqrt{\varepsilon_2}\left\|\gamma \nabla M^\gamma\left(x_k\right)\right\| \cdot \frac{1}{n}\sum_{i=1}^n \left\|\gamma \nabla M_{f_i}^\gamma\left(x_k\right)\right\|.
\end{aligned}$$

Notice that

$$\left\|\gamma \nabla M_{f_i}^\gamma\left(x_k\right)\right\| = \left\|\gamma \nabla M_{f_i}^\gamma\left(x_k\right) - \gamma \nabla M_{f_i}^\gamma\left(x_\star\right)\right\|,$$

holds for any $x_\star$ that is a minimizer of $M^\gamma(x)$ due to interpolation regime assumption. As a result, we can provide an upper bound based on smoothness of each individual $\gamma M_{f_i}^\gamma(x)$ using Fact C.2,

$$\left\|\gamma\nabla M_{f_i}^\gamma(x_k) - \gamma\nabla M_{f_i}^\gamma(x_\star)\right\| \leq \frac{\gamma L_i}{1+\gamma L_i}\|x_k - x_\star\|. \tag{27}$$

Thus,

$$\frac{1}{n}\sum_{i=1}^n \left\|\gamma\nabla M_{f_i}^\gamma(x_k)\right\| \leq \frac{1}{n}\sum_{i=1}^n \frac{\gamma L_i}{1+\gamma L_i}\|x_k - x_\star\| \leq \frac{\gamma L_{\max}}{1+\gamma L_{\max}}\cdot\|x_k - x_\star\|.$$

In addition, we have due to Cauchy-Schwarz inequality and the convexity of $M^\gamma(x)$

$$\|\nabla M^\gamma(x_k)\|\cdot\|x_k - x_\star\| \geq \langle\nabla M^\gamma(x_k), x_k - x_\star\rangle \geq M^\gamma(x_k) - M_{\inf}^\gamma, \tag{28}$$

and due to quadratic growth condition that

$$M^\gamma(x_k) - M_{\inf}^\gamma \geq \frac{\mu}{1+\gamma L_{\max}}\cdot\frac{1}{2}\|x_k - x_\star\|^2. \tag{29}$$

Combining (28) and (29), we have

$$\frac{\mu}{2(1+\gamma L_{\max})}\cdot\|x_k - x_\star\|^2 \overset{(28)+(29)}{\leq} \|\nabla M^\gamma(x_k)\|\cdot\|x_k - x_\star\|.$$

This indicates that

$$\|x_k - x_\star\| \leq \frac{2(1+\gamma L_{\max})}{\mu}\|\nabla M^\gamma(x_k)\|. \tag{30}$$

Combining (27) and (30), we generate the following lower bound

$$\begin{aligned}
\langle\gamma M^\gamma(x_k), \mathbb{E}[g(x_k)]\rangle &\overset{(27)}{\geq} \|\gamma M^\gamma(x_k)\|^2 - \sqrt{\varepsilon_2}\|\gamma M^\gamma(x_k)\|\cdot\frac{\gamma L_{\max}}{1+\gamma L_{\max}}\|x_k - x_\star\| \\
&\overset{(30)}{\geq} \|\gamma M^\gamma(x_k)\|^2 - \sqrt{\varepsilon_2}\cdot\frac{L_{\max}}{1+\gamma L_{\max}}\cdot\frac{2(1+\gamma L_{\max})}{\mu}\|\gamma M^\gamma(x_k)\|^2 \\
&= \left(1 - \sqrt{\varepsilon_2}\cdot\frac{2L_{\max}}{\mu}\right)\cdot\|\gamma M^\gamma(x_k)\|^2.
\end{aligned}$$

Thus, as long as $\varepsilon_2 < \frac{\mu^2}{4L_{\max}^2}$, we have $b = 1 - \sqrt{\varepsilon_2}\cdot\frac{2L_{\max}}{\mu}$, and $c = 0$.

Finding A, B and C: We start with expanding $\|g(x_k)\|^2$,

$$\begin{aligned}
\mathbb{E}\left[\|g(x_k)\|^2\right] &= \left\|\gamma M^\gamma(x_k) - \frac{1}{n}\sum_{i=1}^n\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2 \\
&= \|\gamma M^\gamma(x_k)\|^2 + \underbrace{\left\|\frac{1}{n}\sum_{i=1}^n\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2}_{:=T_2} \\
&\quad \underbrace{-2\left\langle\gamma M^\gamma(x_k), \frac{1}{n}\sum_{i=1}^n\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\rangle}_{:=T_3}.
\end{aligned} \tag{31}$$

It is easy to bound T_2 utilizing the convexity of $\|\cdot\|^2$,

$$\begin{aligned}
T_2 &\leq \frac{1}{n}\sum_{i=1}^n\left\|\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 \\
&\overset{(12)}{\leq} \frac{\varepsilon_2}{n}\sum_{i=1}^n\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 = \frac{\varepsilon_2}{n}\sum_{i=1}^n\left\|\gamma\nabla M_{f_i}^\gamma(x_k)\right\|^2.
\end{aligned}$$

Let $x_\star$ be a minimizer of M^γ, since we assume Assumption 2.7 holds, it is also a minimizer of each $M^\gamma_{f_i}$. As a result,

$$\begin{aligned} T_2 &\leq \frac{\varepsilon_2}{n}\sum_{i=1}^n \left\|\gamma M^\gamma_{f_i}(x_k) - \gamma M^\gamma_{f_i}(x_\star)\right\|^2 \\ &\leq \frac{\varepsilon_2}{n}\sum_{i=1}^n \frac{2\gamma L_i}{1+\gamma L_i}\left(\gamma M^\gamma_{f_i}(x_k) - \gamma M^\gamma_{f_i}(x_\star)\right) \leq \frac{2\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}\cdot\left(\gamma M^\gamma(x_k) - \gamma M^\gamma_{\inf}\right). \end{aligned} \tag{32}$$

We then consider T_3, and start with applying Cauchy-Schwarz inequality

$$T_3 \leq 2\left\|\gamma\nabla M^\gamma(x_k)\right\|\left\|\frac{1}{n}\sum_{i=1}^n\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|. \tag{33}$$

Using the convexity of $\|\cdot\|$, we have

$$\begin{aligned} \left\|\frac{1}{n}\sum_{i=1}^n\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\| &\leq \frac{1}{n}\sum_{i=1}^n\left\|\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right\| \\ &\overset{(12)}{\leq} \frac{\sqrt{\varepsilon_2}}{n}\sum_{i=1}^n\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\| \\ &\overset{(4)}{=} \frac{\sqrt{\varepsilon_2}}{n}\sum_{i=1}^n\left\|\gamma\nabla M^\gamma_{f_i}(x_k) - \gamma\nabla M^\gamma_{f_i}(x_\star)\right\| \\ &\overset{\text{Fact C.2}}{\leq} \frac{\sqrt{\varepsilon_2}}{n}\sum_{i=1}^n\frac{\gamma L_i}{1+\gamma L_i}\left\|x_k - x_\star\right\| \leq \frac{\sqrt{\varepsilon_2}\gamma L_{\max}}{1+\gamma L_{\max}}\cdot\left\|x_k - x_\star\right\|. \end{aligned}$$

Utilizing (30), we have

$$\begin{aligned} \left\|\frac{1}{n}\sum_{i=1}^n\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\| &\leq \frac{\sqrt{\varepsilon_2}\gamma L_{\max}}{1+\gamma L_{\max}}\cdot\frac{2\left(1+\gamma L_{\max}\right)}{\mu}\left\|\nabla M^\gamma(x_k)\right\| \\ &= \frac{2\sqrt{\varepsilon_2}L_{\max}}{\mu}\cdot\left\|\gamma\nabla M^\gamma(x_k)\right\| \end{aligned} \tag{34}$$

Plug the above inequality into (33), we have

$$T_3 \leq \frac{4\sqrt{\varepsilon_2}L_{\max}}{\mu}\cdot\left\|\gamma\nabla M^\gamma(x_k)\right\|^2. \tag{35}$$

Combining (35) and (32), plugging them into (31), we have

$$\mathbb{E}\Big[\left\|g(x_k)\right\|^2\Big] \leq \frac{2\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}\cdot\left(\gamma M^\gamma(x_k) - \gamma M^\gamma_{\inf}\right) + \left(1 + \frac{4\sqrt{\varepsilon_2}L_{\max}}{\mu}\right)\cdot\left\|\gamma\nabla M^\gamma(x_k)\right\|^2.$$

Thus, we have

$$A = \frac{\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}, \quad B = \frac{\mu + 4\sqrt{\varepsilon_2}L_{\max}}{\mu}, \quad C = 0.$$

Applying Theorem D.2: First, we list our the values appeared respectively,

$$A = \frac{\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}, \quad B = \frac{\mu + 4\sqrt{\varepsilon_2}L_{\max}}{\mu}, \quad b = \frac{\mu - 2\sqrt{\varepsilon_2}L_{\max}}{\mu}, \quad C = c = 0.$$

We know that the PL constant of γM^γ is given by $\frac{\gamma\mu}{4(1+\gamma L_{\max})}$ and the corresponding smoothness constant is γL_γ. Applying Theorem D.2, the range of α is given by

$$0 < \alpha < \min\left\{\underbrace{\frac{1}{\gamma L_\gamma}\cdot\frac{\mu - 2\sqrt{\varepsilon_2}L_{\max}}{\mu + 4\sqrt{\varepsilon_2}L_{\max} + 4\varepsilon_2 L_{\max}}}_{:=B_1}, \underbrace{\frac{4\left(1+\gamma L_{\max}\right)}{\gamma\left(\mu - 2\sqrt{\varepsilon_2}L_{\max}\right)}}_{:=B_2}\right\}. \tag{36}$$

Now notice that actually we can prove that for $\varepsilon_2 < \frac{\mu^2}{4L_{\max}^2}$, we have $B_2 > B_1$, and we can simplify the range of α to

$$0 < \alpha \leq \frac{1}{\gamma L_\gamma} \cdot \frac{\mu - 2\sqrt{\varepsilon_2} L_{\max}}{\mu + 4\sqrt{\varepsilon_2} L_{\max} + 4\varepsilon_2 L_{\max}}.$$

Proof of $B_2 > B_1$: It is easy to verify that the above inequality ($B_2 > B_1$) can be equivalently written as

$$4L_\gamma \left(1 + \gamma L_{\max}\right)\left(\mu + 4\sqrt{\varepsilon_2} L_{\max} + 4\varepsilon_2 L_{\max}\right) > \left(\mu - 2\sqrt{\varepsilon_2} L_{\max}\right)^2,$$

since when $\sqrt{\varepsilon_2} < \frac{\mu}{2L_{\max}}$, we have $\mu - 2\sqrt{\varepsilon_2} L_{\max} > 0$. We expand the right-hand side and obtain:

$$\left(\mu - 2\sqrt{\varepsilon_2} L_{\max}\right)^2 = \mu^2 - 4\sqrt{\varepsilon_2} L_{\max} + 4\varepsilon_2 L_{\max}^2 < 2\mu^2 - 4\sqrt{\varepsilon_2} L_{\max} < 2\mu^2.$$

For the left-hand side, as we have already shown in 26, we have

$$4L_\gamma \left(1 + \gamma L_{\max}\right)\left(\mu + 4\sqrt{\varepsilon_2} L_{\max} + 4\varepsilon_2 L_{\max}\right) \geq 4\mu\left(\mu + 4\sqrt{\varepsilon_2} L_{\max} + 2\varepsilon_2 L_{\max}\right) > 4\mu^2.$$

Combining the above inequality we arrive at $B_2 > B_1$.

The convergence guarantee : Given that we select α properly, we have

$$\mathcal{E}_K \leq \left(1 - \alpha \cdot \frac{\gamma\left(\mu - 2\sqrt{\varepsilon_2} L_{\max}\right)}{4\left(1 + \gamma L_{\max}\right)}\right)^K \mathcal{E}_0,$$

where $\mathcal{E}_k = \gamma M^\gamma\left(x_k\right) - \gamma M_{\inf}^\gamma$. We do not have expectation here since we are in the full batch case. Specifically, if we choose the largest α possible, we have

$$\mathcal{E}_K \leq \left(1 - \frac{\mu}{4L_\gamma\left(1 + \gamma L_{\max}\right)} \cdot S\left(\varepsilon_2\right)\right)^k \mathcal{E}_0,$$

where

$$S(\varepsilon_2) = \frac{\left(\mu - 2\sqrt{\varepsilon_2} L_{\max}\right)\left(1 - 2\sqrt{\varepsilon_2}\frac{L_{\max}}{\mu}\right)}{\mu + 4\sqrt{\varepsilon_2} L_{\max} + 4\varepsilon_2 L_{\max}},$$

satisfies $0 < S(\varepsilon_2) \leq 1$ is the factor of slowing down due to inexact proximity operator evaluation. Using smoothness of γL_γ, if we denote $\Delta_k = \|x_k - x_\star\|^2$ where $x_\star$ is a minimizer of both M^γ and f since we assume we are in the interpolation regime (Assumption 2.7), we have

$$\mathcal{E}_0 \leq \frac{\gamma L_\gamma}{2}\Delta_0.$$

Using star strong convexity (quadratic growth property), we have

$$\mathcal{E}_K \geq \frac{\gamma\mu}{2\left(1 + \gamma L_{\max}\right)}\Delta_K.$$

As a result, we can transform the above convergence guarantee into

$$\Delta_K \leq \left(1 - \frac{\mu}{4L_\gamma\left(1 + \gamma L_{\max}\right)} \cdot S\left(\varepsilon_2\right)\right)^K \cdot \frac{L\gamma\left(1 + \gamma L_{\max}\right)}{\mu}\Delta_0.$$

This completes the proof.

I.4 Proof of Theorem 4.3

We start with formalizing the problem. Using (10) and (11), we can write the update rule of Algorithm 1 as

$$x_{k+1} = x_k - \alpha \cdot \left(\frac{1}{n}\sum_{i=1}^n \gamma\nabla M_{f_i}^\gamma\left(x_k\right) - \frac{1}{n}\sum_{i=1}^n \left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}\left(x_k\right)\right)\right). \tag{37}$$

Since by Definition 4.1, we have $\left\|\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}\left(x_k\right)\right\|^2 \leq \varepsilon_2 \left\|\gamma\nabla M_{f_i}^\gamma\left(x_k\right)\right\|^2$, we can view the left hand side as a compressed version of the true gradient. Specifically, there are two possible perspectives:

(I). Let $\mathcal{C}_i(\cdot)$ be the compressing mapping with the i-th client, $i \in \{1, 2, \ldots, n\}$, defined as

$$\mathcal{C}_i\left(\gamma \nabla M_{f_i}^{\gamma}(x_k)\right) := \gamma \nabla M_{f_i}^{\gamma}(x_k) - \left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right).$$

In this way, we reformulate (37) as

$$x_{k+1} = x_k - \alpha \cdot \frac{1}{n} \sum_{i=1}^{n} \mathcal{C}_i\left(\gamma \nabla M_{f_i}^{\gamma}(x_k)\right). \tag{38}$$

(38) is exactly DCGD with biased compression. We can easily prove that

$$\begin{aligned}
&\mathcal{C}_i \in \mathbb{B}^1\left(\alpha = 1 - 2\sqrt{\varepsilon_2}, \beta = \frac{1-\sqrt{\varepsilon_2}}{1+\varepsilon_2}\right)\\
&\mathcal{C}_i \in \mathbb{B}^2\left(\xi = 1 - \sqrt{\varepsilon_2}, \beta = \frac{1-\sqrt{\varepsilon_2}}{1+\varepsilon_2}\right)\\
&\mathcal{C}_i \in \mathbb{B}^3\left(\delta = \frac{1}{1-\varepsilon_2}\right).
\end{aligned}$$

However, DCGD with biased compression may fail to converge even if the above formulation of compression mapping seems quite nice. For an example of such failure, we refer the readers to (Beznosikov et al., 2023, Example 1). This limitation can be circumvented by employing an error feedback mechanism; however, this approach requires modifications to the original algorithm. We therefore leave it as a future research direction.

(II). We can also view it as if we are in the single node case. Let $\mathcal{C}(\cdot)$ be the compressing mapping defined as

$$\begin{aligned}
\mathcal{C}\left(\nabla M^{\gamma}(x_k)\right) &:= \frac{1}{n}\sum_{i=1}^{n} \gamma \nabla M_{f_i}^{\gamma}(x_k) - \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\\
&= \gamma \nabla M^{\gamma}(x_k) - \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right).
\end{aligned} \tag{39}$$

This formulation leads us to the convergence guarantee appeared in Theorem 4.3, as we illustrate below.

Let us first analyze $\mathcal{C}$ defined in (39). We will verify it belongs to $\mathbb{B}^3(\delta)$. The inequality we want to prove can be written equivalently as

$$\left\|\gamma \nabla M^{\gamma}(x_k) - \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right) - \gamma \nabla M^{\gamma}(x_k)\right\|^2 \leq \left(1 - \frac{1}{\delta}\right)\left\|\gamma \nabla M^{\gamma}(x_k)\right\|^2, \tag{40}$$

which is exactly

$$\left\|\frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2 \leq \left\|\gamma \nabla M^{\gamma}(x_k)\right\|^2$$

For the left-hand side, using the convexity of $\|\cdot\|^2$ in combination with Definition 4.1, we obtain

$$\begin{aligned}
\left\|\frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2 &\leq \frac{1}{n}\sum_{i=1}^{n}\left\|\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2\\
&\leq \frac{\varepsilon_2}{n}\sum_{i=1}^{n}\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2.
\end{aligned}$$

Let $x_\star$ be a minimizer of f, since we assume Assumption 2.7 holds, by Fact C.7, it is also a minimizer of γM^γ,

$$\begin{aligned}
\frac{\varepsilon_2}{n}\sum_{i=1}^n \left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 &\overset{(4)}{=} \frac{\varepsilon_2}{n}\sum_{i=1}^n \left\|\gamma\nabla M_{f_i}^\gamma(x_k)\right\|^2 \\
&= \frac{\varepsilon_2}{n}\sum_{i=1}^n \left\|\gamma\nabla M_{f_i}^\gamma(x_k) - \gamma\nabla M_{f_i}^\gamma(x_\star)\right\|^2 \\
&\overset{\text{Fact C.2}}{\leq} \frac{2\varepsilon_2}{n}\sum_{i=1}^n \frac{\gamma L_i}{1+\gamma L_i}\left(\gamma M_{f_i}^\gamma(x_k) - \gamma M_{f_i}^\gamma(x_\star)\right) \\
&\leq \frac{2\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}\left(\gamma M^\gamma(x_k) - \gamma M^\gamma(x_\star)\right).
\end{aligned}$$

We then notice that as it is illustrated by Lemma C.9, we have

$$\left(1-\frac{1}{\delta}\right)\|\gamma\nabla M^\gamma(x_k)\|^2 \geq \left(1-\frac{1}{\delta}\right)\frac{\gamma\mu}{2(1+\gamma L_{\max})}\left(\gamma M^\gamma(x_k) - \gamma M^\gamma(x_\star)\right).$$

Combining the above two inequalities, we know that the following inequality is a sufficient condition for (40),

$$\frac{2\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}\left(\gamma M^\gamma(x_k) - \gamma M^\gamma(x_\star)\right) \leq \left(1-\frac{1}{\delta}\right)\frac{\gamma\mu}{2(1+\gamma L_{\max})}\left(\gamma M^\gamma(x_k) - \gamma M^\gamma(x_\star)\right).$$

It is easy to check that if we pick

$$\delta = \frac{\mu}{\mu - 4\varepsilon_2 L_{\max}} > 0, \tag{41}$$

the condition is met. However, for this to hold, we must ensure that $\varepsilon_2 < \frac{\mu}{4L_{\max}}$.

As we mentioned in Appendix E, Beznosikov et al. (2023) provided the theory of CGD with biased compressor belongs to $\mathbb{B}^3(\delta)$. We have already shown that $\mathcal{C} \in \mathbb{B}^3\left(\delta = \frac{\mu}{\mu - 4\varepsilon_2 L_{\max}}\right)$, when $\varepsilon_2 < \frac{4L_{\max}}{\mu}$. Notice that our objective γM^γ is γL_γ-smooth and $\frac{\gamma\mu}{1+\gamma L_{\max}}$-PL.[5] Therefore, as long as $0 < \alpha \leq \frac{1}{\gamma L_\gamma}$ and $\varepsilon_2 < \frac{\mu}{4L_{\max}}$, we have

$$\mathcal{E}_K \leq \left(1 - \frac{\mu - 4\varepsilon_2 L_{\max}}{\mu}\cdot\frac{\gamma\mu}{4(1+\gamma L_{\max})}\cdot\alpha\right)^K \mathcal{E}_0,$$

Taking $\alpha = \frac{1}{\gamma L_\gamma}$, which is the largest step size possible, we can further simplify the above convergence into

$$M^\gamma(x_k) - M_\star^\gamma \leq \left(1 - \left(1 - \frac{4\varepsilon_2 L_{\max}}{\mu}\right)\cdot\frac{\mu}{4L_\gamma(1+\gamma L_{\max})}\right)^K \left(M^\gamma(x_0) - M^{\gamma_\star}\right).$$

Using smoothness of γL_γ, if we denote $\Delta_k = \|x_k - x_\star\|^2$ where $x_\star$ is a minimizer of both M^γ and f since we assume we are in the interpolation regime (Assumption 2.7), we have

$$\mathcal{E}_0 \leq \frac{\gamma L_\gamma}{2}\Delta_0.$$

Using star strong convexity (quadratic growth property), we have

$$\mathcal{E}_K \geq \frac{\gamma\mu}{2(1+\gamma L_{\max})}\Delta_K.$$

As a result, we can transform the above convergence guarantee into

$$\Delta_K \leq \left(1 - \left(1 - \frac{4\varepsilon_2 L_{\max}}{\mu}\right)\cdot\frac{\mu}{4L_\gamma(1+\gamma L_{\max})}\right)^K \cdot \frac{L_\gamma(1+\gamma L_{\max})}{\mu}\Delta_0.$$

This completes the proof.

[5] Theorem E.4 remains valid if we replace f being strongly convex with PL.

I.5 Proof of Theorem 5.1

Notice that we assume each f_i is L_i-smooth and convex. The local optimization of each client can be written as

$$\min_{z\in\mathbb{R}^d}\left\{A_{k,i}^{\gamma}(z)=f_i(z)+\frac{1}{2\gamma}\|z-x_k\|^2\right\},$$

It is easy to see that $A_{k,i}^{\gamma}(z)$ is $L_i+\frac{1}{\gamma}$-smooth and $\frac{1}{\gamma}$-strongly convex. We first provide the convergence theory of GD for reference.

Theory of GD: For a $\widehat{\mu}$-strongly convex, $\widehat{L}$-smooth function ϕ, the algorithm can be formulated as

$$z_{t+1}=z_t-\eta\nabla\phi(z_t), \tag{GD}$$

where z_t is the iterate in the t-th iteration, and $\eta>0$ is the step size. GD with step size $\eta\in(0,\frac{1}{\widehat{L}}]$ generates iterates that satisfy

$$\|z_t-z_\star\|^2\leq(1-\eta\widehat{\mu})^t\|z_0-z_\star\|^2,$$

where $z_\star$ is a minimizer of ϕ, t is the number of iterations (number of gradient evaluations).

Approximation satisfying Definition 3.1: Notice that $\mathrm{prox}_{\gamma f_i}(x_k)$ is the minimizer of $A_{k,i}^{\gamma}(z)$ and $z_0=x_k$. As a result, if we run GD with the largest step size $\frac{\gamma}{1+\gamma L_i}$,

$$\left\|z_t-\mathrm{prox}_{\gamma f_i}(x_k)\right\|^2\leq\left(1-\frac{1}{1+\gamma L_i}\right)^t\left\|x_k-\mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 \tag{42}$$

We have

$$t=\mathcal{O}\left((1+\gamma L_i)\log\left(\frac{\left\|x_k-\mathrm{prox}_{\gamma f_i}(x_k)\right\|^2}{\varepsilon_1}\right)\right).$$

The unknown term $\left\|x_k-\mathrm{prox}_{\gamma f_i}(x_k)\right\|^2$ within the log can be bounded by

$$\begin{aligned}\left\|x_k-\mathrm{prox}_{\gamma f_i}(x_k)\right\|^2&=\|z_0-z_\star\|^2\\&\leq\gamma^2\left\|\nabla A_{k,i}^{\gamma}(z_0)-\nabla A_{k,i}^{\gamma}(z_\star)\right\|^2=\|\gamma\nabla f_i(x_k)\|^2,\end{aligned} \tag{43}$$

which can be easily calculated.

Approximation satisfying Definition 4.1: According to (42), we have

$$t=\mathcal{O}\left((1+\gamma L_i)\log\left(\frac{1}{\varepsilon_2}\right)\right).$$

This completes the proof.

I.6 Proof of Theorem 5.2

We first provide the theory of AGD (Nesterov, 2004).

Theory of AGD: For a $\widehat{\mu}$-strongly convex, $\widehat{L}$-smooth function ϕ, the algorithm can be formulated as

$$\begin{aligned}y_{t+1}&=z_t+\alpha(z_t-z_{t-1})\\z_{t+1}&=y_{t+1}-\eta\nabla\phi(y_{t+1}),\end{aligned} \tag{AGD}$$

where z_t,y_t are iterates, $\eta>0$ is the step size, $\alpha>0$ is the momentum parameter. AGD with step size $\eta=\frac{1}{\widehat{L}}$, momentum $\alpha=\frac{\sqrt{\widehat{L}}-\sqrt{\widehat{\mu}}}{\sqrt{\widehat{L}}+\sqrt{\widehat{\mu}}}$ generates iterates that satisfy

$$\|z_t-z_\star\|^2\leq\frac{2\widehat{L}}{\widehat{\mu}}\cdot\left(1-\sqrt{\frac{\widehat{\mu}}{\widehat{L}}}\right)^t\|z_0-z_\star\|^2,$$

where $z_\star$ is a minimizer of ϕ, t is the number of iterations (number of gradient evaluations).

Approximation satisfying Definition 3.1: Notice that $\mathrm{prox}_{\gamma f_i}(x_k)$ is the minimizer of $A^{\gamma}_{k,i}(z)$ and $z_0 = x_k$. As a result, if we run AGD with the step size $\frac{\gamma}{1+\gamma L_i}$ and momentum $\alpha = \frac{\sqrt{1+\gamma L_i}-1}{\sqrt{1+\gamma L_i}+1}$,

$$\left\|z_t - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 \leq 2 \cdot (1+\gamma L_i)\left(1 - \frac{1}{\sqrt{1+\gamma L_i}}\right)^t \left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2. \tag{44}$$

We have

$$t = \mathcal{O}\left(\sqrt{1+\gamma L_i}\log\left(\frac{(1+\gamma L_i)\cdot\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2}{\varepsilon_1}\right)\right)$$

Similar to the proof of Theorem 5.1, since we have according to (43),

$$\left\|x_k - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 \leq \left\|\gamma \nabla f_i(x_k)\right\|^2,$$

it is straightforward to determine the number of local iterations needed.

Approximation satisfying Definition 4.1: Using (44), we have

$$t = \mathcal{O}\left(\sqrt{1+\gamma L_i}\log\left(\frac{1+\gamma L_i}{\varepsilon_2}\right)\right).$$

I.7 Proof of Theorem F.1

Here we do not assume the interpolation condition. Consequently, the true objective $M^{\gamma}(x)$ and $f(x)$ do not necessarily share the same minimizers. Denoting

$$\mathcal{X}_f := \arg\min_{x\in\mathbb{R}^d}\left\{f(x) = \frac{1}{n}\sum_{i=1}^n f_i(x)\right\}, \quad \mathcal{X}_M := \arg\min_{x\in\mathbb{R}^d}\left\{M^{\gamma}(x) = \frac{1}{n}\sum_{i=1}^n M^{\gamma}_{f_i}(x)\right\},$$

we have $\mathcal{X}_f \neq \mathcal{X}_M$ in general. Pick $x^M_\star \in \mathcal{X}_M$ and $x^\star_f \in \mathcal{X}_f$ in the following derivation. Following the proof of Theorem 3.2, even without interpolation, we may still apply Theorem D.2. For notational convenience, let μ_γ denote the strong convexity constant of M^{γ} and set $\Delta^M_k := \left\|x_k - x^M_\star\right\|^2$, $\Delta^f_k := \left\|x_k - x^f_\star\right\|^2$. Then, for $0 < \alpha \leq \frac{1}{4\gamma L_\gamma}$, we obtain

$$\mathcal{E}_K \leq \left(1 - \frac{\mu_\gamma}{32 L_\gamma}\right)^K \mathcal{E}_0 + \frac{6\varepsilon_1}{\gamma\mu_\gamma}.$$

To relate this to distances, note that by smoothness and convexity,

$$\mathcal{E}_0 = M^{\gamma}(x_0) - \inf M^{\gamma} \leq \frac{\gamma L_\gamma}{2}\left\|x_0 - x^M_\star\right\|^2, \qquad \mathcal{E}_K \geq \frac{\gamma\mu_\gamma}{2}\left\|x_K - x^M_\star\right\|^2.$$

This yields

$$\Delta^M_K \leq \left(1 - \frac{\mu_\gamma}{32 L_\gamma}\right)^K \frac{L_\gamma}{\mu_\gamma}\Delta^M_0 + \frac{6\varepsilon_1}{\gamma\mu_\gamma^2}.$$

However, this result is with respect to $x^M_\star$, the minimizer of M^{γ}. Since we are ultimately interested in $x^f_\star$, we use Young's inequality to relate them:

$$\Delta^f_K \leq 2\Delta^M_K + 2\left\|x^f_\star - x^M_\star\right\|^2.$$

Substituting the bound on Δ^M_K gives

$$\Delta^f_K \leq \left(1 - \frac{\mu_\gamma}{32 L_\gamma}\right)^K \frac{2L_\gamma}{\mu_\gamma}\Delta^f_0 + \frac{12\varepsilon_1}{\gamma\mu_\gamma^2} + 2\left\|x^f_\star - x^M_\star\right\|^2,$$

which completes the proof.

I.8 Proof of Remark F.3

We denote $y_i := \mathrm{prox}_{\gamma f_i}(x)$ for $x, y \in \mathbb{R}^d$. From the optimality condition of the proximal subproblem,

$$\nabla f_i(y_i) = \tfrac{1}{\gamma}(x - y_i).$$

Using the L_i-smoothness of f_i, we obtain

$$f_i(x) \le f_i(y_i) + \gamma\|\nabla f_i(y_i)\|^2 + \tfrac{\gamma^2 L_i}{2}\|\nabla f_i(y_i)\|^2.$$

Hence,

$$f_i(x) - M^{\gamma}_{f_i}(x) \le \tfrac{\gamma}{2}(1 + \gamma L_i)\|\nabla f_i(y_i)\|^2.$$

To relate $\|\nabla f_i(y_i)\|$ to $\|\nabla f_i(x)\|$, smoothness further yields

$$\|\nabla f_i(y_i)\| - \|\nabla f_i(x)\| \le \|\nabla f_i(y_i) - \nabla f_i(x)\| \le L_i\|x - y_i\| = \gamma L_i\|\nabla f_i(y_i)\|.$$

If γ is sufficiently small so that $\gamma L_i \le 1$, this implies

$$\|\nabla f_i(y_i)\|^2 \le \frac{1}{(1-\gamma L_i)^2}\|\nabla f_i(x)\|^2.$$

Substituting back, we obtain

$$f_i(x) - M^{\gamma}_{f_i}(x) \le \frac{\gamma(1+\gamma L_i)}{2(1-\gamma L_i)^2}\|\nabla f_i(x)\|^2.$$

Summing over $i \in [n]$,

$$f(x) - M^{\gamma}(x) \le \frac{1}{n}\sum_{i=1}^{n}\frac{\gamma(1+\gamma L_i)}{2(1-\gamma L_i)^2}\|\nabla f_i(x)\|^2.$$

Evaluating at $x^f_\star$, and applying μ-strong convexity of f, we get

$$\|x^M_\star - x^f_\star\|^2 \le \frac{2}{\mu}\big(f(x^M_\star) - f(x^f_\star)\big) \le \frac{2}{\mu}\big(f(x^f_\star) - M^{\gamma}(x^f_\star)\big).$$

Thus,

$$\|x^M_\star - x^f_\star\|^2 \le \frac{1}{n}\sum_{i=1}^{n}\frac{\gamma(1+\gamma L_i)}{\mu(1-\gamma L_i)^2}\|\nabla f_i(x^f_\star)\|^2 \le \frac{\gamma(1+\gamma L_{\max})}{\mu(1-\gamma L_{\max})^2}\cdot\sigma^2_\star,$$

where

$$\sigma^2_\star := \tfrac{1}{n}\sum_{i=1}^{n}\|\nabla f_i(x^f_\star)\|^2.$$

Hence, the neighborhood size $\|x^M_\star - x^f_\star\|^2$ depends both on the heterogeneity measure $\sigma^2_\star$ and the proximal parameter γ. It vanishes either when $\gamma \to 0$ or in the interpolation regime where $\sigma^2_\star = 0$.

I.9 Proof of Theorem F.5

We follow the proof of Theorem 4.3. Let us first analyze the compressor $\mathcal{C}$. We will verify it belongs to $\mathbb{B}^3(\delta)$. The inequality we want to prove can be written equivalently as

$$\left\|\gamma\nabla M^{\gamma}(x_k) - \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right) - \gamma\nabla M^{\gamma}(x_k)\right\|^2 \le \left(1 - \frac{1}{\delta}\right)\|\gamma\nabla M^{\gamma}(x_k)\|^2,$$

which is exactly

$$\left\|\frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2 \le \|\gamma\nabla M^{\gamma}(x_k)\|^2$$

For the left-hand side, using the convexity of $\|\cdot\|^2$ in combination with Definition 4.1, we obtain

$$\begin{aligned}\left\|\frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1}-\operatorname{prox}_{\gamma f_i}\left(x_k\right)\right)\right\|^2 &\leq \frac{1}{n}\sum_{i=1}^{n}\left\|\tilde{x}_{i,k+1}-\operatorname{prox}_{\gamma f_i}\left(x_k\right)\right\|^2 \\ &\leq \frac{\varepsilon_2}{n}\sum_{i=1}^{n}\left\|x_k-\operatorname{prox}_{\gamma f_i}\left(x_k\right)\right\|^2.\end{aligned}$$

Let $x_\star^M$ be a minimizer of M^γ, notice that it is also a minimizer of γM^γ since γ here can be viewed as a constant,

$$\begin{aligned}\frac{\varepsilon_2}{n}\sum_{i=1}^{n}\left\|x_k-\operatorname{prox}_{\gamma f_i}\left(x_k\right)\right\|^2 &\overset{(4)}{=} \frac{\varepsilon_2}{n}\sum_{i=1}^{n}\left\|\gamma\nabla M_{f_i}^{\gamma}\left(x_k\right)\right\|^2 \\ &= \frac{\varepsilon_2}{n}\sum_{i=1}^{n}\left\|\gamma\nabla M_{f_i}^{\gamma}\left(x_k\right)-\gamma\nabla M_{f_i}^{\gamma}\left(x_\star^M\right)\right\|^2 \\ &\leq \frac{2\varepsilon_2}{n}\sum_{i=1}^{n}\frac{\gamma L_i}{1+\gamma L_i}\left(\gamma M_{f_i}^{\gamma}\left(x_k\right)-\gamma M_{f_i}^{\gamma}\left(x_\star^M\right)\right) \\ &\leq \frac{2\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}\left(\gamma M^{\gamma}\left(x_k\right)-\gamma M^{\gamma}\left(x_\star^M\right)\right).\end{aligned}$$

By convexity, we have

$$\left(1-\frac{1}{\delta}\right)\left\|\gamma\nabla M^{\gamma}\left(x_k\right)\right\|^2 \geq \left(1-\frac{1}{\delta}\right)\frac{\gamma\mu_\gamma}{2}\left(\gamma M^{\gamma}\left(x_k\right)-\gamma M^{\gamma}\left(x_\star^M\right)\right).$$

The above inequalities allows us to get,

$$\frac{2\varepsilon_2\gamma L_{\max}}{1+\gamma L_{\max}}\left(\gamma M^{\gamma}\left(x_k\right)-\gamma M^{\gamma}\left(x_\star^M\right)\right) \leq \left(1-\frac{1}{\delta}\right)\frac{\gamma\mu_\gamma}{2}\left(\gamma M^{\gamma}\left(x_k\right)-\gamma M^{\gamma}\left(x_\star^M\right)\right).$$

We may verify that if we pick

$$\delta=\frac{\mu_\gamma}{\mu_\gamma-\frac{4\varepsilon_2 L_{\max}}{1+\gamma L_{\max}}}>0,$$

the condition is met. For this to hold, we require that $\varepsilon_2<\frac{\mu_\gamma(\gamma+\frac{1}{L_{\max}})}{4}$. Notice that our objective γM^γ is γL_γ-smooth and $\gamma\mu_\gamma$-convex. As long as $\alpha=\frac{1}{\gamma L_\gamma}$ and $\varepsilon_2<\frac{\mu_\gamma(\gamma+\frac{1}{L_{\max}})}{4}$, we have

$$\gamma M^{\gamma}\left(x_k\right)-\inf\gamma M^{\gamma} \leq \left(1-\frac{\mu_\gamma}{L_\gamma}\left(1-\frac{4\varepsilon_2 L_{\max}}{\mu_\gamma\left(1+\gamma L_{\max}\right)}\right)\right)^K\left(\gamma M^{\gamma}\left(x_0\right)-\inf\gamma M^{\gamma}\right).$$

Now we need to translate the result into squared distance, similar to Theorem F.1, notice that we will use $\Delta_k^M=\left\|x_k-x_\star^M\right\|^2$,

$$\mathcal{E}_0 \leq \frac{\gamma L_\gamma}{2}\Delta_0^M, \qquad \mathcal{E}_K \geq \frac{\gamma\mu_\gamma}{2}\Delta_K^M.$$

As a result, we can transform the above convergence guarantee into

$$\Delta_K^M \leq \left(1-\frac{\mu_\gamma}{L_\gamma}\left(1-\frac{4\varepsilon_2 L_{\max}}{\mu_\gamma\left(1+\gamma L_{\max}\right)}\right)\right)^K\cdot\frac{L_\gamma}{\mu_\gamma}\Delta_0^M.$$

We encounter the same issue as in Theorem F.1, namely that the descent is expressed in terms of Δ_k^M. Adopting a similar approach, we obtain

$$\Delta_K^f \leq 2\Delta_K^M+2\left\|x_\star^f-x_\star^M\right\|^2.$$

Substituting the bound on Δ_K^M gives

$$\Delta_K^f \leq \left(1-\frac{\mu_\gamma}{L_\gamma}\left(1-\frac{4\varepsilon_2 L_{\max}}{\mu_\gamma\left(1+\gamma L_{\max}\right)}\right)\right)^K\cdot\frac{2L_\gamma}{\mu_\gamma}\Delta_0^f+2\left\|x_\star^f-x_\star^M\right\|^2,$$

which completes the proof.

I.10 PROOF OF THEOREM H.1

In this case, the gradient estimator is defined as

$$g(x_k) = \frac{1}{\tau}\sum_{i\in S_k}\left(\gamma\nabla M_{f_i}^{\gamma}(x_k) - \left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right). \tag{45}$$

Notice that we have

$$\begin{aligned}
&\langle \gamma\nabla M^{\gamma}(x_k), \mathbb{E}[g(x_k)]\rangle \\
&= \left\langle \gamma\nabla M^{\gamma}(x_k), \mathbb{E}\left[\frac{1}{\tau}\sum_{i\in S_k}\gamma\nabla M_{f_i}^{\gamma}(x_k) - \frac{1}{\tau}\sum_{i\in S_k}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right]\right\rangle \\
&= \left\langle \gamma\nabla M^{\gamma}(x_k), \gamma\nabla M^{\gamma}(x_k) - \frac{1}{n}\sum_{i=1}^{n}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\rangle.
\end{aligned}$$

Using the same technique in the proof of Theorem 4.2, we are able to obtain that

$$\langle \gamma\nabla M^{\gamma}(x_k), \mathbb{E}[g(x_k)]\rangle \geq \left(1 - \frac{2\sqrt{\varepsilon_2}L_{\max}}{\mu}\right)\cdot\left\|\gamma\nabla M^{\gamma}(x_k)\right\|^2.$$

Thus, as long as we pick $\varepsilon_2 < \frac{\mu^2}{4L_{\max}^2}$, we can pick $b = 1 - \sqrt{\varepsilon_2}\cdot\frac{2L_{\max}}{\mu}$ and $c = 0$. We then compute $\mathbb{E}\left[\|g(x_k)\|^2\right]$,

$$\begin{aligned}
\mathbb{E}\left[\|g(x_k)\|^2\right] &= \mathbb{E}\left[\left\|\frac{1}{\tau}\sum_{i\in S_k}\gamma\nabla M_{f_i}^{\gamma}(x_k) - \frac{1}{\tau}\sum_{i\in S_k}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2\right] \\
&= \underbrace{\mathbb{E}\left[\left\|\frac{1}{\tau}\sum_{i\in S_k}\gamma\nabla M_{f_i}^{\gamma}(x_k)\right\|^2\right]}_{:=T_1} + \underbrace{\mathbb{E}\left[\left\|\frac{1}{\tau}\sum_{i\in S_k}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\|^2\right]}_{:=T_2} \\
&\quad \underbrace{-2\mathbb{E}\left[\left\langle\frac{1}{\tau}\sum_{i\in S_k}\gamma\nabla M_{f_i}^{\gamma}(x_k), \frac{1}{\tau}\sum_{i\in S_k}\left(\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right)\right\rangle\right]}_{:=T_3}.
\end{aligned}$$

We try to provide upper bounds for those terms separately.

Term T_1: We have

$$T_1 = \frac{n-\tau}{\tau(n-1)}\cdot\frac{1}{n}\sum_{i=1}^{n}\left\|\gamma\nabla M_{f_i}^{\gamma}(x_k)\right\|^2 + \frac{n(\tau-1)}{\tau(n-1)}\cdot\left\|\gamma\nabla M^{\gamma}(x_k)\right\|^2.$$

Using smoothness of $\gamma M_{f_i}^{\gamma}$ and the fact that we are in the interpolation regime, we have

$$\begin{aligned}
T_1 &= \frac{n-\tau}{\tau(n-1)}\cdot\frac{1}{n}\sum_{i=1}^{n}\left\|\gamma\nabla M_{f_i}^{\gamma}(x_k) - \gamma\nabla M_{f_i}^{\gamma}(x_\star)\right\|^2 + \frac{n(\tau-1)}{\tau(n-1)}\cdot\left\|\gamma\nabla M^{\gamma}(x_k)\right\|^2 \\
&\leq \frac{n-\tau}{\tau(n-1)}\cdot\frac{1}{n}\sum_{i=1}^{n}\frac{2\gamma L_i}{1+\gamma L_i}\cdot\left(\gamma M_{f_i}^{\gamma}(x_k) - \gamma\left(M_{f_i}^{\gamma}\right)_{\inf}\right) + \frac{n(\tau-1)}{\tau(n-1)}\cdot\left\|\gamma\nabla M^{\gamma}(x_k)\right\|^2 \\
&\leq \frac{n-\tau}{\tau(n-1)}\cdot\frac{2\gamma L_{\max}}{1+\gamma L_{\max}}\cdot\left(\gamma M^{\gamma}(x_k) - \gamma M_{\inf}^{\gamma}\right) + \frac{n(\tau-1)}{\tau(n-1)}\cdot\left\|\gamma\nabla M^{\gamma}(x_k)\right\|^2.
\end{aligned} \tag{46}$$

Term T_2: It is easy to see that using convexity of the squared Euclidean norm, we have

$$\begin{aligned}
T_2 &\leq \mathbb{E}\left[\frac{1}{\tau}\sum_{i\in S_k}\left\|\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2\right] \\
&= \frac{1}{n}\sum_{i=1}^{n}\left\|\tilde{x}_{i,k+1} - \mathrm{prox}_{\gamma f_i}(x_k)\right\|^2 \overset{(12)}{\leq} \frac{\varepsilon_2}{n}\sum_{i=1}^{n}\left\|\gamma\nabla M_{f_i}^{\gamma}(x_k)\right\|^2.
\end{aligned}$$

Using smoothness of each individual $\gamma M_{f_i}^{\gamma}(x_k)$ and the fact we are in the interpolation regime, we have

$$T_2 \leq \frac{2\varepsilon_2 \gamma L_{\max}}{1+\gamma L_{\max}} \left(\gamma M^{\gamma}(x_k) - \gamma M_{\inf}^{\gamma}\right). \tag{47}$$

Term T_3: We have

$$\begin{aligned} T_3 = &-2 \cdot \frac{n-\tau}{\tau(n-1)} \cdot \frac{1}{n} \sum_{i=1}^{n} \left\langle \gamma \nabla M_{f_i}^{\gamma}(x_k), \tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k) \right\rangle \\ &- 2 \cdot \frac{n(\tau-1)}{\tau(n-1)} \cdot \left\langle \gamma \nabla M^{\gamma}(x_k), \frac{1}{n} \sum_{i=1}^{n} \left(\tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k)\right) \right\rangle. \end{aligned}$$

Using Cauchy-Schwarz inequality and convexity, we further obtain

$$\begin{aligned} T_3 \leq &\, 2 \cdot \frac{n-\tau}{\tau(n-1)} \cdot \frac{1}{n} \sum_{i=1}^{n} \left\| \gamma \nabla M_{f_i}^{\gamma}(x_k) \right\| \left\| \tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k) \right\| \\ &+ 2 \cdot \frac{n(\tau-1)}{\tau(n-1)} \left\| \gamma \nabla M^{\gamma}(x_k) \right\| \cdot \frac{1}{n} \sum_{i=1}^{n} \left\| \tilde{x}_{i,k+1} - \operatorname{prox}_{\gamma f_i}(x_k) \right\|. \end{aligned}$$

Using similar approaches in the previous paragraphs, we have

$$\begin{aligned} T_3 &\overset{(12)}{\leq} \frac{2(n-\tau)}{\tau(n-1)} \cdot \frac{\sqrt{\varepsilon_2}}{n} \sum_{i=1}^{n} \left\| \gamma \nabla M_{f_i}^{\gamma}(x_k) \right\|^2 \\ &\quad + \frac{2n(\tau-1)}{\tau(n-1)} \left\| \gamma M^{\gamma}(x_k) \right\| \frac{\sqrt{\varepsilon_2}}{n} \cdot \sum_{i=1}^{n} \left\| \gamma \nabla M_{f_i}^{\gamma}(x_k) \right\| \\ &\leq \frac{2(n-\tau)}{\tau(n-1)} \cdot \frac{\sqrt{\varepsilon_2}}{n} \sum_{i=1}^{n} \left\| \gamma \nabla M_{f_i}^{\gamma}(x_k) - \gamma \nabla M_{f_i}^{\gamma}(x_\star) \right\|^2 \\ &\quad + \frac{2n(\tau-1)}{\tau(n-1)} \left\| \gamma M^{\gamma}(x_k) \right\| \frac{\sqrt{\varepsilon_2}}{n} \cdot \sum_{i=1}^{n} \left\| \gamma \nabla M_{f_i}^{\gamma}(x_k) - \gamma \nabla M_{f_i}^{\gamma}(x_k) \right\| \\ &\leq \frac{4\sqrt{\varepsilon_2}(n-\tau)}{\tau(n-1)} \cdot \frac{\gamma L_{\max}}{1+\gamma L_{\max}} \left(\gamma M^{\gamma}(x_k) - \gamma M_{\inf}^{\gamma}\right) \\ &\quad + \frac{4\sqrt{\varepsilon_2} n(\tau-1)}{\tau(n-1)} \cdot \frac{\gamma L_{\max}}{1+\gamma L_{\max}} \left\| x_k - x_\star \right\| \left\| \gamma \nabla M^{\gamma}(x_k) \right\| \\ &\overset{(24)}{\leq} \frac{4\sqrt{\varepsilon_2}(n-\tau)}{\tau(n-1)} \cdot \frac{\gamma L_{\max}}{1+\gamma L_{\max}} \left(\gamma M^{\gamma}(x_k) - \gamma M_{\inf}^{\gamma}\right) \\ &\quad + \frac{4\sqrt{\varepsilon_2} n(\tau-1)}{\tau(n-1)} \cdot \frac{L_{\max}}{\mu} \left\| \gamma \nabla M^{\gamma}(x_k) \right\|^2. \end{aligned} \tag{48}$$

Combining (46), (47) and (48), we have

$$\begin{aligned} \sum_{i=1}^{3} T_i \leq &\, 2\left(\varepsilon_2 + \frac{2\sqrt{\varepsilon_2}(n-\tau)}{\tau(n-1)} + \frac{(n-\tau)}{\tau(n-1)}\right) \cdot \frac{\gamma L_{\max}}{1+\gamma L_{\max}} \cdot \left(\gamma M^{\gamma}(x_k) - \gamma M_{\inf}^{\gamma}\right) \\ &+ \left(\frac{n(\tau-1)}{\tau(n-1)} + \frac{4\sqrt{\varepsilon_2} n(\tau-1)}{\tau(n-1)}\right) \cdot \frac{L_{\max}}{\mu} \cdot \left\| \gamma M^{\gamma}(x_k) \right\|^2. \end{aligned} \tag{49}$$

Therefore, it is easy to see that we can pick

$$\begin{aligned} A &= \left(\varepsilon_2 + \frac{2\sqrt{\varepsilon_2}(n-\tau)}{\tau(n-1)} + \frac{(n-\tau)}{\tau(n-1)}\right) \cdot \frac{\gamma L_{\max}}{1+\gamma L_{\max}} \\ B &= \left(\frac{n(\tau-1)}{\tau(n-1)} + \frac{4\sqrt{\varepsilon_2} n(\tau-1)}{\tau(n-1)}\right) \cdot \frac{L_{\max}}{\mu}, \qquad C = 0. \end{aligned}$$

Applying Theorem 4 of Demidovich et al. (2024), we list the corresponding values of $A, B, C, b, c \geq 0$ below,

$$A = \frac{\gamma L_{\max}}{1+\gamma L_{\max}}\left(\varepsilon_2 + \frac{2\sqrt{\varepsilon_2}\,(n-\tau)}{\tau\,(n-1)} + \frac{(n-\tau)}{\tau\,(n-1)}\right)$$
$$B = \frac{n\,(\tau-1)}{\tau\,(n-1)}\left(1+\frac{4\sqrt{\varepsilon_2}L_{\max}}{\mu}\right), \quad C = 0$$
$$b = \frac{\mu - 2\sqrt{\varepsilon_2}L_{\max}}{\mu}, \quad c = 0.$$

We know that the PL constant of γM^γ is given by $\frac{\gamma\mu}{4(1+\gamma L_{\max})}$ and the corresponding smoothness constant is γL_γ. As a result, when $\alpha > 0$ satisfies

$$\alpha < \underbrace{\frac{1}{\gamma L_\gamma}\cdot\frac{\mu - 2\sqrt{\varepsilon_2}L_{\max}}{\mu + 4\varepsilon_2 L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} + \frac{n-\tau}{\tau(n-1)}\cdot\left(4L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} - \mu\right)}}_{:=B_1'},$$

and

$$\alpha < \underbrace{\frac{4\,(1+\gamma L_{\max})}{\gamma\left(\mu - 2\sqrt{\varepsilon_2}L_{\max}\right)}}_{=B_2},$$

we can obtain a convergence guarantee for the algorithm. Notice that $B_1' \leq B_1 < B_2$[6], thus we can further simplify the range of α to

$$\alpha \leq \underbrace{\frac{1}{\gamma L_\gamma}\cdot\frac{\mu - 2\sqrt{\varepsilon_2}L_{\max}}{\mu + 4\varepsilon_2 L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} + \frac{n-\tau}{\tau(n-1)}\cdot\left(4L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} - \mu\right)}}_{:=B_1'}.$$

Given that we select α properly, we have

$$\mathbb{E}[\mathcal{E}_K] \leq \left(1 - \alpha\cdot\frac{\gamma\left(\mu - 2\sqrt{\varepsilon_2}L_{\max}\right)}{4\,(1+\gamma L_{\max})}\right)^K \mathcal{E}_0.$$

Specifically, if we choose the largest α possible, we have

$$\mathbb{E}[\mathcal{E}_K] \leq \left(1 - \frac{\mu}{4L_\gamma\,(1+\gamma L_{\max})}\cdot S\left(\varepsilon_2, \tau\right)\right)^K \mathcal{E}_0,$$

where $S\left(\varepsilon_2, \tau\right)$ is defined as

$$S\left(\varepsilon_2, \tau\right) = \frac{\left(\mu - 2\sqrt{\varepsilon_2}L_{\max}\right)\left(1 - 2\sqrt{\varepsilon_2}\frac{L_{\max}}{\mu}\right)}{\mu + 4\varepsilon_2 L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} + \frac{n-\tau}{\tau(n-1)}\cdot\left(4L_{\max} + 4\sqrt{\varepsilon_2}L_{\max} - \mu\right)},$$

satisfying

$$0 < S\left(\varepsilon_2, \tau\right) \leq 1.$$

Using smoothness of γL_γ, if we denote $\Delta_k = \|x_k - x_\star\|^2$ where $x_\star$ is a minimizer of both M^γ and f since we assume we are in the interpolation regime (Assumption 2.7), we have

$$\mathcal{E}_0 \leq \frac{\gamma L_\gamma}{2}\Delta_0.$$

Using star strong convexity (quadratic growth property), we have

$$\mathcal{E}_K \geq \frac{\gamma\mu}{2\,(1+\gamma L_{\max})}\Delta_K.$$

As a result, we can transform the above convergence guarantee into

$$\mathbb{E}[\Delta_K] \leq \left(1 - \frac{\mu}{4L_\gamma\,(1+\gamma L_{\max})}\cdot S\left(\varepsilon_2, \tau\right)\right)^K \cdot \frac{L\gamma\,(1+\gamma L_{\max})}{\mu}\Delta_0.$$

This completes the proof.

[6]The definition of B_1 is given in (36)

J EXPERIMENTS

In this section, we report results from both controlled numerical experiments and large-scale deep learning experiments, designed to validate our theoretical findings and demonstrate the practical effectiveness of the proposed method. Code availability: `https://anonymous.4open.science/r/Inexact-FedExProx-code-E783/`.

J.1 NUMERICAL EXPERIMENTS

We first describe the settings for the numerical experiments and the corresponding results to validate our theoretical findings. We are interested in the following optimization problem in the distributed setting,

$$\min_{x\in\mathbb{R}^d}\left\{f(x)=\frac{1}{n}\sum_{i=1}^{n}f_i\left(x\right)\right\}.$$

Here n denotes the number of clients, d is the dimension, each function $f_i:\mathbb{R}^d\mapsto\mathbb{R}$ has the following form

$$f_i(x)=\frac{1}{2}x^\top\boldsymbol{A}_ix+b_i^\top x+c_i,$$

where $\boldsymbol{A}_i\in\mathbb{S}_+^d, b_i\in\mathbb{R}^d, c_i\in\mathbb{R}$. Specifically, we pick $n=20$ and $d=300$ for the experiments. Notice that we have

$$\nabla f_i(x)=\boldsymbol{A}_ix-b_i;\qquad \nabla^2 f_i(x)=\boldsymbol{A}_i\succeq\boldsymbol{O}_d,$$

which suggests that each f_i is convex and smooth. We can easily compute that in this case, we have

$$\mathrm{prox}_{\gamma f_i}\left(x\right)=\left(\boldsymbol{A}_i+\frac{1}{\gamma}\boldsymbol{I}_d\right)^{-1}\left(\frac{1}{\gamma}x-b_i\right).$$

Comparison of FedProx, FedExProx, FedExProx with absolute and relative approximation. In this section, we compare the convergence of FedProx, FedExProx, and FedExProx with absolute and relative approximations. For FedProx, the server extrapolation is fixed at 1, while for FedExProx we set $\alpha=\frac{1}{\gamma L_\gamma}$, both assuming exact proximal evaluations. For the inexact variants, we fix ε_1 and ε_2 to prescribed values and set α to the corresponding optimal choice. To study the impact of the local step size, we vary $\gamma\in\{1000,100,10,1,0.1,0.01\}$, first with $\varepsilon_1=10^{-3},\varepsilon_2=10^{-2}$, and then with $\varepsilon_1=10^{-6},\varepsilon_2=10^{-3}$.

Notably in Figure 2 and Figure 3, in all cases, FedExProx with absolute approximation exhibits the poorest performance and converges only to a neighborhood of the solution. This is expected, since the bias in this case does not go to zero as the algorithm progresses. It is worth mentioning that as the local step size γ increases, the size of the neighborhood decreases, which supports our claim in Theorem 3.2. As anticipated, in all cases, FedExProx outperforms FedProx due to server extrapolation. The performance of FedExProx with relative approximation is surprisingly good, outperforming FedProx in several cases.

Comparison of FedExProx with absolute approximation under different inaccuracies. In this section, we compare FedExProx with absolute approximations under varying inaccuracies. We fix $\gamma\in\{0.1,1,10\}$ and, for each value, test $\varepsilon_1\in\{0.001,0.005,0.01,0.05,0.1\}$.

As observed in Figure 4, the size of the neighborhood increases with ε_1, further corroborating our theoretical findings in Theorem 3.2.

Comparison of FedExProx with relative approximation under different inaccuracies. In this section, we compare FedExProx with relative approximations under varying inaccuracies. We fix $\gamma\in\{0.1,0.05,0.01\}$ and, for each value, test $\varepsilon_2\in\{0.001,0.005,0.01,0.05,0.1\}$.

As observed in Figure 5, in all cases, a smaller ε_2 corresponds to faster convergence of the algorithm. All the tested algorithm converges to the exact solution linearly, which validates the effectiveness of the proposed technique of relative approximation to reduce the bias term.

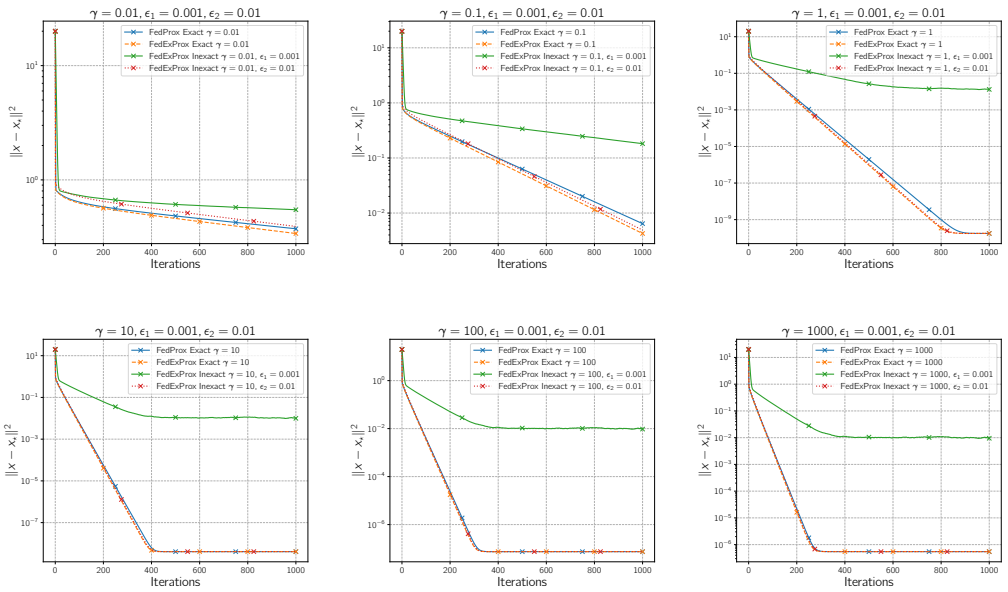

Figure 2: Comparison of FedProx, FedExProx with exact proximal evaluations, FedExProx with ε_1-absolute approximation and FedExProx with ε_2-relative approximation. In this case, we fix $\varepsilon_1 = 0.001$, $\varepsilon_2 = 0.01$ and pick the local step size $\gamma \in \{1000, 100, 10, 1, 0.1.0.01\}$. The y-axis is the squared distance to the minimizer of f, and the x-axis denotes the iterations.

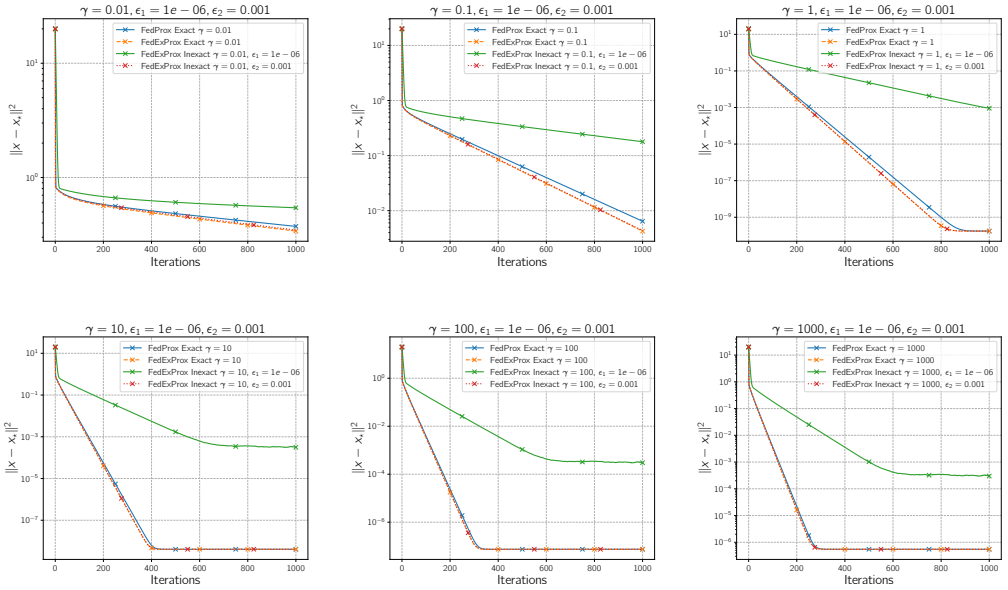

Figure 3: Comparison of FedProx, FedExProx with exact proximal evaluations, FedExProx with ε_1-absolute approximation and FedExProx with ε_2-relative approximation. In this case, we fix $\varepsilon_1 = 1e-6$, $\varepsilon_2 = 0.001$ and pick the local step size $\gamma \in \{1000, 100, 10, 1, 0.1.0.01\}$. The y-axis is the squared distance to the minimizer of f, and the x-axis denotes the iterations.

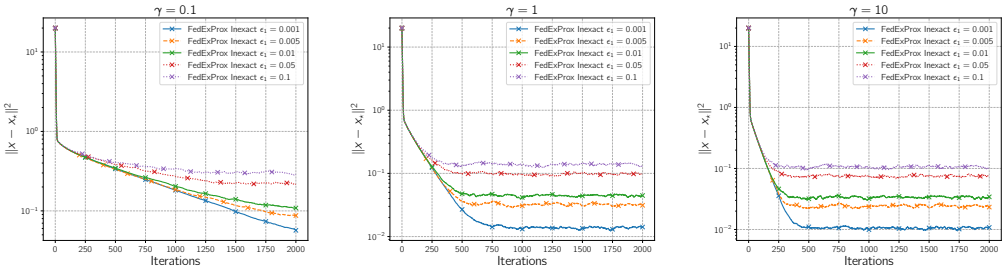

Figure 4: Comparison of FedExProx with ε_1-absolute approximation under different level of inexactness. We select γ from the set $\{0.1, 1, 10\}$ and for each choice of γ, we select ε_1 from the set $\{0.001, 0.005, 0.01, 0.05, 0.1\}$. The y-axis denotes the squared distance to the minimizer and the x-axis is the number of iterations.

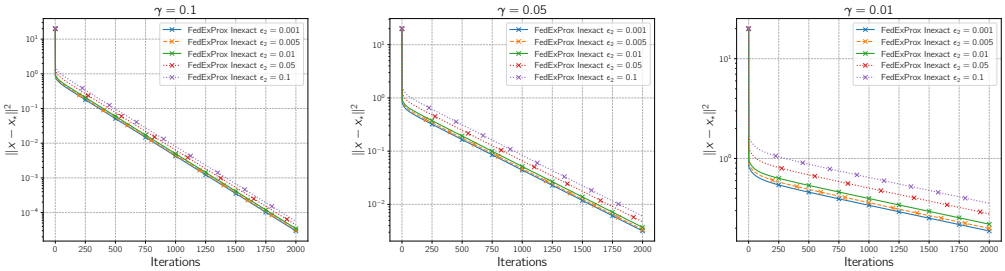

Figure 5: Comparison of FedExProx with ε_2-relative approximation under different level of inexactness. We select γ from the set $\{0.01, 0.05, 0.1\}$ and for each choice of γ, we select ε_2 from the set $\{0.001, 0.005, 0.01, 0.05, 0.1\}$. The y-axis denotes the squared distance to the minimizer and the x-axis is the number of iterations.

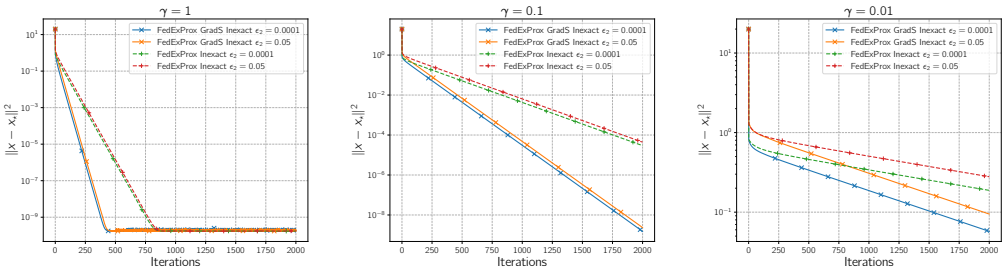

Figure 6: Comparison of FedExProx with ε_2-relative approximation under different level of inexactness using gradient diversity based extrapolation (FedExProx-GraDS). we select γ from the set $\{1, 0.1, 0.01\}$ and for each choice of γ, we select ε_2 from the set $\{0.0001, 0.05\}$. The y-axis denotes the squared distance to the minimizer and the x-axis is the number of iterations.

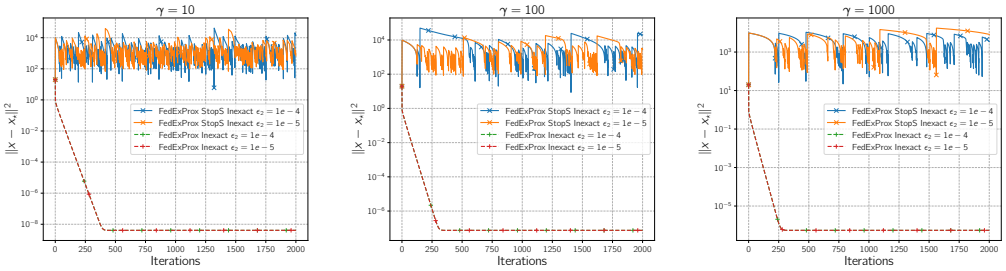

Figure 7: Comparison of FedExProx with ε_2-relative approximation under different level of inexactness using Polyak step size based extrapolation (FedExProx-StoPS). we select γ from the set $\{10, 100, 1000\}$ and for each choice of γ, we select ε_2 from the set $\{1e-4, 1e-5\}$. The y-axis denotes the squared distance to the minimizer and the x-axis is the number of iterations.

Adaptive extrapolation for inexact proximal evaluations In this section, we study adaptive extrapolation for FedExProx with relative approximations. We exclude the absolute case, since it converges only to a neighborhood and interacts poorly with adaptive step sizes such as gradient diversity and Polyak steps.

We adopt the following definitions. Gradient diversity based extrapolation:

$$\alpha_k = \alpha_{k,G} := \frac{1 + \gamma L_{\max}}{\gamma L_{\max}} \cdot \frac{\frac{1}{n}\sum_{i=1}^{n} \left\| x_k - \operatorname{prox}_{\gamma f_i}(x_k) \right\|^2}{\left\| \frac{1}{n}\sum_{i=1}^{n} \left(x_k - \operatorname{prox}_{\gamma f_i}(x_k) \right) \right\|^2}, \tag{GradS}$$

and Polyak-type extrapolation:

$$\alpha_k = \alpha_{k,S} := \frac{\frac{1}{n}\sum_{i=1}^{n} \left(M_{f_i}^{\gamma}(x_k) - \inf M^{\gamma} f_i \right)}{\gamma \left\| \frac{1}{n} \sum i = 1^n \nabla M_{f_i}^{\gamma}(x_k) \right\|^2}. \tag{StoPS}$$

As shown in Figure 6, gradient-diversity-based extrapolation consistently accelerates convergence, suggesting potential for adaptive strategies in our methods. In contrast, Figure 7 shows that Polyak-type extrapolation leads to divergence, highlighting the complexity of designing adaptive step sizes for SGD with biased updates or CGD with biased compression. To our knowledge, this remains an open problem deserving further study.

J.2 Deep Learning Experiments

To further demonstrate the practical relevance of our approach, we extend our evaluation to standard federated learning benchmarks. In particular, we train ResNet18 on CIFAR-10, partitioning the

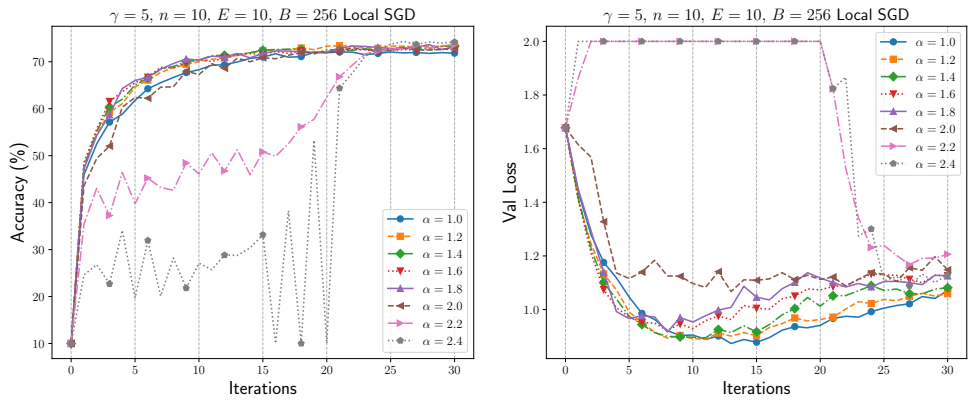

Figure 8: Comparison of FedProx ($\alpha = 1.0$) and FedExProx with extrapolation parameters $\alpha \in \{1.2, \ldots, 2.4\}$. The local step size is 5, with each client performing 10 local SGD steps per iteration to approximate the proximal solution. Validation loss is truncated at 2 for clarity.

Table 3: Comparison of FedProx ($\alpha = 1.0$) and FedExProx with varying extrapolation parameters. This table corresponds to Figure 8. The column "Iter @ 73% Acc" reports the number of iterations required to reach 73% accuracy, "Percentage" indicates the proportion relative to the baseline, and "Acc @ Iter 50" gives the accuracy at iteration 50.

Method	**Iter @ 73% Acc**	**Percentage**	**Acc @ Iter 50**
$\alpha = 1.0$ (FedProx)	43	100.00%	73.24%
$\alpha = 1.2$	21	48.84%	73.84%
$\alpha = 1.4$	27	62.79%	73.50%
$\alpha = 1.6$	35	81.40%	73.88%
$\alpha = 1.8$	23	53.49%	73.92%
$\alpha = 2.0$	31	72.09%	74.19%
$\alpha = 2.2$	28	65.12%	74.18%
$\alpha = 2.4$	26	60.47%	75.04%

training set randomly across n clients. Local subproblems are solved approximately using either Adam or SGD, with each client performing a limited number of local steps E to reflect realistic computational budgets. Throughout the experiments, we vary the number of local steps E, the step size γ, the extrapolation parameter α, and the local batch size B for each client. All experiments are implemented in PyTorch and executed on NVIDIA A100 GPUs.

As shown in Figure 8 and Table 3, extrapolation consistently accelerates the convergence of FedProx, with the best performance observed for α in the range 1.4 to 1.8. Larger values of α may cause temporary instability, yet they still converge after sufficient iterations, highlighting the robustness of the method. The intuition is that in the early stages of training, local models are highly inconsistent across clients, and extrapolation amplifies these discrepancies, causing the global model to oscillate. As training progresses and the local models become more aligned, extrapolation yields more aggressive and effective updates, driving faster convergence toward the global solution. This observation suggests the use of an extrapolation scheduler, starting with small α to maintain stability and gradually increasing it to accelerate convergence.

As a contrast, in Figure 9 we reduce the local step size from $\gamma = 5.0$ to $\gamma = 0.5$, which effectively increases the weight of the quadratic regularizer in the proximal operator. In this regime, client updates become more conservative, making smaller adjustments at each iteration. Under this constraint, the extrapolated variant consistently achieves clear gains, significantly outperforming the standard FedProx baseline. Interestingly, in this setting the optimal extrapolation parameter appears to be larger than 2, as performance continues to improve when α is increased from 1 to 1.8.

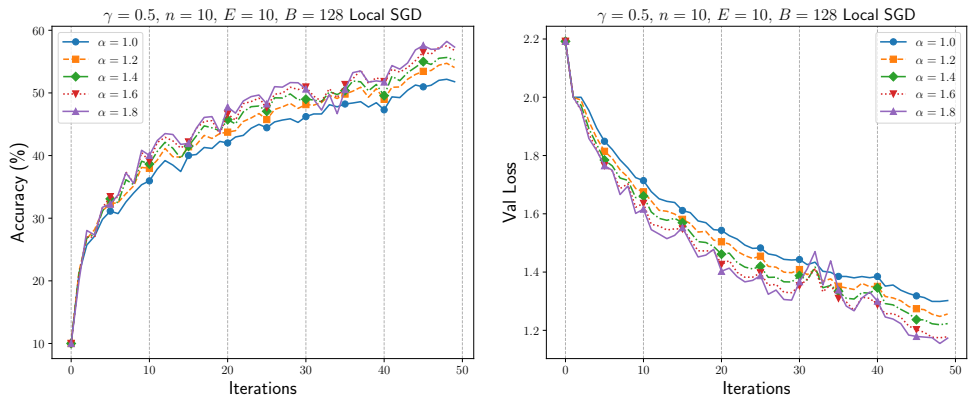

Figure 9: Comparison of FedProx ($\alpha = 1.0$) and FedExProx with extrapolation parameters $\alpha \in \{1.2, \ldots, 1.8\}$. The local step size is 0.01, with each client performing 10 local SGD steps per iteration to approximate the proximal solution. In this case, we are reducing the proximal stepsize γ from 5.0 to 0.5.

Table 4: Comparison of FedProx ($\alpha = 1.0$) and FedExProx with varying extrapolation parameters. This table corresponds to Figure 9. The columns have the same meaning as Table 3

Method	**Iter @ 50% Acc**	**Percentage**	**Acc @ Iter 50**
$\alpha = 1.0$ (FedProx)	44	100.00%	73.24%
$\alpha = 1.2$	37	84.09%	73.84%
$\alpha = 1.4$	31	70.45%	73.50%
$\alpha = 1.6$	29	65.91%	73.88%
$\alpha = 1.8$	28	63.64%	73.92%

We also replace the local SGD optimizer on each client with Adam, as shown in Figure 10. In this setting, the benefits of using an extrapolation parameter in the range $[1.2, 1.8]$ remain consistent, whereas larger values may lead to divergence.

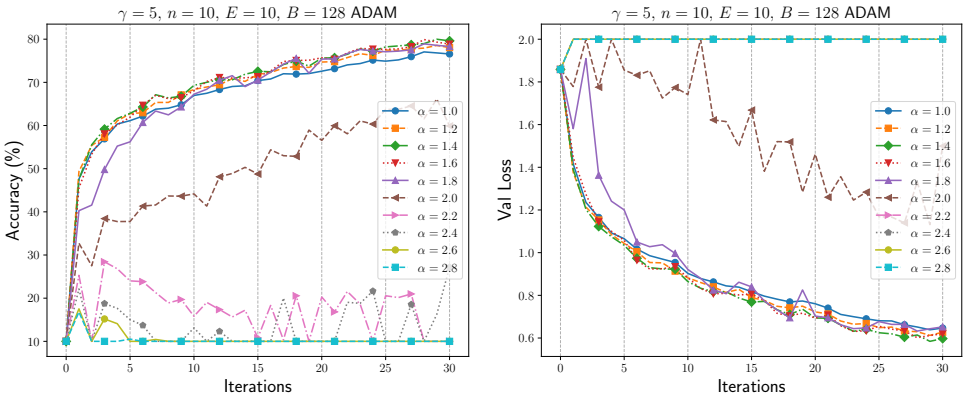

Figure 10: Comparison of FedProx ($\alpha = 1.0$) and FedExProx with extrapolation parameters $\alpha \in \{1.2, \ldots, 2.8\}$. The local step size is 0.001, with each client performing 10 Adam steps per iteration to approximate the proximal solution. Validation loss is truncated at 2 for clarity.