BRIDGING THE GAP BETWEEN HOMOGENEOUS AND HETEROGENEOUS ASYNCHRONOUS OPTIMIZATION IS SURPRISINGLY DIFFICULT

Anonymous authors
Paper under double-blind review

ABSTRACT

Modern large-scale machine learning tasks often require multiple workers, devices, CPUs, or GPUs to compute stochastic gradients in parallel and asynchronously to train model weights. Theoretical results typically distinguish between two settings: (i) the homogeneous setting, where all workers have access to the *same* data distribution, and (ii) the heterogeneous setting, where each worker operates on *different* data distributions. Known optimal time complexities in these settings reveal a significant gap, with far more pessimistic guarantees in the heterogeneous case. In this work, we investigate whether these pessimistic optimal time complexities can be overcome under different assumptions. Surprisingly, we show that improvement is provably impossible under widely used first- and second-order similarity assumptions for any randomized algorithm. We then turn to the interpolation regime and demonstrate that the weak interpolation assumption alone is also insufficient. Finally, we introduce a minimal combination of irreducible assumptions, strong interpolation and the local Polyak-Łojasiewicz condition, to derive a new time complexity bound that matches the dependence on worker computation times in the best-known result in the homogeneous setting, without requiring identical data distributions.

1 INTRODUCTION

We consider optimization problems described by

$$\min_{x\in\mathbb{R}^d}\left\{f(x) := \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}_{\xi_i\sim\mathcal{D}_i}\left[f_i(x;\xi_i)\right]\right\}, \tag{1}$$

where $f_i : \mathbb{R}^d \times \mathbb{S}_{\xi_i} \to \mathbb{R}$ and ξ_i is a random variable with distribution $\mathcal{D}_i$ on $\mathbb{S}_{\xi_i}$ for all $i \in [n]$. Let us denote $f_i(x) := \mathbb{E}_{\xi_i\sim\mathcal{D}_i}\left[f_i(x;\xi_i)\right]$. In our setup, we have n workers/clients/CPUs/GPUs working in parallel and asynchronously, and each worker i has access only to the stochastic gradient $\nabla f_i(x;\xi_i)$ of the function f_i for all $x \in \mathbb{R}^d$. We concentrate on the standard convergence metric and want to find a (possibly random) point $\bar{x}$ such that $\mathbb{E}[\|\bar{x} - x_*\|^2] \leq \varepsilon$, where x_* is a solution of (1). Such a problem arises in many machine learning (ML), deep learning, federated learning (FL), and data science problems (Konečný et al., 2016; McMahan et al., 2017; Goodfellow et al., 2016). In general, we use the following standard assumptions from convex stochastic optimization, but each result states which assumptions it requires.

Assumption 1.1 (Global smoothness)**.** The function f is differentiable and L–smooth, i.e., $\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|$ for all $x, y \in \mathbb{R}^d$.

Assumption 1.2 (Local smoothness)**.** The functions f_i are differentiable and L_i–smooth. We also define $L_{\max} := \max_{i\in[n]} L_i$. Note that $L \leq L_{\max}$ (This is why we distinguish Assum. 1.1 and 1.2).

Assumption 1.3 (Convexity)**.** The functions f_i are convex for all $i \in [n]$. The function f attains a minimum at a (non-unique) point $x_* \in \mathbb{R}^d$.

Assumption 1.4 (Unbiased and σ^2-variance-bounded noise)**.** For all $x \in \mathbb{R}^d$, stochastic gradients $\nabla f_i(x;\xi)$ are unbiased and σ^2-variance-bounded, i.e., $\mathbb{E}_{\xi_i}[\nabla f_i(x;\xi_i)] = \nabla f_i(x)$ and $\mathbb{E}_{\xi_i}[\|\nabla f_i(x;\xi_i) - \nabla f_i(x)\|^2] \leq \sigma^2$ for all $i \in [n]$, where $\sigma^2 \geq 0$.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

We also consider the case where f satisfies the PŁ-condition, which is a much weaker assumption than μ–strong convexity (Karimi et al., 2016):

Assumption 1.5 (Global Polyak-Łojasiewicz condition)**.** There exists $\mu > 0$ such that $\|\nabla f(x)\|^2 \geq 2\mu\left(f(x) - f^*\right)$ for all $x \in \mathbb{R}^d$, where f^* is the finite optimal function value of f.

We focus on the modern setup where many workers work together in a distributed environment, where the workers can have *arbitrarily computation behaviors* due to hardware delays or network connectivity problems. Most previous works typically assume that the workers have the same performance that does not change over time. In contrast, our focus is on the setting where the *computation times are heterogeneous* and non-constant.

In the literature, the optimization problem (1) in the asynchronous environment is considered in two regimes: i) *heterogeneous setting*, where the functions f_i can be arbitrarily different; in the context of ML and FL, it means the workers have access to different datasets. ii) *homogeneous setting*, where the functions f_i are equal; in the context of ML and FL, it means the workers have access to the same dataset (Koloskova et al., 2022; Mishchenko et al., 2022; Feyzmahdavian & Johansson, 2023).

Notations. $[n] := \{1, \ldots, n\}$; $\mathbb{N}_0 := \{0, 1, 2, \ldots\}$; $\|\cdot\|$ is the standard Euclidean norm; $\langle \cdot, \cdot \rangle$ is the standard dot product; $g = \mathcal{O}(f)$: exist $C > 0$ such that $g(z) \leq C \times f(z)$ for all $z \in \mathcal{Z}$; $g = \Omega(f)$: exist $C > 0$ such that $g(z) \geq C \times f(z)$ for all $z \in \mathcal{Z}$; $g = \Theta(f)$: $g = \mathcal{O}(f)$ and $g = \Omega(f)$; $g = \widetilde{\Theta}(f)$: the same as $g = \Theta(f)$ but up to logarithmic factors.

1.1 Previous work

Oracle complexity. In the classical optimization theory (Nemirovskij & Yudin, 1983), algorithms are compared in terms of *oracle calls*. Assume that the number of workers is one and we work with nonconvex functions and Assumptions 1.1 and 1.4. It is well known (Arjevani et al., 2022; Carmon et al., 2020) that the optimal oracle complexity is $\mathcal{O}\left(L\Delta/\varepsilon + \sigma^2 L\Delta/\varepsilon^2\right)$ to find $\bar{x} \in \mathbb{R}^d$ such that $\mathbb{E}[\|\nabla f(\bar{x})\|^2] \leq \varepsilon$. It is attained by the vanilla SGD method: $x^{k+1} = x^k - \gamma \nabla f(x^k; \xi^k)$, where ξ^k are i.i.d. random samples, $\Delta := f(x^0) - f^*$, $x^0 \in \mathbb{R}^d$ is a starting point, and $\gamma = \Theta\left(\min\{1/L, \varepsilon/L\sigma^2\}\right)$ is a step size. In the convex setting (Assumption 1.3), the optimal oracle complexity is $\Theta\left(\sqrt{L}R/\sqrt{\varepsilon} + \sigma^2 R^2/\varepsilon^2\right)$ (Lan, 2020; Nemirovskij & Yudin, 1983) to find $\bar{x} \in \mathbb{R}^d$ such that $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$, where $R := \left\|x^0 - x_*\right\|$. In the μ–strongly convex setting, the optimal complexity $\widetilde{\Theta}\left(\sqrt{L}/\sqrt{\mu} + \sigma^2/\mu^2\varepsilon\right)$ is to find $\bar{x} \in \mathbb{R}^d$ such that $\mathbb{E}[\|\bar{x} - x_*\|^2] \leq \varepsilon$ (up to logarithmic factors).

Oracle complexity with many workers. Many works discovered oracle complexities with multiple workers. Arjevani & Shamir (2015); Scaman et al. (2017) analyze the heterogeneous convex setting and provide lower bounds when the workers are synchronized. Lu & De Sa (2021) consider the similar setup but in the nonconvex setting. Arjevani et al. (2020) analyze settings where methods receive delayed stochastic gradients. Woodworth et al. (2018) provide lower bounds for parallel setups with intermittent communications and delayed updates. The primary limitation of these results is the assumption that all workers have consistent computational performance, without accounting for individual delays, random lags, or variations in performance over time.

Time complexity. To address the problem of analyzing methods with workers having different computation capabilities and performances, Mishchenko et al. (2022) proposed to consider the *fixed computation model*. In this model, it is assumed that

worker i requires at most τ_i seconds to calculate one stochastic gradient.

Without loss of generality, we assume that the times are sorted: $\tau_1 \leq \cdots \leq \tau_n$. One of the most popular methods is Asynchronous SGD (Lian et al., 2015; Zhang et al., 2015; Feyzmahdavian et al., 2016; Sra et al., 2016; Dutta et al., 2018; Stich & Karimireddy, 2020; Wu et al., 2022; Islamov et al., 2024; Maranjyan et al., 2025; Maranjyan & Richtárik, 2026). In the *homogeneous setting*, Mishchenko et al. (2022); Koloskova et al. (2022); Cohen et al. (2021) showed that Asynchronous SGD and Picky SGD can provably improve the performance of the synchronized Minibatch SGD method that does the steps $x^{k+1} = x^k - \gamma/n \sum_{i=1}^n \nabla f(x^k; \xi_i^k)$, where γ is a stepsize, ξ_i^k are i.i.d. samples, and $\nabla f(x^k; \xi_i^k)$ are calculated in parallel in n workers. Minibatch SGD requires $\mathcal{O}\left(L\Delta/\varepsilon + \sigma^2 L\Delta/n\varepsilon^2\right)$ iterations (Cotter et al., 2011; Goyal et al., 2017; Gower et al., 2019) in the nonconvex setting.

Algorithm 1 Malenia SGD or Rennala SGD when $w_i^k = {}^1/_{B_i^k}$ or $w_i^k = {}^n/_{\sum_{i=1}^n B_i^k}$, respectively

Input: point x^0, stepsize γ, parameter S, weights $\{w_i^k\}$
for $k = 0, 1, \ldots, K-1$ **do**
  Ask all workers to calculate stochastic gradients at x^k; init $g_i^k = 0$ and $B_i^k = 0\ \forall i \in [n]$
  while $\left(\frac{1}{n}\sum_{i=1}^n (w_i^k)^2 B_i^k\right)^{-1} \leq \frac{S}{n}$ **do**
    Wait for the next worker j
    Update $B_j^k = B_j^k + 1$
    Receive stochastic gradient $\nabla f_j(x^k; \xi^k_{j,B_j^k})$ and update $g_j^k = g_j^k + \nabla f_j(x^k; \xi^k_{j,B_j^k})$
    Ask this worker to calculate a stochastic gradient at x^k
  end while
  $g_w^k := \frac{1}{n}\sum_{i=1}^n w_i^k g_i^k = \frac{1}{n}\sum_{i=1}^n w_i^k \sum_{j=1}^{B_i^k} \nabla f_i(x^k; \xi_{ij}^k)$
  $x^{k+1} = x^k - \gamma g_w^k$
  Stop all the workers' calculations (or ignore the unfinished calculations in the next iterations)
end for

Moreover, Minibatch SGD converges after $\mathcal{O}\left(\max_{i\in[n]} \tau_i \times \left({}^{L\Delta}/_{\varepsilon} + {}^{\sigma^2 L\Delta}/_{n\varepsilon^2}\right)\right)$ seconds because it waits for the slowest worker with $\max_{i\in[n]} \tau_i$ in every iteration. Asynchronous SGD, methods with the step $x^{k+1} = x^k - {}^{\gamma^k}/_n \sum_{i=1}^n \nabla f(x^{k-\delta_k}; \xi_i^{k-\delta_k})$ and δ_k–delayed stochastic gradients, improve this time complexity to $\mathcal{O}\left(\left({}^1/_n \sum_{i=1}^n {}^1/_{\tau_i}\right)^{-1}\left({}^{L\Delta}/_{\varepsilon} + {}^{\sigma^2 L\Delta}/_{n\varepsilon^2}\right)\right)$.

Optimal time complexities in the heterogeneous and homogeneous settings. Surprisingly, the time complexity can be further improved. In the nonconvex setup (under Assumptions 1.1, and 1.4), Tyurin & Richtárik (2023) formalized the notion of time complexities and showed that the *optimal time complexity* is

$$T_{\text{homog}} := \Theta\left(\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^m \frac{1}{\tau_i}\right)^{-1}\left(\frac{L\Delta}{\varepsilon} + \frac{\sigma^2 L\Delta}{m\varepsilon^2}\right)\right]\right) \tag{2}$$

seconds *in the homogeneous setup* to find an ε–stationary point, achieved by the Rennala SGD method, where, without loss of generality, the times are sorted: $\tau_1 \leq \cdots \leq \tau_n$. *In the heterogeneous setup*, the optimal time complexity is

$$T_{\text{heter}} := \Theta\left(\tau_n \frac{L\Delta}{\varepsilon} + \left(\frac{1}{n}\sum_{i=1}^n \tau_i\right)\frac{\sigma^2 L\Delta}{n\varepsilon^2}\right), \tag{3}$$

achieved by the Malenia SGD method.

A unifying perspective on Rennala SGD **and** Malenia SGD**.** Let us look closer to the Rennala SGD and Malenia SGD methods (see Algorithm 1) that achieve the optimal time complexities (2) and (3) in the homogeneous and heterogeneous setting, accordingly. We now recall how they work. In every iteration, Rennala SGD and Malenia SGD ask all workers to calculate stochastic gradients asynchronously at *the same iterate* x^k. Assume that worker i has calculated B_i^k stochastic gradients for all $i \in [n]$ at the iteration k. Then the methods do the steps

$$x^{k+1} = x^k - \gamma g_{\mathsf{R}}^k, \quad g_{\mathsf{R}}^k := \frac{1}{\sum_{i=1}^n B_i^k}\sum_{i=1}^n\sum_{j=1}^{B_i^k} \nabla f_i(x^k; \xi_{ij}^k) \qquad \text{(Rennala SGD)}$$

and

$$x^{k+1} = x^k - \gamma g_{\mathsf{M}}^k, \quad g_{\mathsf{M}}^k := \frac{1}{n}\sum_{i=1}^n \frac{1}{B_i^k}\sum_{j=1}^{B_i^k} \nabla f_i(x^k; \xi_{ij}^k), \qquad \text{(Malenia SGD)}$$

accordingly. Rennala SGD and Malenia SGD ask all workers calculating stochastic gradients until $\frac{1}{n}\sum_{i=1}^n B_i^k > {}^S/_n$ and $\left(\frac{1}{n}\sum_{i=1}^n {}^1/_{B_i^k}\right)^{-1} > {}^S/_n$ correspondingly, where S is a parameter. Hence, both methods asynchronously collect and aggregate stochastic gradients to compute g_{R}^k and g_{M}^k, and then perform a descent step. However, the way the methods aggregate is both different and important. It turns out the variance of the Rennala SGD's update is smaller. Indeed, one can easily show that

$$\mathbb{E}\left[\left\|g_{\mathsf{R}}^k - \mathbb{E}\left[g_{\mathsf{R}}^k\right]\right\|^2\right] \leq \frac{\sigma^2}{n}\left(\frac{1}{n}\sum_{i=1}^n B_i^k\right)^{-1} \text{ and } \mathbb{E}\left[\left\|g_{\mathsf{M}}^k - \mathbb{E}\left[g_{\mathsf{M}}^k\right]\right\|^2\right] \leq \frac{\sigma^2}{n}\left(\frac{n}{\sum_{i=1}^n \frac{1}{B_i^k}}\right)^{-1}.$$

Thus, the variance of Rennala SGD improves with the *arithmetic mean* of B_i^k, while the variance of Malenia SGD improves with the *harmonic mean* of B_i^k, which can be much smaller. Why wouldn't we use Rennala SGD in all scenarios if it is better? Because g_{R}^k is biased if $\{f_i\}$ are non-homogeneous. In general, $\mathbb{E}\left[g_{\mathsf{R}}^k\right] \neq \nabla f(x^k)$, while it is always true that $\mathbb{E}\left[g_{\mathsf{M}}^k\right] = \nabla f(x^k)$. Both methods can be generalized into $x^{k+1} = x^k - \gamma g_w^k$, $g_w^k := \frac{1}{n}\sum_{i=1}^n w_i^k \sum_{j=1}^{B_i^k} \nabla f_i(x^k; \xi_{ij}^k)$, where the weights $\{w_i^k\}$ are free parameters. If we take $w_i^k = {}^n/_{\sum_{i=1}^n B_i^k}$ for all $i \in [n]$, we get Rennala SGD with small variance. If we take $w_i^k = {}^1/_{B_i^k}$, we get Malenia SGD with high variance but with an unbiased estimator. The weights enable interpolation between the methods.

Difference between the two settings. Using the inequality of arithmetic and harmonic means, one can easily show that $T_{\text{homog}} \leq T_{\text{heter}}$ (ignoring constant factors). At the same time, the gap between the complexities can be arbitrarily huge. Indeed, when the performance τ_1 of the fastest worker tends to 0, one can easily show that $T_{\text{homog}} \to 0$ and $T_{\text{heter}} \to \Theta\left({}^{\tau_n L\Delta}/_{\varepsilon} + \left(\frac{1}{n}\sum_{i=2}^n \tau_i\right) {}^{\sigma^2 L\Delta}/_{n\varepsilon^2}\right)$, and T_{heter} improves by at most $\sum_{i=1}^n \tau_i / \sum_{i=2}^n \tau_i \leq 2$. While the improvement in the homogeneous setup is ∞. Consider another example when the performance τ_n of the slowest worker (straggler) tends to ∞. Then $T_{\text{heter}} \to \infty$ and $T_{\text{homog}} \to \Theta(\min_{m\in[n-1]}[({}^1/_m \sum_{i=1}^m {}^1/_{\tau_i})^{-1}\left({}^{L\Delta}/_{\varepsilon} + {}^{\sigma^2 L\Delta}/_{m\varepsilon^2}\right)])$, so the complexity T_{homog} is robust to stragglers unlike T_{heter}.

Arbitrarily computation dynamics. The previous discussion explain that a significant gap appears between homogeneous and heterogeneous problems under the fixed computation model. This "arithmetic mean vs harmonic mean gap" was also observed in (Tyurin, 2025), where the author generalizes the fixed computation model to the *universal computation model*, accounting for potential disruptions caused by hardware or network delays, and any variations in computation speeds. For simplicity, in this work, we will continue working with the fixed computation model, but we also show how our final results translate to the universal computation model in Section A.

Convex world. When we want to find a point $\bar{x}$ such that $\mathbb{E}\left[f(\bar{x})\right] - f^* \leq \varepsilon$ in the convex setup, the gap is similar. The optimal time complexity *in the homogeneous setup* is

$$\Theta\left(\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^m \frac{1}{\tau_i}\right)^{-1}\left(\frac{\sqrt{L}R}{\sqrt{\varepsilon}} + \frac{\sigma^2 R^2}{m\varepsilon^2}\right)\right]\right) \tag{4}$$

seconds (Tyurin & Richtárik, 2023). While the optimal time complexity *in the heterogeneous setup* is

$$\Theta\left(\tau_n \frac{\sqrt{L}R}{\sqrt{\varepsilon}} + \left(\frac{1}{n}\sum_{i=1}^n \tau_i\right)\frac{\sigma^2 R^2}{n\varepsilon^2}\right) \tag{5}$$

seconds under Assumptions 1.1, 1.3, and 1.4 (**our new contribution**, Theorem E.4; the final puzzle piece needed to reveal the systematic gap between the two settings). Both complexities are achieved by the accelerated versions of Rennala SGD and Malenia SGD accordingly.

Strongly convex world. Assume additionally that the function f is μ–strongly convex. Using reduction (Woodworth & Srebro, 2016), up to logarithmic factors, we can obtain the optimal time complexity

$$\widetilde{\Theta}\left(\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^m \frac{1}{\tau_i}\right)^{-1}\left(\sqrt{\frac{L}{\mu}} + \frac{\sigma^2}{m\varepsilon\mu}\right)\right]\right) \tag{6}$$

in the homogeneous setting and the optimal time complexity

$$\widetilde{\Theta}\left(\tau_n\sqrt{\frac{L}{\mu}} + \left(\frac{1}{n}\sum_{i=1}^n \tau_i\right)\frac{\sigma^2}{n\varepsilon\mu}\right) \tag{7}$$

in the heterogeneous setting when we want to find a point $\bar{x}$ such that $\mathbb{E}\left[f(\bar{x})\right] - f^* \leq \varepsilon$. Here we also observe a large gap between the settings. Note that the complexities (3), (5), and (7) can only be improved under additional assumptions because they are optimal.

Main question: Having the systematic gap between the homogeneous and heterogeneous setups, the goal of this work is to identify theoretical assumptions that are as weak as possible to improve the results of asynchronous methods in heterogeneous scenarios. Under which assumptions can we improve the dependence on the arithmetic mean of $\{\tau_i\}$ (see (3), (5), and (7)) to the dependence on the harmonic mean of $\{\tau_i\}$ (see (2), (4), and (6))? Right now, the only possible way is to assume that the functions $\{f_i\}$ are equal—an assumption we clearly want to avoid in the heterogeneous setting. Is there any chance to relax this assumption?

1.2 Contributions

To the best of our knowledge, this is the first work to address the main question in any setting; our analysis considers the convex setting under standard Assumptions 1.1, 1.2, 1.3, 1.4, and 1.5.

Analysis of first- and second-order similarity. First, we consider the celebrated *first- and second-order similarity* and, surprisingly, prove that even under these assumptions—no matter how close the functions $\{f_i\}$ are—*any randomized algorithm* cannot converge before $\Omega\left(\left(\frac{1}{n}\sum_{i=1}^{n}\tau_i\right)\frac{\sigma^2}{n\mu^2\varepsilon}\right)$ seconds for small ε (Theorem 2.3). Thus, it is infeasible to break the dependence on the arithmetic mean of $\{\tau_i\}$ under these assumptions.

Investigation of the interpolation assumption. Inspired by Theorem 2.3, which provides a construction with local functions having different minimizers, we decided to go in another direction and consider the *interpolation* assumption. Thus, we introduce two additional assumptions, strong interpolation and the local Polyak-Łojasiewicz condition, and prove that it is impossible to drop either of these assumptions for improvement (Theorems 3.6 and 3.7).

Bridging the gap. By identifying this minimal set of assumptions, we derive a new time complexity result that matches the dependence on worker computation times in the best-known bound in the *homogeneous* setting (Theorem 3.8), but without requiring the functions f_i to be identical. Our theoretical results are validated numerically in Section H.

To bridge the gap in Section 3.2, we need to introduce Assumptions 3.3 and 3.4. However, our primary goal was to illustrate and prove that these assumptions are indeed necessary. Merely stating the assumptions might not be convincing; this is why the central part of our paper investigates different assumptions and shows that most of them do not allow us to bridge the gap. While previous work noted the existence of the gap, our contribution goes further by systematically investigating which assumptions are sufficient and which are insufficient to eliminate it.

2 First-Order and Second-Order Similarity Don't Help

The main problem with the *arithmetic mean dependence* in the heterogeneous setting is that this setting considers a worst-case scenario with arbitrarily heterogeneous functions. Due to the fact that Malenia SGD is optimal, we have to introduce *assumptions* to obtain faster convergence. One of the most popular assumptions in the literature is *first-order and second-order similarity of the functions* (Arjevani & Shamir, 2015; Szlendak et al., 2021; Mishchenko et al., 2022):

Assumption 2.1 (First-Order Similarity)**.** The functions f_i satisfy $\max_{i,j\in[n]}\|\nabla f_i(x)-\nabla f_j(x)\|^2 \leq \delta_1$ for all $x\in\mathbb{R}^d$ for some $\delta_1\geq 0$. It implies $\frac{1}{n}\sum_{i=1}^{n}\|\nabla f_i(x)-\nabla f(x)\|^2\leq\delta_1$ for all $x\in\mathbb{R}^d$.

Assumption 2.2 (Second-Order Similarity)**.** The functions f_i satisfy $\max_{i,j\in[n]}\left\|\nabla^2 f_i(x)-\nabla^2 f_j(x)\right\|^2 \leq \delta_2$ for all $x\in\mathbb{R}^d$ for some $\delta_2\geq 0$. It implies $\frac{1}{n}\sum_{i=1}^{n}\left\|\nabla^2 f_i(x)-\nabla^2 f(x)\right\|^2\leq\delta_2$ for all $x\in\mathbb{R}^d$.

One might expect that when both δ_1 and δ_2 are small, it would be possible to exploit the similarity and design a method with smaller variance and better dependence on $\{\tau_i\}$. Surprisingly, this is not the case: for any $\delta_1>0$ and $\delta_2\geq 0$, one can construct a problem for which the convergence speed of Malenia SGD (Theorem F.2) cannot be improved, up to logarithmic factors, in small-ε regimes:

Theorem 2.3 (Lower Bound). *Consider stochastic gradients $\nabla f_i(x;\xi_i) = \nabla f_i(x) + \xi_i e_i$ with $\xi_i \sim \mathcal{N}(0,\sigma^2)$ for all $i \in [n]$, and $x \in \mathbb{R}^n$. Consider any randomized algorithm that has access only to the stochastic gradients (randomized method), which starts at $x^0 = 0$, under the fixed computation model and any $R, \mu, \beta, \sigma, \varepsilon > 0$ such that $0 < \varepsilon \leq \frac{c\beta^2}{\mu^2 n^2}$ and $R^2 \geq \frac{\beta^2}{\mu^2 n}$, where $c > 0$ is a universal constant. For any time budget $t \leq c_0 \left(\frac{1}{n}\sum_{i=1}^{n} \tau_i\right) \frac{\sigma^2}{n\mu^2\varepsilon}$, where c_0 is a universal constant, there exist $f_i(x) : \mathbb{R}^n \to \mathbb{R}$ such that $f_i(x) = \frac{\mu}{2}\|x\|^2 - \beta\varphi_i\langle x, e_i\rangle$ and $\varphi_i \in [-1,1]$. Assumptions 1.1, 1.2, 1.3, 1.4, and 1.5 hold. Moreover, Assumption 2.1 (the first-order similarity) is satisfied with $\delta_1 = 2\beta^2$, and Assumption 2.2 (the second-order similarity) is satisfied with $\delta_2 = 0$, $\left\|x^0 - x^*\right\|^2 \leq R^2$, and the method cannot produce a point $\bar{x}$ such that $\mathbb{E}\left[\|\bar{x} - x^*\|^2\right] \leq \varepsilon$ within t seconds, where x^* is the minimizer of f.*

Hence, for any small $\delta_1 > 0$ and $\delta_2 \geq 0$, the convergence speed cannot be improved over that of Malenia SGD up to logarithmic factors in small-ε regimes (compare to Theorem F.2). Due to the construction in Theorem 2.3, we can choose any $\beta > 0$, and hence any $\delta_1 > 0$. No matter how close the functions are to each other, the lower bound does not allow us to break the pessimistic time complexity. In view of this, additional assumptions about the first- and second-order similarity will not help to improve the time complexity of Malenia SGD.

Remark 2.4. For the construction in Theorem 2.3, we can also show that $\|\nabla f_i(x)\|^2 \leq 2\|\nabla f(x)\|^2 + 2\beta^2 d$ for all $i \in [n]$, which corresponds to the *ρ–strong growth* condition when $\beta = 0$ and $\rho = 2$ (Schmidt & Roux, 2013). Since Theorem 2.3 holds for all $\beta > 0$, we have proved the result for a "slightly" broader class of problems and have "almost" established that, even under the *strong growth* condition, the convergence speed of Malenia SGD cannot be improved up to logarithmic factors. Whether a similar result holds for the class of problems satisfying $\max_{i\in[n]}\|\nabla f_i(x)\|^2 \leq 2\|\nabla f(x)\|^2$ for all $x \in \mathbb{R}^d$ remains an important open research question.

Takeaway 1: Even with first-order and second-order similarity, for any randomized algorithm, there is still no hope of improving upon the convergence speed of Malenia SGD, up to logarithmic factors, in small-ε regimes.

3 Understanding the Gap via Interpolation Assumptions

Looking at Takeaway 1, we see that a different similarity assumption is required to close the gap between the heterogeneous and homogeneous results. Recall Theorem 2.3. The local minima of the functions f_i are not the same. This motivates us to explore an alternative assumption known as the *interpolation* assumption (Vaswani et al., 2019). This assumption provides another way to capture the similarity among the functions f_i by requiring that they share the same set of minimizers as the function f.

Assumption 3.1 (Weak Interpolation). If x^* is a minimizer of f, that is, $\nabla f(x^*) = 0$, then x^* is also a minimizer of each f_i for all $i \in [n]$.

Interpolation is a property of the solutions of f_i, whereas the heterogeneity assumptions, Assumptions 2.1 and 2.2, concern the gradients and Hessians. These are different characteristics of f_i (see Remark 3.5). Assumption 3.1 is considered practical in modern optimization literature, as there is evidence that it holds for large deep learning models (Zou & Gu, 2019; Zhang et al., 2021). However, as we show next, this assumption alone is not sufficient to achieve improved time complexity, leading to yet another pessimistic result:

Theorem 3.2 (Lower Bound). *Consider any randomized method under the fixed computation model and assume that $n \geq 2$. Let us fix any $\varepsilon, L_{\max}, R, \mu, \sigma^2 > 0$ such that $\mu < L_{\max}/(2n)$, $\varepsilon < 0.01$, and $R > 10$. For any time budget $t \leq c_0 \left(\frac{1}{n}\sum_{i=1}^{n} \tau_i\right) \frac{\sigma^2}{\varepsilon n\mu^2}$, where c_0 is a universal constant, there exist functions $\{f_i\}$ and stochastic gradients $\{\nabla f_i(\cdot;\cdot)\}$ such that $\{f_i\}$ satisfy Assumptions 1.2, 1.3, and 3.1, f satisfies Assumptions 1.1 and 1.5 with $L = L_{\max}$, $\{\nabla f_i(\cdot;\cdot)\}$ satisfy Assumption 1.4 such that the method cannot find ε–solution in terms of distances to the solution set after t seconds, when the method starts at a point in a distance less or equal to R to the closest solution.*

Table 1: The summary of our results and the time complexities (up to logarithmic factors) to get a point $\bar{x}$ such that $\mathbb{E}[\|\bar{x} - x_*\|^2] \leq \varepsilon$ under the *fixed computation model* (worker i requires at most τ_i seconds to calculate one stochastic gradient; $\tau_1 \leq \cdots \leq \tau_n$) and Assumptions 1.1, 1.2, 1.3, 1.4, and 1.5, where $\bar{x}_*$ is the closest solution to $\bar{x}$. The table compares methods in the fully heterogeneous setting and lists the extra assumptions the methods require to work.

Method	**Time Complexity Guarantees**	Additional Assumptions
Minibatch SGD	$\tau_n \left(\frac{L}{\mu} + \frac{\sigma^2}{n\varepsilon\mu^2}\right)$	—
Asynchronous SGD (Mishchenko et al., 2022)	$\left(\frac{1}{n}\sum_{i=1}^{n}\frac{1}{\tau_i}\right)^{-1}\left(\frac{L}{\mu} + \frac{\sigma^2}{n\varepsilon\mu^2}\right)$	$\{f_i\}$ are equal μ–strong convexity
Malenia SGD (Tyurin & Richtárik, 2023) (Theorem F.2)	$\tau_n\frac{L}{\mu} + \left(\frac{1}{n}\sum_{i=1}^{n}\tau_i\right)\frac{\sigma^2}{n\varepsilon\mu^2}$	—
Rennala SGD (Tyurin & Richtárik, 2023) (Theorem F.1)	$\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^{m}\frac{1}{\tau_i}\right)^{-1}\left(\frac{L}{\mu} + \frac{\sigma^2}{m\varepsilon\mu^2}\right)\right]$	$\{f_i\}$ are equal
Lower Bounds (new results)		
Under the first-order and second-order similarity, the following results state that it is infeasible to improve Malenia SGD in small-ε regimes:		
Any randomized method (Theorem 2.3)	$\geq \left(\frac{1}{n}\sum_{i=1}^{n}\tau_i\right)\frac{\sigma^2}{n\varepsilon\mu^2}$	Assumptions 2.1 and 2.2 (first-order and second-order similarity don't help)
Under weak interpolation, the following result states that it is infeasible to improve Malenia SGD in small-ε regimes:		
Any randomized method (Theorem 3.2)	$\geq \left(\frac{1}{n}\sum_{i=1}^{n}\tau_i\right)\frac{\sigma^2}{n\varepsilon\mu^2}$	Assumption 3.1
The following results state that any randomized method can not improve Malenia SGD for small ε if we discard Assumption 3.3 or 3.4:		
Any randomized method (Theorem 3.6)	$\geq \left(\frac{1}{n}\sum_{i=1}^{n}\tau_i\right)\frac{\sigma^2}{n\varepsilon\mu^2}$	Assumptions 3.1 and 3.4 (weak interpolation is not enough)
Any randomized method (Theorem 3.7)	$\geq \left(\frac{1}{n}\sum_{i=1}^{n}\tau_i\right)\frac{\sigma^2}{n\varepsilon\mu^2}$	Assumption 3.3
Upper Bound (new result)		
The following results state that under Assumption 3.3 and 3.4 it is possible to improve Malenia SGD:		
Rennala SGD (Theorem 3.8)	$\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^{m}\frac{1}{\tau_i}\right)^{-1}\left(\frac{L_{\max}}{\mu} + \frac{\sigma^2}{m\varepsilon\mu^2}\right)\right]$	Assumptions 3.3 and 3.4 (weaker than the equality of functions $\{f_i\}$)

Thus, even under Assumption 3.1, we can not improve the arithmetic mean dependence on $\{\tau_i\}$.

Takeaway 2: Using the weak interpolation assumption, which captures the similarity of the functions in a different way compared to first-order and second-order similarity, it is still infeasible to improve the pessimistic dependence on $\{\tau_i\}$ achieved by Malenia SGD using any randomized algorithm.

3.1 Strong Interpolation and Local PŁ Condition are Both Required

Once again, we need to go deeper and introduce additional assumptions to break the lower bound from Theorem 3.2. To further investigate the problem, we now turn to two related assumptions.

Assumption 3.3 (Strong Interpolation)**.** For all $i \in [n]$, a point x^* is a minimizer of f, that is, $\nabla f(x^*) = 0$, *if and only if* it is also a minimizer of f_i.

This assumption is clearly stronger than the weak interpolation assumption since it requires all the functions to share the set of minimizers (see Remark 3.9).

Assumption 3.4 (Local Polyak-Łojasiewicz condition)**.** There exists μ such that $\|\nabla f_i(x)\|^2 \geq 2\mu\left(f_i(x) - f_i^*\right)$ for all $x \in \mathbb{R}^d$ and for all $i \in [n]$, where f_i^* is the finite optimal function value of f_i.

This assumption, unlike Assumption 1.5, requires each function to satisfy PŁ condition.

Remark 3.5. The similarity and interpolation assumptions are neither disjoint nor does one imply the other. For example, consider $f_i(x) = \frac{1}{2}x^2 + c_i(1 - \cos x)$, $0 \leq c_i < 1$. These functions have the same unique minimizer $x^\star = 0$ and are strongly convex and smooth, and thus satisfy strong interpolation and the local PŁ condition. At the same time, $|f_i'(x) - f_j'(x)| = |c_i - c_j||\sin x|$, $|f_i''(x) - f_j''(x)| = |c_i - c_j||\cos x|$, so the first- and second-order similarity assumptions also hold. Conversely, the construction in Theorem 2.3 satisfies the similarity assumptions while interpolation fails. In the other direction, the functions $f_i(x) = a_i x^2/2$,$a_i > 0$, have the same minimizer

and satisfy the local PŁ condition, whereas for $a_i \neq a_j$ their gradient difference $|(a_i - a_j)x|$ is unbounded. Thus, interpolation-type and similarity assumptions describe different, overlapping forms of heterogeneity.

It turns out again that if we do not assume *both* Assumption 3.3 and Assumption 3.4, then it is infeasible for any randomized method to get a time complexity faster than in Malenia SGD (Theorem F.2) for ε small enough. This statement is formalized in the following two theorems.

Theorem 3.6 (Lower Bound). *Consider any randomized method under the fixed computation model and assume that $n \geq 2$. Let us fix any $\varepsilon, L_{\max}, R, \mu, \sigma^2 > 0$ such that $\mu < L_{\max}/(2n)$, $\varepsilon < 0.01$, and $R > 10$. For any time budget $t \leq c_0 \left(\frac{1}{n}\sum_{i=1}^{n} \tau_i\right) \frac{\sigma^2}{\varepsilon n \mu^2}$, where c_0 is a universal constant, there exist functions $\{f_i\}$ and stochastic gradients $\{\nabla f_i(\cdot;\cdot)\}$ such that $\{f_i\}$ satisfy Assumptions 1.2, 1.3, 3.1, and 3.4 (**Assumption 3.3 is not imposed and may or may not hold**), f satisfies Assumptions 1.1 and 1.5 with $L = L_{\max}$, $\{\nabla f_i(\cdot;\cdot)\}$ satisfy Assumption 1.4 such that the method cannot find ε–solution in terms of distances to the solution set after t seconds, when the method starts at a point in a distance less or equal to R to the closest solution.*

Theorem 3.7 (Lower Bound). *Consider any randomized method under the fixed computation model and assume that $n \geq 2$. Let us fix any $\varepsilon, L_{\max}, R, \mu, \sigma^2 > 0$ such that $\mu < L_{\max}/(2n)$, $\varepsilon < 0.01$, and $R > 10$. For any time budget $t \leq c_0 \left(\frac{1}{n}\sum_{i=1}^{n} \tau_i\right) \frac{\sigma^2}{\varepsilon n \mu^2}$, where c_0 is a universal constant, there exist functions $\{f_i\}$ and stochastic gradients $\{\nabla f_i(\cdot;\cdot)\}$ such that $\{f_i\}$ satisfy Assumptions 1.2, 1.3, 3.1, and 3.3 (**Assumption 3.4 is not imposed and may or may not hold with parameter μ**), f satisfy Assumptions 1.1 and 1.5 with $L = L_{\max}$, $\{\nabla f_i(\cdot;\cdot)\}$ satisfy Assumption 1.4, such that the method cannot find ε–solution in terms of distances to the solution set after t seconds, when the method starts at a point in a distance less or equal to R to the closest solution.*

Takeaway 3: Even when the weak interpolation assumption is combined with only one of Assumptions 3.3 and 3.4, we still obtain only the arithmetic mean dependence on $\{\tau_i\}$.

Once we drop either Assumption 3.3 or Assumption 3.4, it becomes possible to construct a "bad" function (see the proof of theorems) that provides no room for any randomized method to improve.

3.2 Finally bridging the gap

However, if assume that both Assumption 3.3 and Assumption 3.4 hold, then, finally, we can proof the convergence with harmonic-like dependence on $\{\tau_i\}$:

Theorem 3.8 (Upper Bound). *Let Assumptions 1.2, 1.3, 1.4, 3.3, 3.4 hold[1]. We choose $w_i^k = n/\sum_{i=1}^n B_i^k$ for all $k \geq 0, i \in [n]$ in Algorithm 1 (reduces to Rennala SGD). We take $\gamma = 1/L_{\max}$, $S = 4\sigma^2/\mu L_{\max}\varepsilon$, and run Rennala SGD for $k \geq \Omega\left(\frac{L_{\max}}{\mu}\log\frac{R^2}{\varepsilon}\right)$ iterations, then $\mathbb{E}\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \varepsilon$, where x_*^{k+1} is the closest solution to x^{k+1}. Moreover, under the fixed computation model, the method requires $\mathcal{O}\left(\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^{m}\frac{1}{\tau_i}\right)^{-1}\left(\frac{L_{\max}}{\mu} + \frac{\sigma^2}{m\varepsilon\mu^2}\right)\right]\log\frac{R^2}{\varepsilon}\right)$ seconds.*

Under weaker assumptions, without requiring the equality of the functions $\{f_i\}$, this theorem yields time complexity guarantees with a "harmonic"-like dependence on the times $\{\tau_i\}$ for the Rennala SGD method, improving upon the previous theoretical results in Theorem F.1 and (Tyurin & Richtárik, 2023). Notice that the method in Theorem 3.8 is still biased because $\mathbb{E}\left[\sum_{i=1}^n \sum_{j=1}^{B_i^k} \nabla f_i(x;\xi_{ij}^k)/\sum_{i=1}^n B_i^k\right] \neq \nabla f(x)$ in general. That said, we can successfully prove the theorem under this constraint. One of the primary reasons for this is the right choice of *convergence metric*. Initially, we aimed to analyze the biased gradient estimator in terms of function values and gradient norms, trying to prove that the method returns a point $\bar{x}$ such that $\mathbb{E}[f(\bar{x})] - f^* \leq \varepsilon$ or $\mathbb{E}[\|\nabla f(\bar{x})\|^2] \leq \varepsilon$. However, the more appropriate approach is to show $\mathbb{E}[\|\bar{x} - x_*\|^2] \leq \varepsilon$. Using this convergence metric allows us to analyze the biased gradient estimator. This observation can

[1] It is well-know that Assumption 1.2 implies Assumption 1.1. In Section G, we prove that Assumptions 1.3, 3.3 and Assumption 3.4 with constant μ imply Assumption 1.5 with constant $\mu/4$.

be important on its own. Note that we can get convergence in terms of $\mathbb{E}[f(\bar{x})] - f^* \leq \varepsilon$ using L–smoothness, but the result would be loose. One interesting observation is that we do not observe a regime where any other method or strategy improves upon both Malenia SGD and Rennala SGD.

Takeaway 4: Improving the pessimistic dependence in Malenia SGD is possible with Rennala SGD and the additional assumptions, Assumption 3.3 and Assumption 3.4, in convex optimization.

Remark 3.9. Theorem 3.8 is proved under the strong interpolation assumption. This assumption is essential for our result: even relaxing strong interpolation to weak interpolation is insufficient to improve the pessimistic time complexity achieved by Malenia SGD. At the same time, strong interpolation does not require the local functions f_i to be identical. For example, consider the least-squares problems $f_i(x) = \frac{1}{2}\|\mathbf{A}_i(x - x^\star)\|^2$, where the matrices $\mathbf{A}_i$ can be different but have a common kernel N. Then $\operatorname{argmin} f_i = x^\star + N$ for every worker i, so the strong interpolation assumption holds. Thus, this setting provides a concrete machine-learning example with heterogeneous, non-identical local functions covered by Theorem 3.8.

3.3 Extension to nonconvex and general convex optimization

Although in this paper we focus on the convergence metric $\mathbb{E}[\|x^k - x_*\|^2]$ under Assumptions 1.1, 1.2, 1.3, 1.4, and 1.5, we briefly sketch how the constructions behind Theorems 2.3, 3.2, 3.6, and 3.7 may extend to general convex optimization in terms of function values and to nonconvex optimization in terms of gradient norms. Indeed, in the general convex optimization setting, we typically consider Assumptions 1.1, 1.3, and 1.4. The same constructions can be considered with $f(x) - f^* = \frac{\mu}{2}\|x - x^*\|^2$, where $\mu = L$. For sufficiently small ε, applying the construction in the proof of Theorem 3.7 with squared-distance accuracy $2\varepsilon/\mu$ suggests a lower bound of order $\Omega\big(\big(\frac{1}{n}\sum_{i=1}^n \tau_i\big)\frac{\sigma^2}{\varepsilon n L}\big)$ seconds for finding $\bar{x}$ such that $\mathbb{E}\left[f(\bar{x})\right] - f^* \leq \varepsilon$. Similarly, applying the construction with squared-distance accuracy ε/L^2 suggests a lower bound of order $\Omega\big(\big(\frac{1}{n}\sum_{i=1}^n \tau_i\big)\frac{\sigma^2}{\varepsilon n}\big)$ seconds for finding $\bar{x}$ such that $\mathbb{E}[\|\nabla f(\bar{x})\|^2] \leq \varepsilon$. In nonconvex optimization, we consider Assumptions 1.1 and 1.4, and may use the same (convex) construction to prove the lower bounds. These observations suggest that the importance of Assumptions 3.3 and 3.4 for improving *the arithmetic-mean dependence* may extend to other optimization settings. However, obtaining tight dependence on other parameters such as L and ε would require a different construction and is left for future work.

4 Conclusions

In this work, we investigated various assumptions and setups in an effort to break the pessimistic dependence on $\{\tau_i\}$ achieved by Malenia SGD. We considered the first- and second-order similarity, strong growth, and interpolation assumptions. We proved that under the first- and second-order similarity assumptions, it is infeasible to improve the dependence on the arithmetic mean of $\{\tau_i\}$ with any randomized algorithm. We also showed that under weak interpolation (Assumption 3.1), it is likewise not possible for any randomized algorithm to improve upon the result of Malenia SGD. Subsequently, we presented new theoretical results that provide improved time complexity guarantees in the heterogeneous setting, without assuming that the functions f_i are identical (Theorem 3.8). These results are obtained under the standard assumptions of convex optimization, together with Assumptions 3.3 and 3.4. Importantly, we have not merely introduced these assumptions to close the gap, but have shown that neither Assumption 3.3 nor Assumption 3.4 can be dropped in general within our setting, highlighting the fundamental limits of heterogeneous stochastic optimization. At the same time, we acknowledge that strong interpolation is a restrictive assumption in heterogeneous settings, as it requires the local functions to share the same set of minimizers. Identifying this fundamental difficulty is one of the main goals of our work.

There are many unexplored directions that can build on our initial results and observations. While we focused on the most common assumptions in federated and distributed learning, our findings may inspire the development of new assumptions and settings where it is possible to improve upon Malenia SGD. Moreover, our upper bounds and lower bounds were investigated in terms of $\mathbb{E}[\|x^k - x_*\|^2]$ convergence, and the lower bounds are only tight up to logarithmic factors and in small-ε regimes. It would be interesting to see whether tight lower bounds can be obtained in terms of $\mathbb{E}[\|\nabla f(x^k)\|^2]$ in the non-convex setting, and in terms of $\mathbb{E}[f(x^k)] - f^*$ for convex functions.

AI USE STATEMENT

In this work, we used generative AI tools to check for and correct grammatical mistakes and check mathematical proofs. We have not used generative AI tools for other tasks that require disclosure. We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REFERENCES

Yossi Arjevani and Ohad Shamir. Communication complexity of distributed convex learning and optimization. *Advances in Neural Information Processing Systems*, 28, 2015.

Yossi Arjevani, Ohad Shamir, and Nathan Srebro. A tight convergence analysis for stochastic gradient descent with delayed updates. In *Algorithmic Learning Theory*, pp. 111–132. PMLR, 2020.

Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, pp. 1–50, 2022.

Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points i. *Mathematical Programming*, 184(1):71–120, 2020.

Alon Cohen, Amit Daniely, Yoel Drori, Tomer Koren, and Mariano Schain. Asynchronous stochastic optimization robust to arbitrary delays. *Advances in Neural Information Processing Systems*, 34: 9024–9035, 2021.

Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. *Advances in Neural Information Processing Systems*, 24, 2011.

Sanghamitra Dutta, Gauri Joshi, Soumyadip Ghosh, Parijat Dube, and Priya Nagpurkar. Slow and stale gradients can win the race: Error-runtime trade-offs in distributed SGD. In *International Conference on Artificial Intelligence and Statistics*, pp. 803–812. PMLR, 2018.

Hamid Reza Feyzmahdavian and Mikael Johansson. Asynchronous iterations in optimization: New sequence results and sharper algorithmic guarantees. *Journal of Machine Learning Research*, 24 (158):1–75, 2023.

Hamid Reza Feyzmahdavian, Arda Aytekin, and Mikael Johansson. An asynchronous mini-batch algorithm for regularized stochastic optimization. *IEEE Transactions on Automatic Control*, 61 (12):3740–3754, 2016.

Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.

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.

Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch SGD: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.

Xinmeng Huang, Yiming Chen, Wotao Yin, and Kun Yuan. Lower bounds and nearly optimal algorithms in distributed learning with communication compression. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Rustem Islamov, Mher Safaryan, and Dan Alistarh. AsGrad: A sharp unified analysis of asynchronous-SGD algorithms. In *International Conference on Artificial Intelligence and Statistics*, pp. 649–657. PMLR, 2024.

Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*, pp. 795–811. Springer, 2016.

Anastasia Koloskova, Sebastian U Stich, and Martin Jaggi. Sharper convergence guarantees for asynchronous SGD for distributed and federated learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

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*, 2016.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, Toronto, 2009.

Guanghui Lan. *First-order and stochastic optimization methods for machine learning*. Springer, 2020.

Xiangru Lian, Yijun Huang, Yuncheng Li, and Ji Liu. Asynchronous parallel stochastic gradient for nonconvex optimization. *Advances in Neural Information Processing Systems*, 28, 2015.

Yucheng Lu and Christopher De Sa. Optimal complexity in decentralized training. In *International Conference on Machine Learning*, pp. 7111–7123. PMLR, 2021.

Artavazd Maranjyan and Peter Richtárik. Ringleader asgd: The first asynchronous sgd with optimal time complexity under data heterogeneity. In *International Conference on Learning Representations*, volume 2026, pp. 4738–4768, 2026.

Artavazd Maranjyan, Alexander Tyurin, and Peter Richtárik. Ringmaster ASGD: The first asynchronous SGD with optimal time complexity. In *International Conference on Machine Learning*, 2025.

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, Francis Bach, Mathieu Even, and Blake Woodworth. Asynchronous SGD beats minibatch SGD under arbitrary delays. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Arkadij Semenovič Nemirovskij and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.

Yurii Nesterov. *Lectures on convex optimization*, volume 137. Springer, 2018.

Kevin Scaman, Francis Bach, Sébastien Bubeck, Yin Tat Lee, and Laurent Massoulié. Optimal algorithms for smooth and strongly convex distributed optimization in networks. In *International Conference on Machine Learning*, pp. 3027–3036. PMLR, 2017.

Mark Schmidt and Nicolas Le Roux. Fast convergence of stochastic gradient descent under a strong growth condition. *arXiv preprint arXiv:1308.6370*, 2013.

Suvrit Sra, Adams Wei Yu, Mu Li, and Alex Smola. Adadelay: Delay adaptive distributed stochastic optimization. In *Artificial Intelligence and Statistics*, pp. 957–965. PMLR, 2016.

Sebastian U Stich and Sai Praneeth Karimireddy. The error-feedback framework: SGD with delayed gradients. *Journal of Machine Learning Research*, 21(237):1–36, 2020.

Rafał Szlendak, Alexander Tyurin, and Peter Richtárik. Permutation compressors for provably faster distributed nonconvex optimization. In *International Conference on Learning Representations*, 2021.

Alexander Tyurin. Tight time complexities in parallel stochastic optimization with arbitrary computation dynamics. In *International Conference on Learning Representations (ICLR)*, 2025.

Alexander Tyurin and Peter Richtárik. Optimal time complexities of parallel stochastic optimization methods under a fixed computation model. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Alexander Tyurin, Kaja Gruntkowska, and Peter Richtárik. Freya PAGE: First optimal time complexity for large-scale nonconvex finite-sum optimization with heterogeneous asynchronous computations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Sharan Vaswani, Aaron Mishkin, Issam Laradji, Mark Schmidt, Gauthier Gidel, and Simon Lacoste-Julien. Painless stochastic gradient: Interpolation, line-search, and convergence rates. *Advances in neural information processing systems*, 32, 2019.

Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.

Blake E Woodworth and Nati Srebro. Tight complexity bounds for optimizing composite objectives. *Advances in Neural Information Processing Systems*, 29, 2016.

Blake E Woodworth, Jialei Wang, Adam Smith, Brendan McMahan, and Nati Srebro. Graph oracle models, lower bounds, and gaps for parallel stochastic optimization. *Advances in Neural Information Processing Systems*, 31, 2018.

Xuyang Wu, Sindri Magnusson, Hamid Reza Feyzmahdavian, and Mikael Johansson. Delay-adaptive step-sizes for asynchronous learning. *arXiv preprint arXiv:2202.08550*, 2022.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.

Wei Zhang, Suyog Gupta, Xiangru Lian, and Ji Liu. Staleness-aware async-sgd for distributed deep learning. *arXiv preprint arXiv:1511.05950*, 2015.

Difan Zou and Quanquan Gu. An improved analysis of training over-parameterized deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

CONTENTS

1 Introduction **1**
1.1 Previous work . . . 2
1.2 Contributions . . . 5

2 First-Order and Second-Order Similarity Don't Help **5**

3 Understanding the Gap via Interpolation Assumptions **6**
3.1 Strong interpolation and local PŁ condition are both required . . . 7
3.2 Finally bridging the gap . . . 8
3.3 Extension to nonconvex and general convex optimization . . . 9

4 Conclusions **9**

A Arbitrarily computation dynamics **14**

B Proof of the Main Results **14**

C Proof of Lower Bounds **17**

D Auxiliary Results **20**

E Lower Bound in the Heterogeneous Convex Setting **20**

F Proof of Theorems F.1 and F.2 **23**

G Assumptions 1.3, 3.3 and 3.4 imply Assumption 1.5 **26**

H Experiments **27**
H.1 Without interpolation . . . 27
H.2 With interpolation . . . 27
H.3 ResNet-18 and CIFAR-10 . . . 28

I Experiments Details **29**
I.1 Quadratic optimization task generation procedure . . . 29
I.2 Experiments with ResNet and CIFAR-10 . . . 29

A ARBITRARILY COMPUTATION DYNAMICS

Our new result can be readily extended to *the universal computation model*. To encompass virtually all computation scenarios, assume that each worker i performs computations based on a *computation power* function $v_i : \mathbb{R}_+ \to \mathbb{R}_+$. Then the number of stochastic gradients that worker i can calculate from a time t_0 to a time t_1 is an integral of the computation power v_i followed by the floor operation:

$$\text{“\# of stoch. grad. in } [t_0, t_1]\text{”} = \left\lfloor \int_{t_0}^{t_1} v_i(\tau) d\tau \right\rfloor. \tag{8}$$

For instance, if worker i is inactive for the first t seconds and then active again, it would mean $v_i(\tau) = 0$ for all $\tau \leq t$ and $v_i(\tau) > 0$ for all $\tau > t$. Using the universal computation model, we can prove the theorem:

Theorem A.1. *Consider the assumptions, algorithm, and parameters from Theorem 3.8. Then,* Rennala SGD *converges after at most* $\bar{t}_{\left\lceil c \times \frac{L_{\max}}{\mu} \log \frac{R^2}{\varepsilon} \right\rceil}$ *seconds, where the sequence* $\{\bar{t}_k\}$ *is defined recursively as* $\bar{t}_k :=$

$$\min \left\{ t \geq 0 : \sum_{i=1}^{n} \left\lfloor \int_{\bar{t}_{k-1}}^{t} v_i(\tau) d\tau \right\rfloor \geq \max \left\{ \lceil 2S \rceil, 1 \right\} \right\} \tag{9}$$

for all $k \geq 1$ ($\bar{t}_0 \equiv 0$), *and* c *is a universal constant.*

A similar result was obtained in (Tyurin, 2025). However, Tyurin (2025) requires the equality of the functions $\{f_i\}$.

B PROOF OF THE MAIN RESULTS

Theorem 3.8 (Upper Bound). *Let Assumptions 1.2, 1.3, 1.4, 3.3, 3.4 hold*[2]. *We choose* $w_i^k = {}^{n}/_{\sum_{i=1}^{n} B_i^k}$ *for all* $k \geq 0, i \in [n]$ *in Algorithm 1 (reduces to* Rennala SGD*). We take* $\gamma = {}^{1}/_{L_{\max}}$, $S = {}^{4\sigma^2}/_{\mu L_{\max} \varepsilon}$, *and run* Rennala SGD *for* $k \geq \Omega\left(\frac{L_{\max}}{\mu} \log \frac{R^2}{\varepsilon}\right)$ *iterations, then* $\mathbb{E}\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \varepsilon$, *where* x_*^{k+1} *is the closest solution to* x^{k+1}. *Moreover, under the fixed computation model, the method requires* $\mathcal{O}\left(\min_{m \in [n]} \left[\left(\frac{1}{m} \sum_{i=1}^{m} \frac{1}{\tau_i}\right)^{-1} \left(\frac{L_{\max}}{\mu} + \frac{\sigma^2}{m\varepsilon\mu^2}\right)\right] \log \frac{R^2}{\varepsilon}\right)$ *seconds.*

Proof. Let us define x_*^k as an euclidean projection of the point x^{k+1} on to the solution set of the main problem (1), and take the condition expectation $\mathbb{E}_k[\cdot]$ w.r.t. the randomness from the iteration k only. Then we have

$$\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k}\right\|^2\right]$$

due to the projection's properties. Then

$$\begin{aligned}
&\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \\
&\leq \mathbb{E}_k\left[\left\|x^k - \gamma \frac{1}{n} \sum_{i=1}^{n} w_i^k \sum_{j=1}^{B_i^k} \nabla f_i(x^k; \xi_{ij}^k) - x_*^k\right\|^2\right] \\
&= \left\|x^k - x_*^k\right\|^2 - 2\gamma \mathbb{E}_k\left[\left\langle x^k - x_*^k, \frac{1}{n} \sum_{i=1}^{n} w_i^k \sum_{j=1}^{B_i^k} \nabla f_i(x^k; \xi_{ij}^k) \right\rangle\right] + \gamma^2 \mathbb{E}_k\left[\left\|\frac{1}{n} \sum_{i=1}^{n} w_i^k \sum_{j=1}^{B_i^k} \nabla f_i(x^k; \xi_{ij}^k)\right\|^2\right].
\end{aligned}$$

Using unbiasedness (Assumption 1.4) and the variance decomposition equality, we get

$$\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right]$$

[2] It is well-know that Assumption 1.2 implies Assumption 1.1. In Section G, we prove that Assumptions 1.3, 3.3 and Assumption 3.4 with constant μ imply Assumption 1.5 with constant $\mu/4$.

$$\leq \left\|x^k - x_*^k\right\|^2 - 2\gamma \left\langle x^k - x_*^k, \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \nabla f_i(x^k) \right\rangle$$
$$+ \gamma^2 \left\| \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \nabla f_i(x^k) \right\|^2 + \gamma^2 \mathbb{E}_k \left[\left\| \frac{1}{n}\sum_{i=1}^n w_i^k \sum_{j=1}^{B_i^k} \left(\nabla f_i(x^k;\xi_{ij}^k) - \nabla f_i(x^k)\right) \right\|^2 \right].$$

Consider the last term, due to the independence of stochastic gradients and Assumption 1.4, we ensure that

$$\mathbb{E}_k \left[\left\| \frac{1}{n}\sum_{i=1}^n w_i^k \sum_{j=1}^{B_i^k} \left(\nabla f_i(x^k;\xi_{ij}^k) - \nabla f_i(x^k)\right) \right\|^2 \right]$$
$$= \frac{1}{n^2}\sum_{i=1}^n (w_i^k)^2 \sum_{j=1}^{B_i^k} \mathbb{E}_k \left[\left\| \nabla f_i(x^k;\xi_{ij}^k) - \nabla f_i(x^k) \right\|^2 \right] \leq \frac{1}{n^2}\sum_{i=1}^n (w_i^k)^2 B_i^k \sigma^2.$$

Thus

$$\mathbb{E}_k \left[\left\|x^{k+1} - x_*^{k+1}\right\|^2 \right] \tag{10}$$
$$\leq \left\|x^k - x_*^k\right\|^2 - \frac{2\gamma}{n}\sum_{i=1}^n w_i^k B_i^k \left\langle x^k - x_*^k, \nabla f_i(x^k)\right\rangle + \gamma^2 \left\| \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \nabla f_i(x^k) \right\|^2 + \frac{\gamma^2}{n^2}\sum_{i=1}^n (w_i^k)^2 B_i^k \sigma^2.$$

We now consider the second and the third term. Since $\frac{1}{n}\sum_{i=1}^n w_i^k B_i^k = 1$, using Jensen's inequality, we get

$$-2\gamma \left\langle x^k - x_*^k, \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \nabla f_i(x^k) \right\rangle + \gamma^2 \left\| \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \nabla f_i(x^k) \right\|^2$$
$$\leq -2\gamma \left\langle x^k - x_*^k, \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \nabla f_i(x^k) \right\rangle + \gamma^2 \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \left\|\nabla f_i(x^k)\right\|^2.$$

Due to Assumption 3.3, we get

$$\frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \left\|\nabla f_i(x^k)\right\|^2 = \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \left\|\nabla f_i(x^k) - \nabla f_i(x_*^k)\right\|^2$$
$$\leq \frac{1}{n}\sum_{i=1}^n w_i^k B_i^k L_i \left\langle x^k - x_*^k, \nabla f_i(x^k) - \nabla f_i(x_*^k)\right\rangle$$
$$\leq L_{\max}\frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \left\langle x^k - x_*^k, \nabla f_i(x^k) - \nabla f_i(x_*^k)\right\rangle$$
$$= L_{\max}\frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \left\langle x^k - x_*^k, \nabla f_i(x^k)\right\rangle.$$

In the first inequality, we use Lemma D.1 under Assumption 1.2 and convexity (Assumption 1.3). In the second inequality, we use the bound $L_i \leq L_{\max}$ for all $i \in [n]$. Taking $\gamma \leq 1/L_{\max}$ and substituting the last inequality to (10), we obtain

$$\mathbb{E}_k \left[\left\|x^{k+1} - x_*^{k+1}\right\|^2 \right]$$
$$\leq \left\|x^k - x_*^k\right\|^2 - (2\gamma - L_{\max}\gamma^2)\frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \left\langle x^k - x_*^k, \nabla f_i(x^k)\right\rangle + \frac{\gamma^2}{n^2}\sum_{i=1}^n (w_i^k)^2 B_i^k \sigma^2$$
$$\leq \left\|x^k - x_*^k\right\|^2 - \gamma\frac{1}{n}\sum_{i=1}^n w_i^k B_i^k \left\langle x^k - x_*^k, \nabla f_i(x^k)\right\rangle + \frac{\gamma^2}{n^2}\sum_{i=1}^n (w_i^k)^2 B_i^k \sigma^2.$$

Using the convexity, Assumption 3.4, and Lemma D.2, we get

$$\mathbb{E}_k\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right]$$
$$\leq \left\|x^k-x_*^k\right\|^2-\frac{\gamma\mu}{2}\frac{1}{n}\sum_{i=1}^n w_i^k B_i^k\left\|x^k-x_*^k\right\|^2+\frac{\gamma^2}{n^2}\sum_{i=1}^n (w_i^k)^2 B_i^k\sigma^2.$$

We take $w_i^k = {}^n/_{\sum_{i=1}^n B_i^k}$ in the theorem for all $i\in[n]$. Thus

$$\mathbb{E}_k\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right]\leq\left(1-\frac{\gamma\mu}{2}\right)\left\|x^k-x_*^k\right\|^2+\frac{\gamma^2\sigma^2}{\sum_{i=1}^n B_i^k}.$$

In Algorithm 1, with the chosen weights $\{w_i^k\}$, we wait for the moment when $\sum_{i=1}^n B_i^k > S$. Thus

$$\mathbb{E}_k\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right]\leq\left(1-\frac{\gamma\mu}{2}\right)\left\|x^k-x_*^k\right\|^2+\frac{\gamma^2\sigma^2}{S}.$$

Unrolling the recursion and taking the full expectation, we obtain

$$\mathbb{E}\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right]\leq\left(1-\frac{\gamma\mu}{2}\right)^{k+1}\left\|x^0-x_*^0\right\|^2+\sum_{j=0}^k\left(1-\frac{\gamma\mu}{2}\right)^j\frac{\gamma^2\sigma^2}{S}$$
$$\leq\left(1-\frac{\gamma\mu}{2}\right)^{k+1}\left\|x^0-x_*^0\right\|^2+\frac{2\gamma\sigma^2}{\mu S}.$$

Due the choice of γ, S, and the condition on k, we have $\mathbb{E}\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right]\leq\varepsilon$.

It is sufficient to run the method for

$$\mathcal{O}\left(\frac{L_{\max}}{\mu}\log\frac{R^2}{\varepsilon}\right)$$

iterations. In each iteration, the method has to ensure that $\sum_{i=1}^n B_i^k > S$. A sufficient time for that is

$$2\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^m\frac{1}{\tau_i}\right)^{-1}\left(1+\frac{4\sigma^2}{mL_{\max}\varepsilon\mu}\right)\right].$$

under the fixed computation model (see Theorem 11 in (Tyurin et al., 2024)). □

Theorem A.1. *Consider the assumptions, algorithm, and parameters from Theorem 3.8. Then,* Rennala SGD *converges after at most* $\bar{t}_{\left\lceil c\times\frac{L_{\max}}{\mu}\log\frac{R^2}{\varepsilon}\right\rceil}$ *seconds, where the sequence* $\{\bar{t}_k\}$ *is defined recursively as* $\bar{t}_k :=$

$$\min\left\{t\geq 0:\sum_{i=1}^n\left\lfloor\int_{\bar{t}_{k-1}}^t v_i(\tau)d\tau\right\rfloor\geq\max\left\{\lceil 2S\rceil,1\right\}\right\}\tag{9}$$

for all $k\geq 1$ $(\bar{t}_0\equiv 0)$, *and* c *is a universal constant.*

Proof. From the proof of Theorem 3.8, we know that it is sufficient to run the method for

$$c\times\frac{L_{\max}}{\mu}\log\frac{R^2}{\varepsilon}$$

iterations, where c is a universal constant. The method waits the moment when $\sum_{i=1}^n B_i^k > S$ in each iteration. The workers work in parallel, and for all $i\in[n]$, will calculate

$$\left\lfloor\int_0^t v_i(\tau)d\tau\right\rfloor$$

stochastic gradients after t seconds. In total, all workers will calculate $\sum_{i=1}^{n} \left\lfloor \int_0^t v_i(\tau)d\tau \right\rfloor$ stochastic gradients. Hence, the first iteration will end by

$$\bar{t}_1 := \min \left\{ t \geq 0 \,:\, \sum_{i=1}^{n} \left\lfloor \int_0^t v_i(\tau)d\tau \right\rfloor \geq \max\left\{\lceil 2S \rceil, 1\right\} \right\},$$

seconds. After that, the second iteration starts before time $\bar{t}_1$ and ends no later than time

$$\bar{t}_2 := \min \left\{ t \geq 0 \,:\, \sum_{i=1}^{n} \left\lfloor \int_{\bar{t}_1}^t v_i(\tau)d\tau \right\rfloor \geq \max\left\{\lceil 2S \rceil, 1\right\} \right\},$$

because worker i can calculate at least

$$\left\lfloor \int_{\bar{t}_1}^t v_i(\tau)d\tau \right\rfloor$$

stochastic gradients between the end of the first iteration and a time t. Using the same reasoning, we can recursively define

$$\bar{t}_3, \ldots, \bar{t}_{\left\lceil c \times \frac{L_{\max}}{\mu} \log \frac{R^2}{\varepsilon} \right\rceil}.$$

The algorithm will converge by $\bar{t}_{\left\lceil c \times \frac{L_{\max}}{\mu} \log \frac{R^2}{\varepsilon} \right\rceil}$ seconds due to the discussion at the beginning of the theorem. □

C Proof of Lower Bounds

Theorem 2.3 (Lower Bound). *Consider stochastic gradients $\nabla f_i(x;\xi_i) = \nabla f_i(x) + \xi_i e_i$ with $\xi_i \sim \mathcal{N}(0,\sigma^2)$ for all $i \in [n]$, and $x \in \mathbb{R}^n$. Consider any randomized algorithm that has access only to the stochastic gradients (randomized method), which starts at $x^0 = 0$, under the fixed computation model and any $R, \mu, \beta, \sigma, \varepsilon > 0$ such that $0 < \varepsilon \leq \frac{c\beta^2}{\mu^2 n^2}$ and $R^2 \geq \frac{\beta^2}{\mu^2 n}$, where $c > 0$ is a universal constant. For any time budget $t \leq c_0 \left(\frac{1}{n} \sum_{i=1}^{n} \tau_i \right) \frac{\sigma^2}{n\mu^2\varepsilon}$, where c_0 is a universal constant, there exist $f_i(x) : \mathbb{R}^n \to \mathbb{R}$ such that $f_i(x) = \frac{\mu}{2}\|x\|^2 - \beta\varphi_i \langle x, e_i \rangle$ and $\varphi_i \in [-1,1]$. Assumptions 1.1, 1.2, 1.3, 1.4, and 1.5 hold. Moreover, Assumption 2.1 (the first-order similarity) is satisfied with $\delta_1 = 2\beta^2$, and Assumption 2.2 (the second-order similarity) is satisfied with $\delta_2 = 0$, $\left\|x^0 - x^*\right\|^2 \leq R^2$, and the method cannot produce a point $\bar{x}$ such that $\mathbb{E}\left[\left\|\bar{x} - x^*\right\|^2\right] \leq \varepsilon$ within t seconds, where x^* is the minimizer of f.*

Proof. We assume that $\varphi_i \in [-1,1]$ for all $i \in [n]$ and define them later. The first-order similarity of these functions is

$$\max_{x \in \mathbb{R}^n} \max_{i,j \in [n]} \|\nabla f_i(x) - \nabla f_j(x)\|^2 \leq 2\beta^2.$$

Thus, the parameter β from the construction controls this similarity. Taking β small, we increase similarity between the functions. Notice that the second-order similarity between the functions is zero since $\nabla^2 f_i(x) = \mu \mathbf{I}$ for all $i \in [n]$. The optimal point is x^* such that $x_i^* = \frac{\beta\varphi_i}{\mu n}$ and $\left\|x^0 - x^*\right\|^2 \leq \frac{\beta^2}{\mu^2 n} \leq R^2$.

Let $N_i(t) := \left\lfloor \frac{t}{\tau_i} \right\rfloor$ denote an upper bound on the number of stochastic gradients returned by worker i by time t. For every stochastic gradient returned by worker i, the method gets

$$\nabla f_i(\bar{x};\xi) = \mu\bar{x} - \beta\varphi_i e_i + \xi e_i = \mu\bar{x} - \beta\left(\varphi_i - \frac{\xi}{\beta}\right) e_i \tag{11}$$

where $\bar{x}$ is a query point and $\xi \sim \mathcal{N}(0,\sigma^2)$. Therefore, all the information obtained from worker i about φ_i consists of at most $N_i(t)$ independent observations from $\mathcal{N}(\varphi_i, \frac{\sigma^2}{\beta^2})$.

Let $\bar{x}_t$ be the point returned by the algorithm at time t for the objective parametrized by φ. Then,

$$\|\bar{x}_t - x^*\|^2 = \sum_{i=1}^n \left((\bar{x}_t)_i - \frac{\beta\varphi_i}{\mu n}\right)^2 = \frac{\beta^2}{\mu^2 n^2}\sum_{i=1}^n \left(\frac{\mu n}{\beta}(\bar{x}_t)_i - \varphi_i\right)^2.$$

Thus, we reduced the problem to estimating φ_i using the observations $\mathcal{N}(\varphi_i, \frac{\sigma^2}{\beta^2})$. Using the classical statistical result (e.g., Example 15.4 from (Wainwright, 2019)), for a universal constant $c_2 > 0$,

$$\sup_{\varphi\in[-1,1]^n} \mathbb{E}\left[\|\bar{x}_t - x^*\|^2\right] \geq c_2 \frac{\beta^2}{\mu^2 n^2}\sum_{i=1}^n \min\left\{1, \frac{\sigma^2}{\beta^2 N_i(t)}\right\},$$

where $\bar{x}_t$ depends on φ. Thus,

$$\sup_{\varphi\in[-1,1]^n} \mathbb{E}\left[\|\bar{x}_t - x^*\|^2\right] \geq c_2 \frac{\beta^2}{\mu^2 n^2}\sum_{i=1}^n \min\left\{1, \frac{\sigma^2\tau_i}{\beta^2 t}\right\}.$$

We take

$$T = c_1 \frac{\sigma^2}{\mu^2 n^2 \varepsilon}\sum_{i=1}^n \tau_i,$$

where c_1 is a universal constant. Then,

$$\sup_{\varphi\in[-1,1]^n} \mathbb{E}\left[\|\bar{x}_t - x^*\|^2\right] \geq c_2 \frac{\beta^2}{\mu^2 n^2}\sum_{i=1}^n \min\left\{1, \frac{\sigma^2\tau_i}{\beta^2 T}\right\}$$

due to the bound on t. The condition on ε ensures that $\frac{\sigma^2\tau_i}{\beta^2 T} \leq 1$ for all $i \in [n]$. Therefore,

$$\sup_{\varphi\in[-1,1]^n} \mathbb{E}\left[\|\bar{x}_t - x^*\|^2\right] \geq 2\varepsilon. \tag{12}$$

Finally, $\mathbb{E}\left[\|\bar{x}_t - x^*\|^2\right] > \varepsilon$ after t seconds for some φ due to (12). The adversary can choose this φ. □

Theorem 3.6 (Lower Bound). *Consider any randomized method under the fixed computation model and assume that $n \geq 2$. Let us fix any $\varepsilon, L_{\max}, R, \mu, \sigma^2 > 0$ such that $\mu < L_{\max}/(2n)$, $\varepsilon < 0.01$, and $R > 10$. For any time budget $t \leq c_0\left(\frac{1}{n}\sum_{i=1}^n \tau_i\right)\frac{\sigma^2}{\varepsilon n\mu^2}$, where c_0 is a universal constant, there exist functions $\{f_i\}$ and stochastic gradients $\{\nabla f_i(\cdot;\cdot)\}$ such that $\{f_i\}$ satisfy Assumptions 1.2, 1.3, 3.1, and 3.4 (**Assumption 3.3 is not imposed and may or may not hold**), f satisfies Assumptions 1.1 and 1.5 with $L = L_{\max}$, $\{\nabla f_i(\cdot;\cdot)\}$ satisfy Assumption 1.4 such that the method cannot find ε–solution in terms of distances to the solution set after t seconds, when the method starts at a point in a distance less or equal to R to the closest solution.*

Proof. Without loss of generality, assume that the method starts at 0. Let $S_\tau = \sum_{i=1}^n \tau_i$ and $a_i = R\sqrt{\tau_i/S_\tau}$. For all $i \in [n]$, $\varphi_i \in [-a_i, a_i]$ is defined later, and

$$f_i(x) = \frac{n\mu}{2}\langle x - \varphi, e_i\rangle^2, \qquad \nabla f_i(x;\xi_i) = n\mu\langle x - \varphi, e_i\rangle e_i + \xi_i e_i,$$

where $\xi_i \sim \mathcal{N}(0,\sigma^2)$. The average function is

$$f(x) = \frac{\mu}{2}\|x - \varphi\|^2,$$

and its unique minimizer is $x^* = \varphi$. Moreover, $\|\varphi\| \leq R$.

Each f_i is convex and smooth with parameter $n\mu \leq L_{\max}$, and

$$\|\nabla f_i(x)\|^2 = 2n\mu f_i(x) \geq 2\mu f_i(x).$$

Hence Assumptions 1.3, 1.2, and 3.4 hold. The function f is smooth with parameter $\mu \leq L_{\max}$ and satisfies the global PŁ condition with constant μ. Assumption 3.1 holds because φ minimizes every f_i, while Assumption 3.3 does not hold in general because

$$\arg\min f_i = \{x \in \mathbb{R}^n : x_i = \varphi_i\} \neq \{\varphi\} = \arg\min f.$$

Finally, the stochastic gradients are unbiased and have variance σ^2.

The rest of the proof is the same as the proof of Theorem 2.3. Every stochastic gradient returned by worker i gives an observation from $\mathcal{N}(\varphi_i, \sigma^2/(n^2\mu^2))$. Let $N_i(t) := \lfloor t/\tau_i \rfloor$ be an upper bound on the number of stochastic gradients calculated by worker i by time t. The same Gaussian mean-estimation lower bound gives

$$\begin{aligned}\sup_{\substack{\varphi_i \in [-a_i, a_i] \\ i \in [n]}} \mathbb{E}\left[\left\|\bar{x}_t - \varphi\right\|^2\right] &\geq c_1 \sum_{i=1}^n \min\left\{a_i^2, \frac{\sigma^2}{n^2\mu^2 N_i(t)}\right\} \\ &\geq c_1 \min\left\{R^2, \frac{\sigma^2 S_\tau}{n^2\mu^2 t}\right\} \geq 2\varepsilon\end{aligned}$$

for a universal constant $c_1 > 0$, where $\bar{x}_t$ any possible query point by time t for the objective with φ. The second inequality follows from $N_i(t) \leq t/\tau_i$ and $a_i^2 = R^2\tau_i/S_\tau$, which imply

$$\min\left\{a_i^2, \frac{\sigma^2}{n^2\mu^2 N_i(t)}\right\} \geq \frac{\tau_i}{S_\tau} \min\left\{R^2, \frac{\sigma^2 S_\tau}{n^2\mu^2 t}\right\},$$

where remains to sum over i and use $\sum_{i=1}^n \tau_i/S_\tau = 1$. The third inequality follows from $\varepsilon < 0.01$, $R > 10$, and the condition on t. Finally,

$$\mathbb{E}\left[\left\|\bar{x}_t - \varphi\right\|^2\right] > \varepsilon$$

for some $\varphi_i \in [-a_i, a_i]$, which the adversary can choose. □

Theorem 3.2 (Lower Bound). *Consider any randomized method under the fixed computation model and assume that $n \geq 2$. Let us fix any $\varepsilon, L_{\max}, R, \mu, \sigma^2 > 0$ such that $\mu < L_{\max}/(2n)$, $\varepsilon < 0.01$, and $R > 10$. For any time budget $t \leq c_0 \left(\frac{1}{n}\sum_{i=1}^n \tau_i\right) \frac{\sigma^2}{\varepsilon n \mu^2}$, where c_0 is a universal constant, there exist functions $\{f_i\}$ and stochastic gradients $\{\nabla f_i(\cdot;\cdot)\}$ such that $\{f_i\}$ satisfy Assumptions 1.2, 1.3, and 3.1, f satisfies Assumptions 1.1 and 1.5 with $L = L_{\max}$, $\{\nabla f_i(\cdot;\cdot)\}$ satisfy Assumption 1.4 such that the method cannot find ε–solution in terms of distances to the solution set after t seconds, when the method starts at a point in a distance less or equal to R to the closest solution.*

Proof. The theorem is a simple corollary of Theorem 3.6 because Theorem 3.6 is stated under more strict assumptions on the class of the functions and stochastic gradients. The result of Theorem 3.6 holds even under additional Assumption 3.4. □

Theorem 3.7 (Lower Bound). *Consider any randomized method under the fixed computation model and assume that $n \geq 2$. Let us fix any $\varepsilon, L_{\max}, R, \mu, \sigma^2 > 0$ such that $\mu < L_{\max}/(2n)$, $\varepsilon < 0.01$, and $R > 10$. For any time budget $t \leq c_0 \left(\frac{1}{n}\sum_{i=1}^n \tau_i\right) \frac{\sigma^2}{\varepsilon n \mu^2}$, where c_0 is a universal constant, there exist functions $\{f_i\}$ and stochastic gradients $\{\nabla f_i(\cdot;\cdot)\}$ such that $\{f_i\}$ satisfy Assumptions 1.2, 1.3, 3.1, and 3.3 (**Assumption 3.4 is not imposed and may or may not hold with parameter** μ), f satisfy Assumptions 1.1 and 1.5 with $L = L_{\max}$, $\{\nabla f_i(\cdot;\cdot)\}$ satisfy Assumption 1.4, such that the method cannot find ε–solution in terms of distances to the solution set after t seconds, when the method starts at a point in a distance less or equal to R to the closest solution.*

Proof. Without loss of generalization, assume that the method starts at 0. Let $S_\tau = \sum_{i=1}^n \tau_i$, take

$$a_i = \frac{n\mu\tau_i}{S_\tau}, \qquad f_i(x) = \frac{a_i}{2}(x-\varphi)^2, \qquad \nabla f_i(x;\xi_i) = a_i(x-\varphi) + \xi_i,$$

where $\xi_i \sim \mathcal{N}(0, \sigma^2)$ and $\varphi \in [-R, R]$ is defined later. Since $\frac{1}{n}\sum_{i=1}^n a_i = \mu$,

$$f(x) = \frac{\mu}{2}(x - \varphi)^2.$$

Thus, f satisfies the global PŁ condition with constant μ, and every f_i is convex and a_i–smooth with $a_i \leq n\mu \leq L_{\max}$. Moreover, all the functions have the same unique minimizer φ, so Assumptions 3.1 and 3.3 hold. Notice that $\|\nabla f_i(x)\|^2 = 2a_i(f_i(x) - f_i^*)$. Finally, the stochastic gradients are unbiased and have variance σ^2.

Every stochastic gradient returned by worker i gives an observation from $\mathcal{N}(\varphi, \sigma^2/a_i^2)$. Let $N_i(t) := \lfloor t/\tau_i \rfloor$ be an upper bound on the number of stochastic gradients computed by worker i by time t. Let P_φ^t be the joint distribution of these observations by time t. For Gaussians with the same variance, the KL divergence $D_{\mathrm{KL}}(\mathcal{N}(m_1, v) \,\|\, \mathcal{N}(m_0, v)) = (m_1 - m_0)^2/(2v)$. By the chain rule for KL divergence,

$$D_{\mathrm{KL}}\left(P_\varphi^t \,\middle\|\, P_{-\varphi}^t\right) \leq \sum_{i=1}^n \frac{2\varphi^2 a_i^2}{\sigma^2} \times \frac{t}{\tau_i} = \frac{2t\varphi^2 n^2 \mu^2}{\sigma^2 S_\tau}$$

because $N_i(t) \leq \frac{t}{\tau_i}$. Since $t \leq \frac{1}{100}\left(\frac{1}{n}\sum_{i=1}^n \tau_i\right)\frac{\sigma^2}{\varepsilon n \mu^2}$, we get

$$D_{\mathrm{KL}}\left(P_\varphi^t \,\middle\|\, P_{-\varphi}^t\right) \leq \frac{\varphi^2}{32\varepsilon}.$$

Choosing any

$$\varphi \in \{-4\sqrt{\varepsilon}, 4\sqrt{\varepsilon}\},$$

we ensure that $-\varphi, \varphi \in [-R, R]$ (since $\varepsilon < 0.01$ and $R > 10$) and $D_{\mathrm{KL}}\left(P_\varphi^t \,\middle\|\, P_{-\varphi}^t\right) \leq \frac{1}{2}$. Pinsker's inequality gives $\left\|P_\varphi^t - P_{-\varphi}^t\right\|_{\mathrm{TV}} \leq 1/2$. Hence, the two-point Le Cam bound (Wainwright, 2019) gives

$$\max_{\varphi \in \{-4\sqrt{\varepsilon}, 4\sqrt{\varepsilon}\}} \mathbb{E}\left[|\bar{x}_t - \varphi|^2\right] \geq 4\varepsilon,$$

where $\bar{x}_t$ is the random variable that the random algorithm can produce based on all available information up to time t for the objective with parameter φ. Thus, the adversary can take $\varphi = \arg\max_{\varphi \in \{-4\sqrt{\varepsilon}, 4\sqrt{\varepsilon}\}} \mathbb{E}\left[|\bar{x}_t - \varphi|^2\right]$ to get $\mathbb{E}\left[|\bar{x}_t - \varphi|^2\right] > \varepsilon$. □

D Auxiliary Results

In this section, we present well-known results from optimization.

Lemma D.1 (Nesterov (2018)). *Let $f : \mathbb{R}^d \to \mathbb{R}$ be a function, which L–smooth and convex. Then for all $x, y \in \mathbb{R}^d$ we have:*

$$\|\nabla f(x) - \nabla f(y)\|^2 \leq L \langle \nabla f(x) - \nabla f(y), x - y \rangle . \tag{13}$$

Lemma D.2 (Karimi et al. (2016)). *Let $f : \mathbb{R}^d \to \mathbb{R}$ be a convex function, which satisfies PŁ condition with a parameter μ (Assumption 1.5). Then, for all $x \in \mathbb{R}^d$, we have*

$$\langle \nabla f(x), x - \bar{x}_* \rangle \geq \frac{\mu}{2} \|\bar{x}_* - x\|^2 \tag{14}$$

where $\bar{x}_*$ is the projection of x onto the solution set of $\min_{x \in \mathbb{R}^d} f(x)$.

E Lower Bound in the Heterogeneous Convex Setting

This section complements the results from (Tyurin & Richtárik, 2023), where the authors only prove the optimal time complexities in the *homogeneous nonconvex*, *heterogeneous nonconvex*, and *homogeneous convex* settings. Here, we resolve the last piece, the *heterogeneous convex* setting.

Protocol 2 Time Multiple Oracles Protocol

Input: function(s) $f \in \mathcal{F}$, oracles and distributions $((O_1, ..., O_n), (\mathcal{D}_1, ..., \mathcal{D}_n)) \in \mathcal{O}(f)$, algorithm $A \in \mathcal{A}$
$s_i^0 = 0$ for all $i \in [n]$
for $k = 0, \ldots, \infty$ **do**
  $(t^{k+1}, i^{k+1}, x^k) = A^k(g^1, \ldots, g^k)$, $\triangleright\ t^{k+1} \geq t^k$
  $(s_{i^{k+1}}^{k+1}, g^{k+1}) = O_{i^{k+1}}(t^{k+1}, x^k, s_{i^{k+1}}^k, \xi^{k+1})$, $\xi^{k+1} \sim \mathcal{D}_{i^{k+1}}$ $\triangleright\ s_j^{k+1} = s_j^k \quad \forall j \neq i^{k+1}$
end for

Following Tyurin & Richtárik (2023), we have to formalize and introduce the following protocol and classes.

We investigate the optimization problem (1) when the function f is convex. For the convex case, using Protocol 2, we use the complexity measure

$$\mathfrak{m}_{\text{time}}(\mathcal{A}, \mathcal{F}) := \inf_{A \in \mathcal{A}} \inf \left\{ t \geq 0 \,\middle|\, \sup_{f \in \mathcal{F}} \sup_{(O,\mathcal{D}) \in \mathcal{O}(f)} \left(\mathbb{E}\left[f(x^{k(t)}) \right] - \inf_{x \in Q} f(x) \right) \leq \varepsilon \right\}, \tag{15}$$

where x^k is generated by Protocol 2, $k(t)$ is the largest index such that $t^{k(t)} \leq t$, and Q is a convex set. Let us take any set Q, and consider the following class of convex functions.

Definition E.1 (Function Class $\mathcal{F}_{Q,L}^{\text{conv}}$)**.**
We assume that a function $f : \mathbb{R}^d \to \mathbb{R}$ is convex, differentiable, L-smooth on the set Q, i.e.,

$$\|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\| \quad \forall x, y \in Q.$$

A set of all functions with such properties we define as $\mathcal{F}_{Q,L}^{\text{conv}}$.

Definition E.2 (Algorithm Class $\mathcal{A}_{\text{zr}}$)**.**
An algorithm $A = \{A^k\}_{k=0}^{\infty}$ is a sequence such that

$$A^k : \underbrace{\mathbb{R}^d \times \cdots \times \mathbb{R}^d}_{k \text{ times}} \to \mathbb{R}_{\geq 0} \times \mathbb{R}^d \quad \forall k \geq 1, A^0 \in \mathbb{R}_{\geq 0} \times \mathbb{R}^d,$$

and, for all $k \geq 1$ and $g^1, \ldots, g^k \in \mathbb{R}^d$, $t^{k+1} \geq t^k$, where t^{k+1} and t^k are defined as $(t^{k+1}, \cdot) = A^k(g^1, \ldots, g^k)$ and $(t^k, \cdot) = A^{k-1}(g^1, \ldots, g^{k-1})$. Moreover, $x^k \in Q$ for all $k \geq 0$.

The following oracle helps to formalize the fixed computation model.

$$O_\tau^{\nabla f} : \underbrace{\mathbb{R}_{\geq 0}}_{\text{time}} \times \underbrace{\mathbb{R}^d}_{\text{point}} \times \underbrace{(\mathbb{R}_{\geq 0} \times \mathbb{R}^d \times \{0,1\})}_{\text{input state}} \times \mathbb{S}_\xi \to \underbrace{(\mathbb{R}_{\geq 0} \times \mathbb{R}^d \times \{0,1\})}_{\text{output state}} \times \mathbb{R}^d$$

$$\text{such that } O_\tau^{\nabla f}(t, x, (s_t, s_x, s_q), \xi) = \begin{cases} ((t, x, 1), & 0), & s_q = 0, \\ ((s_t, s_x, 1), & 0), & s_q = 1 \text{ and } t < s_t + \tau, \\ ((0, 0, 0), & \nabla f(s_x; \xi)), & s_q = 1 \text{ and } t \geq s_t + \tau, \end{cases} \tag{16}$$

and $\nabla f(\cdot; \cdot)$ is a stochastic mapping.

Definition E.3 (Oracle Class $\mathcal{O}_{\tau_1,\ldots,\tau_n}^{\text{conv},\sigma^2}$)**.**
Let us consider an oracle class such that, for any $f \in \mathcal{F}_{Q,L}^{\text{conv}}$, it returns oracles $O_i = O_{\tau_i}^{\nabla f_i}$ and distributions $\mathcal{D}_i$ for all $i \in [n]$, where $\nabla f_i(\cdot; \cdot)$ is an unbiased σ^2-variance-bounded mapping on the set Q of the gradient of the local function in worker i. The oracles $O_{\tau_i}^{\nabla f_i}$ are defined in (16). We define such oracle class as $\mathcal{O}_{\tau_1,\ldots,\tau_n}^{\text{conv},\sigma^2}$. Without loss of generality, we assume that $0 < \tau_1 \leq \cdots \leq \tau_n$.

Notice that this oracle class differs from the oracle class for convex functions in (Tyurin & Richtárik, 2023) because we consider the heterogeneous setting where the oracles return unbiased stochastic gradients of the local functions f_i, which can be different. We refer the reader to (Tyurin & Richtárik, 2023) for additional details about the time complexities formalization. We are now ready to state the theorem.

Theorem E.4 (Informal theorem (see the formal Theorem E.5))**.** *Let Assumptions 1.3, 1.1, and 1.4 hold. It is impossible to converge faster than*

$$\Theta\left(\tau_n \sqrt{L}R/\sqrt{\varepsilon} + \left(1/n \sum_{i=1}^{n} \tau_i\right) \sigma^2 R^2/n\varepsilon^2\right)$$

seconds under the fixed computation model.

Theorem E.5. *Let us consider the oracle class $\mathcal{O}^{\text{conv},\sigma^2}_{\tau_1,\ldots,\tau_n}$ for some $\sigma^2 > 0$ and $0 < \tau_1 \leq \cdots \leq \tau_n$. We fix any $R, L, \varepsilon > 0$ such that $\sqrt{L}R > c_1\sqrt{\varepsilon} > 0$. For any*

$$t \leq c \times \left[\tau_n \frac{\sqrt{L}R}{\sqrt{\varepsilon}} + \left(\frac{1}{n}\sum_{i=1}^{n}\tau_i\right)\frac{\sigma^2 R^2}{n\varepsilon^2}\right],$$

in the view Protocol 2, for any algorithm $A \in \mathcal{A}_{\text{zr}}$, there exists a set Q, a function $f \in \mathcal{F}^{\text{conv}}_{Q,L}$ and oracles and distributions $((O_1,\ldots,O_n),(\mathcal{D}_1,\ldots,\mathcal{D}_n)) \in \mathcal{O}^{\text{conv},\sigma^2}_{\tau_1,\ldots,\tau_n}(f)$ such that

$$\mathbb{E}\left[f(x^{k(t)})\right] - \inf_{x\in Q} f(x) > \varepsilon,$$

where $k(t)$ is the largest index such that $t^{k(t)} \leq t$, and R is the euclidean distance between 0 (starting point) and the closest solution $x_* \in Q$. The quantities c_1, and c are universal constants.

Proof. First term. It is easy to prove the dependence on the first term $\tau_n \frac{\sqrt{L}R}{\sqrt{\varepsilon}}$ using the same idea as in (Lu & De Sa, 2021; Tyurin & Richtárik, 2023; Huang et al., 2022). It is sufficient to put a "hard" convex function (Nesterov, 2018; Woodworth et al., 2018) to the slowest worker corresponding with the time $\tau_n = \max_{i\in[n]} \tau_i$. In particular, we can consider the "hard" quadratic function $\bar{f}$ from (Nesterov, 2018)[Section 2.1.2] and take the functions

$$f_i(x) = \begin{cases} n \times \bar{f}(x), & i = n \\ 0, & i < n. \end{cases}$$

for all $x \in \mathbb{R}^d$. The function $f = \frac{1}{n}\sum_{i=1}^{n} f_i = \bar{f}$ belongs to the class $\mathcal{F}^{\text{conv}}_{Q,L}$. We take the stochastic gradients without noise, i.e., $\nabla f_i(x;\xi_i) = \nabla f_i(x)$ deterministically for all $x \in \mathbb{R}^d$, $\xi_i \in \mathbb{S}_{\xi_i}$, and $i \in [n]$. It is clear that the only worker that can solve the problem is worker n, and it takes τ_n seconds to find one gradient by the oracle construction. Thus, the required time complexity is $\Theta\left(\tau_n \frac{\sqrt{L}R}{\sqrt{\varepsilon}}\right)$ since the required oracle complexity is $\Theta\left(\frac{\sqrt{L}R}{\sqrt{\varepsilon}}\right)$ (Nesterov, 2018).

Second term. The proof of the second term is slightly trickier and uses the construction from (Woodworth et al., 2018). Let us fix any algorithm. We use the proof of Lemma 10 from (Woodworth et al., 2018) that has the following result. For any σ^2, $B > 0$ and any algorithm, it is possible to construct a *one dimensional linear* function $g : \mathbb{R} \to \mathbb{R}$ on the domain $\{x \in \mathbb{R} : |x| \leq B\}$, a stochastic gradient mapping $\nabla g : \mathbb{R} \times \mathbb{S}_\xi \to \mathbb{R}$, and a distribution $\mathcal{D}$ such that

$$\mathbb{E}\left[g(x^N)\right] - \min_{|x|\leq B} g(x) \geq \frac{\sigma B}{8\sqrt{N}} \tag{17}$$

after N queries of the oracle, where ∇g is unbiased and σ^2-variance-bounded.

The idea is to put a function g_i to each worker but with different domain sizes. In particular, for all $i \in [n]$, we take the function $f_i : \mathbb{R}^n \to \mathbb{R}$ such that

$$f_i(x) = g_i(x_i), \tag{18}$$

where g_i is the function from Lemma 10 of (Woodworth et al., 2018) applied independently with $B = R_i$ and $N = N_i(\bar{t})$, and x_i is the ith coordinate of a vector x. For all $i \in [n]$, we consider the function f_i on the domain $\{x_i \in \mathbb{R} \mid |x_i| \leq R_i\}$, where $R_i := R \times \frac{\sqrt{\tau_i}}{\sqrt{\sum_{i=1}^{n}\tau_i}}$. One can see that f

is convex, 0–smooth (because g_i is linear). The distance between 0 and the optimal point is less or equal to R because

$$\sum_{i=1}^{n} R_i^2 = \sum_{i=1}^{n} R^2 \frac{\tau_i}{\sum_{i=1}^{n} \tau_i} = R^2$$

and the optimal point for the problem $g_i(x_i) \to \min_{|x_i| \leq R_i}$ is either R_i or $-R_i$. We take

$$Q = \{x \in \mathbb{R}^n \, : \, |x_i| \leq R_i \quad \forall i \in [n]\}.$$

Let us define the time

$$\bar{t} := \frac{\sigma^2 R^2}{256 n \varepsilon^2} \left(\frac{1}{n} \sum_{i=1}^{n} \tau_i \right). \tag{19}$$

By the time $\bar{t}$, worker i can calculate at most

$$N_i(\bar{t}) := \left\lfloor \frac{\bar{t}}{\tau_i} \right\rfloor \tag{20}$$

stochastic gradients. Therefore,

$$\begin{aligned}
\mathbb{E}\left[f(\bar{x})\right] - \min_{x \in Q} f(x) &= \frac{1}{n} \sum_{i=1}^{n} \left(\mathbb{E}\left[g_i(\bar{x}_i)\right] - \min_{|x_i| \leq R_i} g_i(x_i) \right) \overset{(17)}{\geq} \frac{1}{n} \sum_{i=1}^{n} \frac{\sigma R_i}{8\sqrt{N_i(\bar{t})}} \\
&= \frac{1}{n} \sum_{i=1}^{n} \frac{\sigma R \sqrt{\tau_i}}{8\sqrt{N_i(\bar{t})}\sqrt{\sum_{i=1}^{n} \tau_i}} \overset{(19),(20)}{\geq} \sum_{i=1}^{n} \frac{2\varepsilon \tau_i}{\sum_{i=1}^{n} \tau_i} = 2\varepsilon.
\end{aligned}$$

where $\bar{x}$ is any possible output of the algorithm before the time $\bar{t}$.

□

F Proof of Theorems F.1 and F.2

Theorem F.1. *Let Assumptions 1.1, 1.3, 1.4, and 1.5 hold,* ***and the functions*** $\{f_i\}$ ***are equal****. Let us take* $\gamma = 1/L$ *and* $S = 4\sigma^2/\mu L\varepsilon$*, then* Rennala SGD *(Algorithm 1 with* $w_i^k = n/\sum_{i=1}^{n} B_i^k$*) finds* x^{k+1} *such that* $\mathbb{E}\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \varepsilon$ *after*

$$\mathcal{O}\left(\min_{m \in [n]} \left[\left(\frac{1}{m} \sum_{i=1}^{m} \frac{1}{\tau_i} \right)^{-1} \left(\frac{L}{\mu} + \frac{\sigma^2}{m\varepsilon\mu^2} \right) \right] \log \frac{R^2}{\varepsilon} \right) \tag{21}$$

seconds, where x_*^{k+1} *is the closest solution to* x^{k+1}*.*

Proof. Since the functions are equal, Rennala SGD is equivalent to

$$x^{k+1} = x^k - \gamma g_{\mathsf{R}}^k,$$

$$g_{\mathsf{R}}^k := \frac{1}{\sum_{i=1}^{n} B_i^k} \sum_{i=1}^{n} \sum_{j=1}^{B_i^k} \nabla f(x^k; \xi_{ij}^k).$$

Clearly, g_{R}^k is unbiased and

$$\mathbb{E}_k\left[\left\|g_{\mathsf{R}}^k - \nabla f(x^k)\right\|^2\right] = \left(\sum_{i=1}^{n} B_i^k\right)^{-2} \sum_{i=1}^{n} \sum_{j=1}^{B_i^k} \mathbb{E}_k\left[\left\|\nabla f(x^k; \xi_{ij}^k) - \nabla f(x^k)\right\|\right]^2 \leq \sigma^2 \left(\sum_{i=1}^{n} B_i^k\right)^{-1}.$$

Rennala SGD waits for the moment when $\sum_{i=1}^{n} B_i^k > S$ (see Alg. 1 with $w_i^k = n/\sum_{i=1}^{n} B_i^k$). Thus

$$\mathbb{E}_k\left[\left\|g_{\mathsf{R}}^k - \nabla f(x^k)\right\|^2\right] \leq \frac{\sigma^2}{S} \leq \frac{\mu L \varepsilon}{4}$$

We can use Theorem F.3 to get

$$\mathbb{E}\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right] \leq \left(1-\frac{\gamma\mu}{2}\right)^{k+1}\left\|x^0-x_*^k\right\|^2+\frac{\gamma L\varepsilon}{2}.$$

Since $\gamma=\frac{1}{L}$, we obtain

$$\mathbb{E}\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right] \leq \left(1-\frac{\mu}{2L}\right)^{k+1}\left\|x^0-x_*^k\right\|^2+\frac{\varepsilon}{2}.$$

The last inequality ensure that the method finds an ε–solution after

$$\mathcal{O}\left(\frac{L}{\mu}\log\frac{R^2}{\varepsilon}\right)$$

iterations. In each iteration, the method has to ensure that $\sum_{i=1}^n B_i^k > S$. A sufficient time for that is

$$2\min_{m\in[n]}\left[\left(\frac{1}{m}\sum_{i=1}^m\frac{1}{\tau_i}\right)^{-1}(1+S)\right].$$

under the fixed computation model (see Theorem 11 in (Tyurin et al., 2024)). It is left to multiply this time by $\mathcal{O}\left(\frac{L}{\mu}\log\frac{R^2}{\varepsilon}\right)$. □

Theorem F.2. *Let Assumptions 1.1, 1.3, 1.4, and 1.5 hold. Let us take $\gamma=1/L$ and $S=4\sigma^2/\mu L\varepsilon$, then* Malenia SGD *(Algorithm 1 with $w_i^k=1/B_i^k$) finds x^{k+1} such that $\mathbb{E}\left[\left\|x^{k+1}-x_*^{k+1}\right\|^2\right]\leq\varepsilon$ after*

$$\mathcal{O}\left(\left[\tau_n\frac{L}{\mu}+\left(\frac{1}{n}\sum_{i=1}^n\tau_i\right)\frac{\sigma^2}{n\varepsilon\mu^2}\right]\log\frac{R^2}{\varepsilon}\right) \tag{22}$$

seconds, where x_*^{k+1} is the closest solution to x^{k+1}.

Proof. The proof of this theorem almost repeats the proof of Theorem F.1. The variance of Malenia SGD is

$$\mathbb{E}_k\left[\left\|g_{\mathsf{M}}^k-\nabla f(x^k)\right\|^2\right]=\mathbb{E}_k\left[\left\|\frac{1}{n}\sum_{i=1}^n\frac{1}{B_i^k}\sum_{j=1}^{B_i^k}\nabla f_i(x^k;\xi_{ij}^k)-\nabla f(x^k)\right\|^2\right]\leq\frac{\sigma^2}{n}\left(\frac{1}{n}\sum_{j=1}^n\frac{1}{B_i^k}\right).$$

The method waits for the moment when $\left(\frac{1}{n}\sum_{i=1}^n 1/B_i^k\right)^{-1}>\frac{S}{n}$. Therefore

$$\mathbb{E}_k\left[\left\|g_{\mathsf{M}}^k-\nabla f(x^k)\right\|^2\right]\leq\frac{\sigma^2}{S}.$$

Using the same reasoning, the method finds an ε–solution after

$$\mathcal{O}\left(\frac{L}{\mu}\log\frac{R^2}{\varepsilon}\right)$$

iterations. In each iteration, the method has to ensure that $\left(\frac{1}{n}\sum_{i=1}^n 1/B_i^k\right)^{-1}>\frac{S}{n}$. A sufficient time for that is

$$\bar{t}=2\left(\tau_n+\left(\frac{1}{n}\sum_{i=1}^n\tau_i\right)\frac{S}{n}\right)$$

under the fixed computation model because the number of computed stochastic gradients $B_i^k\geq\left\lfloor\frac{\bar{t}}{\tau_i}\right\rfloor$, and

$$\frac{1}{n}\sum_{i=1}^n\frac{1}{B_i^k}\leq\frac{1}{n}\sum_{i=1}^n\frac{1}{\left\lfloor\frac{\bar{t}}{\tau_i}\right\rfloor}\leq\frac{1}{n}\sum_{i=1}^n\frac{2\tau_i}{\bar{t}}<\frac{n}{S},$$

where we use $\lfloor x\rfloor\geq\frac{x}{2}$ for all $x\geq 1$. Multiplying $\bar{t}$ by $\mathcal{O}\left(\frac{L}{\mu}\log\frac{R^2}{\varepsilon}\right)$, we get the result. □

Theorem F.3. *Consider the method*

$$x^{k+1} = x^k - \gamma \nabla f(x^k; \xi^k), \tag{23}$$

where $\mathbb{E}_k\left[\nabla f(x^k;\xi^k)\right] = \nabla f(x^k)$, $\mathbb{E}_k\left[\left\|\nabla f(x^k;\xi^k) - \nabla f(x^k)\right\|^2\right] \leq \sigma^2$, *and* $\sigma^2 > 0$. *Let Assumptions 1.3, 1.5, and 1.1 hold. Let us take* $\gamma = 1/L$ *and* $S = 2\sigma^2/\mu L\varepsilon$, *then the method finds* x^{k+1} *such that*

$$\mathbb{E}\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \left(1 - \frac{\gamma\mu}{2}\right)^{k+1} \left\|x^0 - x_*^k\right\|^2 + \frac{2\gamma\sigma^2}{\mu},$$

where x_*^{k+1} *is the closest solution of* $\min\limits_{x\in\mathbb{R}^d} f(x)$ *to* x^{k+1}.

Proof. Using the properties of the projection and (23), we have

$$\begin{aligned}
\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] &\leq \mathbb{E}_k\left[\left\|x^{k+1} - x_*^k\right\|^2\right] \\
&= \mathbb{E}_k\left[\left\|x^k - \gamma\nabla f(x^k;\xi^k) - x_*^k\right\|^2\right] \\
&= \mathbb{E}_k\left[\left\|x^k - x_*^k\right\|^2\right] - 2\gamma\mathbb{E}_k\left[\left\langle \nabla f(x^k;\xi^k), x^k - x_*^k\right\rangle\right] + \gamma^2\mathbb{E}_k\left[\left\|\nabla f(x^k;\xi^k)\right\|^2\right] \\
&= \mathbb{E}_k\left[\left\|x^k - x_*^k\right\|^2\right] - 2\gamma\left\langle \nabla f(x^k), x^k - x_*^k\right\rangle + \gamma^2\mathbb{E}_k\left[\left\|\nabla f(x^k;\xi^k)\right\|^2\right].
\end{aligned}$$

In the last equality, we use the unbiasedness. Due the variance decomposition equality, we get

$$\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \mathbb{E}_k\left[\left\|x^k - x_*^k\right\|^2\right] - 2\gamma\left\langle \nabla f(x^k), x^k - x_*^k\right\rangle + \gamma^2\left\|\nabla f(x^k)\right\|^2 + \gamma^2\mathbb{E}_k\left[\left\|\nabla f(x^k;\xi^k) - \nabla f(x^k)\right\|^2\right] \tag{24}$$

Since the function f is L–smooth and $\nabla f(x_*^k) = 0$, we obtain

$$\begin{aligned}
&- 2\gamma\left\langle \nabla f(x^k), x^k - x_*^k\right\rangle + \gamma^2\left\|\nabla f(x^k)\right\|^2 \\
&= -2\gamma\left\langle \nabla f(x^k) - \nabla f(x_*), x^k - x_*^k\right\rangle + \gamma^2\left\|\nabla f(x^k) - \nabla f(x_*)\right\|^2 \\
&\leq -2\gamma\left\langle \nabla f(x^k) - \nabla f(x_*), x^k - x_*^k\right\rangle + L\gamma^2\left\langle \nabla f(x^k) - \nabla f(x_*), x^k - x_*^k\right\rangle \\
&= \gamma\left(L\gamma - 2\right)\left\langle \nabla f(x^k) - \nabla f(x_*), x^k - x_*^k\right\rangle.
\end{aligned}$$

Taking $\gamma \leq \frac{1}{L}$ and substituting the inequality to (24), we get

$$\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \mathbb{E}_k\left[\left\|x^k - x_*^k\right\|^2\right] - \gamma\left\langle \nabla f(x^k), x^k - x_*^k\right\rangle + \gamma^2\mathbb{E}_k\left[\left\|\nabla f(x^k;\xi^k) - \nabla f(x^k)\right\|^2\right].$$

The σ^2–variance bounded ensures that

$$\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \mathbb{E}_k\left[\left\|x^k - x_*^k\right\|^2\right] - \gamma\left\langle \nabla f(x^k), x^k - x_*^k\right\rangle + \gamma^2\sigma^2.$$

Due to convexity and Assumption 1.5, we can use Lemma D.2, which yields

$$\begin{aligned}
\mathbb{E}_k\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] &\leq \mathbb{E}_k\left[\left\|x^k - x_*^k\right\|^2\right] - \frac{\gamma\mu}{2}\left\|x^k - x_*^k\right\| + \gamma^2\sigma^2 \\
&= \left(1 - \frac{\gamma\mu}{2}\right)\mathbb{E}_k\left[\left\|x^k - x_*^k\right\|^2\right] + \gamma^2\sigma^2.
\end{aligned}$$

Unrolling the recursion and taking the full expectation, we obtain

$$\mathbb{E}\left[\left\|x^{k+1} - x_*^{k+1}\right\|^2\right] \leq \left(1 - \frac{\gamma\mu}{2}\right)^{k+1}\left\|x^0 - x_*^0\right\|^2 + \frac{2\gamma\sigma^2}{\mu}$$

□

G ASSUMPTIONS 1.3, 3.3 AND 3.4 IMPLY ASSUMPTION 1.5

Theorem G.1. *Let $\{f_i\}$ satisfy Assumption 1.3, 3.3, and Assumption 3.4 with constant μ, then f satisfies Assumption 1.5 with constant $\frac{\mu}{4}$.*

Proof. We fix $x \in \mathbb{R}^d$. Since Assumption 3.3 hold, then the functions share the closest solution x_* to x. Assumption 3.4 ensures that

$$f_i(x) - f_i(x_*) \geq \frac{\mu}{2} \|x - x_*\|^2 .$$

for all $i \in [n]$ (Karimi et al., 2016). Thus

$$f(x) - f(x_*) \geq \frac{\mu}{2} \|x - x_*\|^2 .$$

Due to convexity, we get

$$f(x_*) \geq f(x) + \langle \nabla f(x), x_* - x \rangle .$$

Therefore

$$f(x) - f(x_*) \leq \langle \nabla f(x), x - x_* \rangle \leq \|\nabla f(x)\| \|x - x_*\| \leq \|\nabla f(x)\| \sqrt{\frac{2}{\mu}} \sqrt{f(x) - f(x_*)}$$

and

$$\frac{\mu}{4} (f(x) - f(x_*)) \leq \frac{1}{2} \|\nabla f(x)\|^2 ,$$

which is Assumption 1.5 with constant $\frac{\mu}{4}$. □

H EXPERIMENTS

We conduct a comparison between Rennala SGD and Malenia SGD on both stochastic quadratic optimization tasks and real-world machine learning problems. These are standard quadratic optimization and computer vision problems, the design of which we explain in Section I. We developed a library that simulates the behavior of $n = 100$ workers. Both methods have two hyperparameters: step size γ and parameter S. We do a grid search for both methods and find the best pairs in all setups. We start with synthetic quadratic optimization problems, which are generated *without and with the interpolation regime*. The procedure is described in Section I.1.

H.1 WITHOUT INTERPOLATION

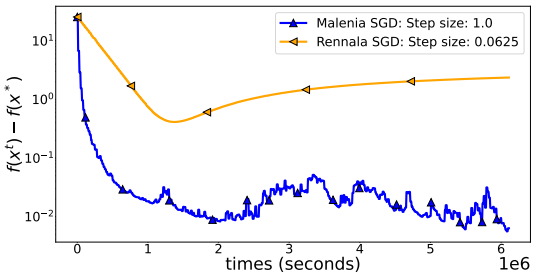

Figure 1: Comparison of the methods on a quadratic optimization problem *without interpolation*. We take the computation time $\tau_i = i^2$ for all $i \in [n]$.

In Figure 1, we present results without interpolation. The plots concur with the theory from Section 3, where we explain that it is essential to have interpolation to break the time complexity of Malenia SGD. Rennala SGD has biased gradient estimators and does not converge to a minimum of the quadratic optimization problem in Figure 1.

H.2 WITH INTERPOLATION

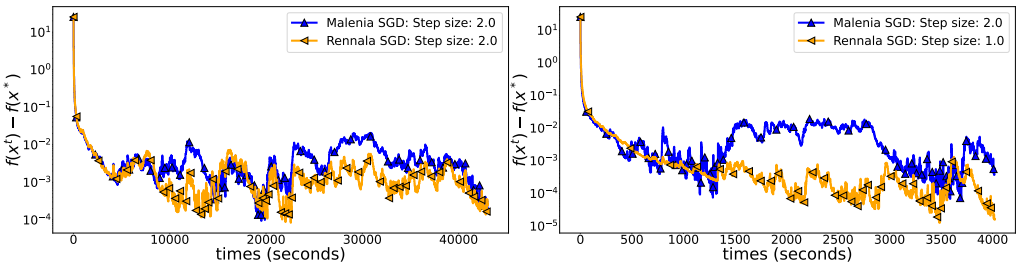

Figure 2: Comparison of the methods on quadratic optimization problems *with interpolation*. Times $\{\tau_i\}$ less diverse: *Left plot:* $\tau_i = \sqrt{i}$ for all $i \in [n]$. *Right plot:* $\tau_1 = 0.01, \tau_2 = 1, \ldots, \tau_n = 1$.

In Figures 2 and 3, we consider the methods in the interpolation regime. As expected, according to Section 3.2, Rennala SGD outperforms Malenia SGD in all experiments. We compare the methods with different $\{\tau_i\}$. In Figures 2, the times $\{\tau_i\}$ are less diverse, so the difference between the methods is less profound. In Figures 3, $\{\tau_i\}$ are more different; thus, we can see that Rennala SGD converges much faster to low function values because it has much less variance in the corresponding gradient estimator.

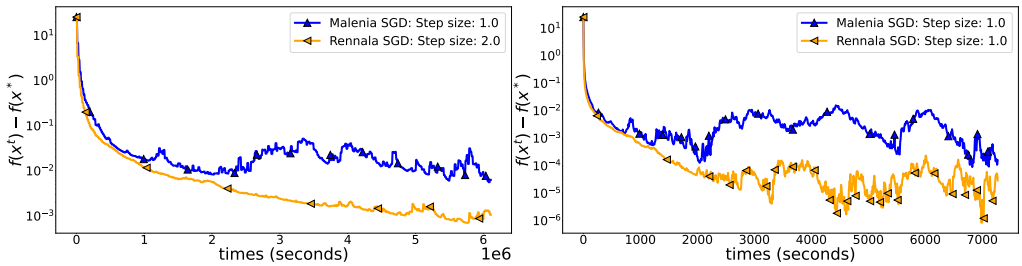

Figure 3: Comparison of the methods on quadratic optimization problems *with interpolation*. Times $\{\tau_i\}$ more diverse: *Left plot:* $\tau_i = i^2$ for all $i \in [n]$. *Right plot:* $\tau_1 = 0.001, \tau_2 = 1, \ldots, \tau_n = 1$.

H.3 ResNet-18 and CIFAR-10

We also verify how Rennala SGD and Malenia SGD work with ResNet-18 and the CIFAR-10 classification problem (Krizhevsky et al., 2009) (License: MIT). Both algorithms take step size $\gamma = 0.25$, sample a batch of size 128, and the smallest S such that all workers calculate at least one batch. The dataset CIFAR-10 is split between the workers, so we consider the heterogeneous setting; all workers access different samples. The results of the experiments are presented in Figure 4. One can see that Rennala SGD converges faster in terms of accuracy, which might be explained by the fact that neural networks work in the interpolation regime. Note that this is an empirical observation in the nonconvex setup, and explaining it from the theoretical point of view is an important future work.

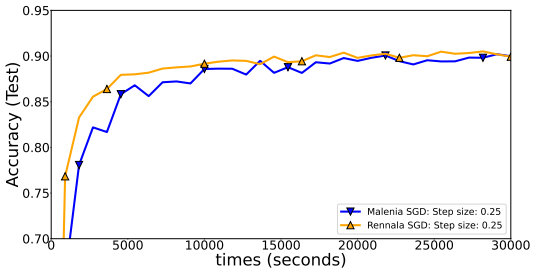

Figure 4: Comparison of the methods on the CIFAR-10 classification problem with ResNet-18. We take the computation time $\tau_i = i^2$.

I EXPERIMENTS DETAILS

The experiments were run in Python 3 using an Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz.

I.1 QUADRATIC OPTIMIZATION TASK GENERATION PROCEDURE

In Section H, we perform experiments using synthetic quadratic optimization problems

$$\min_{x\in\mathbb{R}^d} \frac{1}{n}\sum_{i=1}^{n}\left(\frac{1}{2}x^\top \mathbf{A}_i x - x^\top b_i\right).$$

Below, we present the algorithm, based on (Szlendak et al., 2021), that generates these problems. In all experiments, we take $s = 3$ to ensure that the generated matrices are diverse. We take $n = 100$, $d = 100$, and $\lambda = 0.001$. The stochastic gradients are equal to the true gradients plus standard Gaussian noise added to the coordinates to emulate stochasticity.

With these parameters and procedures, we run the experiments from Section H.1. To conduct the experiments from Section H.2 in the interpolation regime, we take the matrices $\mathbf{A}_1, \cdots, \mathbf{A}_n$, vectors $b_1, \cdots, b_n$ returned by Algorithm 3. Let $\bar{x}_*$ be the solution of the quadratic optimization problem $\frac{1}{n}\sum_{i=1}^{n}\mathbf{A}_i\bar{x}_* = \frac{1}{n}\sum_{i=1}^{n} b_i$. Then, we redefine the vectors $\{b_i\}$ as $b_i = \mathbf{A}_i\bar{x}_*$ to ensure that we are working in the interpolation regime. With this strategy, the matrices are still different, and the functions $\{f_i\}$ are not equal.

Algorithm 3 Generate quadratic optimization tasks

Parameters: number nodes n, dimension d, regularizer λ, and noise scale s.
for $i = 1, \ldots, n$ **do**
  Generate random noises $\eta_i^s = 1 + s\zeta_i^s$ and $\eta_i^b = s\zeta_i^b$, i.i.d. $\zeta_i^s, \zeta_i^b \sim \mathcal{N}(0,1)$
  Take vector $b_i = \frac{\eta_i^s}{4}(-1 + \eta_i^b, 0, \cdots, 0) \in \mathbb{R}^d$
  Take the initial tridiagonal matrix

$$\mathbf{A}_i = \frac{\eta_i^s}{4}\begin{pmatrix} 2 & -1 & & 0 \\ -1 & \ddots & \ddots & \\ & \ddots & \ddots & -1 \\ 0 & & -1 & 2 \end{pmatrix} \in \mathbb{R}^{d\times d}$$

end for
Take the mean of matrices $\mathbf{A} = \frac{1}{n}\sum_{i=1}^{n}\mathbf{A}_i$
Find the minimum eigenvalue $\lambda_{\min}(\mathbf{A})$
for $i = 1, \ldots, n$ **do**
  Update matrix $\mathbf{A}_i = \mathbf{A}_i + (\lambda - \lambda_{\min}(\mathbf{A}))\mathbf{I}$
end for
Take starting point $x^0 = (\sqrt{d}, 0, \cdots, 0)$
Output: matrices $\mathbf{A}_1, \cdots, \mathbf{A}_n$, vectors $b_1, \cdots, b_n$, starting point x^0

I.2 EXPERIMENTS WITH RESNET AND CIFAR-10

In Section H.3, we consider the standard computer vision classification problem with ResNet-18 (He et al., 2016) and CIFAR-10 (Krizhevsky et al., 2009). We conduct the experiments using PyTorch and implement both Rennala SGD and Malenia SGD optimizers. For reproducibility, we use the default ResNet-18 architecture provided in PyTorch and split randomly and evenly the CIFAR-10 dataset across multiple workers to create a heterogeneous data distribution scenario. We use standard preprocessing techniques for CIFAR-10, including normalization and random cropping, and train the network for a fixed number of epochs. The performance metrics include top-1 accuracy. In total, we solve the optimization problem

$$\min_{x\in\mathbb{R}^d} \frac{1}{n}\sum_{i=1}^{n}\left(\frac{1}{m}\sum_{j=1}^{m}\text{loss}(\text{ResNet}(a_{ij}; x), y_{ij})\right),$$

where "loss" is the standard cross-entropy loss, $\{a_{ij}, y_{ij}\}$ are samples from CIFAR-10 splitted between the workers.