

DET-CGD: COMPRESSED GRADIENT DESCENT WITH MATRIX STEPSIZES FOR NON-CONVEX OPTIMIZATION

Hanmin Li Avetik Karagulyan Peter Richtárik

King Abdullah University of Science and Technology

{hanmin.li, avetik.karagulyan, peter.richtarik}@kaust.edu.sa

ABSTRACT

This paper introduces a new method for minimizing matrix-smooth non-convex objectives through the use of novel Compressed Gradient Descent (CGD) algorithms enhanced with a matrix-valued stepsize. The proposed algorithms are theoretically analyzed first in the single-node and subsequently in the distributed settings. Our theoretical results reveal that the matrix stepsize in CGD can capture the objective's structure and lead to faster convergence compared to a scalar stepsize. As a byproduct of our general results, we emphasize the importance of selecting the compression mechanism and the matrix stepsize in a layer-wise manner, taking advantage of model structure. Moreover, we provide theoretical guarantees for free compression, by designing specific layer-wise compressors for the non-convex matrix smooth objectives. Our findings are supported with empirical evidence.

1 INTRODUCTION

The minimization of smooth and non-convex functions is a fundamental problem in various domains of applied mathematics. Most machine learning algorithms rely on solving optimization problems for training and inference, often with structural constraints or non-convex objectives to accurately capture the learning and prediction problems in high-dimensional or non-linear spaces. However, non-convex problems are typically NP-hard to solve, leading to the popular approach of relaxing them to convex problems and using traditional methods. Direct approaches to non-convex optimization have shown success but their convergence and properties are not well understood, making them challenging for large scale optimization. While its convex alternative has been extensively studied and is generally an easier problem, the non-convex setting is of greater practical interest often being the computational bottleneck in many applications.

In this paper, we consider the general minimization problem:

$$\min_{x \in \mathbb{R}^d} f(x), \tag{1}$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a differentiable function. In order for this problem to have a finite solution we will assume throughout the paper that f is bounded from below.

Assumption 1. *There exists $f^{\text{inf}} \in \mathbb{R}$ such that $f(x) \geq f^{\text{inf}}$ for all $x \in \mathbb{R}^d$.*

The stochastic gradient descent (SGD) algorithm (Moulines & Bach, 2011; Bubeck et al., 2015; Gower et al., 2019) is one of the most common algorithms to solve this problem. In its most general form, it can be written as

$$x^{k+1} = x^k - \gamma g(x^k), \tag{2}$$

where $g(x^k)$ is a stochastic estimator of $\nabla f(x^k)$ and $\gamma > 0$ is a positive scalar stepsize. A particular case of interest is the compressed gradient descent (CGD) algorithm (Khirirat et al., 2018), where the estimator g is taken as a compressed alternative of the initial gradient:

$$g(x^k) = \mathcal{C}(\nabla f(x^k)), \tag{3}$$

and the compressor \mathcal{C} is chosen to be a "sparser" estimator that aims to reduce the communication overhead in distributed or federated settings. This is crucial, as highlighted in the seminal paper by

Konečný et al. (2016), which showed that the bottleneck of distributed optimization algorithms is the communication complexity. In order to deal with the limited resources of current devices, there are various compression objectives that are practical to achieve. These include also compressing the model broadcasted from server to clients for local training, and reducing the computational burden of local training. These objectives are mostly complementary, but compressing gradients has the potential for the greatest practical impact due to slower upload speeds of client connections and the benefits of averaging Kairouz et al. (2021). In this paper we will focus on this latter problem.

An important subclass of compressors are the sketches. Sketches are linear operators defined on \mathbb{R}^d , i.e., $\mathcal{C}(y) = \mathbf{S}y$ for every $y \in \mathbb{R}^d$, where \mathbf{S} is a random matrix. A standard example of such a compressor is the Rand- k compressor, which randomly chooses k entries of its argument and scales them with a scalar multiplier to make the estimator unbiased. Instead of communicating all d coordinates of the gradient, one communicates only a subset of size k , thus reducing the number of communicated bits by a factor of d/k . Formally, Rand- k is defined as follows: $\mathbf{S} = \frac{d}{k} \sum_{j=1}^k e_{i_j} e_{i_j}^\top$, where e_{i_j} is the i_j -th standard basis vector in \mathbb{R}^d . We refer the reader to (Safaryan et al., 2022) for an overview on compressions.

Besides the assumption that function f is bounded from below, we also assume that it is \mathbf{L} matrix smooth, as we are trying to take advantage of the entire information contained in the smoothness matrix \mathbf{L} and the stepsize matrix \mathbf{D} .

Assumption 2 (Matrix smoothness). *There exists $\mathbf{L} \in \mathbb{S}_+^d$ such that*

$$f(x) \leq f(y) + \langle \nabla f(y), x - y \rangle + \frac{1}{2} \langle \mathbf{L}(x - y), x - y \rangle \quad (4)$$

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

The assumption of matrix smoothness, which is a generalization of scalar smoothness, has been shown to be a more powerful tool for improving supervised model training. In Safaryan et al. (2021), the authors proposed using smoothness matrices and suggested a novel communication sparsification strategy to reduce communication complexity in distributed optimization for convex objectives. The technique was adapted to three distributed optimization algorithms in the convex setting, resulting in significant communication complexity savings and consistently outperforming the baselines. The results of this study demonstrate the efficacy of the matrix smoothness assumption in improving distributed optimization algorithms.

The case of block-diagonal smoothness matrices is particularly relevant in various applications, such as neural networks (NN). In this setting, each block corresponds to a layer of the network, and we characterize the smoothness with respect to nodes in the i -th layer by a corresponding matrix \mathbf{L}_i . Unlike in the scalar setting, we favor the similarity of certain entries of the argument over the others. This is because the information carried by the layers becomes more complex, while the nodes in the same layers are similar. This phenomenon has been observed visually in various studies, such as those by Yosinski et al. (2015) and Zintgraf et al. (2017).

We study two matrix stepsized CGD-type algorithms and analyze their convergence properties for non-convex matrix-smooth functions. As mentioned earlier, we put special emphasis on the block-diagonal case. We design our sketches and stepsizes in a way that leverages this structure, and we show that in certain cases, we can achieve compression without losing in the overall communication complexity.

1.1 RELATED WORK

Many successful convex optimization techniques have been adapted for use in the non-convex setting. Here is a non-exhaustive list: adaptivity (Dvinskikh et al., 2019; Zhang et al., 2020), variance reduction (J Reddi et al., 2016; Li et al., 2021), and acceleration (Guminov et al., 2019). A paper of particular importance for our work is that of Khaled & Richtárik (2020), which proposes a unified scheme for analyzing stochastic gradient descent in the non-convex regime. A comprehensive overview of non-convex optimization can be found in (Jain et al., 2017; Danilova et al., 2022).

A classical example of a matrix stepsized method is Newton’s method. This method has been popular in the optimization community for a long time (Gragg & Tapia, 1974; Miel, 1980; Yamamoto, 1987).

However, computing the stepsize as the inverse Hessian of the current iteration results in significant computational complexity. Instead, quasi-Newton methods use an easily computable estimator to replace the inverse of the Hessian (Broyden, 1965; Dennis & Moré, 1977; Al-Baali & Khalfan, 2007; Al-Baali et al., 2014). An example is the Newton-Star algorithm (Islamov et al., 2021), which we discuss in Section 2.

Gower & Richtárik (2015) analyzed sketched gradient descent by making the compressors unbiased with a sketch-and-project trick. They provided an analysis of the resulting algorithm for the linear feasibility problem. Later, Hanzely et al. (2018) proposed a variance-reduced version of this method. Sketches are also of independent interest. In particular, Song et al. (2023) described a way of designing the distribution of sketch matrices, while Lee et al. (2019); Qin et al. (2023) used sketches in solving empirical risk minimization problems.

Leveraging the layer-wise structure of neural networks has been widely studied for optimizing the training loss function. For example, (Zheng et al., 2019) propose SGD with different scalar stepsizes for each layer, (Yu et al., 2017; Ginsburg et al., 2019) propose layer-wise normalization for Stochastic Normalized Gradient Descent, and (Dutta et al., 2020; Wang et al., 2022) propose layer-wise compression in the distributed setting.

DCGD, proposed by Khirirat et al. (2018), has since been improved in various ways, such as in (Horvath et al., 2019; Li et al., 2020). There is also a large body of literature on other federated learning algorithms with unbiased compressors (Alistarh et al., 2017; Mishchenko et al., 2019; Gorbunov et al., 2021; Mishchenko et al., 2022; Maranjyan et al., 2022; Horváth et al., 2023).

1.2 CONTRIBUTIONS

Our paper contributes in the following ways:

- We propose two novel matrix stepsize sketch CGD algorithms in Section 2, which, to the best of our knowledge, are the first attempts to analyze a fixed matrix stepsize for non-convex optimization. We present a unified theorem in Section 3 that guarantees stationarity for minimizing matrix-smooth non-convex functions. The results show that taking our algorithms improve on their scalar alternatives. The complexities are summarized in Table 1 for some particular cases.
- We design our algorithms’ sketches and stepsize to take advantage of the layer-wise structure of neural networks, assuming that the smoothness matrix is block-diagonal. In Section 4, we prove that our algorithms achieve better convergence than classical methods.
- Assuming that the server-to-client communication is less expensive Konečný et al. (2016); Kairouz et al. (2021), we propose distributed versions of our algorithms in Section 5, following the standard FL scheme, and prove weighted stationarity guarantees. Our theorem recovers the result for DCGD in the scalar case and improves it in general.
- We validate our theoretical results with experiments. The plots and framework are provided in the Appendix.

1.3 PRELIMINARIES

The usual Euclidean norm on \mathbb{R}^d is defined as $\|\cdot\|$. We use bold capital letters to denote matrices. By \mathbf{I}_d we denote the $d \times d$ identity matrix, and by \mathbf{O}_d we denote the $d \times d$ zero matrix. Let \mathbb{S}_{++}^d (resp. \mathbb{S}_+^d) be the set of $d \times d$ symmetric positive definite (resp. semi-definite) matrices. Given $\mathbf{Q} \in \mathbb{S}_{++}^d$ and $x \in \mathbb{R}^d$, we write $\|x\|_{\mathbf{Q}} := \sqrt{\langle \mathbf{Q}x, x \rangle}$, where $\langle \cdot, \cdot \rangle$ is the standard Euclidean inner product on \mathbb{R}^d . For a matrix $\mathbf{A} \in \mathbb{S}_{++}^d$, we define by $\lambda_{\max}(\mathbf{A})$ (resp. $\lambda_{\min}(\mathbf{A})$) the largest (resp. smallest) eigenvalue of the matrix \mathbf{A} . Let $\mathbf{A}_i \in \mathbb{R}^{d_i \times d_i}$ and $d = d_1 + \dots + d_\ell$. Then the matrix $\mathbf{A} = \text{Diag}(\mathbf{A}_1, \dots, \mathbf{A}_\ell)$ is defined as a block diagonal $d \times d$ matrix where the i -th block is equal to \mathbf{A}_i . We will use $\text{diag}(\mathbf{A}) \in \mathbb{R}^{d \times d}$ to denote the diagonal of any matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$. Given a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, its gradient and its Hessian at point $x \in \mathbb{R}^d$ are respectively denoted as $\nabla f(x)$ and $\nabla^2 f(x)$. A random vector $x \in \mathbb{R}^d$ is an ε -stationary point if $\mathbb{E} \left[\|\nabla f(x)\|^2 \right] \leq \varepsilon^2$, where the expectation is over the randomness of the algorithm.

2 THE ALGORITHMS

Below we define our two main algorithms:

$$x^{k+1} = x^k - \mathbf{D}\mathbf{S}^k\nabla f(x^k), \quad (\text{det-CGD1})$$

and

$$x^{k+1} = x^k - \mathbf{T}^k\mathbf{D}\nabla f(x^k). \quad (\text{det-CGD2})$$

Here, $\mathbf{D} \in \mathbb{S}_{++}^d$ is the fixed stepsize matrix. The sequences of random matrices \mathbf{S}^k and \mathbf{T}^k satisfy the following assumption.

Assumption 3. *We will assume that the random sketches that appear in our algorithms are i.i.d., unbiased, symmetric and positive semi-definite for each algorithm. That is*

$$\begin{aligned} \mathbf{S}^k, \mathbf{T}^k &\in \mathbb{S}_+^d, \quad \mathbf{S}^k \stackrel{iid}{\sim} \mathcal{S} \quad \text{and} \quad \mathbf{T}^k \stackrel{iid}{\sim} \mathcal{T} \\ \mathbb{E}[\mathbf{S}^k] &= \mathbb{E}[\mathbf{T}^k] = \mathbf{I}_d, \quad \text{for every } k \in \mathbb{N}, \end{aligned}$$

where \mathcal{S} and \mathcal{T} are probability distributions over \mathbb{S}_+^d .

A simple instance of [det-CGD1](#) and [det-CGD2](#) is the vanilla GD. Indeed, if $\mathbf{S}^k = \mathbf{T}^k = \mathbf{I}_d$ and $\mathbf{D} = \gamma\mathbf{I}_d$, then $x^{k+1} = x^k - \gamma\nabla f(x^k)$. In general, one may view these algorithms as Newton-type methods. In particular, our setting includes the Newton Star (NS) algorithm by [Islamov et al. \(2021\)](#):

$$x^{k+1} = x^k - (\nabla^2 f(x^{\text{inf}}))^{-1} \nabla f(x^k). \quad (\text{NS})$$

The authors prove that in the convex case it converges to the unique solution x^{inf} locally quadratically, provided certain assumptions are met. However, it is not a practical method as it requires knowledge of the Hessian at the optimal point. This method, nevertheless, hints that constant matrix stepsize can yield fast convergence guarantees. Our results allow us to choose the \mathbf{D} depending on the smoothness matrix \mathbf{L} . The latter can be seen as a uniform upper bound on the Hessian.

The difference between [det-CGD1](#) and [det-CGD2](#) is the update rule. In particular, the order of the sketch and the stepsize is interchanged. When the sketch \mathbf{S} and the stepsize \mathbf{D} are commutative w.r.t. matrix product, the algorithms become equivalent. In general, a simple calculation shows that if we take

$$\mathbf{T}^k = \mathbf{D}\mathbf{S}^k\mathbf{D}^{-1}, \quad (5)$$

then [det-CGD1](#) and [det-CGD2](#) are the same. Defining \mathbf{T}^k according to (5), we recover the unbiasedness condition:

$$\mathbb{E}[\mathbf{T}^k] = \mathbf{D}\mathbb{E}[\mathbf{S}^k]\mathbf{D}^{-1} = \mathbf{I}_d. \quad (6)$$

However, in general $\mathbf{D}\mathbb{E}[\mathbf{S}^k]\mathbf{D}^{-1}$ is not necessarily symmetric, which contradicts to Assumption 3. Thus, [det-CGD1](#) and [det-CGD2](#) are not equivalent for our purposes.

3 MAIN RESULTS

Before we state the main result, we present a stepsize condition for [det-CGD1](#) and [det-CGD2](#), respectively:

$$\mathbb{E}[\mathbf{S}^k\mathbf{D}\mathbf{L}\mathbf{D}\mathbf{S}^k] \preceq \mathbf{D}, \quad (7)$$

and

$$\mathbb{E}[\mathbf{D}\mathbf{T}^k\mathbf{L}\mathbf{T}^k\mathbf{D}] \preceq \mathbf{D}. \quad (8)$$

In the case of vanilla GD (7) and (8) become $\gamma < L^{-1}$, which is the standard condition for convergence. Below is the main convergence theorem for both algorithms in the single-node regime.

Theorem 1. *Suppose that Assumptions 1-3 are satisfied. Then, for each $k \geq 0$*

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left[\|\nabla f(x^k)\|_{\mathbf{D}}^2 \right] \leq \frac{2(f(x^0) - f^{\text{inf}})}{K}, \quad (9)$$

if one of the below conditions is true:

- i) The vectors x^k are the iterates of *det-CGD1* and \mathbf{D} satisfies (7);
- ii) The vectors x^k are the iterates of *det-CGD2* and \mathbf{D} satisfies (8).

It is important to note that Theorem 1 yields the same convergence rate for any $\mathbf{D} \in \mathbb{S}_{++}^d$, despite the fact that the matrix norms on the left-hand side cannot be compared for different weight matrices. To ensure comparability of the right-hand side of (9), it is necessary to normalize the weight matrix \mathbf{D} that is used to measure the gradient norm. We propose using determinant normalization, which involves dividing both sides of (9) by $\det(\mathbf{D})^{1/d}$, yielding the following:

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left[\left\| \nabla f(x^k) \right\|_{\frac{\mathbf{D}}{\det(\mathbf{D})^{1/d}}}^2 \right] \leq \frac{2(f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} K}. \quad (10)$$

This normalization is meaningful because adjusting the weight matrix to $\frac{\mathbf{D}}{\det(\mathbf{D})^{1/d}}$ allows its determinant to be 1, making the norm on the left-hand side comparable to the standard Euclidean norm. It is important to note that the volume of the normalized ellipsoid $\{x \in \mathbb{R}^d : \|x\|_{\mathbf{D}/\det(\mathbf{D})^{1/d}}^2 \leq 1\}$ does not depend on the choice of $\mathbf{D} \in \mathbb{S}_{++}^d$. Therefore, the results of (9) are comparable across different \mathbf{D} in the sense that the right-hand side of (9) measures the volume of the ellipsoid containing the gradient.

3.1 OPTIMAL MATRIX STEPSIZE

In this section, we describe how to choose the optimal stepsize that minimizes the iteration complexity. The problem is easier for *det-CGD2*. We notice that (8) can be explicitly solved. Specifically, it is equivalent to

$$\mathbf{D} \preceq (\mathbb{E} [\mathbf{T}^k \mathbf{L} \mathbf{T}^k])^{-1}. \quad (11)$$

We want to emphasize that the RHS matrix is invertible despite the sketches not being so. Indeed. The map $h : \mathbf{T} \rightarrow \mathbf{T} \mathbf{L} \mathbf{T}$ is convex on \mathbb{S}_+^d . Therefore, Jensen's inequality implies

$$\mathbb{E} [\mathbf{T}^k \mathbf{L} \mathbf{T}^k] \succeq \mathbb{E} [\mathbf{T}^k] \mathbf{L} \mathbb{E} [\mathbf{T}^k] = \mathbf{L} \succ \mathbf{O}_d.$$

This explicit condition on \mathbf{D} can assist in determining the optimal stepsize. Since both \mathbf{D} and $(\mathbb{E} [\mathbf{T}^k \mathbf{L} \mathbf{T}^k])^{-1}$ are positive definite, then the right-hand side of (10) is minimized exactly when

$$\mathbf{D} = (\mathbb{E} [\mathbf{T}^k \mathbf{L} \mathbf{T}^k])^{-1}. \quad (12)$$

Note that the explicit solution of \mathbf{D} needs to be calculated only once, at the beginning of the algorithm. It is then fixed for all iterations. The situation is different for *det-CGD1*. According to (10), the optimal \mathbf{D} is defined as the solution of the following constrained optimization problem:

$$\begin{aligned} & \text{minimize} && \log \det(\mathbf{D}^{-1}) \\ & \text{subject to} && \mathbb{E} [\mathbf{S}^k \mathbf{D} \mathbf{L} \mathbf{D} \mathbf{S}^k] \preceq \mathbf{D} \\ & && \mathbf{D} \in \mathbb{S}_{++}^d. \end{aligned} \quad (13)$$

Proposition 1. *The optimization problem (13) with respect to stepsize matrix $\mathbf{D} \in \mathbb{S}_{++}^d$, is a convex optimization problem with a convex constraint.*

The proof of this proposition can be found in the Appendix. It is based on the reformulation of the constraint to its equivalent quadratic form inequality. Using the trace trick, we can prove that for every vector chosen in the quadratic form, it is convex. Since the intersection of convex sets is convex, we conclude the proof.

One could consider using the CVXPY (Diamond & Boyd, 2016) package to solve (13), provided that it is first transformed into a Disciplined Convex Programming (DCP) form (Grant et al., 2006). Nevertheless, (7) is not recognized as a DCP constraint in the general case. To make CVXPY applicable, additional steps tailored to the problem at hand must be taken.

Table 1: Summary of communication complexities of **det-CGD1** and **det-CGD2** with different sketches and stepsize matrices. The \mathbf{D}_i here for **det-CGD1** is \mathbf{W}_i with the optimal scaling determined using Theorem 2, for **det-CGD2** it is the optimal stepsize matrix defined in (12). The constant $2(f(x^0) - f^{\text{inf}})/\varepsilon^2$ is hidden, ℓ is the number of layers, k_i is the mini-batch size for the i -th layer if we use the rand- k sketch. The notation $\tilde{\mathbf{L}}_{i,k}$ is defined as $\frac{d-k}{d-1} \text{diag}(\mathbf{L}_i) + \frac{k-1}{d-1} \mathbf{L}_i$.

No.	The method	$(\mathbf{S}_i^k, \mathbf{D}_i)$	$l \geq 1, d_i, k_i, \sum_{i=1}^{\ell} k_i = k$, layer structure	$l = 1, k_i = k$, general structure
1.	det-CGD1	$(\mathbf{I}_d, \gamma \mathbf{L}_i^{-1})$	$d \cdot \det(\mathbf{L})^{1/d}$	$d \cdot \det(\mathbf{L})^{1/d}$
2.	det-CGD1	$(\mathbf{I}_d, \gamma \text{diag}^{-1}(\mathbf{L}_i))$	$d \cdot \det(\text{diag}(\mathbf{L}))^{1/d}$	$d \cdot \det(\text{diag}(\mathbf{L}))^{1/d}$
3.	det-CGD1	$(\mathbf{I}_d, \gamma \mathbf{I}_{d_i})$	$d \cdot \left(\prod_{i=1}^{\ell} \lambda_{\max}^{d_i}(\mathbf{L}_i)\right)^{1/d}$	$d \cdot \lambda_{\max}(\mathbf{L})$
4.	det-CGD1	$(\text{rand-1}, \gamma \mathbf{I}_{d_i})$	$\ell \cdot \left(\prod_{i=1}^{\ell} d_i^{d_i} (\max_j (\mathbf{L}_i)_{jj})^{d_i}\right)^{1/d}$	$d \cdot \max_j (\mathbf{L}_{jj})$
5.	det-CGD1	$(\text{rand-1}, \gamma \mathbf{L}_i^{-1})$	$\ell \cdot \left(\frac{\prod_{i=1}^{\ell} d_i^{d_i} \lambda_{\max}^{d_i}(\mathbf{L}_i^{\frac{1}{2}} \text{diag}(\mathbf{L}_i^{-1}) \mathbf{L}_i^{\frac{1}{2}})}{\prod_{i=1}^{\ell} \det(\mathbf{L}_i^{-1})}\right)^{1/d}$	$\frac{d \lambda_{\max}(\mathbf{L}^{\frac{1}{2}} \text{diag}(\mathbf{L}^{-1}) \mathbf{L}^{\frac{1}{2}})}{\det(\mathbf{L}^{-1})^{1/d}}$
6.	det-CGD1	$(\text{rand-1}, \gamma \mathbf{L}_i^{-1/2})$	$\ell \cdot \left(\frac{\prod_{i=1}^{\ell} d_i^{d_i} \lambda_{\max}^{d_i}(\mathbf{L}_i^{1/2})}{\prod_{i=1}^{\ell} \det(\mathbf{L}_i^{-1/2})}\right)^{1/d}$	$d \cdot \lambda_{\max}^{1/2}(\mathbf{L}) \det(\mathbf{L})^{1/(2d)}$
7.	det-CGD1	$(\text{rand-1}, \gamma \text{diag}^{-1}(\mathbf{L}_i))$	$\ell \cdot \left(\frac{\prod_{i=1}^{\ell} d_i^{d_i}}{\prod_{j=1}^d (\mathbf{L}_{jj}^{-1})}\right)^{1/d}$	$d \cdot \det(\text{diag}(\mathbf{L}))^{1/d}$
8.	det-CGD1	$(\text{rand-}k_i, \gamma \text{diag}^{-1}(\mathbf{L}_i))$	$k \cdot \left(\prod_{i=1}^{\ell} \left(\frac{d_i}{k_i}\right)^{d_i} \det(\text{diag}(\mathbf{L}))\right)^{1/d}$	$d \cdot \det(\text{diag}(\mathbf{L}))^{1/d}$
9.	det-CGD2	$(\mathbf{I}_d, \mathbf{L}_i^{-1})$	$d \cdot \det(\mathbf{L})^{1/d}$	$d \cdot \det(\mathbf{L})^{1/d}$
10.	det-CGD2	$(\text{rand-1}, \frac{\text{diag}^{-1}(\mathbf{L}_i)}{d_i})$	$\ell \cdot \left(\prod_{i=1}^{\ell} d_i^{d_i}\right)^{1/d} \det(\text{diag} \mathbf{L})^{1/d}$	$d \cdot \det(\text{diag}(\mathbf{L}))^{1/d}$
11.	det-CGD2	$(\text{rand-}k, \frac{k_i}{d_i} \tilde{\mathbf{L}}_{i,k_i}^{-1})$	$k \cdot \left(\prod_{i=1}^{\ell} \left(\frac{d_i}{k_i}\right)^{d_i}\right) \left(\prod_{i=1}^{\ell} \det(\tilde{\mathbf{L}}_{i,k_i})\right)^{1/d}$	$d \cdot \det(\tilde{\mathbf{L}}_{1,k})$
12.	det-CGD2	$(\text{Bern-}q_i, q_i \mathbf{L}_i^{-1})$	$\left(\sum_{i=1}^{\ell} q_i d_i\right) \cdot \prod_{i=1}^{\ell} \left(\frac{d_i}{q_i}\right)^{d_i} \det(\mathbf{L})^{1/d}$	$d \cdot \det(\mathbf{L})^{1/d}$
13.	GD	$(\mathbf{I}_d, \lambda_{\max}^{-1}(\mathbf{L}) \mathbf{I}_d)$	N/A	$d \cdot \lambda_{\max}(\mathbf{L})$

4 LEVERAGING THE LAYER-WISE STRUCTURE

In this section we focus on the block-diagonal case of \mathbf{L} for both **det-CGD1** and **det-CGD2**. In particular, we propose hyper-parameters of **det-CGD1** designed specifically for training NNs. Let us assume that $\mathbf{L} = \text{Diag}(\mathbf{L}_1, \dots, \mathbf{L}_{\ell})$, where $\mathbf{L}_i \in \mathbb{S}_{++}^{d_i}$. This setting is a generalization of the classical smoothness condition, as in the latter case $\mathbf{L}_i = \mathbf{L} \mathbf{I}_{d_i}$ for all $i = 1, \dots, \ell$. Respectively, we choose both the sketches and the stepsize to be block diagonal: $\mathbf{D} = \text{Diag}(\mathbf{D}_1, \dots, \mathbf{D}_{\ell})$ and $\mathbf{S}^k = \text{Diag}(\mathbf{S}_1^k, \dots, \mathbf{S}_{\ell}^k)$, where $\mathbf{D}_i, \mathbf{S}_i^k \in \mathbb{S}_{++}^{d_i}$.

Let us notice that the left hand side of the inequality constraint in (13) has quadratic dependence on \mathbf{D} , while the right hand side is linear. Thus, for every matrix $\mathbf{W} \in \mathbb{S}_{++}^d$, there exists $\gamma > 0$ such that

$$\gamma^2 \lambda_{\max}(\mathbb{E}[\mathbf{S}^k \mathbf{W} \mathbf{L} \mathbf{W} \mathbf{S}^k]) \leq \gamma \lambda_{\min}(\mathbf{W}).$$

Therefore, for $\gamma \mathbf{W}$ we deduce

$$\mathbb{E}[\mathbf{S}^k (\gamma \mathbf{W}) \mathbf{L} (\gamma \mathbf{W}) \mathbf{S}^k] \preceq \gamma^2 \lambda_{\max}(\mathbb{E}[\mathbf{S}^k \mathbf{W} \mathbf{L} \mathbf{W} \mathbf{S}^k]) \mathbf{I}_d \preceq \gamma \lambda_{\min}(\mathbf{W}) \mathbf{I}_d \preceq \gamma \mathbf{W}. \quad (14)$$

The following theorem is based on this simple fact applied to the corresponding blocks of the matrices $\mathbf{D}, \mathbf{L}, \mathbf{S}^k$ for **det-CGD1**.

Theorem 2. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfy Assumptions 1 and 2, with \mathbf{L} admitting the layer-separable structure $\mathbf{L} = \text{Diag}(\mathbf{L}_1, \dots, \mathbf{L}_{\ell})$, where $\mathbf{L}_1, \dots, \mathbf{L}_{\ell} \in \mathbb{S}_{++}^{d_i}$. Choose random matrices $\mathbf{S}_1^k, \dots, \mathbf{S}_{\ell}^k \in \mathbb{S}_{++}^{d_i}$ to satisfy Assumption 3 for all $i \in [\ell]$, and let $\mathbf{S}^k := \text{Diag}(\mathbf{S}_1^k, \dots, \mathbf{S}_{\ell}^k)$. Furthermore, choose matrices $\mathbf{W}_1, \dots, \mathbf{W}_{\ell} \in \mathbb{S}_{++}^{d_i}$ and scalars $\gamma_1, \dots, \gamma_{\ell} > 0$ such that

$$\gamma_i \leq \lambda_{\max}^{-1} \left(\mathbb{E} \left[\mathbf{W}_i^{-1/2} \mathbf{S}_i^k \mathbf{W}_i \mathbf{L}_i \mathbf{W}_i \mathbf{S}_i^k \mathbf{W}_i^{-1/2} \right] \right) \quad \forall i \in [\ell]. \quad (15)$$

Letting $\mathbf{W} := \text{Diag}(\mathbf{W}_1, \dots, \mathbf{W}_{\ell})$, $\Gamma := \text{Diag}(\gamma_1 \mathbf{I}_{d_1}, \dots, \gamma_{\ell} \mathbf{I}_{d_{\ell}})$ and $\mathbf{D} := \Gamma \mathbf{W}$, we get

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left[\left\| \nabla f(x^k) \right\|_{\frac{\Gamma \mathbf{W}}{\det(\Gamma \mathbf{W})^{1/d}}}^2 \right] \leq \frac{2(f(x^0) - f^{\text{inf}})}{\det(\Gamma \mathbf{W})^{1/d} K}. \quad (16)$$

In particular, if the scalars $\{\gamma_i\}$ are chosen to be equal to their maximum allowed values from (15), then the convergence factor of (16) is equal to

$$\det(\Gamma\mathbf{W})^{-\frac{1}{d}} = \left[\prod_{i=1}^{\ell} \lambda_{\max}^{d_i} \left(\mathbb{E} \left[\mathbf{W}_i^{-\frac{1}{2}} \mathbf{S}_i^k \mathbf{W}_i \mathbf{L}_i \mathbf{W}_i \mathbf{S}_i^k \mathbf{W}_i^{-\frac{1}{2}} \right] \right) \right]^{\frac{1}{d}} \det(\mathbf{W}^{-1})^{\frac{1}{d}}.$$

Table 1 contains the (expected) communication complexities of **det-CGD1**, **det-CGD2** and GD for several choices of \mathbf{W} , \mathbf{D} and \mathbf{S}^k . Here are a few comments about the table. We deduce that taking a matrix stepsize without compression (row 1) we improve GD (row 13). A careful analysis reveals that the result in row 5 is always worse than row 7 in terms of both communication and iteration complexity. However, the results in row 6 and row 7 are not comparable in general, meaning that neither of them is universally better. More discussion on this table can be found in the Appendix.

Compression for free. Now, let us focus on row 12, which corresponds to a sampling scheme where the i -th layer is independently selected with probability q_i . Mathematically, it goes as follows:

$$\mathbf{T}_i^k = \frac{\eta_i}{q_i} \mathbf{I}_{d_i}, \quad \text{where } \eta_i \sim \text{Bernoulli}(q_i). \quad (17)$$

Jensen’s inequality implies that

$$\left(\sum_{i=1}^{\ell} q_i d_i \right) \cdot \prod_{i=1}^{\ell} \left(\frac{1}{q_i} \right)^{\frac{d_i}{d}} \geq d. \quad (18)$$

The equality is attained when $q_i = q$ for all $i \in [\ell]$. The expected bits transferred per iteration of this algorithm is then equal to $k_{\text{exp}} = qd$ and the communication complexity equals $d \det(\mathbf{L})^{1/d}$. Comparing with the results for **det-CGD2** with $\text{rand-}k_{\text{exp}}$ on row 11 and using the fact that $\det(\mathbf{L}) \leq \det(\text{diag}(\mathbf{L}))$, we deduce that the Bernoulli scheme is better than the uniform sampling scheme. Notice also, the communication complexity matches the one for the uncompressed **det-CGD2** displayed on row 9. This, in particular means that using the Bern- q sketches we can compress the gradients for free. The latter means that we reduce the number of bits broadcasted at each iteration without losing in the total communication complexity. In particular, when all the layers have the same width d_i , the number of broadcasted bits for each iteration is reduced by a factor of q .

5 DISTRIBUTED SETTING

In this section we describe the distributed versions of our algorithms and present convergence guarantees for them. Let us consider an objective function that is sum decomposable:

$$f(x) := \frac{1}{n} \sum_{i=1}^n f_i(x),$$

where each $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is a differentiable function. We assume that f satisfies Assumption 1 and the component functions satisfy the below condition.

Assumption 4. *Each component function f_i is L_i -smooth and is bounded from below: $f_i(x) \geq f_i^{\text{inf}}$ for all $x \in \mathbb{R}^d$.*

This assumption also implies that f is of matrix smoothness with $\bar{\mathbf{L}} \in \mathbb{S}_{++}^d$, where $\bar{\mathbf{L}} = \frac{1}{n} \sum_{i=1}^n \mathbf{L}_i$. Following the standard FL framework (Konečný et al., 2016; McMahan et al., 2017; Khirirat et al., 2018), we assume that the i -th component function f_i is stored on the i -th client. At each iteration, the clients in parallel compute and compress the local gradient ∇f_i and communicate it to the central server. The server, then aggregates the compressed gradients, computes the next iterate, and in parallel broadcasts it to the clients. See the below pseudo-codes for the details.

Theorem 3. *Let $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfy Assumption 4 and let f satisfy Assumption 1 and Assumption 2 with smoothness matrix \mathbf{L} . If the stepsize satisfies*

$$DLD \preceq D, \quad (19)$$

then the following convergence bound is true for the iterates of Algorithm 1:

$$\min_{0 \leq k \leq K-1} \mathbb{E} \left[\left\| \nabla f(x^k) \right\|_{\frac{\mathbf{D}}{\det(\mathbf{D})^{1/d}}}^2 \right] \leq \frac{2(1 + \frac{\lambda_{\mathbf{D}}}{n})^K (f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} K} + \frac{2\lambda_{\mathbf{D}}\Delta^{\text{inf}}}{\det(\mathbf{D})^{1/d} n}, \quad (20)$$

where $\Delta^{\text{inf}} := f^{\text{inf}} - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}}$ and

$$\lambda_{\mathbf{D}} := \max_i \left\{ \lambda_{\max} \left(\mathbb{E} \left[\mathbf{L}_i^{\frac{1}{2}} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{L}_i^{\frac{1}{2}} \right] \right) \right\}.$$

Algorithm 1 Distributed **det-CGD1**

1: **Input:** Starting point x^0 , stepsize matrix \mathbf{D} , number of iterations K
2: **for** $k = 0, 1, 2, \dots, K - 1$ **do**
3: The devices in parallel:
4: sample $\mathbf{S}_i^k \sim \mathcal{S}$;
5: compute $\mathbf{S}_i^k \nabla f_i(x^k)$;
6: broadcast $\mathbf{S}_i^k \nabla f_i(x^k)$.
7: **The server:**
8: combines $g^k = \frac{\mathbf{D}}{n} \sum_{i=1}^n \mathbf{S}_i^k \nabla f_i(x^k)$;
9: computes $x^{k+1} = x^k - g^k$;
10: broadcasts x^{k+1} .
11: **end for**
12: **Return:** x^K

Algorithm 2 Distributed **det-CGD2**

1: **Input:** Starting point x^0 , stepsize matrix \mathbf{D} , number of iterations K
2: **for** $k = 0, 1, 2, \dots, K - 1$ **do**
3: The devices in parallel:
4: sample $\mathbf{T}_i^k \sim \mathcal{T}$;
5: compute $\mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k)$;
6: broadcast $\mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k)$.
7: **The server:**
8: combines $g^k = \frac{1}{n} \sum_i \mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k)$;
9: computes $x^{k+1} = x^k - g^k$;
10: broadcasts x^{k+1} .
11: **end for**
12: **Return:** x^K

The same result is true for Algorithm 2 with a different constant $\lambda_{\mathbf{D}}$. The proof of Theorem 3 and its analogue for Algorithm 2 are presented in the Appendix. The analysis is largely inspired by (Khaled & Richtárik, 2020, Theorem 1). Now, let us examine the right-hand side of (20). We start by observing that the first term has exponential dependence in K . However, the term inside the brackets, $1 + \lambda_{\mathbf{D}}/n$, depends on the stepsize \mathbf{D} . Furthermore, it has a second-order dependence on \mathbf{D} , implying that $\lambda_{\alpha \mathbf{D}} = \alpha^2 \lambda_{\mathbf{D}}$, as opposed to $\det(\alpha \mathbf{D})^{1/d}$, which is linear in α . Therefore, we can choose a small enough coefficient α to ensure that $\lambda_{\mathbf{D}}$ is of order n/K . This means that for a fixed number of iterations K , we choose the matrix stepsize to be "small enough" to guarantee that the numerator of the first term is bounded. The following corollary summarizes these arguments, and its proof can be found in the Appendix.

Corollary 1. We reach an ε -stationarity, that is the right-hand side of (20) is upper bounded by ε^2 , if the following conditions are satisfied:

$$\mathbf{D} \mathbf{L} \mathbf{D} \preceq \mathbf{D}, \quad \lambda_{\mathbf{D}} \leq \min \left\{ \frac{n}{K}, \frac{n\varepsilon^2}{4\Delta^{\text{inf}} \det(\mathbf{D})^{1/d}} \right\}, \quad K \geq \frac{12(f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} \varepsilon^2}. \quad (21)$$

Proposition 3 in the Appendix proves that these conditions with respect to \mathbf{D} are convex. In order to minimize the iteration complexity for getting ε^2 error, one needs to solve the following optimization problem

$$\text{minimize} \quad \log \det(\mathbf{D}^{-1})$$

$$\text{subject to} \quad \mathbf{D} \text{ satisfies (21).}$$

Choosing the optimal stepsize for Algorithm 1 is analogous to solving (13). One can formulate the distributed counterpart of Theorem 2 and attempt to solve it for different sketches. Furthermore, this leads to a convex matrix minimization problem involving \mathbf{D} . We provide a formal proof of this property in the Appendix. Similar to the single-node case, computational methods can be employed using the CVXPY package. However, some additional effort is required to transform (21) into the disciplined convex programming (DCP) format.

The second term in (20) corresponds to the convergence neighborhood of the algorithm. It does not depend on the number of iteration, thus it remains unchanged, after we choose the stepsize. Nevertheless, it depends on the number of clients n . In general, the term Δ^{inf}/n can be unbounded, when $n \rightarrow +\infty$. However, per Corollary 1, we require $\lambda_{\mathbf{D}}$ to be upper-bounded by n/K . Thus,




Figure 1: Comparison of standard DCGD, DCGD with matrix smoothness, D-det-CGD1 and D-det-CGD2 with optimal diagonal stepsizes under rand-1 sketch. The stepsize for standard DCGD is determined using (Khaled & Richtárik, 2020, Proposition 4), the stepsize for DCGD with matrix smoothness along with D_1 , D_2 is determined using Corollary 1, the error level is set to be $\varepsilon^2 = 0.0001$. Here $G_{K,D} := \frac{1}{K} \left(\sum_{k=0}^{K-1} \|\nabla f(x^k)\|_{D/\det(D)^{1/d}}^2 \right)$.

the neighborhood term will indeed converge to zero when $K \rightarrow +\infty$, if we choose the stepsize accordingly.

We compare our results with the existing results for DCGD. In particular we use the technique from Khaled & Richtárik (2020) for the scalar smooth DCGD with scalar stepsizes with the results from (Khaled & Richtárik, 2020, Corollary 1). See the Appendix for the details on the analysis of Khaled & Richtárik (2020). Finally, we back up our theoretical findings with experiments. See Figure 1 for a simple experiment confirming that Algorithms 1 and 2 have better iteration and communication complexity compared to scalar stepped DCGD. The graphs of the two proposed algorithms coincide, as the diagonal stepsize and the diagonal sketch commute, resulting in the same method. For more details on the experiments we refer the reader to the corresponding section in the Appendix.

6 CONCLUSION

In this paper, we enhance compressed gradient descent method with matrix-valued stepsize for general non-convex objectives. Convergence guarantees are provided for the algorithms both in the single node case and the distributed setting. By considering the layer-wise structure of models such as neural networks, we are able to design compression mechanisms that achieve compression for free. This is the first time matrix stepsize is used and analyzed together with compression in the non-convex case. Our theoretical findings are supported with abundant numerical experiments.

6.1 LIMITATIONS

It is worth noting that every point in \mathbb{R}^d can be enclosed within some volume 1 ellipsoid. Therefore, having the average D -norm of the gradient bounded by a small number does not guarantee that the average Euclidean norm is small. However, for a fixed D , the standard Euclidean norm is equivalent to the weighted D -norm. This is due to

$$\frac{\lambda_{\min}(D)}{\det(D)^{1/d}} \cdot \|\nabla f(x)\|^2 \leq \|\nabla f(x)\|_{D/(\det(D))^{1/d}}^2 \leq \frac{\lambda_{\max}(D)}{\det(D)^{1/d}} \cdot \|\nabla f(x)\|^2. \quad (22)$$

This relation is further validated by our experiments described in the Appendix.

REFERENCES

- Mehiddin Al-Baali and H Khalfan. An overview of some practical quasi-newton methods for unconstrained optimization. *Sultan Qaboos University Journal for Science [SQUJS]*, 12(2): 199–209, 2007.
- Mehiddin Al-Baali, Emilio Spedicato, and Francesca Maggioni. Broyden’s quasi-Newton methods for a nonlinear system of equations and unconstrained optimization: a review and open problems. *Optimization Methods and Software*, 29(5):937–954, 2014.
- Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. QSGD: Communication-efficient SGD via gradient quantization and encoding. *Advances in neural information processing systems*, 30, 2017.
- Charles G Broyden. A class of methods for solving nonlinear simultaneous equations. *Mathematics of computation*, 19(92):577–593, 1965.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, 2011.
- Marina Danilova, Pavel Dvurechensky, Alexander Gasnikov, Eduard Gorbunov, Sergey Guminov, Dmitry Kamzolov, and Innokentiy Shibaev. Recent theoretical advances in non-convex optimization. In *High-Dimensional Optimization and Probability: With a View Towards Data Science*, pp. 79–163. Springer, 2022.
- John E Dennis, Jr and Jorge J Moré. Quasi-Newton methods, motivation and theory. *SIAM review*, 19(1):46–89, 1977.
- Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research*, 17(1):2909–2913, 2016.
- Aritra Dutta, El Houcine Bergou, Ahmed M Abdelmoniem, Chen-Yu Ho, Atal Narayan Sahu, Marco Canini, and Panos Kalnis. On the discrepancy between the theoretical analysis and practical implementations of compressed communication for distributed deep learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 3817–3824, 2020.
- Darina Dvinskikh, Aleksandr Ogaltsov, Alexander Gasnikov, Pavel Dvurechensky, Alexander Tyurin, and Vladimir Spokoiny. Adaptive gradient descent for convex and non-convex stochastic optimization. *arXiv preprint arXiv:1911.08380*, 2019.
- Boris Ginsburg, Patrice Castonguay, Oleksii Hrinchuk, Oleksii Kuchaiev, Vitaly Lavrukhin, Ryan Leary, Jason Li, Huyen Nguyen, Yang Zhang, and Jonathan M Cohen. Stochastic gradient methods with layer-wise adaptive moments for training of deep networks. *arXiv preprint arXiv:1905.11286*, 2019.
- Eduard Gorbunov, Konstantin P Burlachenko, Zhize Li, and Peter Richtárik. Marina: Faster non-convex distributed learning with compression. In *International Conference on Machine Learning*, pp. 3788–3798. PMLR, 2021.
- Robert M Gower and Peter Richtárik. Randomized iterative methods for linear systems. *SIAM Journal on Matrix Analysis and Applications*, 36(4):1660–1690, 2015.
- 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.
- William B Gragg and Richard A Tapia. Optimal error bounds for the Newton–Kantorovich theorem. *SIAM Journal on Numerical Analysis*, 11(1):10–13, 1974.
- Michael Grant, Stephen Boyd, and Yinyu Ye. Disciplined convex programming. *Global optimization: From theory to implementation*, pp. 155–210, 2006.

- SV Guminov, Yu E Nesterov, PE Dvurechensky, and AV Gasnikov. Accelerated primal-dual gradient descent with linesearch for convex, nonconvex, and nonsmooth optimization problems. In *Doklady Mathematics*, volume 99, pp. 125–128. Springer, 2019.
- Filip Hanzely, Konstantin Mishchenko, and Peter Richtárik. SEGA: Variance reduction via gradient sketching. *Advances in Neural Information Processing Systems*, 31, 2018.
- Samuel Horvath, Chen-Yu Ho, Ludovit Horvath, Atal Narayan Sahu, Marco Canini, and Peter Richtárik. Natural compression for distributed deep learning. *arXiv preprint arXiv:1905.10988*, 2019.
- Samuel Horváth, Dmitry Kovalev, Konstantin Mishchenko, Peter Richtárik, and Sebastian Stich. Stochastic distributed learning with gradient quantization and double-variance reduction. *Optimization Methods and Software*, 38(1):91–106, 2023.
- Rustem Islamov, Xun Qian, and Peter Richtárik. Distributed second order methods with fast rates and compressed communication. In *International conference on machine learning*, pp. 4617–4628. PMLR, 2021.
- Sashank J Reddi, Suvrit Sra, Barnabas Poczos, and Alexander J Smola. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. *Advances in neural information processing systems*, 29, 2016.
- Prateek Jain, Purushottam Kar, et al. Non-convex optimization for machine learning. *Foundations and Trends® in Machine Learning*, 10(3-4):142–363, 2017.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
- Ahmed Khaled and Peter Richtárik. Better theory for sgd in the nonconvex world. *arXiv preprint arXiv:2002.03329*, 2020.
- Sarit Khirirat, Hamid Reza Feysmahdavian, and Mikael Johansson. Distributed learning with compressed gradients. *arXiv preprint arXiv:1806.06573*, 2018.
- Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- Yin Tat Lee, Zhao Song, and Qiuyi Zhang. Solving empirical risk minimization in the current matrix multiplication time. In *Conference on Learning Theory*, pp. 2140–2157. PMLR, 2019.
- Zhize Li, Dmitry Kovalev, Xun Qian, and Peter Richtárik. Acceleration for compressed gradient descent in distributed and federated optimization. *arXiv preprint arXiv:2002.11364*, 2020.
- Zhize Li, Hongyan Bao, Xiangliang Zhang, and Peter Richtárik. PAGE: A simple and optimal probabilistic gradient estimator for nonconvex optimization. In *International conference on machine learning*, pp. 6286–6295. PMLR, 2021.
- Artavazd Maranjyan, Mher Safaryan, and Peter Richtárik. GradSkip: Communication-Accelerated Local Gradient Methods with Better Computational Complexity. *arXiv preprint arXiv:2210.16402*, 2022.
- H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017. URL <http://arxiv.org/abs/1602.05629>.
- George J Miel. Majorizing sequences and error bounds for iterative methods. *Mathematics of Computation*, 34(149):185–202, 1980.

- Konstantin Mishchenko, Eduard Gorbunov, Martin Takáč, and Peter Richtárik. Distributed learning with compressed gradient differences. *arXiv preprint arXiv:1901.09269*, 2019.
- Konstantin Mishchenko, Grigory Malinovsky, Sebastian Stich, and Peter Richtarik. ProxSkip: Yes! Local gradient steps provably lead to communication acceleration! Finally! In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 15750–15769. PMLR, 17–23 Jul 2022.
- Eric Moulines and Francis Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. *Advances in neural information processing systems*, 24, 2011.
- Lianke Qin, Zhao Song, Lichen Zhang, and Danyang Zhuo. An online and unified algorithm for projection matrix vector multiplication with application to empirical risk minimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 101–156. PMLR, 2023.
- Peter Richtárik, Igor Sokolov, and Ilyas Fatkhullin. EF21: A new, simpler, theoretically better, and practically faster error feedback. *Advances in Neural Information Processing Systems*, 34: 4384–4396, 2021.
- Mher Safaryan, Filip Hanzely, and Peter Richtárik. Smoothness matrices beat smoothness constants: Better communication compression techniques for distributed optimization. *Advances in Neural Information Processing Systems*, 34:25688–25702, 2021.
- Mher Safaryan, Egor Shulgin, and Peter Richtárik. Uncertainty principle for communication compression in distributed and federated learning and the search for an optimal compressor. *Information and Inference: A Journal of the IMA*, 11(2):557–580, 2022.
- Zhao Song, Yitan Wang, Zheng Yu, and Lichen Zhang. Sketching for first order method: Efficient algorithm for low-bandwidth channel and vulnerability. In *International Conference on Machine Learning*, pp. 32365–32417. PMLR, 2023.
- Sebastian U Stich. Unified optimal analysis of the (stochastic) gradient method. *arXiv preprint arXiv:1907.04232*, 2019.
- Bokun Wang, Mher Safaryan, and Peter Richtárik. Theoretically better and numerically faster distributed optimization with smoothness-aware quantization techniques. *Advances in Neural Information Processing Systems*, 35:9841–9852, 2022.
- Tetsuro Yamamoto. A convergence theorem for Newton-like methods in Banach spaces. *Numerische Mathematik*, 51:545–557, 1987.
- Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*, 2015.
- Adams Wei Yu, Lei Huang, Qihang Lin, Ruslan Salakhutdinov, and Jaime Carbonell. Block-normalized gradient method: An empirical study for training deep neural network. *arXiv preprint arXiv:1707.04822*, 2017.
- Jingzhao Zhang, Sai Praneeth Karimireddy, Andreas Veit, Seungyeon Kim, Sashank Reddi, Sanjiv Kumar, and Suvrit Sra. Why are adaptive methods good for attention models? *Advances in Neural Information Processing Systems*, 33:15383–15393, 2020.
- Qinghe Zheng, Xinyu Tian, Nan Jiang, and Mingqiang Yang. Layer-wise learning based stochastic gradient descent method for the optimization of deep convolutional neural network. *Journal of Intelligent & Fuzzy Systems*, 37(4):5641–5654, 2019.
- Luisa M Zintgraf, Taco S Cohen, Tameem Adel, and Max Welling. Visualizing deep neural network decisions: Prediction difference analysis. *arXiv preprint arXiv:1702.04595*, 2017.

CONTENTS

A	Single node case	14
A.1	Proof of Theorem 1	14
A.2	Proof of Proposition 1	15
B	Layer-wise case	16
B.1	Proof of Theorem 2	16
B.2	Bernoulli- q sketch for det-CGD2	16
B.3	General cases for det-CGD1	18
B.4	General cases for det-CGD2	18
B.5	Interpretations of Table 1	18
B.5.1	Comparison of row 5 and 7	18
B.5.2	Comparison of row 6 and 7	19
C	Distributed case	19
C.1	Future work	19
C.2	Proof of Theorem 3	19
C.3	Convexity of the constraints	22
C.3.1	Proof of Corollary 1	24
C.4	Distributed det-CGD2	24
C.4.1	Analysis of distributed det-CGD2	24
C.4.2	Optimal stepsize	26
C.5	DCGD with constant stepsize	26
D	Proofs of technical lemmas	27
D.1	Proof of Lemma 1	27
D.2	Proof of Lemma 2	28
D.3	Proof of Lemma 3	28
D.4	Proof of Lemma 4	28
D.5	Proof of Lemma 5	28
D.6	Proof of Lemma 6	29
D.7	Proof of Lemma 7	29
E	Experiments	29
E.1	Single node case	29
E.1.1	Comparison to CGD with scalar stepsize, scalar smoothness constant	30
E.1.2	Comparison of the two algorithms under the same stepsize	31
E.2	Distributed case	31
E.2.1	Comparison to standard DCGD in the distributed case	33

A SINGLE NODE CASE

A.1 PROOF OF THEOREM 1

i) Using Assumption 2 with $x = x^{k+1} = x^k - \mathbf{D}\mathbf{S}^k \nabla f(x^k)$ and $y = x^k$, we get

$$\begin{aligned} \mathbb{E}[f(x^{k+1}) | x^k] &\leq \mathbb{E}\left[f(x^k) + \langle \nabla f(x^k), -\mathbf{D}\mathbf{S}^k \nabla f(x^k) \rangle\right. \\ &\quad \left. + \frac{1}{2} \langle \mathbf{L}(-\mathbf{D}\mathbf{S}^k \nabla f(x^k)), -\mathbf{D}\mathbf{S}^k \nabla f(x^k) \rangle | x^k\right] \\ &= f(x^k) - \langle \nabla f(x^k), \mathbf{D}\mathbb{E}[\mathbf{S}^k] \nabla f(x^k) \rangle + \frac{1}{2} \langle \mathbb{E}[\mathbf{S}^k \mathbf{D}\mathbf{L}\mathbf{D}\mathbf{S}^k] \nabla f(x^k), \nabla f(x^k) \rangle. \end{aligned}$$

From the unbiasedness of the sketch \mathbf{S}^k

$$\begin{aligned} \mathbb{E}[f(x^{k+1}) | x^k] &\leq f(x^k) - \langle \nabla f(x^k), \mathbf{D}\nabla f(x^k) \rangle + \frac{1}{2} \langle \mathbb{E}[\mathbf{S}^k \mathbf{D}\mathbf{L}\mathbf{D}\mathbf{S}^k] \nabla f(x^k), \nabla f(x^k) \rangle \\ &\stackrel{(7)}{\leq} f(x^k) - \langle \nabla f(x^k), \mathbf{D}\nabla f(x^k) \rangle + \frac{1}{2} \langle \mathbf{D}\nabla f(x^k), \nabla f(x^k) \rangle \\ &= f(x^k) - \frac{1}{2} \langle \nabla f(x^k), \mathbf{D}\nabla f(x^k) \rangle \\ &= f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_{\mathbf{D}}^2. \end{aligned} \tag{23}$$

Next, by subtracting f^{inf} from both sides of (23), taking expectation and applying the tower property, we get

$$\begin{aligned} \mathbb{E}[f(x^{k+1})] - f^{\text{inf}} &= \mathbb{E}[\mathbb{E}[f(x^{k+1}) | x^k]] - f^{\text{inf}} \\ &\stackrel{(23)}{\leq} \mathbb{E}\left[f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_{\mathbf{D}}^2\right] - f^{\text{inf}} \\ &= \mathbb{E}[f(x^k)] - f^{\text{inf}} - \frac{1}{2} \mathbb{E}[\|\nabla f(x^k)\|_{\mathbf{D}}^2]. \end{aligned}$$

Letting $\Delta^k := \mathbb{E}[f(x^k)] - f^{\text{inf}}$, the last inequality can be written as $\Delta^{k+1} \leq \Delta^k - \frac{1}{2} \mathbb{E}[\|\nabla f(x^k)\|_{\mathbf{D}}^2]$. Summing these inequalities for $k = 0, 1, \dots, K-1$, we get a telescoping effect leading to

$$\Delta^K \leq \Delta^0 - \frac{1}{2} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla f(x^k)\|_{\mathbf{D}}^2].$$

It remains to rearrange the terms of this inequality, divide both sides by $K \det(\mathbf{D})^{1/d}$, and use the inequality $\Delta^K \geq 0$.

ii) Similar to the previous case, using matrix smoothness for $x = x^{k+1} = x^k - \mathbf{T}^k \mathbf{D} \nabla f(x^k)$ and $y = x^k$, we get

$$\mathbb{E}[f(x^{k+1}) | x^k] \leq \mathbb{E}\left[f(x^k) + \langle \nabla f(x^k), -\mathbf{T}^k \mathbf{D} \nabla f(x^k) \rangle\right] \tag{24}$$

$$\begin{aligned} &+ \frac{1}{2} \langle \mathbf{L}(-\mathbf{T}^k \mathbf{D} \nabla f(x^k)), -\mathbf{T}^k \mathbf{D} \nabla f(x^k) \rangle | x^k \\ &= f(x^k) - \langle \nabla f(x^k), \mathbb{E}[\mathbf{T}^k] \mathbf{D} \nabla f(x^k) \rangle \\ &+ \frac{1}{2} \langle \mathbb{E}[\mathbf{D}(\mathbf{T}^k)^\top \mathbf{L}\mathbf{T}^k \mathbf{D}] \nabla f(x^k), \nabla f(x^k) \rangle. \end{aligned} \tag{25}$$

From Assumption 3 and condition (8) we deduce

$$\begin{aligned} \mathbb{E}[f(x^{k+1}) | x^k] &\leq f(x^k) - \langle \nabla f(x^k), \mathbf{D}\nabla f(x^k) \rangle + \frac{1}{2} \langle \mathbf{D}\nabla f(x^k), \nabla f(x^k) \rangle \\ &= f(x^k) - \frac{1}{2} \langle \nabla f(x^k), \mathbf{D}\nabla f(x^k) \rangle \\ &= f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_{\mathbf{D}}^2. \end{aligned} \tag{26}$$

Thus, we obtain the same upper bound on $\mathbb{E} [f(x^{k+1}) | x^k]$ as in (23). Following the steps from the first part, we conclude the proof.

A.2 PROOF OF PROPOSITION 1

Let us rewrite (7) using quadratic forms. That is for every non-zero $v \in \mathbb{R}^d$, the following inequality must be true:

$$v^\top \mathbb{E} [\mathbf{S}^k \mathbf{D} \mathbf{L} \mathbf{D} \mathbf{S}^k] v \leq v^\top \mathbf{D} v, \quad \forall v \neq 0$$

Notice that both sides of this inequality are real numbers, thus can be written equivalently as

$$\text{tr}(v^\top \mathbb{E} [\mathbf{S}^k \mathbf{D} \mathbf{L} \mathbf{D} \mathbf{S}^k] v) \leq \text{tr}(v^\top \mathbf{D} v), \quad \forall v \neq 0$$

The LHS can be modified in the following way

$$\begin{aligned} \text{tr}(v^\top \mathbb{E} [\mathbf{S}^k \mathbf{D} \mathbf{L} \mathbf{D} \mathbf{S}^k] v) &\stackrel{\text{I}}{=} \text{tr}(\mathbb{E} [v^\top \mathbf{S}^k \mathbf{D} \mathbf{L} \mathbf{D} \mathbf{S}^k v]) \\ &\stackrel{\text{II}}{=} \mathbb{E} [\text{tr}(v^\top \mathbf{S}^k \mathbf{D} \mathbf{L} \mathbf{D} \mathbf{S}^k v)] \\ &\stackrel{\text{III}}{=} \mathbb{E} [\text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbf{S}^k v v^\top \mathbf{S}^k \mathbf{D} \mathbf{L}^{\frac{1}{2}})] \\ &\stackrel{\text{IV}}{=} \text{tr}(\mathbb{E} [\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbf{S}^k v v^\top \mathbf{S}^k \mathbf{D} \mathbf{L}^{\frac{1}{2}}]) \\ &\stackrel{\text{V}}{=} \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbb{E} [\mathbf{S}^k v v^\top \mathbf{S}^k] \mathbf{D} \mathbf{L}^{\frac{1}{2}}), \end{aligned}$$

where I, V are due to the linearity of expectation, II, IV are due to the linearity of trace operator, III is obtained using the cyclic property of trace. Therefore, we can write the condition (7) equivalently as

$$\text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbb{E} [\mathbf{S}^k v v^\top \mathbf{S}^k] \mathbf{D} \mathbf{L}^{\frac{1}{2}}) \leq \text{tr}(v v^\top \mathbf{D}), \quad \forall v \neq 0.$$

We then define function $g_v : \mathbb{S}_{++}^d \rightarrow \mathbb{R}$ for some fixed $v \neq 0$ as

$$g_v(\mathbf{D}) := \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbb{E} [\mathbf{S}^k v v^\top \mathbf{S}^k] \mathbf{D} \mathbf{L}^{\frac{1}{2}}) - \text{tr}(v v^\top \mathbf{D}). \quad (27)$$

We want to show that for every fixed $v \neq 0$, g is a convex function w.r.t \mathbf{D} , so that in this case, the sub-level set $\{\mathbf{D} \in \mathbb{S}_{++}^d \mid g_v(\mathbf{D}) \leq 0\}$ is convex.

- Notice that $v v^\top$ is a rank-1 matrix whose eigenvalues are all zero except one of them is $\|v\|^2 > 0$. We also have $(v v^\top)^\top = (v^\top)^\top v^\top = v v^\top$, so it is also a symmetric matrix. Thus we conclude that $v v^\top \in \mathbb{S}_+^d$ for every choice of v , we use $\mathbf{V} = v v^\top$ to denote it.
- If $\mathbf{S}^k = \mathbf{O}_d$, then the first term is equal to \mathbf{O}_d and the function $g_v(\mathbf{D})$ is linear, thus, also convex. Now, let us assume \mathbf{S}^k is nonzero. Similarly $\mathbf{S}^k v v^\top \mathbf{S}^k = \mathbf{S}^k v (\mathbf{S}^k v)^\top$ is also a symmetric positive semi-definite matrix whose eigenvalues are all 0 except one of them is $\|\mathbf{S}^k v\|^2$, this tells us that its expectation over \mathbf{S}^k is still a symmetric positive semi-definite matrix, we use $\mathbf{R} = \mathbb{E} [\mathbf{S}^k v v^\top \mathbf{S}^k]$ to denote it.

Now we can write function g_v as

$$g_v(\mathbf{D}) = \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbf{R} \mathbf{D} \mathbf{L}^{\frac{1}{2}}) - \text{tr}(\mathbf{V} \mathbf{D}).$$

We present the following lemma that guarantees the convexity of the first term.

Lemma 1. For every matrix $\mathbf{R} \in \mathbb{S}_+^d$, we define

$$f(\mathbf{D}) = \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbf{R} \mathbf{D} \mathbf{L}^{\frac{1}{2}}), \quad (28)$$

where $\mathbf{L}, \mathbf{D} \in \mathbb{S}_{++}^d$. Then function $f : \mathbb{S}_{++}^d \rightarrow \mathbb{R}$ is a convex function.

The proof can be found in Appendix D.1. According to Lemma 1, the first term of $g_v(\mathbf{D})$ is a convex function, and we know that the second term is linear in \mathbf{D} . As a result, $g_v(\mathbf{D})$ is a convex function w.r.t. \mathbf{D} for every $v \neq 0$, thus the sub-level set $\{\mathbf{D} \in \mathbb{S}_{++}^d \mid g_v(\mathbf{D}) \leq 0\}$ is a convex set for every $v \neq 0$. The intersection of all those convex sets corresponding to every $v \neq 0$ is still a convex set, which tells us the original condition (7) is convex. This concludes the proof of the proposition.

B LAYER-WISE CASE

In this section, we provide interpretations about some of the results and conclusions we had in Section 4.

B.1 PROOF OF THEOREM 2

Note that $\mathbb{E} [S^k D L D S^k] = \text{Diag} (\mathbf{Q}_1^k, \dots, \mathbf{Q}_\ell^k)$, where $\mathbf{Q}_i^k := \gamma_i^2 \mathbb{E} [S_i^k \mathbf{W}_i \mathbf{L}_i \mathbf{D}_i S_i^k]$. In other words,

$$\mathbb{E} [S^k D L D S^k] = \begin{pmatrix} \mathbf{Q}_1^k & 0 & \dots & 0 \\ 0 & \mathbf{Q}_2^k & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathbf{Q}_\ell^k \end{pmatrix},$$

which means that (7) holds if and only if $\mathbf{Q}_i^k \preceq \gamma_i \mathbf{W}_i$ for all $i \in [\ell]$, which holds if and only if (15) holds. Therefore, Theorem 1 applies, and we conclude that

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} [\|\nabla f(x^k)\|_{\Gamma \mathbf{W}}^2] \leq \frac{2(f(x^0) - f^{\text{inf}})}{K}. \quad (29)$$

To obtain (16), it remains to multiply both sides of (29) by $\frac{1}{\det(\Gamma \mathbf{W})^{1/d}}$.

B.2 BERNOULLI- q SKETCH FOR `DET-CGD2`

The following corollary of Theorem 2 computes the communication complexity of `det-CGD2` in the block diagonal setting with Bernoulli- q .

Corollary 2. Let \mathbf{T}_i^k for the i -th layer in `det-CGD2` be the Bern- q_i sketch which is defined as

$$\mathbf{T}_i^k = \frac{\eta_i}{q_i} \mathbf{I}_{d_i}, \quad \text{where } \eta_i \sim \text{Bernoulli}(q_i). \quad (30)$$

Then, the communication complexity of `det-CGD2` is given by,

$$\frac{2(f(x^0) - f^{\text{inf}})}{\varepsilon^2} \left(\sum_{i=1}^{\ell} q_i d_i \right) \prod_{i=1}^{\ell} \left(\frac{1}{q_i} \right)^{\frac{d_i}{d}} \det(\mathbf{L})^{\frac{1}{d}}. \quad (31)$$

Furthermore, the communication complexity is minimized if the probabilities when $q_i = q$, $\forall i \in [\ell]$ and the minimum value is equal to

$$\frac{2(f(x^0) - f^{\text{inf}}) \cdot dx \det(\mathbf{L})^{\frac{1}{d}}}{\varepsilon^2}. \quad (32)$$

Proof. For `det-CGD2`, its convergence requires (11). We are using Bernoulli sketch here, so we deduce that

$$\begin{aligned} \mathbb{E} [\mathbf{T}^k \mathbf{L} \mathbf{T}^k] &= \mathbb{E} [\text{Diag}(\mathbf{T}_1^k \mathbf{L}_1 \mathbf{T}_1^k, \dots, \mathbf{T}_\ell^k \mathbf{L}_\ell \mathbf{T}_\ell^k)] \\ &= \text{Diag} (\mathbb{E} [\mathbf{T}_1^k \mathbf{L}_1 \mathbf{T}_1^k], \dots, \mathbb{E} [\mathbf{T}_\ell^k \mathbf{L}_\ell \mathbf{T}_\ell^k]). \end{aligned}$$

Using the fact that for each block, we have

$$\mathbb{E} [\mathbf{T}_i^k \mathbf{L}_i \mathbf{T}_i^k] = (1 - q_i) \mathbf{O}_{d_i} \mathbf{L}_i \mathbf{O}_{d_i} + q_i \cdot \frac{1}{q_i^2} \mathbf{I}_{d_i} \mathbf{L}_i \mathbf{I}_{d_i} = \frac{\mathbf{L}_i}{q_i},$$

we obtain

$$\mathbb{E} [\mathbf{T}^k \mathbf{L} \mathbf{T}^k] = \text{Diag} \left(\frac{\mathbf{L}_1}{q_1}, \dots, \frac{\mathbf{L}_\ell}{q_\ell} \right).$$

Recalling (12), the best stepsize possible is therefore given by

$$\begin{aligned} \mathbf{D} &= (\mathbb{E} [\mathbf{T}^k \mathbf{L} \mathbf{T}^k])^{-1} \\ &= \text{Diag}^{-1} \left(\frac{\mathbf{L}_1}{q_1}, \dots, \frac{\mathbf{L}_\ell}{q_\ell} \right) \\ &= \text{Diag} (q_1 \mathbf{L}_1^{-1}, \dots, q_\ell \mathbf{L}_\ell^{-1}). \end{aligned}$$

From (10), we know that in order for **det-CGD2** to converge to ε^2 -error level (i.e. ε -stationarity), we need

$$\frac{2(f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{\frac{1}{d}} K} \leq \varepsilon^2,$$

which means that we need

$$K \geq \frac{2(f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{\frac{1}{d}} \varepsilon^2} = \frac{1}{\det(\mathbf{D})^{\frac{1}{d}}} \cdot \frac{2(f(x^0) - f^{\text{inf}})}{\varepsilon^2},$$

iterations. For each iteration, the number of bits sent in expectation is equal to $\sum_{i=1}^{\ell} q_i d_i$. As a result, the communication complexity is given by, if we leave out the constant factor $2(f(x^0) - f^{\text{inf}})/\varepsilon^2$,

$$\begin{aligned} \left(\sum_{i=1}^{\ell} q_i d_i \right) \cdot \frac{1}{\det(\mathbf{D})^{\frac{1}{d}}} &= \left(\sum_{i=1}^{\ell} q_i d_i \right) \cdot \det(\mathbf{D}^{-1})^{\frac{1}{d}} \\ &= \left(\sum_{i=1}^{\ell} q_i d_i \right) \cdot \left(\prod_{i=1}^{\ell} \det\left(\frac{\mathbf{L}_i}{q_i}\right) \right)^{\frac{1}{d}} \\ &= \left(\sum_{i=1}^{\ell} q_i d_i \right) \cdot \prod_{i=1}^{\ell} \left(\frac{1}{q_i}\right)^{\frac{d_i}{d}} \left(\prod_{i=1}^{\ell} \det(\mathbf{L}_i) \right)^{\frac{1}{d}} \\ &= \left(\sum_{i=1}^{\ell} q_i d_i \right) \cdot \prod_{i=1}^{\ell} \left(\frac{1}{q_i}\right)^{\frac{d_i}{d}} \cdot \det(\mathbf{L})^{\frac{1}{d}}. \end{aligned}$$

To obtain the optimal probability q_i , we can do the following transformation

$$\left(\sum_{i=1}^{\ell} q_i d_i \right) \cdot \frac{1}{\det(\mathbf{D})^{\frac{1}{d}}} = \left(\sum_{i=1}^{\ell} q_i \frac{d_i}{d} \right) \cdot \prod_{i=1}^{\ell} \left(\frac{1}{q_i}\right)^{\frac{d_i}{d}} \cdot d \det(\mathbf{L})^{\frac{1}{d}}.$$

Therefore, it is equivalent to minimizing the coefficient

$$\left(\sum_{i=1}^{\ell} q_i \frac{d_i}{d} \right) \cdot \prod_{i=1}^{\ell} \left(\frac{1}{q_i}\right)^{\frac{d_i}{d}}.$$

If we denote $\alpha_i = \frac{d_i}{d}$, then we know that $\alpha_i \in (0, 1]$ and $\sum_{i=1}^{\ell} \alpha_i = 1$, the above coefficient turns into

$$\left(\sum_{i=1}^{\ell} \alpha_i q_i \right) \prod_{i=1}^{\ell} \left(\frac{1}{q_i}\right)^{\alpha_i}.$$

From the strict log-concavity of the $\log(\cdot)$ function and Jensen's inequality we have

$$\left(\sum_{i=1}^{\ell} \alpha_i q_i \right) \geq \prod_{i=1}^{\ell} q_i^{\alpha_i}.$$

The identity is obtained if and only if $q_i = q_j$, for all $i \neq j$. Thus, we get

$$\left(\sum_{i=1}^{\ell} \alpha_i q_i \right) \prod_{i=1}^{\ell} \left(\frac{1}{q_i}\right)^{\alpha_i} \geq 1,$$

which in its turn implies that the minimum of expected communication complexity is equal to $d \cdot \det(\mathbf{L})^{\frac{1}{d}}$. The equality is achieved when the probabilities are equal. This concludes the proof. \square

By utilizing the block diagonal structure, we are able to design special sketches that allow us to compress for free. This can be seen from row 12, where the communication complexity of using Bernoulli compressor with equal probabilities for **det-CGD2** in expectation is the same with GD, but the number of bits sent per iteration is reduced.

B.3 GENERAL CASES FOR **DET-CGD1**

The first part (row 1 to row 8) of Table 1 records the communication complexities of **det-CGD1** in the block diagonal setting and in the general setting. Depending on the types of sketches \mathcal{S}_i^k and matrices \mathbf{W}_i we are using, we can calculate the optimal scaling factor γ_i using Theorem 2. According to (10), in order to reach an error level of ε^2 , we need

$$K \geq \frac{1}{\det(\mathbf{D})^{\frac{1}{d}}} \cdot \frac{2(f(x^0) - f^{\text{inf}})}{\varepsilon^2}, \quad (33)$$

where K is the number of iterations in total. We can then obtain the communication complexity taking into account the number of bits transferred in each iteration in the block diagonal case. The same applies to the general case which can be viewed as a special case of the block diagonal setting where there is only 1 block.

B.4 GENERAL CASES FOR **DET-CGD2**

The second part of Table 1 (row 9 to row 12) records the communication complexities of **det-CGD2**. Unlike **det-CGD1**, we can always obtain the best stepsize matrix \mathbf{D} here if the sketch \mathcal{S}^k is given. The communication complexity can then be obtained in the same way as in the previous case using (33) combined with the number of bits sent per iteration.

B.5 INTERPRETATIONS OF TABLE 1

The communication complexity of the Gradient Descent algorithm (row 13) in the general non-convex setting is equal to $d\lambda_{\max}(\mathbf{L})$, where $\lambda_{\max}(\mathbf{L})$ serves as the smoothness constant of the function. Compared to the GD, **det-CGD1** and **det-CGD2** that use matrix stepsize without compression (row 1 and 9) are better in terms of both iteration and communication complexity. There are some results in the table that need careful analysis and we then present below. In the remainder of the section we will omit the constant multiplier $2(f(x^0) - f^{\text{inf}})/\varepsilon^2$ in communication complexity, as it appears for every setting in the table and thus is redundant for comparison purposes.

B.5.1 COMPARISON OF ROW 5 AND 7

Here we show that the communication complexity given in row 5 is always worse than that of row 7. This can be seen from the following proposition.

Proposition 2. *For any matrix $\mathbf{L} \in \mathbb{S}_{++}^d$, the following inequality holds*

$$\lambda_{\max} \left(\mathbf{L}^{\frac{1}{2}} \text{diag}(\mathbf{L}^{-1}) \mathbf{L}^{\frac{1}{2}} \right) \cdot \det(\mathbf{L})^{\frac{1}{d}} \geq \det(\text{diag}(\mathbf{L}))^{\frac{1}{d}}.$$

Proof. The inequality given in Proposition 2 can be reformulated as

$$\lambda_{\max}(\mathbf{L} \text{diag}(\mathbf{L}^{-1})) \geq \det(\mathbf{L}^{-1} \text{diag}(\mathbf{L}))^{\frac{1}{d}}.$$

We use the notation

$$\mathbf{M}_1 = \mathbf{L} \text{diag}(\mathbf{L}^{-1}), \quad \mathbf{M}_2 = \mathbf{L}^{-1} \text{diag}(\mathbf{L}),$$

and notice that for any $i \in [d]$, we have

$$(\mathbf{M}_1)_{ii} = (\mathbf{L})_{ii} \cdot (\mathbf{L}^{-1})_{ii} = (\mathbf{M}_2)_{ii}.$$

Here the notation $(\mathbf{A})_{ij}$ refers to the entry (i, j) of matrix \mathbf{A} . As a result

$$\lambda_{\max}(\mathbf{M}_1) \geq \left(\prod_{i=1}^d (\mathbf{M}_1)_{ii} \right)^{\frac{1}{d}} = \left(\prod_{i=1}^d (\mathbf{M}_2)_{ii} \right)^{\frac{1}{d}} \geq \det(\mathbf{M}_2)^{\frac{1}{d}},$$

where the first inequality is due to the fact that each diagonal element is upper-bounded by the maximum eigenvalue value, while the second one is obtained using the fact that the product of the diagonal elements is an upper bound of the determinant. \square

From Proposition 2, it immediately follows that the result in row 7 is better than row 5 in terms of both communication and iteration complexity.

B.5.2 COMPARISON OF ROW 6 AND 7

In this section we bring an examples of matrices \mathbf{L} which show that rows 6 and 7 are not comparable in general. Let $d = 2$ and $\mathbf{L} \in \mathbb{S}_{++}^2$. If we pick

$$\mathbf{L} = \begin{pmatrix} 16 & 0 \\ 0 & 1 \end{pmatrix},$$

then

$$\begin{aligned} \det(\text{diag}(\mathbf{L}))^{\frac{1}{d}} &= 4; \\ \lambda_{\max}^{\frac{1}{2}}(\mathbf{L}) \det(\mathbf{L})^{\frac{1}{2d}} &= 8. \end{aligned}$$

However, if we pick

$$\mathbf{L} = \begin{pmatrix} 16 & 3.9 \\ 3.9 & 1 \end{pmatrix},$$

then

$$\begin{aligned} \det(\text{diag}(\mathbf{L}))^{\frac{1}{d}} &= 4; \\ \lambda_{\max}^{\frac{1}{2}}(\mathbf{L}) \det(\mathbf{L})^{\frac{1}{2d}} &\simeq 3.88. \end{aligned}$$

From this example, we can see that the relation between the results in row 6 and 7 may vary depending on the value of \mathbf{L} .

C DISTRIBUTED CASE

C.1 FUTURE WORK

Matrix stepsize gradient methods are still not well studied and require further analysis. Although many important algorithms have been proposed using scalar stepsizes and are known to have good performance, their matrix analogs have yet to be thoroughly examined. The distributed algorithms proposed in Section 5 follow the structure of DCGD by Khirirat et al. (2018). However, other federated learning mechanisms such as MARINA, which has variance reduction (Gorbunov et al., 2021), or EF21 by Richtárik et al. (2021), which has powerful practical performance, should also be explored.

C.2 PROOF OF THEOREM 3

We first present some simple technical lemmas whose proofs are deferred to Appendix D. Let us recall that $\mathbf{D} \in \mathbb{S}_{++}^d$ is the stepsize matrix, $\mathbf{L}, \mathbf{L}_i \in \mathbb{S}_{++}^d$ are the smoothness matrices for f and f_i , respectively.

Lemma 2 (Variance Decomposition). *For any random vector $x \in \mathbb{R}^d$, and any matrix $\mathbf{M} \in \mathbb{S}_{++}^d$, the following identity holds*

$$\mathbb{E} \left[\|x - \mathbb{E}[x]\|_{\mathbf{M}}^2 \right] = \mathbb{E} \left[\|x\|_{\mathbf{M}}^2 \right] - \|\mathbb{E}[x]\|_{\mathbf{M}}^2. \quad (34)$$

Lemma 3. *Assume $\{a_i\}_{i=1}^n$ is a set of independent random vectors in \mathbb{R}^d , which satisfy*

$$\mathbb{E}[a_i] = 0, \quad \forall i \in [n].$$

Then, for any $\mathbf{M} \in \mathbb{S}_{++}^d$, we have

$$\mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n a_i \right\|_{\mathbf{M}}^2 \right] = \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\|a_i\|_{\mathbf{M}}^2 \right]. \quad (35)$$

Lemma 4. For any vector $x \in \mathbb{R}^d$, and sketch matrix $\mathbf{S} \in \mathbb{S}_+^d$ taken from some distribution \mathcal{S} over \mathbb{S}_+^d , which satisfies

$$\mathbb{E}[\mathbf{S}] = \mathbf{I}_d.$$

Then for any matrix $\mathbf{M} \in \mathbb{S}_{++}^d$, we have the following identity holds,

$$\mathbb{E} \left[\|\mathbf{S}x - x\|_{\mathbf{M}}^2 \right] = \|x\|_{\mathbb{E}[\mathbf{S}\mathbf{M}\mathbf{S}] - \mathbf{M}}^2. \quad (36)$$

Lemma 5. If we have a differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, that is \mathbf{L} matrix smooth and lower bounded by f^{inf} , if we assume $\mathbf{L} \in \mathbb{S}_{++}^d$, then the following inequality holds

$$\langle \nabla f(x), \mathbf{L}^{-1} \nabla f(x) \rangle \leq 2(f(x) - f^{\text{inf}}). \quad (37)$$

Let the gradient estimator of our algorithm be defined as

$$g(x) := \frac{1}{n} \sum_{i=1}^n \mathbf{S}_i^k \nabla f_i(x), \quad (38)$$

as a result, **det-CGD1** in the distributed case can then be written as

$$x^{k+1} = x^k - \mathbf{D}g(x^k).$$

Notice that we have

$$\mathbb{E}[g(x^k) | x^k] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathbf{S}_i^k] \nabla f_i(x^k) = \nabla f(x^k). \quad (39)$$

We start with applying the \mathbf{L} -matrix smoothness of f :

$$\begin{aligned} f(x^{k+1}) &\leq f(x^k) + \langle \nabla f(x^k), x^{k+1} - x^k \rangle + \frac{1}{2} \langle \mathbf{L}(x^{k+1} - x^k), x^{k+1} - x^k \rangle \\ &= f(x^k) + \langle \nabla f(x^k), -\mathbf{D}g(x^k) \rangle + \frac{1}{2} \langle \mathbf{L}(-\mathbf{D}g(x^k)), -\mathbf{D}g(x^k) \rangle \\ &= f(x^k) - \langle \nabla f(x^k), \mathbf{D}g(x^k) \rangle + \frac{1}{2} \langle \mathbf{L}\mathbf{D}g(x^k), \mathbf{D}g(x^k) \rangle. \end{aligned}$$

Taking expectation conditioned on x^k , we get

$$\begin{aligned} \mathbb{E}[f(x^{k+1}) | x^k] &\leq f(x^k) - \langle \nabla f(x^k), \mathbf{D}\mathbb{E}[g(x^k) | x^k] \rangle + \frac{1}{2} \mathbb{E}[\langle \mathbf{L}\mathbf{D}g(x^k), \mathbf{D}g(x^k) \rangle | x^k] \\ &\stackrel{(39)}{=} f(x^k) - \langle \nabla f(x^k), \mathbf{D}\nabla f(x^k) \rangle + \frac{1}{2} \mathbb{E}[\langle \mathbf{L}\mathbf{D}g(x^k), \mathbf{D}g(x^k) \rangle | x^k] \\ &= f(x^k) - \|\nabla f(x^k)\|_{\mathbf{D}}^2 + \frac{1}{2} \underbrace{\mathbb{E}[\langle \mathbf{L}\mathbf{D}g(x^k), \mathbf{D}g(x^k) \rangle | x^k]}_{:=T}. \end{aligned} \quad (40)$$

Applying Lemma 2 to the term T we obtain

$$\begin{aligned} T &= \mathbb{E} \left[\|g(x^k)\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2 | x^k \right] \\ &\stackrel{(34)}{=} \mathbb{E} \left[\|g(x^k) - \mathbb{E}[g(x^k) | x^k]\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2 | x^k \right] + \|\mathbb{E}[g(x^k) | x^k]\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2. \end{aligned}$$

From the unbiasedness of the sketches, we have $\mathbb{E}[g(x^k) | x^k] = \nabla f(x^k)$, which yields

$$\begin{aligned} T &= \mathbb{E} \left[\|g(x^k) - \nabla f(x^k)\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2 | x^k \right] + \|\nabla f(x^k)\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2 \\ &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n (\mathbf{S}_i^k \nabla f_i(x^k) - \nabla f_i(x^k)) \right\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2 | x^k \right] + \|\nabla f(x^k)\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2. \end{aligned}$$

Using Lemma 3, we have

$$\begin{aligned} T &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\|\mathbf{S}_i^k \nabla f_i(x^k) - \nabla f_i(x^k)\|_{DLD}^2 \mid x^k \right] + \|\nabla f(x^k)\|_{DLD}^2 \\ &\leq \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\|\mathbf{S}_i^k \nabla f_i(x^k) - \nabla f_i(x^k)\|_{DLD}^2 \mid x^k \right] + \|\nabla f(x^k)\|_D^2, \end{aligned} \quad (41)$$

where the last inequality holds due to the inequality $DLD \preceq D$.

Lemma 6. Let \mathbf{S} be an unbiased ($\mathbb{E}[\mathbf{S}] = \mathbf{I}_d$) sketch drawn randomly from some distribution \mathcal{S} over \mathbb{S}_+^d . The following bound holds for any $x \in \mathbb{R}^d$ and any matrix \mathbf{A} ,

$$\mathbb{E} \left[\|\mathbf{S}x - x\|_{DLD}^2 \right] \leq \lambda_{\max} \left(\mathbf{A}^{\frac{1}{2}} \mathbb{E}[(\mathbf{S} - \mathbf{I}_d) DLD (\mathbf{S} - \mathbf{I}_d)] \mathbf{A}^{\frac{1}{2}} \right) \cdot \|x\|_{\mathbf{A}^{-1}}^2. \quad (42)$$

Plugging (41) into (40) and applying Lemmas 5 and 6 we deduce

$$\begin{aligned} \mathbb{E} [f(x^{k+1}) \mid x^k] &\leq f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_D^2 \\ &\quad + \frac{1}{2n^2} \sum_{i=1}^n \mathbb{E} \left[\|\mathbf{S}_i^k \nabla f_i(x^k) - \nabla f_i(x^k)\|_{DLD}^2 \mid x^k \right]. \\ &\stackrel{(42)}{\leq} f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_D^2 \\ &\quad + \frac{1}{2n^2} \sum_{i=1}^n \lambda_{\max} \left(\mathbb{E} \left[\mathbf{L}_i^{\frac{1}{2}} (\mathbf{S}_i^k - \mathbf{I}_d) DLD (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{L}_i^{\frac{1}{2}} \right] \right) \|\nabla f_i(x^k)\|_{\mathbf{L}_i^{-1}}^2 \\ &\stackrel{(37)}{\leq} f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_D^2 \\ &\quad + \frac{1}{n^2} \sum_{i=1}^n \lambda_{\max} \left(\mathbb{E} \left[\mathbf{L}_i^{\frac{1}{2}} (\mathbf{S}_i^k - \mathbf{I}_d) DLD (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{L}_i^{\frac{1}{2}} \right] \right) (f_i(x^k) - f_i^{\text{inf}}). \end{aligned}$$

Recalling the definition of λ_D , we bound $f(x^{k+1})$ by

$$\begin{aligned} \mathbb{E} [f(x^{k+1}) \mid x^k] &\leq f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_D^2 + \frac{1}{n^2} \sum_{i=1}^n \lambda_D (f_i(x^k) - f_i^{\text{inf}}) \\ &= f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_D^2 + \frac{\lambda_D}{n} \left(\frac{1}{n} \sum_{i=1}^n f_i(x^k) - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}} \right) \\ &= f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_D^2 + \frac{\lambda_D}{n} (f(x^k) - f^{\text{inf}}) + \frac{\lambda_D}{n} \left(f^{\text{inf}} - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}} \right). \end{aligned}$$

Subtracting f^{inf} from both sides, we get

$$\begin{aligned} \mathbb{E} [f(x^{k+1}) - f^{\text{inf}} \mid x^k] &\leq f(x^k) - f^{\text{inf}} - \frac{1}{2} \|\nabla f(x^k)\|_D^2 + \frac{\lambda_D}{n} (f(x^k) - f^{\text{inf}}) \\ &\quad + \frac{\lambda_D}{n} \left(f^{\text{inf}} - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}} \right). \end{aligned}$$

Taking expectation, applying tower property and rearranging terms, we get

$$\begin{aligned} \mathbb{E} [f(x^{k+1}) - f^{\text{inf}}] &\leq \left(1 + \frac{\lambda_D}{n} \right) \mathbb{E} [f(x^k) - f^{\text{inf}}] - \frac{1}{2} \mathbb{E} \left[\|\nabla f(x^k)\|_D^2 \right] \\ &\quad + \frac{\lambda_D}{n} \left(f^{\text{inf}} - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}} \right). \end{aligned} \quad (43)$$

If we denote

$$\delta^k = \mathbb{E} [f(x^k) - f^{\text{inf}}], \quad r^k = \mathbb{E} [\|\nabla f(x^k)\|_{\mathbf{D}}^2], \quad \Delta^{\text{inf}} = f^{\text{inf}} - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}},$$

then (43) becomes

$$\frac{1}{2} r^k \leq \left(1 + \frac{\lambda_{\mathbf{D}}}{n}\right) \delta^k - \delta^{k+1} + \frac{\lambda_{\mathbf{D}} \Delta^{\text{inf}}}{n}. \quad (44)$$

In order to approach the final result, we now follow [Stich \(2019\)](#), [Khaled & Richtárik \(2020\)](#) and define an exponentially decaying weighting sequence $\{w_k\}_{k=-1}^K$, where K is the total number of iterations. We fix $w_{-1} > 0$ and define

$$w_k = \frac{w_{k-1}}{1 + \lambda_{\mathbf{D}}/n}, \quad \text{for all } k \geq 0.$$

By multiplying both sides of the recursion (44) by w_k , we get

$$\frac{1}{2} w_k r^k \leq w_{k-1} \delta^k - w_k \delta^{k+1} + \frac{\lambda_{\mathbf{D}} \Delta^{\text{inf}}}{n} w_k.$$

Summing up the inequalities from $k = 0, \dots, K-1$, we get

$$\frac{1}{2} \sum_{k=0}^{K-1} w_k r^k \leq w_{-1} \delta^0 - w_{K-1} \delta^K + \frac{\lambda_{\mathbf{D}} \Delta^{\text{inf}}}{n} \sum_{k=0}^{K-1} w_k.$$

Define $W_K = \sum_{k=0}^{K-1} w_k$, and divide both sides by W_K , we get

$$\frac{1}{2} \min_{0 \leq k \leq K-1} r^k \leq \frac{1}{2} \frac{\sum_{k=0}^{K-1} w_k r^k}{W_K} \leq \frac{w_{-1} \delta^0}{W_K} + \frac{\lambda_{\mathbf{D}} \Delta^{\text{inf}}}{n}.$$

Notice that from the definition of w_k , we know that the following inequality holds,

$$\frac{w_{-1}}{W_K} \leq \frac{w_{-1}}{K w_{K-1}} = \frac{(1 + \frac{\lambda_{\mathbf{D}}}{n})^K}{K}.$$

As a result, we have

$$\min_{0 \leq k \leq K-1} r^k \leq \frac{2(1 + \frac{\lambda_{\mathbf{D}}}{n})^K}{K} \delta^0 + \frac{2\lambda_{\mathbf{D}} \Delta^{\text{inf}}}{n}.$$

Recalling the definition for r^k and δ^k , we get the following result,

$$\min_{0 \leq k \leq K-1} \mathbb{E} [\|\nabla f(x^k)\|_{\mathbf{D}}^2] \leq \frac{2(1 + \frac{\lambda_{\mathbf{D}}}{n})^K (f(x^0) - f^{\text{inf}})}{K} + \frac{2\lambda_{\mathbf{D}} \Delta^{\text{inf}}}{n}.$$

Finally, we apply determinant normalization and get

$$\min_{0 \leq k \leq K-1} \mathbb{E} [\|\nabla f(x^k)\|_{\mathbf{D}/\det(\mathbf{D})^{1/d}}^2] \leq \frac{2(1 + \frac{\lambda_{\mathbf{D}}}{n})^K (f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} K} + \frac{2\lambda_{\mathbf{D}} \Delta^{\text{inf}}}{\det(\mathbf{D})^{1/d} n}. \quad (45)$$

This concludes the proof.

C.3 CONVEXITY OF THE CONSTRAINTS

Proposition 3. *The set of matrices \mathbf{D} that satisfy (21) is convex.*

Proof. The first inequality in (21) can be reformulated into

$$\mathbf{D} \preceq \mathbf{L}^{-1},$$

which is linear in \mathbf{D} and, therefore, is convex. For the second constraint in (21),

$$\max_i \left\{ \lambda_{\max} \left(\mathbb{E} \left[\mathbf{L}_i^{\frac{1}{2}} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L}_i \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{L}_i^{\frac{1}{2}} \right] \right) \right\} \leq \frac{n}{K}, \quad (46)$$

we can reformulate it into n constraints, one for each client i :

$$\begin{aligned} \lambda_{\max} \left(\mathbb{E} \left[\mathbf{L}_i^{\frac{1}{2}} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{L}_i^{\frac{1}{2}} \right] \right) &\leq \frac{n}{K}, \quad \forall i \in [\ell] \\ \Leftrightarrow \mathbb{E} \left[\mathbf{L}_i^{\frac{1}{2}} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{L}_i^{\frac{1}{2}} \right] &\preceq \frac{n}{K} \mathbf{I}_d, \quad \forall i \in [\ell] \\ \Leftrightarrow \mathbf{L}_i^{\frac{1}{2}} \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \right] \mathbf{L}_i^{\frac{1}{2}} &\preceq \frac{n}{K} \mathbf{I}_d, \quad \forall i \in [\ell] \\ \Leftrightarrow \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \right] &\preceq \frac{n}{K} \mathbf{L}_i^{-1}, \quad \forall i \in [\ell]. \end{aligned}$$

We then look at the individual condition for one client i ,

$$\mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \right] \preceq \frac{n}{K} \mathbf{L}_i^{-1}, \quad (47)$$

that is for any vector $u \in \mathbb{R}^d$, we require

$$\begin{aligned} u^\top \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \right] u &\leq \frac{n}{K} u^\top \mathbf{L}_i^{-1} u, \\ \Leftrightarrow \text{tr} \left(u^\top \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \right] u \right) &\leq \frac{n}{K} \text{tr} \left(u^\top \mathbf{L}_i^{-1} u \right), \\ \Leftrightarrow \mathbb{E} \left[\text{tr} \left(u^\top (\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) u \right) \right] &\leq \text{tr} \left(u^\top \mathbf{L}_i^{-1} u \right), \\ \Leftrightarrow \text{tr} \left(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) u u^\top (\mathbf{S}_i^k - \mathbf{I}_d) \right] \mathbf{D} \mathbf{L}^{\frac{1}{2}} \right) &\leq \text{tr} \left(u^\top \mathbf{L}_i^{-1} u \right). \end{aligned}$$

We now define function $g_u : \mathbb{S}_{++}^d \rightarrow \mathbb{R}$ for every fixed $u \neq 0$,

$$g_u(\mathbf{D}) = \text{tr} \left(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) u u^\top (\mathbf{S}_i^k - \mathbf{I}_d) \right] \mathbf{D} \mathbf{L}^{\frac{1}{2}} \right), \quad (48)$$

notice that $u u^\top$ is a rank-1 matrix that is positive semi-definite, so for every $y \in \mathbb{R}^d$,

$$\left((\mathbf{S}_i^k - \mathbf{I}_d) y \right)^\top u u^\top \left((\mathbf{S}_i^k - \mathbf{I}_d) y \right) \geq 0,$$

which means that $(\mathbf{S}_i^k - \mathbf{I}_d) u u^\top (\mathbf{S}_i^k - \mathbf{I}_d) \in \mathbb{S}_+^d$, and thus $\mathbf{R} := \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) u u^\top (\mathbf{S}_i^k - \mathbf{I}_d) \right] \in \mathbb{S}_+^d$ as well. Using Lemma 1, we know that $g_u(\mathbf{D})$ is a convex function for every $0 \neq u \in \mathbb{R}^d$, thus its sub-level set $\{\mathbf{D} \in \mathbb{S}_{++}^d \mid g_u(\mathbf{D}) \leq \text{tr}(u^\top \mathbf{L}_i^{-1} u)\}$ is a convex set. The intersection of those convex sets corresponding to the individual constraint (47) of client i is convex. Again the intersection of those convex sets for each client i , which corresponds to (46), is still convex.

For the third constraint in (21), we can transform it using similar steps as we obtain (46) into

$$\mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \right] \preceq \frac{n\varepsilon^2}{4\Delta_{\text{inf}}} \det(\mathbf{D})^{1/d} \mathbf{L}_i^{-1}, \quad \forall i. \quad (49)$$

If we look at each individual constraint, we can write in quadratic forms for any $0 \neq u \in \mathbb{R}^d$,

$$u^\top \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L} \mathbf{D} (\mathbf{S}_i^k - \mathbf{I}_d) \right] u \leq \frac{n\varepsilon^2}{4\Delta_{\text{inf}}} \det(\mathbf{D})^{1/d} \cdot u^\top \mathbf{L}_i^{-1} u, \quad \forall u \neq 0.$$

Using the linearity of expectation and the trace operator with the trace trick, we can transform the above condition into,

$$\text{tr} \left(\mathbf{L}^{\frac{1}{2}} \mathbf{D} \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) u u^\top (\mathbf{S}_i^k - \mathbf{I}_d) \right] \mathbf{D} \mathbf{L}^{\frac{1}{2}} \right) \leq \frac{n\varepsilon^2}{4\Delta_{\text{inf}}} \det(\mathbf{D})^{\frac{1}{d}} \text{tr} \left(u^\top \mathbf{L}_i^{-1} u \right) \quad \forall u \neq 0.$$

notice that we have already shown that $\mathbf{R} = \mathbb{E} \left[(\mathbf{S}_i^k - \mathbf{I}_d) u u^\top (\mathbf{S}_i^k - \mathbf{I}_d) \right] \in \mathbb{S}_+^d$. Thus if we apply Lemma 1, we know that the left-hand side of the previous inequality is convex w.r.t. \mathbf{D} . On the other hand we know that $\det(\mathbf{D})^{\frac{1}{d}}$ is a concave function for symmetric positive definite matrices \mathbf{D} . So the set of \mathbf{D} satisfying the constraint here for every $u \in \mathbb{R}^d$ is convex, thus their intersection is convex as well. Which means that the set of \mathbf{D} satisfying the constraint for each client i is convex. Thus the intersection of those convex sets corresponding to different clients, which corresponds to (49), is still convex. Now we know that the set of \mathbf{D} satisfying each of the three constraints in (21) is convex, thus the intersection of them is convex as well. This concludes the proof. \square

C.3.1 PROOF OF COROLLARY 1

For the first term in the RHS of convergence bound (20) under condition (21), we know that

$$2 \left(1 + \frac{\lambda_{\mathbf{D}}}{n}\right)^K \leq 2 \cdot \exp\left(\lambda_{\mathbf{D}} \cdot \frac{K}{n}\right) \leq 2 \cdot \exp(1) \leq 6,$$

thus

$$\begin{aligned} \frac{2(1 + \frac{\lambda_{\mathbf{D}}}{n})^K (f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} K} &\leq \frac{6 (f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} K} \\ &\leq \frac{6 (f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d}} \cdot \frac{\varepsilon^2 \det(\mathbf{D})^{\frac{1}{d}}}{12 (f(x^0) - f^{\text{inf}})} \\ &= \frac{\varepsilon^2}{2}. \end{aligned}$$

While for the second term of RHS in (20), we have

$$\frac{2\lambda_{\mathbf{D}}\Delta^{\text{inf}}}{\det(\mathbf{D})^{1/d} n} \leq \frac{2\Delta^{\text{inf}}}{\det(\mathbf{D})^{1/d} n} \cdot \frac{\varepsilon^2 (\det(\mathbf{D}))^{1/d} n}{4\Delta^{\text{inf}}} \leq \frac{\varepsilon^2}{2}.$$

Thus we know that the left hand side of (20) is upper bounded by

$$\min_{0 \leq k \leq K-1} \mathbb{E} \left[\left\| \nabla f(x^k) \right\|_{\frac{\mathbf{D}}{\det(\mathbf{D})^{1/d}}}^2 \right] \leq \frac{\varepsilon^2}{2} + \frac{\varepsilon^2}{2} = \varepsilon^2.$$

This concludes the proof.

C.4 DISTRIBUTED DET-CGD2

We also extend `det-CGD2` to the distributed case. Consider the method

$$x^{k+1} = x^k - \frac{1}{n} \sum_{i=1}^n \mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k), \quad (50)$$

where $\mathbf{D} \in \mathbb{S}_{++}^d$ is the stepsize matrix, and each \mathbf{T}_i^k is a sequence of sketch matrices drawn randomly from some distribution \mathcal{T} over \mathbb{S}_+^d independent of each other, satisfying

$$\mathbb{E} [\mathbf{T}_i^k] = \mathbf{I}_d. \quad (51)$$

C.4.1 ANALYSIS OF DISTRIBUTED DET-CGD2

In this section, we present the theory for Algorithm 2, which is an analogous to what we have seen for Algorithm 1. We first present the following lemma which is necessary for our analysis.

Lemma 7. *For any sketch \mathbf{T}_i^k of client i drawn randomly from some distribution \mathcal{T} over \mathbb{S}_+^d which satisfies*

$$\mathbb{E} [\mathbf{T}_i^k] = \mathbf{I}_d,$$

the following inequality holds for any $x \in \mathbb{R}^d$ for each client i ,

$$\mathbb{E} \left[\left\| \mathbf{T}_i^k \mathbf{D} x - \mathbf{D} x \right\|_{\mathbf{L}}^2 \right] \leq \lambda_{\max} \left(\mathbf{L}_i^{\frac{1}{2}} \mathbf{D} \mathbb{E} [(\mathbf{T}_i^k - \mathbf{I}_d) \mathbf{L} (\mathbf{T}_i^k - \mathbf{I}_d)] \mathbf{D} \mathbf{L}_i^{\frac{1}{2}} \right) \cdot \|x\|_{\mathbf{L}_i^{-1}}^2. \quad (52)$$

Theorem 4. *Let $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfy Assumption 4 and let f satisfy Assumptions 1 and 2 with a smoothness matrix \mathbf{L} . If the stepsize satisfies,*

$$\mathbf{D} \mathbf{L} \mathbf{D} \preceq \mathbf{D}, \quad (53)$$

then the following convergence bound is true for the iteration of Algorithm 2

$$\min_{0 \leq k \leq K-1} \mathbb{E} \left[\left\| \nabla f(x^k) \right\|_{\frac{\mathbf{D}}{\det(\mathbf{D})^{1/d}}}^2 \right] \leq \frac{2(1 + \frac{\lambda'_{\mathbf{D}}}{n})^K (f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} K} + \frac{2\lambda'_{\mathbf{D}}\Delta^{\text{inf}}}{\det(\mathbf{D})^{1/d} n}, \quad (54)$$

where $\Delta^{\text{inf}} := f^{\text{inf}} - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}}$ and

$$\lambda'_{\mathbf{D}} := \max_i \left\{ \lambda_{\max} \left(\mathbb{E} \left[\mathbf{L}_i^{\frac{1}{2}} \mathbf{D} (\mathbf{T}_i^k - \mathbf{I}_d) \mathbf{L} (\mathbf{T}_i^k - \mathbf{I}_d) \mathbf{D} \mathbf{L}_i^{\frac{1}{2}} \right] \right) \right\}.$$

Proof. We first define function $g(x)$ as follows,

$$g(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k).$$

As a result, Algorithm 2 can be written as

$$x^{k+1} = x^k - g(x^k).$$

Notice that

$$\mathbb{E}[g(x)] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathbf{T}_i^k] \mathbf{D} \nabla f_i(x) = \mathbf{D} \nabla f(x). \quad (55)$$

We then start with the \mathbf{L} matrix smoothness of function f ,

$$\begin{aligned} f(x^{k+1}) &\leq f(x^k) + \langle \nabla f(x^k), x^{k+1} - x^k \rangle + \frac{1}{2} \langle \mathbf{L}(x^{k+1} - x^k), x^{k+1} - x^k \rangle \\ &= f(x^k) + \langle \nabla f(x^k), -g(x^k) \rangle + \frac{1}{2} \langle \mathbf{L}(-g(x^k)), -g(x^k) \rangle \\ &= f(x^k) - \langle \nabla f(x^k), g(x^k) \rangle + \frac{1}{2} \langle \mathbf{L}g(x^k), g(x^k) \rangle. \end{aligned}$$

We then take expectation conditioned on x^k ,

$$\begin{aligned} \mathbb{E}[f(x^{k+1}) | x^k] &\leq f(x^k) - \langle \nabla f(x^k), \mathbb{E}[g(x^k) | x^k] \rangle + \frac{1}{2} \mathbb{E}[\langle \mathbf{L}g(x^k), g(x^k) \rangle | x^k] \\ &= f(x^k) - \langle \nabla f(x^k), \mathbf{D} \nabla f(x^k) \rangle + \frac{1}{2} \underbrace{\mathbb{E}[\langle \mathbf{L}g(x^k), g(x^k) \rangle | x^k]}_{:=T}. \end{aligned} \quad (56)$$

Lemma 2 yields

$$\begin{aligned} T &= \mathbb{E}[\|g(x^k)\|_{\mathbf{L}}^2 | x^k] \\ &\stackrel{(34)}{=} \mathbb{E}[\|g(x^k) - \mathbb{E}[g(x^k) | x^k]\|_{\mathbf{L}}^2 | x^k] + \|\mathbb{E}[g(x^k) | x^k]\|_{\mathbf{L}}^2. \end{aligned}$$

From (55) we deduce

$$\begin{aligned} T &= \mathbb{E}[\|g(x^k) - \mathbf{D} \nabla f(x^k)\|_{\mathbf{L}}^2 | x^k] + \|\mathbf{D} \nabla f(x^k)\|_{\mathbf{L}}^2 \\ &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k) - \frac{\mathbf{D}}{n} \sum_{i=1}^n \nabla f_i(x^k) \right\|_{\mathbf{L}}^2 | x^k \right] + \|\nabla f(x^k)\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2 \\ &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n (\mathbf{T}_i^k \mathbf{D} - \mathbf{D}) \nabla f_i(x^k) \right\|_{\mathbf{L}}^2 | x^k \right] + \|\nabla f(x^k)\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2. \end{aligned}$$

Recalling Lemma 3 we obtain

$$\begin{aligned} T &\stackrel{(35)}{=} \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\|\mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k) - \mathbf{D} \nabla f_i(x^k)\|_{\mathbf{L}}^2 | x^k \right] + \|\nabla f(x^k)\|_{\mathbf{D}\mathbf{L}\mathbf{D}}^2 \\ &\stackrel{(53)}{\leq} \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\|\mathbf{T}_i^k \mathbf{D} \nabla f_i(x^k) - \mathbf{D} \nabla f_i(x^k)\|_{\mathbf{L}}^2 | x^k \right] + \|\nabla f(x^k)\|_{\mathbf{D}}^2. \end{aligned}$$

By applying Lemma 7, we get

$$\begin{aligned} T &\leq \frac{1}{n^2} \sum_{i=1}^n \lambda_{\max} \left(\mathbf{L}_i^{\frac{1}{2}} \mathbf{D} \mathbb{E}[(\mathbf{T}_i^k - \mathbf{I}_d) \mathbf{L} (\mathbf{T}_i^k - \mathbf{I}_d)] \mathbf{D} \mathbf{L}_i^{\frac{1}{2}} \right) \|\nabla f_i(x^k)\|_{\mathbf{L}_i^{-1}}^2 + \|\nabla f(x^k)\|_{\mathbf{D}}^2 \\ &\stackrel{(37)}{\leq} \lambda_{\mathbf{D}}' \cdot \frac{2}{n} \left(f(x^k) - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}} \right) + \|\nabla f(x^k)\|_{\mathbf{D}}^2. \end{aligned}$$

Then we plug the upper bound of T back into (56), we get

$$\begin{aligned} & \mathbb{E} [f(x^{k+1}) \mid x^k] \\ & \leq f(x^k) - \frac{1}{2} \|\nabla f(x^k)\|_{\mathbf{D}}^2 + \frac{\lambda'_{\mathbf{D}}}{n} (f(x^k) - f^{\text{inf}}) + \frac{\lambda'_{\mathbf{D}}}{n} (f^{\text{inf}} - \frac{1}{n} \sum_{i=1}^n f_i^{\text{inf}}). \end{aligned}$$

Taking expectation, subtracting f^{inf} from both sides, and using tower property, we get

$$\begin{aligned} & \mathbb{E} [f(x^{k+1}) - f^{\text{inf}}] \\ & \leq \mathbb{E} [f(x^k) - f^{\text{inf}}] - \frac{1}{2} \mathbb{E} [\|\nabla f(x^k)\|_{\mathbf{D}}^2] + \frac{\lambda'_{\mathbf{D}}}{n} \mathbb{E} [f(x^k) - f^{\text{inf}}] + \frac{\lambda'_{\mathbf{D}}}{n} \Delta^{\text{inf}}. \end{aligned}$$

Then following similar steps as in the proof of Theorem 3, we are able to get

$$\min_{0 \leq k \leq K-1} \mathbb{E} \left[\|\nabla f(x^k)\|_{\frac{\mathbf{D}}{\det(\mathbf{D})^{1/d}}}^2 \right] \leq \frac{2(1 + \frac{\lambda'_{\mathbf{D}}}{n})^K (f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} K} + \frac{2\lambda'_{\mathbf{D}} \Delta^{\text{inf}}}{\det(\mathbf{D})^{1/d} n}.$$

This concludes the proof. \square

Similar to Algorithm 1, we can choose the parameters of the algorithm to avoid the exponential blow-up in convergence bound (54). The following corollary sums up the convergence conditions for Algorithm 2.

Corollary 3. *We reach an ε^2 -stationary point in (54) if the following conditions are satisfied:*

$$DLD \preceq \mathbf{D}, \quad \lambda'_{\mathbf{D}} \leq \min \left\{ \frac{n}{K}, \frac{n\varepsilon^2}{4\Delta^{\text{inf}} \det(\mathbf{D})^{1/d}} \right\}, \quad K \geq \frac{12(f(x^0) - f^{\text{inf}})}{\det(\mathbf{D})^{1/d} \varepsilon^2}. \quad (57)$$

The proof of this corollary is exactly the same as for Corollary 1.

C.4.2 OPTIMAL STEPSIZE

In order to minimize the iteration complexity for Algorithm 2, the following optimization problem needs to be solved

$$\begin{aligned} & \min \quad \log \det(\mathbf{D}^{-1}) \\ & \text{subject to} \quad \mathbf{D} \text{ satisfies (57)} \end{aligned}$$

Following similar techniques in the proof of Proposition 3, we are able to prove that the above optimization problem is still a convex optimization problem. One simple way to find stepsize matrices is to follow the scheme suggested for solving (13). That is we first fix $\mathbf{W} \in \mathbb{S}_{++}^d$ and we find the optimal $0 < \gamma \in \mathbb{R}$, such that $\mathbf{D} = \gamma \mathbf{W}$ satisfies (57).

C.5 DCGD WITH CONSTANT STEPSIZE

In this section we describe the convergence result for DCGD from Khaled & Richtárik (2020). We assume that the component functions f_i satisfy Assumption 4 with $\mathbf{L}_i = L_i \mathbf{I}_d$ and f satisfies Assumption 1 and 2 with $\mathbf{L} = L \mathbf{I}_d$. Khaled & Richtárik (2020) proposed a unified analysis for non-convex optimization algorithms based on a generic upper bound on the second moment of the gradient estimator $g(x^k)$:

$$\mathbb{E} \left[\|g(x^k)\|^2 \right] \leq 2A (f(x^k) - f^{\text{inf}}) + B \|\nabla f(x^k)\|^2 + C, \quad (58)$$

In our case the gradient estimator is defined as follows

$$g_{\text{DCGD}}(x^k) = \frac{1}{n} \sum_{i=1}^n \mathbf{S}_i^k \nabla f_i(x^k). \quad (59)$$

Here each \mathbf{S}_i^k is the sketch matrix on the i -th client at the k -th iteration. One may check that g_{DCGD} satisfies (58) with the following constants:

$$A = \frac{\omega L_{\max}}{n}, \quad B = 1, \quad C = \frac{2\omega L_{\max}}{n} \Delta^{\text{inf}}. \quad (60)$$

The constant L_{\max} is defined as the maximum of all L_i and $\omega = \lambda_{\max} \left(\mathbb{E} \left[(\mathbf{S}_i^k)^\top \mathbf{S}_i^k \right] \right) - 1$. Applying Corollary 1 from Khaled & Richtárik (2020), we deduce the following. If

$$\gamma \leq \min \left\{ \frac{1}{L}, \frac{\sqrt{n}}{\sqrt{\omega L L_{\max} K}}, \frac{n \varepsilon^2}{4 L L_{\max} \omega \Delta^{\text{inf}}} \right\} \text{ and } \gamma K \geq \frac{12 (f(x^0) - f^{\text{inf}})}{\varepsilon^2}, \quad (61)$$

then

$$\min_{k=0, \dots, K-1} \mathbb{E} \left[\|\nabla f(x^k)\|^2 \right] \leq \varepsilon^2. \quad (62)$$

D PROOFS OF TECHNICAL LEMMAS

D.1 PROOF OF LEMMA 1

Let us pick any two matrices $\mathbf{D}_1, \mathbf{D}_2 \in \mathbb{S}_{++}^d$, scalar α satisfying $0 \leq \alpha \leq 1$ and show that the following inequality holds regardless of the choice of \mathbf{R} ,

$$f(\alpha \mathbf{D}_1 + (1 - \alpha) \mathbf{D}_2) \leq \alpha f(\mathbf{D}_1) + (1 - \alpha) f(\mathbf{D}_2). \quad (63)$$

For the LHS, we have

$$\begin{aligned} & f(\alpha \mathbf{D}_1 + (1 - \alpha) \mathbf{D}_2) \\ &= \text{tr}(\mathbf{L}^{\frac{1}{2}} (\alpha \mathbf{D}_1 + (1 - \alpha) \mathbf{D}_2) \mathbf{R} (\alpha \mathbf{D}_1 + (1 - \alpha) \mathbf{D}_2) \mathbf{L}^{\frac{1}{2}}) \\ &= \alpha^2 \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) + (1 - \alpha)^2 \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}) \\ &\quad + \alpha(1 - \alpha) \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}) + \alpha(1 - \alpha) \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}). \end{aligned}$$

and for the RHS, we have

$$\alpha f(\mathbf{D}_1) + (1 - \alpha) f(\mathbf{D}_2) = \alpha \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) + (1 - \alpha) \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}).$$

Thus (63) can be simplified to the following inequality after rearranging terms

$$\begin{aligned} & \alpha(1 - \alpha) \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}) + \alpha(1 - \alpha) \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) \\ & \leq \alpha(1 - \alpha) \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) + \alpha(1 - \alpha) \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}). \end{aligned}$$

This is equivalent to

$$\text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) + \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}) - \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}) - \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) \geq 0.$$

To show that the above inequality holds, we do the following transformation for the LHS

$$\begin{aligned} & \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) + \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}) - \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} \mathbf{D}_2 \mathbf{L}^{\frac{1}{2}}) - \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} \mathbf{D}_1 \mathbf{L}^{\frac{1}{2}}) \\ &= \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_1 \mathbf{R} (\mathbf{D}_1 - \mathbf{D}_2) \mathbf{L}^{\frac{1}{2}}) + \text{tr}(\mathbf{L}^{\frac{1}{2}} \mathbf{D}_2 \mathbf{R} (\mathbf{D}_2 - \mathbf{D}_1) \mathbf{L}^{\frac{1}{2}}) \\ &= \text{tr}(\mathbf{L}^{\frac{1}{2}} (\mathbf{D}_1 - \mathbf{D}_2) \mathbf{R} (\mathbf{D}_1 - \mathbf{D}_2) \mathbf{L}^{\frac{1}{2}}). \end{aligned}$$

Since $\mathbf{R} \in \mathbb{S}_+^d$ and $\mathbf{D}_1 - \mathbf{D}_2, \mathbf{L}$ are symmetric, for any vector $u \in \mathbb{R}^d$

$$u^\top \mathbf{L}^{\frac{1}{2}} (\mathbf{D}_1 - \mathbf{D}_2) \mathbf{R} (\mathbf{D}_1 - \mathbf{D}_2) \mathbf{L}^{\frac{1}{2}} u = \left((\mathbf{D}_1 - \mathbf{D}_2) \mathbf{L}^{\frac{1}{2}} u \right)^\top \mathbf{R} \left((\mathbf{D}_1 - \mathbf{D}_2) \mathbf{L}^{\frac{1}{2}} u \right) \geq 0. \quad (64)$$

Thus, $\mathbf{L}^{\frac{1}{2}} (\mathbf{D}_1 - \mathbf{D}_2) \mathbf{R} (\mathbf{D}_1 - \mathbf{D}_2) \mathbf{L}^{\frac{1}{2}} \in \mathbb{S}_+^d$, which yields the positivity of its trace. Therefore, (63) holds, thus $f(D)$ is a convex function. This concludes the proof.

D.2 PROOF OF LEMMA 2

We have

$$\begin{aligned}
\mathbb{E} \left[\|x - \mathbb{E}[x]\|_{\mathbf{M}}^2 \right] &= \mathbb{E} [\langle x - \mathbb{E}[x], \mathbf{M}(x - \mathbb{E}[x]) \rangle] \\
&= \mathbb{E} \left[(x - \mathbb{E}[x])^\top \mathbf{M}(x - \mathbb{E}[x]) \right] \\
&= \mathbb{E} \left[x^\top \mathbf{M}x - \mathbb{E}[x]^\top \mathbf{M}x - x^\top \mathbf{M}\mathbb{E}[x] + \mathbb{E}[x]^\top \mathbf{M}\mathbb{E}[x] \right] \\
&= \mathbb{E} [x^\top \mathbf{M}x] - 2\mathbb{E}[x]^\top \mathbf{M}\mathbb{E}[x] + \mathbb{E}[x]^\top \mathbf{M}\mathbb{E}[x] \\
&= \mathbb{E} [x^\top \mathbf{M}x] - \mathbb{E}[x]^\top \mathbf{M}\mathbb{E}[x] \\
&= \mathbb{E} \left[\|x\|_{\mathbf{M}}^2 \right] - \|\mathbb{E}[x]\|_{\mathbf{M}}^2,
\end{aligned}$$

which concludes the proof.

D.3 PROOF OF LEMMA 3

Proof. We have

$$\begin{aligned}
\mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n a_i \right\|_{\mathbf{M}}^2 \right] &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} [\langle a_i, \mathbf{M}a_i \rangle] + \frac{1}{n^2} \sum_{i \neq j} \mathbb{E} [\langle a_i, \mathbf{M}a_j \rangle] \\
&= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} [\|a_i\|_{\mathbf{M}}^2] + \frac{1}{n} \sum_{i \neq j} \langle \mathbb{E}[a_i], \mathbf{M}\mathbb{E}[a_j] \rangle \\
&= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} [\|a_i\|_{\mathbf{M}}^2].
\end{aligned}$$

This concludes the proof. \square

D.4 PROOF OF LEMMA 4

Notice that

$$\mathbb{E}[\mathbf{S}x] = \mathbb{E}[\mathbf{S}]x = x.$$

We start with variance decomposition in the matrix norm,

$$\begin{aligned}
\mathbb{E} \left[\|\mathbf{S}x - x\|_{\mathbf{M}}^2 \right] &\stackrel{(34)}{=} \mathbb{E} \left[\|\mathbf{S}x\|_{\mathbf{M}}^2 \right] - \|x\|_{\mathbf{M}}^2 \\
&= \mathbb{E} [\langle \mathbf{S}x, \mathbf{M}\mathbf{S}x \rangle] - \langle x, \mathbf{M}x \rangle \\
&= \langle x, \mathbb{E}[\mathbf{S}\mathbf{M}\mathbf{S}]x \rangle - \langle x, \mathbf{M}x \rangle \\
&= \langle x, (\mathbb{E}[\mathbf{S}\mathbf{M}\mathbf{S}] - \mathbf{M})x \rangle \\
&= \|x\|_{\mathbb{E}[\mathbf{S}\mathbf{M}\mathbf{S}] - \mathbf{M}}^2.
\end{aligned}$$

This concludes the proof.

D.5 PROOF OF LEMMA 5

We follow the definition of \mathbf{L} matrix smoothness of function f , that for any $x^+, x \in \mathbb{R}^d$, we have

$$f(x^+) \leq f(x) + \langle \nabla f(x), x^+ - x \rangle + \frac{1}{2} \langle x^+ - x, \mathbf{L}(x^+ - x) \rangle.$$

We plug in $x^+ = x - \mathbf{L}^{-1}\nabla f(x)$, and get

$$f^{\text{inf}} \leq f(x^+) \leq f(x) - \langle \nabla f(x), \mathbf{L}^{-1}\nabla f(x) \rangle + \frac{1}{2} \langle \nabla f(x), \mathbf{L}^{-1}\nabla f(x) \rangle.$$

Rearranging terms we get

$$\|\nabla f(x)\|_{\mathbf{L}^{-1}}^2 \leq 2(f(x) - f^{\text{inf}}), \tag{65}$$

which completes the proof.

D.6 PROOF OF LEMMA 6

$$\begin{aligned}
\mathbb{E} \left[\|S - x\|_{DLD}^2 \right] &= \mathbb{E} [\langle (S - I_d)x, DLD(S - I_d)x \rangle] \\
&= \mathbb{E} [x^\top (S - I_d)DLD(S - I_d)x] \\
&= x^\top \mathbb{E} [(S - I_d)DLD(S - I_d)] x \\
&= x^\top \mathbf{A}^{-\frac{1}{2}} \left(\mathbf{A}^{\frac{1}{2}} \mathbb{E} [(S - I_d)DLD(S - I_d)] \mathbf{A}^{\frac{1}{2}} \right) \mathbf{A}^{-\frac{1}{2}} x \\
&\leq \lambda_{\max} \left(\mathbf{A}^{\frac{1}{2}} \mathbb{E} [(S - I_d)DLD(S - I_d)] \mathbf{A}^{\frac{1}{2}} \right) \left\| \mathbf{A}^{-\frac{1}{2}} x \right\|^2 \\
&= \lambda_{\max} \left(\mathbf{A}^{\frac{1}{2}} \mathbb{E} [(S - I_d)DLD(S - I_d)] \mathbf{A}^{\frac{1}{2}} \right) \|x\|_{\mathbf{A}^{-1}}^2.
\end{aligned}$$

This completes the proof.

D.7 PROOF OF LEMMA 7

$$\begin{aligned}
\mathbb{E} \left[\left\| \mathbf{T}_i^k \mathbf{D}x - \mathbf{D}x \right\|_{\mathbf{L}}^2 \right] &= \mathbb{E} [\langle (\mathbf{T}_i^k - I_d)\mathbf{D}x, \mathbf{L}(\mathbf{T}_i^k - I_d)\mathbf{D}x \rangle] \\
&= \mathbb{E} [x^\top \mathbf{D}(\mathbf{T}_i^k - I_d)\mathbf{L}(\mathbf{T}_i^k - I_d)\mathbf{D}x] \\
&= x^\top \mathbf{D} \mathbb{E} [(\mathbf{T}_i^k - I_d)\mathbf{L}(\mathbf{T}_i^k - I_d)] \mathbf{D}x \\
&= x^\top \mathbf{L}_i^{-\frac{1}{2}} \left(\mathbf{L}_i^{\frac{1}{2}} \mathbf{D} \mathbb{E} [(\mathbf{T}_i^k - I_d)\mathbf{L}(\mathbf{T}_i^k - I_d)] \mathbf{D} \mathbf{L}_i^{\frac{1}{2}} \right) \mathbf{L}_i^{-\frac{1}{2}} x \\
&\leq \lambda_{\max} \left(\mathbf{L}_i^{\frac{1}{2}} \mathbf{D} \mathbb{E} [(\mathbf{T}_i^k - I_d)\mathbf{L}(\mathbf{T}_i^k - I_d)] \mathbf{D} \mathbf{L}_i^{\frac{1}{2}} \right) \left\| \mathbf{L}_i^{-\frac{1}{2}} x \right\|^2 \\
&= \lambda_{\max} \left(\mathbf{L}_i^{\frac{1}{2}} \mathbf{D} \mathbb{E} [(\mathbf{T}_i^k - I_d)\mathbf{L}(\mathbf{T}_i^k - I_d)] \mathbf{D} \mathbf{L}_i^{\frac{1}{2}} \right) \|x\|_{\mathbf{L}_i^{-1}}^2.
\end{aligned}$$

This completes the proof.

E EXPERIMENTS

In this section, we describe the settings and results of numerical experiments to demonstrate the effectiveness of our method. We perform several experiments under single node case and distributed case. The code is available at https://anonymous.4open.science/r/detCGD_Code-A87D/.

E.1 SINGLE NODE CASE

For the single node case, we study the logistic regression problem with a non-convex regularizer. The objective is given as

$$f(x) = \frac{1}{n} \sum_{i=1}^n \log \left(1 + e^{-b_i \cdot \langle a_i, x \rangle} \right) + \lambda \cdot \sum_{j=1}^d \frac{x_j^2}{1 + x_j^2},$$

where $x \in \mathbb{R}^d$ is the model, $(a_i, b_i) \in \mathbb{R}^d \times \{-1, +1\}$ is one data point in the dataset whose size is n . The constant $\lambda > 0$ is a tunable hyperparameter associated with the regularizer. We conduct numerical experiments using several datasets from the LibSVM repository (Chang & Lin, 2011). We estimate the smoothness matrix of function f here as

$$\mathbf{L} = \frac{1}{n} \sum_{i=1}^n \frac{a_i a_i^\top}{4} + 2\lambda \cdot \mathbf{I}_d.$$

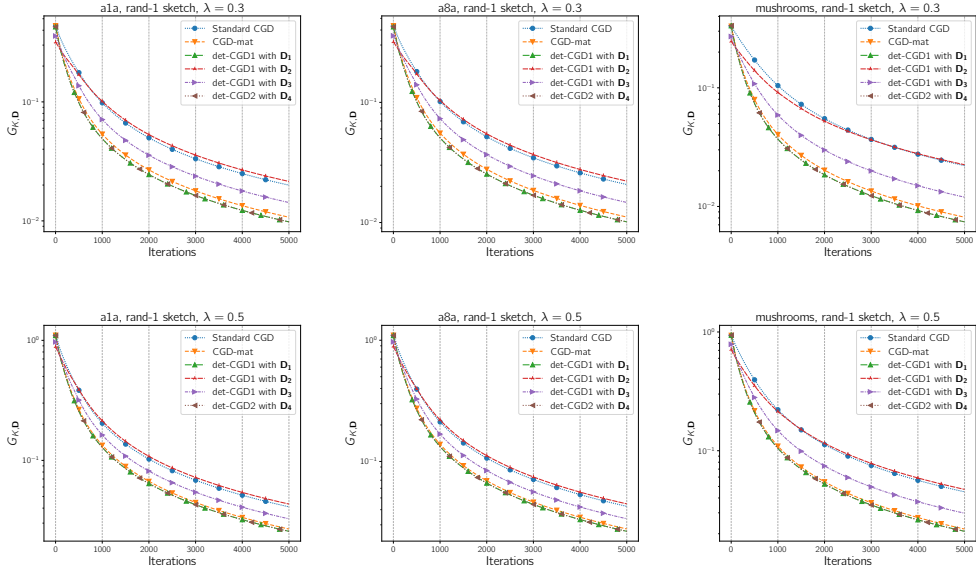


Figure 2: Comparison, in terms of iteration complexity, of standard CGD, CGD-mat, $\det\text{-CGD1}$ with $D_1 = \gamma_1 \cdot \text{diag}^{-1}(L)$, $\det\text{-CGD1}$ with $D_2 = \gamma_2 \cdot L^{-1}$, $\det\text{-CGD1}$ with $D_3 = \gamma_3 \cdot L^{-1/2}$ and $\det\text{-CGD2}$ with $D_4 = \gamma_4 \cdot \text{diag}^{-1}(L)$, where $\gamma_1, \gamma_2, \gamma_3$ are the optimal scaling factors for $\det\text{-CGD1}$ in that case, D_4 is the optimal matrix stepsize for $\det\text{-CGD2}$. Rand-1 sketch is used in all the methods through out the experiments. The x -axis is the number of iterations performed. The notation $G_{K,D}$ in the y -axis is defined in (66).

E.1.1 COMPARISON TO CGD WITH SCALAR STEPSIZE, SCALAR SMOOTHNESS CONSTANT

The purpose of the first experiment is to show that by using matrix stepsize, $\det\text{-CGD1}$ and $\det\text{-CGD2}$ will have better iteration and communication complexities compared to standard CGD. We run a CGD with scalar stepsize γ and a scalar smoothness constant $L = \lambda_{\max}(L)$ and CGD with scalar stepsize $\gamma \cdot I_d$ and smoothness matrix L . We use standard CGD to refer to the CGD with scalar stepsize, scalar smoothness constant, and CGD-mat to refer to CGD with scalar stepsize, smoothness matrix in Figure 2, 3 and 4. The notation $G_{K,D}$ appears in the label of y axis is defined as

$$G_{K,D} := \frac{1}{K} \left(\sum_{k=0}^{K-1} \|\nabla f(x^k)\|_{\frac{D}{\det(D)^{1/d}}}^2 \right), \quad (66)$$

it is the average matrix norm of the gradient of f over the first $K - 1$ iterations in log scale. The weight matrix here has determinant 1, and thus it is comparable to the standard Euclidean norm. The result is meaningful in this sense.

The result presented in Figure 2 and Figure 3 suggest that compared to the standard CGD (Khairat et al., 2018), CGD-mat performs better in terms of both iteration complexity and communication complexity. Furthermore, $\det\text{-CGD1}$ and $\det\text{-CGD2}$ with the best diagonal matrix stepsizes outperform both CGD and CGD-mat which confirms our theory. The scaling factors $\gamma_1, \gamma_2, \gamma_3$ here for $\det\text{-CGD1}$ are determined using Theorem 2 with $\ell = 1$. The matrix stepsize for $\det\text{-CGD2}$ is determined through (12). $\det\text{-CGD1}$ and $\det\text{-CGD2}$ with diagonal matrix stepsizes perform very similarly in the experiment, this is expected since we are using rand-1 sketch, which means that the stepsize matrix and the sketch matrix are commutable since they are both diagonal. We also notice that $\det\text{-CGD1}$ with $D_2 = \gamma_2 \cdot L^{-1}$ is always worse than $D_4 = \gamma_4 \cdot \text{diag}^{-1}(L)$, this is also expected since we mentioned in Appendix B.5.1 that the result row 5 (corresponding to D_2) in Table 1 is always worse than row 7 (corresponding to D_4).




Figure 3: Comparison, in terms of communication complexity, of standard CGD, CGD-mat, **det-CGD1** with $\mathbf{D}_1 = \gamma_1 \cdot \text{diag}^{-1}(\mathbf{L})$, **det-CGD1** with $\mathbf{D}_2 = \gamma_2 \cdot \mathbf{L}^{-1}$, **det-CGD1** with $\mathbf{D}_3 = \gamma_3 \cdot \mathbf{L}^{-1/2}$ and **det-CGD2** with $\mathbf{D}_4 = \gamma_4 \cdot \text{diag}^{-1}(\mathbf{L})$, where $\gamma_1, \gamma_2, \gamma_3$ are the optimal scaling factors for **det-CGD1** in that case, \mathbf{D}_4 is the optimal matrix stepsize for **det-CGD2**. Rand-1 sketch is used in all the methods through out the experiments. The x -axis is the number of bits transmitted. The notation $G_{K,D}$ in the y -axis is defined in (66).

E.1.2 COMPARISON OF THE TWO ALGORITHMS UNDER THE SAME STEPSIZE

The purpose of the second experiment is to compare the performance of **det-CGD1** and **det-CGD2** in terms of iteration complexity and communication complexity. We know the conditions for **det-CGD1** and **det-CGD2** to converge are given by (7) and (8) respectively. As a result, we are able to obtain the optimal matrix stepsize for **det-CGD2** if we are using rand- τ sparsification. It is given by

$$\mathbf{D}_2^* = \frac{\tau}{d} \left(\frac{d-\tau}{d-1} \text{diag}(\mathbf{L}) + \frac{\tau-1}{d-1} \mathbf{L} \right)^{-1},$$

according to (11). The definition of $G_{K,D}$ is given in (66). Parameter τ here for random sparsification is set to be an the integer part $\{\frac{d}{4}, \frac{d}{2}, \frac{3d}{4}\}$, where d is the dimension of the model.

It can be observed from the result presented in Figure 4, that in almost all cases in this experiment, **2** with $\mathbf{D} = \mathbf{D}_2^*$ outperforms the other methods. Compared to standard CGD and CGD with matrix stepsize, **det-CGD1** and **det-CGD2** are always better. This provides numerical evidence in support of our theory. In this case, the stepsize matrix is not diagonal for **det-CGD1** and **det-CGD2**, so we do not expect them to perform similarly. Notice that in dataset `phishing`, the four algorithms behave very similarly, this is because the smoothness matrix \mathbf{L} here has a concentrated spectrum.

E.2 DISTRIBUTED CASE

For the distributed case, we again use the logistic regression problem with a non-convex regularizer as our experiment setting. The objective is given similarly as

$$f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x); \quad f_i(x) = \frac{1}{m_i} \sum_{j=1}^{m_i} \log \left(1 + e^{-b_{i,j} \cdot \langle a_{i,j}, x \rangle} \right) + \lambda \cdot \sum_{t=1}^d \frac{x_t^2}{1 + x_t^2},$$

where $x \in \mathbb{R}^d$ is the model, $(a_{i,j}, b_{i,j}) \in \mathbb{R}^d \times \{-1, +1\}$ is one data point in the dataset of client i whose size is m_i . $\lambda > 0$ is a constant associated with the regularizer. For each dataset used in the distributed setting, we randomly reshuffled the dataset before splitting it equally to each client. We

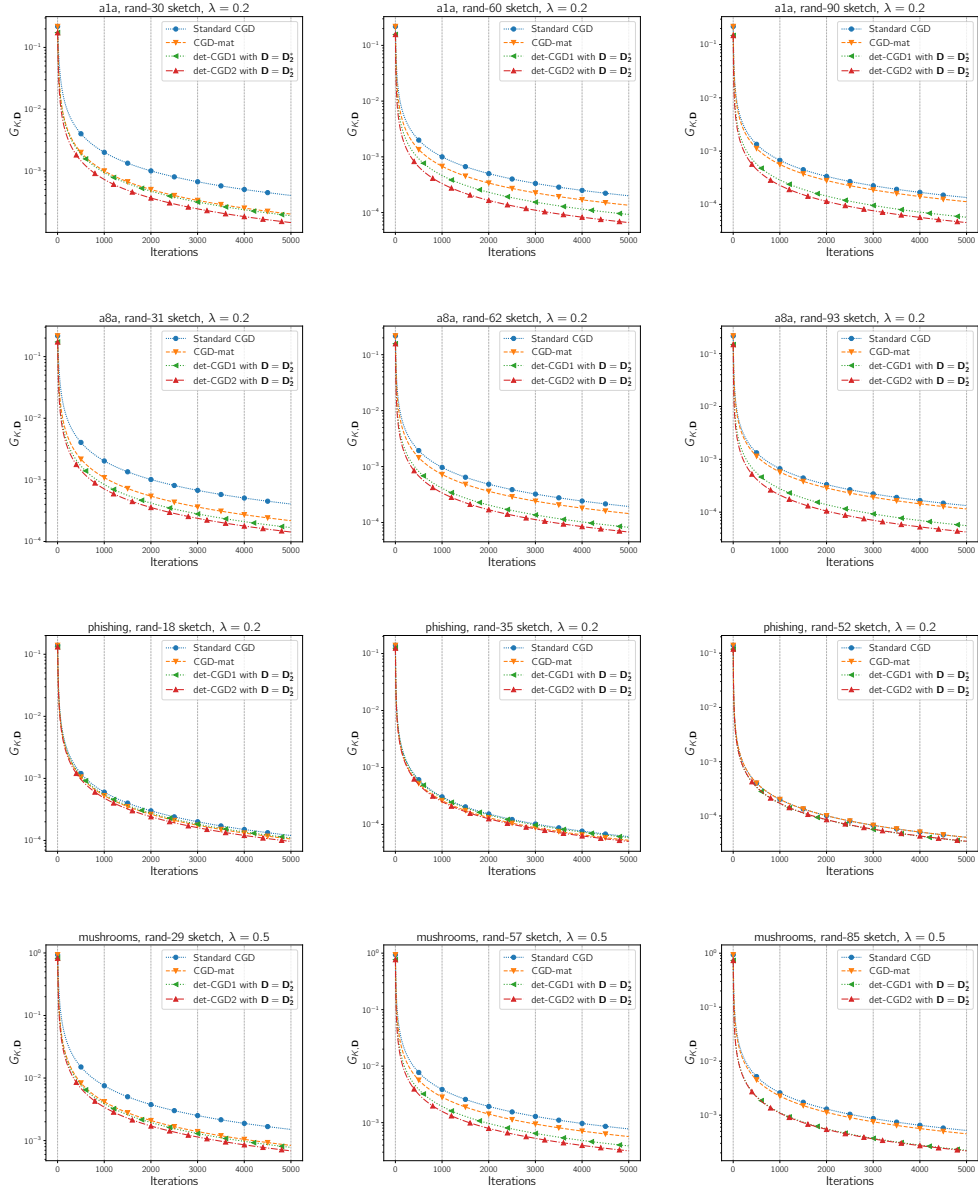


Figure 4: Comparison of standard CGD, CGD-mat \det -CGD1 with stepsize $D = D_2^*$ and \det -CGD2 with stepsize $D = D_2^*$, where D_2^* is the optimal stepsize matrix for \det -CGD2 and the optimal diagonal stepsize matrix for \det -CGD1. Rand- τ sketch is used in all the algorithms throughout the experiments. The notation $G_{K,D}$ in the y -axis is defined in (66).

estimate the smoothness matrices of function f and each individual function f_i here as

$$\begin{aligned} \mathbf{L}_i &= \frac{1}{m_i} \sum_{i=1}^{m_i} \frac{a_i a_i^\top}{4} + 2\lambda \cdot \mathbf{I}_d; \\ \mathbf{L} &= \frac{1}{n} \sum_{i=1}^n \mathbf{L}_i. \end{aligned}$$

The value of Δ^{inf} here is determined in the following way, we first perform gradient descent on f and record the minimum value in the entire run, f^{inf} , as the estimate of its global minimum, then we do the same procedure for each f_i to obtain the estimate of its global minimum f_i^{inf} . After that we estimate Δ^{inf} using its definition.

E.2.1 COMPARISON TO STANDARD DCGD IN THE DISTRIBUTED CASE

To ease the reading of this section we use **D-det-CGD1** (resp. **D-det-CGD2**) to refer to Algorithm 1 (resp. Algorithm 2). This experiment is designed to show that **D-det-CGD1** and **D-det-CGD2** will have better iteration and communication complexity compared to standard DCGD (Khairat et al., 2018) and DCGD with scalar stepsize, smoothness matrix. We will use the standard DCGD here to refer to DCGD with a scalar stepsize and a scalar smoothness constant, and DCGD-mat to refer to the DCGD with a scalar stepsize with smoothness. The Rand-1 sparsifier is used in all the algorithms throughout the experiment. The error level is fixed as $\varepsilon^2 = 0.0001$, the conditions for the standard DCGD to converge can be deduced using Proposition 4 in Khaled & Richtárik (2020), we use the largest possible scalar stepsize here for standard DCGD. The optimal scalar stepsize for DCGD-mat, optimal diagonal matrix stepsize D_1 for **D-det-CGD1** and D_2 for **D-det-CGD2** can be determined using Corollary 1.

From the result of Figure 5 and Figure 6, we are able to see that both **D-det-CGD1** and **D-det-CGD2** outperform standard DCGD and DCGD-mat in terms of iteration complexity and communication complexity, which confirms our theory. Notice that **D-det-CGD1**, **D-det-CGD2** are expected to perform very similarly because the stepsize matrix and sketches are diagonal which means that they are commutable. We also plot the corresponding standard Euclidean norm of iterates of **D-det-CGD1** and **D-det-CGD2** in Figure 7, the E_K here appears in the y -axis is defined as,

$$E_K := \frac{1}{K} \sum_{k=0}^{K-1} \|\nabla f(x^k)\|^2. \quad (67)$$

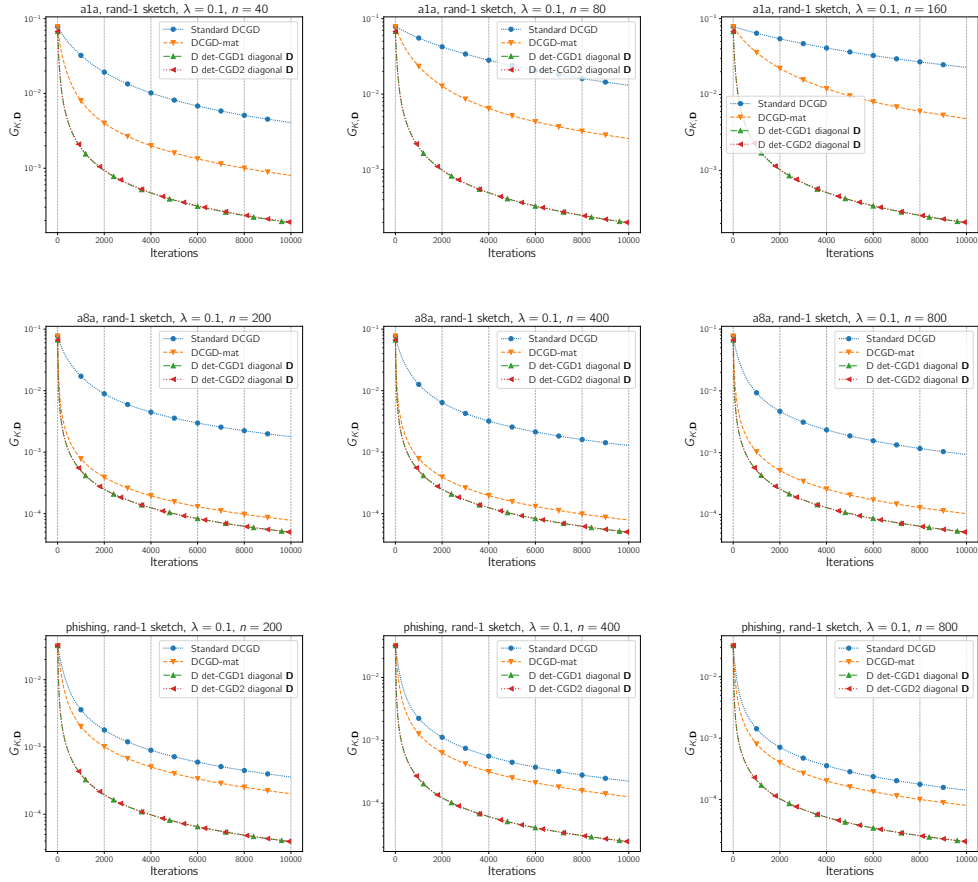


Figure 5: Comparison, in terms of iteration complexity, of standard DCGD, DCGD-mat, **D-det-CGD1** with matrix stepsize D_1 and **D-det-CGD2** with matrix stepsize D_2 , where D_1, D_2 are the optimal diagonal matrix stepsizes for **D-det-CGD1** and **D-det-CGD2** respectively. Rand-1 sketch is used in all the algorithms throughout the experiment. The x -axis is the number of iterations performed. The notation $G_{K,D}$ in the y -axis is defined in (66).

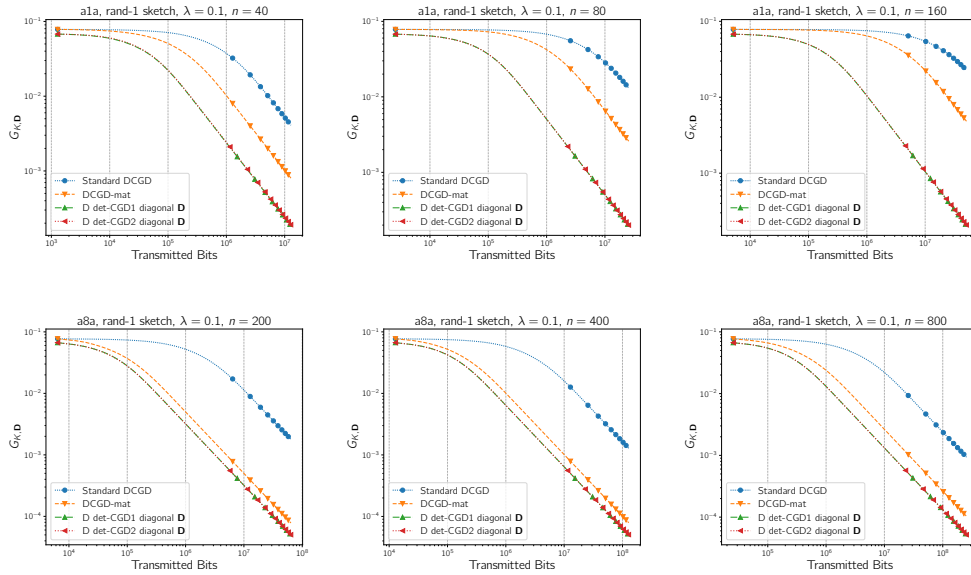


Figure 6: Comparison, in terms of communication complexity, of standard DCGD, DCGD-mat, D-det-CGD1 with matrix stepsize D_1 and D-det-CGD2 with matrix stepsize D_2 , where D_1, D_2 are the optimal diagonal matrix stepsizes for D-det-CGD1 and D-det-CGD2 respectively. Rand-1 sketch is used in all the algorithms throughout the experiment. The x -axis is the number of bits transmitted. The notation $G_{K,D}$ in the y -axis is defined in (66).

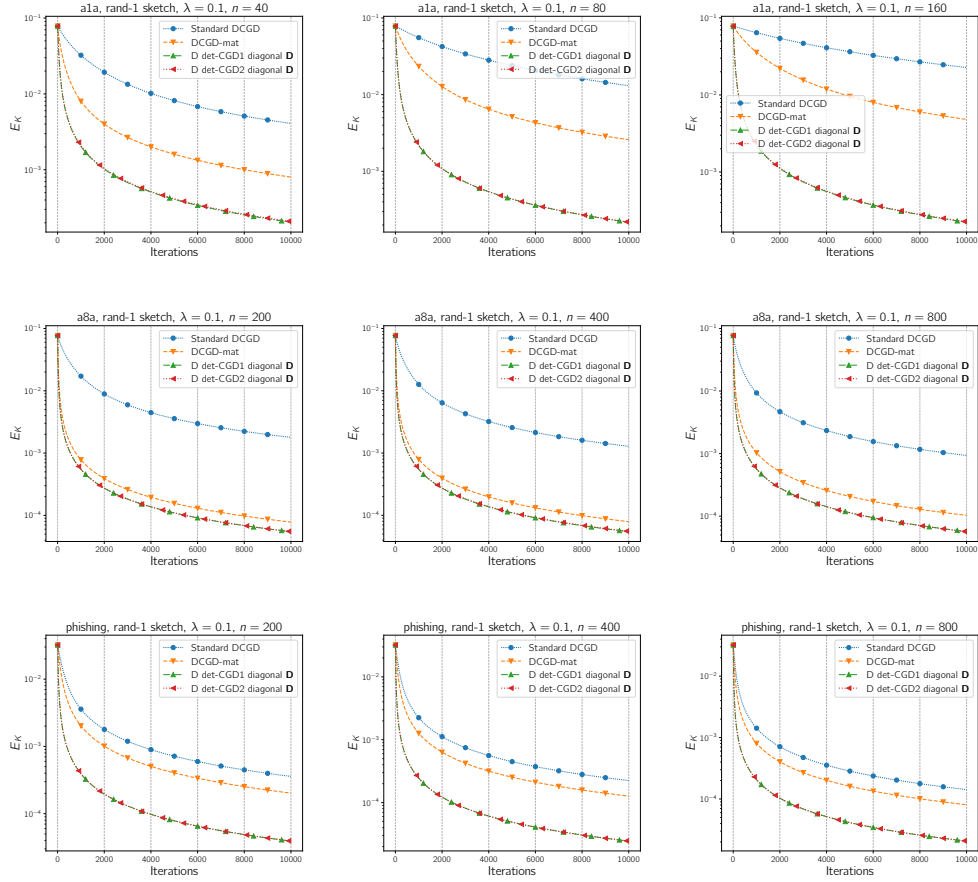


Figure 7: Comparison, in terms of iteration complexity, of standard DCGD, DCGD-mat, **D-det-CGD1** with matrix stepsize D_1 and **D-det-CGD2** with matrix stepsize D_2 , where D_1, D_2 are the optimal diagonal matrix stepsizes for **D-det-CGD1** and **D-det-CGD2** respectively. Rand-1 sketch is used in all the algorithms throughout the experiment. The x -axis is the number of iterations. The y -axis is now standard Euclidean norm defined in (67).