

# On the Optimality of Averaging in Distributed Statistical Learning

JONATHAN ROSENBLATT\*,

\*`jonathan.rosenblatt@weizmann.ac.il`

BOAZ NADLER

`boaz.nadler@weizmann.ac.il`

Department of Computer Science and Applied Mathematics,  
Weizmann Institute of Science,  
Rehovot, Israel

June 11, 2015

A common approach to statistical learning with big-data is to randomly split it among  $m$  machines and learn the parameter of interest by averaging the  $m$  individual estimates. In this paper, focusing on empirical risk minimization, or equivalently M-estimation, we study the statistical error incurred by this strategy. We consider two large-sample settings: First, a classical setting where the number of parameters  $p$  is fixed, and the number of samples per machine  $n \rightarrow \infty$ . Second, a high-dimensional regime where both  $p, n \rightarrow \infty$  with  $p/n \rightarrow \kappa \in (0, 1)$ . For both regimes and under suitable assumptions, we present *asymptotically exact* expressions for this estimation error. In the fixed- $p$  setting, under suitable assumptions, we prove that to leading order averaging is *as accurate as* the centralized solution. We also derive the second order error terms, and show that these can be non-negligible, notably for non-linear models. The high-dimensional setting, in contrast, exhibits a qualitatively different behavior: data splitting incurs a first-order accuracy loss, which to leading order increases linearly with the number of machines. The dependence of our error approximations on the number of machines traces an interesting accuracy-complexity tradeoff, allowing the practitioner an informed choice on the number of machines to deploy. Finally, we confirm our theoretical analysis with several simulations.

# 1. Introduction

The Big-data era, characterized by huge datasets and an appetite for new scientific and business insights, often involves learning statistical models of great complexity. Typically, the storage and analysis of such data cannot be performed on a single machine. Several platforms such as Map-Reduce [Dean and Ghemawat, 2008], Hadoop [Shvachko et al., 2010], and Spark [Zaharia et al., 2010] have thus become standards for distributed learning with big-data.

These platforms allow learning in an “embarrassingly parallel” scheme, whereby a large dataset with  $N$  observations is split to  $m$  machines, each having access to only a subset of  $n = N/m$  samples. Approaches to “embarrassingly-parallel” learning can roughly be categorized along the output of each machine: predictions, parameters or gradients. In this paper we consider the second, whereby each of the  $m$  individual machines fits a model with  $p$  parameters and transmits them to a central node for merging. This split-and-merge strategy, advocated by McDonald et al. [2009] for striking the best balance between accuracy and communication, is both simple to program and communication efficient: only a single round of communication is performed and only to a central node. It is restrictive in that machines do not communicate between themselves, and splitting is done only along observations and not along variables. For an overview of more general distributed learning strategies see for example Bekkerman et al. [2011].

Our focus is on the *statistical properties* of this split-and-merge approach, under the assumption that the data are split uniformly at random among the  $m$  machines. In particular, we study the simplest merging strategy, of averaging the  $m$  individual machine estimates, denoted as the Mixture Weight Method in McDonald et al. [2009]. In this context we ask the following questions: (i) what is the estimation error of simple averaging as compared to a centralized solution? (ii) what is its distribution? (iii) under which criteria, if any, is averaging optimal? and (iv) how many machines to deploy?

McDonald et al. were among the first to study some of these issues for multinomial regression (a.k.a. Conditional Maximum Entropy), deriving finite sample bounds on the expected error of the averaged estimator [McDonald et al., 2009, Theorem 3]. In a follow-up work, Zinkevich et al. [2010] compared the statistical properties of the averaged estimator to the centralized one for more general learning tasks, assuming each machine estimates the model parameters by stochastic gradient descent. More recently, under appropriate conditions and for a large class of loss functions, Zhang et al. [2013b, Theorem 1] derived bounds for the leading order term in the mean squared error (MSE) of the averaged estimator and provided the rates of higher order terms. They further proposed several improvements to the simple averaging strategy that reduce the second order term in the MSE, and reduce the machine-wise run time via modified optimization algorithms.

In this paper we extend and generalize these previous works in several aspects. First, in Section 3 we study the statistical properties of the averaged estimator, when the number of parameters  $p$  is fixed, under conditions similar to those of Zhang et al. [2013b]. Using the classical statistical theory of M-estimators [Vaart, 1998, Rieder, 2012], we provide not only asymptotic bounds on the MSE, but rather an asymptotic expansion

of the error itself. This allows us to derive the exact constants in the MSE expansion, and prove that as  $n \rightarrow \infty$ , the MSE of the averaging strategy in fact equals that of the centralized solution. Put differently, for various learning tasks, when the number of machines  $m$  and their available memory are such that in each machine there are many observations per parameter ( $n \gg p$ ), then averaging machine-wise estimates is as accurate as the centralized solution. Furthermore, if the centralized solution enjoys first-order statistical properties such as efficiency or robustness, then so will the parallelized solution. We remark that for maximum likelihood problems, independently of our work, the asymptotic agreement between centralized and averaged estimators was also noted by Liu and Ihler [2014]. The asymptotic representation of the averaged estimator also readily yields its limiting distribution. This allows to construct confidence intervals, perform hypothesis tests on the unknown parameters and feature selection without the need for computationally intensive procedures such as Bootstrapping.

The first-order equivalence between the averaged and centralized estimators may seem as a free lunch: run-time speedups with no accuracy loss. Distributed estimation via split-and-average, however, does incur an accuracy loss captured in the higher order error terms. The classical theory of M-estimators permits the derivation of these terms, in principle up to an arbitrary order. We do so explicitly up to second order, revealing the accuracy loss of split-and-average schemes.

In Section 4 we consider the statistical effects of data-splitting in a high-dimensional regime, where the model dimension  $p$ , grows with the number of observations  $n$ :  $p, n \rightarrow \infty$  with  $p/n \rightarrow \kappa \in (0, 1)$ . Our motivation comes from modern day data analysis practices, where increasingly complicated models are considered as more data is made available. This is a challenging regime in that typically machine-wise estimates are not only inconsistent, but in fact do not even converge to deterministic quantities. Here, in the absence of a general theory of M-estimators, we restrict our analysis to generative linear models. In contrast to the fixed- $p$  setting, in this high-dimensional regime there is a first order accuracy loss due to the split data, which increases (approximately) linearly with the number of machines. Luckily, in several practical situations, this accuracy loss is moderate. Our analysis builds upon the recent results of El Karoui et al. [2013] and Donoho and Montanari [2013], and to the best of our knowledge, is the first to study the error loss of parallelization in this high-dimensional regime.

In Section 5 we present several simulations both in the fixed- $p$  and in the high-dimensional regime that illustrate the utility but also the limitations of our results. These confirm that when learning linear models with abundant data, random splitting and averaging is attractive both computationally and statistically. In contrast, for non-linear models, the accuracy loss due to data splitting can be considerable.

In Section 6 we attend to practical considerations, such as parallelization vs. sub-sampling and the choice of number of machines,  $m$ . For the latter, we distinguish between parallelization due to memory constraints, and that motivated by run-time speedups. For these two scenarios we formulate the choice of  $m$  as optimization problems constrained on the desired error level. Interestingly, when motivated by run-time speedups, using our approximations for the estimation error, and varying  $m$  traces the accuracy-complexity tradeoff facing the practitioner.

We conclude with a discussion and several further insights in Section 7. All proofs appear in the appendices.

## 2. Problem Setup

We consider the following general statistical learning setup: Let  $Z$  be a random variable defined on an instance space  $\mathcal{Z}$  and having an unknown density  $p_Z$ . Also, let the parameter space  $\Theta \subset \mathbb{R}^p$  be an open convex subset of Euclidean space, and let  $f : \mathcal{Z} \times \Theta \rightarrow \mathbb{R}^+$  denote a loss function. Our interest is to estimate the  $p$ -dimensional parameter  $\theta^* \in \Theta$  that minimizes the population risk

$$R(\theta) = \mathbb{E}_Z[f(Z, \theta)] = \int f(z, \theta) p_Z(z) \, dz. \quad (1)$$

In the following, we assume that  $\theta^*$  exists in  $\Theta$  and is unique. Given  $N$  i.i.d. samples  $\{z_i\}_{i=1}^N$  of the r.v.  $Z$ , a standard approach, known as M-estimation or empirical risk minimization (ERM), is to calculate the estimator  $\hat{\theta}_N \in \Theta$  that minimizes the empirical risk

$$\hat{R}_N(\theta) = \frac{1}{N} \sum_{i=1}^N f(z_i, \theta). \quad (2)$$

This framework covers many common unsupervised and supervised learning tasks. In the latter,  $Z = (X, Y)$  consists of both features  $X$  and labels  $Y$ . There is by now an established theory providing conditions for  $\hat{\theta}_N$  to be a consistent estimator of  $\theta^*$ , and non asymptotic bounds on its finite sample deviation from  $\theta^*$  (see Devroye et al. [1997], Shalev-Shwartz and Ben-David [2014] and references therein).

In this paper we consider a big-data setting, whereby the number of samples  $N$  is so large that instead of minimizing Eq.(2) on a single machine, the data is randomly allocated among  $m$  machines, each having access to only a subset of size  $n := N/m$ . In line with the Map-Reduce workflow, a typical approach in this distributed scenario is that each machine computes its own M-estimator and transmits it to a central node for further processing. In this work we focus on the most common aggregation procedure, namely simple averaging

$$\bar{\theta} := \frac{1}{m} \sum_{j=1}^m \hat{\theta}_n^{(j)} \quad (3)$$

where  $\hat{\theta}_n^{(j)}$  denotes the  $j$ -th machine minimizer of Eq. (2) over its own observed data.

Our questions of interest are: (i) what is the accuracy of  $\bar{\theta}$  vs. that of  $\hat{\theta}_N$ ? (ii) what are the statistical properties of  $\bar{\theta}$ ? (iii) under which criteria, if any, is  $\bar{\theta}$  optimal? and (iv) how many machines to deploy?

## 3. Fixed- $p$ Setting

First, we consider the error of the split-and-average estimator  $\bar{\theta}$  of Eq.(3), when data is abundant and the model dimension  $p$  and number of machines  $m$  are both fixed. In

this setting, bounds on the  $MSE[\bar{\theta}, \theta^*] := \mathbb{E}[\|\bar{\theta} - \theta^*\|^2]$  were derived by both Zhang et al. [2013b] and McDonald et al. [2009]. For the particular case of maximum likelihood estimation, Liu and Ihler [2014, Theorem 4.6] derived the exact asymptotic expression of the first two leading error terms in the MSE, as  $n \rightarrow \infty$ . We take a similar approach but for the more general M-estimators. Instead of focusing on the MSE, we derive an exact asymptotic representation of the first two terms in the error  $\bar{\theta} - \theta^*$  itself.

### 3.1. First Order Statistical Properties of Averaging

We start by analyzing the exact asymptotic expression for the dominant error term. We make the following standard assumptions [Vaart, 1998, Theorem 5.23], similar to those made in Zhang et al. [2013b]:

#### Assumption Set 1.

- A1  $\hat{\theta}_n$  is consistent:  $\hat{\theta}_n = \theta^* + o_P(1)$ .
- A2  $R(\theta)$  admits a second order Taylor expansion at  $\theta^*$  with non singular Hessian  $V_{\theta^*}$ .
- A3  $f(Z, \theta)$  is differentiable at  $\theta^*$  almost surely (a.s.) or in probability.
- A4  $f(Z, \theta)$  is Lipschitz near  $\theta^*$ :  $|f(Z, \theta_1) - f(Z, \theta_2)| \leq M(Z)\|\theta_1 - \theta_2\|$  with Lipschitz coefficient  $M(Z)$  bounded in squared expectation,  $\mathbb{E}[M(Z)^2] < \infty$ .

Our first result, formally stated in the following theorem, is that under Assumption Set 1 averaging machine-wise estimates enjoys the *same first-order statistical properties* as the centralized solution.

**Theorem 1.** *Under Assumption Set 1, as  $n \rightarrow \infty$  with  $p$  fixed, and any norm*

$$\frac{\|\bar{\theta} - \theta^*\|}{\|\hat{\theta}_N - \theta^*\|} = 1 + o_P(1). \quad (4)$$

We say that two estimators are *first-order equivalent* if their leading error terms converge to the same limit at the same rate, with the same limiting distribution. Assumption Set 1 implies that  $\hat{\theta}_n$  converges to  $\theta^*$  at rate  $O(n^{-1/2})$  [Vaart, 1998, Corollary 5.53]. Theorem 1 thus directly implies the following:

**Corollary 1.** *The averaged estimator  $\bar{\theta}$  is first-order equivalent to the centralized solution  $\hat{\theta}_N$ .*

**Remark 1.** *In practice, Eq.(2) is minimized only approximately, typically by some iterative scheme such as gradient descent (GD), stochastic gradient descent (SGD), etc. An important point is that Theorem 1 holds not only for the exact empirical minimizer  $\hat{\theta}_n$  of Eq.(2), but also for any approximate minimizer  $\tilde{\theta}_n$  as long as it satisfies  $\hat{R}_n(\tilde{\theta}_n) \leq \hat{R}_n(\theta_n) + o_p(n^{-1})$  [Vaart, 1998, Theorem 5.23]. In other words, for Corollary 1 to hold, it suffices to minimize the empirical risk up to  $o_p(n^{-1})$  precision.*

Theorem 1 has important implications on the statistical properties of  $\bar{\theta}$ , its optimality and robustness. We discuss these in detail below, but before, let us describe the scope which this theorem covers.

**Scope** As detailed further in Appendix H, the learning tasks covered by Theorem 1 are quite broad, and include: linear or non-linear regression with  $l_2$ , Huber, or log likelihood loss; linear or non-linear quantile regression with continuous predictors; binary regression where  $P(Y = 1|X) = \Psi_\theta(X)$  for any smooth  $\Psi_\theta$  and  $l_2$ , log likelihood or Huberized hinge loss<sup>1</sup>; binary hinge loss regression (i.e. SVM regression) with continuous predictors; unsupervised learning of location and scale. Furthermore, Theorem 1 also covers regularized risk minimization with a *fixed* regularization term  $J(\theta)$ , of the form  $\theta^* := \operatorname{argmin}_\theta \{R(\theta) + J(\theta)\}$ , provided that the modified loss function  $\tilde{f}(Z, \theta) = f(Z, \theta) + J(\theta)$  satisfies the required assumptions.

Some learning problems, however, are not covered by Theorem 1. Examples include: non-uniform allocation of samples to machines; non-convex parameter spaces; a data driven regularization term; non differentiable loss with discrete predictors. Also not covered is the  $n < p$  regime, in which Shamir et al. [2013] showed that averaging (denoted there as One Shot Averaging) can, in general, be unboundedly worse than the centralized solution.

**On the optimality of averaging.** Recall that common notions of asymptotic optimality, such as Best Regular and Local Minimax depend only on the leading order error term [Vaart, 1998, Chapter 8]. Hence, if the centralized estimator  $\hat{\theta}_N$  is optimal w.r.t. any of these criteria, Eq.(4) readily implies that so is the averaged estimate  $\bar{\theta}$ . A notable example, discussed in [Zhang et al., 2013b, Corollary 3] and in Liu and Ihler [2014], is when the loss function is the negative log likelihood of the generative model. The centralized solution, being the maximum-likelihood estimate of  $\theta^*$ , is optimal in several distinct senses. Theorem 1 thus implies that  $\bar{\theta}$  is optimal as well and the factor 1 in Eq.(4) cannot be improved.

**Robustness.** An important question in distributed learning is how to handle potential outliers: should these be dealt with at the machine-level, the aggregation level, or both? Recall that the robustness literature mostly considered the construction of estimators having minimal asymptotic variance, under the constraint of bounded influence of individual observations. For estimating the mean of a Gaussian distribution under possible contamination, Huber derived his famous loss function, and proved it to be optimal. As the Huber-loss yields an M-estimator that satisfies the assumptions of Theorem 1, it thus follows that averaging machine-wise robust estimators is optimal in the same sense.

Hence, if the probability of a high proportion of outliers in any machine is negligible, and machine-failure is not a concern, it suffices to deal with outliers at the machine level alone. In other cases robust aggregation functions should be considered [Hsu and Sabato, 2013, Feng et al., 2014].

---

<sup>1</sup>A smooth version of the Huber loss [Rosset and Zhu, 2007].

**Asymptotic Linearity.** The proof of Theorem 1 relies on the asymptotic linearity of the estimator in some non-linear transformation of the samples. This is known as the *asymptotic linearity property* and the corresponding transformation is the *Influence Function*. Asymptotic linearity holds for several other estimators, including L, R and Minimum Distance. Hence, first-order equivalence of averaging to the centralized solution is rather general. It typically holds for asymptotically Gaussian estimators [Rieder, 2012, Chapter 1,6] and has also been observed in other contexts, such as that of particle filters [Achutegui et al., 2014].

**Limiting Distribution** The asymptotic linearity of  $\bar{\theta}$  in the influence function immediately offers the following limiting Gaussian distribution:

**Corollary 2** (Asymptotic Normality). *Under the assumptions of Theorem 1, when  $n \rightarrow \infty$  with  $p$  fixed, then  $\sqrt{N}(\bar{\theta} - \theta^*)$  converges in distribution to*

$$\mathcal{N}(0, V_{\theta^*}^{-1} \mathbb{E} [\nabla f(\theta^*) \nabla f(\theta^*)'] V_{\theta^*}^{-1}).$$

Corollary 2 allows to construct confidence intervals and test hypotheses on the unknown  $\theta^*$ . To this end, the asymptotic covariance matrix also needs to be estimated. Plugging any  $O(N^{-1/2})$  consistent estimator for the covariance matrix will conserve the asymptotic normality via Slutsky's Theorem.

## 3.2. Second Order Terms

As we show empirically in Section 5, relatively little accuracy is lost when parallelizing a linear model but much can be lost when the model is non-linear. One reason is that the second order error term may be non-negligible. As discussed in Section 6, this term is also imperative when deciding how many machines to deploy, as the first-order approximation of the error does not depend on  $m$  for fixed  $N$ .

Before studying this second order term, let us provide a high level view. Intuitively, the first-order term captures estimation variance, which is reduced by averaging. The second order term captures also bias, which is not reduced by averaging. We would thus expect some second order suboptimality when parallelizing. Indeed, Theorem 2 below shows that the (second order) bias in a parallelized estimator is  $m$  times larger than that of the centralized one. The comparison between the second order MSE matrix of the parallelized and centralized estimators is more complicated. Theorem 3 provides an explicit expression, whose terms ultimately depend on the curvature of the risk  $R(\theta)$  at  $\theta^*$ .

### 3.2.1. Notation and Assumptions

To study the second order error of  $\bar{\theta}$ , we make suitable assumptions that ensure that the machine-wise M-estimator admits the following higher-order expansion,

$$\hat{\theta}_n = \theta^* + \xi_{-1/2}(\hat{\theta}_n) + \xi_{-1}(\hat{\theta}_n) + \xi_{-3/2}(\hat{\theta}_n) + O_P(n^{-2}), \quad (5)$$

where  $\xi_{-\alpha}(\hat{\theta}_n)$  denotes the  $O_P(n^{-\alpha})$  error term in  $\hat{\theta}_n$  and  $\alpha = \{1/2, 1, 3/2, \dots\}$ . The following set of assumptions with  $s = 4$  is sufficient for Eq.(5) to hold, see Rilstone et al. [1996].

**Assumption Set 2.** *There exist a neighborhood of  $\theta^*$  in which all of the following conditions hold:*

*B1 Local differentiability:*  $\nabla^s f(\theta, Z)$  up to order  $s$ , exist a.s. and  $\mathbb{E}[\|\nabla^s f(\theta^*, Z)\|] < \infty$ .

*B2 Bounded empirical Hessian:*  $(\nabla^2 \hat{R}_n(\theta))^{-1} = O_P(1)$ .

*B3 Lipschitz gradients:*  $\|\nabla^s f(\theta, Z) - \nabla^s f(\theta^*, Z)\| \leq M\|\theta - \theta^*\|$ , where  $\mathbb{E}[|M|] \leq C < \infty$ .

For future use, and following the notation in Rilstone et al. [1996], we define the following  $p \times 1$  column vector  $\delta$ , and  $p \times p$  matrices  $\gamma_0, \dots, \gamma_4$ ,

$$\begin{aligned} \mathbb{E}[\xi_{-1}(\hat{\theta}_n)] &= n^{-1}\delta; & \mathbb{E}[\xi_{-1}(\hat{\theta}_n)] \mathbb{E}[\xi'_{-1}(\hat{\theta}_n)] &= n^{-2}\gamma_0 = n^{-2}\delta\delta'; \\ \mathbb{E}[\xi_{-1/2}(\hat{\theta}_n) \xi'_{-1/2}(\hat{\theta}_n)] &= n^{-1}\gamma_1; & \mathbb{E}[\xi_{-1}(\hat{\theta}_n) \xi'_{-1/2}(\hat{\theta}_n)] &= n^{-2}\gamma_2; \\ \mathbb{E}[\xi_{-1}(\hat{\theta}_n) \xi'_{-1}(\hat{\theta}_n)] &= n^{-2}\gamma_3 + o(n^{-2}); & \mathbb{E}[\xi_{-3/2}(\hat{\theta}_n) \xi'_{-1/2}(\hat{\theta}_n)] &= n^{-2}\gamma_4 + o(n^{-2}). \end{aligned} \quad (6)$$

### 3.2.2. Second Order Bias

Let  $B_2(\hat{\theta}_n)$  denote the second order bias of  $\hat{\theta}_n$  w.r.t.  $\theta^*$ :

$$B_2(\hat{\theta}_n) := \mathbb{E}[\xi_{-1/2}(\hat{\theta}_n) + \xi_{-1}(\hat{\theta}_n)]. \quad (7)$$

The following theorem, proven in Appendix B, shows that under our assumptions averaging over  $m$  machines is (up to second order)  $m$  times more biased than the centralized solution.

**Theorem 2** (Second Order Bias). *Under Assumption Set 2 with  $s = 3$ ,  $B_2(\hat{\theta}_N) = \delta/N$  and  $B_2(\bar{\theta}) = \delta/n$ , so that*

$$B_2(\bar{\theta}) = m B_2(\hat{\theta}_N). \quad (8)$$

**Remark 2.** *The second order bias  $B_2(\bar{\theta})$  can be reduced at the cost of a larger first-order error, i.e., trading bias for variance. In general, this should be done with caution, since in extreme cases debiasing may inflate variance infinitely [Doss and Sethuraman, 1989]. Approaches to reduce the second order bias include that of Kim [2006] who modifies the machine-wise loss function, Liu and Ihler [2014] who propose a different aggregation of the  $m$  machine-wise estimates, and Zhang et al. [2013b], whose SAVGM algorithm estimates the machine-wise bias via bootstrap. A different approach is to trade bias for communication. Recent works that reduce the bias by allowing communication between the  $m$  machines include the DDPCA algorithm [Meng et al., 2012], which transfers parts of the inverse Hessian between machines and DANE [Shamir et al., 2013], which transfers gradients.*



### 3.2.3. Second Order MSE

Following Rilstone et al. [1996], for any estimator  $\tilde{\theta}_n$  based on  $n$  samples, we denote by  $M_2(\tilde{\theta}_n)$  its second order MSE matrix,

$$\mathbb{E} \left[ (\tilde{\theta}_n - \theta^*)(\tilde{\theta}_n - \theta^*)' \right] = M_2(\tilde{\theta}_n) + o(n^{-2}). \quad (9)$$

It follows from Rilstone et al. [1996, Proposition 3.4] that under Assumption Set 2 with  $s = 4$

$$M_2(\hat{\theta}_n) = \frac{1}{n}\gamma_1 + \frac{1}{n^2}(\gamma_2 + \gamma'_2 + \gamma_3 + \gamma_4 + \gamma'_4). \quad (10)$$

The following theorem compares between  $M_2(\hat{\theta}_N)$  and  $M_2(\bar{\theta})$ .

**Theorem 3** (Second Order MSE). *Under Assumption Set 2 with  $s = 4$ , the matrix  $M_2(\bar{\theta})$  is given by*

$$M_2(\bar{\theta}) = \frac{m-1}{m} \frac{1}{n^2} \gamma_0 + \frac{1}{mn} \gamma_1 + \frac{1}{mn^2} (\gamma_2 + \gamma'_2 + \gamma_3 + \gamma_4 + \gamma'_4). \quad (11)$$

Furthermore, the excess second order error due to parallelization is given by

$$M_2(\bar{\theta}) - M_2(\hat{\theta}_N) = \frac{m-1}{m} \frac{1}{n^2} \gamma_0 + \frac{m-1}{m^2} \frac{1}{n^2} (\gamma_2 + \gamma'_2 + \gamma_3 + \gamma_4 + \gamma'_4). \quad (12)$$

In general, the second order MSE matrix  $M_2(\hat{\theta}_n)$  of Eq.(10) need not be positive definite (PD) [Rilstone et al., 1996]. Note that since both matrices  $\gamma_0$  and  $\gamma_1$  are PD by definition, a simple condition to ensure that both  $M_2(\hat{\theta}_n)$  and  $M_2(\bar{\theta}) - M_2(\hat{\theta}_N)$  are PD is that  $(\gamma_2 + \gamma'_2 + \gamma_3 + \gamma_4 + \gamma'_4)$  is PD. If this holds, then parallelization indeed deteriorates accuracy, at least up to second order.

**Remark 3.** *Even if the second order MSE matrix is PD, due to higher order terms, parallelization may actually be more accurate than the centralized solution. An example is ridge regression with a fixed penalty and null coefficients (i.e.,  $\theta^* = 0$ ). The regularization term, being fixed, acts more aggressively with  $n$  observations than with  $N$ . The machine-wise estimates are thus more biased towards 0 than the centralized one. As the bias acts in the correct direction,  $\bar{\theta}$  is more accurate than  $\hat{\theta}_N$ . Two remarks are, however, in order: (a) This phenomenon is restricted to particular parameter values and shrinkage estimators. It does not occur uniformly over the parameter space. (b) In practice, the precise regularization penalty may be adapted to account for the parallelization, see for example Zhang et al. [2013a].*

### 3.3. Examples

We now apply our results to two popular learning tasks: ordinary least squares (OLS) and ridge regression, both assuming a generative linear model. We study these two cases not only due to their popularity, but also as they are analytically tractable. As we show below, parallelizing the OLS task incurs no excess bias, but does exhibit excess (second order) MSE. The ridge problem, in contrast, has both excess bias and excess (second order) MSE.

### 3.3.1. OLS

Consider the standard generative linear model  $Y = X'\theta_0 + \varepsilon$ , where the explanatory variable  $X$  satisfies  $\mathbb{E}[X] = 0$ ;  $\text{Var}[X] = \Sigma$ , and the noise  $\varepsilon$  is independent of  $X$  with mean zero and  $\text{Var}[\varepsilon] = \sigma^2$ . The loss is  $f(Y, X; \theta) = \frac{1}{2}(Y - X'\theta)^2$ , whose risk minimizer is the generative parameter,  $\theta^* = \theta_0$ . The following proposition, proved in Appendix D, provides explicit expressions for the second order MSE matrix.

**Proposition 1** (OLS Error Moments). *For the OLS problem, under the above generative linear model,*

$$\begin{aligned}\gamma_0 &= 0, & \gamma_1 &= \sigma^2 \Sigma^{-1}, & \gamma_2 &= -(1+p)\sigma^2 \Sigma^{-1}, \\ \gamma_3 &= (1+p)\sigma^2 \Sigma^{-1}, & \gamma_4 &= (1+p)\sigma^2 \Sigma^{-1}.\end{aligned}$$

Inserting these expressions into Theorem 2 yields that the second order bias vanishes both for the individual machine-wise estimators and for their average, i.e.,  $B_2(\hat{\theta}_n) = B_2(\hat{\theta}_N) = B_2(\bar{\theta}) = 0$ . Combining Proposition 1 with Theorem 3 yields the following expressions for the parallelized second order MSE and the excess error,

$$M_2(\bar{\theta}) = \frac{1}{mn}\sigma^2 \Sigma^{-1} + \frac{1}{mn^2}(1+p)\sigma^2 \Sigma^{-1}; \quad M_2(\bar{\theta}) - M_2(\hat{\theta}_N) = \frac{m-1}{mn^2}(1+p)\sigma^2 \Sigma^{-1}. \quad (13)$$

In OLS, parallelization thus incurs a second order accuracy loss, since  $M_2(\bar{\theta}) - M_2(\hat{\theta}_N)$  is PD.

### 3.3.2. Ridge Regression

Next, we analyze ridge regression under the same generative model  $Y = X'\theta_0 + \varepsilon$ ;  $\mathbb{E}[X] = 0$ ;  $\text{Var}[X] = \Sigma$ , but now with the ridge penalty  $f(Y, X; \theta) = \frac{1}{2}(Y - X'\theta)^2 + \frac{\lambda}{2}\|\theta\|^2$ . The risk minimizer  $\theta^*$  now equals  $(\Sigma + \lambda I)^{-1}\Sigma\theta_0$ .

Adding the simplifying assumption that  $\Sigma = I$ , and denoting  $\lambda_{k,l} := \frac{\lambda^k}{(1+\lambda)^l}$ ,  $B := \theta_0\theta_0'$ , and  $A := \text{Tr}(B)I = \|\theta_0\|^2 I$ , we obtain the following result.

**Proposition 2** (Ridge Error Moments). *For the ridge regression problem, under the above conditions, the matrices  $\gamma_0, \dots, \gamma_4$  that control the second order bias and MSE of Eq.(6) are given by*

$$\begin{aligned}\gamma_0 &= \lambda_{2,6}(1+p)^2 B, & \gamma_1 &= \lambda_{2,4}(B + A) + \lambda_{0,2}\sigma^2 I, \\ \gamma_2 &= -\lambda_{2,5}((4+p)B + (2+p)A) - \lambda_{0,3}\sigma^2(1+p)I, & \gamma_3 &= \lambda_{2,6}((5+p+p^2)B + (2+p)A) + \lambda_{0,4}\sigma^2(1+p)I \\ \gamma_4 &= \lambda_{2,5}((5+2p)B + (3+2p)A) + \lambda_{0,3}\sigma^2(1+p)I.\end{aligned}$$

**Corollary 3** (Ridge Second Order Bias). *Combining Proposition 2 with Theorem 2, under a linear generative model, the second order bias of the parallelized ridge regression estimate is  $B_2(\bar{\theta}) = -\frac{1}{n}\lambda_{1,3}(1+p)\theta_0$ .*

**Corollary 4** (Ridge Second Order MSE). *Combining Proposition 2 with Theorem 3, under a linear generative model, the parallelized second order MSE matrix and excess MSE are given by*

$$M_2(\bar{\theta}) = \frac{1}{mn} (\lambda_{2,4}(B + A) + \lambda_{0,2}\sigma^2 I) + \frac{m-1}{m} \frac{1}{n^2} \lambda_{2,6}(1+p)^2 B + \frac{1}{mn^2} [\lambda_{2,5}2(p+1)(B + A) + \lambda_{2,6}((5+p+p^2)B + (2+p)A) + \lambda_{0,4}(1+p)\sigma^2 I], \quad (14)$$

and

$$M_2(\bar{\theta}) - M_2(\hat{\theta}_N) = \frac{m-1}{m} \frac{1}{n^2} \lambda_{2,6}(1+p)^2 B + \frac{m-1}{m^2} \frac{1}{n^2} [\lambda_{2,5}2(p+1)(B + A) + \lambda_{2,6}((5+p+p^2)B + (2+p)A) + \lambda_{0,4}(1+p)\sigma^2 I].$$

As in the OLS case, since  $B$  is an outer product, and  $A$  is a scaled identity matrix, it follows that both  $M_2(\bar{\theta})$  and  $M_2(\bar{\theta}) - M_2(\hat{\theta}_N)$  are PD matrices. Despite this result, as discussed in Remark 3, it is still possible that  $\mathbb{E}[(\bar{\theta} - \theta^*)(\bar{\theta} - \theta^*)'] - \mathbb{E}[(\hat{\theta}_N - \theta^*)(\hat{\theta}_N - \theta^*)']$  is a negative definite matrix due to higher order error terms, implying that parallelized ridge regression can be more exact than the centralized estimator. This has been confirmed in simulations (not included).

## 4. High-Dimensional Approximation

As reviewed in Section 1, most existing theory on parallelization assumes a fixed- $p$ , independent of  $n$ . This is implied by the assumption that the empirical risk gradients have uniformly bounded moments, independent of  $n$  [e.g. Zhang et al., 2013b, Assumption 3]. However, it is common practice to enrich a model as more data is made available, to the extent that the number of unknown parameters is comparable to the number of samples. If  $p$  is comparable to  $n$ , and both are large, then the approximations of Section 3 may underestimate the parallelization's excess error. To address this setting, we now perform a high-dimensional analysis where  $n, p(n) \rightarrow \infty$  and  $p(n)/n \rightarrow \kappa \in (0, 1)$ .

To the best of our knowledge there is no general theory for the behavior of M-estimators in this regime. To gain insight into the statistical properties of parallelization in this high-dimensional setting, we restrict our focus to *generative linear* models for which the appropriate theory has been developed only recently. Building on the works of Donoho and Montanari [2013] and El Karoui et al. [2013], we thus consider a random variable  $Z = (X, Y)$  consisting of a vector of predictor variables ( $X \in \mathbb{R}^p$ ) and a scalar response variable ( $Y \in \mathbb{R}$ ), which satisfy the following assumptions:

### Assumption Set 3.

*C1 The observed data  $x_i$  are i.i.d. from the random variable  $X \sim \mathcal{N}(0, \Sigma)$ , with invertible  $\Sigma$ .*

C2 *Linear generative model:  $Y = X'\theta_0 + \epsilon$ , where  $\theta_0 \in \mathbb{R}^p$ .*

C3 *The noise random variable  $\epsilon$  has zero mean, finite second moment, and is independent of  $X$ .*

C4 *The loss  $f(Z, \theta) = f(Y - X'\theta)$  is smooth and strongly convex.*

Unlike the fixed- $p$  case, in the high-dimensional regime where  $p, n \rightarrow \infty$  together, each machine-wise estimate is inconsistent. As shown by El Karoui et al. [2013], and Donoho and Montanari [2013], when Assumption Set 3 holds, then as  $n, p(n) \rightarrow \infty$  with  $p(n)/n \rightarrow \kappa \in (0, 1)$ ,

$$\hat{\theta}_n = \theta^* + r(\kappa) \Sigma^{-1/2} \xi (1 + o_P(1)) \quad (15)$$

where  $\xi \sim \mathcal{N}(0, 1/p \times I_p)$  and  $r(\kappa)$  is a deterministic quantity that depends on  $\kappa$ , on the loss function  $f$  and on the distribution of the noise  $\epsilon$ .

Using the above result, we now show that in contrast to the fixed- $p$  setting, averaging is not even first-order equivalent to the centralized solution. The following lemma, proven in Appendix G.1, quantifies this accuracy loss showing that, typically, it is moderate.

**Lemma 1.** *Under Assumption Set 3, as  $\kappa \rightarrow 0$*

$$\frac{\mathbb{E} [\|\bar{\theta} - \theta^*\|^2]}{\mathbb{E} [\|\hat{\theta}_N - \theta^*\|^2]} = 1 + \kappa \frac{r_2}{r_1} \left(1 - \frac{1}{m}\right) + O(\kappa^2), \quad (16)$$

where

$$r_1 = \frac{B_1}{A_2^2}, \quad r_2 = \frac{3B_1T_1}{A_2^4} - \frac{2B_1^2A_4}{A_2^5} + \frac{2B_2}{A_2^3}, \quad (17)$$

and

$$\begin{aligned} A_2 &= \mathbb{E}[f_{[2]}(\epsilon)], & A_4 &= \mathbb{E}[1/2 f_{[4]}(\epsilon)], & T_1 &= \mathbb{E}[f_{[2]}^2(\epsilon) + f_{[1]}(\epsilon)f_{[3]}(\epsilon)], \\ B_1 &= \mathbb{E}[f_{[1]}^2(\epsilon)], & B_2 &= \mathbb{E}[f_{[1]}^2(\epsilon)f_{[2]}(\epsilon)]. \end{aligned} \quad (18)$$

where  $f_{[i]} := \frac{\partial^i}{\partial t^i} f(t)$ .

**Remark 4.** *For simplicity of exposition, we followed the assumptions of Bean et al. [2013, Result 1]. However, many of these can be relaxed, as discussed by El Karoui [2013]. In particular,  $X$  need not be Gaussian provided that it is asymptotically orthogonal and exponentially concentrating;  $f$  need not be strongly convex nor infinitely differentiable and an interplay is possible between assumptions on the tail mass of  $\epsilon$  and  $f_{[1]}$ . Note however, that for our perturbation analysis on the behavior of  $r(\kappa)$  as  $\kappa \rightarrow 0$  to hold, we assume the loss is at-least six times differentiable with bounded sixth derivative.*

**Remark 5.** *There are cases where  $r(\kappa)$  can be evaluated exactly, without recurring to approximations. One such case is least squares loss with arbitrary noise  $\epsilon$ , satisfying C3, in which Wishart theory gives  $r^2(\kappa) = \frac{\kappa}{1-\kappa}\sigma^2$  [El Karoui et al., 2013]. A second order Taylor approximation of this exact result yields  $r_1 = r_2 = \sigma^2$  in accord with our Eqs.(16) and (17). A second case where an exact formula is available is  $l_1$  loss with Gaussian errors: A second order Taylor expansion of the closed form solution derived in [El Karoui et al., 2013, Page 3] gives  $r_2/r_1 = 0.904$ , again consistent with Eq.(16).*

**Typical Accuracy Losses** In classical asymptotics where  $\kappa \rightarrow 0$ , Eq.(16) is consistent with the results of Section 3 in that splitting the data has no (first-order) cost. In practical high-dimensional scenarios, where the practitioner applies the “no less than five observations per parameter” rule of thumb [Huber, 1973], the resulting value of  $\kappa$  is at most 0.2. The accuracy loss of splitting the data is thus small provided that the ratio  $r_2/r_1$  is small. As shown in Table 1, for several loss functions with either Gaussian or Laplace errors, this ratio is approximately one.

| Loss          | $f(t)$   | Gaussian | Laplace |
|---------------|--|----------|---------|
| Squared       | $t^2/2$  | 1        | 1       |
| Pseudo Huber  | $\delta^2(\sqrt{1 + (t/\delta)^2} - 1) ; \delta = \frac{3}{2}$ | 0.92     | 1.3     |
| Absolute Loss | $ t $  | 0.9      | 1.83    |

Table 1: The ratio  $\frac{r_2}{r_1}$  for different loss functions and noise,  $\epsilon$ , distributions (Gaussian and Laplace). From Eq.(16), small values of  $\frac{r_2}{r_1}$  imply a small accuracy loss when parallelizing.

**Limiting Distribution** Similar to the fixed  $p$  regime, using Eq.(15) we can derive the following limiting distribution of  $\bar{\theta}$  in the high-dimensional regime which is immediate from the results of Bean et al. [2013, p.1 in SI], and the fact that  $\bar{\theta}$  has  $m$  times less variance than  $\hat{\theta}_n$ .

**Corollary 5** (Asymptotic Normality). *Under the assumptions of Lemma 1, for a fixed contrast  $v$ , as  $p, n \rightarrow \infty$  with  $p/n \rightarrow \kappa \in (0, 1)$ , then*

$$\frac{v'\bar{\theta} - v'\theta^*}{r(\kappa)\sqrt{\frac{v'\Sigma^{-1}v}{pm}}} \xrightarrow{D} \mathcal{N}(0, 1).$$

## 5. Simulations

We perform several simulations to validate our results and assess their stability in finite samples. For reproducibility, the R simulation code is available at <https://github.com/johnros/ParalSimulat>.




Figure 1: The error ratio  $\|\bar{\theta} - \theta^*\|_2 / \|\hat{\theta}_N - \theta^*\|_2$  as a function of sample size  $n$  in the fixed- $p$  regime. The center point is the median over 500 replications and the bars represent the median absolute deviation. In all four panels  $p = 50$ . Color encodes the number of machines  $m = 10, 20, 40, 100$ . The learning tasks in the four panels are: (a) ordinary least squares; (b) ridge regression; (c) non-linear least squares; (d) logistic regression. Data was generated as follows:  $X \sim \mathcal{N}(0_p, I_{p \times p})$ ;  $\theta_0 = \tilde{\theta} / \|\tilde{\theta}\|$ , and  $\tilde{\theta}_j = j$  for  $j = 1, \dots, p$ ;  $\varepsilon \sim \mathcal{N}(0, 10)$ ; In (a)-(b), the response was drawn from  $Y = X'\theta_0 + \varepsilon$ ; In (b)  $\lambda = 0.1$ ; In (c)  $Y = \exp(X'\theta_0) + \varepsilon$ , whereas in panel (d),  $P(Y = 1|X) = \frac{1}{1 + \exp(-X'\theta_0)}$ .

We start with the fixed- $p$ , large- $n$  regime. Figure 1 shows the empirical median and median absolute deviation of the individual ratios  $\|\bar{\theta} - \theta^*\|_2 / \|\hat{\theta}_N - \theta^*\|_2$  as a function of sample size  $n$ , with  $N = nm$  growing as well. As seen from this figure, and in accord with Theorem 1, for large  $n$   $\bar{\theta}$  is asymptotically equivalent to  $\hat{\theta}_N$  and the error ratio tends to one. We also see that for small to moderate  $n$ , parallelization may incur a non-negligible excess error, in particular for non-linear models.

Figure 2 presents the empirical bias and MSE of  $\bar{\theta}$  in OLS, as a function of number of machines  $m$  with  $N$  fixed, and compares these to their theoretical approximations from Section 3.3.1. In accord with Proposition 1, the parallelized OLS estimate shows no excess bias. In this OLS case, a high-dimensional approximation of the MSE is identical to the fixed- $p$  in panel (b) so the plot is omitted. We thus conclude that both the fixed- $p$  and the high-dim approximations of the MSE are quite accurate for small  $m$  (i.e., large- $n$ ), but underestimate the error as  $p/n$  departs from 0.

Figure 3 is similar to Figure 2, but for ridge regression. We see that, unlike the OLS problem, the ridge problem does have parallelization bias, as predicted by our analysis in Section 3.3.2. While our fixed- $p$  MSE approximation is accurate for small  $m$  (i.e., large- $n$ ), for larger  $m$  the empirical error is smaller than that predicted by our second order analysis. This suggests that higher order error terms in the MSE matrix are negative definite (see also Remark 3).

Next, we consider the high-dimensional regime. Figure 4 shows  $\mathbb{E}[\|\bar{\theta} - \theta^*\|^2] / \mathbb{E}[\|\hat{\theta}_N - \theta^*\|^2]$  as a function of machine-wise sample size  $n$ , while holding  $\kappa$  and  $m$  fixed. In contrast to the fixed- $p$  regime, here there is a first-order accuracy loss, and even for large  $n$  the MSE ratio does not converge to one. In the OLS case, where our high-dimensional approximations are applicable, they are indeed accurate over a wide range of values of  $n$




Figure 2: Bias and MSE in OLS as a function of number of machines  $m$ , with a fixed total number of samples  $N = 50,000$ , for different dimensions  $p = 50, 70, 80, 100$ , averaged over 1,000 replications. Panel (a) shows, in dotted lines, the mean of the empirical bias of an arbitrary coordinate  $j$  in the parallelized estimator  $\mathbb{E}[\bar{\theta}_j - \theta_j^*]$ . The solid line is the theoretical expression for the second order bias, which, as discussed in Section 3.3.1, is precisely zero. Panel (b) shows, in dotted lines, the empirical mean squared error of the parallelized estimator,  $\mathbb{E}[\|\bar{\theta} - \theta^*\|_2^2]$  as a function of  $m$ . The solid lines are the theoretical approximation using the second order fixed- $p$  expansion  $\text{Tr}(M_2(\bar{\theta}))$ , with  $M_2(\bar{\theta})$  from Eq.(13). In (b) the  $y$ -axis is  $\log_{10}$  scaled. Data was generated as follows:  $X \sim \mathcal{N}(0_p, I_{p \times p})$ ;  $\theta_0 = \tilde{\theta}/(\|\tilde{\theta}\|/10)$  where  $\tilde{\theta}_j = j$  for  $j = 1, \dots, p$ .  $Y = X'\theta_0 + \varepsilon$  where  $\varepsilon \sim \mathcal{N}(0, 2)$ .

and  $m$ . As already observed in the fixed- $p$  regime, non-linear models (panels c and d) incur a considerable parallelization excess error.

## 6. Practical Considerations

Parallelization is not necessarily the preferred approach to deal with massive datasets. In principle, when  $N \gg p$  an easy, though potentially not sufficiently accurate solution, is to discard observations by randomly subsampling the data. Parallelization should thus be considered when the accuracy attainable by subsampling is not satisfactory. An important question is then over how many machines should the practitioner distribute the data? When tackling this question, we distinguish between two scaling regimes:  $N$  fixed or  $n$  fixed. Fixed  $n$  captures the single-machine storage constraint: the total available data is virtually infinite and using more machines allows processing of more data, and hence better accuracy, at an obvious financial cost. Fixed  $N$  captures either sampling or computational constraints: here, the total sample size  $N$  is fixed and processing it on a single machine might be too slow. Thus, splitting the data reduces run-time but also decreases the accuracy. In other words, *by parallelizing, we trade accuracy for speed*. Interestingly, when the number of samples  $N$  is fixed, by using our approximations and varying  $m$ , we are able to trace the accuracy-complexity tradeoff facing the practitioner. An informed choice of  $m$  is thus choosing either a desirable run-time, or a desired error level, on this curve.

We now formulate the target functions for choosing the number of machines in these




Figure 3: Bias and MSE in ridge regression as a function of number of machines  $m$ , with a fixed total number of samples  $N = 50,000$ , for different dimensions  $p = 50, 70, 80, 100$ , averaged over 1,000 replications. Panel (a) shows, in dotted lines, the mean of the empirical bias of an arbitrary coordinate  $j$  in the parallelized estimator,  $\mathbb{E} [\bar{\theta}_j - \theta_j^*]$ . The solid line is the theoretical expression for the second order bias from Section 3.3.2. Panel (b) shows, in dotted lines, the empirical MSE of the parallelized estimator,  $\mathbb{E} [\|\bar{\theta} - \theta^*\|_2^2]$  as a function of  $m$ . The solid lines are the theoretical approximation using a second order fixed- $p$  expansion of the error:  $\text{Tr}(M_2(\bar{\theta}))$  where  $M_2(\bar{\theta})$  is given in Eq.(14). In (b) the  $y$ -axis is  $\log_{10}$  scaled. Data was generated as follows:  $\lambda$  is fixed at 1.  $X \sim \mathcal{N}(0_p, I_{p \times p})$ ;  $\theta_0 = \tilde{\theta}/(\|\tilde{\theta}\|/10)$  where  $\tilde{\theta}_j = j$  for  $j = 1, \dots, p$ .  $Y = X'\theta_0 + \varepsilon$  where  $\varepsilon \sim \mathcal{N}(0, 2)$ .

two regimes. For fixed  $n$ , wishing to minimize costs, we analyze what is the *minimal* number of machines that attains a desired accuracy,  $\mathcal{E}(m)$ :

$$\min\{m \quad s.t. \quad \mathcal{E}(m) \leq \epsilon, n \text{ samples per machine}\}. \quad (19)$$

For fixed  $N$ , wishing to minimize runtime, and in the spirit of Shalev-Shwartz and Srebro [2008], we ask what is the *maximal* number of machines so that runtime is minimized while a desired level of accuracy is maintained. Choosing the number of machines in the fixed  $n$  scenario reduces to solving

$$\max\{m \quad s.t. \quad \mathcal{E}(m) \leq \epsilon, N/m \text{ samples per machine}\}. \quad (20)$$

Next, let us study these two optimization problems, Eqs.(19) and (20), when the accuracy measure is  $\mathcal{E}(m) := \mathbb{E} [\|\bar{\theta} - \theta^*\|^2]$ . This is challenging or even infeasible, since in general we do not have explicit expressions for this quantity. Moreover, in the fixed- $p$  regime, approximating the MSE by the asymptotic leading error term yields that this quantity is independent of  $m$ ! As we show below, meaningful and interesting solutions to this optimization problems arise when we approximate  $\mathcal{E}(m)$  by the second order expression in the fixed- $p$  regime. Specifically, using Eq.(11) we approximate  $\mathcal{E}(m) := \text{Tr}(M_2(\bar{\theta}))$ . In the high-dimensional regime, in contrast, the optimization problems (19) and (20) are well posed already when we approximate the MSE by the first order term. Relying on Eq.(G.2) gives  $\mathcal{E}(m) := r^2(\kappa)/m \mathbb{E}[\|\Sigma^{-1/2}\xi\|_2^2]$ .

We now present the optimization problems corresponding to the fixed- $p$  approximation in each scaling scenario:






Figure 4: MSE ratio  $\mathbb{E} [\|\bar{\theta} - \theta^*\|_2^2] / \mathbb{E} [\|\hat{\theta}_N - \theta^*\|_2^2]$  in the high-dimensional regime (with fixed  $\kappa = p/n = 0.2$ ) as a function of sample size  $n$ , averaged over 500 replications. Color encodes different number of machines  $m = 10, 20, 40, 100$ . The four panels depict different learning tasks as in Figure 1. Data was generated as follows:  $X \sim \mathcal{N}(0_p, I_{p \times p})$ ;  $\theta_0 = \tilde{\theta}/\|\tilde{\theta}\|$ , and  $\tilde{\theta}_j = j$  for  $j = 1, \dots, p$ ;  $\varepsilon \sim \mathcal{N}(0, 1)$ . In (a)-(b),  $Y = X'\theta_0 + \varepsilon$ . In (b)  $\lambda = 1$ . In (c)  $Y = \exp(X'\theta_0) + \varepsilon$ , whereas in (d)  $P(Y = 1|X) = \frac{1}{1 + \exp(-X'\theta_0)}$ .

**Fixed-n:**  $\min\{m \text{ s.t. } \frac{m-1}{m} \frac{1}{n^2} \text{Tr}(\gamma_0) + \frac{1}{mn} \text{Tr}(\gamma_1) + \frac{1}{mn^2} \text{Tr}(\gamma_2 + \gamma'_2 + \gamma_3 + \gamma_4 + \gamma'_4) \leq \epsilon\}$ , which stems from Eq.(19) and Eq.(11).

**Fixed-N:**  $\max\{m \text{ s.t. } \frac{(m-1)m}{N^2} \text{Tr}(\gamma_0) + \frac{1}{N} \text{Tr}(\gamma_1) + \frac{m}{N^2} \text{Tr}(\gamma_2 + \gamma'_2 + \gamma_3 + \gamma_4 + \gamma'_4) \leq \epsilon\}$ , which stems from Eq.(20) and Eq.(11) with  $n = N/m$ .

Let us illustrate these formulas in the OLS example from Section 3.3.1. The required quantities for OLS are collected in Appendix D. For example, solving the fixed- $N$  problem, the maximal number of machines that will keep the *per-coordinate* MSE under 0.2, i.e.,  $\varepsilon p = 0.2$ , with  $N = 10^6$ ,  $p = 100$ , and  $\sigma^2 = 10$  is  $m \leq 9,901$ . Alternatively, assuming an abundance of data and a memory limit such that  $n \leq 10^4$ , we solve the fixed- $n$  problem to find that  $m \geq 51$  will satisfy the derived error level.

**Remark 6.** In some cases, the practitioner may wish to control the parallelization error relative to the centralized solution, and not as an absolute value as analyzed above. Namely, the restriction is now  $\mathcal{E}(m) \leq (1 + \varepsilon)\mathcal{E}(1)$ . The scenarios (fixed  $n/N$ ) and approximations previously discussed apply here as well. For example, in our OLS example, solving the Fixed- $N$  problem with  $N = 10^6$ ,  $p = 100$ , and  $\sigma^2 = 10$ , yields that  $m \leq 991$  for  $\bar{\theta}$  to err no more than 10% more than  $\hat{\theta}_N$  ( $\varepsilon = 0.1$ ). On the other hand, For the Fixed- $n$  problem, with  $n = 10^4$ ,  $p = 100$ , and  $\sigma^2 = 10$ , we can parallelize up to  $m \leq 102$  machines, and still maintain the same 10% excess error allowance.

## 7. Discussion

In this work we studied the error of parallelized M-estimators when  $N$  observations are uniformly at random distributed over  $m$  machines. Each machine then learns a  $p$  dimensional model with its  $n = N/m$  observations and the  $m$  machine-wise results are

averaged to a global estimate  $\bar{\theta}$ . We derived several different approximations of the estimation error in  $\bar{\theta}$  with different quantitative and qualitative insights.

**Insights** When  $n \gg p$  not much accuracy is lost by splitting the data. This stands in contrast to other works that demonstrate how, under different assumptions, parallelization combined with averaging may incur a large error [Liu and Ihler, 2014], or even an unbounded one [Shamir et al., 2013]. Our analysis can thus be viewed as providing sufficient conditions for parallelization to be a suitable approach to reduce the overall run-time. A second insight is that if the model is highly non-linear, then the excess parallelization error may be considerably large.

In contrast to the classical fixed- $p$  regime, our high-dimensional analysis, currently confined to generative linear models, showed that splitting the data when there are only few observations per parameter always takes its accuracy toll. The degradation in accuracy due to splitting can still be quantified even though estimates converge to non-degenerate random limits.

**Future Research** At the basis of our work is an attempt to adhere to real-life software and hardware constraints of parallelized learning. The assumption of uniform and random distribution of samples to machines is realistic for some applications, and certainly facilitates the mathematical analysis. It may also be overly restrictive for other applications. A venue for future research is thus the relaxation of this assumption, allowing for some systematic difference between machines. We also aim at analyzing other aggregation schemes. Particularly ones that employ more than the mere machine-wise point estimate, and apply to non convex parameter spaces. An example of such is the Minimum Kullback-Leibler divergence aggregation, proposed by Liu and Ihler [2014]. This may extend the applicability of our results, for example, to image, sound, and graph data.

## Acknowledgments

We thank Derek Bean, Kyoo il Kim, Yaakov Ritov, Saharon Rosset, Ohad Shamir and Yuchen Zhang for fruitful discussions. This research was partly supported by a grant from the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI).

## A. Proof of Theorem 1

Under Assumption Set 1, classical statistical theory guarantees that upon optimizing the empirical risk (2), the resulting estimators,  $\hat{\theta}_n^{(j)}$ , converge in probability to  $\theta^*$  at rate  $n^{-1/2}$ . Moreover, the leading error term is linear in the influence functions

$\nabla f(Z_i, \theta^*)$  [Vaart, 1998, Theorem 5.23]:

$$\begin{aligned}\hat{\theta}_n^{(j)} &= \theta^* - V_{\theta^*}^{-1} \nabla \hat{R}_n^j(\theta^*) + o_P(n^{-1/2}) \\ &= \theta^* - V_{\theta^*}^{-1} \frac{1}{n} \sum_{i \in [j]} \nabla f(Z_i, \theta^*) + o_P(n^{-1/2}).\end{aligned}\tag{A.1}$$

where  $[j]$  denotes the  $n$  indexes of the observations assigned to machine  $j$ . Taking the average of the machine-wise estimators over a fixed number of machines  $\bar{\theta} := \frac{1}{m} \sum_j \hat{\theta}_n^{(j)}$ , and applying Eq.(A.1) yields

$$\begin{aligned}\bar{\theta} &= \theta^* - V_{\theta^*}^{-1} \nabla \hat{R}_N(\theta^*) + o_P(n^{-1/2}) \\ &= \theta^* - V_{\theta^*}^{-1} \frac{1}{N} \sum_{i=1}^N \nabla f(Z_i, \theta^*) + o_P(n^{-1/2}).\end{aligned}\tag{A.2}$$

Similarly, applying Eq.(A.1) to the centralized solution:

$$\hat{\theta}_N = \theta^* - V_{\theta^*}^{-1} \nabla \hat{R}_N(\theta^*) + o_P(N^{-1/2}).$$

Since  $m$  is fixed,  $o_P(N^{-1/2}) = o_P(n^{-1/2})$ . Eq.(4) now follows.  $\square$

## B. Proof of Theorem 2

Under Assumption Set 2, with  $s = 3$ , by Proposition 3.2 in [Rilstone et al., 1996]  $\hat{\theta}_n$  admits the expansion  $\theta^* + \xi_{-1/2}(\hat{\theta}_n) + \xi_{-1}(\hat{\theta}_n) + O_P(n^{-3/2})$ . We can thus decompose

$$B_2(\bar{\theta}) = \mathbb{E} \left[ \frac{1}{m} \sum_j \xi_{-1/2}(\hat{\theta}_n^{(j)}) + \frac{1}{m} \sum_j \xi_{-1}(\hat{\theta}_n^{(j)}) \right].$$

By definition,  $\xi_{-1/2}(\hat{\theta}_n) = -V_{\theta^*}^{-1} \nabla \hat{R}_n(\theta^*)$ . Hence,  $\mathbb{E} [\xi_{-1/2}(\hat{\theta}_n^{(j)})] = 0$ , for all  $j$  and the first term in the equation above vanishes. As for the second term, clearly  $\mathbb{E} [\xi_{-1}(\hat{\theta}_n^{(j)})]$  is independent of  $j$ , thus,  $B_2(\bar{\theta}) = \mathbb{E} [\xi_{-1}(\hat{\theta}_n)]$ .

Again, according to [Rilstone et al., 1996, Proposition 3.2],  $\mathbb{E} [\xi_{-1}(\hat{\theta}_n)] = \delta/n$ , where  $\delta$  depends on the various problem parameters, but is independent of  $n$ . Since  $B_2(\hat{\theta}_N) = \mathbb{E} [\xi_{-1}(\hat{\theta}_N)] = \delta/N$ , Eq.(8) readily follows.  $\square$

## C. Proof of Theorem 3

Under Assumption Set 2, with  $s = 4$ , we can expand  $\bar{\theta}$  as in Eq.(5). Plugging this asymptotic expansion of  $\bar{\theta}$  into the definition of  $M_2(\bar{\theta})$  from Eq.(9), and collecting terms

up to  $O(n^{-2})$  we have

$$M_2(\bar{\theta}) = \mathbb{E} \left[ \xi_{-1/2}(\bar{\theta}) \xi'_{-1/2}(\bar{\theta}) \right] + \mathbb{E} \left[ \xi_{-1}(\bar{\theta}) \xi'_{-1/2}(\bar{\theta}) \right] + \mathbb{E} \left[ \xi_{-1/2}(\bar{\theta}) \xi'_{-1}(\bar{\theta}) \right] + \\ \mathbb{E} \left[ \xi_{-1}(\bar{\theta}) \xi'_{-1}(\bar{\theta}) \right] + \mathbb{E} \left[ \xi_{-3/2}(\bar{\theta}) \xi'_{-1/2}(\bar{\theta}) \right] + \mathbb{E} \left[ \xi_{-1/2}(\bar{\theta}) \xi'_{-3/2}(\bar{\theta}) \right].$$

We now analyze each summand separately.  $\mathbb{E} \left[ \xi_{-1/2}(\bar{\theta}) \xi'_{-1/2}(\bar{\theta}) \right] = \frac{1}{m^2} \sum_{k,l} \mathbb{E} \left[ \xi_{-1/2}(\hat{\theta}_n^{(k)}) \xi'_{-1/2}(\hat{\theta}_n^{(l)}) \right]$ , for  $k \neq l$  then  $\mathbb{E} \left[ \xi_{-1/2}(\hat{\theta}_n^{(k)}) \xi'_{-1/2}(\hat{\theta}_n^{(l)}) \right]$  vanishes. For  $k = l$  we have  $m$  terms which equals  $n^{-1}\gamma_1$  by definition. The same analysis holds for  $\mathbb{E} \left[ \xi_{-1}(\bar{\theta}) \xi'_{-1/2}(\bar{\theta}) \right]$  and  $\mathbb{E} \left[ \xi_{-3/2}(\bar{\theta}) \xi'_{-1/2}(\bar{\theta}) \right]$  denoted  $n^{-2}\gamma_2$  and  $n^{-2}\gamma_4$  respectively. As for  $\mathbb{E} \left[ \xi_{-1}(\bar{\theta}) \xi'_{-1}(\bar{\theta}) \right] = \frac{1}{m^2} \sum_{k,l} \mathbb{E} \left[ \xi_{-1}(\hat{\theta}_n^{(k)}) \xi'_{-1}(\hat{\theta}_n^{(l)}) \right]$  : for  $k \neq l$  we have  $m(m-1)$  terms where  $\mathbb{E} \left[ \xi_{-1}(\hat{\theta}_n^{(k)}) \xi'_{-1}(\hat{\theta}_n^{(l)}) \right] = \mathbb{E} \left[ \xi_{-1}(\hat{\theta}_n^{(k)}) \right] \mathbb{E} \left[ \xi'_{-1}(\hat{\theta}_n^{(l)}) \right]$  which we defined as  $n^{-2}\gamma_0$ . For the remaining  $m$  terms where  $k = l$ , then  $\mathbb{E} \left[ \xi_{-1}(\hat{\theta}_n^{(k)}) \xi'_{-1}(\hat{\theta}_n^{(l)}) \right] = n^{-2}\gamma_3$  by definition. Collecting terms completes the proof.  $\square$

## D. Proof of Proposition 1

Denoting  $A^{-k} := (A^{-1})^k$ ;  $S_i := X_i X_i'$ ;  $\dot{f}_i := \nabla f(X_i, \theta^*)$ ;  $\ddot{f}_i := \nabla^2 f(X_i, \theta^*)$ ;  $V_i := (\ddot{f}_i - V_{\theta^*})$ ;  $d_i := V_{\theta^*}^{-1} \dot{f}_i$ ;  $W_\theta = \nabla^3 R(\theta)$ . We also denote by  $W_\theta(x, y)$  the linear operator in  $\mathbb{R}^p$  returned by evaluating  $W_\theta$  at  $(x, y) \in \mathbb{R}^p \times \mathbb{R}^p$ .

In our OLS setup we have:  $\nabla R(\theta) = -\Sigma(\theta_0 - \theta)$ ;  $V_{\theta^*}^{-1} = \Sigma^{-1}$ ;  $W_\theta \equiv 0$ ;  $\dot{f}_i = -X_i \varepsilon_i$ ;  $\ddot{f}_i = S_i$ ;  $V_i = S_i - \Sigma$ ;  $d_i = -\Sigma^{-1} X_i \varepsilon_i$ . From the proof of Proposition 3.2 in [Rilstone et al., 1996] we have

$$\mathbb{E} \left[ \xi_{-1}(\hat{\theta}_n) \right] = n^{-1} V_{\theta^*}^{-1} (\mathbb{E} [V_1 d_1] - \frac{1}{2} \mathbb{E} [W_{\theta^*}(d_1, d_1)]). \quad (\text{D.1})$$

As  $W_\theta \equiv 0$  and  $\mathbb{E} [\varepsilon_1] = 0$ , independent of  $X_1$ , then  $\mathbb{E} \left[ \xi_{-1}(\hat{\theta}_n) \right] = 0$ , so that from Eq.(6),  $\gamma_0 = 0$ .

Next, we turn to the matrix  $\gamma_1$ . From [Rilstone et al., 1996, page 374] we have

$$\gamma_1 = V_{\theta^*}^{-1} \mathbb{E} \left[ \dot{f}_1 \dot{f}_1' \right] V_{\theta^*}^{-1}.$$

Since  $\dot{f}_1 = -X_1 \varepsilon_1$ , and  $X_1$  independent of  $\varepsilon_1$ , then  $\mathbb{E} \left[ \dot{f}_1 \dot{f}_1' \right] = \sigma^2 \Sigma$ . Now recalling that  $V_{\theta^*} = \Sigma$  we obtain the form of  $\gamma_1$ .

Next, we analyze the matrix  $\gamma_2$ . From [Rilstone et al., 1996, Lemma 3.3] we have

$$\gamma_2 = V_{\theta^*}^{-1} \left( -\mathbb{E} [V_1 d_1 d_1'] + \frac{1}{2} \mathbb{E} [W_{\theta^*}(d_1, d_1) d_1'] \right).$$

The second summand vanishes and

$$\begin{aligned}\mathbb{E}[V_1 d_1 d'_1] &= \mathbb{E}[(S_1 - \Sigma)(\Sigma^{-1} X_1 \varepsilon_1)(\Sigma^{-1} X_1 \varepsilon_1)'] \\ &= \sigma^2 \mathbb{E}[S_1 \Sigma^{-1} S_1 \Sigma^{-1} - \Sigma \Sigma^{-1} S_1 \Sigma^{-1}] = \sigma^2 \mathbb{E}[S_1 \Sigma^{-1} S_1 \Sigma^{-1} - S_1 \Sigma^{-1}].\end{aligned}$$

As  $S_1 \sim \text{Wishart}_p(1, \Sigma)$ , we call upon Wishart theory. We collected the required properties in Appendix F. Applying Theorem F.1 to each summand we have that  $\mathbb{E}[S_1 \Sigma^{-1}] = I$ , and  $\mathbb{E}[S_1 \Sigma^{-1} S_1 \Sigma^{-1}] = (2\Sigma \Sigma^{-1} \Sigma + \text{Tr}(\Sigma \Sigma^{-1})\Sigma)\Sigma^{-1} = (2 + p)I$ .

As for the matrix  $\gamma_3$ , Rilstone et al. [1996, Lemma 3.3] yields

$$\gamma_3 = V_{\theta^*}^{-1} (\mathbb{E}[V_1 d_1 d'_2 V'_2] + \mathbb{E}[V_1 d_2 d'_1 V'_2] + \mathbb{E}[V_1 d_2 d'_2 V'_1]) V_{\theta^*}^{-1}.$$

Because  $\mathbb{E}[\varepsilon_1 \varepsilon_2] = 0$  the first two terms cancel. We now compute  $\mathbb{E}[V_1 d_2 d'_2 V'_1]$ . Recalling that  $S_1$  is independent of  $S_2$

$$\begin{aligned}\mathbb{E}[V_1 d_2 d'_2 V'_1] &= \mathbb{E}[(S_1 - \Sigma)(\Sigma^{-1} X_2 \varepsilon_2^2 X'_2 \Sigma^{-1})(S_1 - \Sigma)] \\ &= \sigma^2 \mathbb{E}[S_1 \Sigma^{-1} S_2 \Sigma^{-1} S_1 - S_1 \Sigma^{-1} S_2 \Sigma^{-1} \Sigma - \Sigma \Sigma^{-1} S_2 \Sigma^{-1} S_1 + \Sigma \Sigma^{-1} S_2 \Sigma^{-1} \Sigma] \\ &= \sigma^2 (\mathbb{E}[S_1 \Sigma^{-1} S_2 \Sigma^{-1} S_1] - \Sigma - \Sigma + \Sigma).\end{aligned}$$

Applying the results in Theorem F.1, then  $\mathbb{E}[S_1 \Sigma^{-1} S_2 \Sigma^{-1} S_1] = (p + 2)\Sigma$ .

Finally, we study the matrix  $\gamma_4$ . Since  $W_\theta$  and higher derivatives, all vanish, then [Rilstone et al., 1996, Lemma 3.3] yields

$$\gamma_4 = V_{\theta^*}^{-1} (\mathbb{E}[V_1 V_{\theta^*}^{-1} V_1 d_2 d'_2] + \mathbb{E}[V_1 V_{\theta^*}^{-1} V_2 d_1 d'_2] + \mathbb{E}[V_1 V_{\theta^*}^{-1} V_2 d_2 d'_1]).$$

Because  $\mathbb{E}[\varepsilon_1 \varepsilon_2] = 0$  the last two terms cancel. We now compute  $\mathbb{E}[-V_1 V_{\theta^*}^{-1} V_1 d_2 d'_2]$ :

$$\begin{aligned}\mathbb{E}[V_1 V_{\theta^*}^{-1} V_1 d_2 d'_2] &= \mathbb{E}[(S_1 - \Sigma)\Sigma^{-1}(S_1 - \Sigma)(\Sigma^{-1} X_2 \varepsilon_2^2 X'_2 \Sigma^{-1})] \\ &= \sigma^2 \mathbb{E}[S_1 \Sigma^{-1} S_1 \Sigma^{-1} S_2 \Sigma^{-1} - \Sigma \Sigma^{-1} S_1 \Sigma^{-1} S_2 \Sigma^{-1} - S_1 \Sigma^{-1} \Sigma \Sigma^{-1} S_2 \Sigma^{-1} + \Sigma \Sigma^{-1} \Sigma \Sigma^{-1} S_2 \Sigma^{-1}] \\ &= \sigma^2 (\mathbb{E}[S_1 \Sigma^{-1} S_1 \Sigma^{-1} S_2 \Sigma^{-1}] - I - I + I).\end{aligned}$$

Applying Theorem F.1 we get  $\mathbb{E}[S_1 \Sigma^{-1} S_1 \Sigma^{-1} S_2 \Sigma^{-1}] = (2 + p)I$ .

□

## E. Proof of Proposition 2

Using the notation from Appendix D we set up some quantities that will be reused throughout the computation. We start with some general results assuming  $\text{Var}[X] = \Sigma$ , but eventually restrict the results to  $\text{Var}[X] = I$  for simplicity.

When  $\text{Var}[X] = \Sigma$ , then  $\nabla R(\theta) = -\Sigma(\theta_0 - \theta) + \lambda\theta$ ;  $V_{\theta^*}^{-1} = (\Sigma + \lambda I)^{-1}$ ;  $W_\theta \equiv 0$ ;  $\dot{f}_i = -S_i(\theta_0 - \theta^*) - X_i \varepsilon_i + \lambda\theta^*$ ;  $\ddot{f}_i = S_i + \lambda I$ . Adding the simplifying assumption that  $\Sigma = I$  then  $\theta^* = \lambda_{0,1}\theta_0$ ;  $V_{\theta^*}^{-1} = \lambda_{0,1}I$ ;  $W_\theta \equiv 0$ ;  $\dot{f}_i = \lambda_{1,1}(I - S_i)\theta_0 - X_i \varepsilon_i$ ;  $\ddot{f}_i = S_i + \lambda I$ ;  $V_i = S_i - I_p$ ;  $d_i = \lambda_{1,2}(I - S_i)\theta_0 - \lambda_{0,1}X_i \varepsilon_i$ .

Starting with the  $\gamma_0$  matrix. Since  $W_\theta \equiv 0$ , then like in the OLS case,  $\delta = V_{\theta^*}^{-1}(\mathbb{E}[V_1 d_1])$ . In our ridge setup, with the Wishart properties in Theorem F.2:

$$\mathbb{E}[V_1 d_1] = \mathbb{E}[(S_1 - I)(\lambda_{1,2}(I - S_1)\theta_0 - \lambda_{0,1}X_1\varepsilon_1)] = \lambda_{1,2}\mathbb{E}[2S_1 - S_1^2 - I]\theta_0 = -\lambda_{1,2}(1 + p)\theta_0.$$

Plugging  $\delta$  in the definition of  $\gamma_0$  in Eq.(6) we get  $\gamma_0 = \lambda_{2,6}(1 + p)^2 B$ .

Moving to the  $\gamma_1$  matrix. From [Rilstone et al., 1996, page 374] we have  $\gamma_1 = V_{\theta^*}^{-1}\mathbb{E}[\dot{f}_1 \dot{f}_1'] V_{\theta^*}^{-1}$ . In our ridge setup, with the Wishart properties in Theorem F.2:

$$\begin{aligned}\mathbb{E}[\dot{f}_1 \dot{f}_1'] &= \mathbb{E}[(\lambda_{1,1}(I - S_1)\theta_0 - X_1\varepsilon_1)(\lambda_{1,1}(I - S_1)\theta_0 - X_1\varepsilon_1)'] \\ &= \mathbb{E}[\lambda_{2,2}(I - S_1)B(I - S_1) + \varepsilon_1^2 S_1] = \lambda_{2,2}(B + A) + \sigma^2 I,\end{aligned}$$

so that  $\gamma_1 = \lambda_{2,4}(B + A) + \lambda_{0,2}\sigma^2 I$ .

Moving to  $\gamma_2$ , from [Rilstone et al., 1996, Lemma 3.3] we have  $\gamma_2 = V_{\theta^*}^{-1}(-\mathbb{E}[V_1 d_1 d_1'])$ , and

$$\begin{aligned}\mathbb{E}[V_1 d_1 d_1'] &= \mathbb{E}[(S_1 - I)(\lambda_{1,2}(I - S_1)\theta_0 - \lambda_{0,1}X_1\varepsilon_1)(\lambda_{1,2}(I - S_1)\theta_0 - \lambda_{0,1}X_1\varepsilon_1)'] \\ &= \mathbb{E}[(S_1 - I)(\lambda_{2,4}(I - S_1)B(I - S_1) + \lambda_{0,2}\varepsilon_1^2 S_1)].\end{aligned}$$

Opening parenthesis and calling upon the Wishart properties in Theorem F.2:  $\mathbb{E}[S_1 B] = B$ ;  $\mathbb{E}[S_1^2 B] = (2 + p)B$ ;  $\mathbb{E}[S_1 B S_1] = 2B + A$ ;  $\mathbb{E}[S_1^2 B S_1] = (8 + 2p)B + (4 + p)A$ . Collecting terms:

$$\gamma_2 = -\lambda_{0,3}\sigma^2(1 + p)I - \lambda_{2,5}((4 + p)B + (3 + p)A).$$

Moving to  $\gamma_3$ . Calling upon [Rilstone et al., 1996, Lemma 3.3]:

$$\gamma_3 = V_{\theta^*}^{-1}(\mathbb{E}[V_1 d_1 d_2' V_2'] + \mathbb{E}[V_1 d_2 d_1' V_2'] + \mathbb{E}[V_1 d_2 d_2' V_1']) V_{\theta^*}^{-1}.$$

In this ridge setup, none of these terms cancel. The computations follow the same lines as for the previous matrices. The same holds for  $\gamma_4$  for which  $W_\theta \equiv 0$  and higher derivatives vanish, thus

$$\gamma_4 = V_{\theta^*}^{-1}(\mathbb{E}[V_1 V_{\theta^*}^{-1} V_1 d_2 d_2'] + \mathbb{E}[V_1 V_{\theta^*}^{-1} V_2 d_1 d_1'] + \mathbb{E}[V_1 V_{\theta^*}^{-1} V_2 d_2 d_1']) V_{\theta^*}^{-1}.$$

□

## F. Properties of Wishart Matrices

In this section we collect some properties of Wishart matrices used in this work. For classical results we provide appropriate references. For results we did not find in the literature, we present their proofs.

**Theorem F.1.** *Let  $S_1$  and  $S_2$  be independent  $\text{Wishart}_p(1, \Sigma)$  distributed, random matrices, and let  $B$  be a fixed symmetric  $p \times p$  matrix. It follows that: (i)  $\mathbb{E}[S_1] = \Sigma$ ; (ii)  $\mathbb{E}[S_1 B S_1] = 2\Sigma B \Sigma + \text{Tr}(\Sigma B)\Sigma$ , and (iii)  $\mathbb{E}[S_1 B S_2 B S_1] = 2\Sigma B \Sigma B \Sigma + \text{Tr}(B \Sigma B \Sigma)\Sigma$ .*

*Proof.* The first two statements are simply an application of [Fujikoshi et al., 2010, Theorem 2.2.5]. To prove (iii) we write  $S_{ij}^1 = (S_1)_{ij}$  and recall the independence between  $S_1$  and  $S_2$  to get

$$\mathbb{E}[(S_1 B S_2 B S_1)_{ij}] = \sum_{stkl} \mathbb{E}[S_{is}^1 B_{st} S_{tk}^2 B_{kl} S_{lj}^1] = \sum_{stkl} B_{st} B_{kl} \mathbb{E}[S_{is}^1 S_{lj}^1] \mathbb{E}[S_{tk}^2]. \quad (\text{F.1})$$

Calling upon [Gupta and Nagar, 1999, Theorem 3.3.3] we can represent  $S$  as  $S = X X'$  where  $X \sim \mathcal{N}_p(0, \Sigma)$ , so that  $\mathbb{E}[S_{is} S_{lj}] = \mathbb{E}[X_i X_s X_l X_j]$  and  $\mathbb{E}[S_{tk}] = \mathbb{E}[X_t X_k]$ . Now calling upon Isserlis' Theorem

$$\mathbb{E}[X_i X_j X_k X_l] = \mathbb{E}[X_i X_k] \mathbb{E}[X_j X_l] + \mathbb{E}[X_i X_l] \mathbb{E}[X_j X_k] + \mathbb{E}[X_i X_j] \mathbb{E}[X_k X_l],$$

which in our case equals

$$\mathbb{E}[X_i X_j X_k X_l] = \Sigma_{ik} \Sigma_{jl} + \Sigma_{il} \Sigma_{jk} + \Sigma_{ij} \Sigma_{kl}. \quad (\text{F.2})$$

Eq.(F.1) thus yields

$$\sum_{stkl} B_{st} B_{kl} \mathbb{E}[S_{is}^1 S_{lj}^1] \mathbb{E}[S_{tk}^2] = \sum_{stkl} B_{st} B_{kl} (\Sigma_{il} \Sigma_{sj} + \Sigma_{ij} \Sigma_{sl} + \Sigma_{is} \Sigma_{lj}) (\Sigma_{tk}).$$

The third statement in the theorem follows by rearranging into matrix notation.  $\square$

**Theorem F.2.** *Let  $S_1$  and  $S_2$  be independent  $\text{Wishart}_p(1, I_p)$  distributed, random matrices, and let  $B$  be a fixed symmetric  $p \times p$  matrix. It follows that:*

- (1)  $\mathbb{E}[S_1 S_2 B S_1 S_2] = (p+6)B + 2 \text{Tr}(B)I.$     (2)  $\mathbb{E}[S_1 S_2 B S_1] = 2B + \text{Tr}(B)I.$
- (3)  $\mathbb{E}[S_1 S_2 B S_2 S_1] = 4B + (4+p) \text{Tr}(B)I.$     (4)  $\mathbb{E}[S_1^2 B S_1] = (8+2p)B + (4+p) \text{Tr}(B)I.$
- (5)  $\mathbb{E}[S_1 S_2 S_1 B S_2] = (6+p)B + \text{Tr}(B)I.$     (6)  $\mathbb{E}[S_1 S_2^2 B S_1] = (4+2p)B + (2+p) \text{Tr}(B)I.$

*Proof.* We present only the proof of the first statement. The others are proved using the same arguments. We also note that these arguments can be used for the more general  $\text{Wishart}_p(n, \Sigma)$  matrices. We now write  $S_{ij}^1 = (S_1)_{ij}$  and recall the independence between  $S_1$  and  $S_2$  to get

$$\mathbb{E}[(S_1 S_2 B S_1 S_2)_{ij}] = \sum_{stkl} \mathbb{E}[S_{is}^1 S_{st}^2 B_{tk} S_{kl}^1 S_{lj}^2] = \sum_{stkl} B_{tk} \mathbb{E}[S_{is}^1 S_{kl}^1] \mathbb{E}[S_{st}^2 S_{lj}^2]. \quad (\text{F.3})$$

The mean  $\mathbb{E}[S_{is} S_{kl}] = \mathbb{E}[X_i X_s X_k X_l]$  is given in Eq.(F.2) which in the case where  $\Sigma = I$  simplifies into

$$\mathbb{E}[X_i X_j X_k X_l] = \delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk} + \delta_{ij} \delta_{kl}.$$

Eq.(F.3) thus yields

$$\sum_{stkl} B_{tk} (\delta_{ik} \delta_{sl} + \delta_{il} \delta_{sk} + \delta_{is} \delta_{kl}) (\delta_{sl} \delta_{tj} + \delta_{sj} \delta_{tl} + \delta_{st} \delta_{lj})$$

The first statement in the theorem is recovered by rearranging into matrix notation.  $\square$

## G. High-Dimensional Regime

### G.1. Proof of Lemma 1

Eq.(15) is simply taken from Bean et al. [2013, Result 1], stated here for completeness. From this equation it follows that the mean squared error of the centralized solution is

$$MSE[\hat{\theta}_N, \theta^*] = \mathbb{E}[\|\hat{\theta}_N - \theta^*\|_2^2] = r^2(\kappa/m) \mathbb{E}[\|\Sigma^{-1/2}\xi\|_2^2] (1 + o(1)). \quad (\text{G.1})$$

In contrast, upon averaging the estimators of  $m$  machines, we obtain

$$MSE[\bar{\theta}, \theta^*] = r^2(\kappa)/m \mathbb{E}[\|\Sigma^{-1/2}\xi\|_2^2] (1 + o(1)). \quad (\text{G.2})$$

By comparing Eq.(G.1) and Eq.(G.2), the accuracy loss of parallelization compared to a centralized estimation is thus given by

$$\frac{MSE[\bar{\theta}, \theta^*]}{MSE[\hat{\theta}_N, \theta^*]} = \frac{r^2(\kappa)/m}{r^2(\kappa/m)} (1 + o(1)). \quad (\text{G.3})$$

Using perturbation analysis, and denoting by  $f_{[i]} := \frac{\partial^i}{\partial t^i} f(t)$  the  $i$ -th derivative of the loss function, we now characterize the behaviour of Eq.(G.3) as  $\kappa$  and  $m$  vary. To this end, and in line with El Karoui et al. [2013] and Donoho and Montanari [2013], we first introduce some more definitions:

$$\xi_\epsilon := \epsilon + r(\kappa) \eta; \quad \eta \sim \mathcal{N}(0, 1), \text{ independent of } \epsilon. \quad (\text{G.4})$$

Furthermore  $f(Y - X'\theta)$  is treated as a univariate function of the residual alone, and

$$\text{prox}_c(z) := \arg \min_x \left\{ f(x) + \frac{1}{2c} \|x - z\|^2 \right\}.$$

Readers familiar with optimization theory will recognize the equation as the standard proximal operator.

As shown by both El Karoui et al. [2013, Corollary 1] and Donoho and Montanari [2013, Theorem 4.1], under the assumptions of Lemma 1 the quantity  $r(\kappa)$  together with a second quantity  $c = c(f, \kappa)$  are the solution of the following set of two coupled non-linear equations

$$\mathbb{E} \left[ \left. \frac{d}{dz} \text{prox}_c(z) \right|_{z=\xi_\epsilon} \right] = 1 - \kappa, \quad (\text{G.5})$$

$$\mathbb{E} \left[ (\xi_\epsilon - \text{prox}_c(\xi_\epsilon))^2 \right] = \kappa r^2(\kappa), \quad (\text{G.6})$$

where the averaging operator in Eqs.(G.5)-(G.6) is with respect to the random variable  $\xi_\epsilon$ .

Since Eqs.(G.5) and (G.6) are solvable analytically only in very specific cases, we study the limiting behaviour of the solution of this set of equations as  $\kappa \rightarrow 0$ . To this end we note that for  $\kappa = 0$ ,  $r(0) = 0$  and similarly  $c(f, 0) = 0$  [El Karoui et al., 2012, Section 4.3]. We first study the behaviour of the proximal operator  $\text{prox}_c(z)$  for small values of  $c$ .



**Lemma 2.** Assume the loss function  $f$  is three times differentiable, with a bounded third derivative,  $|f_3| \leq K$ . Then, as  $c \rightarrow 0$ ,

$$\text{prox}_c(z) = z + \sqrt{c}w^* = z - cf_1(z) + c^2f_1(z)f_2(z) + O(c^3). \quad (\text{G.7})$$

*Proof.* Upon the change of variables  $x = z + \sqrt{c}w$ , Eq.(G.7) becomes

$$\text{prox}_c(z) = z + \sqrt{c} \arg \min_w \left\{ f(z + \sqrt{c}w) + \frac{1}{2}\|w\|^2 \right\}.$$

Denote by  $w^* = w^*(z, c)$  the minimizer in the equation above. Under the assumption that  $f$  is differentiable, it is the solution of

$$\sqrt{c}f_1(z + \sqrt{c}w) + w = 0. \quad (\text{G.8})$$

Next, we make an exact 2-term Taylor expansion of  $f_1$  around the value  $z$ , with a remainder term involving the third derivative of  $f$ . Then at  $w^*$  we have

$$0 = w^* + \sqrt{c} [f_1(z) + \sqrt{c}w^*f_2(z) + \frac{1}{2}c(w^*)^2f_3(\tilde{w})], \quad (\text{G.9})$$

where  $\tilde{w}$  is some point in the interval  $[0, w^*]$ . Eq. (G.9) is a singular quadratic equation in  $w^*$ , although implicit since  $\tilde{w}$  depends on  $w^*$  as well. It has two solutions, one that explodes to  $\infty$  as  $c \rightarrow 0$  and the other, relevant to us, that tends to zero as  $c \rightarrow 0$ . Under the assumption that  $|f_3| \leq K$  we have that

$$\begin{aligned} w^*(z) &= -\frac{\sqrt{c}f_1(z)}{1 + cf_2(z)} + O(c^2\sqrt{c}) \\ &= -\sqrt{c}f_1(z) + c\sqrt{c}f_2(z)f_1(z) + O(c^2\sqrt{c}). \end{aligned} \quad (\text{G.10})$$

Inserting Eq. (G.10) into the definition of the prox function concludes the proof.  $\square$

## G.2. Approximating The Residual Noise Equations

Next, we study the form of Eqs.(G.5) and (G.6) as  $\kappa \rightarrow 0$ . As  $r(\kappa) \rightarrow 0$  the distribution of the random variable  $\xi_\epsilon := \epsilon + r(\kappa)\eta$  converges to that of  $\epsilon$ . The following lemma quantifies how averaging with respect to  $\xi_\epsilon$  is related to averaging with respect to  $\epsilon$ .

**Lemma 3.** Let  $g : \mathbb{R} \rightarrow \mathbb{R}$  be a smooth differentiable function, with bounded fourth derivative,  $|g_4(x)| \leq K$ . Let  $\xi_\epsilon := \epsilon + r\eta$  be defined as in Eq.(G.4). Then, as  $r \rightarrow 0$ ,

$$\mathbb{E}_{\xi_\epsilon}[g(\xi_\epsilon)] = \mathbb{E}_\epsilon[g(\epsilon)] + \frac{1}{2}r^2\mathbb{E}_\epsilon[g_2(\epsilon)] + O(r^4). \quad (\text{G.11})$$

*Proof.* By definition of the random variable  $\xi_\epsilon$  and of the expectation operator,

$$\mathbb{E}_{\xi_\epsilon}[g(\xi_\epsilon)] = \iint g(\epsilon + r\eta)dF_\epsilon dF_\eta, \quad (\text{G.12})$$

where  $dF_\varepsilon$  and  $dF_\eta$  are the CDFs of the random variables  $\varepsilon$ , and  $\eta$ , respectively. Making a Taylor expansion of  $g$  up to fourth order gives

$$g(\epsilon + r\eta) = g(\varepsilon) + r\eta g_1(\epsilon) + \frac{1}{2}r^2\eta^2 g_2(\epsilon) + \frac{1}{6}r^3\eta^3 g_3(\epsilon) + \frac{1}{24}r^4\eta^4 g_4(\tilde{\xi}_\epsilon), \quad (\text{G.13})$$

where  $\tilde{\xi}_\epsilon$  is an intermediate point in the interval  $[\epsilon, \epsilon + r\eta]$ . As  $\eta$  is symmetrically distributed, upon inserting this expansion into Eq.(G.12), odd terms cancel. Also, given that the fourth derivative is bounded, the remainder term is indeed  $O(r^4)$ , and Eq. (G.11) follows.  $\square$

We now arrive at our main result regarding the solution of the system of equations (G.5) and (G.6).

**Theorem G.1.** *As  $\kappa \rightarrow 0$ , the solution of the system of Equations (G.5)-(G.6) admits the following asymptotic form*

$$\begin{aligned} c(\kappa) &= \kappa c_1 + \kappa^2 c_2 + O(\kappa^3), \\ r^2(\kappa) &= \kappa r_1 + \kappa^2 r_2 + O(\kappa^3). \end{aligned} \quad (\text{G.14})$$

The coefficients  $c_1, c_2, r_1, r_2$  are given by

$$c_1 = \frac{1}{A_2}, \quad c_2 = \frac{T_1}{A_2^3} - \frac{B_1 A_4}{A_2^4}, \quad r_1 = \frac{B_1}{A_2^2}, \quad r_2 = \frac{3B_1 T_1}{A_2^4} - \frac{2B_1^2 A_4}{A_2^5} + \frac{2B_2}{A_2^3}, \quad (\text{G.15})$$

where

$$A_2 = \mathbb{E}_\varepsilon[f_{[2]}], \quad A_4 = \mathbb{E}_\varepsilon[1/2 f_{[4]}], \quad T_1 = \mathbb{E}_\varepsilon[f_{[2]}^2 + f_{[1]}f_{[3]}], \quad B_1 = \mathbb{E}_\varepsilon[f_{[1]}^2], \quad B_2 = \mathbb{E}_\varepsilon[f_{[1]}^2 f_{[2]}], \quad (\text{G.16})$$

and expectations are taken only with respect to the distribution of  $\epsilon$ .

*Proof.* As common in applied mathematics, when analyzing the solutions of equations with a small parameter, we assume the solution is of the form (G.14), show that this expansion is consistent with the form of the equations, and extract explicit expressions for the coefficients by comparing terms of equal order in  $\kappa$  [Bender and Orszag, 1999].

To this end, let us first simplify the original set of equations, in the limit as  $\kappa \rightarrow 0$ . Inserting (G.7) into (G.5) yields

$$\kappa = \mathbb{E}_{\xi_\epsilon} [cf_2(\xi_\epsilon) - c^2(f_2^2(\xi_\epsilon) + f_1(\xi_\epsilon)f_3(\xi_\epsilon))] + O(c^3).$$

Since to leading order  $c = O(\kappa)$ , the error in the above equation is  $O(\kappa^3)$ . Next, using Eq. (G.11), the terms on the right hand side may be further approximated as follows,

$$\begin{aligned} \mathbb{E}_{\xi_\epsilon} [cf_2(\xi_\epsilon)] &= c\mathbb{E}_\epsilon[f_2] + \frac{1}{2}cr^2\mathbb{E}[f_4] + O(cr^4), \\ \mathbb{E}_{\xi_\epsilon} [c^2(f_2^2(\xi_\epsilon) + f_1(\xi_\epsilon)f_3(\xi_\epsilon))] &= c^2\mathbb{E}_\epsilon[f_2^2 + f_1f_3] + O(c^2r^2). \end{aligned}$$

Given the assumed expansion for  $c(\kappa)$  and  $r^2(\kappa)$ , both error terms are  $O(\kappa^3)$  and the first equation reads

$$\kappa = c\mathbb{E}_\epsilon[f_2] + \frac{1}{2}cr^2\mathbb{E}_\epsilon[f_4] - c^2\mathbb{E}_\epsilon[f_2^2 + f_1f_3] + O(\kappa^3). \quad (\text{G.17})$$

Next, we simplify the second equation (G.6). To this end, note that as  $\kappa \rightarrow 0$  and also  $c \rightarrow 0$ ,

$$(\xi_\epsilon - \text{prox}_c(\xi_\epsilon))^2 = (cf_1 - c^2f_1f_2 + O(c^3))^2 = c^2f_1^2 - 2c^3f_1^2f_2 + O(\kappa^4). \quad (\text{G.18})$$

Thus, Eq.(G.6) reads

$$\kappa r^2 = \mathbb{E}_{\xi_\epsilon}[(\xi_\epsilon - \text{prox}_c(\xi_\epsilon))^2] = \mathbb{E}_{\xi_\epsilon}[c^2f_1^2(\xi_\epsilon) - 2c^3f_1^2(\xi_\epsilon)f_2(\xi_\epsilon)] + O(\kappa^4). \quad (\text{G.19})$$

Using Eq.(G.11), we further simplify the right hand side, to read

$$\kappa r^2 = c^2\mathbb{E}_\epsilon[f_{[1]}^2] + c^2r^2\mathbb{E}_\epsilon[f_2^2 + f_1f_3] - 2c^3\mathbb{E}_\epsilon[f_1^2f_2] + O(\kappa^4). \quad (\text{G.20})$$

For ease of notation, we thus write the system of Equations (G.17) and (G.20) as follows,

$$\begin{aligned} \kappa &= cA_2 + cr^2A_4 - c^2T_1 + O(\kappa^3), \\ \kappa r^2 &= c^2B_1 + c^2r^2T_1 - 2c^3B_2 + O(\kappa^4), \end{aligned} \quad (\text{G.21})$$

where the coefficients  $A_2, A_4, B_1, B_2$  and  $T_1$  are given in Eq. (G.16).

To determine the coefficients  $c_1, c_2, r_1, r_2$  we now insert the expansion (G.14) into Eq.(G.21), and compare powers of  $\kappa$ . This gives

$$\kappa = A_2c_1\kappa + A_2c_2\kappa^2 + A_4c_1r_1\kappa^2 - T_1c_1^2\kappa^2 + O(\kappa^3), \quad (\text{G.22})$$

$$r_1\kappa^2 + r_2\kappa^3 = B_1(c_1^2\kappa^2 + 2c_1c_2\kappa^3) + T_1c_1^2r_1\kappa^3 - B_2c_1^3\kappa^3 + O(\kappa^4). \quad (\text{G.23})$$

Thus, comparing the  $O(\kappa)$  terms in Eq. (G.22) gives  $c_1 = \frac{1}{A_2}$ . Similarly, comparing the  $O(\kappa^2)$  terms in Eq. (G.23) yields  $r_1 = B_1c_1^2 = \frac{B_1}{A_2^2}$ . Next, comparing the  $O(\kappa^2)$  terms in Eq. (G.22) gives

$$c_2 = \frac{T_1c_1^2 - A_4c_1r_1}{A_2} = \frac{T_1}{A_2^3} - \frac{B_1A_4}{A_2^4}.$$

Finally, comparing the  $O(\kappa^3)$  terms in Eq. (G.23) yields the expression for  $r_2$  in Eq. (G.15).  $\square$

## H. Scope of Theorem 1

The assumptions of Theorem 1 are standard in the M-Estimation literature and apply to many common learning tasks. Here is a non-comprehensive list of examples, relevant for parallelization.

**Example 1** (Non-Linear Least Squares). Here the loss function is  $f(Z, \theta) := (Y - g_\theta(X))^2$ , where  $g_\theta$  is a smooth function indexed by a parameter  $\theta$ . Applicability of Theorem 1 to this example, under suitable regularity conditions on the family  $g_\theta$  is discussed in [Vaart, 1998, Example 5.27].

**Example 2** (Non-Linear Quantile Regression). The loss function corresponding to a quantile level  $\tau$ , is the tilted hinge loss  $f(Z, \theta) := |(Y - g_\theta(X))(\tau - I_{\{Y - g_\theta(X) < 0\}})|$ . This case is similar to Example 1 with the complication that the loss function has a single non differentiability point. Assumption (A3) is still satisfied for continuous  $X$ , since the probability of  $f(Z, \theta)$  to be non-differentiable at  $\theta^*$  is 0.

**Example 3** (Robust Regression). Consider non-linear robust regression with the Huber loss function,

$$f(Z, \theta) := \begin{cases} \frac{1}{2}(Y - g_\theta(X))^2 & \text{if } |Y - g_\theta(X)| \leq \delta \\ \delta (|Y - g_\theta(X)| - \frac{1}{2}\delta) & \text{otherwise} \end{cases}$$

The linear case with  $g_\theta(x) = X'\theta$  is discussed in [Vaart, 1998, Examples 5.28 & 5.29]. Theorem 1 also applies to the non-linear case, under appropriate assumptions on  $g_\theta$ .

**Example 4** (Binary Regression). Here the binary response  $Y$  is assumed to arise from the generative model  $P(Y|X) = \Psi(X'\theta)$ , with a known and sufficiently regular  $\Psi$ . With the loss function taken as the negative log likelihood, this case is discussed in [Vaart, 1998, Example 5.40]. This setting covers popular generalized linear models such as Logistic, Probit, and Complementary Log-Log regression.

**Example 5** (SVM). For  $Y \in \{-1, 1\}$  and  $f(Z, \theta) := \max\{0, 1 - Y X'\theta\} + \frac{\lambda}{2}\|\theta\|_2^2$  we get the SVM problem with  $l_2$  regularization. This problem satisfies the risk regularity assumption (A2) [e.g. Shamir, 2012]. The differentiability of the empirical risk (A3) is settled, for continuous  $X$ , as in Example 2.

## References

- K. Achutegui, D. Crisan, J. Miguez, and G. Rios. A simple scheme for the parallelization of particle filters and its application to the tracking of complex stochastic systems. *ArXiv e-prints*, 2014.
- D. Bean, P. J. Bickel, N. E. Karoui, and B. Yu. Optimal M-estimation in high-dimensional regression. *Proceedings of the National Academy of Sciences*, pages 14563–14568, 2013.
- R. Bekkerman, M. Bilenko, and J. Langford. *Scaling up Machine Learning: Parallel and Distributed Approaches*. Cambridge University Press, Cambridge; New York, 2011.
- C. M. Bender and S. A. Orszag. *Advanced Mathematical Methods for Scientists and Engineers: Asymptotic Methods and Perturbation Theory*. Springer, New York, 1999.

- J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. *Commun. ACM*, 51(1):107–113, 2008.
- L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, New York, 1997.
- D. Donoho and A. Montanari. High Dimensional Robust M-Estimation: Asymptotic Variance via Approximate Message Passing. *ArXiv e-prints*, 2013.
- H. Doss and J. Sethuraman. The Price of Bias Reduction when there is no Unbiased Estimate. *The Annals of Statistics*, 17(1):440–442, 1989.
- N. El Karoui. Asymptotic behavior of unregularized and ridge-regularized high-dimensional robust regression estimators: rigorous results. *ArXiv e-prints*, 1311:2445, 2013.
- N. El Karoui, P. Bickel, D. Bean, C. Lim, and B. Yu. On robust regression with high-dimensional predictors. Technical report, University of California, Berkeley, 2012.
- N. El Karoui, D. Bean, P. J. Bickel, C. Lim, and B. Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, page 1455714562, 2013.
- J. Feng, H. Xu, and S. Mannor. Distributed Robust Learning. *ArXiv e-prints*, 2014.
- Y. Fujikoshi, V. V. Ulyanov, and R. Shimizu. *Multivariate Statistics: High-Dimensional and Large-Sample Approximations*. Wiley, 2010.
- A. K. Gupta and D. K. Nagar. *Matrix Variate Distributions*. Chapman and Hall, Boca Raton, FL, 1999.
- D. Hsu and S. Sabato. Loss minimization and parameter estimation with heavy tails. *ArXiv e-prints*, 2013.
- P. J. Huber. Robust Regression: Asymptotics, Conjectures and Monte Carlo. *The Annals of Statistics*, 1(5):799–821, 1973.
- K. I. Kim. Higher Order Bias Correcting Moment Equation for M-Estimation and Its Higher Order Efficiency. *Research Collection School Of Economics*, 2006.
- Q. Liu and A. T. Ihler. Distributed estimation, information loss and exponential families. In *Advances in Neural Information Processing Systems 27*, pages 1098–1106. Curran Associates, Inc., 2014.
- R. McDonald, M. Mohri, N. Silberman, D. Walker, and G. S. Mann. Efficient Large-Scale Distributed Training of Conditional Maximum Entropy Models. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 1231–1239. Curran Associates, Inc., 2009.

- Z. Meng, A. Wiesel, and A. Hero. Distributed principal component analysis on networks via directed graphical models. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2877–2880, 2012.
- H. Rieder. *Robust Asymptotic Statistics: Volume I*. Springer, New York, NY, 2012.
- P. Rilstone, V. K. Srivastava, and A. Ullah. The second-order bias and mean squared error of nonlinear estimators. *Journal of Econometrics*, 75(2):369–395, 1996.
- S. Rosset and J. Zhu. Piecewise linear regularized solution paths. *The Annals of Statistics*, 35(3):1012–1030, 2007.
- S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, New York, NY, 2014.
- S. Shalev-Shwartz and N. Srebro. SVM Optimization: Inverse Dependence on Training Set Size. *Proceedings of the 25th International Conference on Machine Learning*, pages 928–935, 2008.
- O. Shamir. Is averaging needed for strongly convex stochastic gradient descent. *Open problem presented at COLT*, 2012.
- O. Shamir, N. Srebro, and T. Zhang. Communication Efficient Distributed Optimization using an Approximate Newton-type Method. *ArXiv e-prints*, 2013.
- K. Shvachko, H. Kuang, S. Radia, and R. Chansler. The Hadoop Distributed File System. *Proceedings of IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–10, 2010.
- A. W. v. d. Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: cluster computing with working sets. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, pages 10–10, 2010.
- Y. Zhang, J. Duchi, and M. Wainwright. Divide and conquer kernel ridge regression. In *Conference on Learning Theory*, pages 592–617, 2013a.
- Y. Zhang, J. C. Duchi, and M. J. Wainwright. Communication-efficient Algorithms for Statistical Optimization. *J. Mach. Learn. Res.*, 14(1):3321–3363, 2013b.
- M. Zinkevich, M. Weimer, L. Li, and A. J. Smola. Parallelized stochastic gradient descent. *Advances in Neural Information Processing Systems*, pages 2595–2603, 2010.