

Distributed Stochastic Multi-Task Learning with Graph Regularization

Weiran Wang*, Jialei Wang†, Mladen Kolar‡, and Nathan Srebro*

*Toyota Technological Institute at Chicago, IL, USA

†Department of Computer Science, University of Chicago, IL, USA

‡Booth School of Business, University of Chicago, IL, USA

Abstract

We propose methods for distributed graph-based multi-task learning that are based on weighted averaging of messages from other machines. Uniform averaging or diminishing stepsize in these methods would yield consensus (single task) learning. We show how simply skewing the averaging weights or controlling the stepsize allows learning different, but related, tasks on the different machines.

.ML] 11 Feb 2018

1 Introduction

We consider a distributed learning problem in a multi-task setting: each machine i has access to samples from a different data distribution D_i , with potentially a different optimal predictor, and thus a different learning task, but where we still assume some similarity between different tasks. The goal of each machine is to find a good predictor for its own task, based on its own local data, as well as communicating with the other machines so as to leverage the similarity to other related tasks.

Distributed multi-task learning lies between a homogeneous distributed learning setting (e.g. Shamir and Srebro, 2014), where all machines have data from the same source distribution, and inhomogeneous consensus problems (e.g. Ram et al., 2010; Boyd et al., 2011; Balcan et al., 2012), where each machine sees data from a different source, but the goal is to reach a single consensus predictor. In many distributed learning problems, different machines do indeed see different distributions. For example, machines might serve different geographical regions. In a more extreme “federated learning” (Konecny et al., 2015) scenario, each machine is a single user device, and its data distribution might reflect e.g. the user’s speech, language biases, usage patterns, etc. Such heterogeneity requires departing from a homogeneous model. But if the data distribution on each machine is different, we might as well learn a personalized predictor for each machine, while still leveraging commonalities as in multi-task learning, instead of insisting on consensus. Unlike when seeking consensus, we could learn a predictor entirely locally, ignoring data on other machines. But the premise of multi-task learning is that by communicating with other machines we can improve our predictions, reduce the sample complexity, and hopefully also reduce the computational cost on each machine by distributing the computation.

Central to multi-task learning is the notion of relatedness between tasks. In a high-dimensional setting, with large number of variables, we might expect a small common set of predictive variables, where the form of the dependence on variables in this common set varies between tasks (Turlach et

al., 2005; Obozinski et al., 2011; Lounici et al., 2011; Wang et al., 2015). Another approach is to assume that the predictors lie in a shared lower dimensional subspace (Ando and Zhang, 2005; Yuan et al., 2007; Wang et al., 2016) or all have low-norm under some shared linear representation (Amit et al., 2007; Argyriou et al., 2008). Both the shared sparsity and shared subspaces models have recently been considered in a distributed learning setting (Wang et al., 2015, 2016), and nuclear-norm regularized multi-task learning has been studied from a distributed optimization perspective (Baytas et al., 2016).

In this paper, we consider graph-based multi-task learning, where relatedness between tasks is specified through a weighted graph over the tasks. Neighboring tasks in the graph are expected to be similar, with a penalty for dis-similarity specified by the weight between them (see precise formulation in Section 2) (Maurer, 2006; Evgeniou et al., 2005). This also generalizes a simpler “fully connected” multi-task model where all predictors are close to each other (Evgeniou and Pontil, 2004). A predictor-homogeneous assumption can also be viewed as an extreme case where all weights go to infinity, forcing all predictors to be identical. In distributed multi-task learning, graph-based relatedness is especially appealing if the relatedness graph also matches the graph of network links between machines, as might be the case, e.g. in a geographical setting or with physical sensors. We therefore emphasize and prefer methods with communication only between neighboring tasks on the graph.

In designing methods for graph-based multi-task learning, we are interested in methods that (1) are natural and simple—all our algorithms have a similar and natural structure, involving weighted averaging of messages from neighboring machines and a local gradient or prox calculation; (2) have low communication costs, are sample efficient, and preferably also have low computational cost; and (3) are backed by rigorous guarantees on the amount of communication, samples and computation required.

Graph-based multi-task learning has been recently studied by Vanhaesebrouck et al. (2017) and Liu et al. (2017), both considering the problem as distributed optimization of the multitask regularized *empirical* objective, similar to our approach in Section 3.2). Vanhaesebrouck et al. suggested an asynchronous gossip-type algorithms and an ADMM procedure, while Liu et al. proposed using SDCA, and also considered learning the relatedness graph itself. Neither provides any statistical analysis, nor analysis of the iteration complexity and communication cost based on the methods. We conduct detailed comparison of convergence properties with these methods in Appendix H, providing upper bounds of their iteration complexities when possible; our methods have faster convergence than the guarantees we could obtain for them. Also, neither directly considers the underlying learning problem (minimizing the actual expected errors), and so neither studies stochastic methods (in the flavor of our Section 4).

Here, we show how methods that arise naturally by skewing averaging weights or controlling stepsize of consensus learning methods do yield good guarantees. We also propose stochastic methods which allow reducing the computational cost, and we compare the empirical performance of both our batch and stochastic methods to those of Vanhaesebrouck et al. (2017) and Ma et al. (2015).

Notations In this paper, boldface lower-case letters denote column vectors, boldface capital letters denote matrices, $\text{vec}(\mathbf{U})$ is the vectorial form of a matrix \mathbf{U} which concatenates columns of \mathbf{U} , and

$\mathbf{U} \otimes \mathbf{V}$ is the Kronecker product between two matrices \mathbf{U} and \mathbf{V} . Furthermore, $\mathbf{u}^T \mathbf{v} = \langle \mathbf{u}, \mathbf{v} \rangle$ denotes the inner product of two vectors \mathbf{u} and \mathbf{v} , while $\langle \mathbf{U}, \mathbf{V} \rangle = \text{tr}(\mathbf{U}^T \mathbf{V})$ denotes inner product

of two matrices \mathbf{U} and \mathbf{V} of the same dimensions. We use $\|\mathbf{u}\| = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}$ to denote the length of

a vector \mathbf{u} , $\|\mathbf{U}\|_F = \|\text{vec}(\mathbf{U})\|$ the Frobenius norm of a matrix \mathbf{U} , and $\|\mathbf{U}\|_M =$

$$\text{tr}(\mathbf{U} \mathbf{M} \mathbf{U}^T) =$$

$\langle \mathbf{U}, \mathbf{U} \rangle_M$ the norm of \mathbf{U} with respect to some positive definite matrix \mathbf{M} . A function $f(\mathbf{x})$ is Lipschitz if $|f(\mathbf{x}) - f(\mathbf{y})| \leq L \|\mathbf{x} - \mathbf{y}\|$, $\forall \mathbf{x}, \mathbf{y}$. A convex function $f_2(\mathbf{x})$ is β -smooth and μ -strongly convex if $\frac{1}{2\beta} \|\mathbf{x} - \mathbf{y}\|^2 \leq f_2(\mathbf{x}) - f_2(\mathbf{y}) \leq \frac{1}{2\mu} \|\mathbf{x} - \mathbf{y}\|^2$, $\forall \mathbf{x}, \mathbf{y}$. This definition extends to functions of matrices, by replacing the vector norm with the Frobenius norm in the above inequality.

2 Graph-based multi-task learning

Consider a distributed setting with m machines, where each machine i has access to a data distribution D_i and would like to learn a predictor $\mathbf{w}_i \in \mathbb{R}^d$ for each machines with small expected loss $F_i(\mathbf{w}_i) = \mathbb{E}_{\mathbf{z}_i \sim D_i} [\ell(\mathbf{w}_i, \mathbf{z}_i)]$. A known weighted graph, with known non-negative weights $\{a_{ik}\}$, specifies the relatedness between tasks. Specially, we would like to consider predictor matrices

$\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_m] \in \mathbb{R}^{d \times m}$ from the set

$$\left\{ \mathbf{W} \mid \|\mathbf{w}_i\| \leq \sigma, \sum_k a_{ik} \|\mathbf{w}_i - \mathbf{w}_k\| \leq \tau, \forall i, k \right\}$$

i.e., we would like the norm of each individual predictor to be bounded (so that it has low complexity and generalizes well), and the weighted dis-similarities between related predictors to also be small.

Taking an agnostic PAC-learning approach, our goal is to minimize the overall *population* objective

$$\sum_{i=1}^m F_i(\mathbf{w}_i) + \lambda \|\mathbf{W}\|_M \tag{1}$$

and be competitive with respect to predictors in the set Ω . Denoting $\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W} \in \Omega} F(\mathbf{W})$ the optimal predictor from Ω , and we would like to learn a predictor \mathbf{W} with $F(\mathbf{W}) \leq F(\mathbf{W}^*) + \varepsilon$.

In our analysis, we take the instantaneous loss $\ell(\mathbf{w}, \mathbf{z})$ to be L -Lipschitz continuous, and sometimes also assume it is smooth. In the latter case, we assume machine i 's loss $\ell(\mathbf{w}_i, \mathbf{z}_i)$ is β_i -smooth in \mathbf{w}_i , and so the global loss $F(\mathbf{W}) = \sum_{i=1}^m \ell(\mathbf{w}_i, \mathbf{z}_i)$ is $\beta_F = \max_{i=1, \dots, m} \beta_i$ -smooth in \mathbf{W} .

individual task (i.e. the number of samples from D_i required to ensure $F_i(\mathbf{w}_i) \leq F_i(\mathbf{w}_i^*) + \varepsilon$) is

$\frac{L}{\beta_i}$. That is, with a total of $\frac{L}{\beta_F} \log \frac{1}{\varepsilon}$ samples, we can learn \mathbf{W} with the desired guarantee $F(\mathbf{W}) \leq F(\mathbf{W}^*) + \varepsilon$ without any communication between the machines, by, e.g., solving an independent λ -regularized ERM problem on each machine. This *local* approach is the baseline on which any method involving communication between the machines should improve.

Graph Laplacian

The term $\|\mathbf{L}\|_F$ can be written equivalently using the graph Laplacian. Let $\mathbf{A} = [a_{ik}] \in \mathbb{R}^{m \times m}$ be the adjacency matrix, and $\mathbf{L} = \operatorname{diag}(\mathbf{A}\mathbf{1}) - \mathbf{A}$ be the corresponding

graph Laplacian ($L_{ik} = \sum_{l \neq i} a_{il}$ if $i = k$, and $L_{ik} = -a_{ik}$ otherwise), so that

$\|\mathbf{L}\|_F = \sqrt{\sum_{i,k} L_{ik}^2}$. The eigenvalues of \mathbf{L} will play an impor-

tant role and we denote them by $0 = \lambda_1 \leq \dots \leq \lambda_m$.

Regularized ERM One way for learning regularized empirical risk minimization

the predictors is to solve the (ERM) problem. Let ℓ_i be the local empirical loss of machine i , and let $Z = \{\mathbf{z}_{ij} : i = 1, \dots, m, j = 1, \dots, n\}$ be the sample set. The regularized ERM

objective is

$$= \operatorname{argmin}_{\mathbf{W}} \left\{ \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{w}_i; \mathbf{z}_i) + \frac{\eta}{2m} \sum_{i=1}^m \|\mathbf{w}_i\|^2 + \frac{\tau}{2m} \operatorname{tr}(\mathbf{W}) \right\}$$

where $\eta, \tau \geq 0$ are regularization parameters. Let $\mathbf{W} = \operatorname{argmin}_{\mathbf{W}} \ell(\mathbf{W}) + R(\mathbf{W})$ be the solution

\mathbf{w}

$$\mathbf{W} \mathbf{L} \mathbf{W}^T, \quad (2)$$

to (2).

To understand the statistical property of multi-task learning and facilitate further discussion, we provide same learning guarantee for the solution of a *constrained* ERM problem (i.e., $\operatorname{argmin}_{\mathbf{W} \in \Omega} \ell(\mathbf{W})$), first analyze the generalization error of \mathbf{W} . Inspired by Maurer (2006), who showed essentially the *regularized* ERM rather than *constrained* ERM is that it is easier to solve unconstrained problem using \mathbf{W} we provide guarantee for the *regularized* ERM solution \mathbf{W} . Our motivation for studying regular-

(proximal) gradient methods, and we avoid computing projection onto the constraint set Ω , which is difficult in a distributed setting.¹

While the analysis of Maurer (2006) was based on the Rademacher complexity of Ω (and required the solution to lie in Ω), our proof uses the stability based argument for generalization with strongly convex regularizers (Shalev-Shwartz et al., 2009). Our analysis also reveals a fundamental connection between single- and multi-task learning: to obtain generalization of a single task in the distributed setting, we only need concentration for the sampling process of that task. In our case, we consider strong convexity w.r.t. the $\|\cdot\|_{\mathbf{M}}$ -norm where \mathbf{M} is a positive definite matrix.

Lemma 1. *Assume that the instantaneous loss $\ell(\mathbf{w}, \mathbf{z})$ is L -Lipschitz with respect to \mathbf{w} . Then for the ERM solution defined in (2), we have*

¹ Although for convex optimization, the constrained form and the regularized form are equivalent due to the Lagrange duality, solving the constrained form may still require repeatedly solving the regularized form and searching for the Lagrange multiplier.

Corollary 2. Set $\rho(B, S)$ and $\rho(B, S)$ in (2), where

Then

The quantity $\rho(B, S)$ measures task relatedness and thus the benefit of multi-task learning. It depends on the parameters (B, S) and the graph, but not the data. The value of $\rho(B, S)$ ranges from 0 (when $\rho(B, S) = 0$) to 1 (when $\rho(B, S) = 1$), corresponding to two extreme cases.

- When S is small and the graph is connected with high weights, the predictors are encouraged to be similar to each other (we have a consensus problem if $S = 0$ and the graph is connected), and $\rho(B, S)$ is close to 0. The generalization error is then $O(n^{-1})$, corresponding to that of single task learning using mn samples.
- When S is large or the graph is disconnected, tasks are not very related and $\rho(B, S)$ is close to 1. In this case, the generalization error behaves like $O(n^{-1})$, and we are essentially performing local learning with n samples for each task.

For a fixed number of machines m and graph Laplacian \mathbf{L} , to achieve ε excess population error by the above approach, the number of samples used by each machines is $O(n_L)$. Therefore, when the tasks are related and $\rho(B, S)$ is small, the sample complexity of multi-task learning is significantly smaller than n_L needed by the local approach.

To implement the regularized ERM approach in the distributed setting, we could have each machines send n_c samples to a central machine, and then minimize the regularized empirical loss on that machine. We refer to this baseline as the *centralized* approach—it is sample efficient, but expensive in terms of communication and computation. We are interested in distributed multi-task learning algorithms that are also sample efficient, i.e. use only $O(n_c)$ samples on each machine (or at least, not much more than this), but have low computation and communication costs. This can be done either by low-communication distributed optimization of the regularized empirical error (2).

3 Distributed algorithms for ERM

In this section, we propose efficient distributed algorithms for minimizing the regularized empirical

such updates take the form: \mathbf{b}

objective (2). The simplest approach is perhaps to perform gradient descent on $F(W)$. Interestingly,

$$\mathbf{w}^{t+1} = \left(1 - \alpha^{t+1} \right) \mathbf{w}^t + \alpha^{t+1} \sum_{k \in \mathcal{N}(i)} \mu_{tki+1} \mathbf{w}^k, \quad (3)$$

where $\alpha^{t+1} > 0$ is the stepsize at iteration $t+1$, and the weights for combining neighboring predictors are

$$\mu_{tki+1} = \begin{cases} 1 - \alpha^{t+1} \alpha_t & \text{if } i = k, \\ \alpha^{t+1} \alpha_t \tau_{ik} P_{ko} & \text{otherwise.} \end{cases} \quad (4)$$

With an appropriate step-size schedule (or even a fixed stepsize if the loss is smooth), this method graph, since the update for each machines involves only predictors from neighboring machines converges

to \mathbf{W} . Furthermore, the updates require only communication along the relatedness

(with nonzero affinities). This is already a very natural and intuitive method for distributed multitask learning, and we will return to it later. When the loss is smooth, the method can be accelerated using Nesterov's techniques (Nesterov, 2004, as detailed in Appendix C) without any increase in communication costs nor substantial increase in computation. But first, we suggest two more powerful alternatives.

Taking steps based on the gradients amounts to considering, in each iteration, a linearization of the objective, that is of both the empirical loss $F(\mathbf{W})$ and the regularizer $R(\mathbf{W})$. However, in order to obtain a distributable update, it is sufficient to linearize only one of these components while treating the other more explicitly, since each one of them separately can be efficiently optimized in a distributed way, while $R(\mathbf{W})$ is data independent and could be optimized implicitly in a distributed way: the empirical loss $F(\mathbf{W})$ decomposes over machines, and so can be directly

based on the common knowledge of the relatedness graph. In the following, we consider two distributed schemes, each based on directly handling one of the components, and each preferable in a different regime depending on the relatedness graph and the structure and cost of communication.

3.1 Directly solving the regularizer

We first consider methods which directly handle the regularization term $R(\mathbf{W})$. To do so, we consider the change of variable $\mathbf{U} = \mathbf{M}^{-1} \mathbf{W}$ where $\mathbf{M} = \mathbf{L}$, we can rewrite the ERM objective as

$$\sum_{i=1}^m \ell(\mathbf{U}_i) + \lambda \|\mathbf{U}\|_1 \quad (5)$$

We propose to optimize this objective using gradient descent with respect to \mathbf{U} , which reduces to the updates in the \mathbf{W} -space: for $t = 0, \dots$,

$$\mathbf{W}^{t+1} = \mathbf{M} \mathbf{U}^t \quad (6)$$

where $\alpha^{t+1} > 0$ is the stepsize at iteration $t+1$. In each iteration, machine i performs the following update with $\mu_{tki+1} = \alpha^{t+1} (\mathbf{M}^{-1})_{ki}$:

m

(7)

$$\text{[Redacted Equation]}$$

=1

\mathbf{w} .

This update can be implemented in the distributed setting with a broadcast channel: it requires that each machine has access to gradients of all machines, which can be achieved using one round of global, all-to-all communication (not respecting the graph). We could compute \mathbf{M}^{-1} offline ahead of time, and need not re-calculated at each iteration.

When the loss is smooth, we can accelerate (7) using Nesterov’s techniques without additional communication costs. Setting a constant stepsize η , which is the smoothness parameter of the objective (5) in \mathbf{U}^2 , to achieve ϵ -suboptimality in (2), the iteration complexity of the accelerated algorithm is $\frac{1}{\epsilon}$. To achieve ϵ excess error in the population loss, we set the optimization error $\frac{\epsilon}{m}$ and plug in the choice of η from Corollary 2, yielding $\frac{1}{\epsilon}$ the iteration complexity.

3.2 Directly optimizing the loss

The above algorithm requires dense, broadcast communication for solving the proximal step defined by the graph. In a decentralized setting, it is desired to develop algorithms which use only local,

$$\|\mathbf{M}^{-\frac{1}{2}}\| \cdot \|\nabla_{\text{vec}(\mathbf{W})}^2 \mathfrak{h}(\mathbf{W})\| \cdot \|\mathbf{M}^{-\frac{1}{2}}\| \leq \frac{\beta_F}{m}$$

peer-to-peer communication. This can be achieved by the updates below, where we linearize the graph regularizer but fully optimize over the loss:

$$\mathbf{W}^{t+1} = \underset{\mathbf{W}}{\text{argmin}} \mathfrak{h} \nabla R(\mathbf{W}^t), \mathbf{W} - \alpha^t \mathbf{W}^t$$

\mathbf{w}

$$\text{[Redacted Equation]}$$

(8)

where α^{t+1} is the stepsize at iteration $t+1$. As (8) decouples over machines, machine i independently computes a proximal operation using local data:

$$\mathbf{w}^i = \underset{\mathbf{u}}{\text{argmin}} \text{[Redacted Equation]}$$

$$\text{[Redacted Equation]}$$

By the optimality condition of this update, we have

\mathbf{w}

$$\text{[Redacted Equation]}$$

(9)

² This is because

), and

where the weights for combining neighboring predictors are the same as those in (4). Comparing (9) with the similar update (3) where we linearized both the regularizer and the loss, we observe that (9) is also a form of gradient method, with the gradient of loss evaluated at the “future” point.

The advantage of (9) is that the gradient $\nabla R(\mathbf{W})$ is data-independent and is obtained using only one round of local communication from each machine to its neighbors. Furthermore, the computation decouples over machines, and each machine optimizes the nonlinearized loss without communication. In fact, we need not solve the proximal steps exactly since the (accelerated) proximal gradient method is tolerant to errors in the steps (Schmidt et al., 2011), and sufficiently accurate solutions can often be obtained in time nearly linear in the number of examples processed using variance-reduced finite-sum methods such as SVRG (Johnson and Zhang, 2013). Overall, this is a communication-efficient approach in which each machine tries to spend significant amount of time performing local computations on its own data, and to communicate only infrequently. Note that similar proximal type operations also appear in the ADMM algorithm of Vanhaesebrouck et al. (2017), but the decoupling of tasks is different, because in the local problems of ADMM, each machine optimizes over also a copy of neighboring predictors.

We can again accelerate (9) using Nesterov’s techniques, and set $\tau = \frac{1}{\mu}$, which is the smoothness parameter of $R(\mathbf{W})$ in \mathbf{W} . Then, to achieve ε excess error in the population objective, the number of iterations needed by the accelerated algorithm is, using the choice of η and τ from Corollary 2. We also show that this algorithm is tolerant to delay and analyze its convergence under bounded delay in Appendix G.

4 Stochastic algorithms

In ERM, we collect training samples on each machine ahead of time, and solve a fixed optimization problem defined by them. But in real-world scenarios, we might have access to virtually unlimited data, or a constantly available stream of examples. In this case, it might be statistically wasteful to reuse examples over iterations. Or, even if we do have a finite amount of data, as we shall see, Table 1: Algorithms for distributed stochastic multi-task learning with graph regularization. Here ε is the excess error in the population objective; E denotes the number of edges in the graph. For simplicity, schematic updates ignores acceleration, but the rates are given for the accelerated algorithms. Each cell shall be interpreted as $O_e(\cdot)$ which hides

Algorithms	Communication rounds	Vectors ($\in \mathbb{R}^d$) communicated per machine	Sample complexity per machine	Total Samples processed per machine
local	0	0	$\frac{1}{\varepsilon}$	nL
centralized		nc	$\frac{1}{\varepsilon}$	$m \cdot nc$

poly-logarithmic dependencies.

ERM: directly solving regularizer 1. where(M 2. $\mathbf{w}_{i+1} = \mathbf{w}_i - \mathbf{g}_i$			<i>nc</i>	
ERM: directly optimizing loss 1. where 2. $\mathbf{w}_{i+1} = \mathbf{w}_i \alpha_i$ ($i+1$)			<i>nc</i>	
			<i>nc</i>	<i>nc</i>
Stochastic: directly optimizing loss 1. \mathbf{w}_{kt} 2.			n_s , probably $\in (n_c, n_L)$	<i>ns</i>

e - ∇b

we can get the same communication and statistical guarantee while processing only a minibatch at a time, thus significantly reducing computational cost. We consider stochastic variants of the approaches in Section 3 to directly optimize the population loss $F(\mathbf{W})$, using fresh samples in each update.

4.1 Directly solving the regularizer

Analogous to (7), we could perform minibatch SGD with b samples per machine to approximate the gradient of the population loss: for $t = 0, \dots$,

$$\mathbf{w}_k^{t+1} = \mathbf{w}_k^t - \eta \nabla F(\mathbf{w}_k^t; \mathcal{S}_k^t) \tag{10}$$

where \mathcal{S}_k^t is a set of b samples drawn by machine k at iteration $t + 1$.

We can accelerate (10) using the accelerated stochastic approximation (AC-SA) algorithm of Lan (2012). We provide the detailed accelerated algorithm in both the \mathbf{U} -space and \mathbf{W} -space in Algorithm 2 (Appendix D). We have the following guarantee after running it for T iterations.

Theorem 3. *Set the initialization $\mathbf{W}_0 = \mathbf{0}$ and stepsizes η in Algorithm 2. Then*

Sample complexity Let $n = bT$ be the number of samples used in Algorithm 2. According to Theorem 3, as long as the minibatch size b , the first term in the error bound is dominant and we achieve the generalization error ϵ as in ERM, so we are still sample efficient in the stochastic setting.

Time complexity Algorithm 2 processes the drawn samples only once. While maintaining the sample efficiency, we can set the minibatch size to the largest value $b = b_*$, and this leads to the total number of iterations (and local communication rounds) T , also matching that of ERM. However, since each stochastic gradient uses only $b = o(n)$ samples, the local computation $\nabla F(\mathbf{w}_k^{t+1}; \mathcal{S}_k^t)$ is significantly reduced.

4.2 Directly optimizing the loss

Analogous to (8), we can use the stochastic algorithm where at iteration $t+1$, machine i computes

$$\mathbf{w}_i^{t+1} = \underset{\mathbf{u}}{\operatorname{argmin}} \left[\sum_{\mathcal{S}_i^t} \ell(\mathbf{u}; \mathcal{S}_i^t) + \frac{\lambda}{2} \|\mathbf{u}\|_2^2 \right] \tag{11}$$

For $b = n$, it has the same per iteration computation cost as the ERM counterpart (both process n samples in each iteration). But, intuitively, it would outperform the ERM algorithm for the same number of iterations/communications because it uses more fresh samples. We can prove the convergence of this algorithm, but do not have a satisfactory analysis showing it is sample efficient. We conjecture that its sample complexity per machine, denoted by n_s , is in the range (n_C, n_L) . We

implemented the accelerated version of this simple algorithm and this conjecture seems to be supported by our experiments. In Appendix E, we provide a more complicated algorithm based on the minibatch-prox algorithm of Wang et al. (2017), that is sample efficient and trade off communication and memory costs.

Comparison of the different approaches Table 1 summarizes the communication and computation complexities of the proposed algorithms. Some of our methods require solving local regularized-ERM type problems on each machine. We do not analyze the precise complexity and required accuracy of such local computation, but keep track of the number of samples processed on each machine, i.e. sum of the sizes of the subproblems over the iterations, as the proxy for computational complexity. We emphasize that, despite the simplicity of our ERM methods, their have faster convergence than what we could obtain for previous methods; see detailed discussions in Appendix H. Our stochastic algorithms mirror the ERM algorithms in terms of updates, but can be computationally much more efficient.

5 Connection to consensus learning

The iterations we consider all involve taking a weighted average of messages (iterates or gradients) from other machines and a local gradient or prox computation. These same type of iterates have also been suggested and studied as methods for solving the consensus problem—that is, finding a single consensus predictor \mathbf{w} that is good for all machines and minimizes

But the consensus problem is fundamentally different from our “pluralistic” multi-task problem, with a different optimum. In this section we will understand what makes the same form of updates, namely updates of the form (3), (7), (9) or their stochastic variants, converge to either the consensus solution or to the pluralistic multi-task solution. In particular, we show how consensus methods are obtained as special cases of these updates, or as limits of the multi-task approach.

Averaging gradients Let us begin with the update of the form (7) or its stochastic variant (10), where we take a weighted average of gradients from other machines. When the averaging weights are uniform, i.e. $\mu_{ki}^t = \alpha^t/m$ for all i, k , and as long as all machines start from the same initialization (e.g. $\mathbf{w}^0 = 0$), the iterates will continue to be identical across machines throughout optimization (i.e. we will have $\mathbf{w}_i^t = \mathbf{w}_j^t$ for all i, j, t), thus maintaining consensus. Furthermore, the update (7) then boils down to precisely gradient descent on the empirical consensus objective

$\sum_{i,j} \mu_{ij}^t \ell(\mathbf{w}_i^t, \mathbf{w}_j^t)$, while the stochastic variant (7) is precisely a mini-batch stochastic gradient descent update on the consensus objective, with a mini-batch consisting of the union of the samples used across machines. Indeed, mini-batch SGD is a common approach for solving the distributed consensus problem, or for distributed learning in a homogeneous setting (where we assume the same distribution across machines, or at least the same good predictor). What we saw in Section 3, is that by changing to non-uniform weights, given by $\mu \propto \mathbf{M}^{-1}$, we can allow pluralism and converge to the multi-task solution.

We can furthermore observe how uniform weights (and therefor gradient descent/mini-batch SGD on the consensus problem) are obtained as a limit of the multi-task weights $\mu \propto \mathbf{M}^{-1}$. If the graph is connected, $\lambda_1 = 0$ is the only zero eigenvalue of the Laplacian L with an associated eigenvector of \mathbf{u}

$= [1, \dots, 1]$ (if the graph is not connected, we cannot expect consensus, as each connected component will behave independently). Therefore \mathbf{M} has a leading eigenvalue of 1 of multiplicity one, associated with the eigenvector \mathbf{u} . As $S \rightarrow 0$ and so $\tau \rightarrow \infty$, that is we are demanding increasing similarity between machines, the leading eigenvalue of \mathbf{M}^{-1} remains 1 while all other eigenvalues go to zero, implying that \mathbf{M} and so $\mu_{ki}^t = \alpha^t \mathbf{M}^{-ki} \rightarrow \alpha^t / m$. That is, as we demand increasing similarity between machines, and thus converge to a consensus situation, the updates converge to standard consensus gradient descent or mini-batch SGD updates.

Averaging iterates Let us now turn to updates of the form (3), the related prox updates (9), and their stochastic variants. Nedić and Ozdaglar (2009) proposed updates precisely of the form (3) as a decentralized procedure for the consensus problem. They showed that when the averaging weights $P_k \mu_{ki}^t$ are doubly stochastic and do not vary between iterations (i.e. $\mu_{ki}^t \rightarrow \mu_{ki}, \forall k, P_i \mu_{ki} = 1$ and $\forall_j \mu_{ki} = 1$), and the stepsize on the gradient goes to zero, i.e. $\alpha^t \rightarrow 0$, the updates (3)

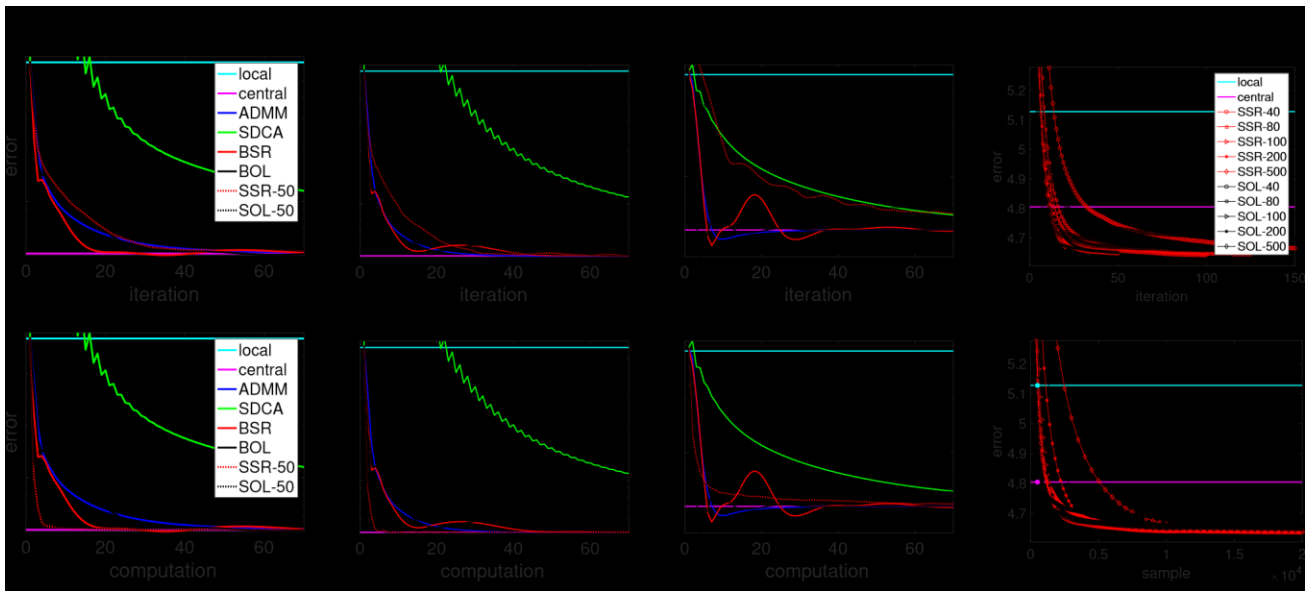


Figure 1: Results for regularized ERM (left panel) and our stochastic methods with different b (right panel).

converge to the consensus solution. In our case, the averaging weights, as defined in (4), deviate from double-stochasticity, since $P_k \mu_{ki}^t = 1 - \alpha^t \eta$. Furthermore, and possibly more significantly, to obtain our convergence guarantees for smooth loss, we do not take α^t to zero. Even if we were to use diminishing

stepsizes in our derivations, we would have $\alpha^t \rightarrow 0$, but in that case the averaging weights would not be fixed over iterations (as is the case in consensus optimization) and we would have $\mu^t \rightarrow \mathbf{I}$.

To see how consensus updates are obtained as a limiting case of our multi-task setting, we again consider a connected graph and study what happens as $S \rightarrow 0$ and so $\tau \rightarrow \infty$, while B and therefore η remain fixed. This corresponds to a fixed amount of local regularization, and increasing expectation that neighboring nodes are similar. Under this scaling, we would indeed have $\alpha = 1/(\eta + \tau\lambda_m) \rightarrow 0$, where $\lambda_m > 0$ since the graph is connected. Furthermore, we have that $\alpha\eta \rightarrow 0$ while $\alpha\tau \rightarrow 1/\lambda_m > 0$. Plugging this scaling into the multi-task averaging weights (4), we obtain the doubly stochastic weights:

$$\begin{aligned}
 & \lambda_{ik} = \alpha \sum_{m=1}^m a_{ik} && : \text{if } i = k, \\
 & \lambda_{ik} = \alpha \sum_{m=1}^m a_{ik} && : \text{otherwise.}
 \end{aligned} \tag{12}$$

To summarize, a significant differentiation between consensus and multi-task learning is therefore in whether α^t diminishes *relative to* $(\mu^t - \mathbf{I})$. When our relatedness constraints approach consensus, α^t can diminish while μ^t is non-trivial and doubly stochastic. In fact, in studying consensus optimization, Yuan et al. (2016) recently noted that when α^t does not diminish, the methods does not converge to the consensus solution but only to a neighborhood of it. In light of our analysis, we now understand that this “neighborhood” corresponds to the multi-task learning solution, which indeed becomes increasingly similar to the consensus solution as $S \rightarrow 0$.

Connection to the decentralized algorithm of Scaman et al. (2017) When the graph is

connected, the consensus constraint $\mathbf{w}_1 = \dots = \mathbf{w}_m$ can be equivalently written as $\mathbf{W}\sqrt{\mathbf{L}} = \mathbf{0}$, since the null space of \mathbf{L} contains only vectors of constants. Then the multi-task formulation (2) is a relaxation of

$$\mathbf{w}_k = \frac{1}{\sqrt{\mathbf{L}}_{k,k}} \sum_{i=1}^m X_{i=1} \mathbf{w}_i \tag{13}$$

with the quadratic term $\frac{1}{2} \text{tr} \mathbf{W}\mathbf{L}\mathbf{W}^T$ penalizing the constraint violation. The quadratic penalty \rightarrow

$\infty \frac{1}{2} \text{tr} \mathbf{W}\mathbf{L}\mathbf{W}^T$ may lead to a large condition number for our algorithm (8) as $\tau \rightarrow \infty$.

Recently, Scaman et al. (2017) proposed an algorithm with optimal iteration/communication complexities for decentralized consensus learning, which performs accelerated gradient descent on the dual problem of (13), with updates (before acceleration):

$$\begin{aligned} \mathbf{W}^{t+1} &= \operatorname{argmax}_{\mathbf{W}} \text{[redacted]}, \\ \mathbf{V}_{t+1} &= \mathbf{V}_t - \alpha \mathbf{W}_{t+1} \mathbf{L}, \end{aligned} \tag{14}$$

where $\mathbf{V}^0 = \mathbf{0}$ and $\alpha > 0$ is the stepsize. It can be seen that their algorithm consists of the same type of basic operations (weighted local average of predictors, and solutions of local subproblems involving non-linearized loss) as ours. As noted by the authors, this is a form of distributed augmented Lagrangian method without the quadratic penalty.

6 Experiments

We examine different graph-based multi-task learning methods on the task of least squares regression using synthetic data. More details of the experiments (including data generation and more results) are given in Appendix I. The tasks are grouped into C clusters and the true predictors within the same cluster are generated from the same Gaussian distribution, thus smaller C implies higher task relatedness. We have input dimension $d = 100$, number of tasks $m = 100$, training set size $n = 500$, and vary number of task clusters C over $\{1, 5, 10, 50\}$. We also generate a dev set of 10000 samples per task for tuning hyper-parameters, and test set of 10000 samples per task for approximately evaluating the population loss. The affinity graph $\mathbf{A} \in \mathbb{R}^{100 \times 100}$ is a (connected) 10-nearest neighbor graph with binary weights built on the true predictors.

The methods compared here are: Local, which solves a local ERM problem (with ℓ_2 -regularization) with n samples for each task; Centralized, which solves the regularized ERM problem (2) with n samples for each task; ADMM, which is the synchronized version of the algorithm of Vanhaesebrouck et al. (2017); SDCA, which is the algorithm used by Liu et al. (2017) for fixed graph; our algorithms are denoted as B/S (batch/stochastic) + SR/OL (solve regularizer/optimize loss).

Empirical risk minimization We first compare the iterative methods on the regularized ERM problem (2), to which the analysis for ADMM and SDCA applies. We tune the ℓ_2 regularization parameter for Local and (η, τ) for Centralized, and then fix the optimal (η, τ) for other methods. We also tune the quadratic penalty parameter for ADMM, the task separability and stepsize parameters for SDCA, and stepsize parameter for BSR/BOL (although the default value based on the smoothness parameter already works well for them). For SSR/SOL, we draw random samples from the fixed training set (with size n), and simply fix the minibatch size to be $n/10$.

Figure 1 (left panel) shows for each method the estimated $F(\mathbf{W})$ over iterations (or rounds of communication) in the top row, and over the amount of computation (measured by the number of passes over the training set) in the bottom row. Observe that all iterative algorithms converge to the same ERM solution, our algorithms tend to consistently outperform ADMM and SDCA.

Stochastic optimization We next demonstrate the efficiency of true stochastic algorithms (using fresh samples for each update) at $C = 10$. We allow the algorithms to process a total of 10000 fresh samples on each machine, and vary the minibatch size b over $\{40,80,100,200,500\}$. The parameters (η, τ) are fixed to those used in the ERM experiments.

Figure 1 (right panel) shows for each method the estimated $F(\mathbf{W})$ over iterations (or rounds of communication) in the left plot, and over the amount of fresh samples processed (or total computation cost) in the right plot. As a reference, the error of Local and Centralized (using $n = 500$ samples per machine) are also given in the plots. We observe that with fresh samples, stochastic algorithms are competitive to ERM algorithms in terms of sample complexity, while being computationally more efficient.

A Proof of Lemma 1

Recall that the ERM problem is defined as

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \sum_{i=1}^m \sum_{j=1}^n \ell(\mathbf{w}; \mathbf{z}_{ij}) + \frac{\eta}{2} \mathbf{w}^T \mathbf{L} \mathbf{w},$$

where $\eta, \tau \geq 0$ are regularization parameters, $Z = \{\mathbf{z}_{ij} : i = 1, \dots, m, j = 1, \dots, n\}$ is the sample set. And recall that $\lambda_i, i = 1, \dots, m$ are the eigenvalues of \mathbf{L} .

Assume that the instantaneous loss $\ell(\mathbf{w}; \mathbf{z})$ is L -Lipschitz in \mathbf{w} . We would like to show that

$$\sum_{i=1}^m \sum_{j=1}^n \ell(\mathbf{w}; \mathbf{z}_{ij}) + \frac{\eta}{2} \mathbf{w}^T \mathbf{L} \mathbf{w} \leq \sum_{i=1}^m \sum_{j=1}^n \ell(\mathbf{w}^*; \mathbf{z}_{ij}) + \frac{\eta}{2} \mathbf{w}^{*T} \mathbf{L} \mathbf{w}^* + \frac{L^2}{2\eta} \sum_{i=1}^m \sum_{j=1}^n \|\mathbf{z}_{ij}\|^2$$

Proof. In the following, we define $\mathbf{M} = \frac{1}{n} \sum_{j=1}^n \mathbf{z}_{ij} \mathbf{z}_{ij}^T$ which is positive definite. Furthermore, perform the following change of variables

$$\mathbf{u} = \mathbf{W} \mathbf{M}^{-1/2} \mathbf{w}, \quad \mathbf{w} = \mathbf{M}^{1/2} \mathbf{u}$$

where \mathbf{e}_i is the i -th standard basis in \mathbb{R}^m .

We can then rewrite the losses using the new variables:

$$\ell(\mathbf{w}; \mathbf{z}_{ij}) = \ell(\mathbf{M}^{1/2} \mathbf{u}; \mathbf{z}_{ij}) = \ell(\mathbf{u}; \mathbf{M}^{-1/2} \mathbf{z}_{ij}), \quad \text{for } i = 1, \dots, m,$$

and the empirical objective as

$$\sum_{i=1}^m \sum_{j=1}^n \ell(\mathbf{w}; \mathbf{z}_{ij}) + \frac{\eta}{2} \mathbf{w}^T \mathbf{L} \mathbf{w} = \sum_{i=1}^m \sum_{j=1}^n \ell(\mathbf{u}; \mathbf{M}^{-1/2} \mathbf{z}_{ij}) + \frac{\eta}{2} \mathbf{u}^T \mathbf{L} \mathbf{u}. \quad (15)$$

We can view (15) as performing ERM in the space of \mathbf{U} , using the instantaneous loss $h_1(\mathbf{U}, \mathbf{z}_1)$ with n independent samples $\{\mathbf{z}_{1j}\}_{j=1, \dots, n}$, and using the term in bracket as the \mathbf{z}_1 -independent regularizer.

Recall that the ERM solution to an objective with Lipschitz loss and strongly convex regularizer is stable. Obviously, the regularization term in (15) is λ -strongly convex in \mathbf{U} . We now bound the Lipschitz constant of $h_1(\mathbf{U}, \mathbf{z}_1)$ in \mathbf{U} . Observe that

$$\|h_1(\mathbf{U}, \mathbf{z}_1) - h_1(\mathbf{U}', \mathbf{z}_1)\| \leq L \|\mathbf{U} - \mathbf{U}'\|$$

and as a result the Lipschitz constant is bounded by

$$L \|\mathbf{U} - \mathbf{U}'\|$$

where we have used the L -Lipschitz continuity of $h_1(\mathbf{w}_1, \mathbf{z}_1)$ which implies $\|\nabla_{\mathbf{w}_1} h_1(\mathbf{w}_1, \mathbf{z}_1)\| \leq L$.

According to Shalev-Shwartz et al. (2009)[Theorem 6], for any fixed \mathbf{z}_1 , it holds for the ERM solution $\mathbf{U}_b = \operatorname{argmin}_{\mathbf{U}} \sum_{j=1}^n h_1(\mathbf{U}, \mathbf{z}_{1j}) + \lambda \|\mathbf{U}\|^2$ that

$$\|\mathbf{U}_b - \mathbf{U}^*\| \leq \frac{L}{\lambda} \sqrt{\frac{2 \log(2/\delta)}{n}}$$

Translating this in terms of the original variables, we have

$$\|\mathbf{w}_b - \mathbf{w}^*\| \leq \frac{L}{\lambda} \sqrt{\frac{2 \log(2/\delta)}{n}}$$

where

By the

convexity of


$$\sum_{j=1}^n h_1(\mathbf{w}_1, \mathbf{z}_{1j}) + \lambda \|\mathbf{w}_1\|^2$$

$$\sum_{j=1}^n h_1(\mathbf{w}_1, \mathbf{z}_{1j})$$

and the Jensen's inequality, this implies

This result shows that, to obtain generalization for a single task, we only need concentration for the sampling process of that task. By the same argument, we obtain similar inequalities regarding stability for losses on each machine.


Finally, we have by the triangle inequality that



which is what we set out to prove.

□


B Proof of Lemma 2

Based on Lemma 1, we now show that by properly setting the regularization parameters in the regularized ERM problem (2), i.e., , we have that



where

Proof. Observe that



where we have used Lemma 1 in the first inequality, and that \mathbf{W} is the empiric risk minimizer in Since

$\mathbf{W}^* \in \Omega$, we can bound the excess error as

the third inequality.

c

$$\text{[Redacted Equation]}$$

(16)

Now, set α and β for some α, β that will be specified later. Continuing from (16) yields

$$\text{[Redacted Equation]}$$

Minimizing the RHS over α gives α^* , and

$$\text{[Redacted Equation]}$$

□

C The accelerated proximal gradient algorithm

We provide the accelerated proximal gradient algorithms in Algorithm 1, which are used to accelerate our ERM algorithms in the main text. The proximal operator is defined as $\text{prox}^{\beta}_h(\mathbf{x}) = \arg\min_{\mathbf{y}} \left(\frac{\beta}{2} \|\mathbf{x} - \mathbf{y}\|^2 + h(\mathbf{y}) \right)$ where $\beta > 0$ and $h(\mathbf{x})$ is convex and possibly non-smooth.

Algorithm 1 ProxGrad(g, h, β, μ): Accelerated proximal gradient descent.

Input: Objective has the form $f(\mathbf{w}) = g(\mathbf{w}) + h(\mathbf{w})$, where $g(\mathbf{w})$ is β -smooth and μ -strongly convex, and $h(\mathbf{w})$ is convex. Initialize $\mathbf{w}^0, \mathbf{y}^1 \leftarrow \mathbf{w}^0$ **for** $t = 1, \dots, T$ **do**

$$\mathbf{w}^t \leftarrow \text{prox}^{\beta}_g(\mathbf{y}^t), \quad \mathbf{y}^{t+1} \leftarrow \text{prox}^{\beta}_h(\mathbf{w}^t)$$

end for

Output: \mathbf{w}^T is the approximate solution.

D Analysis of stochastic optimization by directly solving the regularizer

In each iteration of this algorithm, we draw b samples per machine to approximate the gradient of the population loss and perform minibatch SGD, which amounts to linearizing the loss on a minibatch. The key to being sample efficient is to respect the geometry imposed by the graph Laplacian.

As in Section 3.1, define the change of variable \mathbf{U} where $\mathbf{M} \mathbf{L}$. Our population objective is $\mathcal{L}(\mathbf{U})$, and the predictor \mathbf{U} satisfies the constraint that $\mathbf{U} \mathbf{L} = \mathbf{0}$. We can perform minibatch SGD in the \mathbf{U} -space:

$$\mathbf{U}^{t+1} = \underset{\mathbf{U}}{\operatorname{argmin}} \left[\mathcal{L}(\mathbf{U}) + \frac{\alpha^{t+1}}{2} \mathbf{U}^T \mathbf{L} \mathbf{U} \right], \quad \text{for } t = 0, \dots,$$

where b samples drawn by machine i at iteration $t+1$, and $\alpha^{t+1} > 0$ is a stepsize parameter. In the \mathbf{W} -space, the above update reduces to

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \alpha^{t+1} \nabla F_{b_{t+1}}(\mathbf{W}_t) \cdot \mathbf{M}^{-1},$$

Clearly, this update requires inverting the graph Laplacian.

We can further accelerate this method using the accelerated stochastic approximation (ACSA) algorithm of Lan (2012). We give the detailed stochastic algorithm by directly solving the regularizer (with linearized loss) in Algorithm 2.

Algorithm 2 Accelerated minibatch SGD. This algorithm maintains three iterate sequences: $\{\mathbf{W}_t\}$ is the sequence of prox centers, $\{\mathbf{W}_{mdt}\}$ is the “middle” sequence with which we evaluate the stochastic gradient and build models (approximations) of the objective, and $\{\mathbf{W}_{ag}^t\}$ is the “aggregated” sequence with which we evaluate the objective values.

Input: The stepsize sequences $\{\alpha_t\}$.

Initialize $\mathbf{W}^0 \leftarrow \mathbf{0}$, $\mathbf{W}_{ag}^0 \leftarrow \mathbf{W}^0$ for $t = 0, \dots, T - 1$ do

$\mathbf{W}_{mdt} \leftarrow \mathbf{W}_t$
 $\mathbf{W}_{t+1} \leftarrow \mathbf{W}_t - \alpha^{t+1} \nabla F_{b_{t+1}}(\mathbf{W}_{mdt}) \cdot \mathbf{M}^{-1}$

$\mathbf{W}_{ag}^{t+1} \leftarrow \mathbf{W}_{ag}^t + \alpha^{t+1} (\mathbf{W}_{t+1} - \mathbf{W}_{ag}^t)$

end for

Output: \mathbf{W}_{ag}^T (or equivalently \mathbf{U}_{ag}^T) is the approximate solution.

The key quantity for analyzing the convergence property of minibatch SGD is the variance of stochastic gradients in the \mathbf{U} -space, which we now derive. We can view $\xi = (\mathbf{z}_1, \dots, \mathbf{z}_m)$ as the combined sample, $\mathcal{L}(\mathbf{W}, \xi)$ as the averaged instantaneous loss, so that

$\mathcal{L}(\mathbf{W}, \xi)$ approximates $E_{\xi}[\mathcal{L}_{multi}(\mathbf{W}, \xi)]$ with b combined samples. The lemma below bounds the variance of stochastic gradient estimated with one combined sample.

Lemma 4. *The variance of stochastic gradient in the \mathbf{U} -space is bounded:*

where

Proof. By direct calculation, we have

$$\text{tr} \mathbf{M} = \sum_{i=1}^m \sum_{k=1}^m \text{tr} \mathbf{M}_{ik} = \sum_{i=1}^m \sum_{k=1}^m \text{tr} \mathbf{M}_{ii} + \sum_{i \neq k} \text{tr} \mathbf{M}_{ik}$$

$$\text{tr} \mathbf{M}_{ii} = \sum_{k=1}^m \text{tr} \mathbf{M}_{ik} = \sum_{k=1}^m \text{tr} \mathbf{M}_{ki}$$

where we have used the independence between \mathbf{z}_i and \mathbf{z}_k for $i \neq k$ so that the cross terms vanishes in (17), and the triangle inequality and that $\|\nabla_{\mathbf{w}_i} \ell(\mathbf{w}_i, \mathbf{z}_i)\| \leq L$ in the inequality. \square

Averaging the b independent stochastic gradients on a minibatch reduces the gradient variance to σ^2/b (see, e.g., Dekel et al., 2012, eqn 7). Note that μ is the smoothness parameter of \mathcal{L} w.r.t. \mathbf{U} , and the distance generating function \mathcal{D} is 1-strongly convex w.r.t. the $\|\cdot\|_F$ -norm. Plugging these problem parameters into (Lan, 2012)(Corollary 1) yields Theorem 3.

E A sample-efficient stochastic algorithm by directly optimizing the loss

The key to sample efficiency in the stochastic setting is to couple the individual learning tasks with the graph, and respect the geometry of the \mathbf{U} -space (e.g., in deriving the generalization performance in Lemma 1, we rely on strong convexity in the norm $\|\cdot\|_F$). This motivates us to derive a sample-efficient stochastic algorithm based on the minibatch-prox method (Wang et al., 2017). The minibatch-prox method solves a subproblem involving nonlinearized loss on a minibatch in each iteration, and was shown to have the optimal sample complexity for stochastic convex optimization regardless of the minibatch size (recall from Section 4.1 that minibatch SGD achieves the optimal sample complexity only for small enough minibatch size), and it was the basis for developing communication- and memory-efficient algorithm for distributed stochastic consensus learning in Wang et al. (2017).

We detail the minibatch-prox based algorithm in Algorithm 3, which consists of two nested loops. In the outer loop, we perform minibatch-prox in the space of \mathbf{U} ; in each iteration of the outer loop we use b samples per machines to approximate the nonlinearized loss, and approximately solves a subproblem involving the full Laplacian in the \mathbf{W} -space. The solutions to the subproblems (which is then a small ERM problem with fixed samples) are computed approximately by the inner loops, where we perform accelerated gradient descent in the space of \mathbf{W} .

Algorithm 3 Distributed minibatch prox.

Initialize $\mathbf{W}^0 \leftarrow \mathbf{0}$ **for** $t = 0, \dots, T - 1$ **do**
 Approximately solve
 $\mathbf{W}^{t+1} \approx \mathbf{W}^{c^{t+1}} = \operatorname{argmin}_{\mathbf{W}} \mathbf{M}(\mathbf{W} - \mathbf{W}^t) \mathbf{M}(\mathbf{W} - \mathbf{W}^t)$
 to ζ_{t+1} -suboptimality using the accelerated proximal gradient algorithm
 ProxGrad **end**
for
Output: $\bar{\mathbf{W}} = \frac{1}{T} \sum_{t=1}^T \mathbf{W}^t$ is the approximate solution.

The minibatch-prox algorithm for minimizing $F(\mathbf{U} \mathbf{M}^{-\frac{1}{2}})$ works as follows:

$$\mathbf{U}^{t+1} \approx \mathbf{U}^b = \operatorname{argmin}_{\mathbf{U}} \mathbf{M}(\mathbf{U} - \mathbf{U}^t) \mathbf{M}(\mathbf{U} - \mathbf{U}^t), \quad \text{for } t = 0, \dots, \quad (18)$$

Note that we allow inexact solutions to the objective in (18). The where in each iteration we draw b fresh samples per machine to approximate $F(\mathbf{W})$ by $F^{t+1}(\mathbf{W}) =$ corresponding update of (18) in the \mathbf{W} -space is $\mathbf{W}^{t+1} \approx \mathbf{W}^{c^{t+1}} = \operatorname{argmin}_{\mathbf{W}} f^{t+1}(\mathbf{W})$ where

$$\mathbf{W}^{t+1} \approx \mathbf{W}^{c^{t+1}} = \operatorname{argmin}_{\mathbf{W}} \mathbf{M}(\mathbf{W} - \mathbf{W}^t) \mathbf{M}(\mathbf{W} - \mathbf{W}^t) + \frac{\zeta_{t+1}}{2} \|\mathbf{W} - \mathbf{W}^t\|_F^2 \quad (19)$$

We provide the learning guarantee of the minibatch-prox algorithm in the following theorem.

Theorem 5. Suppose that we initialize Algorithm 3 with $\mathbf{W} = \mathbf{0}$ and set $\zeta_t = \frac{1}{t}$. Assume that for all $t \geq 0$, the error in minimizing (19) satisfies

$$\mathbf{M}(\mathbf{W}^{t+1} - \mathbf{W}^t) \mathbf{M}(\mathbf{W}^{t+1} - \mathbf{W}^t) \leq \frac{\zeta_{t+1}}{2} \|\mathbf{W}^{t+1} - \mathbf{W}^t\|_F^2$$

—T

Then for \mathbf{W} , [REDACTED] we [REDACTED] have.

Proof. Let [REDACTED] where $\text{tr} \mathbf{M}$ [REDACTED]). By an analysis similar to that of Lemma 1 (and essentially due to $f^{t+1}(\mathbf{W})$'s γ -strong convexity w.r.t. the norm $\|\cdot\|_{\mathbf{M}}$), we obtain the

“stability” of the exact minimizer to (19), i.e., [REDACTED]. γ -strong convexity of $f^{t+1}(\mathbf{W})$ w.r.t. the Euclidean norm, we have [REDACTED]. Furthermore, if the suboptimality of \mathbf{W}^{t+1} satisfies [REDACTED] $f^{t+1}(\mathbf{W}^{t+1}) - f^{t+1}(\mathbf{W}^{t+1}) \leq \zeta_{t+1}$, by the

$$, \quad \text{for } i = 1, \dots, m,$$

and consequently by the Lipschitz continuity of the loss, we have

$$[REDACTED]$$

This reconstructs the essential lemma required by the minibatch-prox analysis (Wang et al., 2017, Lemma 2). We can then invoke the learning guarantee of minibatch-prox (Wang et al., 2017, Theorem 7), by using our L_U in place of their L , and our [REDACTED] $_{\text{tr}(\mathbf{M})}$ in place of their η_t . In the end, we have

$$[REDACTED]$$

□

For fixed $n = bT$, minibatch-prox attains the generalization error [REDACTED] for any minibatch size b . Though the error in solving each subproblem (19) seems stringent as it decreases over iterations, we can apply the linearly convergent accelerated proximal gradient method in the inner loops to the subproblems. For any minibatch size b , the number of outer iterations is $T = n/b$, and the number of inner iterations for each outer iteration (the initial [REDACTED] error for the subproblems are bounded with a warm-start, see Appendix F) is, so [REDACTED] the total number of communication rounds is the multiplication [REDACTED]

$$[REDACTED]$$

This algorithm allows us to trade off communication and memory: We could use small number of samples b in each outer iteration (limited by the local memory), but the total number communication rounds increase with [REDACTED]. The most communication-efficient setting is $b = n$, in which case we are essentially solving one ERM problem with mn samples (by linearizing the regularizer). Finally, we note that each update of the simple algorithm (11) (without the outer+inner loop structure) and a single inner iteration of the minibatch-prox subproblem (19) have the same communication/computation costs.

F Warm start when directly optimizing the loss

Lemma 6. Consider the objective of the proximal operator

$$\mathbf{x} = \operatorname{argmin}_{\mathbf{y}} \left\{ \beta \|\mathbf{y} - \mathbf{x}\| + h(\mathbf{y}) \right\}$$

where $h(\mathbf{y})$ is L -Lipschitz, and let $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{y}} f(\mathbf{y})$. Then we have

$$\|\mathbf{x}^* - \mathbf{x}\| \leq L/\beta,$$

and the suboptimality of \mathbf{x} is bounded

$$f(\mathbf{x}) - f(\mathbf{x}^*) \leq L^2/\beta.$$

Proof. By the first-order optimality of \mathbf{x}^* , we have

$$\mathbf{0} = \beta(\mathbf{x}^* - \mathbf{x}) + \nabla h(\mathbf{x}^*)$$

where $\nabla h(\mathbf{x}^*)$ is a subgradient of h at \mathbf{x}^* . By the assumption that $h(\mathbf{y})$ is L -Lipschitz, we have $\|\nabla h(\mathbf{x}^*)\| \leq L$ and consequently $\|\mathbf{x}^* - \mathbf{x}\| = \|\nabla h(\mathbf{x}^*)\|/\beta \leq L/\beta$.

For the suboptimality of \mathbf{x} , it follows again from the Lipschitz continuity of h that

$$f(\mathbf{x}) - f(\mathbf{x}^*) = \beta \|\mathbf{x}^* - \mathbf{x}\| + h(\mathbf{x}) - h(\mathbf{x}^*) \leq \beta \|\mathbf{x}^* - \mathbf{x}\| + L \|\mathbf{x} - \mathbf{x}^*\| \leq L \|\mathbf{x}^* - \mathbf{x}\| \leq L^2/\beta.$$

□

This lemma indicates that for solving the local objectives when directly optimizing the loss, e.g., (8), we can initialize from $\mathbf{W}^t - \beta \nabla R(\mathbf{W}^t)$ which mixes the local predictor with those of the neighbors, and the initial suboptimality of this warm start is bounded by L^2/β .

A similar result holds when the distance term is defined by other non-Euclidean norms. For example, in Section 4.1, we need to solve subproblems of the form (19), where the distance in the \mathbf{W} -space is defined by the $\|\mathbf{W}\|_{\mathbf{M}}$ -norm. By an analysis similar to that of Lemma 6 and noting that $\|\mathbf{W}^t - \beta \nabla R(\mathbf{W}^t)\|_{\mathbf{M}} \leq 1$, we obtain the distance between \mathbf{w}_t and the optimal solution \mathbf{w}^* is at most L/γ . As a result, the suboptimality of solving (19) when initialized from \mathbf{W}^t is at most L^2/γ .

G Directly optimizing the loss with bounded delays

When directly optimizing the loss (while linearizing the regularizer), consider the case where the synchronization step is not perfect. Instead of waiting for neighboring machines to finish their local proximal step and sending in their new weight parameters, each machine can use the stale parameters for neighboring machines. Can we still solve the original ERM problem in this case?

Consider the iteration $t + 1$ on machine i (with delays, t is now considered a local iteration counter). Let the set of neighboring machines be N_i . Due to delay in communication, we have a noisy

gradient $\sum_{k \in N_i} \beta \nabla \ell_k(\mathbf{w}_k(t - d_{ik}(t)))$.

Here $d_{ik}(t) \in [0, \Gamma]$ is the delay of machine k relative to machine i (at iteration $t + 1$): Machine i is using the weight of machine k from $d_{ik}(t)$ steps ago. In this section, we allow the delay to vary over time, as long as it is upper bounded by Γ .

Based on this noisy $\sum_{k \in N_i} \beta \nabla \ell_k(\mathbf{w}_k(t - d_{ik}(t)))$ gradient, machine i computes the following proximal gradient step

$$\mathbf{w}_i(t+1) = \text{prox}_{\beta \gamma}(\mathbf{w}_i(t) + \beta \sum_{k \in N_i} \nabla \ell_k(\mathbf{w}_k(t - d_{ik}(t))))$$

with some stepsize $\beta > 0$. We need to analyze the convergence of the proximal gradient method with errors in the gradient, as done by Schmidt et al. (2011). The difference from their work is that the error in our gradients comes from delay (stale weight parameters). Comparing with the case without delay, we have the “error” in the local gradient:

$$\sum_{k \in N_i} \beta \nabla \ell_k(\mathbf{w}_k(t - d_{ik}(t))) - \sum_{k \in N_i} \beta \nabla \ell_k(\mathbf{w}_k(t)).$$

From iteration $t - d_{ik}(t)$ to iteration t , the k -th machine has performed $d_{ik}(t)$ gradient proximal operations. The intuition is that, by the non-expansiveness of the proximal operator, the error in gradient would not cause too much error in the iterates, and then by the smoothness of the objective, this would in turn only results in small error in gradient of the next step. It is important to note that, all machines are influenced by each other and the local errors are propagated to the entire graph.

Based on the non-expansive property of the proximal operator and the additional assumption of the adjacency matrix being doubly-stochastic, it is straightforward to show the following convergence guarantee for the (non-accelerated) proximal gradient algorithm. The algorithm converges at a slower linear rate than without delays.

Theorem 7. Assume that the affinity matrix \mathbf{A} is doubly-stochastic, i.e., $\sum_{k \in N_i} a_{ik} = 1$ for all i , and the delay in the update rule (20) has delay bounded by Γ . Set the inverse stepsize α

Then after $t \geq 1$ iterations of the algorithm, we have

$$\| \mathbf{W}^t \mathbf{c} - \mathbf{W}^t \mathbf{c}^* \| \leq \alpha \Gamma \sum_{k=1}^t \alpha^k$$

Proof. Since $\mathbf{W} \mathbf{c}^*$ is the optimal solution to the ERM problem, we have that

$$\mathbf{W} \mathbf{c}^* \in \arg \min_{\mathbf{c}} \sum_{k=1}^t \alpha^k \ell_k(\mathbf{c})$$

Then, by the non-expansiveness of the proximal operator, we obtain where

$$\| \mathbf{W}^t \mathbf{c} - \mathbf{W}^t \mathbf{c}^* \| \leq \alpha \sum_{k=1}^t \alpha^k \sum_{i=1}^n \sum_{j=1}^n a_{ij} \| \mathbf{c}_i - \mathbf{c}_j \|^2$$

we have used the triangle inequality in the second inequality.

Assume that the affinity matrix \mathbf{A} is doubly-stochastic, so that $\sum_{k \in N_i} a_{ik} = 1$ for all i . De-

note [REDACTED]. Then (21) implies that [REDACTED] βm^F
 $\max_{t-\Gamma \leq t_0 \leq t} V(t_0)$ holds for all i , and as a result

$$[REDACTED]$$

As long as [REDACTED], we have [REDACTED] 1]. Then according to Feyzmhdavian et al. (2014, Lemma 3), we have

$$[REDACTED]$$

Setting β to be the smallest possible value [REDACTED] yields the desired result. □

H Comparisons with previous distributed multi-task learning algorithms

We now provide upper bounds of the iteration complexities for the distributed multi-task learning algorithms of Vanhaesebrouck et al. (2017) and Liu et al. (2017) in the ERM setting. We convert their notations into ours to be consistent.

H.1 Iteration complexity of the algorithm of Liu et al. (2017)

The full algorithm of Liu et al. (2017) performs alternating optimization over the task relationship and the local predictors on each machine. In order to compare their algorithm with ours on the efficiency of learning predictors, we consider a fixed task correlation matrix \mathbf{M} [REDACTED] \mathbf{L} in their objective (corresponding to Ω in eqn (1) of their paper).

With fixed \mathbf{M} , their algorithm performs distributed SDCA (Ma et al., 2015) for optimizing over the predictors. In each round of distributed SDCA, one constructs an upper bound of the objective that is separable over the machines (predictors), so that each machine solves a subproblem defined by its local data, and then one around of communication is used to aggregate local updates.

When the instantaneous losses are β_F -smooth and each local subproblem is solved exactly (i.e., we set $\Theta = 0$ in their analysis), the number of global (communication) rounds needed for obtaining an approximate solution is, according to Liu et al. (2017, Lemma 7 and Theorem 8), of the order (ignoring the logarithmic factor on final optimization error)

$$[REDACTED]$$

Here, the first term measures the “task separability” with value in $[1, m]$ (see the definitions of \mathbf{K} and $\alpha_{[i]}$ in their Theorem 1, and the discussion of separability in Section 6.3). On the other hand, we have $\max [REDACTED]$ 1. As a result, the iteration complexity of distributed SDCA is

(task separability in $[1, m]$).

This iteration complexity is similar to that of our ERM algorithm by directly solving the regularizer (q), but has worse dependence on the condition number and an unclear multiplicative constant on the tasks separability.

H.2 Comparison with the collaborative algorithm of Vanhaesebrouck et al. (2017)

We now compare with the collaborative learning algorithm of Vanhaesebrouck et al. (2017) in the synchronous and decentralized setting. In their algorithm, each machine augments its local optimization parameters to include a copy of predictor from each neighboring machine. Let Θ_i be the set of $|N_i| + 1$ variables \mathbf{w}_k for $k \in N_i \cup \{i\}$, and Θ_i^k is the copy of \mathbf{w}_k on machine i . We can reformulate the global objective (2) as

$$\begin{aligned} & \min_{\Theta} \sum_{i=1}^m H_i(\Theta_i) \\ & \text{where } \Theta = \bigcup_{i=1}^m \Theta_i \\ & \text{subject to } \Theta_i \cap \Theta_j = \emptyset \end{aligned} \quad (22)$$

Vanhaesebrouck et al. (2017) then introduce variables associated with each edge (4 set of variables per edge) and apply ADMM to the resulting problem. An advantage of ADMM is that it allows decoupling of the local problems when updating primal variables, where the local problem involves the nonlinearized loss function.

Although Vanhaesebrouck et al. (2017) suggest that the convergence results of synchronous decentralized ADMM (Wei and Ozdaglar, 2013; Shi et al., 2014) apply to this formulation (see their Appendix D), we note however that (22) is not in the standard form covered by these results. In particular, the classical decentralized consensus problem has the form $\mathbf{x} \in \mathbb{N}$

$$\mathbf{x}_i = \mathbf{x}_j \quad \text{for all } (i,j) \text{ where } j \in N_i$$

Here, neighboring machines share the same set of optimization parameters and they would like to reach complete consensus, whereas in (22) neighboring machines can have different set of variables and they only try to achieve consensus on the shared parameters. As a result, it is nontrivial to derive the iteration complexity of the collaborative learning algorithm of Vanhaesebrouck et al. (2017) based on the same quantities used in the analysis of our algorithms.

I Experiments

In this section we examine the empirical performance of the proposed algorithms. We consider the problem of linear regression on synthetic data. For the i -th task, we generate data from

where ξ_i is noise drawn from the Normal distribution $N(0,3)$, $\mathbf{x} \in \mathbb{R}^d$ is drawn from a multivariate Normal distribution with mean zero and covariance matrix Σ where $\Sigma_{ij} = 2^{-|i-j|/3}$, and \mathbf{w}_{i^*} is a coefficient vector for the i -th task generated from the following clustered multi-task structure. Each \mathbf{w}_{i^*} is drawn from a mixture of C clusters; there is a reference model \mathbf{r}_j for each cluster $j = 1, \dots, C$, and the task specific model \mathbf{w}_{i^*} is a small perturbation of the corresponding cluster reference model: $\mathbf{w}_{i^*} = \mathbf{r}_j + \xi_i$, if \mathbf{w}_{i^*} is drawn from cluster j .

The cluster reference model \mathbf{r}_j is generated by sampling each entry i.i.d. from $Unif[-0.5,0.5]$, while the perturbation vector ξ_i is generated by sampling each entry i.i.d. from $Unif[-0.05,0.05]$. This construction gives us task specific models which are similar to each other when they belong to the same cluster. The corresponding similarity graph is a 10-nearest neighbor graph (so the graph is connected) with binary weights built on $\{\mathbf{w}_{ij}\}_{i=1, \dots, m}$, i.e., each task is connected to 10 other tasks whose models are most similar.

We tested a few graph-based multi-task learning methods.

- Local: solves a local ERM problem (with only λ_2 regularization) with n samples for each task.
- Centralized: solves the graph-regularized ERM problem (2) with n samples for each task.
- ADMM: the synchronized version of the ADMM algorithm of Vanhaesebrouck et al. (2017).
- SDCA: the distributed SDCA algorithm of Liu et al. (2017) for fixed graph.
- Our algorithms: denoted as B/S (batch/stochastic) + SR/OL (solve regularizer/optimize loss).




Figure 2: Performance of different methods for regularized empirical risk minimization.

In the experiments below, we have problem dimension $d = 100$, number of tasks $m = 100$, training set size $n = 500$, and vary number of task clusters C over $\{1,5,10,50\}$ (smaller C implies overall stronger task similarity). We also generate a dev set of 10000 samples per task for tuning hyper-parameters, and test set of 10000 samples per task for approximately evaluating the population loss.

Empirical risk minimization We first compare the iterative methods on the regularized ERM problem (2), to which the analysis for ADMM and SDCA applies. We tune the ℓ_2 regularization parameter for Local and (η, τ) for Centralized, and then fix the optimal (η, τ) for other methods. We also tune the quadratic penalty parameter for ADMM, the task separability and stepsize parameters for SDCA, and stepsize parameter for BSR/BOL (although the default value based on the smoothness parameter already works well for them). For SSR/SOL, we draw random samples from the fixed training set (with size n), and simply fix the minibatch size to be $n/10$.

Figure 2 shows for each method the estimated $F(\mathbf{W})$ over iterations (or rounds of communication) in the top row, and over the amount of computation (measured by the number of passes over the training set) in the bottom row. Observe that all iterative algorithms converge to the same ERM solution, our algorithms tend to consistently outperform ADMM and SDCA.

Stochastic optimization We next demonstrate the efficiency of true stochastic algorithms (using fresh samples for each update) at $C = 10$. We allow the algorithms to process a total of 10000 fresh samples on each machine, and vary the minibatch size b over $\{40,80,100,200,500\}$. The parameters (η, τ) are fixed to those used in the ERM experiments.

Figure 3 shows for each method the estimated $F(\mathbf{W})$ over iterations (or rounds of communication) in the left plot, and over the amount of fresh samples processed (or total computation cost) in the right plot. As a reference, the error of Local and Centralized (using $n = 500$ samples per machine) are also given in the plots. We observe that with fresh samples, stochastic algorithms are

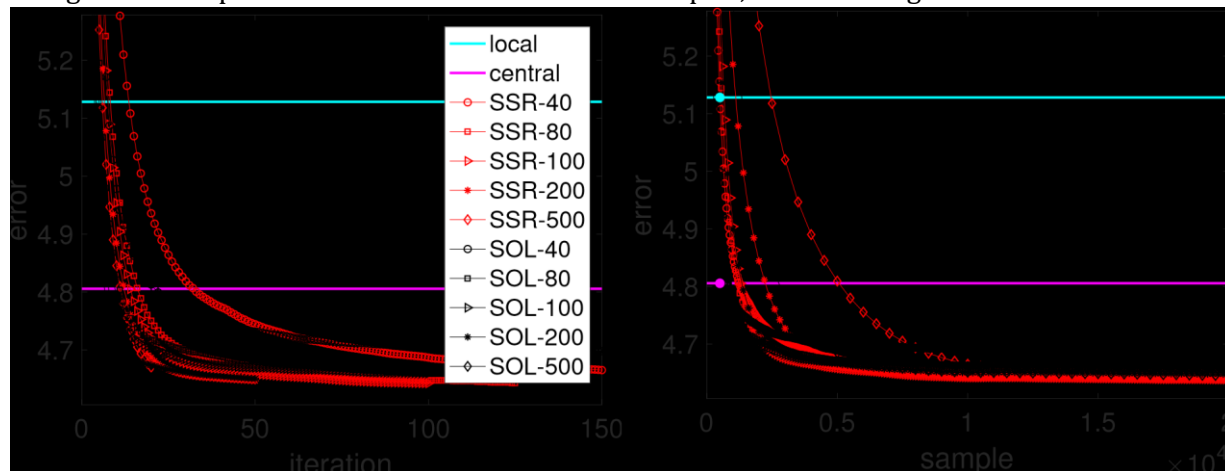


Figure 3: Performance of stochastic algorithms with various minibatch sizes. Here $C = 10$.

competitive to ERM algorithms in terms of sample complexity, while being computationally more efficient.

Acknowledgements

This work was supported by NSF-BSF award 1718970.

References

- Y. Amit, M. Fink, N. Srebro, and S. Ullman. Uncovering shared structures in multiclass classification. In *Proceedings of the 24th international conference on Machine learning*, pages 17–24. ACM, 2007.
- R. K. Ando and T. Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data. *J. Mach. Learn. Res.*, 6:1817–1853, 2005.
- A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Mach. Learn.*, 73(3):243–272, 2008.
- M.-F. Balcan, A. Blum, S. Fine, and Y. Mansour. Distributed learning, communication complexity and privacy. In S. Mannor, N. Srebro, and R. Williamson, editors, *JMLR W&CP 23: COLT 2012*, volume 23, pages 26.1–26.22, 2012.
- I. M. Baytas, M. Yan, A. K. Jain, and J. Zhou. Asynchronous multi-task learning. In *IEEE International Conference on Data Mining (ICDM)*, 2016.
- S. P. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1): 1–122, 2011.
- O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13:165–202, 2012.
- T. Evgeniou and M. Pontil. Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 109–117. ACM, 2004.
- T. Evgeniou, C. A. Micchelli, and M. Pontil. Learning multiple tasks with kernel methods. In *J. Mach. Learn. Res.*, pages 615–637, 2005.
- H. R. Feyzmahdavian, A. Aytakin, and M. Johansson. A delayed proximal gradient method with linear convergence rate. In *2014 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2014.
- R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems (NIPS)*, pages 315–323, 2013.
- J. Konecny, B. McMahan, and D. Ramage. Federated optimization: Distributed optimization beyond the datacenter. arXiv:1511.03575 [cs.LG], 2015.
- G. Lan. An optimal method for stochastic composite optimization. *Math. Prog.*, 133(1–2):365–397, 2012.
- S. Liu, S. J. Pan, and Q. Ho. Distributed multi-task relationship learning. arXiv:1612.04022 [cs.LG], 2017.

- K. Lounici, M. Pontil, A. B. Tsybakov, and S. A. van de Geer. Oracle inequalities and optimal inference under group sparsity. *Ann. Stat.*, 39:2164–204, 2011.
- C. Ma, V. Smith, M. Jaggi, M. I. Jordan, P. Richtarik, and M. Takac. Adding vs. averaging in distributed primal-dual optimization. In *Proc. of the 32st (ICML 2015)*, 2015.
- A. Maurer. The Rademacher complexity of linear transformation classes. In G. Lugosi and H.-U. Simon, editors, *Annual Conference on Learning Theory*, pages 65–78, 2006.
- A. Nedić and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Trans. Automat. Contr.*, 54(1):48–61, 2009.
- Y. Nesterov. *Introductory Lectures on Convex Optimization. A Basic Course*. Number 87. SpringerVerlag, 2004.
- G. Obozinski, M. J. Wainwright, and M. I. Jordan. Support union recovery in high-dimensional multivariate regression. *Ann. Stat.*, 39(1):1–47, 2011.
- S. S. Ram, A. Nedić, and V. V. Veeravalli. Distributed stochastic subgradient projection algorithms for convex optimization. *Journal of optimization theory and applications*, 147(3):516–545, 2010.
- K. Scaman, F. Bach, S. Bubeck, Y. T. Lee, and L. Massoulié. Optimal algorithms for smooth and strongly convex distributed optimization in networks. arXiv:1702.08704 [math.OC], 2017.
- M. Schmidt, N. L. Roux, and F. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. pages 1458–1466, 2011.
- S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Stochastic convex optimization. In S. Dasgupta and A. Klivans, editors, *Proc. of the 22th Annual Conference on Learning Theory (COLT'09)*, Montreal, Quebec, 2009.
- O. Shamir and N. Srebro. Distributed stochastic optimization and learning. In *52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2014*, pages 850–857. IEEE, 2014.
- W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin. On the linear convergence of the ADMM in decentralized consensus optimization. *IEEE Trans. Signal Processing*, 62(7):1750–1761, 2014.
- B. A. Turlach, W. N. Venables, and S. J. Wright. Simultaneous variable selection. *Technometrics*, 47(3):349–363, 2005.
- P. Vanhaesebrouck, A. Bellet, and M. Tommasi. Decentralized collaborative learning of personalized models over networks. In *Int. Workshop on Artificial Intelligence and Statistics*, pages 509–517, 2017.
- J. Wang, M. Kolar, and N. Srebro. Distributed multitask learning. *ArXiv e-prints*, arXiv:1510.00633, 2015, arXiv:1510.00633.
- J. Wang, M. Kolar, and N. Srebro. Distributed multi-task learning with shared representation. 2016, arXiv:1603.02185.

- J. Wang, W. Wang, and N. Srebro. Memory and communication efficient distributed stochastic optimization with minibatch prox. In S. Kale and O. Shamir, editors, *Annual Conference on Learning Theory*, Amsterdam, Netherlands, 2017.
- E. Wei and A. Ozdaglar. On the $o(1/k)$ convergence of asynchronous distributed alternating direction method of multipliers. arXiv:1307.8254 [math.OC], 2013.
- K. Yuan, Q. Ling, and W. Yin. On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3):1835–1854, 2016.
- M. Yuan, A. Ekici, Z. Lu, and R. Monteiro. Dimension reduction and coefficient estimation in multivariate linear regression. *J. R. Stat. Soc. B*, 69(3):329–346, 2007.