

Distributed Self-Paced Learning in Alternating Direction Method of Multipliers

Xuchao Zhang¹, Liang Zhao², Zhiqian Chen¹, Chang-Tien Lu¹

¹Discovery Analytics Center, Virginia Tech, Falls Church, VA, USA

²George Mason University, Fairfax, VA, USA

¹{xuczhang, czq, ctlu}@vt.edu, ²lzhao9@gmu.edu

Abstract

Self-paced learning (*SPL*) mimics the cognitive process of humans, who generally learn from easy samples to hard ones. One key issue in *SPL* is the training process required for each instance weight depends on the other samples and thus cannot easily be run in a distributed manner in a large-scale dataset. In this paper, we reformulate the self-paced learning problem into a distributed setting and propose a novel Distributed Self-Paced Learning method (*DSPL*) to handle large scale datasets. Specifically, both the model and instance weights can be optimized in parallel for each batch based on a consensus alternating direction method of multipliers. We also prove the convergence of our algorithm under mild conditions. Extensive experiments on both synthetic and real datasets demonstrate that our approach is superior to those of existing methods.

1 Introduction

Inspired by the learning processes used by humans and animals [Bengio *et al.*, 2009], Self-Paced Learning (*SPL*) [Kumar *et al.*, 2010] considers training data in a meaningful order, from easy to hard, to facilitate the learning process. Unlike standard curriculum learning [Bengio *et al.*, 2009], which learns the data in a predefined curriculum design based on prior knowledge, *SPL* learns the training data in an order that is dynamically determined by feedback from the individual learner, which means it can be more extensively utilized in practice. In self-paced learning, given a set of training samples along with their labels, a parameter λ is used to represent the “age” of the *SPL* in order to control the learning pace. When λ is small, “easy” samples with small losses are considered. As λ grows, “harder” samples with larger losses are gradually added to the training set. This type of learning process is modeled on the way human education and cognition functions. For instance, students will start by learning easier concepts (e.g., Linear Equations) before moving on to more complex ones (e.g., Differential Equations) in the mathematics curriculum. Self-paced learning can also be finely explained in a robust learning manner, where uncorrupted data samples are likely to be used for training earlier in the process than corrupted data.

In recent years, self-paced learning [Kumar *et al.*, 2010] has received widespread attention for various applications in machine learning, such as image classification [Jiang *et al.*,

2015], event detection [Jiang *et al.*, 2014a; Zhang *et al.*, 2017c] and object tracking [Supancic and Ramanan, 2013; Zhang *et al.*, 2016]. A wide assortment of *SPL*-based methods [Pi *et al.*, 2016; Ma *et al.*, 2017a] have been developed, including self-paced curriculum learning [Jiang *et al.*, 2015], self-paced learning with diversity [Jiang *et al.*, 2014b], multi-view [Xu *et al.*, 2015] and multi-task [Li *et al.*, 2017; Keerthiram Murugesan, 2017] self-paced learning. In addition, several researchers have conducted theoretical analyses of self-paced learning. [Meng *et al.*, 2015] provides a theoretical analysis of the robustness of *SPL*, revealing that the implicit objective function of *SPL* has a similar configuration to a non-convex regularized penalty. Such regularization restricts the contributions of noisy examples to the objective, and thus enhances the learning robustness. [Ma *et al.*, 2017b] proved that the learning process of *SPL* always converges to critical points of its implicit objective under mild conditions, while [Fan *et al.*, 2017] studied a group of new regularizers, named self-paced implicit regularizers that are derived from convex conjugacy.

Existing self-paced learning approaches typically focus on modeling the entire dataset at once; however, this may introduce a bottleneck in terms of memory and computation, as today’s fast-growing datasets are becoming too large to be handled integrally. For those seeking to address this issue, the main challenges can be summarized as follows: 1) *Computational infeasibility of handling the entire dataset at once*. Traditional self-paced learning approaches gradually grow the training set according to their learning pace. However, this strategy fails when the training set grows too large to be handled due to the limited capacity of the physical machines. Therefore, a scalable algorithm is required to extend the existing self-paced learning algorithm for massive datasets. 2) *Existence of heterogeneously distributed “easy” data*. Due to the unpredictability of data distributions, the “easy” data samples can be arbitrarily distributed across the whole dataset. Considering the entire dataset as a combination of multiple batches, some batches may contain large amount of “hard” samples. Thus, simply applying self-paced learning to each batch and averaging across the trained models is not an ideal approach, as some models will only focus on the “hard” samples and thus degrade the overall performance. 3) *Dependency decoupling across different data batches*. In self-paced learning, the instance weights depend on the model trained on the entire dataset. Also, the trained model depends on the weights assigned to each data instance. If we consider each data batch independently, a model trained in a “hard” batch can mistakenly mark some “hard” samples

as “easy” ones. For example, in robust learning, the corrupted data samples are typically considered as “hard” samples, then more corrupted samples will therefore tend to be involved into the training process when we train data batches independently.

In order to simultaneously address all these technical challenges, this paper presents a novel Distributed Self-Paced Learning (*DSPL*) algorithm. The main contributions of this paper can be summarized as follows: 1) We reformulate the self-paced problem into a distributed setting. Specifically, an auxiliary variable is introduced to decouple the dependency of the model parameters for each data batch. 2) A distributed self-paced learning algorithm based on consensus ADMM is proposed to solve the *SPL* problem in a distributed setting. The algorithm optimizes the model parameters for each batch in parallel and consolidates their values in each iteration. 3) A theoretical analysis is provided for the convergence of our proposed *DSPL* algorithm. The proof shows that our new algorithm will converge under mild assumptions, e.g., the loss function can be non-convex. 4) Extensive experiments have been conducted utilizing both synthetic and real-world data based on a robust regression task. The results demonstrate that the proposed approaches consistently outperform existing methods for multiple data settings. To the best of our knowledge, this is the first work to extend self-paced learning to a distributed setting, making it possible to handle large-scale datasets.

The remainder of this paper is organized as follows. Section 2 gives a formal problem formulation. The proposed distributed self-paced learning algorithm is presented in Section 3 and Section 4 presents a theoretical analysis of the convergence of the proposed method. In Section 5, the experimental results are analyzed and the paper concludes with a summary of our work in Section 6.

2 Problem Formulation

In the context of distributed self-paced learning, we consider the samples to be provided in a sequence of mini batches as $\{(X^{(1)}, \mathbf{y}^{(1)}), \dots, (X^{(m)}, \mathbf{y}^{(m)})\}$, where $X^{(i)} \in \mathbb{R}^{p \times n_i}$ represents the sample data for the i^{th} batch, $\mathbf{y}^{(i)}$ is the corresponding response vector, and n_i is the instance number of the i^{th} batch.

The goal of self-paced learning problem is to infer the model parameter $\mathbf{w} \in \mathbb{R}^p$ for the entire dataset, which can be formally defined as follows:

$$\begin{aligned} \arg \min_{\mathbf{w}, \mathbf{v}} \sum_{i=1}^m f_i(\mathbf{w}, \mathbf{v}_i) + \|\mathbf{w}\|_2^2 \\ \text{s.t. } v_{ij} \in [0, 1], \quad \forall i = 1, \dots, m, \forall j = 1, \dots, n_i \end{aligned} \quad (1)$$

where $\|\mathbf{w}\|_2^2$ is the regularization term for model parameters \mathbf{w} . Variable \mathbf{v}_i represents the instance weight vector for the i^{th} batch and v_{ij} is the weight of the j^{th} instance in the i^{th} batch. The objective function $f_i(\mathbf{w}, \mathbf{v}_i)$ for each mini-batch is defined as follows:

$$f_i(\mathbf{w}, \mathbf{v}_i) = \sum_{j=1}^{n_i} v_{ij} \mathcal{L}(y_{ij}, g(\mathbf{w}, \mathbf{x}_{ij})) - \lambda \sum_{j=1}^{n_i} v_{ij} \quad (2)$$

We denote $\mathbf{x}_{ij} \in \mathbb{R}^p$ and $y_{ij} \in \mathbb{R}$ as the feature vector and its corresponding label for the j^{th} instance in the i^{th} batch. The loss function \mathcal{L} is used to measure the error between label y_{ij} and

Table 1: Mathematical Notations

Notations	Explanations
p	feature number in data matrix $X^{(i)}$
n_i	instance number in the i^{th} data batch
$X^{(i)}$	data matrix of the i^{th} batch
$\mathbf{y}^{(i)}$	the response vector of the i^{th} batch
\mathbf{w}	model parameter of the entire dataset
\mathbf{v}_i	instance weight vector of the i^{th} batch
v_{ij}	weight of the j^{th} instance in the i^{th} batch
λ	parameter to control the learning pace
\mathcal{L}	loss function of estimated model

the estimated value from model g . The term $-\lambda \sum_{j=1}^{n_i} v_{ij}$ is the regularization term for instance weights \mathbf{v}_i , where parameter λ controls the learning pace. The notations used in this paper are summarized in Table 1.

The problem defined above is very challenging in the following three aspects. First, data instances for all m batches can be too large to be handled simultaneously in their entirety, thus requiring the use of a scalable algorithm for large datasets. Second, the instance weight variable \mathbf{v}_i of each batch depends on the optimization result for \mathbf{w} shared by all the data, which means all the batches are inter-dependent and it is thus not feasible to run them in parallel. Third, the objective function of variables \mathbf{w}_i and \mathbf{v}_i are jointly non-convex and it is an NP-hard problem to retrieve the global optimal solution [Gorski *et al.*, 2007]. In next section, we present a distributed self-paced learning algorithm based on consensus ADMM to address all these challenges.

3 The Proposed Methodology

In this section, we propose a distributed self-paced learning algorithm based on the alternating direction method of multipliers (ADMM) to solve the problem defined in Section 2.

The problem defined in Equation (1) cannot be solved in parallel because the model parameter \mathbf{w} is shared in each batch and the result of \mathbf{w} will impact on the instance weight variable \mathbf{v}_i for each batch. In order to decouple the relationships among all the batches, we introduce different model parameters \mathbf{w}_i for each batch and use an auxiliary variable \mathbf{z} to ensure the uniformity of all the model parameters. The problem can now be reformulated as follows:

$$\begin{aligned} \arg \min_{\mathbf{w}_i, \mathbf{v}_i, \mathbf{z}} \sum_{i=1}^m f_i(\mathbf{w}_i, \mathbf{v}_i; \lambda) + \|\mathbf{z}\|_2^2 \\ \text{s.t. } v_{ij} \in [0, 1], \quad \forall i = 1, \dots, m, \forall j = 1, \dots, n_i \\ \mathbf{w}_i - \mathbf{z} = 0, \quad \forall i = 1, \dots, m \end{aligned} \quad (3)$$

where the function $f_i(\mathbf{w}_i, \mathbf{v}_i)$ is defined as follows.

$$f_i(\mathbf{w}_i, \mathbf{v}_i; \lambda) = \sum_{j=1}^{n_i} v_{ij} \mathcal{L}(y_{ij}, g(\mathbf{w}_i, \mathbf{x}_{ij})) - \lambda \sum_{j=1}^{n_i} v_{ij} \quad (4)$$

Unlike the original problem defined in Equation (1), here each batch has its own model parameter \mathbf{w}_i and the constraint $\mathbf{w}_i - \mathbf{z} = 0$ for $\forall i = 1, \dots, m$ ensures the model parameter \mathbf{w}_i has the same value as the auxiliary variable \mathbf{z} . The purpose of the problem reformulation is to optimize the model parameters

w_i in parallel for each batch. It is important to note that the reformulation is *tight* because our new problem is equivalent to the original problem when the constraint is strictly satisfied.

In the new problem, Equation (3) is a bi-convex optimization problem over v_i and w_i for each batch with fixed z , which can be efficiently solved using the Alternate Convex Search (ACS) method [Gorski *et al.*, 2007]. With the variable v fixed, the remaining variables $\{w_i\}$, z and α can be solved by consensus ADMM [Boyd *et al.*, 2011]. As the problem is an NP-hard problem, in which the global optimum requires polynomial time complexity, we propose an alternating algorithm *DSPL* based on ADMM to handle the problem efficiently.

The augmented Lagrangian format of optimization in Equation (3) can be represented as follows:

$$L = \sum_{i=1}^m f_i(w_i, v_i; \lambda) + \|z\|_2^2 + \sum_{i=1}^m \alpha_i^T (w_i - z) + \frac{\rho}{2} \sum_{i=1}^m \|w_i - z\|_2^2 \quad (5)$$

where $\{\alpha_i\}_{i=1}^m$ are the Lagrangian multipliers and ρ is the step size of the dual step.

The process used to update model parameter w_i for the i^{th} batch with the other variables fixed is as follows:

$$w_i^{k+1} = \arg \min_{w_i} f_i(w_i, v_i; \lambda) + [\alpha_i^k]^T (w_i - z^k) + \frac{\rho}{2} \|w_i - z^k\|_2^2 \quad (6)$$

Specifically, if we choose the loss function \mathcal{L} to be a squared loss and model $g(w, x_{ij})$ to be a linear regression $g(w, x_{ij}) = w^T x_{ij}$, we have the following analytical solution for w_i :

$$w_i^{k+1} = \left(2 \sum_{j=1}^{n_i} v_{ij} x_{ij} x_{ij}^T + \rho \cdot I \right)^{-1} \cdot \left(2 \sum_{j=1}^{n_i} v_{ij} x_{ij} y_{ij} - \alpha_i^k + \rho z^k \right) \quad (7)$$

The auxiliary variable z and Lagrangian multipliers α_i can be updated as follows:

$$z^{k+1} = \frac{\rho}{2 + \rho m} \sum_{i=1}^m (w_i^{k+1} + \frac{1}{\rho} \alpha_i^k) \quad (8)$$

$$\alpha_i^{k+1} = \alpha_i^k + \rho (w_i^{k+1} - z^{k+1})$$

The stop condition of consensus ADMM is determined by the (squared) norm of the primal and dual residuals of the k^{th} iteration, which are calculated as follows:

$$\|r_k\|_2^2 = \sum_{i=1}^m \|w_i^k - z^k\|_2^2 \quad (9)$$

$$\|s_k\|_2^2 = m\rho^2 \|z^k - z^{k-1}\|_2^2$$

After the weight parameter w_i for each batch has been updated, the instance weight vector v_i for each batch will be updated based on the fixed w_i by solving the following problem:

$$v_i^{t+1} = \arg \min_{v_i} \sum_{j=1}^{n_i} v_{ij} \mathcal{L}(y_{ij}, g(w_i^{t+1}, x_{ij})) - \lambda \sum_{j=1}^{n_i} v_{ij} \quad (10)$$

Algorithm 1: DSPL ALGORITHM

Input: $X \in \mathbb{R}^{p \times n}$, $y \in \mathbb{R}^n$, $\lambda_0 \in \mathbb{R}$, $\tau_\lambda \in \mathbb{R}$, $\mu \in \mathbb{R}$
Output: solution $w^{(t+1)}$, $v^{(t+1)}$

- 1 Initialize $w_i^0 = \mathbf{1}$, $v_i^0 = \mathbf{1}$
- 2 Choose $\varepsilon_L > 0$, $\varepsilon_r > 0$, $\varepsilon_s > 0$, $\lambda \leftarrow \lambda_0$, $t \leftarrow 0$
- 3 **repeat**
- 4 $k \leftarrow 0$
- 5 **repeat**
- 6 $z^{k+1} \leftarrow \frac{1}{m} \sum_{i=1}^m (w_i^{k+1} + \frac{1}{\rho} \alpha_i^k)$
- 7 Update variables w_i^{k+1} in parallel, for $\forall i = 1 \dots m$
- 8 $w_i^{k+1} \leftarrow \arg \min f_i(w_i, v_i) + [\alpha_i^k]^T (w_i - z^k) + \frac{\rho}{2} \|w_i - z^k\|_2^2$
- 9 Update dual $\alpha_i^{k+1} \leftarrow \alpha_i^k + \rho (w_i^{k+1} - z^{k+1})$ in parallel
- 10 Update primal and dual residuals r^{k+1} and s^{k+1} .
- 11 $k \leftarrow k + 1$
- 12 **until** $\|r^{k+1}\|_2^2 < \varepsilon_r$ and $\|s^{k+1}\|_2^2 < \varepsilon_s$
- 13 $v_i^{t+1} \leftarrow \mathbb{1} \left(\mathcal{L}(y_{ij}, g(w_i^{t+1}, x_{ij})) < \lambda \right)$, for $\forall i = 1 \dots m$
- 14 **if** $\lambda < \tau_\lambda$ **then**
- 15 $\lambda \leftarrow \lambda * \mu$
- 16 **else**
- 17 $\lambda \leftarrow \tau_\lambda$
- 18 $t \leftarrow t + 1$
- 19 **until** $\|L^{t+1} - L^t\|_2 < \varepsilon_L$
- 20 **return** z^{t+1} , v^{t+1}

For the above problem in Equation (10), we always obtain the following closed-form solution:

$$v_i^{t+1} = \mathbb{1} \left(\mathcal{L}(y_{ij}, g(w_i^{t+1}, x_{ij})) < \lambda \right) \quad (11)$$

where $\mathbb{1}(\cdot)$ is the indicator function whose value equals to one when the condition $\mathcal{L}(y_{ij}, g(w_i^{t+1}, x_{ij})) < \lambda$ is satisfied; otherwise, its value is zero.

The details of algorithm *DSPL* are presented in Algorithm 1. In Lines 1-2, the variables and parameters are initialized. With the variables v_i fixed, the other variables are optimized in Lines 5-13 based on the result of consensus ADMM, in which both the model weights w_i and Lagrangian multipliers α_i can be updated in parallel for each batch. Note that if no closed-form can be found for Equation (6), the updating of w_i is the most time-consuming operation in the algorithm. Therefore, updating w_i in parallel can significantly improve the efficiency of the algorithm. The variable v_i for each batch is updated in Line 14, with the variable w_i fixed. In Lines 15-18, the parameter λ is enlarged to include more data instances into the training set. τ_λ is the maximum threshold for λ and μ is the step size. The algorithm will be stopped when the Lagrangian is converged in Line 20. The following two alternative methods can be applied to improve the efficiency of the algorithm: 1) dynamically update the penalty parameter ρ after Line 11. When $r > 10s$, we can update $\rho \leftarrow 2\rho$. When $10r < s$, we have $\rho \leftarrow \rho/2$. 2) Move the update of variable v_i into the consensus ADMM step after Line 9. This ensures that the instance weights are updated every time the model is updated, so that the algorithm quickly converges. However, no theoretical convergence guarantee can be made for the two solutions, although in practice they do always converge.

4 Theoretical Analysis

In this section, we will prove the convergence of the proposed algorithm. Before we start to prove the convergence of Algorithm 1, we make the following assumptions regarding our objective function and penalty parameter ρ :

Assumption 1 (Gradient Lipchitz Continuity). *There exists a positive constant φ_i for objective function $f_i(\mathbf{w}_i)$ of each batch with the following properties:*

$$\|\nabla_{\mathbf{w}_i} f_i(\mathbf{x}_i) - \nabla_{\mathbf{w}_i} f_i(\mathbf{y}_i)\| \leq \varphi_i \|\mathbf{x}_i - \mathbf{y}_i\|, \quad (12)$$

$$\forall \mathbf{x}_i, \mathbf{y}_i, i = 1, \dots, m$$

Assumption 2 (Lower Bound). *The objective function in problem (3) has a lower bound \mathcal{B} as follows:*

$$\mathcal{B} = \min_{\mathbf{w}_i, \mathbf{v}_i, \mathbf{z}} \left\{ \sum_{i=1}^m f_i(\mathbf{w}_i, \mathbf{v}_i) + \|\mathbf{z}\|_2^2 \right\} > -\infty \quad (13)$$

Assumption 3 (Penalty Parameter Constraints). *For $\forall i = 1 \dots m$, the penalty parameter ρ_i for each batch is chosen in accord with the following constraints:*

- For $\forall i$, the subproblem (6) of variable \mathbf{w}_i is strongly convex with modulus $\gamma_i(\rho_i)$.
- For $\forall i$, we have $\rho_i \gamma_i(\rho_i) > 2\varphi_i^2$ and $\rho_i \geq \varphi_i$.

Note that when ρ_i increases, subproblem (6) will be eventually become strongly convex with respect to variable \mathbf{w}_i . For simplicity, we will choose the same penalty parameter ρ for all the batches with $\rho = \max_i(\rho_i)$. Based on these assumptions, we can draw the following conclusions.

Lemma 1. *Assume the augmented Lagrangian format of optimization problem (3) satisfies Assumption 1, the augmented Lagrangian L has the following property:*

$$L(\{\mathbf{w}_i^{k+1}\}, \mathbf{z}^{k+1}, \boldsymbol{\alpha}^{k+1}) \leq L(\{\mathbf{w}_i^k\}, \mathbf{z}^k, \boldsymbol{\alpha}^k) \quad (14)$$

Proof. Since the the objective function $f_i(\mathbf{w}_i)$ for each batch is gradient Lipchitz continuous with a positive constant φ_i , the Lagrangian in Equation (5) has the following property according to Lemma 2.2 in [Hong *et al.*, 2016]:

$$\begin{aligned} & L(\{\mathbf{w}_i^{k+1}\}, \mathbf{z}^{k+1}, \boldsymbol{\alpha}^{k+1}) - L(\{\mathbf{w}_i^k\}, \mathbf{z}^k, \boldsymbol{\alpha}^k) \\ & \leq \sum_{i=1}^m \left(\frac{\varphi_i^2}{\rho} - \frac{\gamma_i(\rho)}{2} \right) \|\mathbf{w}_i^{k+1} - \mathbf{w}_i^k\|_2^2 - \frac{\gamma}{2} \|\mathbf{z}^{k+1} - \mathbf{z}^k\|_2^2 \\ & \stackrel{(a)}{\leq} -\frac{\gamma}{2} \|\mathbf{z}^{k+1} - \mathbf{z}^k\|_2^2 \leq 0 \end{aligned} \quad (15)$$

where $\gamma = m\rho > 0$. The inequality (a) follows from Assumption 2, namely that $\rho\gamma_i(\rho) > 2\varphi_i^2$, so we have $\left(\frac{\varphi_i^2}{\rho} - \frac{\gamma_i(\rho)}{2} \right) < 0$. \square

Lemma 2. *Assume the augmented Lagrangian of problem (3) satisfies Assumptions 1-3, the augmented Lagrangian L is lower bounded as follows:*

$$\lim_{k \rightarrow \infty} L(\{\mathbf{w}_i^{k+1}\}, \mathbf{z}^{k+1}, \boldsymbol{\alpha}^{k+1}) \geq \mathcal{B} \quad (16)$$

where \mathcal{B} is the lower bound of the objective function in problem (3).

Proof. Due to the gradient Lipchitz continuity assumption, we have the following optimality condition for the \mathbf{w}_i update step in Equation (6):

$$\nabla_{\mathbf{w}_i} f_i(\mathbf{w}_i^{k+1}) + \boldsymbol{\alpha}_i^k + \rho(\mathbf{w}_i^{k+1} - \mathbf{z}^{k+1}) = 0, \quad \forall i = 1 \dots m \quad (17)$$

Combined with the update of the Lagrangian multipliers $\boldsymbol{\alpha}_i$ in Equation (8), we have

$$\nabla_{\mathbf{w}_i} f_i(\mathbf{w}_i^{k+1}) = -\boldsymbol{\alpha}_i^{k+1}, \quad \forall i = 1 \dots m \quad (18)$$

The augmented Lagrangian can be represented as:

$$\begin{aligned} & L(\{\mathbf{w}_i^{k+1}\}, \mathbf{z}^{k+1}, \boldsymbol{\alpha}^{k+1}) \\ & \stackrel{(a)}{=} \|\mathbf{z}^{k+1}\|_2^2 + \sum_{i=1}^m \left(f_i(\mathbf{w}_i^{k+1}) + \nabla_{\mathbf{w}_i} f_i(\mathbf{w}_i^{k+1}) \cdot \right. \\ & \quad \left. (\mathbf{z}^{k+1} - \mathbf{w}_i^{k+1}) + \frac{\rho}{2} \|\mathbf{w}_i^{k+1} - \mathbf{z}^{k+1}\|_2^2 \right) \\ & \stackrel{(b)}{\geq} \|\mathbf{z}^{k+1}\|_2^2 + \sum_{i=1}^m f_i(\mathbf{z}^{k+1}) \stackrel{(c)}{\geq} \mathcal{B} \end{aligned} \quad (19)$$

Equation (a) follows from Equation (19) and the inequality (b) comes from the Lipschitz continuity property in Assumption 1. The inequality (c) follows from the lower bound property in Assumption 2. \square

Theorem 3. *The Algorithm 1 converges when Assumptions 1-3 are all satisfied.*

Proof. In Lemmas 1 and 2, we proved that the Lagrangian is monotonically decreasing and has a lower bound during the iterations of ADMM. Now we will prove that the same properties hold for the entire algorithm after updating variable \mathbf{v} and parameter λ .

$$\begin{aligned} & L(\{\mathbf{w}^{t+1}\}, \mathbf{v}^{t+1}, \mathbf{z}^{t+1}, \boldsymbol{\alpha}^{t+1}; \lambda^{t+1}) \\ & \stackrel{(a)}{\leq} L(\{\mathbf{w}^t\}, \mathbf{v}^{t+1}, \mathbf{z}^t, \boldsymbol{\alpha}^t; \lambda^{t+1}) \stackrel{(b)}{\leq} L(\{\mathbf{w}^t\}, \mathbf{v}^t, \mathbf{z}^t, \boldsymbol{\alpha}^t; \lambda^{t+1}) \\ & = L(\{\mathbf{w}^t\}, \mathbf{v}^t, \mathbf{z}^t, \boldsymbol{\alpha}^t; \lambda^t) + (\lambda^t - \lambda^{t+1}) \sum_{i=1}^m \sum_{j=1}^{n_i} v_{ij}^t \\ & \stackrel{(c)}{\leq} L(\{\mathbf{w}^t\}, \mathbf{v}^t, \mathbf{z}^t, \boldsymbol{\alpha}^t; \lambda^t) \end{aligned}$$

Inequality (a) follows Lemma 1 and inequality (b) follows the optimization step in Line 14 in Algorithm 1. Inequality (c) follows from the fact that λ increases monotonically so that $\lambda^t \leq \lambda^{t+1}$. As $L(\{\mathbf{w}^{t+1}\}, \mathbf{z}^{t+1}, \boldsymbol{\alpha}^{t+1})$ for some constant values of \mathbf{v} and λ has a lower bound \mathcal{B} , we can easily prove that $L(\{\mathbf{w}^{t+1}\}, \mathbf{v}^{t+1}, \mathbf{z}^{t+1}, \boldsymbol{\alpha}^{t+1}; \lambda^{t+1}) \geq \mathcal{B} + C - \tau_\lambda n$, where C is a constant and n is the size of the entire dataset. Therefore, the Lagrangian L is convergent. According to the stop condition for Algorithm 1, the algorithm converges when the Lagrangian L is converged. \square

5 Experimental Results

In this section, the performance of the proposed algorithm DSPL is evaluated for both synthetic and real-world datasets in

Table 2: Regression Coefficient Recovery Performance for Different Corrupted Batches

	4/10	5/10	6/10	7/10	8/10	9/10
TORR	0.093	0.109	0.088	0.086	0.079	0.083
TORR25	0.174	0.165	0.189	0.214	0.216	0.241
RLHH	0.635	0.619	0.670	0.907	0.851	0.932
DRLR	0.014	0.131	0.222	0.274	0.304	0.346
SPL	0.038	0.047	0.047	0.044	0.053	0.064
DSPL	0.030	0.034	0.039	0.036	0.041	0.045

robust regression task. After the experimental setup has been introduced in Section 5.1, we present the results for the regression coefficient recovery performance with different settings using synthetic data in Section 5.2, followed by house rental price prediction evaluation using real-world data in Section 5.3. All the experiments were performed on a 64-bit machine with an Intel(R) Core(TM) quad-core processor (i7CPU@3.6GHz) and 32.0GB memory. Details of both the source code and the datasets used in the experiment can be downloaded here¹.

5.1 Experimental Setup

Datasets and Labels

The datasets used for the experimental verification were composed of synthetic and real-world data. The simulation samples were randomly generated according to the model $\mathbf{y}^{(i)} = [X^{(i)}]^T \mathbf{w}_* + \mathbf{u}^{(i)} + \varepsilon^{(i)}$ for each mini-batch, where \mathbf{w}_* represents the ground truth coefficients and $\mathbf{u}^{(i)}$ the adversarial corruption vector. $\varepsilon^{(i)}$ represents the additive dense noise for the i^{th} batch, where $\varepsilon_j^{(i)} \sim \mathcal{N}(0, \sigma^2)$. We sampled the regression coefficients $\mathbf{w}_* \in \mathbb{R}^p$ as a random unit norm vector. The covariance data $X^{(i)}$ for each mini-batch was drawn independently and identically distributed from $x_i \sim \mathcal{N}(0, I_p)$ and the uncorrupted response variables were generated as $\mathbf{y}_*^{(i)} = [X^{(i)}]^T \mathbf{w}_* + \varepsilon^{(i)}$. The corrupted response vector for each mini-batch was generated as $\mathbf{y}^{(i)} = \mathbf{y}_*^{(i)} + \mathbf{u}^{(i)}$, where the corruption vector $\mathbf{u}^{(i)}$ was sampled from the uniform distribution $[-5\|\mathbf{y}_*^{(i)}\|_\infty, 5\|\mathbf{y}_*^{(i)}\|_\infty]$. The set of uncorrupted points was selected as a uniformly random subset of $[n_i]$ for each batch.

The real-world datasets utilized consisted of house rental transaction data from two cities, *New York City* and *Los Angeles* listed on the Airbnb² website from January 2015 to October 2016. These datasets can be publicly downloaded³. For the *New York City* dataset, the first 321,530 data samples from January 2015 to December 2015 were used as the training data and the remaining 329,187 samples from January to October 2016 as the test data. For the *Los Angeles* dataset, the first 106,438 samples from May 2015 to May 2016 were used as training data, and the remaining 103,711 samples as test data. In each dataset, there were 21 features after data preprocessing, including the number of beds and bathrooms, location, and average price in the area.

¹<https://goo.gl/cis7tK>

²<https://www.airbnb.com/>

³<http://insideairbnb.com/get-the-data.html>

Evaluation Metrics

For the synthetic data, we measured the performance of the regression coefficient recovery using the averaged L_2 error $e = \|\hat{\mathbf{w}} - \mathbf{w}_*\|_2$, where $\hat{\mathbf{w}}$ represents the recovered coefficients for each method and \mathbf{w}_* represents the ground truth regression coefficients. For the real-world dataset, we used the mean absolute error (MAE) to evaluate the performance for rental price prediction. Defining $\hat{\mathbf{y}}$ and \mathbf{y} as the predicted price and ground truth price, respectively, the mean absolute error between $\hat{\mathbf{y}}$ and \mathbf{y} can be presented as $\text{MAE}(\hat{\mathbf{y}}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|$, where y_i represents the label of the i^{th} data sample.

Comparison Methods

We used the following methods to compare the performance of the robust regression task: *Torrent* (Abbr. *TORR*) [Bhatia *et al.*, 2015], which is a hard-thresholding based method that includes a parameter for the corruption ratio. As this parameter is hard to estimate in practice, we opted to use a variant, *TORR25*, which represents a corruption parameter that is uniformly distributed across a range of $\pm 25\%$ off the true value. We also used *RLHH* [Zhang *et al.*, 2017b] for the comparison, which applies a recently proposed robust regression method based on heuristic hard thresholding with no additional parameters. This method computes the regression coefficients for each batch, and averages them all. The *DRLR* [Zhang *et al.*, 2017a] algorithm, which is a distributed robust learning method specifically designed to handle large scale data with a distributed robust consolidation. The traditional self-paced learning algorithm (*SPL*) [Kumar *et al.*, 2010] with the parameter $\lambda = 1$ and the step size $\mu = 1.1$ was also compared in our experiment. For *DSPL*, we used the same settings as for *SPL* with the initial $\lambda_0 = 0.1$ and $\tau_\lambda = 1$. All the results from each of these comparison methods were averaged over 10 runs.

5.2 Robust Regression in Synthetic Data

Recovery Coefficients Recovery

Figure 1 shows the coefficient recovery performance for different corruption ratios in uniform distribution. Specifically, Figures 1(a) and 1(b) show the results for a different number of features with a fixed data size. Looking at the results, we can conclude: 1) Of the six methods tested, the *DSPL* method outperformed all the competing methods, including *TORR*, whose corruption ratio parameter uses the ground truth value. 2) Although *DRLR* turned in a competitive performance when the data corruption level was low. However, when the corruption ratio rose to over 40%, the recovery error is increased dramatically. 3) The *TORR* method is highly dependent on the corruption ratio parameter. When the parameter is 25% different from the ground truth, the error for *TORR25* was over 50% compared to *TORR*, which uses the ground truth corruption ratio. 4) When the feature number is increased, the average error for the *SPL* algorithm increased by a factor of four. However, the results obtained for the *DSPL* algorithm varied consistently with the corruption ratio and feature number. The results presented in Figures 1(a) and 1(c) conform that the *DSPL* method consistently outperformed the other methods for larger datasets, while still achieving a close recovery of the ground truth coefficient.

Performance on Different Corrupted Batches

The regression coefficient recovery performance for different numbers of corrupted batches is shown in Table 2, ranging from

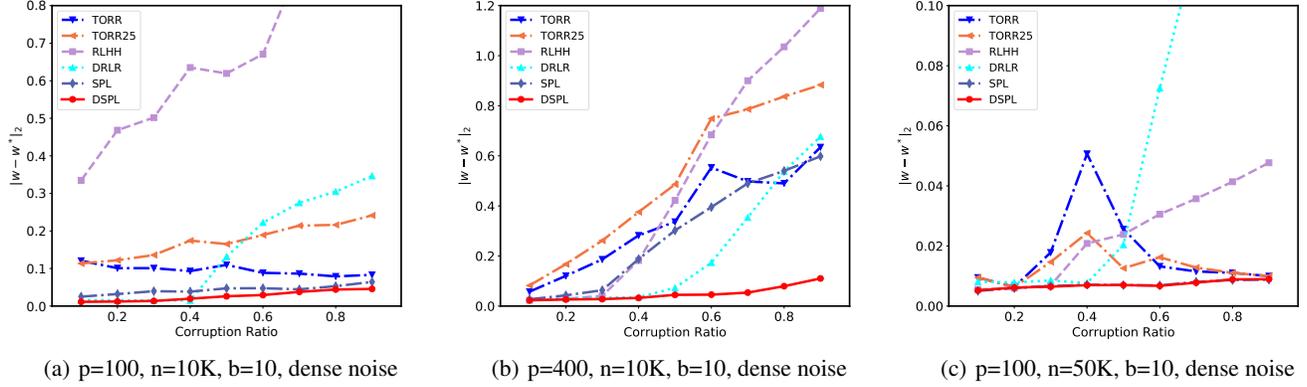


Figure 1: Regression coefficient recovery performance for different corruption ratios.

New York City (Corruption Ratio)						
	10%	30%	50%	70%	90%	Avg.
TORR	3.970±0.007	4.097±0.199	5.377±2.027	7.025±3.379	7.839±3.435	5.662±1.809
TORR25	3.978±0.012	4.207±0.324	5.885±2.615	7.462±3.569	8.369±3.675	5.980±2.039
RLHH	3.965±0.000	4.244±0.544	5.977±2.543	7.525±3.491	8.463±3.646	6.034±2.045
DRLR	3.963±0.000	4.026±0.089	5.884±2.692	7.350±3.469	8.325±3.669	5.908±1.984
SPL	3.979±0.006	4.141±0.199	5.185±1.578	6.413±2.562	7.283±2.892	5.400±1.447
DSPL	3.972±0.007	4.020±0.085	4.123±0.198	5.291±2.086	6.444±2.997	4.770±1.075

Los Angeles (Corruption Ratio)						
	10%	30%	50%	70%	90%	Avg.
TORR	3.991±0.001	4.035±0.067	5.666±2.754	7.569±4.098	8.561±4.170	5.964±2.218
TORR25	3.993±0.003	4.103±0.147	5.986±3.062	7.834±4.181	8.930±4.338	6.169±2.346
RLHH	3.992±0.000	4.023±0.064	6.224±3.198	8.013±4.179	9.091±4.317	6.268±2.352
DRLR	3.990±0.001	4.016±0.031	6.471±3.552	8.147±4.246	9.197±4.341	6.364±2.434
SPL	3.994±0.004	4.135±0.159	5.432±2.097	6.856±3.109	7.857±3.435	5.655±1.761
DSPL	3.992±0.021	4.034±0.137	4.510±0.599	5.717±2.237	6.943±3.194	5.062±1.238

Table 3: Mean Absolute Error for Rental Price Prediction

four to nine corrupted batches out of the total of 10 batches. Each corrupted batch used in the experiment contains 90% corrupted samples and each uncorrupted batch has 10% corrupted samples. The results are shown for the averaged L_2 error in 10 different synthetic datasets with randomly ordered batches. Looking at the results shown in Table 2, we can conclude: 1) When the ratio of corrupted batches exceeds 50%, *DSPL* outperforms all the competing methods with a consistent recovery error. 2) *DRLR* performs the best when the mini-batch is 40% corrupted, although its recovery error increases dramatically when the number of corrupted batch increases. 3) *SPL* turns in a competitive performance for different levels of corrupted batches, but its error almost doubles when the number of corrupted batches increases from four to nine.

Analysis of Parameter λ

Figure 2 show the relationship between the parameter λ and the coefficient recovery error, along with the corresponding Lagrangian L . This result is based on the robust coefficient recovery task for a 90% data corruption setting. Examining the blue line, as the parameter λ increases, the recovery error continues to decrease until it reaches a critical point, after which it increases. These results indicate that the training process will

keep improving the model until the parameter λ becomes so large that some corrupted samples become incorporated into the training data. In the case shown here, the critical point is around 1.0. The red line shows the value of the Lagrangian L in terms of different values of the parameter λ , leading us to conclude that: 1) the Lagrangian monotonically decreases as λ increases. 2) The Lagrangian decreases much faster once λ reaches a critical point, following the same pattern as the recovery error shown in blue line.

5.3 House Rental Price Prediction

To evaluate the effectiveness of our proposed method in a real-world dataset, we compared its performance for rental price prediction for a number of different corruption settings, ranging from 10% to 90%. The additional data corruption was sampled from the uniform distribution $[-0.5y_i, 0.5y_i]$, where y_i denotes the price value of the i^{th} data point. Table 3 shows the mean absolute error for rental price prediction and its corresponding standard deviation over 10 runs for the New York City and Los Angeles datasets. The results indicate that: 1) The *DSPL* method outperforms all the other methods for all the different corruption settings except when the corruption ratio is less than

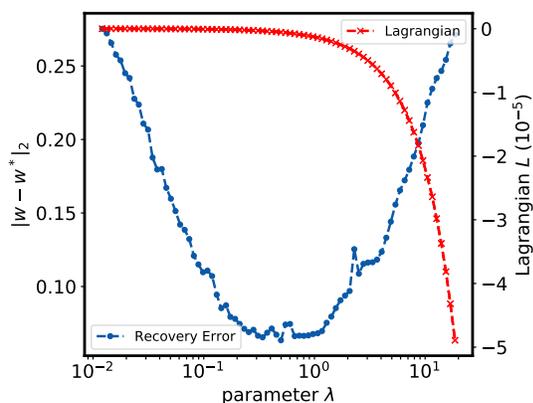


Figure 2: Relationship between parameter λ and coefficient recovery error and the corresponding Lagrangian.

30%, and consistently produced with the most stable results (smallest standard deviation). 2) Although the *DRLR* method performs the best when the corruption ratio is less than 30%, the results of all the methods are very close. Moreover, as the corruption ratio rises, the error for *DRLR* increases dramatically. 3) *SPL* has a very competitive performance for all the corruption settings but is still around 12% worse than the new *DSPL* method proposed here, which indicates that considering the data integrally produces a better performance than can be achieved by breaking up the data into batches and treating them separately.

6 Conclusion

In this paper, a distributed self-paced learning algorithm (*DSPL*) is proposed to extend the traditional *SPL* algorithm to its distributed version for large scale datasets. To achieve this, we reformulated the original *SPL* problem into a distributed setting and optimized the problem of treating different mini-batches in parallel based on consensus *ADMM*. We also proved that our algorithm can be convergent under mild assumptions. Extensive experiments on both synthetic data and real-world rental price data demonstrated that the proposed algorithms are very effective, outperforming the other comparable methods over a range of different data settings.

References

- [Bengio *et al.*, 2009] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM, 2009.
- [Bhatia *et al.*, 2015] Kush Bhatia, Prateek Jain, and Purushottam Kar. Robust regression via hard thresholding. In *Advances in Neural Information Processing Systems*, pages 721–729, 2015.
- [Boyd *et al.*, 2011] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [Fan *et al.*, 2017] Yanbo Fan, Ran He, Jian Liang, and Baogang Hu. Self-paced learning: an implicit regularization perspective. In *AAAI*, 2017.
- [Gorski *et al.*, 2007] Jochen Gorski, Frank Pfeuffer, and Kathrin Klamroth. Biconvex sets and optimization with biconvex functions: a survey and extensions. *Mathematical Methods of Operations Research*, 66(3):373–407, 2007.
- [Hong *et al.*, 2016] Mingyi Hong, Zhi-Quan Luo, and Meisam Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *SIAM Journal on Optimization*, 26(1):337–364, 2016.
- [Jiang *et al.*, 2014a] Lu Jiang, Deyu Meng, Teruko Mitamura, and Alexander G Hauptmann. Easy samples first: Self-paced reranking for zero-example multimedia search. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 547–556. ACM, 2014.
- [Jiang *et al.*, 2014b] Lu Jiang, Deyu Meng, Shou-I Yu, Zhenzhong Lan, Shiguang Shan, and Alexander Hauptmann. Self-paced learning with diversity. In *Advances in Neural Information Processing Systems*, pages 2078–2086, 2014.
- [Jiang *et al.*, 2015] Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander Hauptmann. Self-paced curriculum learning, 2015.
- [Keerthiram Murugesan, 2017] Jaime Carbonell Keerthiram Murugesan. Self-paced multitask learning with shared knowledge. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 2522–2528, 2017.
- [Kumar *et al.*, 2010] M Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *Advances in Neural Information Processing Systems*, pages 1189–1197, 2010.
- [Li *et al.*, 2017] Changsheng Li, Junchi Yan, Fan Wei, Weishan Dong, Qingshan Liu, and Hongyuan Zha. Self-paced multi-task learning. In *AAAI*, pages 2175–2181, 2017.
- [Ma *et al.*, 2017a] Fan Ma, Deyu Meng, Qi Xie, Zina Li, and Xuanyi Dong. Self-paced co-training. In *International Conference on Machine Learning*, pages 2275–2284, 2017.
- [Ma *et al.*, 2017b] Zilu Ma, Shiqi Liu, and Deyu Meng. On convergence property of implicit self-paced objective. *arXiv preprint arXiv:1703.09923*, 2017.
- [Meng *et al.*, 2015] Deyu Meng, Qian Zhao, and Lu Jiang. What objective does self-paced learning indeed optimize? *arXiv preprint arXiv:1511.06049*, 2015.
- [Pi *et al.*, 2016] Te Pi, Xi Li, Zhongfei Zhang, Deyu Meng, Fei Wu, Jun Xiao, and Yueting Zhuang. Self-paced boost learning for classification. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, pages 1932–1938. AAAI Press, 2016.
- [Supancic and Ramanan, 2013] James S Supancic and Deva Ramanan. Self-paced learning for long-term tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2379–2386, 2013.
- [Xu *et al.*, 2015] Chang Xu, Dacheng Tao, and Chao Xu. Multi-view self-paced learning for clustering. In *IJCAI*, pages 3974–3980, 2015.

- [Zhang *et al.*, 2016] Dingwen Zhang, Deyu Meng, Long Zhao, and Junwei Han. Bridging saliency detection to weakly supervised object detection based on self-paced curriculum learning. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, pages 3538–3544. AAAI Press, 2016.
- [Zhang *et al.*, 2017a] Xuchao Zhang, Liang Zhao, Arnold P. Boedihardjo, and Chang-Tien Lu. Online and distributed robust regressions under adversarial data corruption. *2017 IEEE International Conference on Data Mining (ICDM)*, pages 625–634, 2017.
- [Zhang *et al.*, 2017b] Xuchao Zhang, Liang Zhao, Arnold P. Boedihardjo, and Chang-Tien Lu. Robust regression via heuristic hard thresholding. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI'17*. AAAI Press, 2017.
- [Zhang *et al.*, 2017c] Xuchao Zhang, Liang Zhao, Arnold P. Boedihardjo, Chang-Tien Lu, and Naren Ramakrishnan. Spatiotemporal event forecasting from incomplete hyper-local price data. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, pages 507–516, New York, NY, USA, 2017. ACM.