On Accelerating Distributed Convex Optimizations

Kushal Chakrabarti
Department of Electrical and Computer Engineering
University of Maryland
College Park, Maryland 20742, U.S.A.
kchak@terpmail.umd.edu

Nirupam Gupta
École polytechnique fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
nirupam.gupta@epfl.ch

Nikhil Chopra
Department of Mechanical Engineering
University of Maryland
College Park, Maryland 20742, U.S.A.
nchopra@umd.edu

Abstract

This paper studies a distributed multi-agent convex optimization problem. The system comprises multiple agents in this problem, each with a set of local data points and an associated local cost function. The agents are connected to a server, and there is no inter-agent communication. The agents' goal is to learn a parameter vector that optimizes the aggregate of their local costs without revealing their local data points. In principle, the agents can solve this problem by collaborating with the server using the traditional distributed gradient-descent method. However, when the aggregate cost is *ill-conditioned*, the gradient-descent method (i) requires a large number of iterations to converge, and (ii) is highly unstable against process noise. We propose an *iterative pre-conditioning* technique to mitigate the deleterious effects of the cost function's conditioning on the convergence rate of distributed gradient-descent. Unlike the conventional pre-conditioning techniques, the pre-conditioner matrix in our proposed technique updates iteratively to facilitate implementation on the distributed network. In the distributed setting, we provably show that the proposed algorithm converges *linearly* with an improved rate of convergence than the traditional and adaptive gradient-descent methods. Additionally, for the special case when the minimizer of the aggregate cost is unique, our algorithm converges *superlinearly*. We demonstrate our algorithm's superior performance compared to prominent distributed algorithms for solving real logistic regression problems and emulating neural network training via a noisy quadratic model, thereby signifying the proposed algorithm's efficiency for distributively solving non-convex optimization. Moreover, we empirically show that the proposed algorithm results in faster training without compromising the generalization performance.

1 Introduction

In this paper, we consider solving multi-agent distributed convex optimization problems. Precisely, we consider m agents in the system. The agents can only interact with a common server and the overall system is assumed to be synchronous. Each agent $i \in \{1, \ldots, m\}$ has a set of local data points and a *local cost function* $f^i : \mathbb{R}^d \to \mathbb{R}$ associated with its local data points. The aim of the agents is to compute a minimum point of the *aggregate cost function* $f : \mathbb{R}^d \to \mathbb{R}$, taking values $f(x) = \sum_{i=1}^{m} f^i(x)$ for each $x \in \mathbb{R}^d$, across all the agents in the system. Formally, the goal of the

agents is to *distributively* compute a common parameter vector $x^* \in \mathbb{R}^d$ such that

$$x^* \in X^* = \arg\min_{x \in \mathbb{R}^d} \sum_{i=1}^{m} f^i(x). \tag{1}$$

Since each agent only partially knows the collective data points, they collaborate with the server for solving the distributed problem (1). However, the agents <u>cannot</u> share their local data points to the server. An algorithm that enables the agents to jointly solve the above problem in the aforementioned settings is defined as a *distributed algorithm*.

In principle, the optimization problem (1) can be solved using the distributed gradient-descent (GD) algorithm [1]. Built upon the prototypical gradient-descent method, several *adaptive gradient algorithms* have been proposed, which distributively solve (1) [2, 3, 4, 5, 6, 7, 8]. Amongst them, some notable distributed algorithms are Nesterov's accelerated gradient-descent (NAG) [2], heavy-ball method (HBM) [3], adaptive momentum estimation (Adam) [4], BFGS method [5]. All these methods are iterative in nature, wherein the server maintains an estimate of a solution defined in (1) and iteratively refines it using the sum of *local gradients* computed by the agents.

If the aggregate cost function f is convex, the distributed GD method converges at a rate of $\mathcal{O}(1/t)$, where t is the number of iterations [9]. The momentum-based methods, such as NAG, HBM, and Adam, improve upon the convergence rate of GD. In particular, the Adam method has been demonstrated to compare favorably with other optimization algorithms for a wide range of machine learning problems [10, 11, 12]. However, for general convex cost functions f, these algorithms converges at a *sublinear* rate [13, 14]. For the special case of strong convex cost function f, the aforementioned methods converge *linearly* [9, 13, 14]. Quasi-Newton methods, such as BFGS, explore second-order information of the cost functions. If the aggregate cost function's Hessian is positive definite at a solution of (1), then the BFGS method locally converges to a solution at *superlinear* rate [5].

We propose an *adaptive gradient algorithm* for improving the convergence rate of the distributed gradient-descent method when solving the convex optimization problem (1) in distributed networks. The key concept in our proposed method is *iterative pre-conditioning*. The idea of *iterative pre-conditioning* has been originally proposed in [16], wherein the server updates the estimate using the sum of the agents' gradient multiplied with a suitable iterative pre-conditioning matrix. However, [16] considers only quadratic cost functions. Note that the iterative pre-conditioning in [16] does not trivially extend to general cost functions due to non-linearity in the gradients. The proposed algorithm rigorously extends that idea of iterative pre-conditioning to general convex optimization problems (1). Using real-world datasets, we empirically show that the proposed algorithm
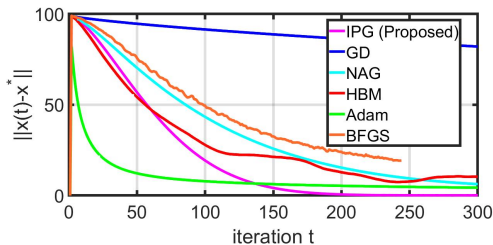

Figure 1: Estimation error $\|x(t) - x^*\|$ of different algorithms for the noisy quadratic model of neural network training [15]. All the algorithms have the same initial estimate $\{x_i(0) \in \mathcal{N}(0, 1),\ i = 1, \dots, d\}$. Other parameters are in the supplemental material.

converges in fewer iterations compared to the aforementioned state-of-the-art methods for solving the distributed convex optimization problem (1). Besides empirical results, we also present a formal analysis of the proposed algorithm's convergence.

Several research works have used quadratic models for approximating the loss functions of neural networks [17, 18, 19, 20, 21, 22]. Quadratic model-based analyses of neural networks have produced vital insights, including learning rate scheduling [18], and the Bayesian viewpoint of SGD with fixed stepsize [23]. The noisy quadratic model (NQM) with carefully chosen model parameters is a proxy for neural network training. The NQM parameters proposed in [15] make predictions that are aligned with deep neural networks in realistic experimental settings. The results in [15] are supported by further evidence from [24], which rigorously show that quadratic loss function model governs infinite width neural network of arbitrary depth. The theoretical NQM in [21] correctly captures the short-horizon bias of learning rates in neural network training. Thus, although the general optimization model of neural networks is non-convex, noisy convex quadratic models such as [15] trim away inessentials features while capturing the key aspects of real neural network training, including

generalization performance [22] and the effects of pre-conditioning on gradients. This motivates us to implement our proposed iterative pre-conditioning scheme on the noisy quadratic model [15] of neural networks and empirically validating it for solving general non-convex optimization problems, including neural network training.

1.1 Summary of our contributions

Our key contributions can be summarized as follows.

- We formally show that our algorithm converges *linearly* for a general class of convex cost functions when the Hessian is non-singular at a solution of (1). In the special case when the solution of (1) is unique, the convergence of our algorithm is *superlinear*. Formal details are presented in Theorem 1 and Theorem 2 in Section 2.2.

- We rigorously show that the convergence rate of our algorithm compares favorably to prominent distributed algorithms, namely the GD, NAG, HBM, and Adam methods.

  - When the solution (1) is unique, our algorithm converges *superlinearly* which is only comparable to the BFGS method [5]. The convergence of the other aforementioned algorithms, on the other hand, is only *linear* [9, 13, 14].
  - Moreover, in the general case, our algorithm converges *linearly*. On the other hand, the convergence of GD, NAG, HBM, and Adam is *sublinear* [9, 13, 14]. Only the BFGS method has *superlinear* rate in this case.

- Though numerical experiments on real-world data analysis problems, we demonstrate the improved convergence rate of our algorithm.

  - The noisy quadratic model in [15] has been claimed to emulate neural network training. Our empirical study shows that the proposed algorithm converges faster than the aforementioned distributed methods on this model, thereby demonstrating the proposed algorithm's efficiency for distributed solution of non-convex optimization problems. Please refer to Section 3.1 for further details.
  - Our empirical results for distributed binary logistic regression problem on the "MNIST" and "CIFAR-10" datasets validate the proposed algorithm's superior convergence rate and comparable test accuracy, particularly under the influence of process noise in the system and the stochastic mini-batch setting. Our algorithm's final estimated training loss is favorable to BFGS and comparable to the other methods. Please refer to Section 3.2 for further details.

2 Proposed algorithm: Iteratively Pre-conditioned Gradient-descent (IPG)

In this section, we present our algorithm. Our algorithm follows the basic prototype of the distributed gradient-descent method. However, a notable difference is that in our algorithm, the server multiplies the aggregate of the gradients received from the agents by a *pre-conditioner* matrix. The server uses the *pre-conditioned* aggregates of the agents' gradients to update its current estimates. This technique is commonly known as *pre-conditioning* [25]. When the aggregate cost f is strongly convex, the best pre-conditioner matrix for the gradient-descent method is the inverse Hessian matrix $\nabla^2 f(x(t))^{-1}$, resulting in Newton's method which converges *superlinearly* [5]. However, $\nabla^2 f(x(t))^{-1}$ cannot be computed directly in a distributed setting as it requires the agents to send their local Hessian matrices $\nabla^2 f^i(x(t))$ to the server, which may require the agents to share their local data points, such as the quadratic optimization problem [16]. Thus, we propose a distributed scheme where the server iteratively updates the pre-conditioner matrix without requiring the agents to share their local Hessian and their local data points in such a way that it linearly converges to $\nabla^2 f(x(t))^{-1}$. Thus, our algorithm eventually converges to Newton's method and has *superlinear* convergence rate for strongly convex problems, as shown later in Section 2.2.

Next, we describe the proposed Iteratively Pre-conditioned Gradient-descent (IPG) algorithm. The algorithm is iterative when in each iteration $t \in \{0, 1, \ldots\}$, the server maintains an estimate $x(t)$ of a minimum point (1), and a pre-conditioner matrix $K(t) \in \mathbb{R}^{d \times d}$. Both the estimate and the pre-conditioner matrix are updated using steps presented below.

Initialization: Before starting the iterative process, the server chooses an initial estimate $x(0)$ and a pre-conditioner matrix $K(0)$ from $\mathbb{R}^d$ and $\mathbb{R}^{d \times d}$, respectively. Further, the server chooses a sequence of non-negative scalar parameters $\{\alpha(t), t \geq 0\}$ and two non-negative scalar real-valued parameters δ, β. The server broadcasts parameter β to all the agents. The specific values of these parameters are presented later in Section 2.2.

2.1 Steps in each iteration t

In each iteration t, the algorithm comprises four steps that are executed collaboratively by the server and the agents.

- *Step 1:* The server sends the estimate $x(t)$ and the matrix $K(t)$ to each agent $i \in \{1, \ldots, m\}$.
- *Step 2:* Each agent $i \in \{1, \ldots, m\}$ computes the gradient $g^i(t)$ of its local cost function, defined as

$$g^i(t) = \nabla f^i(x(t)). \tag{2}$$

  Let I denote the $(d \times d)$-dimensional identity matrix. Let e_j and $k_j(t)$ denote the j-th columns of matrices I and $K(t)$, respectively, so that $K(t) = [k_1(t), \ldots, k_d(t)]$.

  In the same step, each agent i computes a set of vectors $\left\{ R_j^i(t) : j = 1, \ldots, d \right\}$ such that for each j,

$$R_j^i(t) = \left(\nabla^2 f^i(x(t)) + \left(\frac{\beta}{m} \right) I \right) k_j(t) - \left(\frac{1}{m} \right) e_j. \tag{3}$$

- *Step 3:* Each agent i sends gradient $g^i(t)$ and the set of vectors $\left\{ R_j^i(t), j = 1, \ldots, d \right\}$ to the server.
- *Step 4:* The server updates the estimate $x(t)$ to $x(t+1)$ such that

$$x(t+1) = x(t) - \delta K(t) \sum_{i=1}^{m} g^i(t). \tag{4}$$

  The server updates the pre-conditioner matrix $K(t)$ to $K(t+1)$ such that

$$k_j(t+1) = k_j(t) - \alpha(t) \sum_{i=1}^{m} R_j^i(t), \quad j = 1, ..., d. \tag{5}$$

Next, we formally analyze convergence of the proposed IPG algorithm. For analysis, we consider only full batched-data. However, IPG is applicable also to the stochastic setting, as we show later in the experimental results (ref. Section 3).

2.2 Convergence guarantee

We make the following assumptions for our theoretical results. Recall that f denotes the aggregate cost function, i.e., $f = \sum_{i=1}^{m} f^i$, and x^* denotes a minimum point of f, defined in (1).

Assumption 1. *Assume that the minimum of function f exists and is finite, i.e., $|\min_{x \in \mathbb{R}^d} f(x)| < \infty$.*

Assumption 2. *Assume that each local cost function f^i is convex and twice continuously differentiable over a convex domain $\mathcal{D} \subseteq \mathbb{R}^d$ containing the set of minimum points X^*, defined in (1).*

Assumption 3. *Assume that the Hessian $\nabla^2 f$ is Lipschitz continuous over the domain $\mathcal{D}$ with respect to the 2-norm with Lipschitz constant γ. Specifically, for any $x, y \in \mathcal{D}$,*

$$\left\| \nabla^2 f(x) - \nabla^2 f(y) \right\| \leq \gamma \left\| x - y \right\|.$$

Assumption 4. *Assume that the Hessian $\nabla^2 f$ is non-singular at any minimum point $x^* \in X^*$.*

Assumption 2 implies that the aggregate cost function $f = \sum_{i=1}^{m} f^i$, defined in (1), is twice continuously differentiable over $\mathcal{D}$. Thus, its gradient, denoted by ∇f is Lipschitz continuous over $\mathcal{D}$. In other words, there exists a positive real valued constant scalar l such that for any $x, y \in \mathcal{D}$,

$$\left\| \nabla f(x) - \nabla f(y) \right\| \leq l \left\| x - y \right\|. \tag{6}$$

Notation: To formally state our convergence result we introduce some notation below.

- For $\beta > 0$, we define
$$K^* = \left(\nabla^2 f(x^*) + \beta I\right)^{-1}.$$

  Under Assumption 2; the function f is convex, thus $\left(\nabla^2 f(x^*) + \beta I\right)$ is positive definite when $\beta > 0$, and hence K^* is well-defined.

- We let η denote the induced 2-norm of K^*, i.e., $\eta = \|K^*\|$.

- For each iteration $t \geq 0$, we define
$$\rho(t) = \left\| I - \alpha(t)\left(\nabla^2 f(x(t)) + \beta I\right) \right\|.$$

- We let $\lambda_{max}\left[\cdot\right]$ and $\lambda_{min}\left[\cdot\right]$, respectively denote the largest and smallest eigenvalue of a matrix.

Lemma 1 below states a preliminary result about the convergence of the sequence of pre-conditioner matrices $\{K(t), t \geq 0\}$ to K^*. This lemma is important for our key results presented afterward. Proof of Lemma 1 is deferred to Appendix A.2.

Lemma 1. *Consider the IPG Algorithm with parameters $\beta > 0$ and $\alpha(t)$ subject to*

$$0 < \alpha(t) < \frac{1}{\lambda_{max}\left[\nabla^2 f(x(t))\right] + \beta}, \quad \forall t \geq 0. \tag{7}$$

Then, under Assumptions 1 and 2, $\rho(t) \in [0,\ 1)$ for all $t \geq 0$.

Note that under the conditions assumed in Lemma 1,

$$\rho := \sup_{t \geq 0} \rho(t) \text{ exists, and is less than } 1.$$

We now present below in Theorem 1 a key convergence result of our proposed IPG method. Proof of Theorem 1 is deferred to Appendix A.3. Recall that x^* denotes a minimum point of f, defined in (1).

Theorem 1. *Suppose that Assumptions 1, 2, 3 and 4 hold true. Consider the IPG Algorithm with parameters $\beta > 0$, $\delta = 1$ and $\alpha(t)$ satisfying (7) for all $t \geq 0$. Let the parameter β, the initial estimate $x(0) \in \mathcal{D}$ and pre-conditioner matrix $K(0)$ be chosen such that*

$$\frac{\eta\gamma}{2}\left\|x(0) - x^*\right\| + l\left\|K(0) - K^*\right\| + \eta\beta \leq \frac{1}{2\mu} \tag{8}$$

where $\mu \in \left(1,\ \frac{1}{\rho}\right)$ and $\eta = \|K^*\|$. If for $t \geq 0$,

$$\alpha(t) < \min\left\{ \frac{1}{\lambda_{max}\left[\nabla^2 f(x(t))\right] + \beta},\ \frac{\mu^t(1 - \mu\rho)}{2l(1 - (\mu\rho)^{t+1})} \right\}, \tag{9}$$

then we obtain that, for all $t \geq 0$,

$$\left\|x(t+1) - x^*\right\| \leq \frac{1}{\mu}\left\|x(t) - x^*\right\|. \tag{10}$$

Since $\mu > 1$, Theorem 1 implies that the sequence of estimates $\{x(t), t \geq 0\}$ locally converges to a solution $x^* \in X^*$ with a *linear* convergence rate $\frac{1}{\mu}$. To obtain a simpler condition on $\alpha(t)$, compared to (9), we can use a more conservative upper bound. Specifically, let Λ be an upper bound of $\lambda_{max}\left[\nabla^2 f(x(t))\right]$ that holds true for each iteration t. From condition (8) in Theorem 1, we have $0 < \mu\rho < 1$ and $\mu > 1$. Thus,

$$\frac{\mu^t(1 - \mu\rho)}{2l(1 - (\mu\rho)^{t+1})} > \frac{\mu^t(1 - \mu\rho)}{2l}.$$

Ultimately, from above we infer that if

$$\alpha(t) < \min\left\{ \frac{1}{\Lambda + \beta},\ \frac{\mu^t(1 - \mu\rho)}{2l} \right\}, \quad \forall t \geq 0,$$

then condition (9) is implied, and Theorem 1 holds true.

5

The case of strong convexity: We further show that our algorithm attains *superlinear* convergence when the aggregate cost function f is *strongly convex* (however, the local costs may only be convex), as stated in the assumption below.

Assumption 5. *Assume that the aggregate cost function f is strongly convex over the domain $\mathcal{D}$.*

In this case, the Hessian $\nabla^2 f$ is positive definite over the entire domain $\mathcal{D}$. Thus, solution to problem (1) is unique, which we denote by x^*. Under Assumption 5, we show that IPG method with parameter $\beta = 0$ converges *superlinearly*.

Theorem 2. *Consider the IPG Algorithm presented in Section 2.1, with parameters $\beta = 0$, $\delta = 1$, and $\alpha(t)$ for each iteration $t \geq 0$. If Assumptions 1-3, Assumption 5, and the conditions (8)-(9) are satisfied, then the following statement holds true:*

$$\lim_{t \to \infty} \frac{\|x(t+1) - x^*\|}{\|x(t) - x^*\|} = 0. \tag{11}$$

Proof of Theorem 2 is deferred to Appendix A.4.

3 Experimental results

This section presents our results on three different experiments, validating the convergence of our proposed algorithm on real-world problems and its comparison with related methods.

3.1 Distributed noisy quadratic model

In the first experiment, we implement our proposed algorithm for distributively solving (1) in the case of the noisy quadratic model (NQM) [15], which has been shown to agree with the results of training real neural networks. The noisy quadratic model of neural networks consists of a quadratic aggregate cost function $f(x) = \frac{1}{2}x^T H x$, $x \in \mathbb{R}^d$ where H is a $(d \times d)$-dimensional diagonal matrix whose i-th element in the diagonal is $\frac{1}{i}$, and each gradient query is corrupted with independent and identically distributed noise of zero mean and the diagonal covariance matrix H. In the experiments, we choose $d = 10^4$, which is also the condition number of the Hessian H. The distributed implementation of the aforementioned noisy quadratic model is set up as follows. The data points represented by the rows of the matrix H are divided amongst $m = 10$ agents. Thus, each agent $1, \ldots, 10$ has a data matrix of dimension $10^3 \times 10^4$. Since the Hessian matrix H in the noisy quadratic model is positive definite, the optimization problem (1) has a unique solution $x^* = 0_d$.

Table 1: Comparisons between the number of iterations required by different algorithms to attain the specified values for the relative estimation errors ϵ_{tol}.

Dataset	ϵ_{tol}	IPG	GD	NAG	HBM	Adam	BFGS
Noisy Quadratic	10^{-3}	242	$> 10^4$	$> 10^4$	$> 10^4$	$> 10^4$	unstable after 245
MNIST (deterministic)	0	214	$> 10^4$	486	462	851	39
CIFAR-10 (deterministic)	0	196	$> 10^4$	627	634	3750	79

We compare the performance of our proposed IPG algorithm with other algorithms when implemented on the above NQM in the distributed architecture. Specifically, these algorithms are the distributed versions of gradient-descent (GD) [1], Nesterov's accelerated gradient-descent (NAG) [2], heavy-ball method (HBM) [3], adaptive momentum estimation (Adam) [4], and BFGS method [5].

The parameters of the respective distributed algorithms are selected such that each of these methods converges in a fewer number of iterations. Specifically, these parameters are selected as described below. For the GD, NAG, and HBM methods, the algorithm's parameters are set such that each of these methods achieves the optimal (smallest) rate of convergence. The specific definition of these

optimal parameters can be found in [26]. Since the Hessian of the noisy quadratic model is positive definite, Assumption 5 holds. So, we set the parameter $\beta = 0$ for the IPG method. The optimal convergence rate of the linear version of the proposed IPG method is obtained when $\alpha = \frac{2}{\lambda_1 + \lambda_d}$ [16]. Here, $\lambda_1 = 1$ and $\lambda_d = \frac{1}{d}$ respectively denote the largest and the smallest eigenvalue of H. We find that the IPG method implemented for the NQM converges fastest when the parameter α is set similarly as $\alpha(t) = \frac{2}{\lambda_1 + \lambda_d}$. The step-size parameter $\alpha(t)$ of Adam is selected from the set $\{c, \frac{c}{\sqrt{t}}, \frac{c}{t}\}$ where c is from the set $\{0.01, 0.05, 0.1, 0.5, 1, 2\}$. The other parameters of Adam are set at their usual values of $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The step-size of the BFGS method is obtained using the backtrack routine. The best parameter combinations from above are reported in the supplemental material.

The initial estimate $x(0)$ for these algorithms is randomly drawn from the standard normal distribution. The initial pre-conditioner matrix $K(0)$ for the IPG algorithm is the zero matrix of dimension $(d \times d)$. The initial approximation of the Hessian matrix $B(0)$ for the BFGS algorithm is the identity matrix of dimension $(d \times d)$.

We compare the number of iterations needed by these algorithms to reach a *relative estimation error* defined as $\epsilon_{tol} = \frac{\|x(t) - x^*\|}{\|x(0) - x^*\|}$. The results are recorded in Table 1 and Figure 1, from which we find that the proposed IPG algorithm converges fastest among the distributed algorithms implemented on the NQM.

Table 2: Comparisons between the number of iterations required by different algorithms to attain the specified values for relative error in estimated cost ϵ_{tol} when subjected to noise.

Dataset	ϵ_{tol}	IPG	GD	NAG	HBM	Adam	BFGS
MNIST (process noise)	$6e{-}8$	216	$> 10^4$	$> 10^4$	532	878	43
MNIST (mini-batch)	$2e{-}3$	737	$> 10^4$	$> 10^4$	$> 10^4$	$> 10^4$	$> 10^4$
CIFAR-10 (process noise)	$4e{-}6$	289	$> 10^4$	$> 10^4$	350	1191	$> 10^4$
CIFAR-10 (mini-batch)	$2e{-}3$	1960	$> 10^4$	3151	$> 10^4$	$> 10^4$	$> 10^4$

3.2 Distributed logistic regression

In the second and third experiments, we implement the proposed algorithm for distributively solving the logistic regression problem, respectively on the "MNIST" [27] and "CIFAR-10" [28] datasets.

Table 3: Comparisons between the relative final error sse in estimated cost function obtained by different algorithms when subjected to noise.

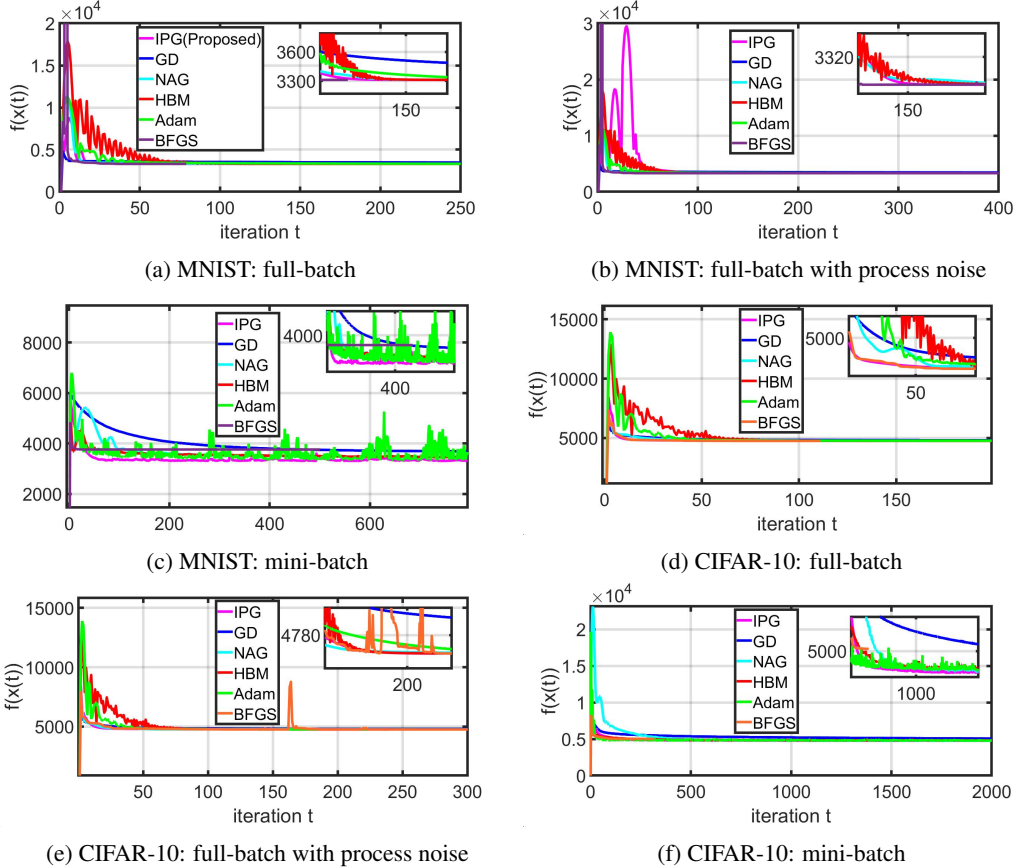| Dataset	IPG	GD	NAG	HBM	Adam	BFGS
MNIST (process noise)	$6e{-}8$	$6e{-}8$	$6e{-}8$	$6e{-}8$	$3e{-}8$	∞
MNIST (mini-batch)	$2e{-}3$	$3e{-}3$	$5e{-}3$	$1e{-}2$	$1e{-}1$	$1e{-}1$
CIFAR-10 (process noise)	$4e{-}6$	$3e{-}6$	$1e{-}5$	0	$2e{-}8$	∞
CIFAR-10 (mini-batch)	$2e{-}3$	$2e{-}2$	$1e{-}3$	$8e{-}3$	$3e{-}3$	∞

Figure 2: Estimated aggregate cost $f(x(t))$ of different algorithms represented by different colors. For all the algorithms, $\{x_i(0) \in \mathcal{N}(0, 0.1), \; i = 1, \ldots, d\}$. Additionally, $K(0) = O_{d \times d}$ for IPG and $B(0) = I_{d \times d}$ for BFGS. The other parameters are enlisted in the supplemental material.

From the training examples of the "MNIST" dataset, we select 10^4 arbitrary instances labeled as either the digit one or the digit five. For each instance, we calculate two quantities, namely the average intensity and the average symmetry of the image [29]. Let the column vectors a_1 and a_2 respectively denote the average intensity and the average symmetry of those 10^4 instances. These two features are then mapped to a second-order polynomial space. Then, our input data matrix before pre-processing is $\begin{bmatrix} a_1 & a_2 & a_1.^2 & a_1.*a_2 & a_2.^2 \end{bmatrix}$. Here, $(.*)$ represents element-wise multiplication and $(.^2)$ represents element-wise squares. This raw data matrix is then pre-processed as follows. Each column is shifted by the mean value of the corresponding column and then divided by the standard deviation of that column. Finally, a 10^4-dimensional column vector of unity is appended to this pre-processed matrix. This is our final input matrix A of dimension $(10^4 \times 6)$. The collective data points (A, B) are then distributed among $m = 10$ agents, in the manner already described in Section 3.1.

Similarly, from the training examples of the "CIFAR-10" dataset, we select 10^4 instances labeled as either "airplane" or "automobile". For each instance, we calculate six quantities, namely the average intensity and the average symmetry of the image for each of the three pixel-colors. Let the column vector $a_1, \ldots, a_6$ respectively denote these six features of all 10^4 instances. We apply a feature transform so that our input matrix before pre-processing is $\begin{bmatrix} a_1 & \ldots a_6 & a_1.^2 & \ldots a_6.^2 \end{bmatrix}$. The pre-processing step and splitting the data among agents is the same as described above.

We compare the performance of the proposed IPG method with other algorithms for distributively solving the logistic regression problem on the "MNIST" and "CIFAR-10" datasets. As in Section 3.1, the algorithm parameters are selected such that each of these methods converges in a fewer number of iterations. Specifically, the parameter selection is described below. The best parameter combinations,

8

Table 4: Comparisons between the test error obtained by training with different algorithms.

Dataset	IPG	GD	NAG	HBM	Adam	BFGS
MNIST (process noise)	0.13	0.13	0.13	0.13	0.13	N/A
MNIST (mini-batch)	0.14	0.14	0.14	0.14	0.14	0.15
CIFAR-10 (process noise)	0.20	0.20	0.20	0.21	0.21	N/A
CIFAR-10 (mini-batch)	0.21	0.21	0.21	0.21	0.21	N/A

for which the respective algorithms converge in a fewer number of iterations, are then reported for each dataset in the supplemental material. The algorithms are initialized as described in Section 3.1.

IPG: The parameter $\alpha(t)$ is selected from the set $\{c \times 10^{-3}, c \times 10^{-4}\}$ where c is from the set $\{1, 2, 5\}$. The parameter δ is chosen from the set $\{1, 0.1, 0.05\}$ and β from the set $\{0, 0.1, 1\}$.

GD: The GD algorithm has only one parameter α, which is selected from the set $\{c \times 10^{-3}, c \times 10^{-4}\}$ where c is from the set $\{1, 2, 5\}$.

NAG and HBM: The parameter α is selected from the set $\{c \times 10^{-3}, c \times 10^{-4}\}$ where c is from the set $\{1, 2, 3, 5\}$. The parameter β is selected from the set $\{0.91, 0.92, \ldots, 0.99\}$.

Adam: The parameters of Adam are selected from the set described in Section 1.

BFGS: The step-size $\alpha(t)$ is either obtained following the backtrack routine or chosen from the set $\{c \times 10^{-p}\}$ where c is from the set $\{1, 2, 5\}$ and p is from $\{2, 3, 4, 5\}$.

For each dataset, we implement the distributed algorithms in three different settings as follows.

Full-batch or deterministic: In this setting, we compare the number of iterations needed by different algorithms to reach a minimum point. Let f^* denote the value of aggregate cost function at this minimum point. The results are recorded in Table 1, Figure 2a, and Figure 2d. We observe that only the BFGS algorithm converges in fewer iterations than IPG.

Full-batch with process noise: In this setting, we add random process noise to each iterated variable of the respective algorithms. Specifically, the additive noise is *uniformly* distributed within $(0, 2.3 \times 10^{-4})$ and $(0, 10^{-4})$, respectively for the MNIST and CIFAR-10 datasets. We compare (i) the number of iterations needed to reach a *relative estimated cost* defined as $\epsilon_{tol} = \frac{f(x(t)) - f^*}{f^*}$, and (ii) the *final error in aggregate cost*, defined as $sse = \lim_{t \to \infty} \frac{f(x(t)) - f^*}{f^*}$, obtained by an algorithm. For the criteria-(i), each iterative algorithm is run (ref. Figure 2b and Figure 2e) until its relative estimated cost does not exceed ϵ_{tol} over a period of 10 consecutive iterations, and the smallest such iteration is reported in Table 2. For the criteria-(ii), each iterative algorithm is run until the changes in its subsequent estimated cost is less than 10^{-4} over 50 consecutive iterations, and the final relative error in estimated cost is reported in Table 3. For the MNIST dataset, we observe that only the BFGS algorithm reaches ϵ_{tol} in fewer iterations compared to IPG (ref. Table 2). However, the BFGS method diverges after 83 iterations. The final estimated cost for IPG is comparable to the other algorithms, except for BFGS, which diverges (ref. Table 3). For the CIFAR-10 dataset, the IPG algorithm needs the least iterations (ref. Table 2). However, its final cost is slightly larger than HBM and Adam and comparable to GD (ref. Table 3).

Mini-batch: Each agent computes its local gradients based on 10 randomly selected local data points at each iteration in this setting. As there are $m = 10$ agents, the server updates the estimate

based on an effective batch size of 100. We compare the algorithms based on criteria-(i) and -(ii) described in the previous setting. From Table 2, the IPG algorithm requires the least number of iterations to reach ϵ_{tol} for both the datasets. For the MNIST dataset, the IPG algorithm has the smallest final estimated cost (ref. Table 3). For the CIFAR-10 dataset, only NAG has a smaller final cost than IPG (ref. Table 3).

We compare the performance of the algorithms on the respective test datasets (ref. Table 4). We observe that the test error of the logistic regression model trained with the proposed IPG algorithm is comparable with the other distributed algorithm, including the GD method. Thus, the IPG algorithm results in faster training without degrading the generalization performance.

4 Summary

We have proposed an adaptive gradient-descent algorithm for distributed learning without requiring the agents to reveal their local data points. The key ingredient of our proposed algorithm is an iterative pre-conditioning technique, which provably improves the convergence rate over the existing adaptive gradient methods in the distributed architecture. We have empirically shown our algorithm's faster convergence, and comparable generalization performance, when compared to the state-of-the-art distributed methods. In this work, we have assumed the distributed network to be synchronous. Our theoretical results have been derived for the full-batch settings and convex cost functions. However, through experiments we have also shown the efficacy of our algorithm in stochastic settings and in training of non-convex neural network via a noisy quadratic model.

References

[1] Dimitri P Bertsekas and John N Tsitsiklis. *Parallel and distributed computation: numerical methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.

[2] Y Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Sov. Math. Doklady*, 27(2):372–376, 1983.

[3] Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.

[4] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[5] Carl T Kelley. *Iterative methods for optimization*. SIAM, 1999.

[6] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.

[7] Matthew D Zeiler. Adadelta: An adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.

[8] Timothy Dozat. Incorporating nesterov momentum into Adam. 2016.

[9] Jeffrey A Fessler. Image reconstruction: Algorithms and analysis. `http://web.eecs.umich.edu/~fessler/book/c-opt.pdf`, 2021. [Online book draft; accessed 27-May-2021].

[10] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.

[11] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

[12] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 2018.

[13] Qianqian Tong, Guannan Liang, and Jinbo Bi. Calibrating the adaptive learning rate to improve convergence of ADAM. *arXiv preprint arXiv:1908.00700*, 2019.

[14] Weijie Su, Stephen Boyd, and Emmanuel Candes. A differential equation for modeling Nesterov's accelerated gradient method: Theory and insights. *Advances in neural information processing systems*, 27:2510–2518, 2014.

[15] Guodong Zhang, Lala Li, Zachary Nado, James Martens, Sushant Sachdeva, George E Dahl, Christopher J Shallue, and Roger Grosse. Which algorithmic choices matter at which batch sizes? Insights from a noisy quadratic model. *arXiv preprint arXiv:1907.04164*, 2019.

[16] Kushal Chakrabarti, Nirupam Gupta, and Nikhil Chopra. Iterative pre-conditioning to expedite the gradient-descent method. In *2020 American Control Conference (ACC)*, pages 3977–3982. IEEE, 2020.

[17] James Lucas, Juhan Bae, Michael R Zhang, Stanislav Fort, Richard Zemel, and Roger Grosse. Analyzing monotonic linear interpolation in neural network loss landscapes. *arXiv preprint arXiv:2104.11044*, 2021.

[18] Tom Schaul, Sixin Zhang, and Yann LeCun. No more pesky learning rates. In *International Conference on Machine Learning*, pages 343–351. PMLR, 2013.

[19] Stanisław Jastrzębski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017.

[20] Michael R Zhang, James Lucas, Geoffrey Hinton, and Jimmy Ba. Lookahead optimizer: k steps forward, 1 step back. *arXiv preprint arXiv:1907.08610*, 2019.

[21] Yuhuai Wu, Mengye Ren, Renjie Liao, and Roger Grosse. Understanding short-horizon bias in stochastic meta-optimization. *arXiv preprint arXiv:1803.02021*, 2018.

[22] Yeming Wen, Kevin Luk, Maxime Gazeau, Guodong Zhang, Harris Chan, and Jimmy Ba. An empirical study of large-batch stochastic gradient descent with structured covariance noise. *arXiv e-prints*, pages arXiv–1902, 2019.

[23] Stephan Mandt, Matthew D Hoffman, and David M Blei. Stochastic gradient descent as approximate Bayesian inference. *arXiv preprint arXiv:1704.04289*, 2017.

[24] Jaehoon Lee, Lechao Xiao, Samuel S Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Journal of Statistical Mechanics: Theory and Experiment*, 2020 (12):124002, 2020.

[25] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.

[26] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.

[27] MNIST in CSV. `https://www.kaggle.com/oddrationale/mnist-in-csv`, 2018. Accessed: 22-May-2021.

[28] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

[29] Yaser S Abu-Mostafa, Malik Magdon-Ismail, and Hsuan-Tien Lin. *Learning from data*, volume 4. AMLBook New York, NY, USA, 2012.

A Supplemental material

A.1 Algorithm parameters used in the experiments

Table 5: The parameters used in different algorithms.*

Dataset	IPG	GD	NAG	HBM	Adam	BFGS
	$\alpha(t), \delta, \beta$	α	α, β	α, β	$\alpha(t), \beta_1, \beta_2, \epsilon$	$\alpha(t)$
Noisy Quadratic	$1.99, 1, 0$	1.99	$1.33, 0.97$	$3.92, 0.96$	$\frac{1}{t}, 0.9, 0.999, 1e{-}8$	according to backtrack [5]
MNIST (deterministic)	$5e{-}4, 1, 0$	$5e{-}4$	$5e{-}4, 0.97$	$1e{-}3, 0.94$	$2, 0.9, 0.999, 1e{-}8$	$1e{-}3$
MNIST (process noise)	$5e{-}4, 1, 0$	$5e{-}4$	$5e{-}4, 0.97$	$1e{-}3, 0.94$	$2, 0.9, 0.999, 1e{-}8$	$1e{-}3$
MNIST (mini-batch)	$1e{-}4, 0.05, 1$	$1e{-}4$	$5e{-}4, 0.97$	$1e{-}3, 0.95$	$1, 0.9, 0.999, 1e{-}8$	0.05
CIFAR-10 (deterministic)	$2e{-}4, 1, 0$	$2e{-}4$	$1e{-}4, 0.95$	$3e{-}4, 0.94$	$\frac{1}{\sqrt{t}}, 0.9, 0.999, 1e{-}8$	$2e{-}4$
CIFAR-10 (process noise)	$2e{-}4, 0.05, 0$	$2e{-}4$	$1e{-}4, 0.93$	$3e{-}4, 0.94$	$0.1, 0.9, 0.999, 1e{-}8$	$1e{-}5$
CIFAR-10 (mini-batch)	$1e{-}3, 0.05, 0$	$2e{-}4$	$2e{-}4, 0.95$	$2e{-}4, 0.92$	$\frac{1}{\sqrt{t}}, 0.9, 0.999, 1e{-}8$	according to backtrack

The argument behind selection of these specific parameter values has been described in Section 3.1, for the noisy quadratic model, and Section 3.2, for the MNIST and CIFAR-10 datasets.

A.2 Proof of Lemma 1

In this section, we prove the result in Lemma 1.

Recall the definition of K^* from Section 2.2. For each iteration $t \geq 0$, define the matrix

$$\widetilde{K}(t) = K(t) - K^*. \tag{12}$$

Consider an arbitrary iteration $t \geq 0$. Upon substituting from (3) and (5) in the definition (12),

$$\begin{aligned}
\widetilde{K}(t+1) =& K(t) - \alpha(t) \left(\left(\nabla^2 f(x(t)) + \beta I \right) K(t) - I \right) - K^* \\
=& K(t) - K^* - \alpha(t) \left(\left(\nabla^2 f(x(t)) + \beta I \right) K(t) - \left(\nabla^2 f(x(t)) + \beta I \right) K^* \right) \\
& - \alpha(t) \left(\left(\nabla^2 f(x(t)) + \beta I \right) K^* - \left(\nabla^2 f(x^*) + \beta I \right) \left(\nabla^2 f(x^*) + \beta I \right)^{-1} \right).
\end{aligned}$$

Upon substituting above from (12) and the definition of K^*,

$$\widetilde{K}(t+1) = \left(I - \alpha(t) \left(\nabla^2 f(x(t)) + \beta I \right) \right) \widetilde{K}(t) - \alpha(t) \left(\nabla^2 f(x(t)) - \nabla^2 f(x^*) \right) K^*.$$

Applying Cauchy-Schwartz inequality above,

$$\left\| \widetilde{K}(t+1) \right\| \leq \left\| I - \alpha(t) \left(\nabla^2 f(x(t)) + \beta I \right) \right\| \left\| \widetilde{K}(t) \right\| + \alpha(t) \left\| \nabla^2 f(x(t)) - \nabla^2 f(x^*) \right\| \left\| K^* \right\|. \tag{13}$$

Under Assumption 2, each local cost function f^i is convex, implying that the aggregate cost function $f = \sum_{i=1}^{m} f^i$ is convex. Thus, $\left(\nabla^2 f(x(t)) + \beta I\right)$ is positive definite for $\beta > 0$. It follows that, if $\alpha(t) < \frac{1}{\lambda_{max}[\nabla^2 f(x(t))] + \beta}$ then $\rho(t) = \left\| I - \alpha(t)\left(\nabla^2 f(x(t)) + \beta I\right) \right\| < 1$. This proves the lemma.

A.3   Proof of Theorem 1

In this section, we present a formal proof of our main result in Theorem 1. We proceed with the proof in three steps as follows.

Before presenting the proof, we define the following notation. Recall the definition of the minimum points X^* in (1). For a minimum points $x^* \in X^*$ and each iteration $t \geq 0$, define the estimation error

$$z(t) = x(t) - x^*. \tag{14}$$

Step 1:   In this step, we derive an upper bound of $\left\| \widetilde{K}(t+1) \right\|$, provided with the iterations from 0 to t. In this step, we use (13) from the proof of Lemma 1.

Consider an arbitrary iteration $t \geq 0$. From Section 2.2, recall that η denotes the induced 2-norm of the matrix K^*. Then,

$$\eta = \|K^*\| = \left\| \left(\nabla^2 f(x^*) + \beta I\right)^{-1} \right\| = \frac{1}{\lambda_{min}\left[\nabla^2 f(x^*)\right] + \beta}. \tag{15}$$

Under Assumption 2 and 4, $\nabla^2 f(x^*)$ is positive definite. Then, $\lambda_{min}\left[\nabla^2 f(x^*)\right] > 0$, and $\eta < \frac{1}{\beta}$. From Section 2.2, recall the definition $\rho = \sup_{t \geq 0} \rho(t)$. From Lemma 1 we have $0 \leq \rho < 1$. Upon substituting from (15) and the definition of ρ in (13),

$$\left\| \widetilde{K}(t+1) \right\| \leq \rho \left\| \widetilde{K}(t) \right\| + \eta\, \alpha(t) \left\| \nabla^2 f(x(t)) - \nabla^2 f(x^*) \right\|. \tag{16}$$

Under Assumption 3, $\left\| \nabla^2 f(x(t)) - \nabla^2 f(x^*) \right\| \leq \gamma \|x(t) - x^*\|$. Upon substituting above from the definition (14), $\left\| \nabla^2 f(x(t)) - \nabla^2 f(x^*) \right\| \leq \gamma \|z(t)\|$. Upon substituting from above in (16),

$$\left\| \widetilde{K}(t+1) \right\| \leq \rho \left\| \widetilde{K}(t) \right\| + \eta\gamma\, \alpha(t) \|z(t)\|.$$

Iterating the above from t to 0 we have

$$\left\| \widetilde{K}(t+1) \right\| \leq \rho^{t+1} \left\| \widetilde{K}(0) \right\| + \eta\gamma\, \alpha(t) \left(\|z(t)\| + \rho \|z(t-1)\| + \ldots + \rho^t \|z(0)\| \right). \tag{17}$$

Step 2:   In this step, we derive an upper bound on the estimation error $\|z(t+1)\|$, provided iteration t.

Upon substituting from (4) in (14), with the parameter $\delta = 1$,

$$z(t+1) = z(t) - K(t) \sum_{i=1}^{m} g^i(t).$$

Upon substituting above from (2) and the definition of aggregate cost function $f = \sum_{i=1}^{m} f^i$,

$$z(t+1) = z(t) - K(t)\nabla f(x(t)).$$

Upon substituting above from the definition of $\widetilde{K}(t)$ in (12),

$$z(t+1) = z(t) - K^*\nabla f(x(t)) - \widetilde{K}(t)\nabla f(t)$$
$$= -K^*\left(\nabla f(x(t)) - (K^*)^{-1} z(t)\right) - \widetilde{K}(t)\nabla f(x(t)).$$

Since x^* is a minimum point of the aggregate cost function f in (1), the first order necessary condition states that $\nabla f(x^*) = 0_d$. Here, 0_d denotes the d-dimensional zero vector. Thus, the above equation can be rewritten as

$$z(t+1) = -K^*\left(\nabla f(x(t)) - \nabla f(x^*) - (K^*)^{-1} z(t)\right) - \widetilde{K}(t)\nabla f(x(t)). \tag{18}$$

14

Using the fundamental theorem of calculus [5],

$$\nabla f(x(t)) - \nabla f(x^*) = (x(t) - x^*) \int_0^1 \nabla^2 f(yx(t) + (1-y)x^*)dy.$$

From above and the definition of $K^* = \left(\nabla^2 f(x^*) + \beta I\right)^{-1}$ (see Section 2.2), we have

$$\nabla f(x(t)) - \nabla f(x^*) - (K^*)^{-1} z(t)$$

$$= (x(t) - x^*) \int_0^1 \nabla^2 f(yx(t) + (1-y)x^*)dy - \left(\nabla^2 f(x^*) + \beta I\right) z(t)$$

$$= (x(t) - x^*) \int_0^1 \left(\nabla^2 f(yx(t) + (1-y)x^*) - \nabla^2 f(x^*)\right) dy - \beta z(t).$$

From (14), recall that $z(t) = x(t) - x^*$. Thus, from above we obtain that

$$\left\|\nabla f(x(t)) - \nabla f(x^*) - (K^*)^{-1} z(t)\right\|$$

$$\leq \|z(t)\| \int_0^1 \left\|\nabla^2 f(yx(t) + (1-y)x^*) - \nabla^2 f(x^*)\right\| dy + \beta \|z(t)\|.$$

Under Assumption 3, $\left\|\nabla^2 f(yx(t) + (1-y)x^*) - \nabla^2 f(x^*)\right\| \leq \gamma(1-y) \|z(t)\|$. Thus,

$$\left\|\nabla f(x(t)) - \nabla f(x^*) - (K^*)^{-1} z(t)\right\| \leq \gamma \|z(t)\|^2 \int_0^1 (1-y)dy + \beta \|z(t)\|$$

$$= \frac{\gamma}{2} \|z(t)\|^2 + \beta \|z(t)\|.$$

Upon using Cauchy-Schwartz inequality in (18), and then substituting from above, we obtain that

$$\|z(t+1)\| \leq \eta \frac{\gamma}{2} \|z(t)\|^2 + \eta\beta \|z(t)\| + \left\|\widetilde{K}(t)\right\| \|\nabla f(x(t))\|.$$

Since $\nabla f(x^*) = 0_d$, the above can be rewritten as

$$\|z(t+1)\| \leq \eta \frac{\gamma}{2} \|z(t)\|^2 + \eta\beta \|z(t)\| + \left\|\widetilde{K}(t)\right\| \|\nabla f(x(t)) - \nabla f(x^*)\|.$$

Upon using the Lipschitz property (6) in above,

$$\|z(t+1)\| \leq \eta \frac{\gamma}{2} \|z(t)\|^2 + \eta\beta \|z(t)\| + l \left\|\widetilde{K}(t)\right\| \|z(t)\|. \tag{19}$$

Upon substituting above from (17) in Step 1, we have

$$\|z(t+1)\| \leq \eta \frac{\gamma}{2} \|z(t)\|^2 + \eta\beta \|z(t)\|$$

$$+ l \|z(t)\| \left(\rho^t \left\|\widetilde{K}(0)\right\| + \eta\gamma \, \alpha(t-1) \left(\|z(t-1)\| + \rho \|z(t-2)\| + \ldots + \rho^{t-1} \|z(0)\|\right)\right). \tag{20}$$

Step 3: In this final step of the proof, we prove (10) which is a direct result of the following claim.
Claim. *Given the conditions of Theorem 1, for each iteration $t \geq 0$ the following statement holds true:*

$$\|z(t+1)\| \leq \frac{1}{\mu} \|z(t)\| \quad \text{and} \quad \|z(t)\| < \frac{1}{\mu\eta\gamma}. \tag{21}$$

Proof. We proof the aforementioned claim by the principle of induction. First, we show that the claim is true for iteration $t = 0$. At $t = 0$, from (19) we have

$$\|z(1)\| \leq \|z(0)\| \left(\eta \frac{\gamma}{2} \|z(0)\| + \eta\beta + l \left\|\widetilde{K}(0)\right\|\right).$$

Upon substituting above from the condition (8),

$$\|z(1)\| \leq \frac{1}{2\mu} \|z(0)\|.$$

Since $\mu > 1$, from above we have

$$\|z(1)\| < \frac{1}{\mu}\|z(0)\|.$$

Moreover, (8) implies that $\|z(0)\| < \frac{1}{\mu\eta\gamma}$. Thus, the claim holds for $t = 0$. Now, we assume that the claim is true for the iterations from 0 to t. We need to show that the claim is also true for the iteration $t + 1$. From the above assumption we have

$$\|z(t+1)\| \le \frac{1}{\mu}\|z(t)\| \le \frac{1}{\mu^2}\|z(t-1)\| \le \ldots \le \frac{1}{\mu^{t+1}}\|z(0)\| < \frac{1}{\mu^{t+1}}\frac{1}{\mu\eta\gamma}. \tag{22}$$

Since $\mu > 1$, the above implies that $\|z(t+1)\| < \frac{1}{\mu\eta\gamma}$. Now, consider the expression $(\|z(t)\| + \rho\|z(t)\| + \ldots + \rho^t\|z(0)\|)$ in the R.H.S. of (20) for iteration $t + 1$. Upon substituting from (22),

$$\|z(t)\| + \rho\|z(t)\| + \ldots + \rho^t\|z(0)\| \le \|z(0)\|\left(\frac{1}{\mu^t} + \frac{\rho}{\mu^{t-1}} + \ldots + \rho^t\right) = \|z(0)\|\frac{1-(\mu\rho)^{t+1}}{\mu^t(1-\mu\rho)}.$$

From the condition of Theorem 1 we have that $\mu\rho < 1$. Upon substituting from above in (20) for iteration $t + 1$,

$$\|z(t+2)\| \le \|z(t+1)\|\left(\eta\frac{\gamma}{2}\|z(t+1)\| + \eta\beta + l\rho^{t+1}\left\|\widetilde{K}(0)\right\| + \eta\gamma l\,\alpha(t)\|z(0)\|\frac{1-(\mu\rho)^{t+1}}{\mu^t(1-\mu\rho)}\right). \tag{23}$$

From Step 1 we have $\rho < 1$. Thus, $l\rho^{t+1}\left\|\widetilde{K}(0)\right\| < l\left\|\widetilde{K}(0)\right\|$. Additionally, if $\alpha(t) < \frac{\mu^t(1-\mu\rho)}{2l(1-(\mu\rho)^{t+1})}$, then $\eta\gamma l\,\alpha(t)\|z(0)\|\frac{1-(\mu\rho)^{t+1}}{\mu^t(1-\mu\rho)} < \eta\frac{\gamma}{2}\|z(0)\|$. Using condition (8) then we have

$$\eta\beta + l\rho^{t+1}\left\|\widetilde{K}(0)\right\| + \eta\gamma l\,\alpha(t)\|z(0)\|\frac{1-(\mu\rho)^{t+1}}{\mu^t(1-\mu\rho)} < \frac{1}{2\mu}. \tag{24}$$

We have shown above that $\|z(t+1)\| < \frac{1}{\mu\eta\gamma}$, which implies that $\eta\frac{\gamma}{2}\|z(t+1)\| < \frac{1}{2\mu}$. Upon substituting from above and (24) in (23),

$$\|z(t+2)\| \le \frac{1}{\mu}\|z(t+1)\|.$$

Thus, the claim holds for iteration $t + 1$. Due to the principle of induction, the proof of the claim is complete. □

Equation (10) directly follows from the aforementioned claim, and concludes the proof of Theorem 1.

A.4 Proof of Theorem 2

In this section, we present the proof of Theorem 2. For this proof, we borrow the results (17) and (19) from the proof of Theorem 1.

Under Assumption 5, $\nabla^2 f(x(t))$ is positive definite for $x(t) \in \mathcal{D}$. So, Assumption 5 is a special case of Assumption 4. Moreover, since $\nabla^2 f(x(t))$ is positive definite for $x(t) \in \mathcal{D}$, Theorem 1 is applicable with $\beta = 0$ in this case. To prove (11), consider (19) from the proof of Theorem 1. Upon taking limits on both sides as $t \to \infty$ and substituting $\beta = 0$, we get

$$\lim_{t\to\infty}\frac{\|z(t+1)\|}{\|z(t)\|} \le \lim_{t\to\infty}\left(\eta\frac{\gamma}{2}\|z(t)\| + l\left\|\widetilde{K}(t)\right\|\right). \tag{25}$$

Since $\mu > 1$, from (10) of Theorem 1 we have that $\lim_{t\to\infty}\|z(t)\| = 0$. Since $\rho < 1$ and $\lim_{t\to\infty}\|z(t)\| = 0$, from (17) we have $\lim_{t\to\infty}\left\|\widetilde{K}(t)\right\| = 0$. Upon substituting them in (25) above, we obtain (11). Hence, the proof.