
Theory of Dual-sparse Regularized Randomized Reduction

Tianbao Yang

TIANBAO-YANG@UIOWA.EDU

Department of Computer Science, the University of Iowa, Iowa City, USA

Lijun Zhang

ZHANGLJ@LAMDA.NJU.EDU.CN

National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

Rong Jin

RONGJIN@CSE.MSU.EDU

Department of Computer Science and Engineering, Michigan State University, East Lansing, USA

Institute of Data Science and Technologies at Alibaba Group, Seattle, USA

Shenghuo Zhu

SHENGHUO@GMAIL.COM

Institute of Data Science and Technologies at Alibaba Group, Seattle, USA

Abstract

In this paper, we study *randomized reduction methods*, which reduce high-dimensional features into low-dimensional space by randomized methods (e.g., random projection, random hashing), for large-scale high-dimensional classification. Previous theoretical results on randomized reduction methods hinge on strong assumptions about the data, e.g., low rank of the data matrix or a large separable margin of classification, which hinder their applications in broad domains. To address these limitations, we propose dual-sparse regularized randomized reduction methods that introduce a sparse regularizer into the reduced dual problem. Under a mild condition that the original dual solution is a (nearly) sparse vector, we show that the resulting dual solution is close to the original dual solution and concentrates on its support set. In numerical experiments, we present an empirical study to support the analysis and we also present a novel application of the dual-sparse regularized randomized reduction methods to reducing the communication cost of distributed learning from large-scale high-dimensional data.

1. Introduction

As the scale and dimensionality of data continue to grow in many applications (e.g., bioinformatics, finance, com-

puter vision, medical informatics) (Sánchez et al., 2013; Mitchell et al., 2004; Simianer et al., 2012; Bartz et al., 2011), it becomes critical to develop efficient and effective algorithms to solve big data machine learning problems. Randomized reduction methods for large-scale or high-dimensional data analytics have received a great deal of attention in recent years (Mahoney & Drineas, 2009; Shi et al., 2012; Paul et al., 2013; Weinberger et al., 2009; Mahoney, 2011). By either reducing the dimensionality (*referred to as feature reduction*) or reducing the number of training instances (*referred to as instance reduction*), the resulting problem has a smaller size of training data that is not only memory-efficient but also computation-efficient. While randomized instance reduction has been studied a lot for fast least square regression (Drineas et al., 2008; 2006; 2011; Ma et al., 2014), randomized feature reduction is more popular for linear classification (Blum, 2005; Shi et al., 2012; Paul et al., 2013; Weinberger et al., 2009; Shi et al., 2009a) (e.g., random hashing is a noticeable built-in tool in Vowpal Wabbit¹, a fast learning library, for solving high-dimensional problems.). In this paper, we focus on the latter technique and refer to randomized feature reduction as randomized reduction for short.

Although several theoretical properties have been examined for randomized reduction methods when applied to classification, e.g., generalization performance (Paul et al., 2013), preservation of margin (Blum, 2005; Balcan et al., 2006; Shi et al., 2012) and the recovery error of the model (Zhang et al., 2014), these previous results rely on strong assumptions about the data. For example, both (Paul et al., 2013) and (Zhang et al., 2014) assume the data matrix is of low-rank, and (Blum, 2005; Balcan et al.,

Proceedings of the 32nd International Conference on Machine Learning, Lille, France, 2015. JMLR: W&CP volume 37. Copyright 2015 by the author(s).

¹<http://hunch.net/~vw/>

2006; Shi et al., 2012) make a assumption that all examples in the original space are separated with a positive margin (with a high probability). Another analysis in (Zhang et al., 2014) assumes the weight vector for classification is sparse. These assumptions are too strong to hold in many real applications.

Contributions. To address these limitations, we propose dual-sparse regularized randomized reduction methods referred to as **DSRR** by leveraging the (near) sparsity of dual solutions for large-scale high-dimensional (**LSHD**) classification problems (i.e., the number of (effective) support vectors is small compared to the total number of examples). In particular, we add a dual-sparse regularizer into the reduced dual problem. We present a novel theoretical analysis of the recovery error of the dual variables and the primal variable and study its implication for different randomized reduction methods (e.g., random projection, random hashing and random sampling).

Novelties. Compared with previous works (Blum, 2005; Balcan et al., 2006; Shi et al., 2012; Paul et al., 2013), our theoretical analysis demands a mild assumption about the data and directly provides guarantee on a small recovery error of the obtained model, which is critical for subsequent analysis, e.g., feature selection (Guyon et al., 2002; Brank et al., 2002) and model interpretation (Rätsch et al., 2005; Sonnenburg & Franc, 2010; Rtsch et al., 2005; Sonnenburg et al., 2007; Ben-Hur et al., 2008). For example, when exploiting a linear model to classify people into sick or not sick based on genomic markers, the learned weight vector is important for understanding the effect of different genomic markers on the disease and for designing effective medicine (Jostins & Barrett, 2011; Kang & Cho, 2011). In addition, the recovery could also increase the predictive performance, in particular when there exists noise in the original features (Goldberger et al., 2005).

Compared with (Zhang et al., 2014) that proposes to recover a linear model in the original feature space by dual recovery, i.e., constructing a weight vector using the dual variables learned from the reduced problem and the original feature vectors, our methods are better in that (i) we rely on a more realistic assumption of the sparsity of dual variables (e.g., in support vector machine (SVM)); (ii) we analyze both smooth loss functions and non-smooth loss functions (they focused on smooth functions); (iii) we study different randomized reduction methods in the same framework not just the random projection.

In numerical experiments, we present an empirical study on a real data set to support our analysis and we also demonstrate a novel application of the reduction and recovery framework in distributed learning from LSHD data, which combines the benefits of the two complementary techniques for addressing big data problems. Distributed

learning/optimization recently receives significant interest in solving big data problems (Jaggi et al., 2014; Li et al., 2014; Yang, 2013; Agarwal et al., 2011). However, it is notorious for high communication cost, especially when the dimensionality of data is very high. By solving a dimensionality reduced data problem and using the recovered solution as an initial solution to the distributed optimization on the original data, we can reduce the number of iterations and the communication cost. In practice, we employ the recently developed distributed stochastic dual coordinate ascent algorithm (Yang, 2013), and observe that using the recovered solution as an initial solution we are able to attain almost the same performance with only one or two communications of high dimensional vectors among multiple machines.

2. Preliminaries

Let $(\mathbf{x}_i, y_i), i = 1, \dots, n$ denote a set of training examples, where $\mathbf{x}_i \in \mathbb{R}^d, y_i \in \{1, -1\}$. Assume both n and d are very large. The goal of classification is to solve the following optimization problem:

$$\mathbf{w}_* = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{w}^\top \mathbf{x}_i y_i) + \frac{\lambda}{2} \|\mathbf{w}\|_2^2 \quad (1)$$

where $\ell(z)$ is a convex loss function and λ is a regularization parameter. Using the conjugate function, we can turn the problem into a dual problem:

$$\alpha_* = \arg \max_{\alpha \in \mathbb{R}^n} -\frac{1}{n} \sum_{i=1}^n \ell_i^*(\alpha_i) - \frac{1}{2\lambda n^2} \alpha^\top X^\top X \alpha \quad (2)$$

where $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ is the data matrix and $\ell_i^*(\alpha)$ is the convex conjugate function of $\ell(z y_i)$. Given the optimal dual solution α_* , the optimal primal solution can be computed by $\mathbf{w}_* = -\frac{1}{\lambda n} X \alpha_*$. For LSHD problems, directly solving the primal problem (1) or the dual problem (2) could be very expensive. We aim to address the challenge by randomized reduction methods. Let $A(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^m$ denote a randomized reduction operator that reduces a d -dimensional feature vector into m -dimensional feature vector. Let $\hat{\mathbf{x}} = A(\mathbf{x})$ denote the reduced feature vector. With the reduced feature vectors $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n$ of the training examples, a conventional approach is to solve the following reduced primal problem

$$\mathbf{u}_* = \arg \min_{\mathbf{u} \in \mathbb{R}^m} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{u}^\top \hat{\mathbf{x}}_i y_i) + \frac{\lambda}{2} \|\mathbf{u}\|_2^2 \quad (3)$$

or its the dual problem

$$\hat{\alpha}_* = \arg \max_{\alpha \in \mathbb{R}^n} -\frac{1}{n} \sum_{i=1}^n \ell_i^*(\alpha_i) - \frac{1}{2\lambda n^2} \alpha^\top \hat{X}^\top \hat{X} \alpha \quad (4)$$

where $\hat{X} = (\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n) \in \mathbb{R}^{m \times n}$. Previous studies have analyzed the reduced problems for random projection methods and proved the preservation of margin (Blum, 2005; Shi et al., 2012) and the preservation of minimum

enclosing ball (Paul et al., 2013). Zhang et al. (2014) proposed a dual recovery approach that constructs a recovered solution by $\widehat{\mathbf{w}}_* = -\frac{1}{\lambda n} \sum_{i=1}^n [\widehat{\alpha}_*]_i \mathbf{x}_i$ and proved the recovery error for random projection under the assumption of low-rank data matrix or sparse \mathbf{w}_* . In addition, they also showed that the naive recovery by $A^\top \mathbf{u}_*$ (when $A(\mathbf{x}) = A\mathbf{x}$) has a large recovery error.

One deficiency with the simple dual recovery approach is that due to the reduction in the feature space, many non-support vectors for the original optimization problem will become support vectors, which could result in the corruption in the recovery error. As a result, the original analysis of dual recovery method requires a strong assumption of data (i.e., the low rank assumption). In this work, we plan to address this limitation in a different way, which allows us to relax the assumption significantly.

3. DSRR and its Guarantee

To reduce the number of or the contribution of training instances that are non-support vectors in the original optimization problem and are transformed into support vectors due to the reduction of the feature space, we employ a simple trick that adds a dual-sparse regularization to the reduced dual problem. In particular, we solve the following problem:

$$\tilde{\alpha}_* = \quad (5)$$

$$\arg \max_{\alpha \in \mathbb{R}^n} -\frac{1}{n} \sum_{i=1}^n \ell_i^*(\alpha_i) - \frac{1}{2\lambda n^2} \alpha^T \widehat{X}^\top \widehat{X} \alpha - \frac{1}{n} R(\alpha)$$

where $R(\alpha) = \tau \|\alpha\|_1$, and $\tau > 0$ is a regularization parameter, whose theoretical value will be revealed later.

To further understand the added dual-sparse regularizer, we consider SVM, where the loss function can be either the hinge loss (a non-smooth function) $\ell(zy) = \max(0, 1 - zy)$ or the squared hinge loss (a smooth function) $\ell(zy) = \max(0, 1 - zy)^2$. We first consider the hinge loss, where $\ell_i^*(\alpha_i) = \alpha_i y_i$ for $\alpha_i y_i \in [-1, 0]$. Then the new dual problem is equivalent to

$$\max_{\alpha \circ \mathbf{y} \in [-1, 0]^n} \frac{1}{n} \sum_{i=1}^n -\alpha_i y_i - \frac{1}{2\lambda n^2} \alpha^T \widehat{X}^\top \widehat{X} \alpha - \frac{\tau}{n} \|\alpha\|_1$$

Using variable transformation $-\alpha_i y_i \rightarrow \beta_i$, the above problem is equivalent to

$$\max_{\beta \in [0, 1]^n} \frac{1}{n} \sum_{i=1}^n \beta_i (1 - \tau) - \frac{1}{2\lambda n^2} (\beta \circ \mathbf{y})^T \widehat{X}^\top \widehat{X} (\beta \circ \mathbf{y})$$

Changing into the primal form, we have

$$\max_{\mathbf{u} \in \mathbb{R}^m} \frac{1}{n} \sum_{i=1}^n \ell_{1-\tau}(\mathbf{u}^\top \widehat{\mathbf{x}}_i y_i) + \frac{\lambda}{2} \|\mathbf{u}\|_2^2 \quad (6)$$

where $\ell_\gamma(z) = \max(0, \gamma - z)$ is a max-margin loss with margin given by γ . It can be understood that adding the ℓ_1 regularization in the reduced problem of SVM is equivalent

to using a max-margin loss with a smaller margin, which is intuitive because examples become difficult to separate after dimensionality reduction and is consistent with several previous studies that the margin is reduced in the reduced feature space (Blum, 2005; Shi et al., 2012). Similarly for squared hinge loss, the equivalent primal problem is

$$\max_{\mathbf{u} \in \mathbb{R}^m} \frac{1}{n} \sum_{i=1}^n \ell_{1-\tau}^2(\mathbf{u}^\top \widehat{\mathbf{x}}_i y_i) + \frac{\lambda}{2} \|\mathbf{u}\|_2^2 \quad (7)$$

where $\ell_\gamma^2(z) = \max(0, \gamma - z)^2$.

Although adding a dual-sparse regularizer is intuitive and can be motivated from previous results, we emphasize that the proposed dual-sparse formulation provides a new perspective and bounding the dual recovery error $\|\tilde{\alpha}_* - \alpha_*\|$ is a non-trivial task, which is a major contribution of this paper. We first state our main result in Theorem 1 for smooth loss functions.

Theorem 1. *Let $\tilde{\alpha}_*$ be the optimal dual solution to (5). Assume α_* is s -sparse with the support set given by S . If $\tau \geq \frac{2}{\lambda n} \|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty$, then we have*

$$\|[\tilde{\alpha}_*]_{S^c}\|_1 \leq 3 \|[\tilde{\alpha}_*]_S - [\alpha_*]_S\|_1 \quad (8)$$

Furthermore, if $\ell(z)$ is a L -smooth loss function², we have

$$\|\tilde{\alpha}_* - \alpha_*\|_2 \leq 3\tau L \sqrt{s}, \quad \|\tilde{\alpha}_* - \alpha_*\|_1 \leq 12\tau L s \quad (9)$$

$$\|[\tilde{\alpha}_*]_S - [\alpha_*]_S\|_1 \leq 3\tau L s, \quad \|[\tilde{\alpha}_*]_{S^c}\|_1 \leq 9\tau L s \quad (10)$$

where S^c is the complement of S , and $[\alpha]_S$ is a vector that only contains the elements of α in the set S .

Remark 1: The proof is presented at the end of Section 4. It can be seen that the dual recovery error is proportional to the value of τ which is dependent on $\|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty$, which we can bound without using any assumption about the data matrix or the optimal dual variable α_* . In contrast, previous bounds (Zhang et al., 2013; 2014; Paul et al., 2013) depend on $\|X^\top X - \widehat{X}^\top \widehat{X}\|_2$, which requires the low rank assumption on X . In next section, we provide an upper bound of $\frac{1}{\lambda n} \|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty$ that will allow us to understand how the reduced dimensionality m affects the recovery error. Essentially, the results indicate that for random projection, randomized Hadamard transform and random hashing, $\frac{1}{\lambda n} \|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty \leq O(\sqrt{\frac{\log(n/\delta)}{m}}) \|\mathbf{w}_*\|_2$ with a high probability $1 - \delta$, and thus the recovery error will be scaled as $\sqrt{1/m}$ in terms of m - the same order of recovery error as in (Zhang et al., 2013; 2014) that assumes low rank of the data matrix.

Remark 2: We would like to make a connection with LASSO for sparse signal recovery. In sparse signal recovery under noise measurements $\mathbf{f} = U\mathbf{w}_* + \mathbf{e}$, where \mathbf{e} denotes the noise in measurements, if a LASSO

²A function is L -smooth if its gradient is L -Lipschitz continuous.

$\min_{\mathbf{w}} \frac{1}{2} \|U\mathbf{w} - \mathbf{f}\|_2^2 + \lambda \|\mathbf{w}\|_1$ is solved for the solution, then the regularization parameter λ is required to be larger than the quantity $\|U^\top \mathbf{e}\|_\infty$ that depends on the noise in order to have an accurate recovery (Eldar & Kutyniok, 2012). Similarly in our formulation, the added ℓ_1 regularization $\tau \|\alpha\|_1$ is to counteract the noise in $\widehat{X}\widehat{X}^\top$ as compared with XX^\top and the value of τ is dependent on the noise.

To present the theoretical result on the non-smooth loss functions, we need to introduce restricted eigen-value conditions similar to those used in the sparse recovery analysis for LASSO (Bickel et al., 2009; Xiao & Zhang, 2013). In particular, we introduce the following definition of restricted eigen-value condition.

Definition 2. Given an integer $s > 0$, we define

$$\mathcal{K}_{n,s} = \{\alpha \in \mathbb{R}^n : \|\alpha\|_2 \leq 1, \|\alpha\|_1 \leq \sqrt{s}\}.$$

We say that X satisfies the restricted eigenvalue condition at sparsity level s if there exist positive constants ρ_s^+ and ρ_s^- such that

$$\rho_s^+ = \sup_{\alpha \in \mathcal{K}_{n,s}} \frac{\alpha^\top X^\top X \alpha}{n}, \quad \rho_s^- = \inf_{\alpha \in \mathcal{K}_{n,s}} \frac{\alpha^\top X^\top X \alpha}{n}.$$

We also define another quantity that measures the restricted eigen-value of $X^\top X - \widehat{X}^\top \widehat{X}$, namely

$$\sigma_s = \sup_{\alpha \in \mathcal{K}_{n,s}} \frac{|\alpha^\top (X^\top X - \widehat{X}^\top \widehat{X}) \alpha|}{n}. \quad (11)$$

Theorem 3. Let $\tilde{\alpha}_*$ be the optimal dual solution to (5). Assume α_* is s -sparse with the support set given by \mathcal{S} . If $\tau \geq \frac{2}{\lambda n} \|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty$, then we have

$$\|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1 \leq 3\|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1$$

Assume the data matrix X satisfies the restricted eigen-value condition at sparsity level $16s$ and $\sigma_{16s} < \rho_{16s}^-$, we have

$$\begin{aligned} \|\tilde{\alpha}_* - \alpha_*\|_2 &\leq \frac{3\lambda}{2(\rho_{16s}^- - \sigma_{16s})} \tau \sqrt{s} \\ \|\tilde{\alpha}_* - \alpha_*\|_1 &\leq \frac{6\lambda}{(\rho_{16s}^- - \sigma_{16s})} \tau s \end{aligned}$$

Remark 3: The proof appears in the full version of the paper³. Compared to smooth loss functions, the conditions that guarantee a small recovery for non-smooth loss functions are more restricted. In next section, we will provide a bound on σ_{16s} to further understand the condition of $\sigma_{16s} \leq \rho_{16s}^-$, which essentially implies that $m \geq \Omega\left(\left(\frac{\rho_{16s}^+}{\rho_{16s}^-}\right)^2 s \log(n/s)\right)$.

Last but not least, we provide a theoretical result on the recovery error for the nearly sparse optimal dual variable α_* . We state the result for smooth loss functions. To quantify the near sparsity, we let $\alpha_*^s \in \mathbb{R}^n$ denote a vector that zeros all entries in α_* except for the top- s elements in magnitude

and assume α_*^s satisfies the following condition:

$$\left\| \nabla \ell^*(\alpha_*^s) + \frac{1}{\lambda n} X^\top X \alpha_*^s \right\|_\infty \leq \xi \quad (12)$$

where $\nabla \ell^*(\alpha) = (\nabla \ell_1^*(\alpha_1), \dots, \nabla \ell_n^*(\alpha_n))^\top$. The above condition can be considered as a sub-optimality condition (Boyd & Vandenberghe, 2004) of α_*^s measured in the infinite norm. For the optimal solution α_* , we have $\nabla \ell^*(\alpha_*) + \frac{1}{\lambda n} X^\top X \alpha_* = 0$.

Theorem 4. Let $\tilde{\alpha}_*$ be the optimal dual solution to (5). Assume α_* is nearly s -sparse such that (12) holds with the support set of α_*^s given by \mathcal{S} . If $\tau \geq \frac{2}{\lambda n} \|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty + 2\xi$, then we have

$$\|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1 \leq 3\|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1$$

Furthermore, if $\ell(z)$ is a L -smooth loss function, we have

$$\|\tilde{\alpha}_* - \alpha_*^s\|_2 \leq 3\tau L \sqrt{s}, \quad \|\tilde{\alpha}_* - \alpha_*^s\|_1 \leq 12\tau L s \quad (13)$$

$$\|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1 \leq 3\tau L s, \quad \|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1 \leq 9\tau L s \quad (14)$$

Remark 4: The proof appears in the full version of the paper. Compared to Theorem 1 for exactly sparse optimal dual solution, the dual recovery error bound for nearly sparse optimal dual solution is increased by $6L\sqrt{s}\xi$ for ℓ_2 norm and by $24Ls\xi$ for ℓ_1 norm.

Finally, we note that with the recovery error bound for the dual solution, we can easily derive an error bound for the primal solution $\tilde{\mathbf{w}}_* = -\frac{1}{\lambda n} X \tilde{\alpha}_*$. Below we present a theorem for smooth loss functions. One can easily extend the result to non-smooth loss functions.

Theorem 5. Let $\tilde{\mathbf{w}}_*$ be the recovered primal solution using $\tilde{\alpha}_*$ the optimal dual solution to (5). Assume α_* is s -sparse and $\ell(z)$ is a L -smooth loss function. If $\tau \geq \frac{2}{\lambda n} \|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty$ then we have

$$\|\tilde{\mathbf{w}}_* - \mathbf{w}_*\|_2 \leq \frac{\sigma_1}{\lambda n} 3L\tau \sqrt{s}$$

where σ_1 is the maximum singular value of X . Furthermore if $\frac{1}{n} X^\top X$ has a restricted eigen-value ρ_{16s}^+ at sparsity level $16s$, then

$$\|\tilde{\mathbf{w}}_* - \mathbf{w}_*\|_2 \leq \frac{\sqrt{\rho_{16s}^+}}{\lambda \sqrt{n}} 3L\tau \sqrt{s}$$

Remark 5: Since ρ_{16s}^+ is always less than σ_1^2/n , the second result if the restricted eigen-value condition holds is always better than the first result. With the bound of τ as revealed later, we can see that the error of $\tilde{\mathbf{w}}_*$ scales as $O(\sqrt{\frac{s}{m}} \|\mathbf{w}_*\|_2)$ in terms of sparsity s of α_* , the reduced dimensionality m and the magnitude of \mathbf{w}_* . A similar order of error bound was established in (Zhang et al., 2014) assuming \mathbf{w}_* is s -sparse and X is approximately low rank. In contrast, we do not assume X is approximately low rank.

4. Analysis

In this section, we first provide upper bound analysis of $\frac{2}{\lambda n} \|(X^\top X - \widehat{X}^\top \widehat{X})\alpha_*\|_\infty$ and σ_s , and then present the

³<http://arxiv.org/abs/1504.03991>

proof of Theorem 1 for smooth loss functions. To facilitate our analysis, we define

$$\Delta = \frac{1}{\lambda n} (\widehat{X}^\top \widehat{X} - X^\top X) \alpha_*$$

4.1. Bounding $\|\Delta\|_\infty$

A critical condition in both Theorem 1 and Theorem 3 is $\tau > \|\Delta\|_\infty$. In order to reveal the theoretical value of τ and its implication for various randomized reduction methods, we need to bound $\|\Delta\|_\infty$. We first provide a general analysis and then study its implication for various randomized reduction methods separately. The analysis is based on the following assumption, which essentially is indicated by Johnson-Lindenstrauss (JL)-type lemmas.

Assumption 1 (A1). Let $A(\mathbf{x}) = A\mathbf{x}$ be a linear projection operator where $A \in \mathbb{R}^{m \times d}$ such that for any given $\mathbf{x} \in \mathbb{R}^d$ with a high probability $1 - \delta$, we have

$$\left| \|A\mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \right| \leq \epsilon_{A,\delta} \|\mathbf{x}\|_2^2$$

where $\epsilon_{A,\delta}$ depends on m , δ and possibly d .

With this assumption, we have the following theorem regarding the upper bound of $\|\Delta\|_\infty$.

Theorem 6. Suppose $A \in \mathbb{R}^{m \times d}$ satisfies Assumption A, then with a high probability $1 - 2\delta$ we have

$$\|\Delta\|_\infty \leq R \|\mathbf{w}_*\|_{2\epsilon_{A,\delta}/n}$$

where $R = \max_i \|\mathbf{x}_i\|_2$.

Proof.

$$\begin{aligned} \frac{1}{\lambda n} (\widehat{X}^\top \widehat{X} - X^\top X) \alpha_* &= \frac{1}{\lambda n} (X^\top A^\top A X - X^\top X) \alpha_* \\ &= \frac{1}{\lambda n} X^\top (A^\top A - I) X \alpha_* = X^\top (I - A^\top A) \mathbf{w}_* \end{aligned}$$

where we use the fact $\mathbf{w}_* = -\frac{1}{\lambda n} X \alpha_*$. Then

$$\frac{1}{\lambda n} [(\widehat{X}^\top \widehat{X} - X^\top X) \alpha_*]_i = \mathbf{x}_i^\top (I - A^\top A) \mathbf{w}_*$$

Therefore in order to bound $\|\Delta\|_\infty$, we need to bound $\mathbf{x}_i^\top (I - A^\top A) \mathbf{w}_*$ for all $i \in [n]$. We first bound for individual i and then apply the union bound. Let $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{w}}_*$ be normalized version of \mathbf{x}_i and \mathbf{w}_* , i.e., $\tilde{\mathbf{x}}_i = \mathbf{x}_i / \|\mathbf{x}_i\|_2$ and $\tilde{\mathbf{w}}_* = \mathbf{w}_* / \|\mathbf{w}_*\|_2$. Suppose Assumption A is satisfied, then with a probability $1 - \delta$,

$$\begin{aligned} \tilde{\mathbf{x}}_i^\top A^\top A \tilde{\mathbf{w}}_* - \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{w}}_* &= \frac{\|A(\tilde{\mathbf{x}}_i + \tilde{\mathbf{w}}_*)\|_2^2 - \|A(\tilde{\mathbf{x}}_i - \tilde{\mathbf{w}}_*)\|_2^2}{4} \\ &\quad - \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{w}}_* \leq \frac{\epsilon_{A,\delta}}{2} (\|\tilde{\mathbf{x}}_i\|_2^2 + \|\tilde{\mathbf{w}}_*\|_2^2) \leq \epsilon_{A,\delta} \end{aligned}$$

Similarly with a probability $1 - \delta$,

$$\begin{aligned} \tilde{\mathbf{x}}_i^\top A^\top A \tilde{\mathbf{w}}_* - \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{w}}_* &= \frac{\|A(\tilde{\mathbf{x}}_i + \tilde{\mathbf{w}}_*)\|_2^2 - \|A(\tilde{\mathbf{x}}_i - \tilde{\mathbf{w}}_*)\|_2^2}{4} \\ &\quad - \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{w}}_* \geq -\frac{\epsilon_{A,\delta}}{2} (\|\tilde{\mathbf{x}}_i\|_2^2 + \|\tilde{\mathbf{w}}_*\|_2^2) \geq -\epsilon_{A,\delta} \end{aligned}$$

Therefore with a probability $1 - 2\delta$, we have

$$\begin{aligned} |\mathbf{x}_i^\top A^\top A \mathbf{w}_* - \mathbf{x}_i^\top \mathbf{w}_*| &\leq \|\mathbf{x}_i\|_2 \|\mathbf{w}_*\|_2 |\tilde{\mathbf{x}}_i^\top A^\top A \tilde{\mathbf{w}}_* - \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{w}}_*| \leq \|\mathbf{x}_i\|_2 \|\mathbf{w}_*\|_{2\epsilon_{A,\delta}} \end{aligned}$$

Then applying union bound, we complete the proof.

Next, we discuss four classes of randomized reduction operators, namely random projection, randomized Hadamard transform, random hashing and random sampling, and study the corresponding $\epsilon_{A,\delta}$ and their implications for the recovery error.

Random Projection. Random projection has been employed widely for dimension reduction. The projection operator A is usually sampled from sub-Gaussian distributions with mean 0 and variance $1/m$, e.g., (i) Gaussian distribution: $A_{ij} \sim \mathcal{N}(0, 1/m)$, (ii) Rademacher distribution: $\Pr(A_{ij} = \pm 1/\sqrt{m}) = 0.5$, (iii) discrete distribution: $\Pr(A_{ij} = \pm\sqrt{3/m}) = 1/6$ and $\Pr(A_{ij} = 0) = 2/3$. The last two distributions for dimensionality reduction were proposed and analyzed in (Achlioptas, 2003). The following lemma is the general JL-type lemma for A with sub-Gaussian entries, which reveals the value of $\epsilon_{A,\delta}$ in Assumption A.

Lemma 1. (Nelson) Let $A \in \mathbb{R}^{m \times d}$ be a random matrix with subGaussian entries of mean 0 and variance $1/m$. For any given \mathbf{x} with a probability $1 - \delta$, we have

$$\left| \|A\mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \right| \leq c \sqrt{\frac{\log(1/\delta)}{m}} \|\mathbf{x}\|_2^2$$

where c is some small universal constant.

Randomized Hadamard Transform. Randomized Hadamard transform was introduced to speed-up random projection, reducing the computational time⁴ of random projection from $O(dm)$ to $O(d \log d)$ or even $O(d \log m)$. The projection matrix A is of the form $A = PHD$, where

- $D \in \mathbb{R}^{d \times d}$ is a diagonal matrix with $D_{ii} = \pm 1$ with equal probabilities.
- H is the $d \times d$ Hadamard matrix (assuming d is a power of 2), scaled by $1/\sqrt{d}$.
- $P \in \mathbb{R}^{m \times d}$ is typically a sparse matrix that facilitates computing $P\mathbf{x}$. Several choices of P are possible (Nelson; Ailon & Chazelle, 2009; Tropp, 2011). Below we provide a JL-type lemma for a randomized Hadamard transform with $P \in \mathbb{R}^{m \times d}$ that samples m coordinates from $\sqrt{\frac{d}{m}} HD\mathbf{x}$ with replacement.

Lemma 2. (Nelson) Let $A = \sqrt{\frac{d}{m}} PHD \in \mathbb{R}^{m \times d}$ be a randomized Hadamard transform with P being a random sampling matrix. For any given \mathbf{x} with a probability $1 - \delta$, we have

$$\left| \|A\mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \right| \leq c \sqrt{\frac{\log(1/\delta) \log(d/\delta)}{m}} \|\mathbf{x}\|_2^2$$

where c is some small universal constant.

Remark 6: Compared to random projection, there is an additional $\sqrt{\log(d/\delta)}$ factor in $\epsilon_{A,\delta}$. However, it can

⁴refers to the running time of computing $A\mathbf{x}$.

be removed by applying an additional random projection. In particular, if we let $A = \sqrt{\frac{d}{m}} P' P H D \in \mathbb{R}^{m \times d}$, where $P \in \mathbb{R}^{t \times d}$ is a random sampling matrix with $t = m \log(d/\delta)$ and $P' \in \mathbb{R}^{m \times t}$ is a random projection matrix that satisfies Lemma 1, then we have the same order of $\epsilon_{A,\delta}$. Please refer to (Nelson) for more details.

Random Hashing. Another line of work to speed-up random projection is random hashing which makes the projection matrix A much sparser and takes advantage of the sparsity of feature vectors. It was introduced in (Shi et al., 2009b) for dimensionality reduction and later was improved to an unbiased version by (Weinberger et al., 2009) with some theoretical analysis. Dasgupta et al. (2010) provided a rigorous analysis of the unbiased random hashing. Recently, Kane & Nelson (2014) proposed two new random hashing algorithms with a slightly sparser random matrix A . Here we provide a JL-type lemma for the random hashing algorithm in (Weinberger et al., 2009; Dasgupta et al., 2010). Let $h : \mathbb{N} \rightarrow [m]$ denote a random hashing function, and $\xi = (\xi_1, \dots, \xi_d)$ denote a Rademacher random variable, i.e., $\xi_i, i = 1, \dots, d$ are independent and $\xi_i \in \{1, -1\}$ with equal probabilities. The projection matrix A can be written as $A = HD$, where $D \in \mathbb{R}^{d \times d}$ is a diagonal matrix with $D_{jj} = \xi_j$, and $H \in \mathbb{R}^{m \times d}$ with $H_{ij} = \delta_{i,h(j)}$ ⁵. Under the random matrix A , the feature vector $\mathbf{x} \in \mathbb{R}^d$ is reduced to $\hat{\mathbf{x}} \in \mathbb{R}^m$, where $[\hat{\mathbf{x}}]_i = \sum_{j:h(j)=i} [\mathbf{x}]_j \xi_j$. The following JL-type Lemma is a basic result from (Dasgupta et al., 2010) with a rephrasing.

Lemma 3. *Let $A = HD \in \mathbb{R}^{m \times d}$ be a random hashing matrix. For any given vector $\mathbf{x} \in \mathbb{R}^d$ such that $\frac{\|\mathbf{x}\|_\infty}{\|\mathbf{x}\|_2} \leq \frac{1}{\sqrt{c}}$, for $\delta < 0.1$, with a probability $1 - 3\delta$, we have*

$$\|A\mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \leq \sqrt{\frac{12 \log(1/\delta)}{m}} \|\mathbf{x}\|_2^2$$

where $c = 8\sqrt{m/3} \log^{1/2}(1/\delta) \log^2(m/\delta)$.

Remark 7: Compared to random projection, there is an additional condition on the feature vector $\|\mathbf{x}\|_\infty \leq \frac{\|\mathbf{x}\|_2}{\sqrt{c}}$. However, it can be removed by applying an extra preconditioner P to \mathbf{x} before applying the projection matrix A , i.e., $\hat{\mathbf{x}} = HDP\mathbf{x}$. Two preconditioners were discussed in (Dasgupta et al., 2010), with one corresponding to duplicating \mathbf{x} c times and scaling it by $1/\sqrt{c}$ and another one given by $P \in \mathbb{R}^{d \times d}$ which consists of d/b diagonal blocks of $b \times b$ randomized Hadamard matrix, where $b = 6c \log(3c/\delta)$. The running time of the reduction using the later preconditioner is $O(d \log c \log \log c)$.

Random Sampling. Last we discuss random sampling and compare with the aforementioned randomized reduction methods. In fact, the JL-type lemma for random sam-

pling is implicit in the proof of Lemma 2. We make it explicit in the following lemma.

Lemma 4. *Let $A = \sqrt{\frac{d}{m}} P \in \mathbb{R}^{m \times d}$ be a scaled random sampling matrix where $P \in \mathbb{R}^{m \times d}$ samples m coordinates with replacement. Then with a probability $1 - \delta$, we have*

$$\|A\mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \leq \frac{\|\mathbf{x}\|_\infty}{\|\mathbf{x}\|_2} \sqrt{\frac{3d \log(1/\delta)}{m}} \|\mathbf{x}\|_2^2$$

Remark 8: Compared with other three randomized reduction methods, there is an additional $\frac{\|\mathbf{x}\|_\infty}{\|\mathbf{x}\|_2} \sqrt{d}$ factor in $\epsilon_{A,\delta}$, which could result in a much larger $\epsilon_{A,\delta}$ and consequently a larger recovery error. That is why the randomized Hadamard transform was introduced to make this additional factor close to a constant.

From the above discussions, we can conclude that with random projection, randomized Hadamard transform and random hashing, with a probability $1 - \delta$ we have,

$$\begin{aligned} \|\Delta\|_\infty &= \max_i |\mathbf{x}_i^\top (I - A^\top A) \mathbf{w}_*| \\ &\leq cR \sqrt{\frac{\log(n/\delta)}{m}} \|\mathbf{w}_*\|_2. \end{aligned}$$

which essentially indicates that $\tau \geq 2cR \sqrt{\frac{\log(n/\delta)}{m}} \|\mathbf{w}_*\|_2$.

4.2. Bounding σ_s for non-smooth case

Another condition in Theorem 3 is to require $\sigma_{16s} \leq \rho_{16s}^-$. Since ρ_{16s}^- is dependent on the data, we provide an upper bound of σ_{16s} to further understand the condition. In the following analysis, we assume $\epsilon_{A,\delta} = O(\sqrt{\frac{\log(1/\delta)}{m}})$. Recall the definition of σ_s :

$$\sigma_s = \sup_{\alpha \in \mathcal{K}_{n,s}} \frac{|\alpha^\top (X^\top X - \hat{X}^\top \hat{X}) \alpha|}{n}. \quad (15)$$

We provide a bound of σ_s below.

The key idea is to use the convex relaxation of $\mathcal{K}_{n,s}$. Define $\mathcal{S}_{n,s} = \{\alpha \in \mathbb{R}^n : \|\alpha\|_2 \leq 1, \|\alpha\|_0 \leq s\}$. It was shown in (Plan & Vershynin, 2011) that $\text{conv}(\mathcal{S}_{n,s}) \subset \mathcal{K}_{n,s} \subset 2\text{conv}(\mathcal{S}_{n,s})$, where $\text{conv}(\mathcal{S})$ is the convex hull of the set \mathcal{S} . It is not difficult to show that (see the supplement)

$$\begin{aligned} &\max_{\alpha \in \mathcal{K}_{n,s}} |(X\alpha)^\top (I - A^\top A)(X\alpha)| \\ &\leq 4 \max_{\alpha_1, \alpha_2 \in \mathcal{S}_{n,s}} |(X\alpha_1)^\top (I - A^\top A)(X\alpha_2)| \end{aligned}$$

Let $\mathbf{u}_1 = X\alpha_1$ and $\mathbf{u}_2 = X\alpha_2$. For any fixed $\alpha_1, \alpha_2 \in \mathcal{S}_{n,s}$, with a probability $1 - \delta$ we can have

$$\frac{1}{n} |(X\alpha_1)^\top (I - A^\top A)(X\alpha_2)| = O\left(\rho_s^+ \sqrt{\frac{\log(1/\delta)}{m}}\right)$$

where we use

$$\max_{\alpha \in \mathcal{S}_{n,s}} \frac{\|X\alpha\|_2^2}{n} \leq \max_{\alpha \in \mathcal{K}_{n,s}} \frac{\|X\alpha\|_2^2}{n} = \rho_s^+$$

Then by using Lemma 3.3 in (Plan & Vershynin, 2011) about the entropy of $\mathcal{S}_{n,s}$ and the union bound, we can ar-

⁵ $\delta_{ij} = 1$ if $i = j$, and 0 otherwise.

rive at the following upper bound for σ_s .

Theorem 7. *With a probability $1 - \delta$, we have*

$$\sigma_s \leq O\left(\rho_s^+ \sqrt{\frac{(\log(1/\delta) + s \log(n/s))}{m}}\right)$$

Remark 9: With above result, we can further understand the condition $\sigma_{16s} \leq \rho_{16s}^-$, which amounts to

$$O\left(\rho_{16s}^+ \sqrt{\frac{(\log(1/\delta) + s \log(n/s))}{m}}\right) \leq \rho_{16s}^-$$

i.e., $m \geq \Omega(\kappa_{16s}^2 (\log(1/\delta) + s \log(n/s)))$ where $\kappa_{16s} = \rho_{16s}^+ / \rho_{16s}^-$ is the restricted condition number of the data matrix.

4.3. Proof Sketch of Theorem 1

We present a proof sketch of Theorem 1. Due to limitation of space, other proofs are provided in the supplement. Let $\hat{F}(\alpha)$ be defined as

$$\hat{F}(\alpha) = \frac{1}{n} \sum_{i=1}^n \ell_i^*(\alpha_i) + \frac{1}{2\lambda n^2} \alpha^T \hat{X}^T \hat{X} \alpha + \frac{\tau}{n} \|\alpha\|_1$$

Since $\tilde{\alpha}_* = \arg \min \hat{F}(\alpha)$ therefore for any $g_* \in \partial \|\alpha_*\|_1$

$$\begin{aligned} 0 &\geq \hat{F}(\tilde{\alpha}_*) - \hat{F}(\alpha_*) \\ &\geq (\tilde{\alpha}_* - \alpha_*)^\top \left(\frac{1}{n} \nabla \ell^*(\alpha_*) + \frac{1}{\lambda n^2} \hat{X}^\top \hat{X} \alpha_* \right) \\ &\quad + \frac{\tau}{n} (\tilde{\alpha}_* - \alpha_*)^\top g_* + \frac{1}{2nL} \|\tilde{\alpha}_* - \alpha_*\|_2^2 \end{aligned}$$

where we used the strong convexity of ℓ_i^* and its strong convexity modulus $1/L$. By the optimality condition of α_* , we can have

$$0 \geq (\alpha_* - \tilde{\alpha}_*)^\top \left(\frac{1}{n} \nabla \ell^*(\alpha_*) + \frac{1}{\lambda n^2} X^\top X \alpha_* \right) \quad (16)$$

Combining the above two inequalities we have

$$0 \geq (\tilde{\alpha}_* - \alpha_*)^\top \frac{1}{n} \Delta + \frac{\tau}{n} (\tilde{\alpha}_* - \alpha_*)^\top g_* + \frac{1}{2nL} \|\tilde{\alpha}_* - \alpha_*\|_2^2$$

Since the above inequality holds for any $g_* \in \partial \|\alpha_*\|_1$, if we choose $[g_*]_i = \text{sign}([\tilde{\alpha}_*]_i)$, $i \in \mathcal{S}^c$, then we have

$$(\tilde{\alpha}_* - \alpha_*)^\top g_* \geq -\|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1 + \|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1$$

Combining the above inequalities leads to

$$\begin{aligned} (\tau + \|\Delta\|_\infty) \|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1 &\geq (\tau - \|\Delta\|_\infty) \|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1 \\ &\quad + \frac{1}{2L} \|\tilde{\alpha}_* - \alpha_*\|_2^2 \end{aligned}$$

Assuming $\tau \geq 2\|\Delta\|_\infty$, we have

$$\begin{aligned} \|\tilde{\alpha}_* - \alpha_*\|_2^2 &\leq 3\tau L \|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1 \\ \|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1 &\leq 3\|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1 \end{aligned} \quad (17)$$

Therefore,

$$\|[\tilde{\alpha}_* - \alpha_*]_{\mathcal{S}}\|_1^2 \leq s \|\tilde{\alpha}_* - \alpha_*\|_2^2 \leq 3\tau L s \|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1$$

leading to the result

$$\|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1 \leq 3\tau L s.$$

Combing this inequality with inequalities in (17) we have

$$\|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1 \leq 9\tau L s, \quad \|\tilde{\alpha}_* - \alpha_*\|_2 \leq 3\tau L \sqrt{s}.$$

5. Numerical Experiments

In this section, we provide a case study in support of DSRR and the theoretical analysis, and a demonstration of the application of DSRR to distributed optimization.

A case study on text classification. We use the RCV1-binary data (Lewis et al., 2004) to conduct a case study. The data contains 697, 641 documents and 47, 236 features. We use a splitting 677, 399/20, 242 for training and testing. The feature vectors were normalized such that the ℓ_2 norm is equal to 1. We only report the results using random hashing since it is the most efficient, while other randomized reduction methods (except for random sampling) have similar performance. For the loss function, we use both the squared hinge loss (smooth) and the hinge loss (non-smooth). We aim to examine two questions related to our analysis and motivation (i) how does the value of τ affect the recovery error? (ii) how does the number of samples m affect the recovery error?

We vary the value of τ among 0, 0.1, 0.2, \dots , 0.9, the value of m among 1024, 2048, 4096, 8192, and the value of λ among 0.001, 0.00001. Note that $\tau = 0$ corresponds to the randomized reduction approach without the sparse regularizer. The results averaged over 5 random trials are shown in Figure 1 for the squared hinge loss and in Figure 2 for the hinge loss. We first analyze the results in Figure 1. We can observe that when τ increases the ratio of $\frac{\|[\tilde{\alpha}_*]_{\mathcal{S}^c}\|_1}{\|[\tilde{\alpha}_*]_{\mathcal{S}} - [\alpha_*]_{\mathcal{S}}\|_1}$ decreases indicating that the magnitude of dual variables for the original non-support vectors decreases. This is intuitive and consistent with our motivation. The recovery error of the dual solution (middle) first decreases and then increases. This can be partially explained by the theoretical result in Theorem 1. When the value of τ becomes larger than a certain threshold making $\tau > \|\Delta\|_\infty$ hold, then Theorem 1 implies that a larger τ will lead to a larger error. On the other hand, when τ is less than the threshold, the dual recovery error will decrease as τ increases. In addition, the figures exhibit that the thresholds for larger m are smaller which is consistent with our analysis of $\|\Delta\|_\infty = O(\sqrt{1/m})$. The difference between $\lambda = 0.001$ and $\lambda = 0.00001$ is because that smaller λ will lead to larger $\|\mathbf{w}_*\|_2$. In terms of the hinge loss, we observe similar trends, however, the recovery is much more difficult than that for squared hinge loss especially when the value of λ is small.

An application to distributed learning. Although in some cases the solution learned in the reduced space can provide sufficiently good performance, it usually performs worse than the optimal solution that solves the original problem and sometimes the performance gap between them can not be ignored as seen in following experiments. To address this issue, we combine the benefits of distributed learning and the proposed randomized reduction methods for solv-

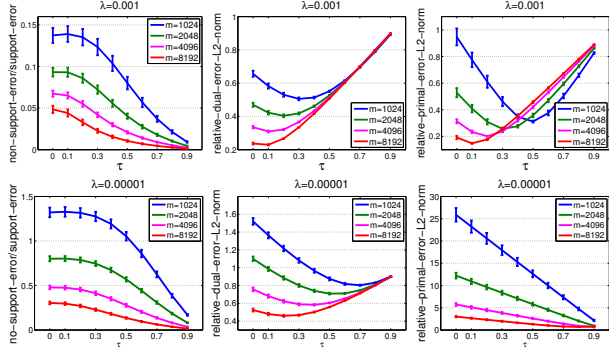


Figure 1. Recovery error for squared hinge loss. From left to right: $\frac{\|[\tilde{\alpha}_*]_{S^c}\|_1}{\|[\tilde{\alpha}_*]_S - [\alpha_*]_S\|_1}$ vs τ , $\frac{\|\tilde{\alpha}_* - \alpha_*\|_2}{\|\alpha_*\|_2}$ vs τ , and $\frac{\|\tilde{\mathbf{w}}_* - \mathbf{w}_*\|_2}{\|\mathbf{w}_*\|_2}$ vs τ .

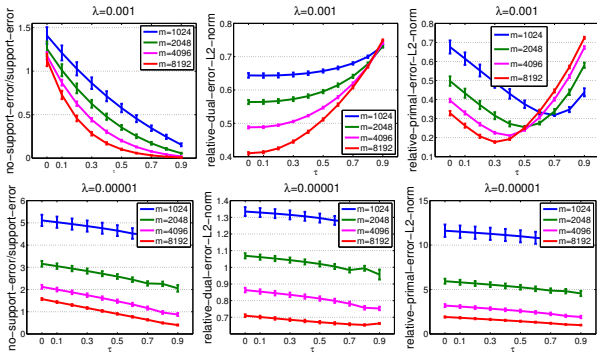


Figure 2. Same curves as above but for non-smooth hinge loss.

ing big data problems. When data is too large and sits on multiple machines, distributed learning can be employed to solve the optimization problem. In distributed learning, individual machines iteratively solve sub-problems associated with the subset of data on them and communicate some global variables (e.g., the primal solution $\mathbf{w} \in \mathbb{R}^d$) among them. When the dimensionality d is very large, the total communication cost could be very high. To reduce the total communication cost, we propose to first solve the reduced data problem and then use the found solution as the initial solution to the distributed learning for the original data.

Below, we demonstrate the effectiveness of DSRR for the recently proposed distributed stochastic dual coordinate ascent (DisDCA) algorithm (Yang, 2013). The procedure is (1) reduce original high-dimensional data to very low dimensional space on individual machines; (2) use DisDCA to solve the reduced problem; (3) use the optimal dual solution to the reduce problem as an initial solution to DisDCA for solving the original problem. We record the running time for randomized reduction in step 1 and optimization of the reduced problem in step 2, and the optimization of the original problem in step 3. We compare the performance of four methods (i) the **DSRR** method that uses the model of the reduced problem solved by DisDCA to make predictions, (ii) the method that uses the recovered model

Table 1. Statistics of datasets

Name	#Training	#Testing	#Features	#Nodes
RCV1	677,399	20,242	47,236	5
KDD	8,407,752	748,401	29,890,095	10

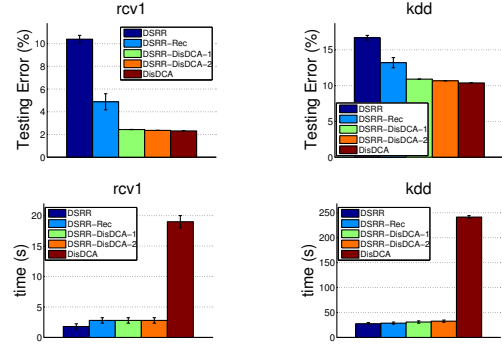


Figure 3. Top: Testing error for different methods. Bottom: Training time for different methods. The value of $\lambda = 10^{-5}$ and the value of $\tau = 0.9$. The high-dimensional features are reduced to $m = 1024$ -dimensional space using random hashing. The loss function is the squared hinge loss.

in the original space, referred to as **DSRR-Rec**; (iii) the method that uses the dual solution to the reduced problem as an initial solution of DisDCA and runs it for the original problem with $k = 1$ or 2 communications (the number of updates before each communication is set to the number of examples in each machine), referred to as **DSRR-DisDCA- k** ; and (iv) the distributed method that directly solves the original problem by **DisDCA**. For DisDCA to solve the original problem, we stop running when its performance on the testing data does not improve. Two data sets are used, namely RCV1-binary, KDD 2010 Cup data. For KDD 2010 Cup data, we use the one available on LibSVM data website. The statistics of the two data sets are summarized in Table 1. The results averaged over 5 trials are shown in Figure 3, which exhibit that the performance of DSRR-DisDCA-1/2 is remarkable in the sense that it achieves almost the same performance of directly training on the original data (DisDCA) and uses much less training time. In addition, DSRR-DisDCA performs much better than DSRR and has small computational overhead.

6. Conclusions

In this paper, we have proposed dual-sparse regularized randomized reduction methods for classification. We presented rigorous theoretical analysis of the proposed dual-sparse randomized reduction methods in terms of recovery error under a mild condition that the optimal dual variable is (nearly) sparse for both smooth and non-smooth loss functions, and for various randomized reduction approaches. The numerical experiments validate our theoretical analysis and also demonstrate that the proposed reduction and recovery framework can benefit distributed optimization by providing a good initial solution.

Acknowledgements

The authors would like to thank the anonymous reviewers for their helpful and insightful comments. T. Yang was supported in part by NSF (IIS-1463988). R. Jin was partially supported by NSF IIS-1251031 and ONR N000141410631.

References

- Achlioptas, Dimitris. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of Computer and System Sciences.*, 66:671–687, 2003.
- Agarwal, Alekh, Chappelle, Olivier, Dudk, Miroslav, and Langford, John. A reliable effective terascale linear learning system. *CoRR*, 2011.
- Ailon, Nir and Chazelle, Bernard. The fast johnson–lindenstrauss transform and approximate nearest neighbors. *SIAM J. Comput.*, 39(1):302–322, 2009.
- Balcan, Maria-Florina, Blum, Avrim, and Vempala, Santosh. Kernels as features: On kernels, margins, and low-dimensional mappings. *Machine Learning*, 65:79–94, 2006.
- Bartz, Daniel, Hatrick, Kerr, Hesse, Christian W., Müller, Klaus-Robert, and Lemm, Steven. Directional variance adjustment: improving covariance estimates for high-dimensional portfolio optimization. *arXiv:1109.3069*, 2011.
- Ben-Hur, Asa, Ong, Cheng Soon, Sonnenburg, Sören, Schölkopf, Bernhard, and Rätsch, Gunnar. Support vector machines and kernels for computational biology. *PLoS Comput Biology*, 4:e1000173, 2008.
- Bickel, Peter J., Ritov, Ya’acov, and Tsybakov, Alexandre B. Simultaneous analysis of lasso and dantzig selector. *ANNALS OF STATISTICS*, 37(4), 2009.
- Blum, Avrim. Random projection, margins, kernels, and feature-selection. In *Proceedings of the 2005 International Conference on Subspace, Latent Structure and Feature Selection*, volume 3940, pp. 52–68. Springer, 2005.
- Boyd, Stephen and Vandenberghe, Lieven. *Convex Optimization*. Cambridge University Press, 2004.
- Brank, Janez, Grobelnik, Marko, Milić-Frayling, Natasa, and Mladenić, D. Feature selection using support vector machines. In *Proceedings of the International Conference on Data Mining Methods and Databases for Engineering, Finance, and Other Fields*, pp. 84–89, 2002.
- Dasgupta, Anirban, Kumar, Ravi, and Sarlós, Tamás. A sparse johnson–lindenstrauss transform. In *Proceedings of the 42nd ACM symposium on Theory of computing, STOC ’10*, pp. 341–350, 2010.
- Drineas, Petros, Mahoney, Michael W., and Muthukrishnan, S. Sampling algorithms for l2 regression and applications. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1127–1136, 2006.
- Drineas, Petros, Mahoney, Michael W., and Muthukrishnan, S. Relative-error cur matrix decompositions. *SIAM Journal Matrix Analysis Applications*, 30:844–881, 2008.
- Drineas, Petros, Mahoney, Michael W., Muthukrishnan, S., and Sarlós, Tamás. Faster least squares approximation. *Numerische Mathematik*, 117(2):219–249, February 2011.
- Eldar, Yonina C. and Kutyniok, Gitta. *Compressed Sensing: Theory and Applications*. Compressed Sensing: Theory and Applications. Cambridge University Press, 2012. ISBN 9781107005587.
- Goldberger, Jacob, Roweis, Sam, Hinton, Geoffrey, and Salakhutdinov, Ruslan. Neighbourhood components analysis. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 513–520, 2005.
- Guyon, Isabelle, Weston, Jason, Barnhill, Stephen, and Vapnik, Vladimir. Gene selection for cancer classification using support vector machines. *Machine Learning (ML)*, 46:389–422, 2002.
- Jaggi, Martin, Smith, Virginia, Takác, Martin, Terhorst, Jonathan, Krishnan, Sanjay, Hofmann, Thomas, and Jordan, Michael I. Communication-efficient distributed dual coordinate ascent. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 3068–3076, 2014.
- Jostins, Luke and Barrett, Jeffrey C. Genetic risk prediction in complex disease. *Human Molecular Genetics*, 20 (R2):R182–R188, 2011.
- Kane, Daniel M. and Nelson, Jelani. Sparser johnson–lindenstrauss transforms. *Journal of the ACM*, 61:4:1–4:23, 2014.
- Kang, J., Kugathasan S. Georges M. Zhao H. and Cho, J.H. Improved risk prediction for crohn’s disease with a multi-locus approach. *Human Molecular Genetics*, 20: 2435–2442, 2011.
- Lewis, David D., Yang, Yiming, Rose, Tony G., and Li, Fan. Rcv1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research (JMLR)*, 5:361–397, 2004.

- Li, Mu, Andersen, David G, Smola, Alex J, and Yu, Kai. Communication efficient distributed machine learning with the parameter server. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 19–27. 2014.
- Ma, Ping, Mahoney, Michael W., and Yu, Bin. A statistical perspective on algorithmic leveraging. In *Proceedings of the 31th International Conference on Machine Learning (ICML)*, pp. 91–99, 2014.
- Mahoney, Michael W. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2):123–224, 2011.
- Mahoney, Michael W. and Drineas, Petros. Cur matrix decompositions for improved data analysis. *Proceedings of the National Academy of Sciences*, 106(3):697–702, 2009.
- Mitchell, Tom M., Hutchinson, Rebecca, Niculescu, Radu S., Pereira, Francisco, Wang, Xuerui, Just, Marcel, and Newman, Sharlene. Learning to decode cognitive states from brain images. *Machine Learning*, 57(1-2): 145–175, 2004.
- Nelson, Jelani. Johnson-lindenstrauss notes. Technical report, MIT.
- Paul, Saurabh, Boutsidis, Christos, Magdon-Ismael, Malik, and Drineas, Petros. Random projections for support vector machines. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 498–506, 2013.
- Plan, Yaniv and Vershynin, Roman. One-bit compressed sensing by linear programming. *CoRR*, abs/1109.4299, 2011.
- Räsch, G., Sonnenburg, S., and Schölkopf, B. RASE: recognition of alternatively spliced exons in c.elegans. *Bioinformatics*, 21:i369–i377, 2005.
- Rtsch, Gunnar, Sonnenburg, Sren, and Schölkopf, Bernhard. Rase: recognition of alternatively spliced exons in c.elegans. In *Proceedings of the International Conference on Intelligent Systems for Molecular Biology (ISMB) (Supplement of Bioinformatics)*, pp. 369–377, 2005.
- Sánchez, Jorge, Perronnin, Florent, Mensink, Thomas, and Verbeek, Jakob J. Image classification with the fisher vector: Theory and practice. *International Journal of Computer Vision*, 105(3):222–245, 2013.
- Shi, Qinfeng, Petterson, James, Dror, Gideon, Langford, John, Smola, Alex, and Vishwanathan, S.V.N. Hash kernels for structured data. *Journal of Machine Learning Research (JMLR)*, 10:2615–2637, 2009a.
- Shi, Qinfeng, Petterson, James, Dror, Gideon, Langford, John, Smola, Alexander J., Strehl, Alexander L., and Vishwanathan, Vishy. Hash kernels. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 496–503, 2009b.
- Shi, Qinfeng, Shen, Chunhua, Hill, Rhys, and van den Hengel, Anton. Is margin preserved after random projection? In *Proceedings of the International Conference on Machine Learning (ICML)*, 2012.
- Simianer, Patrick, Riezler, Stefan, and Dyer, Chris. Joint feature selection in distributed stochastic learning for large-scale discriminative training in smt. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 11–21, 2012.
- Sonnenburg, Sren and Franc, Vojtech. Coffin: A computational framework for linear svms. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 999–1006, 2010.
- Sonnenburg, Sren, Schweikert, Gabriele, Philips, Petra, Behr, Jonas, and Rtsch, Gunnar. Accurate splice site prediction using support vector machines. *BMC Bioinformatics*, 8, 2007.
- Tropp, Joel A. Improved analysis of the subsampled randomized hadamard transform. *Advances in Adaptive Data Analysis*, 3(1-2):115–126, 2011.
- Weinberger, Kilian Q., Dasgupta, Anirban, Langford, John, Smola, Alexander J., and Attenberg, Josh. Feature hashing for large scale multitask learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1113–1120, 2009.
- Xiao, Lin and Zhang, Tong. A proximal-gradient homotopy method for the sparse least-squares problem. *SIAM Journal on Optimization*, 23(2):1062–1091, 2013.
- Yang, Tianbao. Trading computation for communication: Distributed stochastic dual coordinate ascent. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 629–637, 2013.
- Zhang, Lijun, Mahdavi, Mehrdad, Jin, Rong, Yang, Tianbao, and Zhu, Shenghuo. Recovering the optimal solution by dual random projection. In *Proceedings of Conference of Learning Theory*, pp. 135–157, 2013.
- Zhang, Lijun, Mahdavi, Mehrdad, Jin, Rong, Yang, Tianbao, and Zhu, Shenghuo. Random projections for classification: A recovery approach. *IEEE Transactions on Information Theory (IEEE TIT)*, 60(11):7300–7316, 2014.