

# Communication-Efficient Accurate Statistical Estimation

Jianqing Fan\*

Department of ORFE, Princeton University

Yongyi Guo

Department of ORFE, Princeton University  
and

Kaizheng Wang

Department of IEOR, Columbia University

February 20, 2022

## Abstract

When the data are stored in a distributed manner, direct applications of traditional statistical inference procedures are often prohibitive due to communication costs and privacy concerns. This paper develops and investigates two Communication-Efficient Accurate Statistical Estimators (CEASE), implemented through iterative algorithms for distributed optimization. In each iteration, node machines carry out computation in parallel and communicate with the central processor, which then broadcasts aggregated information to node machines for new updates. The algorithms adapt to the similarity among loss functions on node machines, and converge rapidly when each node machine has large enough sample size. Moreover, they do not require good initialization and enjoy linear converge guarantees under general conditions. The contraction rate of optimization errors is presented explicitly, with dependence on the local sample size unveiled. In addition, the improved statistical accuracy per iteration is derived. By regarding the proposed method as a multi-step statistical estimator, we show that statistical efficiency can be achieved in finite steps in typical statistical applications. In addition, we give the conditions under which the one-step CEASE estimator is statistically efficient. Extensive numerical experiments on both synthetic and real data validate the theoretical results and demonstrate the superior performance of our algorithms.

---

\*We gratefully acknowledge *NSF grants DMS-1662139, DMS-1712591, DMS-2053832, DMS-2052926, NIH grant 2R01-GM072611-15, and ONR grant N00014-19-1-2120*. We acknowledge computing resources from Columbia University’s Shared Research Computing Facility project, which is supported by *NIH Research Facility Improvement Grant 1G20-RR030893-01*, and associated funds from the New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract C090171, both awarded April 15, 2010.



# 1 Introduction

Statistical inference in modern era faces tremendous challenge on computation and storage. The exceedingly large size of data often makes it impossible to store all of them on a single machine. Moreover, many applications have individual agents (e.g. local governments, research labs, hospitals, smart phones) collecting data independently. Communication is prohibitively expensive due to the limited bandwidth, and direct data sharing raises concerns in privacy and loss of ownership. These constraints make it necessary to develop methodologies for distributed systems, solving statistical problems with divide-and-conquer procedures and communicating only certain summary statistics.

Distributed statistical inference has received considerable attention recently, covering a wide spectrum of topics including  $M$ -estimation (Zhang et al., 2013; Chen and Xie, 2014; Shamir et al., 2014; Rosenblatt and Nadler, 2016; Wang et al., 2017a; Lee et al., 2017b; Battey et al., 2018; Wang et al., 2018; Shi et al., 2018; Banerjee et al., 2019), principal component analysis (Fan et al., 2019; Garber et al., 2017), nonparametric regression (Shang and Cheng, 2017; Szabó and Van Zanten, 2019; Han et al., 2018), quantile regression (Volgushev et al., 2019; Chen et al., 2021), bootstrap (Kleiner et al., 2014), confidence intervals (Jordan et al., 2019; Chen et al., 2021), Bayesian methods (Wang and Dunson, 2013; Jordan et al., 2019), etc. In the commonly-used setting, the overall dataset is partitioned and stored on  $m$  node machines connected to a central processor. Most of the approaches only require one round of communication: the node machines work in parallel and send their results to the central processor, which then aggregates the information to get a final result. As typical examples, Zhang et al. (2013) average the  $M$ -estimators on node machines; Battey et al. (2018) average debiased estimators; and Fan et al. (2019) average subspaces via eigen-decomposition. While these one-shot methods are communication-efficient, they

only work with a small number of node machines (e.g.  $m = o(\sqrt{N})$ , where  $N$  is the total sample size) and require large sample on each, as their theories heavily rely on asymptotic expansions of estimators. Such conditions are easily violated in practice.

Multi-round procedures come as a remedy, which alternate between local computations and global aggregations. It is possible to achieve optimal statistical precision after a few rounds of communications, under broader settings than those for one-shot procedures.

Shamir et al. (2014) propose a Distributed Approximate NEwton (DANE) algorithm where, in each iteration, each node machine minimizes a modified loss function based on its own samples and the gradient information from all other machines obtained through communication. However, for non-quadratic losses, the analysis in Shamir et al. (2014) does not imply any advantage of DANE in terms of communication over distributed implementation of gradient descent. Other approximate Newton algorithms include Zhang and Xiao (2015), Wang et al. (2018), Chen et al. (2021) and Crane and Roosta (2019). Jordan et al. (2019) develop a Communication-efficient Surrogate Likelihood (CSL) framework for estimation and inference in regular parametric models, penalized regression, and Bayesian statistics. A similar method also appears in Wang et al. (2017a). These methods no longer have restrictions on the number of machines such as  $m = o(\sqrt{N})$ .

Due to the nature of Newton-type methods, existing theories for these algorithms heavily rely on good initialization or even self-concordance assumption on loss functions. They essentially focus on improving an initial estimator that is already consistent but not efficient, whose ideas coincide with the classical one-step estimator (Bickel, 1975). Such initialization itself needs additional efforts and assumptions. Moreover, current results still require each machine to have sufficiently many samples so that loss functions on different machines are similar to each other. These all make the proposed methods unreliable in

practice.

Aside from distributed statistical inference, there has also been a vast literature in distributed optimization. The ADMM (Boyd et al., 2011) is a celebrated example among the numerous algorithms that handle deterministic optimization problems with minimum structural assumption. Yet, the convergence can be quite slow and it cannot fully utilize the similarity among loss functions on node machines.

In this paper, we develop and study two Communication-Efficient Accurate Statistical Estimators (CEASE) based on multi-round algorithms for distributed statistical estimation. Our new algorithms extend the DANE algorithm (Shamir et al., 2014) to regularized empirical risk minimization. Moreover, we provide sharp convergence guarantees for general scenarios, even if the local loss functions are dissimilar and regularization is nonsmooth.

We assume that all the  $m$  node machines have the same sample size  $n$ . Each has a regularized empirical risk function  $f_k + g$  defined by the samples stored there, and the goal is to compute the minimizer of the overall regularized risk function  $\frac{1}{m} \sum_{k=1}^m f_k + g$  to statistical precision. When  $n$  is sufficiently large, their rates of convergence are better than or comparable to existing methods designed for this large-sample regime. For moderate or small  $n$ , they are still guaranteed to converge linearly even without good initialization, while other statistical methods fail. In addition, our algorithms take advantage of the similarity among  $\{f_k\}_{k=1}^m$  and thus improve over general-purpose algorithms like ADMM. They interpolate between distributed algorithms for statistical estimation and general deterministic problems. Theoretical findings are verified by extensive numerical experiments. From a technical point of view, our algorithms use the proximal point algorithm (Rockafellar, 1976) as the backbone and obtain inexact updates in a distributed manner. This turns out to be crucial for proving convergence under general conditions. Our techniques are

potentially useful for studying other distributed algorithms.

The rest of this paper is organized as follows. Section 2 introduces the algorithms. Section 3 presents deterministic convergence results. Section 4 provides guarantees in statistical problems. Section 5 shows numerical results on both synthetic and real data. Section 6 concludes the paper and discusses possible future directions.

Here we list the notations used throughout the paper. We denote by  $[n]$  the set  $\{1, 2, \dots, n\}$ . We write  $a_n = O(b_n)$  or  $a_n \lesssim b_n$  if there exists a constant  $C > 0$  such that  $a_n \leq Cb_n$  holds for sufficiently large  $n$ ; and  $a_n \asymp b_n$  if  $a_n = O(b_n)$  and  $b_n = O(a_n)$ . Given  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^k$  and  $r > 0$ , we define  $B(\mathbf{x}, r) = \{\mathbf{z} \in \mathbb{R}^k : \|\mathbf{z} - \mathbf{x}\|_2 \leq r\}$  and  $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{j=1}^k x_j y_j$ . For a convex function  $h$  on  $\mathbb{R}^k$ , we let  $\partial h(\mathbf{x})$  be its sub-differential set at  $\mathbf{x} \in \mathbb{R}^k$ , and  $\text{argmin}_{\mathbf{x} \in \mathbb{R}^k} h(\mathbf{x})$  be the set of its minimizers if  $\inf_{\mathbf{x} \in \mathbb{R}^k} h(\mathbf{x}) > -\infty$ . We use  $\|\cdot\|_2$  to denote the  $\ell_2$  norm of a vector or operator norm of a matrix. For two sequences of random variables  $\{X_n\}_{n=1}^\infty$  and  $\{Y_n\}_{n=1}^\infty$  where  $Y_n \geq 0$ , we write  $X_n = O_{\mathbb{P}}(Y_n)$  if for any  $\varepsilon > 0$  there exists  $C > 0$  such that  $\mathbb{P}(|X_n| \geq CY_n) \leq \varepsilon$  for sufficiently large  $n$ . We use  $\|X\|_{\psi_2} = \sup_{p \geq 1} \frac{1}{\sqrt{p}} \mathbb{E}^{1/p} |X|^p$  to refer to the sub-Gaussian norm of random variable  $X$ , and  $\|\mathbf{X}\|_{\psi_2} = \sup_{\|u\|_2=1} \|\langle u, \mathbf{X} \rangle\|_{\psi_2}$  to denote the sub-Gaussian norm of random vector  $\mathbf{X}$ .

## 2 The CEASE algorithm

### 2.1 Problem setup

Let  $\mathcal{P}$  be an unknown probability distribution over some sample space  $\mathcal{X}$ . For any parameter  $\boldsymbol{\theta} \in \mathbb{R}^p$ , define its population risk  $F(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{X} \sim \mathcal{P}} \ell(\boldsymbol{\theta}; \mathbf{X})$  based on a loss function  $\ell : \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ . In parametric inference problems,  $\ell$  is often chosen as the negative log-likelihood function of some parametric family. Under mild conditions,  $F$  is well-defined

and has a unique minimizer  $\boldsymbol{\theta}^*$ . A ubiquitous problem in statistics and machine learning is to estimate  $\boldsymbol{\theta}^*$  given i.i.d. samples  $\{\mathbf{X}_i\}_{i=1}^N$  from  $\mathcal{P}$ , and the minimizer of the empirical risk  $f(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N \ell(\boldsymbol{\theta}; \mathbf{X}_i)$  becomes a natural candidate. To achieve desirable precision in high-dimensional problems, it is often necessary to incorporate prior knowledge of  $\boldsymbol{\theta}^*$ . A principled approach is the regularized empirical risk minimization

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^p} \{f(\boldsymbol{\theta}) + g(\boldsymbol{\theta})\}, \quad (2.1)$$

where  $g(\boldsymbol{\theta})$  is a deterministic penalty function. Common choices for  $g(\boldsymbol{\theta})$  include the  $\ell_2$  penalty  $\lambda \|\boldsymbol{\theta}\|_2^2$  (Hoerl and Kennard, 1970), the  $\ell_1$  penalty  $\lambda \|\boldsymbol{\theta}\|_1$  (Tibshirani, 1996), and a family of folded concave penalty functions  $\|p_\lambda(|\boldsymbol{\theta}|)\|_1$  such as SCAD (Fan and Li, 2001) and MCP (Zhang, 2010), where  $\lambda > 0$  is a regularization parameter. Throughout the paper, we assume that both  $\ell$  and  $g$  are convex in  $\boldsymbol{\theta}$ , and  $\ell$  is twice continuously differentiable in  $\boldsymbol{\theta}$ . We allow  $g$  to be non-smooth (e.g. the  $\ell_1$  penalty).

Consider the distributed setting where the  $N$  samples are stored on  $m$  machines connected to a central processor. Denote by  $\mathcal{I}_k$  the index set of samples on the  $k$ th machine and  $f_k(\boldsymbol{\theta}) = \frac{1}{|\mathcal{I}_k|} \sum_{i \in \mathcal{I}_k} \ell(\boldsymbol{\theta}; \mathbf{X}_i)$ . For simplicity, we assume that  $\{\mathcal{I}_k\}_{k=1}^m$  are disjoint,  $N$  is a multiple of  $m$ , and  $|\mathcal{I}_k| = n = N/m$  for all  $k \in [m]$ . Then (2.1) can be rewritten as

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^p} \{f(\boldsymbol{\theta}) + g(\boldsymbol{\theta})\}, \quad f(\boldsymbol{\theta}) = \frac{1}{m} \sum_{k=1}^m f_k(\boldsymbol{\theta}). \quad (2.2)$$

Each machine  $k$  only has access to its local data and hence local loss function  $f_k$  and the penalty  $g$ . We aim to solve (2.2) in a distributed manner with both statistical efficiency and communication-efficiency.

## 2.2 Adaptive gradient enhancements and distributed algorithms in large-sample regimes

In the large-sample regime, we drop the regularization term for now and consider the empirical risk minimization problem  $\min_{\boldsymbol{\theta} \in \mathbb{R}^p} f(\boldsymbol{\theta})$  for estimating  $\boldsymbol{\theta}^* = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^p} F(\boldsymbol{\theta})$ . In some problems, direct minimization of  $f$  is costly, while it is easy to obtain some rough estimate  $\bar{\boldsymbol{\theta}}$  that is close to  $\boldsymbol{\theta}^*$  but not as accurate as the global minimizer  $\hat{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^p} f(\boldsymbol{\theta})$ .

[Bickel \(1975\)](#) proposes the one-step estimator based on the local quadratic approximation and shows that it is as efficient as  $\hat{\boldsymbol{\theta}}$  if the initial estimator  $\bar{\boldsymbol{\theta}}$  is accurate enough. Iterating this further results in multiple-step estimators that improve the optimization error and hence statistical errors when the initial estimator is not good enough ([Robinson, 1988](#)).

This inspires us to refine an existing estimator using some proxy of  $f$ .

In the distributed environment, starting from an initial estimator  $\bar{\boldsymbol{\theta}}$ , the gradient vector  $\nabla f(\bar{\boldsymbol{\theta}})$  can easily be communicated. Construct a linear function  $f^{(1)}(\boldsymbol{\theta}) = f(\bar{\boldsymbol{\theta}}) + \langle \nabla f(\bar{\boldsymbol{\theta}}), \boldsymbol{\theta} - \bar{\boldsymbol{\theta}} \rangle$ , the first-order Taylor expansion of  $f$  around  $\bar{\boldsymbol{\theta}}$ . The object function to be minimized can be written as

$$f(\boldsymbol{\theta}) = f^{(1)}(\boldsymbol{\theta}) + R(\boldsymbol{\theta}), \quad \text{where} \quad R(\boldsymbol{\theta}) = f(\boldsymbol{\theta}) - f^{(1)}(\boldsymbol{\theta}).$$

Since the linear function  $f^{(1)}(\boldsymbol{\theta})$  can easily be communicated to each node machine whereas  $R(\cdot)$  can not, the latter is naturally replaced by its subsampled version at node  $k$ :

$$R_k(\boldsymbol{\theta}) = f_k(\boldsymbol{\theta}) - [f_k(\bar{\boldsymbol{\theta}}) + \langle \nabla f_k(\bar{\boldsymbol{\theta}}), \boldsymbol{\theta} - \bar{\boldsymbol{\theta}} \rangle],$$

where  $f_k(\boldsymbol{\theta})$  is the loss function based on the data at node  $k$ . With this replacement, the

---

**Algorithm 1** CSL (Jordan et al., 2019)

---

**Input:** Initial value  $\boldsymbol{\theta}_0$ , number of iterations  $T$ .

**For**  $t = 0, 1, 2, \dots, T - 1$ :

- Each machine evaluates  $\nabla f_k(\boldsymbol{\theta}_t)$  and sends to the 1st machine;
- The 1st machine computes  $\nabla f(\boldsymbol{\theta}_t) = \frac{1}{m} \sum_{k=1}^m \nabla f_k(\boldsymbol{\theta}_t)$  and

$$\boldsymbol{\theta}_{t+1} = \operatorname{argmin}_{\boldsymbol{\theta}} \{f_1(\boldsymbol{\theta}) + g(\boldsymbol{\theta}) - \langle \nabla f_1(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t), \boldsymbol{\theta} \rangle\}$$

and broadcasts to other machines.

**Output:**  $\boldsymbol{\theta}_T$ .

---

target of optimization at node  $k$  becomes  $f^{(1)}(\boldsymbol{\theta}) + R_k(\boldsymbol{\theta})$ , which equals to

$$f_k(\boldsymbol{\theta}) - \langle \nabla f_k(\bar{\boldsymbol{\theta}}) - \nabla f(\bar{\boldsymbol{\theta}}), \boldsymbol{\theta} \rangle$$

up to an additive constant. This function will be called gradient-enhanced loss function, in which the gradient at point  $\bar{\boldsymbol{\theta}}$  based on the local data is replaced by the global one. This function has one very nice fixed point at the global minimum  $\hat{\boldsymbol{\theta}}$ : the minimizer of the adaptive gradient-enhanced function at  $\bar{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}$  is still  $\hat{\boldsymbol{\theta}}$ . This can easily be seen by verifying that the gradient at the point  $\hat{\boldsymbol{\theta}}$  is zero.

The idea of using such an adaptive gradient-enhanced function has been proposed in Shamir et al. (2014) and Jordan et al. (2019), though the motivations are different. Jordan et al. (2019) develop a Communication-efficient Surrogate Likelihood (CSL) method using the gradient-enhanced loss function  $f_1(\boldsymbol{\theta}) - \langle \nabla f_1(\bar{\boldsymbol{\theta}}) - \nabla f(\bar{\boldsymbol{\theta}}), \boldsymbol{\theta} \rangle$  on the first machine, uses the minimizer on that machine as a new estimate, and iterates these steps until convergence. In the presence of a regularizer  $g$  in (2.1), one simply adds  $g$  to the gradient-enhanced loss; see the Algorithm 1 below.

Note that in Algorithm 1, only the first machine solves optimization problems and others just evaluate gradients. These machines are idling while the first one is working

---

**Algorithm 2** Distributed estimation using gradient-enhanced loss

---

**Input:** Initial value  $\boldsymbol{\theta}_0$ , number of iterations  $T$ .

**For**  $t = 0, 1, 2, \dots, T - 1$ :

- Each machine evaluates  $\nabla f_k(\boldsymbol{\theta}_t)$  and sends to the central processor;
- The central processor computes  $\nabla f(\boldsymbol{\theta}_t) = \frac{1}{m} \sum_{k=1}^m \nabla f_k(\boldsymbol{\theta}_t)$  and broadcasts to machines;
- Each machine computes

$$\boldsymbol{\theta}_{t,k} = \operatorname{argmin}_{\boldsymbol{\theta}} \{f_k(\boldsymbol{\theta}) + g(\boldsymbol{\theta}) - \langle \nabla f_k(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t), \boldsymbol{\theta} \rangle\}$$

and sends to the central processor;

- The central processor computes  $\boldsymbol{\theta}_{t+1} = \frac{1}{m} \sum_{k=1}^m \boldsymbol{\theta}_{t,k}$  and broadcasts to machines.

**Output:**  $\boldsymbol{\theta}_T$ .

---

hard. To fully utilize the computing power of machines and accelerate convergence, all the machines can optimize their corresponding gradient-enhanced loss functions in parallel and the central processor then aggregates the results. This is motivated by the Distributed Approximate NEWton (DANE) algorithm (Shamir et al., 2014). Algorithm 2 describes the procedure in detail. Intuitively, the averaging step requires little computation but helps reduce the variance of estimators on node machines and enhance the accuracy.

We now illustrate Algorithm 2 in the context of linear regression. Given samples  $\{(\mathbf{x}_i, y_i)\}_{i \in [N]}$ , the  $k$ th machine defines a quadratic loss function

$$f_k(\boldsymbol{\theta}) = \frac{1}{2n} \sum_{i \in \mathcal{I}_k} (y_i - \mathbf{x}_i^\top \boldsymbol{\theta})^2 = \frac{1}{2} \boldsymbol{\theta}^\top \widehat{\boldsymbol{\Sigma}}_k \boldsymbol{\theta} - \widehat{\mathbf{w}}_k^\top \boldsymbol{\theta} + \frac{1}{2n} \sum_{i \in \mathcal{I}_k} y_i^2.$$

Here,  $\widehat{\boldsymbol{\Sigma}}_k = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{x}_i \mathbf{x}_i^\top$  and  $\widehat{\mathbf{w}}_k = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{x}_i y_i$ . The overall loss function is  $f(\boldsymbol{\theta}) = \frac{1}{m} \sum_{k=1}^m f_k(\boldsymbol{\theta}) = \frac{1}{2} \boldsymbol{\theta}^\top \widehat{\boldsymbol{\Sigma}} \boldsymbol{\theta} - \widehat{\mathbf{w}}^\top \boldsymbol{\theta}$ , where  $\widehat{\boldsymbol{\Sigma}} = \frac{1}{m} \sum_{k=1}^m \widehat{\boldsymbol{\Sigma}}_k$ ,  $\widehat{\mathbf{w}} = \frac{1}{m} \sum_{k=1}^m \widehat{\mathbf{w}}_k$ . Then the

update of Algorithm 2 in one iteration is

$$\boldsymbol{\theta}_{t+1,k} = (\mathbf{I} - \widehat{\boldsymbol{\Sigma}}_k^{-1} \widehat{\boldsymbol{\Sigma}}) \boldsymbol{\theta}_t + \widehat{\boldsymbol{\Sigma}}_k^{-1} \widehat{\mathbf{w}}, \quad (2.3)$$

$$\boldsymbol{\theta}_{t+1} = \left( \mathbf{I} - \frac{1}{m} \sum_{k=1}^m \widehat{\boldsymbol{\Sigma}}_k^{-1} \widehat{\boldsymbol{\Sigma}} \right) \boldsymbol{\theta}_t + \frac{1}{m} \sum_{k=1}^m \widehat{\boldsymbol{\Sigma}}_k^{-1} \widehat{\mathbf{w}}. \quad (2.4)$$

Intuitively, this is a form of contraction towards the global minimizer  $\widehat{\boldsymbol{\theta}}$ . As for the logistic regression, we can also write out the corresponding enhanced losses and minimize them using Newton's method. Due to space limitations, we refer to Appendix C for details.

### 2.3 The CEASE Algorithm in general regimes

Algorithms 1 and 2 are built upon large-sample regimes, with sufficiently strong convexity of  $\{f_k + g\}_{k=1}^m$  and small discrepancy between them. This requires the local sample size  $n$  to be large enough, which may not be the case in practice. Even worse, the required local sample size depends on structural parameters, making such a condition unverifiable. In fact, our numerical experiments confirm the instability of Algorithms 1 and 2 even for moderate  $n$ . A naive method of remedy is to add strict convex quadratic regularization  $q(\boldsymbol{\theta})$ . While this remedy can make the algorithm converge rapidly, the nonadaptive nature of  $q(\boldsymbol{\theta})$  will lead to a wrong target. Instead of using a fixed  $q$ , we will adjust it according to current solutions. The idea stems from the proximal point algorithm (Rockafellar, 1976).

**Definition 2.1.** For any convex function  $h : \mathbb{R}^p \rightarrow \mathbb{R}$ , define the proximal mapping  $\text{prox}_h : \mathbb{R}^p \rightarrow \mathbb{R}^p$ ,  $\mathbf{x} \mapsto \text{argmin}_{\mathbf{y} \in \mathbb{R}^p} \{h(\mathbf{y}) + \|\mathbf{y} - \mathbf{x}\|_2^2/2\}$ .

For a given  $\alpha > 0$ , the proximal point algorithm for minimizing  $h$  iteratively computes

$$\mathbf{x}_{t+1} = \text{prox}_{\alpha^{-1}h}(\mathbf{x}_t) = \text{argmin}_{\mathbf{x} \in \mathbb{R}^p} \{h(\mathbf{x}) + (\alpha/2)\|\mathbf{x} - \mathbf{x}_t\|_2^2\}, \quad \forall t \geq 0,$$

starting from some initial value  $\mathbf{x}_0$ . It is a strongly convex optimization, shrinking towards the current value  $\mathbf{x}_t$ . Under mild conditions,  $\{\mathbf{x}_t\}_{t=0}^\infty$  converges linearly to some  $\hat{\mathbf{x}} \in \operatorname{argmin}_{\mathbb{R}^p} h(\mathbf{x})$  (Rockafellar, 1976).

Now we take  $h = f + g$  and write the proximal point iteration for our problem (2.2):

$$\boldsymbol{\theta}_{t+1} = \operatorname{prox}_{\alpha^{-1}(f+g)}(\boldsymbol{\theta}_t) = \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^p} \left\{ f(\boldsymbol{\theta}) + g(\boldsymbol{\theta}) + \frac{\alpha}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_2^2 \right\}. \quad (2.5)$$

Each iteration (2.5) is a distributed optimization problem, whose object function is not available to node machines. But it can be solved by Algorithms 1 and 2. Specifically, suppose we have already obtained  $\boldsymbol{\theta}_t$  and aim for  $\boldsymbol{\theta}_{t+1}$  in (2.5). Letting  $\tilde{g}(\boldsymbol{\theta}) = g(\boldsymbol{\theta}) + (\alpha/2)\|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_2^2$ , Algorithm 2 starting from  $\tilde{\boldsymbol{\theta}}_0 = \boldsymbol{\theta}_t$  produces iterations over  $s = 0, 1, \dots$

$$\begin{aligned} \tilde{\boldsymbol{\theta}}_{s,k} &= \operatorname{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^p} \left\{ f_k(\boldsymbol{\theta}) + \tilde{g}(\boldsymbol{\theta}) + \langle \nabla f_k(\tilde{\boldsymbol{\theta}}_s) - \nabla f(\tilde{\boldsymbol{\theta}}_s), \boldsymbol{\theta} \rangle \right\}, \quad k \in [m], \\ \tilde{\boldsymbol{\theta}}_{s+1} &= \frac{1}{m} \sum_{k=1}^m \tilde{\boldsymbol{\theta}}_{s,k}. \end{aligned}$$

When  $\alpha + \rho_0 > \delta$ ,  $\{\tilde{\boldsymbol{\theta}}_s\}_{s=0}^\infty$  converges  $Q$ -linearly <sup>[\*]</sup> to  $\boldsymbol{\theta}_{t+1}$ . On the other hand, there is no need to solve (2.5) exactly, as  $\operatorname{prox}_{\alpha^{-1}(f+g)}(\boldsymbol{\theta}_t)$  is merely an intermediate quantity for computing  $\hat{\boldsymbol{\theta}}$ . We therefore only run one iteration of Algorithm 2 and use the resulting approximate solution as  $\boldsymbol{\theta}_{t+1}$ . This considerably simplifies the algorithm, reducing double loops to a single loop, and enhances statistical interpretation of the method as a multi-step estimator. However, it makes technical analysis more challenging. Similarly, we may also use one step of Algorithm 1 to compute the inexact proximal update.

The above discussions lead us to propose two Communication-Efficient Accurate Statistical Estimators (CEASE) in Algorithms 3 and 4, which use the proximal point algorithm

---

<sup>\*</sup>According to Nocedal and Wright (2006), a sequence  $\{\mathbf{x}_n\}_{n=1}^\infty$  in  $\mathbb{R}^p$  is said to converge  $Q$ -linearly to  $\mathbf{x}^* \in \mathbb{R}^p$  if there exists  $r \in (0, 1)$  such that  $\|\mathbf{x}_{n+1} - \mathbf{x}^*\|_2 \leq r\|\mathbf{x}_n - \mathbf{x}^*\|_2$  for  $n$  sufficiently large.

---

**Algorithm 3** Communication-Efficient Accurate Statistical Estimators (CEASE)

---

**Input:** Initial value  $\boldsymbol{\theta}_0$ , regularizer  $\alpha \geq 0$ , number of iterations  $T$ .

**For**  $t = 0, 1, 2, \dots, T - 1$ :

- Each machine evaluates  $\nabla f_k(\boldsymbol{\theta}_t)$  and sends to the 1<sup>st</sup> machine;
- The 1<sup>st</sup> machine computes  $\nabla f(\boldsymbol{\theta}_t) = \frac{1}{m} \sum_{k=1}^m \nabla f_k(\boldsymbol{\theta}_t)$  and

$$\boldsymbol{\theta}_{t+1} = \operatorname{argmin}_{\boldsymbol{\theta}} \left\{ f_1(\boldsymbol{\theta}) + g(\boldsymbol{\theta}) - \langle \nabla f_k(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t), \boldsymbol{\theta} \rangle + \frac{\alpha}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_2^2 \right\},$$

and broadcasts to other machines.

**Output:**  $\boldsymbol{\theta}_T$ .

---

**Algorithm 4** CEASE with averaging

---

**Input:** Initial value  $\boldsymbol{\theta}_0$ , regularizer  $\alpha \geq 0$ , number of iterations  $T$ .

**For**  $t = 0, 1, 2, \dots, T - 1$ :

- Each machine evaluates  $\nabla f_k(\boldsymbol{\theta}_t)$  and sends to the central processor;
- The central processor computes  $\nabla f(\boldsymbol{\theta}_t) = \frac{1}{m} \sum_{k=1}^m \nabla f_k(\boldsymbol{\theta}_t)$  and broadcasts to machines;
- Each machine computes

$$\boldsymbol{\theta}_{t,k} = \operatorname{argmin}_{\boldsymbol{\theta}} \left\{ f_k(\boldsymbol{\theta}) + g(\boldsymbol{\theta}) - \langle \nabla f_k(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t), \boldsymbol{\theta} \rangle + \frac{\alpha}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_2^2 \right\}$$

and sends to the central processor;

- The central processor computes  $\boldsymbol{\theta}_{t+1} = \frac{1}{m} \sum_{k=1}^m \boldsymbol{\theta}_{t,k}$  and broadcasts to machines.

**Output:**  $\boldsymbol{\theta}_T$ .

---

as the backbone and obtain inexact updates in a distributed manner. They are regularized versions of Algorithms 1 and 2, with an additional proximal term in the objective functions.

That term reduces relative differences of the local loss functions on individual machines, and is crucial for convergence when  $\{f_k\}_{k=1}^m$  are not similar enough. In Appendix A we introduce a variant of Algorithm 4 which also stabilizes Algorithm 2.

In each iteration, Algorithm 3 has one round of communication and one optimization problem to solve. Although Algorithm 4 has two rounds of communication per iteration, only one round involves parallel optimization and the other is simply averaging. We will

compare their theoretical guarantees as well as practical performances in the sequel.

Algorithm 4 is an extension of the DANE algorithm in Shamir et al. (2014) to regularized empirical risk minimization. While DANE is originally motivated by mirror descent, we view it as a *distributed implementation of the proximal point algorithm*. The new perspective helps us obtain stronger convergence guarantees. Ideas from the proximal point algorithm have appeared in the literature of distributed stochastic optimization for different purposes such as accelerating first-order algorithms (Lee et al., 2017a) and regularizing sizes of updates (Wang et al., 2017b).

### 3 Deterministic analysis

We first present in Section 3.1 the deterministic (almost sure) results for Algorithms 3 and 4 based on high-level structural assumptions. As special cases of these algorithms with  $\alpha = 0$ , Algorithms 1 and 2 will be analyzed in Section 3.2.

#### 3.1 Deterministic analysis of the CEASE algorithm

**Definition 3.1.** Let  $h : \mathbb{R}^p \rightarrow \mathbb{R}$  be a convex function,  $\Omega \subseteq \mathbb{R}^p$  be a convex set, and  $\rho \geq 0$ .

$h$  is  $\rho$ -strongly convex in  $\Omega$  if  $h(\mathbf{y}) \geq h(\mathbf{x}) + \langle \mathbf{g}, \mathbf{y} - \mathbf{x} \rangle + \frac{\rho}{2} \|\mathbf{y} - \mathbf{x}\|_2^2$ ,  $\forall \mathbf{x}, \mathbf{y} \in \Omega$  and  $\mathbf{g} \in \partial h(\mathbf{x})$ .

**Assumption 3.1** (Strong convexity).  $f + g$  has a unique minimizer  $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^p$ , and is  $\rho$ -strongly convex in  $B(\widehat{\boldsymbol{\theta}}, R)$  for some  $R > 0$  and  $\rho > 0$ .

**Assumption 3.2** (Homogeneity).  $\|\nabla^2 f_k(\boldsymbol{\theta}) - \nabla^2 f(\boldsymbol{\theta})\|_2 \leq \delta$ ,  $\forall k \in [m]$ ,  $\boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R)$ .

We will refer to  $\delta$  as a homogeneity parameter. Based on both assumptions, we define

$$\rho_0 = \sup \left\{ c \in [0, \rho] : \{f_k + g\}_{k=1}^m \text{ are } c\text{-strongly convex in } B(\widehat{\boldsymbol{\theta}}, R) \right\}. \quad (3.1)$$

A simple but useful fact is  $\max\{\rho - \delta, 0\} \leq \rho_0 \leq \rho$ . In most interesting problems, the population risk  $F$  is smooth and strongly convex on any compact set. When  $\{\mathbf{X}_i\}_{i=1}^N$  are i.i.d. and the total sample size  $N$  is large, the empirical risk  $f$  concentrates around  $F$  and inherits nice properties from the latter, making Assumption 3.1 hold easily.

On the other hand, since the empirical risk functions  $\{f_k\}_{k=1}^m$  are i.i.d. stochastic approximations of the population risk  $F$ , they should not be too far away from their average  $f$  provided that  $n$  is not too small. Assumption 3.2 is a natural way of characterizing this similarity. It is a generalization of the concept “ $\delta$ -related functions” for quadratic losses in Arjevani and Shamir (2015). With high probability, it holds with reasonably small  $\delta$  and large  $R$  under general conditions. Large  $n$  implies small homogeneity parameter  $\delta$  and thus similar  $\{f_k\}_{k=1}^m$ . Assumption 3.2 always holds with  $\delta = \max_{k \in [m]} \{\sup_{\boldsymbol{\theta} \in B(\hat{\boldsymbol{\theta}}, R)} \|\nabla^2 f_k(\boldsymbol{\theta}) - \mathbb{E} \nabla^2 f_k(\boldsymbol{\theta})\|_2 + \sup_{\boldsymbol{\theta} \in B(\hat{\boldsymbol{\theta}}, R)} \|\nabla^2 f(\boldsymbol{\theta}) - \mathbb{E} \nabla^2 f_k(\boldsymbol{\theta})\|_2\}$ .

The following additional assumption on smoothness of the Hessian matrix of  $f + g$  is not necessary for contraction, but it helps us obtain a much stronger result on the contraction rate of Algorithm 2, justifying the power of the simple averaging step.

**Assumption 3.3** (Smoothness of Hessian).  $g \in C^2(\mathbb{R}^p)$ , and there exists  $M \geq 0$  such that  $\|[\nabla^2 f(\boldsymbol{\theta}') + \nabla^2 g(\boldsymbol{\theta}')] - [\nabla^2 f(\boldsymbol{\theta}'') + \nabla^2 g(\boldsymbol{\theta}'')]\|_2 \leq M \|\boldsymbol{\theta}' - \boldsymbol{\theta}''\|_2$ ,  $\forall \boldsymbol{\theta}', \boldsymbol{\theta}'' \in B(\hat{\boldsymbol{\theta}}, R)$ .

Theorem 3.1 gives contraction guarantees for Algorithms 3 and 4. It is deterministic and non-asymptotic by nature.

**Theorem 3.1.** *Let Assumptions 3.1 and 3.2 hold. Consider the multi-step estimators  $\{\boldsymbol{\theta}_t\}_{t=0}^T$  generated by Algorithm 3 or 4. Suppose that  $\boldsymbol{\theta}_0 \in B(\hat{\boldsymbol{\theta}}, R/2)$  and  $[\delta/(\rho_0 + \alpha)]^2 < \rho/(\rho + 2\alpha)$ .*

(a) For both Algorithms 3 and 4, we have

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \cdot \frac{\frac{\delta}{\rho_0+\alpha} \sqrt{\rho^2 + 2\alpha\rho} + \alpha}{\rho + \alpha}, \quad 0 \leq t \leq T-1; \quad (3.2)$$

(b) If Assumption 3.3 also holds, then for Algorithm 4 we have

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \cdot \frac{\gamma_t \sqrt{\rho^2 + 2\alpha\rho} + \alpha}{\rho + \alpha}, \quad 0 \leq t \leq T-1, \quad (3.3)$$

where we define  $\gamma_t = \frac{\delta}{\rho_0+\alpha} \cdot \min\{1, \frac{\delta}{\rho+\alpha}(1 + \frac{M}{\rho_0+\alpha}\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2)\}$ ;

(c) Both multiplicative factors in (3.2) and (3.3) are strictly less than 1.

In the contraction factor in (3.2), the two summands  $\frac{\delta\sqrt{\rho^2+2\alpha\rho}}{(\rho_0+\alpha)^2}$  and  $\frac{\alpha}{\rho+\alpha}$  come from bounding the inexact proximal update  $\|\boldsymbol{\theta}_{t+1} - \text{prox}_{\alpha^{-1}(f+g)}(\boldsymbol{\theta}_t)\|_2$  and the residual  $\|\text{prox}_{\alpha^{-1}(f+g)}(\boldsymbol{\theta}_t) - \widehat{\boldsymbol{\theta}}\|_2$ , respectively. Similar results hold for (3.3). The condition  $[\delta/(\rho_0 + \alpha)]^2 < \rho/(\rho + 2\alpha)$  ensures that both contraction factors are less than 1. Note that (3.3) requires Assumption 3.3, which forces  $g$  to be smooth.

Theorem 3.1 shows the  $Q$ -linear convergence of the sequence  $\{\boldsymbol{\theta}_t\}_{t=0}^\infty$  generated by both Algorithms 3 and 4 under quite general settings. The contraction rate depends explicitly on the structural parameters and the choice of  $\alpha$ . The local loss functions  $\{f_k\}_{k=1}^m$  just need to be convex and smooth, and the convex penalty  $g$  is allowed to be non-smooth, e.g. the  $\ell_1$  norm. On the contrary, most algorithms for distributed statistical estimation are only designed for smooth problems, and many of them are only rigorously studied when the loss functions are quadratic or self-concordant (Shamir et al., 2014; Zhang and Xiao, 2015; Wang et al., 2017b). This is another important aspect of our contributions.

We immediately see from Theorem 3.1 that Algorithms 3 and 4 converge linearly as long as  $[\delta/(\rho_0 + \alpha)]^2 < \rho/(\rho + 2\alpha)$ , which is guaranteed to hold by choosing sufficiently large

$\alpha$ . By contrast, however, we'll show in Section 3.2 that Algorithms 1 and 2 (corresponding to  $\alpha = 0$ ) hinges on the homogeneity assumption  $\rho_0 > \delta$  in Theorem 3.2, i.e. the functions  $\{f_k\}_{k=1}^m$  must be similar enough. In the statistical setting, this requires the local sample size  $n$  to be large. Therefore, proper regularization provides a safety net for the algorithms under general regimes with potentially insufficient local sample size. Corollary 3.1 below gives a guideline for choosing  $\alpha$  to make Algorithms 3 and 4 converge in general.

**Corollary 3.1.** *Let Assumptions 3.1 and 3.2 hold,  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R/2)$ , and  $\{\boldsymbol{\theta}_t\}_{t=0}^T$  be the iterates of Algorithm 3 or 4. With any  $\alpha \geq 4\delta^2/\rho$ , both algorithms converge with contraction factors in (3.2) and (3.3) bounded by  $(1 - \frac{\rho}{10(\alpha+\rho)})$ .*

On the other hand, consider the case where the local loss functions have small relative difference  $\delta/\rho$ . In this case, Theorem 3.2 states that the contraction factors for unregularized versions ( $\alpha = 0$ ) of Algorithms 3 and 4 are in the same order of  $\delta/\rho$  and  $(\delta/\rho)^2$ , respectively, which are smaller than the contracting factors with  $\alpha > 0$ . The following corollary characterizes the upper bound for  $\alpha$  so that the contraction factors remain at these small orders.

**Corollary 3.2.** *Let Assumptions 3.1 and 3.2 hold,  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R/2)$ , and suppose  $\alpha \leq C\delta^2/\rho$  for some constant  $C$ . There exist constants  $C_1$  and  $C_2$  such that the followings hold when  $\delta/\rho$  is sufficiently small:*

(a) *Algorithms 3 and 4 have the contraction property*

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq C_1(\delta/\rho)\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2, \quad 0 \leq t \leq T-1;$$

(b) If Assumption 3.3 also holds and  $\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \leq \rho/M$ , then for Algorithm 4

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq C_2(\delta/\rho)^2 \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2.$$

Note that the second result above only holds given Assumption 3.3, which requires a smooth regularization  $g$ . Corollary 3.2 reveals that by choosing  $\alpha \asymp \delta^2/\rho$ , the contraction factors are essentially the same as those of the unregularized ( $\alpha = 0$ ) algorithms when  $\delta/\rho$  is small. By combining Corollaries 3.1 and 3.2, we use  $\alpha \asymp \delta^2/\rho$  as a default choice for Algorithms 3 and 4 to become both fast and robust. They are reliable in general cases (Corollary 3.1) and efficient in nice cases (Corollary 3.2 and Theorem 3.1 with  $\alpha = 0$ )

Algorithms 3 and 4 attain communication efficiency by utilizing similarity among local loss functions: The contraction factors in Corollary 3.2 go to zero if  $\delta/\rho$  does. In fact, both algorithms achieve  $\varepsilon$ -accuracy within  $O(\max\{1, (\delta/\rho)^2\} \log(\frac{\|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2}{\varepsilon}))$  rounds of communication. In contrast, the distributed accelerated gradient descent requires  $O(\sqrt{\kappa_0} \log(\frac{\|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2}{\varepsilon}))$  rounds of communication to achieve  $\varepsilon$ -accuracy (Shamir et al., 2014), with  $\kappa_0$  being the condition number of  $(f + g)$ , which does not take advantage of sample size  $n$ . As long as  $\delta/\rho \ll \kappa_0^{1/4}$ , Algorithms 3 and 4 communicate less than the distributed accelerated gradient descent. And again, our general results for Algorithms 3 and 4 also apply to the case with nonsmooth penalty functions while those for distributed accelerated gradient descent do not.

Moreover, if  $(f + g)$  is smooth and  $\boldsymbol{\theta}_t$  is reasonably close to  $\widehat{\boldsymbol{\theta}}$ , Corollary 3.2 shows that each iteration of Algorithm 4 is roughly equivalent to two iterations of Algorithm 3, although the former only has one round of optimization. The averaging step in Algorithm 4 reduces the error as much as the optimization step, while taking much less time. In this case, Algorithm 4 is preferable, and our numerical experiments also confirm this.

For unregularized empirical risk minimization, i.e.  $g = 0$  in (2.2), Algorithm 4 reduces to an extension or a useful case of the DANE algorithm (Shamir et al., 2014). In this case, Theorem 3.1 and its corollaries refine the analysis of DANE (Shamir et al., 2014) in several aspects. On the one hand, our analysis handle both smooth and nonsmooth problems, while in Shamir et al. (2014), the theoretical analysis beyond quadratic loss requires extremal choice of tuning parameters and does not show any advantage over distributed implementation of the gradient descent. On the other hand, as mentioned in Section 2, we derive Algorithm 4 from the proximal point algorithm with a new prospective, which leads to sharp convergence analysis along with suggestions on choosing the tuning parameter  $\alpha$ . As a by-product, we close a gap in the theory of DANE in non-quadratic settings. Our analysis techniques are potentially useful for other distributed optimization algorithms, especially when the loss is not quadratic.

### 3.2 Deterministic analysis in large-sample regimes

In this section, we restrict ourselves to large-sample regimes where the local sample size  $n$  is sufficiently large such that  $\rho_0 > \delta \geq 0$ , where  $\rho_0$  is the strong convexity parameter in (3.1). The following theorem gives deterministic results for Algorithms 1 and 2.

**Theorem 3.2.** *Let Assumptions 3.1 and 3.2 hold, and  $\rho_0 > \delta \geq 0$ . Consider the iterates  $\{\boldsymbol{\theta}_t\}_{t=0}^\infty$  produced by Algorithm 1 or 2, with  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R)$ . Then*

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq (\delta/\rho_0) \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2, \quad \forall t \geq 0.$$

In addition, if Assumption 3.3 also holds, then for Algorithm 2 we have

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq \frac{\delta}{\rho_0} \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \cdot \min \left\{ 1, \frac{\delta}{\rho} \left( 1 + \frac{M}{\rho_0} \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \right) \right\}, \quad \forall t \geq 0.$$

Note that the last inequality requires Assumption 3.3 and thus a smooth regularization. The first part of Theorem 3.2 is a refinement of the analysis in Jordan et al. (2019), since we allow the initial estimator to be inaccurate and we have more explicit rates of contraction of optimization errors. This will be further demonstrated in Section 4.2.

The second part points out benefits of the averaging step, which is a novel result. Similar to the results on Algorithm 4, Theorem 3.2 shows that when  $\boldsymbol{\theta}_t$  is close to  $\widehat{\boldsymbol{\theta}}$ , with an additional standard assumption on Hessian smoothness, the averaging step alone in Algorithm 2 is almost as powerful as an optimization step in terms of contraction: The contracting constant will eventually be  $\frac{\delta}{\rho_0} \frac{\delta}{\rho}$ . With negligible computational cost, averaging significantly improves upon individual solutions  $\{\boldsymbol{\theta}_{t,k}\}_{k=1}^m$  by doubling the speed of convergence.

## 4 Statistical analysis

We further analyze the statistical properties of the above algorithms under a generalized linear model. Essentially, both the CSL methods and the CEASE algorithm are  $T$ -step estimators, starting from the initial estimator  $\boldsymbol{\theta}_0$ . The question here is the effect of iterations in the multiple step estimators and the role of the initial estimator. We start with statistical analysis of Algorithms 3 and 4 in Section 4.1, and then study Algorithms 1 and 2 in Section 4.2. In Section 4.3, we provide practical guidance when implementing the CEASE algorithm based on these analysis.

## 4.1 Multi-step estimators in general regimes

The deterministic analysis in Section 3.2 applies to a wide range of statistical models. Here we consider the generalized linear model with canonical link, where our samples are i.i.d. pairs  $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$  of covariates and responses and the conditional density of  $y_i$  given  $\mathbf{x}_i$  is given by

$$h(y_i; \mathbf{x}_i, \boldsymbol{\theta}^*) = c(\mathbf{x}_i, y_i) \exp \left( \mathbf{y}_i(\mathbf{x}_i^\top \boldsymbol{\theta}^*) - b(\mathbf{x}_i^\top \boldsymbol{\theta}^*) \right).$$

For simplicity, we let the dispersion parameter to be 1 as we do not consider the issue of over-dispersion;  $b(\cdot)$  is some known convex function, and  $c$  is a known function such that  $h$  is a valid probability density function. The negative log-likelihood of the whole data is an affine transformation of  $f(\boldsymbol{\theta}) = \frac{1}{m} \sum_{k=1}^m f_k(\boldsymbol{\theta})$  with

$$f_k(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i \in \mathcal{I}_k} [b(\mathbf{x}_i^\top \boldsymbol{\theta}) - y_i(\mathbf{x}_i^\top \boldsymbol{\theta})].$$

It is easy to verify that

$$\nabla f_k(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i \in \mathcal{I}_k} [b'(\mathbf{x}_i^\top \boldsymbol{\theta}) - y_i] \mathbf{x}_i \quad \text{and} \quad \nabla^2 f_k(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i \in \mathcal{I}_k} b''(\mathbf{x}_i^\top \boldsymbol{\theta}) \mathbf{x}_i \mathbf{x}_i^\top.$$

Assume that  $\mathbf{x}_i = (1, \mathbf{u}_i^\top)^\top \in \mathbb{R}^p$ , where  $\{\mathbf{u}_i\}_{i=1}^N \subseteq \mathbb{R}^{p-1}$  are i.i.d. random covariate vectors with zero mean and covariance matrix  $\boldsymbol{\Sigma}$ . Suppose there exist universal positive constants  $A_1$ ,  $A_2$  and  $A_3$  such that  $A_1 \leq \|\boldsymbol{\Sigma}\|_2 \leq A_2 p^{A_3}$ . Let  $\boldsymbol{\Sigma}^* = \mathbb{E}(\mathbf{x}_i \mathbf{x}_i^\top) = \begin{pmatrix} 1 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma} \end{pmatrix}$ ,  $g$  be a deterministic penalty function, and  $F(\boldsymbol{\theta}) = \mathbb{E}f(\boldsymbol{\theta})$  be the population risk function. Below we impose some standard regularity assumptions.

**Assumption 4.1.**  $\bullet \quad \{\boldsymbol{\Sigma}^{-1/2} \mathbf{u}_i\}_{i=1}^N$  are i.i.d. sub-Gaussian random vectors.

- For all  $x \in \mathbb{R}$ ,  $|b''(x)|$  and  $|b'''(x)|$  are bounded by some constant.
- $\|\boldsymbol{\theta}^*\|_2$  is bounded by some constant.

As in Assumptions 3.1 and 3.3, the following general assumptions are also needed for our analysis. Here  $R$  is some positive quantity that satisfies  $R < A_4 p^{A_5}$  for some universal constants  $A_4$  and  $A_5$ .

**Assumption 4.2.** *There exists a universal constant  $\rho > 0$  such that  $(F + g)$  is  $\rho$ -strongly convex in  $B(\boldsymbol{\theta}^*, 2R)$ .*

The following smoothness assumption is only needed for a part of our theory; it is used to show that the averaging step in Algorithm 4 can significantly enhance the accuracy.

**Assumption 4.3.**  *$g \in C^2(\mathbb{R}^p)$ , and there exists a universal constant  $M \geq 0$  such that  $\|[\nabla^2 F(\boldsymbol{\theta}') + \nabla^2 g(\boldsymbol{\theta}')] - [\nabla^2 F(\boldsymbol{\theta}'') + \nabla^2 g(\boldsymbol{\theta}'')]\|_2 \leq M\|\boldsymbol{\theta}' - \boldsymbol{\theta}''\|_2$ ,  $\forall \boldsymbol{\theta}', \boldsymbol{\theta}'' \in B(\boldsymbol{\theta}^*, 2R)$ .*

Under the model assumptions above, we can explicitly determine the rate of  $\delta$  in Assumption 3.2. In particular, we will show in Lemma E.5 in Appendix E that

$$\max_{k \in [m]} \max_{\boldsymbol{\theta} \in B(\hat{\boldsymbol{\theta}}, R)} \|\nabla^2 f_k(\boldsymbol{\theta}) - \nabla^2 f(\boldsymbol{\theta})\|_2 = O_{\mathbb{P}} \left( \|\boldsymbol{\Sigma}\|_2 \sqrt{\frac{p(\log p + \log N)}{n}} \right),$$

provided that  $n \geq cp$  for an arbitrary positive constant  $c$ . Therefore, with high probability,  $\delta \asymp \|\boldsymbol{\Sigma}\|_2 \sqrt{p(\log p + \log N)/n}$ . Omitting the logarithmic terms, we see that the contraction factor is approximately  $\kappa \sqrt{p/n}$ , where  $\kappa \triangleq \|\boldsymbol{\Sigma}\|_2/\rho$  can be viewed as a condition number. This rate is more explicit on  $p$  and  $\kappa$  than that in [Jordan et al. \(2019\)](#), where finite  $p$  and  $\kappa$  are assumed. In addition, with a smooth regularization, Algorithm 4 benefits from the averaging step in that it improves the contraction rate to approximately  $\kappa^2 p/n$ .

Let  $\boldsymbol{\theta}_t$  be the  $t$ -th iterate of one of the proposed algorithms. It is clear that the statistical error of the estimator  $\boldsymbol{\theta}_t$  is bounded by its optimization error and the statistical error of  $\widehat{\boldsymbol{\theta}}$ :

$$\|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|_2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 + \|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2.$$

The second term is well-studied in statistics, which is of order  $O_{\mathbb{P}}(\sqrt{p/N})$  under mild conditions. In the following theorem, we show that for the first term (i.e. the optimization error), with proper choice of  $\alpha$ , each iteration of Algorithm 3 or 4 makes  $\boldsymbol{\theta}_t$  closer to the global minimum  $\widehat{\boldsymbol{\theta}}$  by some order depending on local sample size. Thus, through finite steps, the optimization errors are eventually negligible in comparison with statistical errors (assuming  $N$  is of order  $(n/p)^a$  for a finite  $a$  in typical applications), and the distributed multi-step estimator will work as well as the global minimum as if the data were aggregated in the central server.

**Theorem 4.1.** *Suppose that Assumptions 4.1 and 4.2 hold, and with high probability the initial value  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R/2)$ . Let  $\eta = \kappa^2(\log N)p/n$  and  $\kappa = \|\boldsymbol{\Sigma}\|_2/\rho$ . For any  $c_1, c_2 > 0$ , there exists  $C > 0$  such that the followings hold with high probability:*

(a) *If  $n \geq c_1 p$  and  $\alpha \geq C\rho\eta$ , then both Algorithms 3 and 4 have linear convergence*

$$\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \leq \left[1 - \frac{\rho}{10(\alpha + \rho)}\right]^t \|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2, \quad \forall t \geq 0;$$

(b) *If  $\eta$  is sufficiently small and  $\alpha \leq c_2\rho\eta$ , then for both algorithms*

$$\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(\eta^{t/2} \|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2), \quad \forall t \geq 0;$$

in addition, if Assumption 4.3 also holds, then for Algorithm 4 we have

$$\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(\eta^{t-t_0} \|\boldsymbol{\theta}_{t_0} - \widehat{\boldsymbol{\theta}}\|_2), \quad \forall t \geq t_0,$$

where  $t_0 = \lceil \frac{2 \log(CMR/\rho)}{\log(1/\eta)} \rceil$ .

In contrast to a fixed contraction derived by Shamir et al. (2014), Theorem 4.1 explains the significant benefits of large local sample size even in the presence of a non-smooth penalty: the optimization error shrinks by a factor that converges to zero explicitly in  $n$ . As a brief illustration, let us consider the case with smooth loss functions, sufficient local sample size and no regularization. Let  $\boldsymbol{\theta}_0$  be the average of individual estimators on node machines. By Corollary 2 in Zhang et al. (2013), this simple divide-and-conquer estimator has accuracy  $\|\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*\|_2 = O_{\mathbb{P}}(\max\{\sqrt{\frac{p}{N}}, \frac{p}{n}, \frac{\kappa\sqrt{p\log p}}{n}\})$ . Using the explicit expression of  $\eta$ , we can easily deduce from Theorem 4.1 (b) that the one-step estimator  $\boldsymbol{\theta}_1$  obtained by Algorithm 3 behaves the same as the global minimizer  $\widehat{\boldsymbol{\theta}}$  if the local sample size satisfies  $n^3 \gg N(\kappa^2 p \log N)(p + \kappa^2 \log p)$ . In this case, the local optimization in Algorithm 3 can further be replaced by using the explicit one-step estimator as in Bickel (1975) and Jordan et al. (2019), since the initial estimator is in a consistent neighborhood. More generally, the  $t$ -step estimator  $\boldsymbol{\theta}_t$  has negligible optimization error under even weaker local sample size requirement:  $n^{t+2} \gg N(\kappa^2 p \log N)^t(p + \kappa^2 \log p)$ . A similar remark applies to Algorithm 4.

Similar to the deterministic results, the averaging step is about as effective as the optimization step when  $g$  is smooth and  $\boldsymbol{\theta}_t$  is sufficiently close to  $\widehat{\boldsymbol{\theta}}$  in that after a finite  $t_0$  iterations. See a simplified example in Section 4.2 with  $\alpha = 0$ .

As for the initialization, the condition  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R/2)$  is mild, since  $\widehat{\boldsymbol{\theta}}$  is usually a consistent estimate and  $\|\boldsymbol{\theta}^*\|_2$  is bounded (Assumption 4.1). In contrast with Jordan et al. (2019), we allow inaccurate initial value such as  $\boldsymbol{\theta}_0 = \mathbf{0}$  and give more explicit rates of

contraction even when  $p$  and  $\kappa$  diverge. On the other hand, the accuracy of the initial estimator  $\boldsymbol{\theta}_0$  does help reduce the number of iterations.

Combining the results to be presented in Theorem 4.2, we'll see that by choosing  $\alpha \asymp \rho\eta$ , Algorithms 3 and 4 inherit all the merits of Algorithms 1 and 2 in the large- $n$  regime – fast linear contraction of rate  $\sqrt{\eta} = \kappa\sqrt{p(\log N)/n}$ , and for Algorithm 4, an even faster rate of  $\eta = \kappa^2 p(\log N)/n$  to  $\widehat{\boldsymbol{\theta}}$  when the loss and the penalty functions are smooth. These facts also guarantee that Algorithms 3 and 4 reach the statistical efficiency in  $O(\frac{\log\|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2 + \log(N/p)}{\log(1/\eta)})$  iterations. On the other hand, compared to Algorithms 1 and 2, Algorithms 3 and 4 overcome the difficulties with a small local sample size  $n$  in that as long as  $n/p$  is bounded away from some small constant (which is reasonable for many big-data problems of interest), shrinkage of optimization error is guaranteed. Moreover, while it is hard to check whether  $n$  is sufficiently large in practice, proper choice of  $\alpha$  always guarantees linear convergence, and the contraction rates adapt to the sample size  $n$ . In this way, Algorithms 3 and 4 perfectly resolve the main issue of their vanilla versions.

We can get stronger results in the specific case of distributed linear regression, where the contraction rate has nearly no dependence on the conditional number  $\kappa$ . Due to space constraints, we put all the details in Appendix B.

## 4.2 Multi-step estimators in large-sample regimes

We now present the contraction of optimization error of Algorithms 1 and 2.

**Theorem 4.2.** *Suppose that Assumptions 4.1 and 4.2 hold, and with probability tending to one,  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R)$  for some  $R > \|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2$ . For Algorithms 1 and 2, we have*

$$\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(\eta^{t/2} \|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2), \quad \forall t \geq 0,$$

where  $\eta = \kappa^2 p(\log N)/n$ . In addition, let Assumption 4.3 also hold. There exists some constant  $C$  such that for Algorithm 2 we have

$$\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(\eta^{t-t_0} \|\boldsymbol{\theta}_{t_0} - \widehat{\boldsymbol{\theta}}\|_2), \quad \forall t \geq t_0, \quad (4.1)$$

where  $t_0 = \lceil \frac{2 \log(CMR/\rho)}{\log(1/\eta)} \rceil$ .

The strengthened result (4.1) requires Assumption 4.3 and thus smooth  $g$ . Theorem 4.2 shows that when  $n$  is sufficiently large, Algorithms 1 and 2 behave similarly as Algorithms 3 and 4 – faster convergence with the larger  $n$ , mild restrictions on initialization, and averaging speeds up contraction given smooth loss. However, there is no convergence guarantee in general regimes. Section 5 further shows that with insufficient local sample size, the practical performance Algorithms 1 and 2 is less satisfactory.

Finally, to see how the averaging step reduces the statistical error (i.e. the distance between the estimator and  $\boldsymbol{\theta}^*$ ), we continue to look at the linear regression example mentioned at the end of Section 2.2. For simplicity, assume  $\{\mathbf{x}_i\}_{i \in [N]}$  are i.i.d. standard normal random vectors. (2.3) and (2.4) can be expressed as

$$\begin{aligned} \widehat{\boldsymbol{\Sigma}}^{1/2}(\boldsymbol{\theta}_{t+1,k} - \widehat{\boldsymbol{\theta}}) &= \left( \mathbf{I} - \widehat{\boldsymbol{\Sigma}}^{1/2} \widehat{\boldsymbol{\Sigma}}_k^{-1} \widehat{\boldsymbol{\Sigma}}^{1/2} \right) \cdot \widehat{\boldsymbol{\Sigma}}^{1/2}(\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}), \\ \widehat{\boldsymbol{\Sigma}}^{1/2}(\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}) &= \left[ \mathbf{I} - \widehat{\boldsymbol{\Sigma}}^{1/2} \left( \frac{1}{m} \sum_{k=1}^m \widehat{\boldsymbol{\Sigma}}_k^{-1} \right) \widehat{\boldsymbol{\Sigma}}^{1/2} \right] \cdot \widehat{\boldsymbol{\Sigma}}^{1/2}(\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}). \end{aligned}$$

If further  $n \gg p$ , then the two contraction factors satisfy

$$\|\mathbf{I} - \widehat{\boldsymbol{\Sigma}}^{1/2} \widehat{\boldsymbol{\Sigma}}_k^{-1} \widehat{\boldsymbol{\Sigma}}^{1/2}\|_2 = O_{\mathbb{P}}(\sqrt{p/n}), \quad (4.2)$$

$$\left\| \mathbf{I} - \widehat{\boldsymbol{\Sigma}}^{1/2} \left( \frac{1}{m} \sum_{k=1}^m \widehat{\boldsymbol{\Sigma}}_k^{-1} \right) \widehat{\boldsymbol{\Sigma}}^{1/2} \right\|_2 = O_{\mathbb{P}}(p/n). \quad (4.3)$$

When the Algorithm 1 applies to this problem, the expression changes slightly to

$$\widehat{\Sigma}^{1/2}(\boldsymbol{\theta}_{t+1,1} - \widehat{\boldsymbol{\theta}}) = \left( \mathbf{I} - \widehat{\Sigma}^{1/2} \widehat{\Sigma}_1^{-1} \widehat{\Sigma}^{1/2} \right) \cdot \widehat{\Sigma}^{1/2}(\boldsymbol{\theta}_{t,1} - \widehat{\boldsymbol{\theta}}).$$

The smaller magnitude in (4.3) is due to the averaging.

Suppose that we initialize the algorithm using the one-shot average  $\boldsymbol{\theta}_0 = \frac{1}{m} \sum_{k=1}^m \widehat{\boldsymbol{\theta}}_k$ , where  $\widehat{\boldsymbol{\theta}}_k$  is the least squares solution on the  $k$ th machine. When  $n = O(\sqrt{Np})$ , we have  $\sqrt{p/N} = O(p/n)$ . Zhang et al. (2013) assert that  $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 = O_{\mathbb{P}}(\sqrt{\frac{p}{N}})$ ,  $\|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(\frac{p}{n})$ . Thus, the initial statistical error is  $\|\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*\|_2 = O_{\mathbb{P}}(\max\{\sqrt{\frac{p}{N}}, \frac{p}{n}\}) = O_{\mathbb{P}}(\frac{p}{n})$ . By the contraction properties in (4.2) and (4.3), the optimization errors are

$$\|\boldsymbol{\theta}_{1,k} - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(\sqrt{p/n})\|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(p^{3/2}/n^{3/2}),$$

$$\|\boldsymbol{\theta}_1 - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(p/n)\|\boldsymbol{\theta}_0 - \widehat{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(p^2/n^2).$$

We can see that when  $p^2/n^2 \ll \sqrt{p/N} \ll p^{3/2}/n^{3/2}$  (or equivalently,  $p^{3/4}N^{1/4} \ll n \ll p^{3/2}N^{1/3}$ ), the optimization error is negligible for  $\boldsymbol{\theta}_1$ , but is not negligible for  $\boldsymbol{\theta}_{1,k}$ . When  $n$  is smaller, even more iterations are needed. A refined analysis of distributed least squares is in Appendix B.

### 4.3 Guidance on practice

We now provide some general guidance on how to implement the CEASE algorithm in practice. First, we recommend choosing an initialization depending on the magnitude of the local sample size  $n$ . In particular, when  $n$  is not very large compared to the dimension  $p$ , a zero initialization would be more robust. On the other hand, given a moderate or large  $n$ , the one-shot average estimator will lead to extremely fast convergence.

Second, according to Theorem 4.1, it suffices to take  $\alpha$  of the order of  $\rho\kappa^2 p \log N/n$ . In practice, setting  $\alpha$  to be a small multiple of  $p/n$  seems suitable in many occasions.

Finally, as is already shown, both Algorithm 3 and 4 reach statistical efficiency in  $O(\frac{\log \|\theta_0 - \bar{\theta}\|_2 + \log(N/p)}{\log(1/\eta)})$  iterations. In both our simulations and real data examples, the CEASE algorithms with a properly chosen  $\alpha$  converge to the centralized estimator within 10 iterations. With a moderate  $n$ , a warm start further boosts the convergence speed.

## 5 Numerical experiments

### 5.1 Synthetic data

We first conduct distributed logistic regression to illustrate the effect of local sample size and initialization on convergence. We keep the total sample size  $N = 10000$  and the dimensionality  $p = 101$  fixed, and generate the i.i.d. data  $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$  as follows:  $\mathbf{x}_i = (1, \mathbf{u}_i^\top)^\top$  with  $\mathbf{u}_i \sim N(\mathbf{0}_{p-1}, \Sigma)$  and  $\Sigma = \text{diag}(10, 5, 2, 1 \cdots 1) \in \mathbb{R}^{(p-1) \times (p-1)}$ ;  $\mathbb{P}(y_i = 1) = 1 - \mathbb{P}(y_i = 0) = 1/(1 + e^{-\mathbf{x}_i^\top \boldsymbol{\theta}^*})$  where  $\boldsymbol{\theta}^* \in \mathbb{R}^p$  is a random vector with norm 3 whose direction is chosen uniformly at random from the sphere. We use the natural logarithm of the estimation error  $\|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|_2$  to measure the performance of different algorithms, including multiple versions of the CEASE algorithms, GIANT (Wang et al., 2018), ADMM (Boyd et al., 2011) and accelerated gradient descent (Nesterov, 1983).

Figure 1 shows how the estimation errors evolve with iterations. The curves show the average values over 100 independent runs; the error bands correspond to one standard deviation. The regimes “large  $n$ ”, “moderate  $n$ ” and “small  $n$ ” refer to  $(n, m) = (2000, 5)$ ,  $(1000, 10)$  and  $(250, 40)$ ; “zero initialization” and “good initialization” refer to  $\boldsymbol{\theta}_0 = \mathbf{0}$  (bottom panel) and  $\bar{\boldsymbol{\theta}}$  (top panel), respectively. Here  $\bar{\boldsymbol{\theta}}$  is the one-shot distributed estimator

(Zhang et al., 2013) that averages the individual estimators on node machines. According to Figure 1, the standard deviation of each iterate is around 0.1. For the “large  $n$ , good initialization” regime, all of the iterates are unsurprisingly very close to the optimal solution. Their error bands will cover up the curves. So we omit the bands in that case for the sake of clarity.

With proper regularization, the two CEASE algorithms are the only ones that converge rapidly in all scenarios. The purely deterministic methods ADMM (Boyd et al., 2011) and accelerated gradient descent (Nesterov, 1983) are also reliable but slow. Other distributed algorithms like unregularized CEASE and GIANT (Wang et al., 2018) easily fail when the local sample size is small or the initialization is uninformative. In addition, the CEASE with averaging (Algorithm 4) is superior to the one without averaging (Algorithm 3). For example, when  $(n, m) = (1000, 10)$ , the averaged CEASE with  $\alpha = 0$  converges while the one without averaging does not. Hence the averaging step leads to better performance.

We also test the efficacy of our algorithms in the distributed  $\ell_1$ -regularized logistic regression, where the penalty  $g$  is nonsmooth (See Appendix D for details). To summarize, our simulations demonstrate several important properties of the CEASE Algorithms:

- In all scenarios, the CEASE Algorithms converge rapidly, usually within several steps, which is consistent with our theory;
- The CEASE Algorithms efficiently utilize statistical structures and similarities among local losses, and benefit from the averaging step with smooth loss functions;
- The CEASE Algorithms are also able to handle the most general situations (e.g. small local sample size, uninformative initialization) with convergence guarantees.

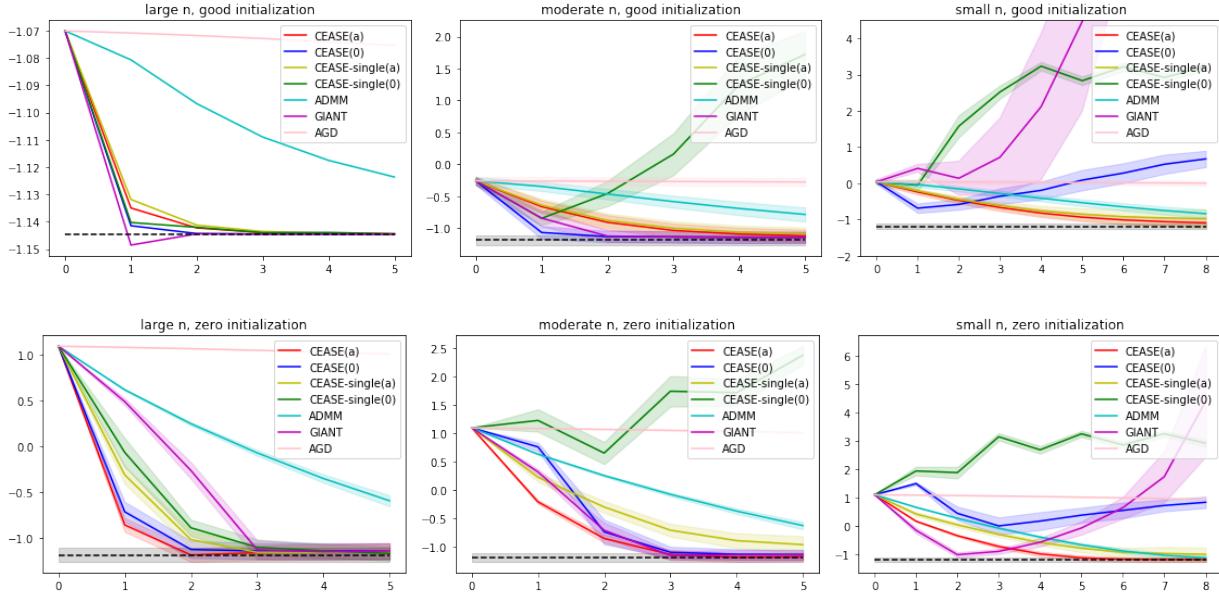


Figure 1: Impacts of local sample size and initialization on convergence. The  $x$ -axis and  $y$ -axis are the number of iterations and  $\log \|\theta_t - \theta^*\|_2$ . The dashed lines show the error of the minimizer of the overall loss function. The top and bottom panels use  $\bar{\theta}$  and  $\mathbf{0}$  for initialization, respectively. CEASE(a) and CEASE(0) refer to Algorithm 4 with  $\alpha = 0.15p/n$  and 0; CEASE-single(a) and CEASE-single(0) refer to Algorithm 3 with  $\alpha = 0.15p/n$  and 0, respectively. In particular, CEASE-single(0) is equivalent to the CSL algorithm in [Jordan et al. \(2019\)](#).

## 5.2 Real data

As a real data example, we choose the Fashion-MNIST dataset ([Xiao et al., 2017](#)) as a testbed for comparison of algorithms. The whole dataset consists of 70000 grayscale images of fashion products in 10 classes, each of which has 6000 training samples and 1000 testing samples. We choose the 7th and 9th classes (Sneakers and Ankle boots) and the goal is to train a classifier that distinguishes them. Each image has  $28 \times 28 = 784$  pixels, represented by a feature vector in  $[0, 1]^{784}$ . The number of training (or testing) samples is  $6000 \times 2 = 12000$  (or  $1000 \times 2 = 2000$ ). We randomly partition the training set and conduct logistic regression in a distributed manner. The performance metric is the classification error on the testing set. Figure 2 shows the average performance of the CEASE algorithms, ADMM, GIANT and AGD based on 100 independent runs, together

with error bars showing one standard deviation. Here ‘‘large  $n$ ’’, ‘‘moderate  $n$ ’’ and ‘‘small  $n$ ’’ refer to  $(n, m) = (1200, 10)$ ,  $(480, 25)$  and  $(240, 50)$ , respectively. All of the iterations are initialized with the one-shot average (Zhang et al., 2013). The experiments on this real data example also support our theoretical findings.

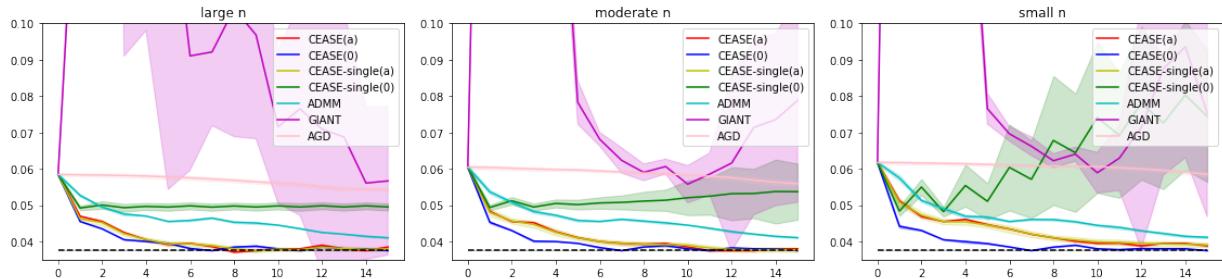


Figure 2: Fashion-MNIST dataset. The  $x$ -axis and  $y$ -axis are the number of iterations and the testing error. The dashed lines show the error of the classifier based on all of the training samples. All of the iterations are initialized with the one-shot average  $\bar{\theta}$ . CEASE(a) and CEASE(0) refer to Algorithm 4 with  $\alpha = 0.15p/n$  and 0; CEASE-single(a) and CEASE-single(0) refer to Algorithm 3 with  $\alpha = 0.15p/n$  and 0, respectively. In particular, CEASE-single(0) is equivalent to the CSL algorithm in Jordan et al. (2019). GIANT and CEASE-single(0) do not converge to the optimal solution.

## 6 Discussions

We have developed two CEASE distributed estimators (Algorithms 3 and 4) for statistical estimation, with theoretical guarantees and superior performance on real data. Several new directions are worth exploring. First, while we assumed exact computation for simplicity, finer analysis should allow for inexact updates in practice. Second, we hope to extend the algorithms to decentralized and asynchronous settings. Third, distributed versions of confidence regions and hypothesis tests are of great importance, and our point estimation strategies may serve as a starting point. Finally, it will be interesting to explore non-convex statistical optimization problems such as mixture models and deep learning. We believe that the idea of gradient-enhanced loss function still plays an important role.

## SUPPLEMENTARY MATERIAL

**Supplementary material:** The file “supplementary.pdf” contains more details and proofs of the results in this paper.

## References

Arjevani, Y. and Shamir, O. (2015). Communication complexity of distributed convex learning and optimization. In *Advances in Neural Information Processing Systems*, pages 1756–1764.

Banerjee, M., Durot, C., and Sen, B. (2019). Divide and conquer in nonstandard problems and the super-efficiency phenomenon. *The Annals of Statistics*, 47(2):720–757.

Battey, H., Fan, J., Liu, H., Lu, J., and Zhu, Z. (2018). Distributed testing and estimation under sparse high dimensional models. *The Annals of Statistics*, 46(3):1352–1382.

Bickel, P. J. (1975). One-step Huber estimates in the linear model. *Journal of the American Statistical Association*, 70(350):428–434.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122.

Chen, X., Liu, W., and Zhang, Y. (2021). First-order Newton-type estimator for distributed estimation and inference. *Journal of the American Statistical Association*, pages 1–17.

Chen, X. and Xie, M.-g. (2014). A split-and-conquer approach for analysis of extraordinarily large data. *Statistica Sinica*, pages 1655–1684.

Crane, R. and Roosta, F. (2019). Dingo: Distributed Newton-type method for gradient-norm optimization. In *Advances in Neural Information Processing Systems*, volume 32.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360.

Fan, J., Wang, D., Wang, K., and Zhu, Z. (2019). Distributed estimation of principal eigenspaces. *The Annals of Statistics*, 47(6):3009.

Garber, D., Shamir, O., and Srebro, N. (2017). Communication-efficient algorithms for distributed stochastic principal component analysis. In *International Conference on Machine Learning*, volume 70, pages 1203–1212.

Han, Y., Mukherjee, P., Ozgur, A., and Weissman, T. (2018). Distributed statistical estimation of high-dimensional and nonparametric distributions. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 506–510. IEEE.

Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67.

Jordan, M. I., Lee, J. D., and Yang, Y. (2019). Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 114(526):668–681.

Kleiner, A., Talwalkar, A., Sarkar, P., and Jordan, M. I. (2014). A scalable bootstrap for massive data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(4):795–816.

Lee, J. D., Lin, Q., Ma, T., and Yang, T. (2017a). Distributed stochastic variance reduced gradient methods by sampling extra data with replacement. *The Journal of Machine Learning Research*, 18(1):4404–4446.

Lee, J. D., Liu, Q., Sun, Y., and Taylor, J. E. (2017b). Communication-efficient sparse

regression. *The Journal of Machine Learning Research*, 18(5):1–30.

Nesterov, Y. E. (1983). A method for solving the convex programming problem with

convergence rate  $O(1/k^2)$ . In *Dokl. Akad. Nauk SSSR*, volume 269, pages 543–547.

Nocedal, J. and Wright, S. J. (2006). *Numerical optimization (Second Edition)*. Springer.

Robinson, P. M. (1988). The stochastic difference between econometric statistics. *Econo-*

*metrica: Journal of the Econometric Society*, pages 531–548.

Rockafellar, R. T. (1976). Monotone operators and the proximal point algorithm. *SIAM*

*Journal on Control and Optimization*, 14(5):877–898.

Rosenblatt, J. D. and Nadler, B. (2016). On the optimality of averaging in distributed

statistical learning. *Information and Inference: A Journal of the IMA*, 5(4):379–404.

Shamir, O., Srebro, N., and Zhang, T. (2014). Communication-efficient distributed opti-

mization using an approximate Newton-type method. In *International Conference on*

*Machine Learning*, pages 1000–1008.

Shang, Z. and Cheng, G. (2017). Computational limits of a distributed algorithm for

smoothing spline. *The Journal of Machine Learning Research*, 18(1):3809–3845.

Shi, C., Lu, W., and Song, R. (2018). A massive data framework for M-estimators with

cubic-rate. *Journal of the American Statistical Association*, pages 1–12.

Szabó, B. and Van Zanten, H. (2019). An asymptotic analysis of distributed nonparametric

methods. *The Journal of Machine Learning Research*, 20:87–1.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the*

*Royal Statistical Society. Series B (Methodological)*, pages 267–288.

Volgushev, S., Chao, S.-K., and Cheng, G. (2019). Distributed inference for quantile regression processes. *The Annals of Statistics*, 47(3):1634–1662.

Wang, J., Kolar, M., Srebro, N., and Zhang, T. (2017a). Efficient distributed learning with sparsity. In *International Conference on Machine Learning*, pages 3636–3645.

Wang, J., Wang, W., and Srebro, N. (2017b). Memory and communication efficient distributed stochastic optimization with minibatch prox. In *Conference on Learning Theory*, pages 1882–1919.

Wang, S., Roosta-Khorasani, F., Xu, P., and Mahoney, M. W. (2018). Giant: Globally improved approximate newton method for distributed optimization. In *Advances in Neural Information Processing Systems*, pages 2338–2348.

Wang, X. and Dunson, D. B. (2013). Parallelizing MCMC via Weierstrass sampler. *arXiv preprint arXiv:1312.4605*.

Xiao, H., Rasul, K., and Vollgraf, R. (2017). Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.

Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942.

Zhang, Y., Duchi, J. C., and Wainwright, M. J. (2013). Communication-efficient algorithms for statistical optimization. *The Journal of Machine Learning Research*, 14:3321–3363.

Zhang, Y. and Xiao, L. (2015). Disco: Distributed optimization for self-concordant empirical loss. In *International Conference on Machine Learning*, pages 362–370.

# Supplement to “Communication-Efficient Accurate Statistical Estimation”

Jianqing Fan\*

Department of ORFE, Princeton University

Yongyi Guo

Department of ORFE, Princeton University  
and

Kaizheng Wang

Department of IEOR, Columbia University

February 20, 2022

Section A introduces a variant of Algorithm 4. Section B outlines the deferred results for distributed linear regression. Section C shows the CEASE iterates for distributed logistic regression. Section D presents the numerical results on the distributed  $\ell_1$ -regularized logistic regression. Section E presents the proofs of the main results. Section F lists a few technical lemmas that are used throughout the proofs.

## A A variant of Algorithm 4

As mentioned in the main text, Algorithm 2 is unstable when the local sample size  $n$  is not sufficiently large. The proximal gradient method Algorithm A.1 was introduced to

---

\*We gratefully acknowledge *NSF grants DMS-1662139 and DMS-1712591, NIH grant 2R01-GM072611-14, and ONR grant N00014-19-1-2120*. We acknowledge computing resources from Columbia University’s Shared Research Computing Facility project, which is supported by *NIH Research Facility Improvement Grant 1G20-RR030893-01*, and associated funds from the New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract C090171, both awarded April 15, 2010.

---

**Algorithm A.1** Distributed estimation using gradient-enhanced loss (small step-sizes)

---

1: **Input:** Initial value  $\boldsymbol{\theta}_0$ , number of iterations  $T$ , step-sizes  $\{\alpha_t\}_{t=0}^{T-1}$ .

2: **For**  $t = 0, 1, 2, \dots, T-1$ :

- Each machine evaluates  $\nabla f_k(\boldsymbol{\theta}_t)$  and sends to the central processor;
- The central processor computes  $\nabla f(\boldsymbol{\theta}_t) = \frac{1}{m} \sum_{k=1}^m \nabla f_k(\boldsymbol{\theta}_t)$  and broadcasts to machines;
- Each machine computes

$$\boldsymbol{\theta}_{t,k} = \operatorname{argmin}_{\boldsymbol{\theta}} \{f_k(\boldsymbol{\theta}) + g(\boldsymbol{\theta}) - \langle \nabla f_k(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t), \boldsymbol{\theta} \rangle\}$$

and sends to the central processor;

- The central processor computes  $\boldsymbol{\theta}_{t+1} = (1 - \alpha_t) \boldsymbol{\theta}_t + \frac{\alpha_t}{m} \sum_{k=1}^m \boldsymbol{\theta}_{t,k}$  and broadcasts to machines.

3: **Output:**  $\boldsymbol{\theta}_T$ .

---

stabilize the solution path by shrinking towards the solution in the previous step. A variant of CEASE that stabilizes Algorithm 2 is to take smaller step-sizes, which we now present. The idea is applicable to stabilize Algorithm 1 too, resulting a variant to Algorithm 3.

The only difference between Algorithms 2 and A.1 lies in the aggregation step. From  $\boldsymbol{\theta}_t$  the former directly jumps to the average of new individual estimators  $\{\boldsymbol{\theta}_{t,k}\}_{k=1}^m$ , while the latter proceeds more cautiously in that direction. Algorithm 2 is a special case of Algorithm A.1 with  $\alpha_t = 1$  for all  $t$ . Choosing  $\alpha_t \in (0, 1)$  helps stabilize the iterates especially when  $n$  is small.

Algorithm A.1 is conceptually simple and easy to implement. To see its performance, we conduct distributed logistic regression on the first set of synthetic data in Section 5.1. The numerical results there show that Algorithm 2 fails to converge when  $(n, m) = (250, 40)$ . We run Algorithm A.1 with constant step-size  $\alpha_0 = \alpha_1 = \dots$  under this setting, where  $\alpha_0 \in \{1, 1/2, 1/4, 1/8\}$ . We also run CEASE for comparison. Figure 1 summarizes all the results. Again, the curves show the average values over 100 independent runs; the

error bands correspond to one standard deviation. In this experiment, the performance of Algorithm A.1 with  $\alpha_t = 1/4$  is similar to that of CEASE;  $\alpha_t = 1/2$  leads to even faster convergence; and  $\alpha_t = 1/8$  slows it down. For simplicity, we take the step-size to be constant over time. It would be interesting to explore decaying schemes such as  $\alpha_t \asymp t^{-\beta}$  for some  $\beta > 0$ .

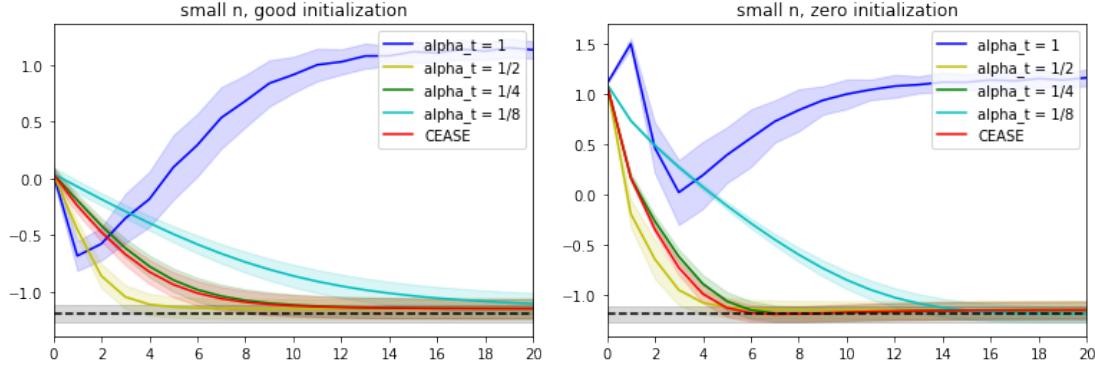


Figure 1: Impacts of step-size  $\alpha_t$  on convergence. The  $x$ -axis and  $y$ -axis are the number of iterations and  $\log \|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|_2$ . The dashed lines show the error of the minimizer of the overall loss function. The left and right plots use  $\bar{\boldsymbol{\theta}}$  and  $\mathbf{0}$  for initialization.  $\alpha_t = 1, 1/2, 1/4$  and  $1/8$  refer to Algorithm A.1 with the corresponding step-sizes. In particular,  $\alpha_t = 1$  is equivalent to Algorithm 2. CEASE refers to Algorithm 4 with  $\alpha = 0.15p/n$ .

The new algorithm also has some drawbacks. For instance, it implicitly assumes that the local sample size is large enough so that each local machine can solve its optimization problems reliably. Consider the example of the distributed logistic regression with  $N$  samples and  $p$  variables in total. When  $m$  is large so that  $n = N/m < p$ , the loss functions  $\{f_k\}_{k=1}^m$  on local machines are no longer strongly convex and thus do not have unique minima. Then Algorithm A.1 needs to be modified because

$$\boldsymbol{\theta}_{t,k} = \operatorname{argmin}_{\boldsymbol{\theta}} \{f_k(\boldsymbol{\theta}) + g(\boldsymbol{\theta}) - \langle \nabla f_k(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t), \boldsymbol{\theta} \rangle\}$$

is not uniquely defined. We need some pivoting rule to choose one optimal solution, such as the one with the minimum Euclidean norm. That complicates the algorithm and it is

not clear how to establish theoretical guarantees then. In contrast, the quadratic proximity term in CEASE always make the objective function strongly convex. It ensures the uniqueness of  $\{\boldsymbol{\theta}_{t,k}\}_{k=1}^m$  and facilitates computation.

## B Distributed linear regression

In distributed linear regression, recall that the  $k^{th}$  machine defines a quadratic loss function

$$\frac{1}{2n} \sum_{i \in \mathcal{I}_k} (y_i - \mathbf{x}_i^\top \boldsymbol{\theta})^2 = \frac{1}{2} \boldsymbol{\theta}^\top \widehat{\boldsymbol{\Sigma}}_k \boldsymbol{\theta} - \widehat{\mathbf{w}}_k^\top \boldsymbol{\theta} + \frac{1}{2n} \sum_{i \in \mathcal{I}_k} y_i^2,$$

where  $\widehat{\boldsymbol{\Sigma}}_k = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{x}_i \mathbf{x}_i^\top$  and  $\widehat{\mathbf{w}}_k = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{x}_i y_i$ . Let  $f(\boldsymbol{\theta}) = \frac{1}{2} \boldsymbol{\theta}^\top \widehat{\boldsymbol{\Sigma}} \boldsymbol{\theta} - \widehat{\mathbf{w}}^\top \boldsymbol{\theta}$ . Without loss of generality we write  $\mathbf{x}_i = (1, \mathbf{u}_i^\top)^\top \in \mathbb{R}^p$ .

**Assumption B.1.** •  $\mathbb{E} \mathbf{u}_i = 0$  and  $\mathbb{E}(\mathbf{u}_i \mathbf{u}_i^\top) = \boldsymbol{\Sigma} \succ 0$ .  $\{\boldsymbol{\Sigma}^{-1/2} \mathbf{u}_i\}_{i=1}^N$  are i.i.d. sub-Gaussian random vectors with bounded  $\|\boldsymbol{\Sigma}^{-1/2} \mathbf{u}_i\|_{\psi_2}$ .

- The minimum eigenvalue  $\lambda_{\min}(\boldsymbol{\Sigma})$  is bounded away from zero.
- $N/\text{Tr}(\boldsymbol{\Sigma}) \geq C > 0$  and  $n/\log m \geq c > 0$  where  $C$  and  $c$  are constants.

For the least-squares, Algorithm 3 admits a close-form:

$$\boldsymbol{\theta}_{t+1} = [\mathbf{I} - (\widehat{\boldsymbol{\Sigma}}_1 + \alpha \mathbf{I})^{-1} \widehat{\boldsymbol{\Sigma}}] \boldsymbol{\theta}_t + (\widehat{\boldsymbol{\Sigma}}_1 + \alpha \mathbf{I})^{-1} \widehat{\mathbf{w}},$$

and so does Algorithm 4:

$$\begin{aligned} \boldsymbol{\theta}_{t+1,k} &= [\mathbf{I} - (\widehat{\boldsymbol{\Sigma}}_k + \alpha \mathbf{I})^{-1} \widehat{\boldsymbol{\Sigma}}] \boldsymbol{\theta}_t + (\widehat{\boldsymbol{\Sigma}}_k + \alpha \mathbf{I})^{-1} \widehat{\mathbf{w}}, \\ \boldsymbol{\theta}_{t+1} &= \left( \mathbf{I} - \frac{1}{m} \sum_{k=1}^m (\widehat{\boldsymbol{\Sigma}}_k + \alpha \mathbf{I})^{-1} \widehat{\boldsymbol{\Sigma}} \right) \boldsymbol{\theta}_t + \frac{1}{m} \sum_{k=1}^m (\widehat{\boldsymbol{\Sigma}}_k + \alpha \mathbf{I})^{-1} \widehat{\mathbf{w}}. \end{aligned}$$

Intuitively, the averaging step in Algorithm 4 reduces variance and accelerates convergence.

Below we study Algorithm 4 with the help of these analytical expressions. In the large sample regime, we achieve a contraction factor of  $O(p/n)$  without any condition number; in the general regime, linear convergence is still guaranteed.

**Theorem B.1.** *Suppose Assumption B.1 holds and  $n/p$  is bounded away from zero. Then, there exist positive constants  $C_1, C_2$  and  $C_3$  such that when (i)  $n \geq C_1 p$  and  $\alpha \geq 0$  or (ii)  $\alpha \geq C_1 \text{Tr}(\Sigma)/n$ , with probability tending to 1,*

$$\|\boldsymbol{\theta}_t - \hat{\boldsymbol{\theta}}\|_2 \leq 2\sqrt{\kappa} \eta^t \|\boldsymbol{\theta}_0 - \hat{\boldsymbol{\theta}}\|_2, \quad \forall t \geq 0, \quad (\text{B.1})$$

where  $\kappa = \lambda_{\max}(\Sigma)/\lambda_{\min}(\Sigma)$  and  $\eta = 1 - \frac{1 - \min\{1/2, C_2 p/n\}}{1 + C_3 \alpha}$ .

Theorem B.1 reveals the following remarkable facts about Algorithm 4: No matter what relationship  $n$  and  $p$  have, proper regularization always guarantees linear convergence, and the rate exhibits a smooth transition as  $p/n$  grows. Hence we can handle the distributed statistical estimation problem without assuming large enough  $n$ , overcoming the difficulty of other algorithms in literature (Zhang et al., 2013; Battey et al., 2018; Jordan et al., 2019).

If  $n/p$  is large enough, the regularization is not necessary, but choosing  $\alpha \asymp p/n$  does not hurt much. This is because we can control the contraction factor as:

$$1 - \frac{1 - C_2 p/n}{1 + C_3 \alpha} = \frac{C_3 \alpha + C_2 p/n}{1 + C_3 \alpha} = O(p/n).$$

When  $n/p$  is not that large, most distributed statistical estimation procedures fail. By choosing  $\alpha = \tilde{C} \text{Tr}(\Sigma)/n$  for  $\tilde{C} > C_1$  (see Condition (ii) of Theorem B.1) we still have

linear convergence with contraction factor at most

$$1 - \frac{1 - 1/2}{1 + C_3\alpha} = 1 - \frac{1}{2 + 2C_3\tilde{C}\text{Tr}(\boldsymbol{\Sigma})/n} < 1.$$

In most situations of interest we have  $\text{Tr}(\boldsymbol{\Sigma}) \asymp p$  (even for pervasive factor models). Therefore we see that  $\alpha \asymp p/n$  is a universal and adaptive choice of regularization over all the possible relation between  $n$  and  $p$ .

Another benefit of the Algorithms is that the condition number  $\kappa$  has only logarithmic effect on the iteration complexity, and the contraction factor in Theorem B.1 does not depend on  $\kappa$  at all. This is in stark contrast to the analysis under the same setting in Shamir et al. (2014), and helps relax the commonly used boundedness assumption on the condition number in Zhang et al. (2013), Battey et al. (2018), Jordan et al. (2019), among others. It is worth mentioning that Wang et al. (2018) derive similar results for distributed linear regression when the local sample size  $n$  is sufficiently large.

## C Distributed logistic regression

In this section, we demonstrate the iterates of CEASE in distributed logistic regression. For illustration purposes, we use Newton's method on each local machine to solve the optimization problem. It is worth pointing out that Newton's method is not the only choice. First-order methods such gradient descent can also efficiently do the job.

Given samples  $\{(\mathbf{x}_i, y_i)\}_{i \in [N]}$ , the loss on the  $k$ th machine is

$$f_k(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \left[ \log(1 + e^{\boldsymbol{\theta}^\top \mathbf{x}_i}) - y_i(\boldsymbol{\theta}^\top \mathbf{x}_i) \right].$$

Thus for Algorithm 4, we have for each iteration

$$\begin{aligned}\boldsymbol{\theta}_{t,k} &= \operatorname{argmin}_{\boldsymbol{\theta}} \left\{ f_k(\boldsymbol{\theta}) - \langle \nabla f_k(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t), \boldsymbol{\theta} \rangle + \frac{\alpha}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_2^2 \right\} \\ &= \operatorname{argmin}_{\boldsymbol{\theta}} \left\{ \frac{1}{n} \sum_{i \in \mathcal{I}_k} \log(1 + e^{\boldsymbol{\theta}^\top \mathbf{x}_i}) - \langle \widehat{\mathbf{w}}_k + \boldsymbol{\eta}_{t,k} + \alpha \boldsymbol{\theta}_t, \boldsymbol{\theta} \rangle + \frac{\alpha}{2} \|\boldsymbol{\theta}\|_2^2 \right\}.\end{aligned}$$

Here  $\widehat{\mathbf{w}}_k = \frac{1}{m} \sum_{i \in \mathcal{I}_k} y_i \mathbf{x}_i$ ,  $\boldsymbol{\eta}_{t,k} = \nabla f_k(\boldsymbol{\theta}_t) - \nabla f(\boldsymbol{\theta}_t)$ . Assume that we start with  $\boldsymbol{\theta}_{t,k}^{(0)} = \boldsymbol{\theta}_t$ .

Then one Newton iteration of the above minimization problem is

$$\boldsymbol{\theta}_{t,k}^{(j+1)} = \boldsymbol{\theta}_{t,k}^{(j)} - \left( \frac{1}{n} \mathbf{X}_k^\top \mathbf{W}_k^{(j)} \mathbf{X}_k + \alpha \mathbf{I} \right)^{-1} \left[ \mathbf{X}_k^\top (\mathbf{p}_k^{(j)} - \mathbf{Y}_k) - \boldsymbol{\eta}_{t,k} + \alpha (\boldsymbol{\theta}_{t,k}^{(j)} - \boldsymbol{\theta}_t) \right].$$

Here,  $\mathbf{X}_k \in \mathbb{R}^{n \times p}$  is a matrix with rows consisting of  $\{\mathbf{x}_i\}_{i \in \mathcal{I}_k}$ ,  $\mathbf{Y}_k \in \mathbb{R}^n$  denotes the vector of  $\{y_i\}_{i \in \mathcal{I}_k}$ ,  $\mathbf{p}_k^{(j)} \in \mathbb{R}^n$  is the vector of fitted probabilities with entry  $i$  equal to  $e^{(\boldsymbol{\theta}_{t,k}^{(j)})^\top \mathbf{x}_i} / (1 + e^{(\boldsymbol{\theta}_{t,k}^{(j)})^\top \mathbf{x}_i})$ , and  $\mathbf{W}_k^{(j)} \in \mathbb{R}^{n \times n}$  is a diagonal matrix with entries  $\{e^{(\boldsymbol{\theta}_{t,k}^{(j)})^\top \mathbf{x}_i} / (1 + e^{(\boldsymbol{\theta}_{t,k}^{(j)})^\top \mathbf{x}_i})^2\}_{i \in \mathcal{I}_k}$ . Running one-step or multi-step Newton's iteration will result in some  $\boldsymbol{\theta}_{t,k}^{(j)}$  sufficiently close to  $\boldsymbol{\theta}_{t,k}$ . After that, we simply average them across all the machines and obtain  $\boldsymbol{\theta}_{t+1}$ .

For Algorithm 2, in each iteration, we simply set  $\alpha = 0$  in the above procedure. For Algorithm 3 and Algorithm 1, we simply omit the averaging step in Algorithm 4 and Algorithm 2 respectively.

## D Numerical results on distributed $\ell_1$ -regularized logistic regression

In this section, we present the performance of our algorithms in the distributed  $\ell_1$ -regularized logistic regression problem where nonsmooth penalty is present. We fix the total sample size  $N = 5000$  and the dimensionality  $p = 1001$ , and generate the i.i.d. data  $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$

as follows:  $\mathbf{x}_i = (1, \mathbf{u}_i^\top)^\top$  with  $\mathbf{u}_i \sim N(\mathbf{0}_{p-1}, \mathbf{I}_{p-1})$ ;  $\mathbb{P}(y_i = 1) = 1 - \mathbb{P}(y_i = 0) = \frac{e^{\mathbf{x}_i^\top \boldsymbol{\theta}^*}}{1 + e^{\mathbf{x}_i^\top \boldsymbol{\theta}^*}}$  where  $\boldsymbol{\theta}^* = (\mathbf{1}_{10}^\top, \mathbf{0}_{991}^\top)^\top / \sqrt{2} \in \mathbb{R}^p$ . We define the penalty function  $g(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1$  with  $\lambda = 0.5 \sqrt{\frac{\log p}{N}}$ , such that the regularized MLE over the whole dataset recovers the nonzeros of  $\boldsymbol{\theta}^*$  accurately. Figure 2 shows the performance of CEASE algorithms and ADMM, where “large  $n$ ”, “moderate  $n$ ” and “small  $n$ ” refer to  $(n, m) = (1000, 5)$ ,  $(500, 10)$  and  $(250, 20)$ , and “zero initialization” and “good initialization” refer to  $\boldsymbol{\theta}_0 = \mathbf{0}$  and  $\bar{\boldsymbol{\theta}}$ , respectively. Again,  $\bar{\boldsymbol{\theta}}$  is the one-shot distributed estimator (Zhang et al., 2013). All the results are average values of 100 independent runs.

Similar to the distributed logistic regression case, the CEASE algorithms with proper regularization (Algorithms 3 and 4) work well in general; without regularization, the CEASE algorithm fails to converge when the local sample size  $n$  is small and the initialization is uninformative. For this nonsmooth problem, the CEASE algorithm with averaging (Algorithm 4) does not seem to have advantage over the single version (Algorithm 3). The ADMM converges quickly to a region near the minimizer but then proceeds quite slowly, which appears to be a common phenomenon in many distributed optimization problems (Boyd et al., 2011).

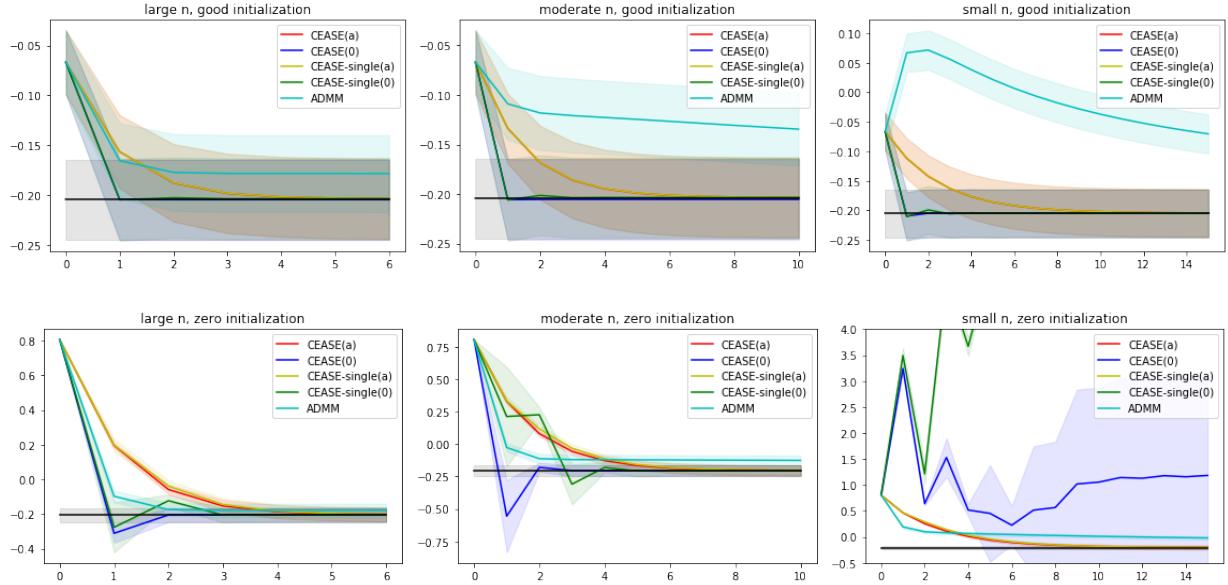


Figure 2: Nonsmooth minimization. The  $x$ -axis and  $y$ -axis are the number of iterations and  $\log \|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|_2$ . The dashed lines show the error of the minimizer of the overall regularized loss function. The top and bottom panels use  $\bar{\boldsymbol{\theta}}$  and  $\mathbf{0}$  for initialization, respectively. CEASE(a) and CEASE(0) refer to Algorithm 4 with  $\alpha = 0.05p/n$  and 0; CEASE-single(a) and CEASE-single(0) refer to Algorithm 3 with  $\alpha = 0.05p/n$  and 0, respectively. In particular, CEASE-single(0) is equivalent to the CSL algorithm in (Jordan et al., 2019).

## E Proofs

### E.1 Proof of Theorem 3.1

**Lemma E.1.** *Let Assumptions 3.1 and 3.2 hold. Consider the iterates  $\{\boldsymbol{\theta}_t\}_{t=0}^T$  generated by Algorithm 4. Define*

$$\gamma_t = \begin{cases} \frac{\delta}{\rho_0 + \alpha} \cdot \min\left\{1, \frac{\delta}{\rho_0 + \alpha} \left(1 + \frac{M}{\rho_0 + \alpha} \|\boldsymbol{\theta}_t - \hat{\boldsymbol{\theta}}\|_2\right)\right\} & , \text{ if Assumption 3.3 holds} \\ \frac{\delta}{\rho_0 + \alpha} & , \text{ otherwise} \end{cases}, \quad 0 \leq t \leq T-1.$$

*If  $0 < \|\boldsymbol{\theta}_t - \hat{\boldsymbol{\theta}}\|_2 < R/2$ ,  $\delta < \rho_0 + \alpha$  and  $\gamma_t^2 < \rho/(\rho + 2\alpha)$ , then*

$$\frac{\|\boldsymbol{\theta}_{t+1} - \hat{\boldsymbol{\theta}}\|_2}{\|\boldsymbol{\theta}_t - \hat{\boldsymbol{\theta}}\|_2} \leq \frac{\gamma_t \sqrt{\rho^2 + 2\alpha\rho} + \alpha}{\rho + \alpha} < 1.$$

Theorem 3.1 directly follows from Lemma E.1 and induction. Below we only prove Lemma E.1 with Assumption 3.3. The other part in Lemma E.1 without Assumption 3.3 can be derived by slightly modifying this proof.

**Proof of Lemma E.1 with Assumption 3.3.** Let  $\boldsymbol{\theta}_t^+ = \text{prox}_{\alpha^{-1}(f+g)}(\boldsymbol{\theta}_t)$ . By the triangle inequality,

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t^+\|_2 + \|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2. \quad (\text{E.1})$$

We first invoke Theorem 3.2 to bound the first term  $\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t^+\|_2$  in (E.1). Define  $\tilde{g}(\boldsymbol{\theta}) = g(\boldsymbol{\theta}) + (\alpha/2)\|\boldsymbol{\theta} - \boldsymbol{\theta}_t\|_2^2$  for  $\boldsymbol{\theta} \in \mathbb{R}^p$ . Then  $\boldsymbol{\theta}_{t+1}$  is the first iterate of Algorithm 2 initialized at  $\boldsymbol{\theta}_t$  for computing  $\boldsymbol{\theta}_t^+ = \text{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^p} \left\{ \frac{1}{m} \sum_{k=1}^m f_k(\boldsymbol{\theta}) + \tilde{g}(\boldsymbol{\theta}) \right\}$ .

From  $\widehat{\boldsymbol{\theta}} = \text{prox}_{\alpha^{-1}(f+g)}(\widehat{\boldsymbol{\theta}})$  and Lemma F.3 we obtain that  $\|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2$ . Then the condition  $\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 < R/2$  leads to  $B(\boldsymbol{\theta}_t^+, R/2) \subseteq B(\widehat{\boldsymbol{\theta}}, R)$ . By Assumptions 3.1 and 3.2,

- in  $B(\boldsymbol{\theta}_t^+, R/2)$ ,  $\{f_k + \tilde{g}\}_{k=1}^m$  are  $(\rho_0 + \alpha)$ -strongly convex and  $(f + \tilde{g})$  is  $(\rho + \alpha)$ -strongly convex;
- $\|\nabla^2 f_k(\boldsymbol{\theta}) - \nabla^2 f(\boldsymbol{\theta})\|_2 \leq \delta$  holds for all  $k \in [m]$  and  $\boldsymbol{\theta} \in B(\boldsymbol{\theta}_t^+, R/2)$ .

Furthermore, Lemma F.3 also yields

$$\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_t^+\|_2^2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2^2 - \|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2^2 = \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2^2 \left(1 - \|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2^2 / \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2^2\right). \quad (\text{E.2})$$

Then  $\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_t^+\|_2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 < R/2$  and  $\tilde{\boldsymbol{\theta}}_0 \in B(\boldsymbol{\theta}_t^+, R/2)$ . Based on these conditions and

$\alpha + \rho_0 > \delta$ , we use Theorem 3.2 to get

$$\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t^+\|_2 \leq \gamma_t \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_t^+\|_2.$$

From here, (E.1) and (E.2) we obtain that

$$\begin{aligned} \|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 &\leq \gamma_t \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_t^+\|_2 + \|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2 \\ &\leq \gamma_t \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \left(1 - \|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2^2 / \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2^2\right)^{1/2} + \|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2 \\ &= \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \cdot h(\|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2 / \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2), \end{aligned}$$

where  $h(x) = \gamma_t \sqrt{1-x^2} + x$ ,  $\forall x \in [0, 1]$ . From  $h'(x) = 1 - \gamma_t x / \sqrt{1-x^2}$  we see that  $h' \geq 0$  on  $[0, 1/\sqrt{1+\gamma_t^2}]$ .

On the one hand, Lemma F.3 asserts that  $\|\boldsymbol{\theta}^+ - \widehat{\boldsymbol{\theta}}\|_2 / \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \leq \alpha / (\rho + \alpha)$ . On the other hand, the assumption  $\gamma_t^2 < \rho / (\rho + 2\alpha)$  forces

$$\frac{1}{\sqrt{1+\gamma_t^2}} > \frac{1}{\sqrt{1+\rho/(\rho+2\alpha)}} = \frac{\sqrt{\rho/2+\alpha}}{\sqrt{\rho+\alpha}} \geq \frac{\rho/2+\alpha}{\rho+\alpha} \geq \alpha / (\rho + \alpha).$$

The proof is completed by computation:

$$\begin{aligned} \frac{\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2}{\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2} &\leq h\left(\frac{\|\boldsymbol{\theta}_t^+ - \widehat{\boldsymbol{\theta}}\|_2}{\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2}\right) \leq h\left(\frac{\alpha}{\rho+\alpha}\right) = \gamma_t \left[1 - \left(\frac{\alpha}{\rho+\alpha}\right)^2\right]^{1/2} + \frac{\alpha}{\rho+\alpha} \\ &= \frac{\gamma_t \sqrt{\rho^2 + 2\rho\alpha} + \alpha}{\rho+\alpha} < \frac{\sqrt{[\rho/(\rho+2\alpha)] \cdot (\rho^2 + 2\rho\alpha)} + \alpha}{\rho+\alpha} = 1, \end{aligned}$$

where we used the assumption  $\gamma_t^2 < \rho / (\rho + 2\alpha)$  again.  $\square$

## E.2 Proof of Corollary 3.1

We claim that  $(\frac{\delta}{\rho_0+\alpha})^2 \leq \frac{7}{9} \cdot \frac{\rho}{\rho+2\alpha}$ . Given this, Corollary 3.1 follows from Theorem 3.1 and

$$\frac{\frac{\delta}{\rho_0+\alpha} \sqrt{\rho^2 + 2\alpha\rho} + \alpha}{\rho + \alpha} \leq \frac{\sqrt{\frac{7}{9} \cdot \frac{\rho}{\rho+2\alpha} \cdot \rho(\rho+2\alpha)} + \alpha}{\rho + \alpha} = 1 - \frac{(1 - \sqrt{7/9})\rho}{\rho + \alpha} \leq 1 - \frac{\rho/10}{\rho + \alpha}.$$

The claim trivially holds if  $\delta = 0$ . When  $\delta > 0$ , let us first assume  $0 < \delta \leq \rho$  and define  $b = \alpha\rho/\delta^2$ . Then  $\alpha = b\delta^2/\rho$ ,  $b \geq 4$  and  $\rho_0 \geq \max\{\rho - \delta, 0\}$  force

$$\rho_0 + \alpha \geq \rho - \delta + b\delta^2/\rho = (\rho/\delta + b\delta/\rho - 1)\delta \geq (2\sqrt{(\rho/\delta) \cdot (b\delta/\rho)} - 1)\delta = (2\sqrt{b} - 1)\delta > \delta.$$

and  $[\delta/(\rho_0 + \alpha)]^2 \leq [\delta/(\rho_0 + \alpha)]^2 \leq 1/h_1(\delta/\rho)$ , where  $h_1(x) = (bx + x^{-1} - 1)^2$ . On the other hand,  $\rho/(\rho + 2\alpha) = 1/(1 + 2\alpha/\rho) = 1/h_2(\delta/\rho)$ , where  $h_2(x) = 1 + 2bx^2$ .

We are going to show  $h_2(x) \leq 7h_1(x)/9$ ,  $\forall x \in (0, 1]$ , which leads to the desired result under  $0 < \delta \leq \rho$ . If  $0 < x \leq \sqrt{3}/2$ , then  $h_1(x) \geq (2\sqrt{b} - 1)^2 \geq (2\sqrt{b} - \sqrt{b}/2)^2 \geq 9b/4$  and

$$h_2(x) \leq 1 + 2b \cdot (3/4) \leq (b/4) + (6b/4) = 7b/4 \leq 7h_1(x)/9.$$

If  $\sqrt{3}/2 < x \leq 1$ , then  $h_1(x) \geq b^2x^2 \geq 3b^2/4$ ,  $h_2(x) \leq 1 + 2b \leq (b/4) + 2b = 9b/4$ , and  $h_2(x)/h_1(x) = 3/b \leq 3/4 \leq 7/9$ .

Suppose now that  $\delta > \rho$ , and define  $b = \alpha\rho/\delta^2$ . Then

$$\begin{aligned} \left(\frac{\delta}{\rho_0 + \alpha}\right)^2 &\leq \left(\frac{\delta}{\alpha}\right)^2 = \left(\frac{1}{b\delta/\rho}\right)^2 = \frac{1}{b^2(\delta/\rho)^2}, \\ \frac{\rho}{\rho + 2\alpha} &= \frac{1}{1 + 2\alpha/\rho} = \frac{1}{1 + 2b(\delta/\rho)^2}. \end{aligned}$$

From  $b \geq 4$  and  $\delta/\rho > 1$  we get  $(\frac{\delta}{\rho_0+\alpha})^2 \leq \frac{7}{9} \cdot \frac{\rho}{\rho+2\alpha}$  from

$$\begin{aligned} b^2(\delta/\rho)^2 - \frac{9}{7}[1 + 2b(\delta/\rho)^2] &= (\delta/\rho)^2 b(b - 18/7) - 9/7 \\ &\geq 1 \cdot 4 \cdot (4^1 - 18/7) - 9/7 = 31/7 > 0. \end{aligned}$$

### E.3 Proof of Corollary 3.2

Throughout the proof we assume that  $\delta/\rho$  is sufficiently small. The regularity conditions in Theorem 3.1 are easily verified as  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R/2)$  and  $[\delta/(\rho_0 + \alpha)]^2 < \rho/(\rho + 2\alpha)$ . Here we used the fact  $\rho_0 \geq \rho - \delta$ .

From  $\rho_0 \geq \rho - \delta$  we get  $\rho_0 + \alpha \geq \rho_0 \geq \rho/2$  and  $\delta/(\rho_0 + \alpha) \leq 2\delta/\rho$ . Also,  $\sqrt{\rho^2 + 2\alpha\rho} \leq \rho\sqrt{1 + 2C(\delta/\rho)^2} \lesssim \rho$ . We control the contraction factor in (3.2):

$$\frac{\frac{\delta}{\rho_0+\alpha}\sqrt{\rho^2+2\alpha\rho}+\alpha}{\rho+\alpha} \lesssim \frac{(2\delta/\rho)\rho + C\delta^2/\rho}{\rho} = \frac{2\delta}{\rho} + \frac{C}{4} \left(\frac{2\delta}{\rho}\right)^2 \lesssim \frac{\delta}{\rho}.$$

Recall that  $\gamma_t = \frac{\delta}{\rho_0+\alpha} \cdot \min\{1, \frac{\delta}{\rho_0+\alpha}(1 + \frac{M}{\rho_0+\alpha}\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2)\}$  in Theorem 3.1. When  $\delta/\rho$  is small and  $\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \leq \rho/M$ , we have  $\rho_0 + \alpha \geq \rho/2$ ,  $\frac{M}{\rho_0+\alpha}\|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \leq 2$ , and  $\gamma_t \leq (2\delta/\rho)^2$ .

This help bound the contraction factor in (3.2):

$$\frac{\gamma_t \cdot \sqrt{\rho^2 + 2\alpha\rho} + \alpha}{\rho + \alpha} \lesssim \frac{(2\delta/\rho)^2\rho + C\delta^2/\rho}{\rho} \lesssim \left(\frac{\delta}{\rho}\right)^2.$$

### E.4 Proof of Theorem 3.2

Theorem 3.2 is a direct summary of the following two lemmas.

**Lemma E.2** (Contraction). *Let Assumptions 3.1 and 3.2 hold, with  $\rho_0 > \delta \geq 0$ . Then*

$$\|\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}\|_2 \leq (\delta/\rho_0)\|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2, \quad \forall \boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R), \quad \forall k \in [m].$$

*Proof.* Fix  $\boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R)$ . By the first order condition of  $\varphi_k(\boldsymbol{\theta})$ , we have that

$$\nabla f_k(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta}) \in \partial\{f_k[\varphi_k(\boldsymbol{\theta})] + g[\varphi_k(\boldsymbol{\theta})]\}. \quad (\text{E.3})$$

Using the fixed point property  $\varphi_k(\widehat{\boldsymbol{\theta}}) = \widehat{\boldsymbol{\theta}}$ , we have  $\nabla f_k(\widehat{\boldsymbol{\theta}}) - \nabla f(\widehat{\boldsymbol{\theta}}) \in \partial[f_k(\widehat{\boldsymbol{\theta}}) + g(\widehat{\boldsymbol{\theta}})]$ . By the Taylor expansion and Assumption 3.2,

$$\begin{aligned} & \|[\nabla f_k(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta})] - [\nabla f_k(\widehat{\boldsymbol{\theta}}) - \nabla f(\widehat{\boldsymbol{\theta}})]\|_2 \\ &= \left\| \int_0^1 \left( \nabla^2 f_k[(1-t)\widehat{\boldsymbol{\theta}} + t\boldsymbol{\theta}] - \nabla^2 f[(1-t)\widehat{\boldsymbol{\theta}} + t\boldsymbol{\theta}] \right) (\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}) dt \right\|_2 \\ &\leq \sup_{\boldsymbol{\zeta} \in B(\widehat{\boldsymbol{\theta}}, R)} \|\nabla^2 f_k(\boldsymbol{\zeta}) - \nabla^2 f(\boldsymbol{\zeta})\|_2 \cdot \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2 \\ &\leq \delta \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2 < \rho_0 R. \end{aligned}$$

From this, (E.3) and Lemma F.2, we obtain that  $\|\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}\|_2 \leq (\delta/\rho_0) \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2$ .

□

**Lemma E.3** (Averaging). *Let Assumptions 3.1, 3.2 and 3.3 hold, with  $\rho_0 > \delta \geq 0$ . We have*

$$\left\| \frac{1}{m} \sum_{k=1}^m \varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}} \right\|_2 \leq \frac{\delta^2}{\rho_0 \rho} (1 + M\rho_0^{-1} \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2) \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2, \quad \forall \boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R).$$

*Proof.* Define  $L_k(\boldsymbol{\theta}) = f_k(\boldsymbol{\theta}) + g(\boldsymbol{\theta})$  and  $L(\boldsymbol{\theta}) = f(\boldsymbol{\theta}) + g(\boldsymbol{\theta})$  for  $\boldsymbol{\theta} \in \mathbb{R}^p$ . Then  $\widehat{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\xi} \in \mathbb{R}^p} L(\boldsymbol{\xi})$  and  $\varphi_k(\boldsymbol{\theta}) = \operatorname{argmin}_{\boldsymbol{\xi} \in \mathbb{R}^p} \{L_k(\boldsymbol{\xi}) - \langle \nabla L_k(\boldsymbol{\theta}) - \nabla L(\boldsymbol{\theta}), \boldsymbol{\xi} \rangle\}$ . By the optimality conditions,

$$\nabla L_k[\varphi_k(\boldsymbol{\theta})] - \nabla L_k(\boldsymbol{\theta}) + \nabla L(\boldsymbol{\theta}) = \mathbf{0} = \nabla L(\widehat{\boldsymbol{\theta}}).$$

After subtracting  $\nabla L_k(\widehat{\boldsymbol{\theta}})$  from both sides and rearranging terms, we get

$$\nabla L_k[\varphi_k(\boldsymbol{\theta})] - \nabla L_k(\widehat{\boldsymbol{\theta}}) = [\nabla L_k(\boldsymbol{\theta}) - \nabla L_k(\widehat{\boldsymbol{\theta}})] - [\nabla L(\boldsymbol{\theta}) - \nabla L(\widehat{\boldsymbol{\theta}})].$$

Note that the average of the right hand side over  $k \in [m]$  is  $\mathbf{0}$ .

Define  $\mathbf{H}_k = \int_0^1 \nabla^2 L_k[(1-t)\widehat{\boldsymbol{\theta}} + t\varphi_k(\boldsymbol{\theta})] dt$  for  $k \in [m]$  and  $\widehat{\mathbf{H}} = \nabla^2 L(\widehat{\boldsymbol{\theta}})$ . Then

$$\nabla L_k[\varphi_k(\boldsymbol{\theta})] - \nabla L_k(\widehat{\boldsymbol{\theta}}) = \mathbf{H}_k(\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}) = \widehat{\mathbf{H}}(\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}) + (\mathbf{H}_k - \widehat{\mathbf{H}})(\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}),$$

$$\mathbf{0} = \frac{1}{m} \sum_{k=1}^m \left( \nabla L_k[\varphi_k(\boldsymbol{\theta})] - \nabla L_k(\widehat{\boldsymbol{\theta}}) \right) = \widehat{\mathbf{H}}[\bar{\varphi}(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}] + \frac{1}{m} \sum_{k=1}^m (\mathbf{H}_k - \widehat{\mathbf{H}})(\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}),$$

where we let  $\bar{\varphi}(\boldsymbol{\theta}) = \frac{1}{m} \sum_{k=1}^m \varphi_k(\boldsymbol{\theta})$ . As a result,

$$\begin{aligned} \|\bar{\varphi}(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}\|_2 &= \left\| \frac{1}{m} \sum_{k=1}^m \widehat{\mathbf{H}}^{-1}(\mathbf{H}_k - \widehat{\mathbf{H}})(\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}) \right\|_2 \\ &\leq \|\widehat{\mathbf{H}}^{-1}\|_2 \max_{k \in [m]} \|\mathbf{H}_k - \widehat{\mathbf{H}}\|_2 \cdot \max_{k \in [m]} \|\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}\|_2. \end{aligned}$$

Lemma E.2 forces  $\max_{k \in [m]} \|\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}\|_2 \leq (\delta/\rho_0) \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2$ , and Assumption 3.1 yields  $\widehat{\mathbf{H}} \succeq \rho \mathbf{I}$  and  $\|\widehat{\mathbf{H}}\|_2 \leq 1/\rho$ . Furthermore, we use Assumptions 3.2 and 3.3 to get

$$\begin{aligned} \|\mathbf{H}_k - \widehat{\mathbf{H}}\|_2 &\leq \left\| \int_0^1 \left( \nabla^2 L_k[(1-t)\widehat{\boldsymbol{\theta}} + t\varphi_k(\boldsymbol{\theta})] - \nabla^2 L[(1-t)\widehat{\boldsymbol{\theta}} + t\varphi_k(\boldsymbol{\theta})] \right) dt \right\|_2 \\ &\quad + \left\| \int_0^1 \left( \nabla^2 L[(1-t)\widehat{\boldsymbol{\theta}} + t\varphi_k(\boldsymbol{\theta})] - \nabla^2 L(\widehat{\boldsymbol{\theta}}) \right) dt \right\|_2 \\ &\leq \delta + M \|\varphi_k(\boldsymbol{\theta}) - \widehat{\boldsymbol{\theta}}\|_2 \leq \delta + M(\delta/\rho_0) \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2. \end{aligned}$$

The proof is finished by combining all the estimates above.  $\square$

## E.5 Proof of Theorem 4.1

The proof is implied by combining proof of Corollary 3.2 with the results of the following two lemmas, the first of which is a direct counterpart of Theorem 3.1 in the stochastic setting, given an additional condition on similarity between local Hessians. The second lemma below specifies the order of Hessian difference in the generalized linear model, hence providing a contraction rate and guiding the choice of  $\alpha$ .

**Lemma E.4.** *Let Assumption 4.2 hold. Denote*

$$\widehat{\delta} := 2 \sup_{k \in [m]} \sup_{\boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R)} \|\nabla^2 f_k(\boldsymbol{\theta}) - \nabla^2 F(\boldsymbol{\theta})\|_2.$$

Consider the iterates  $\{\boldsymbol{\theta}_t\}_{t=0}^T$  generated by Algorithm 3 or Algorithm 4. Suppose that  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R/2)$  and  $[\widehat{\delta}/(\rho_0 + \alpha)]^2 < \rho/(\rho + 2\alpha)$ .

- For both Algorithms 3 and 4, we have

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \cdot \frac{\frac{\widehat{\delta}}{\rho_0 + \alpha} \cdot \sqrt{\rho^2 + 2\alpha\rho} + \alpha}{\rho + \alpha}, \quad 0 \leq t \leq T-1; \quad (\text{E.4})$$

- If in addition, Assumption 4.3 also holds, then for Algorithm 4 we have

$$\|\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2 \cdot \frac{\gamma_t \sqrt{\rho^2 + 2\alpha\rho} + \alpha}{\rho + \alpha}, \quad 0 \leq t \leq T-1, \quad (\text{E.5})$$

where we define  $\gamma_t = \frac{\widehat{\delta}}{\rho_0 + \alpha} \cdot \min\{1, \frac{\widehat{\delta}}{\rho + \alpha} (1 + \frac{M}{\rho_0 + \alpha} \|\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}\|_2)\}$ ;

- Both multiplicative factors in (E.4) and (E.5) are strictly less than 1.

**Proof of Lemma E.4.** We first assume that  $\rho_0 > \widehat{\delta}$  and analyze the vanilla DANE algorithm under the new assumptions. Let Assumption 4.2 hold, and  $\{\boldsymbol{\theta}_t\}_{t=0}^\infty$  be the iter-

ates with  $\boldsymbol{\theta}_0 \in B(\widehat{\boldsymbol{\theta}}, R)$ . For any  $\boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R)$ ,  $\|\nabla^2 f(\boldsymbol{\theta}) - \nabla^2 F(\boldsymbol{\theta})\|_2 \leq \widehat{\delta}/2$  and thus  $\|\nabla^2 f_k(\boldsymbol{\theta}) - \nabla^2 f(\boldsymbol{\theta})\|_2 \leq \widehat{\delta}$  for  $k \in [m]$ . Hence it implies Assumption 3.2 with  $\delta = \widehat{\delta}$ , and Lemma E.2 continues to hold. We can also get the result in Lemma E.3 under Assumption 4.3, by replacing  $\widehat{\mathbf{H}} = \nabla^2 f(\widehat{\boldsymbol{\theta}})$  in the proof of Lemma E.3 by  $\nabla^2 F(\widehat{\boldsymbol{\theta}}) + \nabla^2 g(\boldsymbol{\theta})$ . Then we drop the assumption  $\rho_0 > \widehat{\delta}$  can reproduce the results in Theorem 3.1 under the new setting, by following its original proof.  $\square$

**Lemma E.5.** *Under Assumption 4.1, for an arbitrarily small positive constant  $c$ , there exist universal positive constants  $C_1, C_2$  and  $C_3$  depending only on  $c$  such that as long as  $n \geq cp$ , with probability at least  $1 - 2e^{-C_2n} - Ne^{-C_3p}$ ,*

$$\sup_{k \in [m]} \sup_{\boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R)} \|\nabla^2 f_k(\boldsymbol{\theta}) - \nabla^2 F(\boldsymbol{\theta})\|_2 \leq C_1 \|\boldsymbol{\Sigma}\|_2 \sqrt{\frac{p \max\{1, \log(Np^{1/2} \|\boldsymbol{\Sigma}\|_2 R)\}}{n}}.$$

**Proof of Lemma E.5.** Let  $\tilde{\mathbf{x}}_i = (\boldsymbol{\Sigma}^*)^{-1/2} \mathbf{x}_i$ ,  $\tilde{\boldsymbol{\theta}} = (\boldsymbol{\Sigma}^*)^{1/2} \boldsymbol{\theta}$ , and define a new loss function  $\tilde{l}(\tilde{\boldsymbol{\theta}}, \tilde{\mathbf{x}}_i) = b(\tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\theta}}) - y_i(\tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\theta}})$ . Let  $\widehat{R}_k(\tilde{\boldsymbol{\theta}}) = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \tilde{l}(\tilde{\boldsymbol{\theta}}, \tilde{\mathbf{x}}_i)$  for  $k \in [m]$ . Then we have that  $\nabla^2 \widehat{R}_k(\tilde{\boldsymbol{\theta}}) = \frac{1}{n} \sum_{i \in \mathcal{I}_k} b''(\tilde{\boldsymbol{\theta}}^\top \tilde{\mathbf{x}}_i) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top$  and  $\nabla^2 f_k(\boldsymbol{\theta}) = (\boldsymbol{\Sigma}^*)^{1/2} \nabla^2 \widehat{R}_k(\tilde{\boldsymbol{\theta}}) (\boldsymbol{\Sigma}^*)^{1/2}$ . Similarly we have  $\nabla^2 F(\boldsymbol{\theta}) = (\boldsymbol{\Sigma}^*)^{1/2} \mathbb{E} \nabla^2 \widehat{R}_k(\tilde{\boldsymbol{\theta}}) (\boldsymbol{\Sigma}^*)^{1/2}$ .

Therefore

$$\max_{k \in [m]} \max_{\boldsymbol{\theta} \in B(\widehat{\boldsymbol{\theta}}, R)} \|\nabla^2 f_k(\boldsymbol{\theta}) - \nabla^2 F(\boldsymbol{\theta})\|_2 \leq \|\boldsymbol{\Sigma}^*\|_2 \max_{k \in [m]} \max_{\tilde{\boldsymbol{\theta}} \in B((\boldsymbol{\Sigma}^*)^{1/2} \widehat{\boldsymbol{\theta}}, \tilde{R})} \|\nabla^2 \widehat{R}_k(\tilde{\boldsymbol{\theta}}) - \nabla^2 \mathbb{E} \widehat{R}_k(\tilde{\boldsymbol{\theta}})\|_2 \quad (\text{E.6})$$

and we only need to control the quantity on the right hand side. Here  $\tilde{R} = \|\boldsymbol{\Sigma}^*\|_2^{1/2} R$ .

Define  $\Delta_0 = \max_{k \in [m]} \max_{\tilde{\boldsymbol{\theta}} \in B((\boldsymbol{\Sigma}^*)^{1/2} \widehat{\boldsymbol{\theta}}, \tilde{R})} \|\nabla^2 \widehat{R}_k(\tilde{\boldsymbol{\theta}}) - \nabla^2 \mathbb{E} \widehat{R}_k(\tilde{\boldsymbol{\theta}})\|_2$  and

$$\phi_k(\tilde{\boldsymbol{\theta}}) = \left\| \frac{1}{n} \sum_{i \in \mathcal{I}_k} b''(\tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\theta}}) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top - \mathbb{E} b''(\tilde{\mathbf{X}}^\top \tilde{\boldsymbol{\theta}}) \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top \right\|_2.$$

Here  $\tilde{\mathbf{X}}$  shares the distribution with  $\tilde{\mathbf{x}}_i$ . Firstly we bound  $\phi_k(\tilde{\boldsymbol{\theta}})$  for any fixed  $\tilde{\boldsymbol{\theta}} \in B(\tilde{\boldsymbol{\theta}}^*, 2\tilde{R})$ , where  $\tilde{\boldsymbol{\theta}}^* := (\boldsymbol{\Sigma}^*)^{1/2}\widehat{\boldsymbol{\theta}}$ . For any  $k \in [m]$ , under Assumption 4.1, there exist constants  $c_1, c_2$  such that for any  $\epsilon \geq 0$

$$\mathbb{P}(\phi_k(\tilde{\boldsymbol{\theta}}) \geq \epsilon) \leq 2e^{c_1 p - c_2 \min\{\epsilon, \epsilon^2\}n}. \quad (\text{E.7})$$

To see this, notice that  $\phi_k(\tilde{\boldsymbol{\theta}}) = \sup_{\mathbf{u} \in \mathbb{S}^{p-1}} g_k(\mathbf{u})$ , where  $g_k(\mathbf{u}) = \mathbf{u}^\top \left\{ \frac{1}{n} \sum_{i \in \mathcal{I}_k} b''(\tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\theta}}) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top - \mathbb{E} b''(\tilde{\mathbf{X}}^\top \tilde{\boldsymbol{\theta}}) \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top \right\} \mathbf{u}$ . Let  $\mathcal{N}$  be a  $\frac{1}{4}$ -covering of  $\mathbb{S}^{p-1}$ , and  $|\mathcal{N}| \leq 9^p$ . Denote  $\widehat{\mathbf{u}} = \arg \max_{\mathbf{u}} g_k(\mathbf{u})$ . Find  $\tilde{\mathbf{u}} \in \mathcal{N}$  such that  $\|\widehat{\mathbf{u}} - \tilde{\mathbf{u}}\|_2 \leq \frac{1}{4}$ . Then

$$|g_k(\tilde{\mathbf{u}}) - g_k(\widehat{\mathbf{u}})| = |(\tilde{\mathbf{u}} + \widehat{\mathbf{u}})^\top \left\{ \frac{1}{n} \sum_{i \in \mathcal{I}_k} b''(\tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\theta}}) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top - \mathbb{E} b''(\tilde{\mathbf{X}}^\top \tilde{\boldsymbol{\theta}}) \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top \right\} (\tilde{\mathbf{u}} - \widehat{\mathbf{u}})| \leq \frac{1}{2} g_k(\widehat{\mathbf{u}}),$$

and thus

$$\sup_{\mathbf{u} \in \mathbb{S}^{p-1}} g_k(\mathbf{u}) \leq 2 \sup_{\mathbf{u} \in \mathcal{N}} g_k(\mathbf{u}).$$

On the other hand from Bernstein's inequality we see that there exists a constant  $c_2$  such that for any  $\mathbf{u} \in \mathcal{N}$ ,  $\epsilon \geq 0$ ,  $\mathbb{P}(g_k(\mathbf{u}) \geq \frac{\epsilon}{2}) \leq 2e^{-c_2 \min\{\epsilon, \epsilon^2\}n}$ . Therefore

$$\mathbb{P}(\phi_k(\tilde{\boldsymbol{\theta}}) \geq \epsilon) = \mathbb{P}\left(\sup_{\mathbf{u} \in \mathbb{S}^{p-1}} g_k(\mathbf{u}) \geq \epsilon\right) \leq \mathbb{P}\left(\sup_{\mathbf{u} \in \mathcal{N}} g_k(\mathbf{u}) \geq \frac{\epsilon}{2}\right) \leq |\mathcal{N}| \cdot 2e^{-c_2 \min\{\epsilon, \epsilon^2\}n} \leq 2e^{c_1 p - c_2 \min\{\epsilon, \epsilon^2\}n}$$

where  $c_1 = \log 9$ .

Now for  $t \geq 1$ , define the event  $E_t \triangleq \left\{ \max_{i=1}^N \|\tilde{\mathbf{x}}_i\|_2^3 < (8t)^{3/2} \mathbb{E} \|\tilde{\mathbf{X}}\|_2^3 \right\}$ . Then by Theorem 2.1 in [Hsu et al. (2012)],  $\mathbb{P}(E_t^c) = \mathbb{P}(\max_{i=1}^N \|\tilde{\mathbf{x}}_i\|_2^2 \geq 8t[\mathbb{E} \|\tilde{\mathbf{X}}\|_2^3]^{2/3}) \leq \mathbb{P}(\max_{i=1}^N \|\tilde{\mathbf{x}}_i\|_2^2 \geq 8t \mathbb{E} \|\tilde{\mathbf{X}}\|_2^2) \leq N \mathbb{P}(\|\tilde{\mathbf{X}}\|_2^2 \geq 8t \mathbb{E} \|\tilde{\mathbf{X}}\|_2^2) \leq N e^{-tp}$ . Under the event  $E_t$ , for  $\tilde{\boldsymbol{\theta}}_1, \tilde{\boldsymbol{\theta}}_2 \in B(\tilde{\boldsymbol{\theta}}^*, 2\tilde{R})$ ,

we have

$$\begin{aligned}
|\phi_k(\tilde{\boldsymbol{\theta}}_1) - \phi_k(\tilde{\boldsymbol{\theta}}_2)| &\leq \left\| \frac{1}{n} \sum_{i \in \mathcal{I}_k} [b''(\tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\theta}}_1) - b''(\tilde{\mathbf{x}}_i^\top \tilde{\boldsymbol{\theta}}_2)] \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \right\|_2 + \|\mathbb{E} b''(\tilde{\mathbf{X}}^\top \tilde{\boldsymbol{\theta}}_1) \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top - \mathbb{E} b''(\tilde{\mathbf{X}}^\top \tilde{\boldsymbol{\theta}}_2) \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top\|_2 \\
&\leq B_3 \|\tilde{\boldsymbol{\theta}}_1 - \tilde{\boldsymbol{\theta}}_2\|_2 \cdot (\mathbb{E} \|\tilde{\mathbf{X}}\|_2^3 + \frac{1}{n} \sum_{i=1}^n \|\tilde{\mathbf{x}}_{ki}\|_2^3) \\
&\leq (9t)^{3/2} B_3 \cdot \mathbb{E} \|\mathbf{U}\|_2^3 \cdot \|\tilde{\boldsymbol{\theta}}_1 - \tilde{\boldsymbol{\theta}}_2\|_2 \\
&\leq c_3 (pt)^{3/2} \|\tilde{\boldsymbol{\theta}}_1 - \tilde{\boldsymbol{\theta}}_2\|_2
\end{aligned}$$

for some constant  $c_3$  depending only on  $B_3$ .

Now let  $\mathcal{N}_\delta$  be a  $\delta$ -covering of  $B(\tilde{\boldsymbol{\theta}}^*, 2\tilde{R})$ , where  $\delta = \frac{\epsilon}{c_3(tp)^{3/2}}$ . We can also assume that  $|\mathcal{N}_\delta| \leq (\frac{6\tilde{R}}{\delta})^p$ . Therefore for any  $k \in [m]$ ,

$$\begin{aligned}
\mathbb{P} \left( E_t \cap \left\{ \sup_{\tilde{\boldsymbol{\theta}} \in B(\tilde{\boldsymbol{\theta}}^*, 2\tilde{R})} \phi_k(\tilde{\boldsymbol{\theta}}) \geq 2\epsilon \right\} \right) &\leq \mathbb{P} \left( E_t \cap \left\{ \sup_{\tilde{\boldsymbol{\theta}} \in \mathcal{N}_\delta} \phi_k(\tilde{\boldsymbol{\theta}}) \geq \epsilon \right\} \right) \\
&\leq |\mathcal{N}_\delta| \cdot 2e^{c_1 p - c_2 \min\{\epsilon, \epsilon^2\}n} = 2e^{c_4 p + c_3 p \log \frac{(tp)^{3/2}\tilde{R}}{\epsilon} - c_2 \min\{\epsilon, \epsilon^2\}n}.
\end{aligned}$$

Thus

$$\begin{aligned}
\mathbb{P}(\Delta_0 \geq 2\epsilon) &\leq \mathbb{P} \left( \bigcup_{k \in [m]} \left\{ \sup_{\tilde{\boldsymbol{\theta}} \in B(\tilde{\boldsymbol{\theta}}^*, \tilde{R})} \phi_k(\tilde{\boldsymbol{\theta}}) \geq 2\epsilon \right\} \right) \\
&\leq \mathbb{P}(E_t^c) + \sum_{k=1}^m \mathbb{P} \left( E_t \cap \left\{ \sup_{\tilde{\boldsymbol{\theta}} \in B(\tilde{\boldsymbol{\theta}}^*, \tilde{R})} \phi_k(\tilde{\boldsymbol{\theta}}) \geq 2\epsilon \right\} \right) \\
&\leq Ne^{-tp} + 2m e^{c_4 p + c_3 p \log \frac{(tp)^{3/2}\tilde{R}}{\epsilon} - c_2 \min\{\epsilon, \epsilon^2\}n}.
\end{aligned}$$

It is easily seen that the last expression is no more than  $Ne^{-tp} + 2e^{-C_2 p}$  when

$$\begin{cases} \min\{\epsilon, \epsilon^2\} \geq C'_1 \frac{\max\{p, p \log(t^{3/2}R), \log m\}}{n}, \\ \epsilon \geq 1 \text{ or } \frac{\epsilon^2}{\log \frac{1}{\epsilon}} \geq C'_2 \cdot \frac{p}{n}, \end{cases}$$

which is satisfied if  $t$  is chosen as a suitable constant depending on  $c$ , and that

$$\epsilon = C \sqrt{\frac{\log m + p \max\{1, \log(np^{1/2}\tilde{R})\}}{n}}. \quad (\text{E.8})$$

Here  $C'_i$  and  $C$  is a constant depending only on  $c$ .

Finally, note that

$$\Sigma^* = \text{cov}(\mathbf{x}_i) = \begin{pmatrix} 1 \\ & \Sigma \end{pmatrix}$$

and that  $\|\Sigma\|_2 \geq A_1$  for a universal  $A_1 > 0$ , we have  $\|\Sigma^*\|_2 \leq \max\{1, 1/A_1\}\|\Sigma\|_2$ . Thus combining (E.8) and (E.6) completes the proof.  $\square$

## E.6 Proof of Theorem 4.2

Theorem 4.2 is a special of Theorem 4.1 by taking  $\alpha = 0$ .

## E.7 Proof of Theorem B.1

We first present three lemmas, based on which we build the proof of the main theorem.

**Lemma E.6.** *Suppose that  $\widehat{\Sigma}$  is positive-definite, i.e.  $\lambda_{\min}(\widehat{\Sigma}) > 0$ . Define  $\boldsymbol{\varepsilon}_t = \widehat{\Sigma}^{1/2}(\boldsymbol{\theta}_t - \widehat{\Sigma}^{-1}\widehat{\mathbf{w}})$  for  $t \geq 0$  and*

$$\Delta_k = (\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2}(\widehat{\Sigma}_k - \widehat{\Sigma})(\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2}, \quad \forall k \in [m].$$

If  $\alpha \geq 0$  is appropriately chosen such that  $\Delta = \max_{k \in [m]} \|\Delta_k\|_2 \leq 1/2$ . Then

$$\|\varepsilon_{t+1}\|_2 \leq \frac{2\Delta^2 + \alpha/\lambda_{\min}(\widehat{\Sigma})}{1 + \alpha/\lambda_{\min}(\widehat{\Sigma})} \|\varepsilon_t\|_2, \quad \forall t \geq 0,$$

which guarantees linear convergence of  $\{\varepsilon_t\}_{t=0}^\infty$ .

**Proof of Lemma E.6.** Define  $\varepsilon_{t,k} = \widehat{\Sigma}^{1/2}(\boldsymbol{\theta}_{t,k} - \widehat{\Sigma}^{-1}\widehat{\mathbf{w}})$ . Then

$$\begin{aligned} \varepsilon_{t+1,k} &= \widehat{\Sigma}^{1/2}(\boldsymbol{\theta}_{t+1,k} - \widehat{\Sigma}^{-1}\widehat{\mathbf{w}}) \\ &= \widehat{\Sigma}^{1/2}[\mathbf{I} - (\widehat{\Sigma}_k + \alpha\mathbf{I})^{-1}\widehat{\Sigma}]\boldsymbol{\theta}_t + \widehat{\Sigma}^{1/2}(\widehat{\Sigma}_k + \alpha\mathbf{I})^{-1}\widehat{\mathbf{w}} - \widehat{\Sigma}^{-1/2}\widehat{\mathbf{w}} \\ &= [\mathbf{I} - \widehat{\Sigma}^{1/2}(\widehat{\Sigma}_k + \alpha\mathbf{I})^{-1}\widehat{\Sigma}^{1/2}]\varepsilon_t. \end{aligned}$$

Define  $\widetilde{\Sigma}_k^{(1)} = \widehat{\Sigma}^{-1/2}\widehat{\Sigma}_k\widehat{\Sigma}^{-1/2}$  for  $k \in [m]$ . The fact

$$\widehat{\Sigma}_k + \alpha\mathbf{I} = \widehat{\Sigma}^{1/2}(\widehat{\Sigma}^{-1/2}\widehat{\Sigma}_k\widehat{\Sigma}^{-1/2} + \alpha\widehat{\Sigma}^{-1})\widehat{\Sigma}^{1/2} = \widehat{\Sigma}^{1/2}(\widetilde{\Sigma}_k^{(1)} + \alpha\widehat{\Sigma}^{-1})\widehat{\Sigma}^{1/2}$$

gives  $\widehat{\Sigma}^{1/2}(\widehat{\Sigma}_k + \alpha\mathbf{I})^{-1}\widehat{\Sigma}^{1/2} = (\widetilde{\Sigma}_k^{(1)} + \alpha\widehat{\Sigma}^{-1})^{-1}$  and  $\varepsilon_{t+1,k} = [\mathbf{I} - (\widetilde{\Sigma}_k^{(1)} + \alpha\widehat{\Sigma}^{-1})^{-1}]\varepsilon_t$ .

Define  $\widehat{\mathbf{D}} = (\mathbf{I} + \alpha\widehat{\Sigma}^{-1})^{-1}$  and  $\widetilde{\Sigma}_k = \widehat{\mathbf{D}}^{1/2}\widetilde{\Sigma}_k^{(1)}\widehat{\mathbf{D}}^{1/2}$ . From

$$(\widetilde{\Sigma}_k^{(1)} + \alpha\widehat{\Sigma}^{-1})^{-1} = [\widehat{\mathbf{D}}^{-1} + (\widetilde{\Sigma}_k^{(1)} - \mathbf{I})]^{-1} = \widehat{\mathbf{D}}^{1/2}[\mathbf{I} + (\widetilde{\Sigma}_k - \widehat{\mathbf{D}})]^{-1}\widehat{\mathbf{D}}^{1/2}$$

and

$$\begin{aligned} \widetilde{\Sigma}_k - \widehat{\mathbf{D}} &= \widehat{\mathbf{D}}^{1/2}(\widetilde{\Sigma}_k^{(1)} - \mathbf{I})\widehat{\mathbf{D}}^{1/2} = (\mathbf{I} + \alpha\widehat{\Sigma}^{-1})^{-1/2}\widehat{\Sigma}^{-1/2}(\widehat{\Sigma}_k - \widehat{\Sigma})\widehat{\Sigma}^{-1/2}(\mathbf{I} + \alpha\widehat{\Sigma}^{-1})^{-1/2} \\ &= (\widehat{\Sigma} + \alpha\mathbf{I})^{-1/2}(\widehat{\Sigma}_k - \widehat{\Sigma})(\widehat{\Sigma} + \alpha\mathbf{I})^{-1/2} = \Delta_k, \end{aligned}$$

we get

$$\boldsymbol{\varepsilon}_{t+1,k} = (\mathbf{I} - \widehat{\mathbf{D}}^{1/2} \mathbf{C}_k \widehat{\mathbf{D}}^{1/2}) \boldsymbol{\varepsilon}_t \quad \text{and} \quad \boldsymbol{\varepsilon}_{t+1} = (\mathbf{I} - \widehat{\mathbf{D}}^{1/2} \mathbf{C} \widehat{\mathbf{D}}^{1/2}) \boldsymbol{\varepsilon}_t,$$

where  $\mathbf{C}_k = (\mathbf{I} + \boldsymbol{\Delta}_k)^{-1}$  and  $\mathbf{C} = \frac{1}{m} \sum_{k=1}^m \mathbf{C}_k$ . Let  $\mathbf{R}_k = \mathbf{C}_k - (\mathbf{I} - \boldsymbol{\Delta}_k)$  and  $\mathbf{R} = \frac{1}{m} \sum_{k=1}^m \mathbf{R}_k = \mathbf{C} - \mathbf{I}$ . We have  $\|\mathbf{I} - \widehat{\mathbf{D}}^{1/2} \mathbf{C} \widehat{\mathbf{D}}^{1/2}\|_2 = \|\mathbf{I} - \widehat{\mathbf{D}}^{1/2} (\mathbf{I} + \mathbf{R}) \widehat{\mathbf{D}}^{1/2}\|_2$ . Below we control the right-hand side.

By  $\Delta = \max_{k \in [m]} \|\boldsymbol{\Delta}_k\|_2 \leq 1/2$  and Lemma F.4, we obtain that  $\|\mathbf{C}_k\|_2 \leq \frac{1}{1-\Delta} \leq 2$  and  $\|\mathbf{R}_k\|_2 \leq 2\Delta^2$ . Consequently,  $\|\mathbf{C}\|_2 \leq 2$  and  $\|\mathbf{R}\|_2 \leq 2\Delta^2 \leq 1/2$ . Then we obtain that

$$(1 - 2\Delta^2)\mathbf{I} \preceq \mathbf{I} + \mathbf{R} \preceq (1 + 2\Delta^2)\mathbf{I},$$

$$(1 - 2\Delta^2)\widehat{\mathbf{D}} \preceq \widehat{\mathbf{D}}^{1/2} (\mathbf{I} + \mathbf{R}) \widehat{\mathbf{D}}^{1/2} \preceq (1 + 2\Delta^2)\widehat{\mathbf{D}},$$

$$\mathbf{I} - (1 + 2\Delta^2)\widehat{\mathbf{D}} \preceq \mathbf{I} - \widehat{\mathbf{D}}^{1/2} (\mathbf{I} + \mathbf{R}) \widehat{\mathbf{D}}^{1/2} \preceq \mathbf{I} - (1 - 2\Delta^2)\widehat{\mathbf{D}}.$$

Consequently,

$$\|\mathbf{I} - \widehat{\mathbf{D}}^{1/2} (\mathbf{I} + \mathbf{R}) \widehat{\mathbf{D}}^{1/2}\|_2 \leq \max \left\{ \|\mathbf{I} - (1 - 2\Delta^2)\widehat{\mathbf{D}}\|_2, \|\mathbf{I} - (1 + 2\Delta^2)\widehat{\mathbf{D}}\|_2 \right\}.$$

Let  $\{\widehat{\lambda}_j\}_{j=1}^p$  be the eigenvalues of  $\widehat{\boldsymbol{\Sigma}}$  sorted in descending order. Since  $\widehat{\mathbf{D}}$  has eigenvalues  $\{(1 + \alpha/\widehat{\lambda}_j)^{-1}\}_{j=1}^p \subseteq (0, 1]$ , the eigenvalues of  $\mathbf{I} - (1 \pm 2\Delta^2)\widehat{\mathbf{D}}$  are  $\left\{1 - \frac{1 \pm 2\Delta^2}{1 + \alpha/\widehat{\lambda}_j}\right\}_{j=1}^p$ . Then

$$\|\mathbf{I} - (1 \pm 2\Delta^2)\widehat{\mathbf{D}}\|_2 = \max \left\{ \left|1 - \frac{1 \pm 2\Delta^2}{1 + \alpha/\widehat{\lambda}_1}\right|, \left|1 - \frac{1 \pm 2\Delta^2}{1 + \alpha/\widehat{\lambda}_p}\right| \right\}.$$

By elementary calculation and the fact  $2\Delta^2 \leq 1/2 < 1$  we get

$$\begin{aligned} \left| 1 - \frac{1 + 2\Delta^2}{1 + \alpha/\widehat{\lambda}_1} \right| &= \frac{\left| \alpha/\widehat{\lambda}_1 - 2\Delta^2 \right|}{1 + \alpha/\widehat{\lambda}_1} \leq \frac{\max\{\alpha/\widehat{\lambda}_1, 2\Delta^2\}}{1 + \alpha/\widehat{\lambda}_1} \leq \max \left\{ \frac{\alpha/\widehat{\lambda}_p}{1 + \alpha/\widehat{\lambda}_p}, 2\Delta^2 \right\}, \\ \left| 1 - \frac{1 + 2\Delta^2}{1 + \alpha/\widehat{\lambda}_p} \right| &= \frac{\left| \alpha/\widehat{\lambda}_p - 2\Delta^2 \right|}{1 + \alpha/\widehat{\lambda}_p} \leq \frac{2\Delta^2 + \alpha/\widehat{\lambda}_p}{1 + \alpha/\widehat{\lambda}_p}, \\ 0 \leq 1 - \frac{1 - 2\Delta^2}{1 + \alpha/\widehat{\lambda}_1} &\leq 1 - \frac{1 - 2\Delta^2}{1 + \alpha/\widehat{\lambda}_p} = \frac{2\Delta^2 + \alpha/\widehat{\lambda}_p}{1 + \alpha/\widehat{\lambda}_p}. \end{aligned}$$

Therefore,

$$\|\mathbf{I} - \widehat{\mathbf{D}}^{1/2} \mathbf{C} \widehat{\mathbf{D}}^{1/2}\|_2 \leq \max \left\{ \frac{2\Delta^2 + \alpha/\widehat{\lambda}_p}{1 + \alpha/\widehat{\lambda}_p}, 2\Delta^2 \right\} = \frac{2\Delta^2 + \alpha/\widehat{\lambda}_p}{1 + \alpha/\widehat{\lambda}_p} < 1.$$

□

**Lemma E.7.** Suppose Assumption B.1 hold. Then there exists a constant  $C$  determined by  $\|\mathbf{u}_i\|_{\psi_2}$ , such that  $\mathbb{P}(\Delta \leq 1/2) \geq 1 - 2me^{-n/C}$  holds under either of the two conditions: (i)  $n \geq Cp$  and  $\alpha \geq 0$ ; (ii)  $\alpha \geq C\text{Tr}(\Sigma)/n$ .

**Proof of Lemma E.7.** Define  $\widehat{\mathbf{S}}_k = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{u}_i \mathbf{u}_i^\top$  and  $\bar{\mathbf{u}}_{(k)} = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{u}_i$  for  $k \in [m]$ . Then we have  $\widehat{\Sigma}_k = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{x}_i \mathbf{x}_i^\top = \begin{pmatrix} 1 & \bar{\mathbf{u}}_{(k)}^\top \\ \bar{\mathbf{u}}_{(k)} & \widehat{\mathbf{S}}_k \end{pmatrix}$ . Let  $\Sigma^* = \mathbb{E} \widehat{\Sigma}_k = \begin{pmatrix} 1 & \mathbf{0} \\ \mathbf{0} & \Sigma \end{pmatrix}$  and observe that

$$\Delta_k = (\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2} (\widehat{\Sigma}_k - \Sigma^*) (\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2} - (\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2} (\widehat{\Sigma} - \Sigma^*) (\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2}.$$

Let  $\mathbf{B}_k = (\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2} (\widehat{\Sigma}_k - \Sigma^*) (\widehat{\Sigma} + \alpha \mathbf{I})^{-1/2}$ . Since

$$\max_{k \in [m]} \|\Delta_k\|_2 = \max_{k \in [m]} \|\mathbf{B}_k - \frac{1}{m} \sum_{\ell=1}^m \mathbf{B}_\ell\|_2 \leq 2 \max_{k \in [m]} \|\mathbf{B}_k\|_2, \quad (\text{E.9})$$

it boils down to bound  $\|\mathbf{B}_k\|_2$ . To do this, we let  $\mathbf{A}_0 = (\Sigma^* + \alpha\mathbf{I})^{1/2}(\widehat{\Sigma} + \alpha\mathbf{I})^{-1/2}$  and write

$$\begin{aligned}\|\mathbf{B}_k\|_2 &= \|\mathbf{A}_0^\top(\Sigma^* + \alpha\mathbf{I})^{-1/2}(\widehat{\Sigma}_k - \Sigma^*)(\Sigma^* + \alpha\mathbf{I})^{-1/2}\mathbf{A}_0\|_2 \\ &\leq \|\mathbf{A}_0\|_2^2 \|(\Sigma^* + \alpha\mathbf{I})^{-1/2}(\widehat{\Sigma}_k - \Sigma^*)(\Sigma^* + \alpha\mathbf{I})^{-1/2}\|_2.\end{aligned}$$

Define  $\mathbf{D} = (\mathbf{I} + \alpha(\Sigma^*)^{-1})^{-1}$  and  $\widehat{\Sigma}_k^{(1)} = (\Sigma^* + \alpha\mathbf{I})^{-1/2}\widehat{\Sigma}_k(\Sigma^* + \alpha\mathbf{I})^{-1/2}$ . On the one hand,

$$(\Sigma^* + \alpha\mathbf{I})^{-1/2}(\widehat{\Sigma}_k - \Sigma^*)(\Sigma^* + \alpha\mathbf{I})^{-1/2} = \widehat{\Sigma}_k^{(1)} - \mathbf{D}.$$

On the other hand,

$$\begin{aligned}\|\mathbf{A}_0\|_2^2 &= \|\mathbf{A}_0\mathbf{A}_0^\top\|_2 = \|(\Sigma^* + \alpha\mathbf{I})^{1/2}(\widehat{\Sigma} + \alpha\mathbf{I})^{-1}(\Sigma^* + \alpha\mathbf{I})^{1/2}\|_2 \\ &= \|(\Sigma^* + \alpha\mathbf{I})^{1/2}[(\widehat{\Sigma} - \Sigma^*) + (\Sigma^* + \alpha\mathbf{I})]^{-1}(\Sigma^* + \alpha\mathbf{I})^{1/2}\|_2 \\ &\leq \|[(\Sigma^* + \alpha\mathbf{I})^{-1/2}(\widehat{\Sigma} - \Sigma^*)(\Sigma^* + \alpha\mathbf{I})^{-1/2} + \mathbf{I}]^{-1}\|_2 = \|[\mathbf{I} + \frac{1}{m} \sum_{k=1}^m (\widehat{\Sigma}_k^{(1)} - \mathbf{D})]^{-1}\|_2 \\ &\leq \frac{1}{1 - \|\frac{1}{m} \sum_{k=1}^m (\widehat{\Sigma}_k^{(1)} - \mathbf{D})\|_2} \leq \frac{1}{1 - \max_{k \in [m]} \|\widehat{\Sigma}_k^{(1)} - \mathbf{D}\|_2},\end{aligned}$$

where we used Lemma [F.4](#). Based on these, we have

$$\max_{k \in [m]} \|\mathbf{B}_k\|_2 \leq \frac{\max_{k \in [m]} \|\widehat{\Sigma}_k^{(1)} - \mathbf{D}\|_2}{1 - \max_{k \in [m]} \|\widehat{\Sigma}_k^{(1)} - \mathbf{D}\|_2}, \quad (\text{E.10})$$

and it suffices to prove under the given conditions that

$$\mathbb{P} \left( \max_{k \in [m]} \|\widehat{\Sigma}_k^{(1)} - \mathbf{D}\|_2 \leq 1/5 \right) \geq 1 - 2me^{-n/C} \quad (\text{E.11})$$

holds for some constant  $C$ .

By definition, we have

$$\begin{aligned}
\widehat{\Sigma}_k^{(1)} - \mathbf{D} &= \begin{pmatrix} (1 + \alpha)^{-1/2} & \mathbf{0} \\ \mathbf{0} & (\Sigma + \alpha \mathbf{I})^{-1/2} \end{pmatrix} \begin{pmatrix} \mathbf{0} & \bar{\mathbf{u}}_{(k)}^\top \\ \bar{\mathbf{u}}_{(k)} & \widehat{\mathbf{S}}_k - \Sigma \end{pmatrix} \begin{pmatrix} (1 + \alpha)^{-1/2} & \mathbf{0} \\ \mathbf{0} & (\Sigma + \alpha \mathbf{I})^{-1/2} \end{pmatrix} \\
&= \begin{pmatrix} \mathbf{0} & (1 + \alpha)^{-1/2}[(\Sigma + \alpha \mathbf{I})^{-1/2} \bar{\mathbf{u}}_{(k)}]^\top \\ (1 + \alpha)^{-1/2}(\Sigma + \alpha \mathbf{I})^{-1/2} \bar{\mathbf{u}}_{(k)} & (\Sigma + \alpha \mathbf{I})^{-1/2}(\widehat{\mathbf{S}}_k - \Sigma)(\Sigma + \alpha \mathbf{I})^{-1/2} \end{pmatrix}
\end{aligned}$$

and as a result,

$$\|\widehat{\Sigma}_k^{(1)} - \mathbf{D}\|_2 \leq (1 + \alpha)^{-1/2} \|(\Sigma + \alpha \mathbf{I})^{-1/2} \bar{\mathbf{u}}_{(k)}\|_2 + \|(\Sigma + \alpha \mathbf{I})^{-1/2}(\widehat{\mathbf{S}}_k - \Sigma)(\Sigma + \alpha \mathbf{I})^{-1/2}\|_2. \quad (\text{E.12})$$

Here we used a simple fact that  $\left\| \begin{pmatrix} \mathbf{0} & \mathbf{A}^\top \\ \mathbf{A} & \mathbf{0} \end{pmatrix} \right\|_2 = \|\mathbf{A}\|_2$  for any matrix  $\mathbf{A}$ .

Observe that  $\mathbf{v}_i = (\Sigma + \alpha \mathbf{I})^{-1/2} \mathbf{u}_i$  is a sub-gaussian random variable with zero mean and covariance matrix  $(1 + \alpha \Sigma^{-1})^{-1}$ . On the other hand,  $(\Sigma + \alpha \mathbf{I})^{-1/2}(\widehat{\mathbf{S}}_k - \Sigma)(\Sigma + \alpha \mathbf{I})^{-1/2} = \frac{1}{n} \sum_{i \in \mathcal{I}_k} \mathbf{v}_i \mathbf{v}_i^\top - (1 + \alpha \Sigma^{-1})^{-1}$ . Lemma F.5 forces

$$\begin{aligned}
\mathbb{P} \left( \|(\Sigma + \alpha \mathbf{I})^{-1/2} \bar{\mathbf{u}}_{(k)}\|_2 > 1/10 \right) &\leq e^{-n/C}, \\
\mathbb{P} \left( \|(\Sigma + \alpha \mathbf{I})^{-1/2}(\widehat{\mathbf{S}}_k - \Sigma)(\Sigma + \alpha \mathbf{I})^{-1/2}\|_2 > 1/10 \right) &\leq e^{-n/C},
\end{aligned}$$

where  $C$  is the constant therein. These estimates and (E.12) lead to (E.11). □

Following the similar idea in the proof above, we get the following results.

**Lemma E.8.** *Suppose Assumption B.1 holds with  $C$  being the constant in Lemma E.7.*

*Then*

- $\mathbb{P}(\|(\Sigma^*)^{-1/2}(\widehat{\Sigma} - \Sigma^*)(\Sigma^*)^{-1/2}\|_2 \leq 1/2) \geq 1 - 2me^{-N/C};$
- $\mathbb{P}(\Delta \leq C'_2 \sqrt{p/n}) \geq 1 - 2e^{-C'_1 p}$  holds for some constants  $C'_1$  and  $C'_2$ .

Now we come back to the main proof. We will use the three lemmas above to show that with high probability,  $\lambda_{\max}(\widehat{\Sigma})/\lambda_{\min}(\widehat{\Sigma}) \leq 3\kappa$  and

$$\left\| \widehat{\Sigma}^{1/2} (\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}) \right\|_2 \leq \left( 1 - \frac{1 - \min\{1/2, C_2 p/n\}}{1 + C_3 \alpha} \right) \left\| \widehat{\Sigma}^{1/2} (\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}) \right\|_2, \quad \forall t \geq 0. \quad (\text{E.13})$$

Then we conclude the proof by induction and some simple linear algebra.

First, Lemma E.8 asserts that with high probability,  $\|(\Sigma^*)^{-1/2}(\widehat{\Sigma} - \Sigma^*)(\Sigma^*)^{-1/2}\|_2 \leq 1/2$ . On this event, we have  $-\frac{1}{2}\Sigma^* \preceq \widehat{\Sigma} - \Sigma^* \preceq \frac{1}{2}\Sigma^*$  and thus  $\frac{1}{2}\Sigma^* \preceq \widehat{\Sigma} \preceq \frac{3}{2}\Sigma^*$ . Hence

$$\frac{1}{2}\lambda_{\min}(\Sigma^*) \leq \lambda_{\min}(\widehat{\Sigma}) \leq \lambda_{\max}(\widehat{\Sigma}) \leq \frac{3}{2}\lambda_{\max}(\Sigma^*)$$

and  $\lambda_{\max}(\widehat{\Sigma})/\lambda_{\min}(\widehat{\Sigma}) \leq 3\lambda_{\max}(\Sigma^*)/\lambda_{\min}(\Sigma^*) = 3\kappa$ .

Second, Lemma E.6 forces

$$\left\| \widehat{\Sigma}^{1/2} (\boldsymbol{\theta}_{t+1} - \widehat{\boldsymbol{\theta}}) \right\|_2 \leq \frac{2\Delta^2 + \alpha/\lambda_{\min}(\widehat{\Sigma})}{1 + \alpha/\lambda_{\min}(\widehat{\Sigma})} \left\| \widehat{\Sigma}^{1/2} (\boldsymbol{\theta}_t - \widehat{\boldsymbol{\theta}}) \right\|_2, \quad \forall t \geq 0. \quad (\text{E.14})$$

Lemmas E.7 and E.8 imply that  $\Delta \leq 1/2$ ,  $\Delta \leq \tilde{C} \sqrt{p/n}$ , and  $\lambda_{\min}(\widehat{\Sigma}) \geq 1$  hold simultaneously with high probability, where  $\tilde{C}$  is some constant. On this event, we have

$$\frac{2\Delta^2 + \alpha/\lambda_{\min}(\widehat{\Sigma})}{1 + \alpha/\lambda_{\min}(\widehat{\Sigma})} = 1 - \frac{1 - 2\Delta^2}{1 + \alpha/\lambda_{\min}(\widehat{\Sigma})} \leq 1 - \frac{1 - \min\{1/2, C_2 p/n\}}{1 + C_3 \alpha}$$

for some constants  $C_2$  and  $C_3$ . Then we get (E.13) from the estimates above and complete the proof.

## F Technical lemmas

The following lemma lists basic properties of strongly convex functions, which can be found in standard textbooks on convex optimization (Nesterov, 2013).

**Lemma F.1.** *Suppose  $f$  is a convex function defined on some convex open set  $\Omega \subseteq \mathbb{R}^p$ , and  $\partial f(\mathbf{x})$  denotes its subdifferential set at  $\mathbf{x} \in \Omega$ . The followings are equivalent:*

- $f$  is  $\rho$ -strongly convex in  $\Omega$ ;
- $f[(1-t)\mathbf{x} + t\mathbf{y}] \leq (1-t)f(\mathbf{x}) + t f(\mathbf{y}) - (\rho/2)t(1-t)\|\mathbf{y} - \mathbf{x}\|_2^2$ ,  $\forall \mathbf{x}, \mathbf{y} \in \Omega$  and  $t \in [0, 1]$ ;
- $\langle \mathbf{h} - \mathbf{g}, \mathbf{y} - \mathbf{x} \rangle \geq \rho\|\mathbf{y} - \mathbf{x}\|_2^2$ ,  $\forall \mathbf{x}, \mathbf{y} \in \Omega$ ,  $\mathbf{g} \in \partial f(\mathbf{x})$  and  $\mathbf{h} \in \partial f(\mathbf{y})$ .

If any of the above holds,  $f$  is said to be  $\rho$ -strongly convex.

**Lemma F.2.** *Let  $f : \mathbb{R}^p \rightarrow \mathbb{R}$  be a convex function. Suppose there exists  $\mathbf{x} \in \mathbb{R}^p$  and  $r > 0$  such that  $f$  is  $\rho$ -strongly convex in  $B(\mathbf{x}, r)$ . If  $\|\mathbf{h} - \mathbf{g}\|_2 < \rho r$  holds for some  $\mathbf{g} \in \partial f(\mathbf{x})$  and  $\mathbf{h} \in \partial f(\mathbf{y})$ , then  $\|\mathbf{y} - \mathbf{x}\|_2 \leq \|\mathbf{h} - \mathbf{g}\|_2/\rho \leq r$ .*

*Proof.* If we know a priori that  $\|\mathbf{y} - \mathbf{x}\|_2 \leq r$ , then we use the strong convexity of  $f$  in  $B(\mathbf{x}, r)$  and Cauchy-Schwarz inequality to obtain

$$\rho\|\mathbf{y} - \mathbf{x}\|_2^2 \leq \langle \mathbf{h} - \mathbf{g}, \mathbf{y} - \mathbf{x} \rangle \leq \|\mathbf{h} - \mathbf{g}\|_2\|\mathbf{y} - \mathbf{x}\|_2,$$

and get the desired result. Suppose on the contrary that  $\|\mathbf{y} - \mathbf{x}\|_2 > r$ , and define  $\bar{\mathbf{y}} = \mathbf{x} + r(\mathbf{y} - \mathbf{x})/\|\mathbf{y} - \mathbf{x}\|_2$ . Then  $\|\bar{\mathbf{y}} - \mathbf{x}\|_2 = r$ . The strong convexity of  $f$  in  $B(\mathbf{x}, r)$  and Lemma F.1 yield

$$\langle \mathbf{s} - \mathbf{g}, \bar{\mathbf{y}} - \mathbf{x} \rangle \geq \rho\|\bar{\mathbf{y}} - \mathbf{x}\|_2^2, \quad \forall \mathbf{s} \in \partial f(\bar{\mathbf{y}}).$$

By the convexity of  $f$ , we always have

$$\langle \mathbf{h} - \mathbf{s}, \bar{\mathbf{y}} - \mathbf{x} \rangle = \frac{r}{\|\mathbf{y} - \mathbf{x}\|_2 - r} \langle \mathbf{h} - \mathbf{s}, \mathbf{y} - \bar{\mathbf{y}} \rangle \geq 0, \quad \forall \mathbf{s} \in \partial f(\bar{\mathbf{y}}).$$

Summing up the two inequalities above, we get

$$\rho \|\bar{\mathbf{y}} - \mathbf{x}\|_2^2 \leq \langle \mathbf{h} - \mathbf{g}, \bar{\mathbf{y}} - \mathbf{x} \rangle \leq \|\mathbf{h} - \mathbf{g}\|_2 \|\bar{\mathbf{y}} - \mathbf{x}\|_2,$$

where we also used the Cauchy-Schwarz inequality. Then  $\|\mathbf{h} - \mathbf{g}\|_2 \geq \rho \|\bar{\mathbf{y}} - \mathbf{x}\|_2 = \rho r$  leads to contradiction. Hence, we must have only the case  $\|\mathbf{y} - \mathbf{x}\|_2 \leq r$ .  $\square$

**Lemma F.3.** *Let  $f : \mathbb{R}^p \rightarrow \mathbb{R}$  be a convex function. For any  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$ , we have*

$$\|\text{prox}_f(\mathbf{x}) - \text{prox}_f(\mathbf{y})\|_2^2 \leq \langle \mathbf{x} - \mathbf{y}, \text{prox}_f(\mathbf{x}) - \text{prox}_f(\mathbf{y}) \rangle.$$

*If  $\inf_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) > -\infty$ , then  $\|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2 \leq \|\mathbf{x} - \mathbf{x}^*\|_2$  and*

$$\|\text{prox}_f(\mathbf{x}) - \mathbf{x}\|_2^2 \leq \|\mathbf{x} - \mathbf{x}^*\|_2^2 - \|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2^2$$

*hold for any  $\mathbf{x}^* \in \text{argmin}_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x})$ .*

*If  $f$  is  $\rho$ -strongly convex in  $B(\mathbf{x}^*, r)$  for some  $r > 0$  and  $\mathbf{x}^* = \text{argmin}_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x})$ , then*

$$\|\text{prox}_{\alpha^{-1}f}(\mathbf{x}) - \mathbf{x}^*\|_2 \leq \frac{\alpha}{\alpha + \rho} \|\mathbf{x} - \mathbf{x}^*\|_2, \quad \forall \mathbf{x} \in B(\mathbf{x}^*, r) \text{ and } \alpha > 0.$$

**Proof of Lemma F.3.** The first claim is the well-known ‘‘firm non-expansiveness’’ property of the proximal mapping (Parikh and Boyd, 2014).

If  $\inf_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x}) > -\infty$ , then any  $\mathbf{x}^* \in \text{argmin}_{\mathbf{x} \in \mathbb{R}^p} f(\mathbf{x})$  is a fixed point of  $\text{prox}_f$ . The

firm non-expansiveness with  $\mathbf{y} = \mathbf{x}^*$  yields

$$\|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2^2 \leq \langle \mathbf{x} - \mathbf{x}^*, \text{prox}_f(\mathbf{x}) - \mathbf{x}^* \rangle \quad (\text{F.1})$$

and then  $\|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2 \leq \|\mathbf{x} - \mathbf{x}^*\|_2$ . The next claim is proved by

$$\begin{aligned} \|\text{prox}_f(\mathbf{x}) - \mathbf{x}\|_2^2 &= \|[\text{prox}_f(\mathbf{x}) - \mathbf{x}^*] - (\mathbf{x} - \mathbf{x}^*)\|_2^2 \\ &= \|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2^2 + \|\mathbf{x} - \mathbf{x}^*\|_2^2 - 2\langle \text{prox}_f(\mathbf{x}) - \mathbf{x}^*, \mathbf{x} - \mathbf{x}^* \rangle \\ &\leq \|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2^2 + \|\mathbf{x} - \mathbf{x}^*\|_2^2 - 2\|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2^2 \\ &= \|\mathbf{x} - \mathbf{x}^*\|_2^2 - \|\text{prox}_f(\mathbf{x}) - \mathbf{x}^*\|_2^2, \end{aligned}$$

where the inequality follows from (F.1).

For the last claim, we fix any  $\alpha > 0$  and  $\mathbf{x} \in B(\mathbf{x}^*, r)$  and define  $\mathbf{x}^+ = \text{prox}_{\alpha^{-1}f}(\mathbf{x})$ .

Then  $\|\mathbf{x}^+ - \mathbf{x}^*\|_2 \leq \|\mathbf{x} - \mathbf{x}^*\|_2 < r$ . The optimality conditions for  $\mathbf{x}^* = \text{argmin}_{\mathbf{y} \in \mathbb{R}^p} f(\mathbf{y})$  and  $\mathbf{x}^+ = \text{argmin}_{\mathbf{y} \in \mathbb{R}^p} \{f(\mathbf{y}) + (\alpha/2)\|\mathbf{y} - \mathbf{x}\|_2^2\}$  imply that  $\mathbf{0} \in \partial f(\mathbf{x}^*)$  and  $-\alpha(\mathbf{x}^+ - \mathbf{x}) \in \partial f(\mathbf{x}^+)$ .

Since  $f$  is  $\rho$ -strongly convex in  $B(\mathbf{x}^*, r)$ , Lemma F.1 forces

$$\begin{aligned} \rho\|\mathbf{x}^+ - \mathbf{x}^*\|_2^2 &\leq \langle -\alpha(\mathbf{x}^+ - \mathbf{x}) - \mathbf{0}, \mathbf{x}^+ - \mathbf{x}^* \rangle = -\alpha\|\mathbf{x}^+ - \mathbf{x}^*\|_2^2 - \alpha\langle \mathbf{x}^* - \mathbf{x}, \mathbf{x}^+ - \mathbf{x}^* \rangle \\ &\leq -\alpha\|\mathbf{x}^+ - \mathbf{x}^*\|_2^2 + \alpha\|\mathbf{x}^* - \mathbf{x}\|_2\|\mathbf{x}^+ - \mathbf{x}^*\|_2 \end{aligned}$$

and thus  $\|\mathbf{x}^+ - \mathbf{x}^*\|_2 \leq \frac{\alpha}{\alpha + \rho}\|\mathbf{x} - \mathbf{x}^*\|_2$ . □

**Lemma F.4** (Neumann expansion). *Let  $\|\cdot\|$  be a submultiplicative matrix norm with  $\|\mathbf{I}\| = 1$ .*

*1. When  $\|\mathbf{M}\| < 1$ , we have  $(\mathbf{I} - \mathbf{M})^{-1} = \sum_{j=0}^{\infty} \mathbf{M}^j = \mathbf{I} + \mathbf{M} + \mathbf{M}(\mathbf{I} - \mathbf{M})^{-1}\mathbf{M}$ ,  $\|\mathbf{I} - \mathbf{M}\| \leq 1/(1 - \|\mathbf{M}\|)$  and  $\|(\mathbf{I} - \mathbf{M})^{-1} - (\mathbf{I} + \mathbf{M})\| \leq \|\mathbf{M}\|^2/(1 - \|\mathbf{M}\|)$ .*

**Lemma F.5.** *Let  $\mathbf{S} \succ 0$  and  $\alpha \geq 0$  be deterministic,  $\mathbf{A} = (\mathbf{I} + \alpha\mathbf{S}^{-1})^{-1}$ ,  $\{\mathbf{u}_i\}_{i=1}^n \subseteq \mathbb{R}^d$  be*

i.i.d. sub-gaussian random vectors with zero mean and covariance matrix  $\mathbf{A}$ ,  $\bar{\mathbf{u}} = \frac{1}{n} \sum_{i=1}^n \mathbf{u}_i$  and  $\hat{\mathbf{A}} = \frac{1}{n} \sum_{i=1}^n \mathbf{u}_i \mathbf{u}_i^T$ . Then the following hold:

1. There exists some positive constant  $C$  that only depends on  $\|\mathbf{u}_1\|_{\psi_2}$ , such that  $\mathbb{P}(\|\bar{\mathbf{u}}\|_2 > 1/10) \leq e^{-n/C}$  and  $\mathbb{P}(\|\hat{\mathbf{A}} - \mathbf{A}\|_2 > 1/10) \leq e^{-n/C}$  hold under any one of the two conditions holds: (i)  $n \geq Cd$ ; (ii)  $\alpha \geq C\text{Tr}(\mathbf{S})/n$ .
2. Suppose that  $n \geq cd$  for some constant  $c > 0$ . There exist positive constants  $C'_1$  and  $C'_2$  such that  $\mathbb{P}(\max\{\|\bar{\mathbf{u}}\|_2, \|\hat{\mathbf{A}} - \mathbf{A}\|_2\} \geq C'_2 \sqrt{d/n}) \leq 2e^{-C'_1 d}$ .

**Proof of Lemma F.5.** Let  $\lambda_1 > \dots > \lambda_d > 0$  be the eigenvalues of  $\mathbf{S}$ , and  $C > 0$  be some constant to be determined. When  $n \geq Cd$ , the fact  $\|\mathbf{A}\|_2 = 1/(1 + \alpha/\lambda_1) \leq 1$  implies  $n \geq C\text{Tr}(\mathbf{A})$ . Also, if  $\alpha \geq C\text{Tr}(\mathbf{S})/n$ , then the crude estimate

$$\text{Tr}(\mathbf{A}) = \sum_{j=1}^d \frac{1}{1 + \alpha/\lambda_j} \leq \sum_{j=1}^d \frac{1}{\alpha/\lambda_j} = \sum_{j=1}^d \frac{\lambda_j}{\alpha} = \frac{\text{Tr}(\mathbf{S})}{\alpha} \leq \frac{n}{C_0}$$

also leads to  $n \geq C\text{Tr}(\mathbf{A})$ . Hence it suffices to find some proper  $C$  and show the desired results given  $n \geq C\text{Tr}(\mathbf{A})$ .

Now we prove the first statement. We first study concentration of the sample mean vector  $\bar{\mathbf{u}}$ . Since  $\bar{\mathbf{u}}$  is a sub-gaussian random vector with covariance matrix  $n^{-1}\mathbf{A}$ , Theorem 2.1 in [Hsu et al. \(2012\)](#) asserts the existence of a constant  $c_1 > 0$  such that

$$\mathbb{P} \left[ \|\bar{\mathbf{u}}\|_2^2 \leq c_1 n^{-1} \left( \text{Tr}(\mathbf{A}) + 2\sqrt{\text{Tr}(\mathbf{A}^2)} + 2\|\mathbf{A}\|_2 t \right) \right] \geq 1 - e^{-t}, \quad \forall t > 0. \quad (\text{F.2})$$

Choose any constant  $C_1 \geq 500c_1$ . Let  $t = n/C_1$ , and suppose that  $n \geq C_1\text{Tr}(\mathbf{A})$ . Using

$\|\mathbf{A}\|_2 \leq 1$  and  $\text{Tr}(\mathbf{A}^2) \leq \text{Tr}(\mathbf{A})\|\mathbf{A}\|_2 \leq \text{Tr}(\mathbf{A})$ , we get

$$\begin{aligned} c_1 n^{-1} \left( \text{Tr}(\mathbf{A}) + 2\sqrt{\text{Tr}(\mathbf{A}^2)t} + 2\|\mathbf{A}\|_2 t \right) &\leq \frac{c_1}{C_1} + \frac{2c_1\sqrt{\text{Tr}(\mathbf{A})t}}{n} + \frac{2c_1 t}{n} \\ &= \frac{c_1}{C_1} + \frac{2c_1\sqrt{\text{Tr}(\mathbf{A})n/C_1}}{n} + \frac{2c_1(n/C_1)}{n} = \frac{3c_1}{C_1} + 2c_1\sqrt{\frac{\text{Tr}(\mathbf{A})}{C_1 n}} \leq \frac{5c_1}{C_1} \leq \frac{1}{10^2}. \end{aligned}$$

Hence  $\mathbb{P}(\|\bar{\mathbf{u}}\|_2 > 1/10) \geq e^{-n/C_1}$ .

Now we come to concentration of the sample covariance matrix  $\hat{\mathbf{A}}$ . Let  $r(\mathbf{A}) = \text{Tr}(\mathbf{A})/\|\mathbf{A}\|_2$ . According to Theorem 9 in [Koltchinskii et al. \(2017\)](#), there exists a constant  $c_2 \geq 1$  such that the following holds: for any  $t \geq 1$ , with probability at least  $1 - e^{-t}$  we have

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq c_2 \|\mathbf{A}\|_2 \max \left\{ \sqrt{\frac{r(\mathbf{A})}{n}}, \frac{r(\mathbf{A})}{n}, \sqrt{\frac{t}{n}}, \frac{t}{n} \right\}. \quad (\text{F.3})$$

Note that the upper bound above can be rewritten as

$$c_2 \max \left\{ \sqrt{\frac{\|\mathbf{A}\|_2 \text{Tr}(\mathbf{A})}{n}}, \frac{\text{Tr}(\mathbf{A})}{n}, \|\mathbf{A}\|_2 \sqrt{\frac{t}{n}}, \|\mathbf{A}\|_2 \frac{t}{n} \right\} \leq c_2 \max \left\{ \sqrt{\frac{\text{Tr}(\mathbf{A})}{n}}, \frac{\text{Tr}(\mathbf{A})}{n}, \sqrt{\frac{t}{n}}, \frac{t}{n} \right\}. \quad (\text{F.4})$$

Let  $C_2 = 100c_2^2$ . When  $n \geq C_2 \text{Tr}(\mathbf{A})$ , by taking  $t = n/C_2$  we get  $\mathbb{P}(\|\hat{\mathbf{A}} - \mathbf{A}\|_2 > 1/10) \leq e^{-n/C_2}$ . The proof of the first statement is then finished by taking  $C = \max\{C_1, C_2\}$ .

We proceed to prove the second statement. Let  $t = C'_1 d$  for some constant  $C'_1$ . Note that  $d \geq \text{Tr}(\mathbf{A})$  and  $t \geq C'_1 \text{Tr}(\mathbf{A})$ . According to [\(F.2\)](#), [\(F.3\)](#) and [\(F.4\)](#), we obtain that with probability at least  $1 - 2e^{-C'_1 d}$ ,  $\|\bar{\mathbf{u}}\|_2 \leq \tilde{C} \sqrt{d/n}$  and  $\|\hat{\mathbf{A}} - \mathbf{A}\|_2 \leq \tilde{C} \max \left\{ \sqrt{d/n}, d/n \right\}$  hold with some constant  $\tilde{C}$ . Since  $n \geq cd$ , we have

$$\max \left\{ \sqrt{d/n}, d/n \right\} \leq \max\{1, 1/\sqrt{c}\} \sqrt{d/n}.$$

By combining the inequalities above, we obtain that

$$\mathbb{P}(\max\{\|\bar{\mathbf{u}}\|_2, \|\hat{\mathbf{A}} - \mathbf{A}\|_2\} \geq C'_2 \sqrt{d/n}) \leq 2e^{-C'_1 d}$$

with  $C'_2 = \tilde{C} \max\{1, 1/\sqrt{c}\}$ . □

## References

Battey, H., Fan, J., Liu, H., Lu, J., and Zhu, Z. (2018). Distributed testing and estimation under sparse high dimensional models. *The Annals of Statistics*, 46(3):1352–1382.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122.

Hsu, D., Kakade, S., and Zhang, T. (2012). A tail inequality for quadratic forms of sub-gaussian random vectors. *Electronic Communications in Probability*, 17.

Jordan, M. I., Lee, J. D., and Yang, Y. (2019). Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 114(526):668–681.

Koltchinskii, V., Lounici, K., et al. (2017). Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23(1):110–133.

Nesterov, Y. (2013). *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media.

Parikh, N. and Boyd, S. (2014). Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239.

Shamir, O., Srebro, N., and Zhang, T. (2014). Communication-efficient distributed optimization using an approximate Newton-type method. In *International Conference on Machine Learning*, pages 1000–1008.

Wang, S., Roosta-Khorasani, F., Xu, P., and Mahoney, M. W. (2018). Giant: Globally improved approximate newton method for distributed optimization. In *Advances in Neural Information Processing Systems*, pages 2338–2348.

Zhang, Y., Duchi, J. C., and Wainwright, M. J. (2013). Communication-efficient algorithms for statistical optimization. *The Journal of Machine Learning Research*, 14:3321–3363.