

Recovering the Optimal Solution by Dual Random Projection

Lijun Zhang

Mehrdad Mahdavi

Rong Jin

Department of Computer Science and Engineering

Michigan State University, East Lansing, MI 48824, USA

Tianbao Yang

GE Global Research, San Ramon, CA 94583, USA

Shenghuo Zhu

NEC Laboratories America, Cupertino, CA 95014, USA

ZHANGLIJ@MSU.EDU

MAHDAVIM@MSU.EDU

RONGJIN@CSE.MSU.EDU

TYANG@GE.COM

ZSH@NEC-LABS.COM

Abstract

Random projection has been widely used in data classification. It maps high-dimensional data into a low-dimensional subspace in order to reduce the computational cost in solving the related optimization problem. While previous studies are focused on analyzing the classification performance of using random projection, in this work, we consider the **recovery problem**, i.e., how to accurately recover the optimal solution to the original optimization problem in the high-dimensional space based on the solution learned from the subspace spanned by random projections. We present a simple algorithm, termed **Dual Random Projection**, that uses the dual solution of the low-dimensional optimization problem to recover the optimal solution to the original problem. Our theoretical analysis shows that with a high probability, the proposed algorithm is able to accurately recover the optimal solution to the original problem, provided that the data matrix is of low rank or can be well approximated by a low rank matrix.

Keywords: Random projection, Primal solution, Dual solution, Low rank

1. Introduction

Random projection is a simple yet powerful dimensionality reduction technique that projects the original high-dimensional data onto a low-dimensional subspace using a random matrix (Kaski, 1998; Bingham and Mannila, 2001). It has been successfully applied to many machine learning tasks, including classification (Fradkin and Madigan, 2003; Vempala, 2004; Rahimi and Recht, 2008), regression (Maillard and Munos, 2012), clustering (Fern and Brodley, 2003; Boutsidis et al., 2010), manifold learning (Dasgupta and Freund, 2008; Freund et al., 2008), and information retrieval (Goel et al., 2005).

In this work, we focus on random projection for classification. While previous studies were devoted to analyzing the classification performance using random projection (Arriaga and Vempala, 1999; Balcan et al., 2006; Paul et al., 2012; Shi et al., 2012), we examine the effect of random projection from a very different aspect. In particular, we are interested in accurately recovering the optimal solution to the original high-dimensional optimization

problem using random projection. This is particularly useful for feature selection (Guyon and Elisseeff, 2003), where important features are often selected based on their weights in the linear prediction model learned from the training data. In order to ensure that similar features are selected, the prediction model based on random projection needs to be close to the model obtained by solving the original optimization problem directly.

The proposed algorithm for recovering the optimal solution consists of two simple steps. In the first step, similar to previous studies, we apply random projection to reducing the dimensionality of the data, and then solve a low-dimensional optimization problem. In the second step, we construct the dual solution of the low-dimensional problem from its primal solution, and then use it to recover the optimal solution to the original high-dimensional problem. Our analysis reveals that with a high probability, we are able to recover the optimal solution with a small error by using $\Omega(\mathbf{r} \log \mathbf{r})$ projections, where \mathbf{r} is the rank of the data matrix. A similar result also holds when the data matrix can be well approximated by a low rank matrix. We further show that the proposed algorithm can be applied iteratively to recovering the optimal solution with a relative error by using $\mathbf{O}(\log 1/\epsilon)$ iterations.

The rest of the paper is arranged as follows. Section 2 describes the problem of recovering optimal solution by random projection, the theme of this work. Section 3 describes the dual random projection approach for recovering the optimal solution. Section 4 presents the main theoretical results for the proposed algorithm. Section 5 presents the proof for the theorems stated in Section 4. Section 6 concludes this work.

2. The Problem

Let $(\mathbf{x}_i, \mathbf{y}_i), i = 1, \dots, n$ be a set of training examples, where $\mathbf{x}_i \in \mathbb{R}^d$ is a vector of d dimension and $\mathbf{y}_i \in \{-1, +1\}$ is the binary class assignment for \mathbf{x}_i . Let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ and $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)^\top$ include input patterns and the class assignments of all training examples. A classifier $\mathbf{w} \in \mathbb{R}^d$ is learned by solving the following optimization problem:

$$\min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{w}\|^2 + \sum_{i=1}^n (\mathbf{y}_i \mathbf{x}_i^\top \mathbf{w}), \quad (1)$$

where $\phi(\mathbf{z})$ is a convex loss function that is differentiable¹. By writing $\phi(\mathbf{z})$ in its convex conjugate form, i.e.,

$$\phi(\mathbf{z}) = \max_{\alpha \in \Omega} \mathbf{z}^\top \alpha - \psi(\alpha),$$

where $\psi(\alpha)$ is the convex conjugate of $\phi(\mathbf{z})$ and Ω is the domain of the dual variable, we get the dual optimization problem:

$$\max_{\alpha \in \Omega^n} - \sum_{i=1}^n \psi(\alpha_i) - \frac{1}{2} \alpha^\top \mathbf{G} \alpha, \quad (2)$$

where $\alpha = (\alpha_1, \dots, \alpha_n)^\top$ and $\mathbf{D}(\mathbf{y}) = \text{diag}(\mathbf{y})$ and \mathbf{G} is the Gram matrix given by

$$\mathbf{G} = \mathbf{D}(\mathbf{y}) \mathbf{X}^\top \mathbf{X} \mathbf{D}(\mathbf{y}). \quad (3)$$

1. For non-differentiable loss functions such as hinge loss, we could apply the smoothing technique described in (Nesterov, 2005) to make it differentiable.

In the following, we denote by $\mathbf{w}_* \in \mathbb{R}^d$ the optimal primal solution to (1), and by $\boldsymbol{\alpha}_* \in \mathbb{R}^n$ the optimal dual solution to (2). The following proposition connects \mathbf{w}_* and $\boldsymbol{\alpha}_*$.

Proposition 1 *Let $\mathbf{w}_* \in \mathbb{R}^d$ be the optimal primal solution to (1), and $\boldsymbol{\alpha}_* \in \mathbb{R}^n$ be the optimal dual solution to (2). We have*

$$\mathbf{w}_* = -\frac{1}{\mathbf{D}(\mathbf{y})} \mathbf{X} \mathbf{D}(\mathbf{y}) \boldsymbol{\alpha}_*, \text{ and } [\boldsymbol{\alpha}_*]_i = \nabla \left(\mathbf{y}_i \mathbf{x}_i^\top \mathbf{w}_* \right), \quad i = 1, \dots, n. \quad (4)$$

The proof of Proposition 1 and other omitted proofs are deferred to the Appendix. When the dimensionality \mathbf{d} is high and the number of training examples \mathbf{n} is large, solving either the primal problem in (1) or the dual problem in (2) can be computationally expensive. To reduce the computational cost, one common approach is to significantly reduce the dimensionality by random projection. Let $\mathbf{R} \in \mathbb{R}^{d \times m}$ be a Gaussian random matrix, where each entry $\mathbf{R}_{i,j}$ is independently drawn from a Gaussian distribution $\mathcal{N}(0, 1)$ and \mathbf{m} is significantly smaller than \mathbf{d} . Using the random matrix \mathbf{R} , we generate a low-dimensional representation for each input example by

$$\hat{\mathbf{x}}_i = \frac{1}{\sqrt{\mathbf{m}}} \mathbf{R}^\top \mathbf{x}_i, \quad (5)$$

and solve the following low-dimensional optimization problem:

$$\min_{\mathbf{z} \in \mathbb{R}^m} \frac{1}{2} \|\mathbf{z}\|^2 + \sum_{i=1}^n (\mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i). \quad (6)$$

The corresponding dual problem is written as

$$\max_{\boldsymbol{\alpha} \in \Omega^n} - \sum_{i=1}^n \mathbf{y}_i \hat{\mathbf{x}}_i^\top \boldsymbol{\alpha} - \frac{1}{2} \boldsymbol{\alpha}^\top \hat{\mathbf{G}} \boldsymbol{\alpha}, \quad (7)$$

where

$$\hat{\mathbf{G}} = \mathbf{D}(\mathbf{y}) \mathbf{X}^\top \frac{\mathbf{R} \mathbf{R}^\top}{\mathbf{m}} \mathbf{X} \mathbf{D}(\mathbf{y}). \quad (8)$$

Intuitively, the choice of Gaussian random matrix \mathbf{R} is justified by the expectation of the dot-product between any two examples in the projected space is equal to the dot-product in the original space, i.e.,

$$\mathbb{E}[\hat{\mathbf{x}}_i^\top \hat{\mathbf{x}}_j] = \mathbf{x}_i^\top \mathbb{E} \left[\frac{1}{\mathbf{m}} \mathbf{R} \mathbf{R}^\top \right] \mathbf{x}_j = \mathbf{x}_i^\top \mathbf{x}_j,$$

where the last equality follows from $\mathbb{E}[\mathbf{R} \mathbf{R}^\top / \mathbf{m}] = \mathbf{I}$. Thus, $\mathbf{G} = \hat{\mathbf{G}}$ holds in expectation.

Let $\mathbf{z}_* \in \mathbb{R}^m$ denote the optimal primal solution to the low-dimensional problem (6), and $\hat{\boldsymbol{\alpha}}_* \in \mathbb{R}^n$ denote the optimal dual solution to (7). Similar to Proposition 1, the following proposition connects \mathbf{z}_* and $\hat{\boldsymbol{\alpha}}_*$.

Proposition 2 *We have*

$$\mathbf{z}_* = -\frac{1}{\sqrt{\mathbf{m}}} \mathbf{R}^\top \mathbf{X} \mathbf{D}(\mathbf{y}) \hat{\boldsymbol{\alpha}}_*, \text{ and } [\hat{\boldsymbol{\alpha}}_*]_i = \nabla \left(\frac{\mathbf{y}_i}{\sqrt{\mathbf{m}}} \mathbf{x}_i^\top \mathbf{R} \mathbf{z}_* \right), \quad i = 1, \dots, n. \quad (9)$$

Given the optimal solution $\mathbf{z}_* \in \mathbb{R}^m$, the data point $\mathbf{x} \in \mathbb{R}^d$ is classified by $\mathbf{x}^\top \mathbf{R}\mathbf{z}_*/\sqrt{\mathbf{m}}$, which is equivalent to defining a new solution $\hat{\mathbf{w}} \in \mathbb{R}^d$ given below, which we refer to as the naive solution

$$\hat{\mathbf{w}} = \frac{1}{\sqrt{\mathbf{m}}} \mathbf{R}\mathbf{z}_*. \quad (10)$$

The classification performance of $\hat{\mathbf{w}}$ has been examined by many studies (Arriaga and Vempala, 1999; Balcan et al., 2006; Paul et al., 2012; Shi et al., 2012). The general conclusion is that when the original data is linearly separable with a large margin, the classification error for $\hat{\mathbf{w}}$ is usually small.

Although these studies show that $\hat{\mathbf{w}}$ can achieve a small classification error under appropriate assumptions, it is unclear whether $\hat{\mathbf{w}}$ is a good approximation of the optimal solution \mathbf{w}_* . In fact, as we will see in Section 4, the naive solution is almost guaranteed to be a BAD approximation of the optimal solution, that is, $\|\hat{\mathbf{w}} - \mathbf{w}_*\|_2 = \Omega(\sqrt{d/m}\|\mathbf{w}_*\|_2)$. This observation leads to an interesting question: **is it possible to accurately recover the optimal solution \mathbf{w}_* based on \mathbf{z}_* , the optimal solution to the low-dimensional optimization problem?**

Relationship to Compressive Sensing The proposed problem is closely related to compressive sensing (Donoho, 2006; Candès and Wakin, 2008) where the goal is to recover a high-dimensional but sparse vector using a small number of random measurements. The key difference between our work and compressive sensing is that we do not have the direct access to the random measurement of the target vector (which in our case is $\mathbf{w}_* \in \mathbb{R}^d$). Instead, $\mathbf{z}_* \in \mathbb{R}^m$ is the optimal solution to (6), the primal problem using random projection. However, the following theorem shows that \mathbf{z}_* is a good approximation of $\mathbf{R}^\top \mathbf{w}_*/\sqrt{\mathbf{m}}$, which includes \mathbf{m} random measurements of \mathbf{w}_* , if the data matrix \mathbf{X} is of low rank and the number of random measurements \mathbf{m} is sufficiently large.

Theorem 1 *For any $0 < \epsilon \leq 1/2$, with a probability at least $1 - \exp(-\mathbf{m}/32)$, we have*

$$\|\sqrt{\mathbf{m}}\mathbf{z}_* - \mathbf{R}^\top \mathbf{w}_*\|_2 \leq \frac{\sqrt{2}}{\sqrt{1-\epsilon}} \|\mathbf{R}^\top \mathbf{w}_*\|_2,$$

provided

$$\mathbf{m} \geq \frac{(\mathbf{r} + 1) \log(2\mathbf{r}/\epsilon)}{\mathbf{c}^2},$$

where constant \mathbf{c} is at least $1/4$, and \mathbf{r} is the rank of \mathbf{X} .

Given the approximation bound in Theorem 1, it is appealing to reconstruct \mathbf{w}_* using the compressive sensing algorithm provided that \mathbf{w}_* is sparse to certain bases. We note that the low rank assumption for data matrix \mathbf{X} implies that \mathbf{w}_* is sparse with respect to the singular vectors of \mathbf{X} . However, since \mathbf{z}_* only provides an approximation to the random measurements of \mathbf{w}_* , running the compressive sensing algorithm will not be able to perfectly recover \mathbf{w}_* from \mathbf{z}_* . In Section 3, we present an algorithm, that recovers \mathbf{w}_* with a small error, provided that the data matrix \mathbf{X} is of low rank. Compared to the compressive sensing algorithm, the main advantage of the proposed algorithm is its computational simplicity; it neither computes the singular vectors of \mathbf{X} nor solves an optimization problem that minimizes the ℓ_1 norm.

Algorithm 1 A Dual Random Projection Approach for Recovering Optimal Solution

- 1: **Input:** input patterns $\mathbf{X} \in \mathbb{R}^{d \times n}$, binary class assignment $\mathbf{y} \in \{-1, +1\}^n$, and sample size \mathbf{m}
- 2: Sample a Gaussian random matrix $\mathbf{R} \in \mathbb{R}^{d \times m}$ and compute $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n] = \mathbf{R}^\top \mathbf{X} / \sqrt{\mathbf{m}}$
- 3: Obtain the primal solution $\mathbf{z}_* \in \mathbb{R}^m$ by solving the optimization problem in (6)
- 4: Construct the dual solution $\hat{\alpha}_* \in \mathbb{R}^n$ by Proposition 2, i.e.,

$$[\hat{\alpha}_*]_i = \nabla \left(\frac{\mathbf{y}_i}{\sqrt{\mathbf{m}}} \mathbf{x}_i^\top \mathbf{R} \mathbf{z}_* \right), \quad i = 1, \dots, n$$

- 5: Compute $\tilde{\mathbf{w}} \in \mathbb{R}^d$ according to (12), i.e., $\tilde{\mathbf{w}} = -\mathbf{X} \mathbf{D}(\mathbf{y}) \hat{\alpha}_* /$
 - 6: **Output:** the recovered solution $\tilde{\mathbf{w}}$
-

3. Algorithm

To motivate our algorithm, let us revisit the optimal primal solution \mathbf{w}_* to (1), which is given in Proposition 1, i.e.,

$$\mathbf{w}_* = -\frac{1}{\mathbf{m}} \mathbf{X} \mathbf{D}(\mathbf{y}) \alpha_*, \quad (11)$$

where α_* is the optimal solution to the dual problem (2). Given the projected data $\hat{\mathbf{x}} = \mathbf{R}^\top \mathbf{x} / \sqrt{\mathbf{m}}$, we have reached an approximate dual problem in (7). Comparing it with the dual problem in (2), the only difference is that the Gram matrix $\mathbf{G} = \mathbf{D}(\mathbf{y}) \mathbf{X}^\top \mathbf{X} \mathbf{D}(\mathbf{y})$ in (2) is replaced with $\hat{\mathbf{G}} = \mathbf{D}(\mathbf{y}) \mathbf{X}^\top \mathbf{R} \mathbf{R}^\top \mathbf{X} \mathbf{D}(\mathbf{y}) / \mathbf{m}$ in (7). Recall that $\mathbb{E}[\mathbf{R} \mathbf{R}^\top / \mathbf{m}] = \mathbf{I}$. Thus, when the number of random projections \mathbf{m} is sufficiently large, $\hat{\mathbf{G}}$ will be close to the \mathbf{G} and we would also expect $\hat{\alpha}_*$ to be close to α_* . As a result, we can use $\hat{\alpha}_*$ to approximate α_* in (11), which yields a recovered prediction model given below:

$$\tilde{\mathbf{w}} = -\frac{1}{\mathbf{m}} \mathbf{X} \mathbf{D}(\mathbf{y}) \hat{\alpha}_* = -\sum_{i=1}^n \frac{1}{\mathbf{m}} \mathbf{y}_i [\hat{\alpha}_*]_i \mathbf{x}_i. \quad (12)$$

Note that the key difference between the recovered solution $\tilde{\mathbf{w}}$ and the naive solution $\hat{\mathbf{w}}$ is that $\hat{\mathbf{w}}$ is computed by mapping the optimal primal solution \mathbf{z}_* in the projected space back to the original space via the random matrix \mathbf{R} , while $\tilde{\mathbf{w}}$ is computed directly in the original space using the approximate dual solution $\hat{\alpha}_*$. Therefore, the naive solution $\hat{\mathbf{w}}$ lies in the subspace spanned by the column vectors in the random matrix \mathbf{R} (denoted by \mathcal{A}_R), while the recovered solution $\tilde{\mathbf{w}}$ lies in the subspace that also contains the optimal solution \mathbf{w}_* , i.e., the subspace spanned by columns of \mathbf{X} (denoted by \mathcal{A}). The mismatch between spaces \mathcal{A}_R and \mathcal{A} leads to the large approximation error for $\hat{\mathbf{w}}$.

According to Proposition 2, we can construct the dual solution $\hat{\alpha}_*$ from the primal solution \mathbf{z}_* . Thus, we do not need to solve the dual problem in (7) to obtain $\hat{\alpha}_*$. Instead, we solve the low-dimensional optimization problem in (6) to get \mathbf{z}_* and construct $\hat{\alpha}_*$ from it. Algorithm 1 shows the details of the proposed method. We note that although dual variables have been widely used in the analysis of convex optimization (Boyd and Vandenberghe, 2004; Hazan et al., 2011) and online learning (Shalev-Shwartz and Singer, 2006), to the best of

Algorithm 2 An Iterative Dual Random Projection Approach for Recovering Optimal Solution

- 1: **Input:** input patterns $\mathbf{X} \in \mathbb{R}^{d \times n}$, binary class assignment $\mathbf{y} \in \{-1, +1\}^n$, sample size \mathbf{m} , and number of iterations \mathbf{T}
- 2: Sample a Gaussian random matrix $\mathbf{R} \in \mathbb{R}^{d \times m}$ and compute $\hat{\mathbf{X}} = \mathbf{R}^\top \mathbf{X} / \sqrt{\mathbf{m}}$
- 3: Initialize $\tilde{\mathbf{w}}^0 = \mathbf{0}$
- 4: **for** $\mathbf{t} = 1, \dots, \mathbf{T}$ **do**
- 5: Obtain $\mathbf{z}_*^t \in \mathbb{R}^m$ by solving the following optimization problem

$$\min_{\mathbf{z} \in \mathbb{R}^m} \frac{1}{2} \left\| \mathbf{z} + \frac{1}{\sqrt{\mathbf{m}}} \mathbf{R}^\top \tilde{\mathbf{w}}^{t-1} \right\|_2^2 + \sum_{i=1}^n \left(\mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) \quad (13)$$

- 6: Construct the dual solution $\hat{\boldsymbol{\alpha}}_*^t \in \mathbb{R}^n$ using

$$[\hat{\boldsymbol{\alpha}}_*^t]_i = \nabla \left(\mathbf{y}_i \hat{\mathbf{x}}_i^\top \mathbf{z}_*^t + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right), \quad i = 1, \dots, n$$

- 7: Update the solution by $\tilde{\mathbf{w}}^t = -\mathbf{X} \mathbf{D}(\mathbf{y}) \hat{\boldsymbol{\alpha}}_*^t /$
 - 8: **end for**
 - 9: **Output** the recovered solution $\tilde{\mathbf{w}}^T$
-

our knowledge, this is the first time that dual variables are used in conjunction with random projection for recovering the optimal solution.

To further reduce the recovery error, we develop an iterative method shown in Algorithm 2. The idea comes from that fact that if $\|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2 \leq \|\mathbf{w}_*\|_2$ with a small ϵ , we can apply the same dual random projection algorithm again to recover $\Delta \mathbf{w} = \mathbf{w}_* - \tilde{\mathbf{w}}$, which will result in a recovery error of $\|\Delta \mathbf{w}\|_2 \leq \epsilon^2 \|\mathbf{w}_*\|_2$. If we repeat the above process with \mathbf{T} iterations, we should be able to obtain a solution with a recovery error of $\epsilon^T \|\mathbf{w}_*\|_2$. This simple intuition leads to an iterative method shown in Algorithm 2. At the \mathbf{t} -th iteration, given the recovered solution $\tilde{\mathbf{w}}^{t-1}$ obtained from the previous iteration, we solve the optimization problem in (13) that is designed to recover $\mathbf{w}_* - \tilde{\mathbf{w}}^{t-1}$. The detailed derivation of Algorithm 2 is provided in Section 5.2.

It is important to note that although Algorithm 2 consists of multiple iterations, the random projection of the data matrix is only computed once before the start of the iterations. This important feature makes the iterative algorithm computationally attractive as calculating random projections of a large data matrix is computationally expensive and has been the subject of many studies, e.g., (Achlioptas, 2003; Liberty et al., 2008; Braverman et al., 2010). However, it is worth noting that at each iteration in Algorithm 2, we need to compute the dot-product $[\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i$ for all training data in the original space. We also note that Algorithm 2 is related to the Epoch gradient descent algorithm (Hazan and Kale, 2011) for stochastic optimization in the sense that the solution obtained from the previous iteration serves as the starting point to the optimization problem at the current iteration. Unlike the algorithm in (Hazan and Kale, 2011), we do not shrink the domain size over the iterations in Algorithm 2.

Application to the Square Loss In the following, we take the square loss $(\mathbf{z}) = \frac{1}{2}(1 - \mathbf{z})^2$ as an example to illustrate the recovery procedure. The original optimization problem, which is referred to as the ridge regression (Hastie et al., 2009), is

$$\min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{w}\|^2 + \frac{1}{2} \sum_{i=1}^n (1 - \mathbf{y}_i \mathbf{x}_i^\top \mathbf{w})^2 \stackrel{y_i \in \{\pm 1\}}{=} \frac{1}{2} \|\mathbf{w}\|^2 + \frac{1}{2} \sum_{i=1}^n (\mathbf{y}_i - \mathbf{x}_i^\top \mathbf{w})^2.$$

Setting the derivative to $\mathbf{0}$, we obtain the optimal solution

$$\mathbf{w}_* = \left(\mathbf{I} + \mathbf{X} \mathbf{X}^\top \right)^{-1} \mathbf{X} \mathbf{y} = \mathbf{X} \left(\mathbf{I} + \mathbf{X}^\top \mathbf{X} \right)^{-1} \mathbf{y}. \quad (14)$$

where the last equality follows from the Woodbury matrix identity (Golub and Van Loan, 1996). Thus, the computational cost is either $\mathcal{O}(\mathbf{d}^2 \mathbf{n} + \mathbf{d}^3)$ or $\mathcal{O}(\mathbf{n}^2 \mathbf{d} + \mathbf{n}^3)$.

Following the dual random projection algorithm, we first solve the low-dimensional problem in (6), whose solution is $\mathbf{z}_* = (\mathbf{I} + \widehat{\mathbf{X}} \widehat{\mathbf{X}}^\top)^{-1} \widehat{\mathbf{X}} \mathbf{y}$, where $\widehat{\mathbf{X}} = \mathbf{R}^\top \mathbf{X} / \sqrt{\mathbf{m}}$. Then, we construct the dual solution $\widehat{\boldsymbol{\alpha}}_* = \mathbf{D}(\mathbf{y}) \widehat{\mathbf{X}}^\top \mathbf{z}_* - \mathbf{1}$. Finally, we recover the optimal solution $\widetilde{\mathbf{w}} = -\frac{1}{\lambda} \mathbf{X} \mathbf{D}(\mathbf{y}) \widehat{\boldsymbol{\alpha}}_*$. It is straightforward to check that computational cost of our algorithm is $\mathcal{O}(\mathbf{m} \mathbf{n} \mathbf{d} + \mathbf{m}^2 \mathbf{n} + \mathbf{m}^3)$, which is significantly smaller than that of (14) when both \mathbf{d} and \mathbf{n} are large.

After some algebraic manipulation, we can show that

$$\widetilde{\mathbf{w}} = \mathbf{X} \left(\mathbf{I} + \mathbf{X}^\top \frac{\mathbf{S} \mathbf{S}^\top}{\mathbf{m}} \mathbf{X} \right)^{-1} \mathbf{y}. \quad (15)$$

Comparing (14) with (15), we can see the difference between \mathbf{w}_* and $\widetilde{\mathbf{w}}$ comes from the Gram matrix. When \mathbf{m} is large enough, $\mathbf{X}^\top \mathbf{R} \mathbf{R}^\top \mathbf{X} / \mathbf{m}$ is close to the $\mathbf{X}^\top \mathbf{X}$, and as a result $\widetilde{\mathbf{w}}$ is also close to \mathbf{w}_* .

4. Main Results

In this section, we will bound the recovery error $\|\mathbf{w}_* - \widetilde{\mathbf{w}}\|_2$ of dual random projection. We first assume \mathbf{X} is of low rank, and then extend the results to the full rank case.

4.1. Low Rank

The low rank assumption is closely related to the sparsity assumption made in compressive sensing. This is because \mathbf{w}_* lies in the subspace spanned by the column vectors of \mathbf{X} and the low rank assumption directly implies that \mathbf{w}_* is sparse with respect to the singular vectors of \mathbf{X} .

We denote by \mathbf{r} the rank of matrix \mathbf{X} . The following theorem shows that the recovery error of Algorithm 1 is small provided that (1) \mathbf{X} is of low rank (i.e., $\mathbf{r} \ll \min(\mathbf{d}, \mathbf{n})$), and (2) the number of random projections is sufficiently large.

Theorem 2 *Let \mathbf{w}_* be the optimal solution to (1) and let $\widetilde{\mathbf{w}}$ be the solution recovered by Algorithm 1. For any $0 < \epsilon \leq 1/2$, with a probability at least $1 - \epsilon$, we have*

$$\|\widetilde{\mathbf{w}} - \mathbf{w}_*\|_2 \leq \frac{\epsilon}{1 - \epsilon} \|\mathbf{w}_*\|_2,$$

provided

$$\mathbf{m} \geq \frac{(\mathbf{r} + 1) \log(2\mathbf{r}/\epsilon)}{\epsilon^2},$$

where constant ϵ is at least $1/4$.

According to Theorem 2, the number of required random projections is $\Omega(\mathbf{r} \log \mathbf{r})$. This is similar to a compressive sensing result if we view rank \mathbf{r} as the sparsity measure used in compressive sensing. Following the same arguments as compressive sensing, it may be possible to argue that $\Omega(\mathbf{r} \log \mathbf{r})$ is optimal due to the result of the coupon collector's problem (Mowani and Raghavan, 1995), although the rigorous analysis remains to be developed.

As a comparison, the following theorem shows that with a high probability, the naive solution $\hat{\mathbf{w}}$ given in (10) does not accurately recover the true optimal solution \mathbf{w}_* .

Theorem 3 *For any $0 < \epsilon \leq 1/3$, with a probability at least $1 - \exp(-(\mathbf{d} - \mathbf{r})/32) - \exp(-\mathbf{m}/32) - \epsilon$, we have*

$$\|\hat{\mathbf{w}} - \mathbf{w}_*\|_2 \geq \frac{1}{2} \sqrt{\frac{\mathbf{d} - \mathbf{r}}{\mathbf{m}}} \left(1 - \frac{\sqrt{2(1 + \epsilon)}}{1 - \epsilon} \right) \|\mathbf{w}_*\|_2,$$

provided the condition on \mathbf{m} in Theorem 2 holds.

As indicated by Theorem 3, when \mathbf{m} is sufficiently larger than \mathbf{r} but significantly smaller than \mathbf{d} , we have $\|\hat{\mathbf{w}} - \mathbf{w}_*\|_2 = \Omega(\sqrt{\mathbf{d}/\mathbf{m}} \|\mathbf{w}_*\|_2)$, indicating that $\hat{\mathbf{w}}$ does not approximate \mathbf{w}_* well.

It is important to note that Theorem 3 does not contradict the previous results showing that the random projection based method could result in a small classification error if the data is linearly separable with a large margin. This is because, to decide whether $\hat{\mathbf{w}}$ carries a similar classification performance to \mathbf{w}_* , we need to measure the following term

$$\max_{\mathbf{x} \in \text{span}(\mathbf{X}), \|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top (\hat{\mathbf{w}} - \mathbf{w}_*). \quad (16)$$

Since $\|\hat{\mathbf{w}} - \mathbf{w}_*\|_2$ can also be written as

$$\|\hat{\mathbf{w}} - \mathbf{w}_*\|_2 = \max_{\|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top (\hat{\mathbf{w}} - \mathbf{w}_*),$$

the quantity defined in (16) could be significantly smaller than $\|\hat{\mathbf{w}} - \mathbf{w}_*\|_2$ if the data matrix \mathbf{X} is of low rank. The following theorem quantifies this statement.

Theorem 4 *For any $0 < \epsilon \leq 1/2$, with a probability at least $1 - \epsilon$, we have*

$$\max_{\mathbf{x} \in \text{span}(\mathbf{X}), \|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top (\hat{\mathbf{w}} - \mathbf{w}_*) \leq \left(1 + \frac{1}{1 - \epsilon} \right) \|\mathbf{w}_*\|_2,$$

provided the condition on \mathbf{m} in Theorem 2 holds.

We note that Theorem 4 directly implies the result of margin classification error for random projection (Blum, 2006; Shi et al., 2012). This is because when a data point $(\mathbf{x}_i, \mathbf{y}_i)$ can be separated by \mathbf{w}_* with a margin γ , i.e., $\mathbf{y}_i \mathbf{w}_*^\top \mathbf{x}_i \geq \gamma \|\mathbf{w}_*\|_2$, it will be classified a mb

Theorem 5 *Let \mathbf{w}_* be the optimal solution to (1) and let $\tilde{\mathbf{w}}^T$ be the solution recovered by Algorithm 2. For any $0 < \varepsilon < 1/2$, with a probability at least $1 - \varepsilon$, we have*

$$\|\tilde{\mathbf{w}}^T - \mathbf{w}_*\|_2 \leq \left(\frac{\varepsilon}{1 - \varepsilon}\right)^T \|\mathbf{w}_*\|_2,$$

provided the condition on \mathbf{m} in Theorem 2 holds.

Notice that the number of random projection \mathbf{m} does not depend on the number of iterations T . That is because we only apply random projection once to reducing the dimensionality of the data. Theorem 5 implies that we can recover the optimal solution with a relative error ε , i.e., $\|\mathbf{w}_* - \tilde{\mathbf{w}}^T\|_2 \leq \varepsilon \|\mathbf{w}_*\|_2$, by using $\log_{(1-\varepsilon)/\varepsilon} 1/\varepsilon$ iterations.

4.2. Full Rank

If \mathbf{X} has full rank, we established the following theorem to bound the recovery error.

Theorem 6 *Assume \mathbf{w}_* lies in the subspace spanned by the first k left singular vectors of \mathbf{X} , and the loss $\ell(\cdot)$ is γ -smooth. For any $0 < \varepsilon \leq 1$, with a probability at least $1 - \varepsilon$, we have*

$$\|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2 \leq \frac{\varepsilon}{1 - \varepsilon} \left(1 + \frac{\sqrt{\lambda}}{\sqrt{\gamma} k}\right) \|\mathbf{w}_*\|_2,$$

provided

$$\mathbf{m} \geq \frac{\mathbf{r}^2}{c^2 \left(\frac{1}{\gamma} + \frac{1}{\lambda}\right)} \log \frac{2\mathbf{d}}{\varepsilon},$$

where λ_i is the i -th singular value of \mathbf{X} , $\mathbf{r} = \sum_{i=1}^d \frac{\sigma_i^2}{\lambda/\gamma + \sigma_i^2}$, and the constant c is at least $1/32$.

The above theorem implies the number of required random projections is $\Omega(\mathbf{r} \log \mathbf{d})$, which can be significantly smaller than \mathbf{d} . The number \mathbf{r} is closely related to the numerical $\sqrt{\frac{\lambda}{\gamma}}$ -rank of \mathbf{X} (Hansen, 1998). We say that \mathbf{X} has numerical $\sqrt{\frac{\lambda}{\gamma}}$ -rank \mathbf{r}_ν if

$$r_\nu > \sum_{i=r_\nu+1}^d \frac{\sigma_i^2}{\lambda/\gamma + \sigma_i^2}.$$

Using the notation of numerical rank, we have

$$\mathbf{r} \leq \mathbf{r}_{\sqrt{\lambda/\gamma}} + \sum_{i=\mathbf{r}_{\sqrt{\lambda/\gamma}}+1}^d \frac{\sigma_i^2}{\lambda/\gamma + \sigma_i^2}.$$

Thus, when the singular value $\lambda_i \ll \sqrt{\lambda/\gamma}$ for $i > \mathbf{r}_{\sqrt{\lambda/\gamma}}$, which means that \mathbf{X} can be well approximated by a rank $\mathbf{r}_{\sqrt{\lambda/\gamma}}$ matrix, we have $\mathbf{r} = \mathbf{O}(\mathbf{r}_{\sqrt{\lambda/\gamma}})$.

One remarking property of our approach is it enjoys a multiplicative bound even in the full rank case. Thus, as long as $\frac{\varepsilon}{1-\varepsilon} \left(1 + \frac{\sqrt{\lambda}}{\sqrt{\gamma} k}\right) < 1$, we can use Algorithm 2 to reduce the reconstruction error exponentially over the iterations. In contrast, the random projection based algorithm for SVD (Halko et al., 2011), although is able to accurately recover the

eigen-space when the matrix is of low rank, it will result in a significant error in uncovering the subspace spanned by the top singular vectors when applied to matrices of full rank, and therefore is unable to recover the optimal solution \mathbf{w}_* accurately.

Finally, we note that the assumption that the optimal solution lies in the subspace spanned by the top singular vectors has been used in kernel learning (Guo and Zhou, 2012) and semi-supervised learning (Ji et al., 2012).

5. Analysis

Due to the limitation of space, we just provided the analysis for the low rank case. Before presenting the analysis, we first establish some notations and facts. Let the SVD of \mathbf{X} be

$$\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^\top,$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$, $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_r)$, $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_r)$, σ_i is the i -th singular value of \mathbf{X} , $\mathbf{u}_i \in \mathbb{R}^d$ and $\mathbf{v}_i \in \mathbb{R}^n$ are the corresponding left and right singular vectors of \mathbf{X} . We define

$$\gamma_* = \Sigma\mathbf{V}^\top \mathbf{D}(\mathbf{y})\alpha_*, \text{ and } \tilde{\gamma} = \Sigma\mathbf{V}^\top \mathbf{D}(\mathbf{y})\hat{\alpha}_*. \quad (17)$$

It is straightforward to show that

$$\mathbf{w}_* = -\frac{1}{\Sigma} \mathbf{U}\Sigma\mathbf{V}^\top \mathbf{D}(\mathbf{y})\alpha_* = -\frac{1}{\Sigma} \mathbf{U} \gamma_*, \text{ and } \tilde{\mathbf{w}} = -\frac{1}{\Sigma} \mathbf{U}\Sigma\mathbf{V}^\top \mathbf{D}(\mathbf{y})\hat{\alpha}_* = -\frac{1}{\Sigma} \mathbf{U} \tilde{\gamma}.$$

Since \mathbf{U} is an orthogonal matrix, we have

$$\|\mathbf{w}_*\|_2 = \|\gamma_*\|_2, \|\tilde{\mathbf{w}}\|_2 = \|\tilde{\gamma}\|_2, \text{ and } \|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2 = \|\tilde{\gamma} - \gamma_*\|_2. \quad (18)$$

Let us define

$$\mathbf{A} = \mathbf{U}^\top \mathbf{R} \in \mathbb{R}^{r \times m}.$$

It is easy to verify that \mathbf{A} is a Gaussian matrix of size $\mathbf{r} \times \mathbf{m}$.

5.1. Proof of Theorem 2

We first introduce the following concentration inequality for Gaussian random matrix, which serves the key to our analysis.

Corollary 7 *Let $\mathbf{A} \in \mathbb{R}^{r \times m}$ be a standard Gaussian random matrix. For any $0 < \epsilon \leq 1/2$, with a probability at least $1 - \epsilon$, we have*

$$\left\| \frac{1}{\mathbf{m}} \mathbf{A} \mathbf{A}^\top - \mathbf{I} \right\|_2 \leq \epsilon,$$

provided

$$\mathbf{m} \geq \frac{(\mathbf{r} + 1) \log(2\mathbf{r}/\epsilon)}{\epsilon^2},$$

where $\|\cdot\|_2$

Define $\mathbf{L}(\boldsymbol{\alpha})$ and $\widehat{\mathbf{L}}(\boldsymbol{\alpha})$ as

$$\mathbf{L}(\boldsymbol{\alpha}) = -\sum_{i=1}^n \ast(\mathbf{y}_i - \mathbf{A}_i \boldsymbol{\alpha}) - \frac{1}{2} \boldsymbol{\alpha}^\top \mathbf{G} \boldsymbol{\alpha}, \text{ and } \widehat{\mathbf{L}}(\boldsymbol{\alpha}) = -\sum_{i=1}^n \ast(\mathbf{y}_i - \widehat{\mathbf{A}}_i \boldsymbol{\alpha}) - \frac{1}{2} \boldsymbol{\alpha}^\top \widehat{\mathbf{G}} \boldsymbol{\alpha}.$$

Since $\widehat{\boldsymbol{\alpha}}_*$ maximizes $\widehat{\mathbf{L}}(\boldsymbol{\alpha})$ over the domain Ω^n , we have

$$\widehat{\mathbf{L}}(\widehat{\boldsymbol{\alpha}}_*) \geq \widehat{\mathbf{L}}(\boldsymbol{\alpha}_*) + \frac{1}{2} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \widehat{\mathbf{G}} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*). \quad (19)$$

Using the concaveness of $\widehat{\mathbf{L}}(\boldsymbol{\alpha})$, we have

$$\begin{aligned} & \widehat{\mathbf{L}}(\widehat{\boldsymbol{\alpha}}_*) + \frac{1}{2} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \widehat{\mathbf{G}} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*) \\ & \leq \widehat{\mathbf{L}}(\boldsymbol{\alpha}_*) + (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \left(\nabla \widehat{\mathbf{L}}(\boldsymbol{\alpha}_*) - \nabla \mathbf{L}(\boldsymbol{\alpha}_*) + \nabla \mathbf{L}(\boldsymbol{\alpha}_*) \right) \\ & \leq \widehat{\mathbf{L}}(\boldsymbol{\alpha}_*) + \frac{1}{2} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top (\mathbf{G} - \widehat{\mathbf{G}}) (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*), \end{aligned} \quad (20)$$

where the last inequality follows from the fact that $(\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \nabla \mathbf{L}(\boldsymbol{\alpha}_*) \leq 0$ since $\boldsymbol{\alpha}_*$ maximizes $\mathbf{L}(\boldsymbol{\alpha})$ over the domain Ω^n . Combining the inequalities in (19) and (20), we have

$$\frac{1}{2} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top (\mathbf{G} - \widehat{\mathbf{G}}) (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*) \geq \frac{1}{2} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \widehat{\mathbf{G}} (\widehat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*).$$

We rewrite \mathbf{G} and $\widehat{\mathbf{G}}$ as

$$\begin{aligned} \mathbf{G} &= \mathbf{D}(\mathbf{y}) \mathbf{V} \Sigma \mathbf{U}^\top \mathbf{U} \Sigma \mathbf{V}^\top \mathbf{D}(\mathbf{y}) = \mathbf{D}(\mathbf{y}) \mathbf{V} \Sigma \Sigma \mathbf{V}^\top \mathbf{D}(\mathbf{y}), \\ \widehat{\mathbf{G}} &= \mathbf{D}(\mathbf{y}) \mathbf{V} \Sigma \mathbf{U}^\top \frac{\mathbf{R} \mathbf{R}^\top}{\mathbf{m}} \mathbf{U} \Sigma \mathbf{V}^\top \mathbf{D}(\mathbf{y}) = \mathbf{D}(\mathbf{y}) \mathbf{V} \Sigma \frac{\mathbf{A} \mathbf{A}^\top}{\mathbf{m}} \Sigma \mathbf{V}^\top \mathbf{D}(\mathbf{y}). \end{aligned}$$

Using the definitions of $\boldsymbol{\gamma}_*$ and $\widetilde{\boldsymbol{\gamma}}$ in (17), we obtain

$$(\widetilde{\boldsymbol{\gamma}} - \boldsymbol{\gamma}_*)^\top \left(\mathbf{I} - \frac{\mathbf{A} \mathbf{A}^\top}{\mathbf{m}} \right) \boldsymbol{\gamma}_* \geq (\widetilde{\boldsymbol{\gamma}} - \boldsymbol{\gamma}_*)^\top \frac{\mathbf{A} \mathbf{A}^\top}{\mathbf{m}} (\widetilde{\boldsymbol{\gamma}} - \boldsymbol{\gamma}_*). \quad (21)$$

From Corollary 7, with a probability at least $1 - \delta$, we have $\left\| \mathbf{I} - \frac{1}{m} \mathbf{A} \mathbf{A}^\top \right\|_2 \leq \delta$, under the given condition on \mathbf{m} . Therefore, we obtain

$$(1 - \delta) \|\widetilde{\boldsymbol{\gamma}} - \boldsymbol{\gamma}_*\|_2 \leq \|\boldsymbol{\gamma}_*\|_2.$$

We complete the proof by using the equalities given in (18).

5.2. Proof of Theorem 5

At the t -th iteration, we consider the following optimization problem:

$$\min_{\mathbf{w} \in \mathbb{R}^d} \mathbf{L}^t(\mathbf{w}; \mathbf{X}, \mathbf{y}) = \frac{1}{2} \|\mathbf{w} + \widetilde{\mathbf{w}}^{t-1}\|_2^2 + \sum_{i=1}^n \left(\mathbf{y}_i (\mathbf{w} + \widetilde{\mathbf{w}}^{t-1})^\top \mathbf{x}_i \right), \quad (22)$$

where $\tilde{\mathbf{w}}^{t-1}$ is the solution obtained from the $\mathbf{t}-1$ -th iteration. It is straightforward to show that $\Delta_*^t = \mathbf{w}_* - \tilde{\mathbf{w}}^{t-1}$ is the optimal solution to (22). Then we can use the dual random projection approach to recover Δ_*^t by $\tilde{\Delta}^t$. If we can similarly show that

$$\|\tilde{\Delta}^t - \Delta_*^t\|_2 \leq \frac{1}{1-} \|\Delta_*^t\|_2,$$

then we update the recovered solution by $\tilde{\mathbf{w}}^t = \tilde{\mathbf{w}}^{t-1} + \tilde{\Delta}^t$ and have

$$\|\tilde{\mathbf{w}}^t - \mathbf{w}_*\|_2 = \|\tilde{\Delta}^t - \Delta_*^t\|_2 \leq \frac{1}{1-} \|\Delta_*^t\|_2 = \frac{1}{1-} \|\tilde{\mathbf{w}}^{t-1} - \mathbf{w}_*\|_2.$$

As a result, if we repeat the above process for $\mathbf{t} = 1, \dots, \mathbf{T}$, the recovery error of the last solution $\tilde{\mathbf{w}}^T$ is upper bounded by

$$\|\tilde{\mathbf{w}}^T - \mathbf{w}_*\|_2 \leq \left(\frac{1}{1-}\right)^T \|\tilde{\mathbf{w}}^0 - \mathbf{w}_*\|_2 = \left(\frac{1}{1-}\right)^T \|\mathbf{w}_*\|_2,$$

where we assume $\tilde{\mathbf{w}}^0 = \mathbf{0}$.

The remaining question is how to compute the $\tilde{\Delta}^t$ using the dual random projection approach. In order to make the previous analysis remain valid for the recovered solution $\tilde{\Delta}^t$ to the problem (22), we need to write the primal optimization problem in the same form as in (1). To this end, we first note that $\tilde{\mathbf{w}}^{t-1}$ lies in the subspace spanned by $\mathbf{x}_1, \dots, \mathbf{x}_n$, thus we write $\tilde{\mathbf{w}}^{t-1}$ as

$$\tilde{\mathbf{w}}^{t-1} = -\frac{1}{\mathbf{D}(\mathbf{y})} \hat{\alpha}_*^{t-1} = -\frac{1}{\mathbf{D}(\mathbf{y})} \sum_{i=1}^n [\hat{\alpha}_*^{t-1}]_i \mathbf{y}_i \mathbf{x}_i.$$

Then, $\mathbf{L}^t(\mathbf{w}; \mathbf{X}, \mathbf{y})$ can be written as

$$\begin{aligned} \mathbf{L}^t(\mathbf{w}; \mathbf{X}, \mathbf{y}) &= \frac{1}{2} \|\tilde{\mathbf{w}}^{t-1}\|_2^2 + \frac{1}{2} \|\mathbf{w}\|_2^2 + \mathbf{w}^\top \tilde{\mathbf{w}}^{t-1} + \sum_{i=1}^n \left(\mathbf{y}_i \mathbf{w}^\top \mathbf{x}_i + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) \\ &= \frac{1}{2} \|\tilde{\mathbf{w}}^{t-1}\|_2^2 + \frac{1}{2} \|\mathbf{w}\|_2^2 + \sum_{i=1}^n \left(\mathbf{y}_i \mathbf{w}^\top \mathbf{x}_i + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) - [\hat{\alpha}_*^{t-1}]_i \mathbf{y}_i \mathbf{w}^\top \mathbf{x}_i \\ &= \frac{1}{2} \|\tilde{\mathbf{w}}^{t-1}\|_2^2 + \frac{1}{2} \|\mathbf{w}\|_2^2 + \sum_{i=1}^n \tilde{\ell}_i^t(\mathbf{z}), \end{aligned}$$

where the new loss function $\tilde{\ell}_i^t(\mathbf{z}), i = 1, \dots, n$ is defined as

$$\tilde{\ell}_i^t(\mathbf{z}) = \left(\mathbf{z} + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) - [\hat{\alpha}_*^{t-1}]_i \mathbf{z}. \quad (23)$$

Therefore, Δ_*^t is the solution to the following problem:

$$\min_{\mathbf{w} \in \mathbb{R}^d}$$

To apply the dual random projection approach to recover Δ_*^t , we solve the following low-dimensional optimization problem:

$$\min_{\mathbf{z} \in \mathbb{R}^m} \frac{1}{2} \|\mathbf{z}\|_2^2 + \sum_{i=1}^n \frac{t}{i} \left(\mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i \right),$$

where $\hat{\mathbf{x}}_i \in \mathbb{R}^m$ is the low-dimensional representation for example $\mathbf{x}_i \in \mathbb{R}^d$. The following derivation signifies that the above problem is equivalent to the problem in (13).

$$\begin{aligned} & \frac{1}{2} \|\mathbf{z}\|_2^2 + \sum_{i=1}^n \frac{t}{i} \left(\mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i \right) \\ = & \frac{1}{2} \|\mathbf{z}\|_2^2 + \sum_{i=1}^n \left(\mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) - [\hat{\boldsymbol{\alpha}}_*^{t-1}]_i \mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i \\ = & \frac{1}{2} \|\mathbf{z}\|_2^2 + \frac{1}{\sqrt{\mathbf{m}}} \mathbf{z}^\top (\mathbf{R}^\top \tilde{\mathbf{w}}^{t-1}) + \sum_{i=1}^n \left(\mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) \\ = & \frac{1}{2} \left\| \mathbf{z} + \frac{1}{\sqrt{\mathbf{m}}} \mathbf{R}^\top \tilde{\mathbf{w}}^{t-1} \right\|_2^2 + \sum_{i=1}^n \left(\mathbf{y}_i \mathbf{z}^\top \hat{\mathbf{x}}_i + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) - \frac{1}{2} \left\| \frac{1}{\sqrt{\mathbf{m}}} \mathbf{R}^\top \tilde{\mathbf{w}}^{t-1} \right\|_2^2, \end{aligned}$$

where in the third line we use the fact that $\hat{\mathbf{x}}_i = \mathbf{R}^\top \mathbf{x}_i / \sqrt{\mathbf{m}}$ and $\tilde{\mathbf{w}}^{t-1} = -\sum_i [\hat{\boldsymbol{\alpha}}_*^{t-1}]_i \mathbf{y}_i \mathbf{x}_i /$. Given the optimal solution \mathbf{z}_*^t to the above problem, we can recover Δ_*^t by

$$\tilde{\Delta}^t = -\frac{1}{\mathbf{m}} \mathbf{X} \mathbf{D}(\mathbf{y}) \hat{\mathbf{a}}_*^t,$$

where $\hat{\mathbf{a}}_*^t$ is computed by

$$[\hat{\mathbf{a}}_*^t]_i = \nabla \frac{t}{i} \left(\mathbf{y}_i \hat{\mathbf{x}}_i^\top \mathbf{z}_*^t \right) = \nabla \left(\mathbf{y}_i \hat{\mathbf{x}}_i^\top \mathbf{z}_*^t + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right) - [\hat{\boldsymbol{\alpha}}_*^{t-1}]_i, \quad \mathbf{i} = 1, \dots, \mathbf{n}.$$

The updated solution $\tilde{\mathbf{w}}^t$ is computed by

$$\tilde{\mathbf{w}}^t = \tilde{\mathbf{w}}^{t-1} + \tilde{\Delta}^t = -\frac{1}{\mathbf{m}} \mathbf{X} \mathbf{D}(\mathbf{y}) (\hat{\boldsymbol{\alpha}}_*^{t-1} + \hat{\mathbf{a}}_*^t) = -\frac{1}{\mathbf{m}} \mathbf{X} \mathbf{D}(\mathbf{y}) \hat{\boldsymbol{\alpha}}_*^t,$$

where $[\hat{\boldsymbol{\alpha}}_*^t]_i = [\hat{\boldsymbol{\alpha}}_*^{t-1}]_i + [\hat{\mathbf{a}}_*^t]_i = \nabla \left(\mathbf{y}_i \hat{\mathbf{x}}_i^\top \mathbf{z}_*^t + \mathbf{y}_i [\tilde{\mathbf{w}}^{t-1}]^\top \mathbf{x}_i \right), \mathbf{i} = 1, \dots, \mathbf{n}.$

6. Conclusion

In this paper, we consider the problem of recovering the optimal solution \mathbf{w}_* to the original high-dimensional optimization problem using random projection. To this end, we propose to use the dual solution $\hat{\boldsymbol{\alpha}}_*$ to the low-dimensional optimization problem to recover \mathbf{w}_* . Our analysis shows that with a high probability, the solution $\tilde{\mathbf{w}}$ returned by our proposed method approximates the optimal solution \mathbf{w}_* with small error.

Acknowledgments

This work is partially supported by Office of Navy Research (ONR Award N00014-09-1-0663 and N000141210431).

References

- Dimitris Achlioptas. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of Computer and System Sciences*, 66(4):671 – 687, 2003.
- Rosa I. Arriaga and Santosh Vempala. An algorithmic theory of learning: robust concepts and random projection. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, pages 616–623, 1999.
- Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. Kernels as features: On kernels, margins, and low-dimensional mappings. *Machine Learning*, 65(1):79–94, 2006.
- Ella Bingham and Heikki Mannila. Random projection in dimensionality reduction: applications to image and text data. In *Proceedings of the 7th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 245–250, 2001.
- Avrim Blum. Random projection, margins, kernels, and feature-selection. In *Proceedings of the 2005 international conference on Subspace, Latent Structure and Feature Selection*, pages 52–68, 2006.
- J.M. Borwein, A.S. Lewis, J. Borwein, and AS Lewis. *Convex analysis and nonlinear optimization: theory and examples*. Springer New York, 2006.
- Christos Boutsidis, Anastasios Zouzias, and Petros Drineas. Random projections for k -means clustering. In *Advances in Neural Information Processing Systems 23*, pages 298–306, 2010.
- Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- Vladimir Braverman, Rafail Ostrovsky, and Yuval Rabani. Rademacher chaos, random eulerian graphs and the sparse johnson-lindenstrauss transform. *ArXiv e-prints*, arXiv:1011.2590, 2010.
- Emmanuel J. Candès and Michael B. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21 –30, 2008.
- Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Sanjoy Dasgupta and Yoav Freund. Random projection trees and low dimensional manifolds. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 537–546, 2008.
- David L. Donoho. Compressed sensing. *IEEE Transaction on Information Theory*, 52: 1289–1306, 2006.
- Xiaoli Zhang Fern and Carla E. Brodley. Random projection for high dimensional data clustering: a cluster ensemble approach. In *Proceedings of the 20th International Conference on Machine Learning*, pages 186–193, 2003.

- Dmitriy Fradkin and David Madigan. Experiments with random projections for machine learning. In *Proceedings of the 9th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 517–522, 2003.
- Yoav Freund, Sanjoy Dasgupta, Mayank Kabra, and Nakul Verma. Learning the structure of manifolds using random projections. In *Advances in Neural Information Processing Systems 20*, pages 473–480, 2008.
- Alex Gittens and Joel A. Tropp. Tail bounds for all eigenvalues of a sum of random matrices. *ArXiv e-prints*, arXiv:1104.4513, 2011.
- Navin Goel, George Bebis, and Ara Nefian. Face recognition experiments with random projection. In *Proceedings of SPIE*, pages 426–437, 2005.
- Gene H. Golub and Charles F. Van Loan. *Matrix computations, 3rd Edition*. Johns Hopkins University Press, 1996.
- Xin Guo and Ding-Xuan Zhou. An empirical feature-based learning algorithm producing sparse approximations. *Applied and Computational Harmonic Analysis*, 32(3):389–400, 2012.
- Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.
- N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for cona

- Edo Liberty, Nir Ailon, and Amit Singer. Dense fast random projections and lean walsh transforms. In *In Proceedings of the 12th International Workshop on Randomization and Computation (RANDOM)*, pages 512–522, 2008.
- Oldalric-Ambrym Maillard and Remi Munos. Linear regression with random projections. *Journal of Machine Learning Research*, 13:2735–2772, 2012.
- Rajeev Mowani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- Yu. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- Saurabh Paul, Christos Boutsidis, Malik Magdon-Ismael, and Petros Drineas. Random projections for support vector machines. *ArXiv e-prints*, arXiv:1211.6085, 2012.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems 20*, pages 1177–1184, 2008.
- Shai Shalev-Shwartz and Yoram Singer. Online learning meets optimization in the dual. In *Proceedings of 19th Annual Conference on Learning Theory (COLT)*, pages 423–437, 2006.
- Qinfeng Shi, Chunhua Shen, Rhys Hill, and Anton van den Hengel. Is margin preserved after random projection? In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- Santosh S. Vempala. *The Random Projection Method*. American Mathematical Society, 2004.
- Shenghuo Zhu. A short note on the tail bound of wishart distribution. *ArXiv e-prints*, arXiv:1212.5860, 2012.

Appendix A. Proof of Proposition 1 and Proposition 2

Since the two propositions can be proved similarly, we only present the proof of Proposition 1. First, if α_* is the optimal dual solution, by replacing (\cdot) in (1) with its conjugate form, the optimal primal solution can be solved by

$$\mathbf{w}_* = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{w}\|_2^2 + \sum_{i=1}^n [\alpha_*]_i \mathbf{y}_i \mathbf{x}_i^\top \mathbf{w}.$$

Setting the gradient with respect to \mathbf{w} to zero, we obtain

$$\mathbf{w}_* = -\frac{1}{2} \sum_{i=1}^n [\alpha_*]_i \mathbf{y}_i \mathbf{x}_i = -\frac{1}{2} \mathbf{X} \mathbf{D}(\mathbf{y}) \alpha_*.$$

Second, let's consider how to obtain the dual solution α_* from the primal solution \mathbf{w}_* . Note that

$$(\mathbf{y}_i \mathbf{x}_i^\top \mathbf{w}_*) = [\alpha_*]_i \left(\mathbf{y}_i \mathbf{x}_i^\top \mathbf{w}_* \right) - *([\alpha_*]_i).$$

By the Fenchel conjugate theory (Borwein et al., 2006; Cesa-Bianchi and Lugosi, 2006), we have α_* satisfying

$$[\alpha_*]_i = \nabla \left(\mathbf{y}_i \mathbf{x}_i^\top \mathbf{w}_* \right), \quad i = 1, \dots, n.$$

Appendix B. Proof of Corollary 7

In the proof, we make use of the recent development in tail bounds for the eigenvalues of a sum of random matrices (Gittens and Tropp, 2011; Zhu, 2012).

Theorem 8 (Theorem 1 (Zhu, 2012)) *Let $\{\xi_j : j = 1, \dots, n\}$ be i.i.d. samples drawn from a multivariate Gaussian distribution $\mathcal{N}(0, \mathbf{C})$, where $\mathbf{C} \in \mathbb{R}^{d \times d}$. Define*

$$\hat{\mathbf{C}}_n = \frac{1}{n} \sum_{j=1}^n \xi_j \xi_j^\top.$$

We denote the trace of \mathbf{X} by $\text{tr}(\mathbf{X})$, and the spectral norm of \mathbf{X} by $\|\mathbf{X}\|$. Then, for any ≥ 0

$$\Pr \left\{ \left\| \hat{\mathbf{C}}_n - \mathbf{C} \right\| \geq \left(\sqrt{\frac{2(\mathbf{k} + 1)}{n}} + \frac{2\mathbf{k}}{n} \right) \|\mathbf{C}\| \right\} \leq 2\mathbf{d} \exp(-),$$

where $\mathbf{k} = \text{tr}(\mathbf{C})/\|\mathbf{C}\|$.

We write $\mathbf{A} = (\xi_1, \dots, \xi_m)$, where $\xi_i \in \mathbb{R}^r$ is i.i.d. sampled from the Gaussian distribution $\mathcal{N}(0, \mathbf{I})$, and write $\mathbf{A}\mathbf{A}^\top/m$ as

$$\frac{1}{m} \mathbf{A}\mathbf{A}^\top = \frac{1}{m} \sum_{i=1}^m \eta_i \eta_i^\top.$$

Following Theorem 8, we have, with a probability at least $1 - 2\mathbf{r} \exp(-)$

$$\left\| \frac{1}{m} \mathbf{A}\mathbf{A}^\top - \mathbf{I} \right\| \leq \sqrt{\quad}$$

Appendix C. Proof of Theorem 1

Before presenting our analysis, we first state a version of Johnson-Lindenstrauss theorem that is useful to our analysis.

Theorem 9 (Theorem 2 (Blum, 2006)) *Let $\mathbf{x} \in \mathbb{R}^d$, and $\hat{\mathbf{x}} = \mathbf{R}^\top \mathbf{x} / \sqrt{\mathbf{m}}$, where $\mathbf{R} \in \mathbb{R}^{d \times m}$ is a random matrix whose entries are chosen independently from $\mathcal{N}(0, 1)$. Then*

$$\Pr \left\{ (1 - \epsilon) \|\mathbf{x}\|_2^2 \leq \|\hat{\mathbf{x}}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x}\|_2^2 \right\} \geq 1 - 2 \exp \left(-\frac{\mathbf{m}}{4} (\epsilon^2 - \epsilon^3) \right).$$

According to (21) in the proof of Theorem 2, we have

$$(\tilde{\gamma} - \gamma_*)^\top \left(\mathbf{I} - \frac{\mathbf{A}\mathbf{A}^\top}{\mathbf{m}} \right) \gamma_* \geq (\tilde{\gamma} - \gamma_*)^\top \frac{\mathbf{A}\mathbf{A}^\top}{\mathbf{m}} (\tilde{\gamma} - \gamma_*).$$

Notice that

$$\begin{aligned} \sqrt{\mathbf{m}} \mathbf{z}_* &= -\frac{1}{\mathbf{m}} \mathbf{R}^\top \mathbf{X} \mathbf{D}(\mathbf{y}) \hat{\alpha}_* = -\frac{1}{\mathbf{m}} \mathbf{R}^\top \mathbf{U} \Sigma \mathbf{V}^\top \mathbf{D}(\mathbf{y}) \hat{\alpha}_* = -\frac{1}{\mathbf{m}} \mathbf{R}^\top \mathbf{U} \tilde{\gamma} = -\frac{1}{\mathbf{m}} \mathbf{A}^\top \tilde{\gamma}, \\ \mathbf{R}^\top \mathbf{w}_* &= -\frac{1}{\mathbf{m}} \mathbf{R}^\top \mathbf{U} \mathbf{w}_* = -\frac{1}{\mathbf{m}} \mathbf{A}^\top \mathbf{w}_*. \end{aligned}$$

Then, we have

$$\frac{2}{\mathbf{m}} \|\sqrt{\mathbf{m}} \mathbf{z}_* - \mathbf{R}^\top \mathbf{w}_*\|_2^2 \leq (\tilde{\gamma} - \gamma_*)^\top \left(\mathbf{I} - \frac{\mathbf{A}\mathbf{A}^\top}{\mathbf{m}} \right) \gamma_*.$$

Using Corollary 7, with a probability at least $1 - \epsilon$, we have

$$\frac{1}{\mathbf{m}} \|\sqrt{\mathbf{m}} \mathbf{z}_* - \mathbf{R}^\top \mathbf{w}_*\|_2^2 \leq \|\mathbf{w}_*\|_2 \|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2.$$

Following Theorem 2, with a probability at least $1 - \epsilon$, we have

$$\frac{1}{\mathbf{m}} \|\sqrt{\mathbf{m}} \mathbf{z}_* - \mathbf{R}^\top \mathbf{w}_*\|_2^2 \leq \frac{2}{1 - \epsilon} \|\mathbf{w}_*\|_2^2. \quad (24)$$

To replace \mathbf{w}_* on R. H. S. of the above inequality with $\mathbf{R}^\top \mathbf{w}_*$, we make use of Theorem 9. With a probability at least $1 - \exp(-(\epsilon^2 - \epsilon^3)\mathbf{m}/4)$, we have

$$(1 - \epsilon) \|\mathbf{w}_*\|_2^2 \leq \frac{1}{\mathbf{m}} \|\mathbf{R}^\top \mathbf{w}_*\|_2^2.$$

By choosing $\epsilon = 1/2$, we have, with a probability at least $1 - \exp(-\mathbf{m}/32)$

$$\frac{1}{2} \|\mathbf{w}_*\|_2^2 \leq \frac{1}{\mathbf{m}} \|\mathbf{R}^\top \mathbf{w}_*\|_2^2. \quad (25)$$

We complete the proof by combining the two inequalities in (24) and (25).

Appendix D. Proof of Theorem 3

As discussed before, the key reason for the large difference between $\widehat{\mathbf{w}}$ and \mathbf{w}_* is because they do not lie in the same subspace: \mathbf{w}_* lies in the subspace spanned by the columns in \mathbf{U} while $\widehat{\mathbf{w}}$ lies in the subspace spanned by the column vectors in a random matrix.

In the subspace orthogonal to $\mathbf{u}_1, \dots, \mathbf{u}_r$, we randomly choose a subset of $\mathbf{d}-\mathbf{r}$ orthogonal bases, denoted by $\mathbf{u}_{r+1}, \dots, \mathbf{u}_d$. Let $\mathbf{U}_\perp = (\mathbf{u}_{r+1}, \dots, \mathbf{u}_d)$. Since

$$\|\widehat{\mathbf{w}} - \mathbf{w}_*\|_2 = \max_{\|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top (\widehat{\mathbf{w}} - \mathbf{w}_*),$$

to facilitate our analysis, we restrict the choice of \mathbf{x} to the subspace spanned by $\mathbf{u}_{r+1}, \dots, \mathbf{u}_d$ and have

$$\|\widehat{\mathbf{w}} - \mathbf{w}_*\|_2 \geq \max_{\mathbf{x} \in \text{span}(\mathbf{u}_{r+1}, \dots, \mathbf{u}_d), \|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top \widehat{\mathbf{w}},$$

where we use the fact $\mathbf{w}_* \perp \text{span}(\mathbf{u}_{r+1}, \dots, \mathbf{u}_d)$. From Proposition 2, we can express $\widehat{\mathbf{w}}$ as

$$\widehat{\mathbf{w}} = \frac{1}{\sqrt{\mathbf{m}}} \mathbf{R} \mathbf{z}_* = -\frac{1}{\mathbf{m}} \mathbf{R} \mathbf{R}^\top \mathbf{X} \mathbf{D}(\mathbf{y}) \widehat{\alpha}_* = -\frac{1}{\mathbf{m}} \mathbf{R} \mathbf{R}^\top \mathbf{U} \Sigma \mathbf{V}^\top \mathbf{D}(\mathbf{y}) \widehat{\alpha}_* = -\frac{1}{\mathbf{m}} \mathbf{R} \mathbf{R}^\top \mathbf{U} \widetilde{\gamma},$$

where $\widetilde{\gamma}$ is defined in (17). Write \mathbf{x} as $\mathbf{x} = \mathbf{U}_\perp \mathbf{a}$, where $\mathbf{a} \in \mathbb{R}^{d-r}$. Define

$$\Lambda = \mathbf{U}_\perp^\top \mathbf{R} \in \mathbb{R}^{(d-r) \times m}.$$

As a result, we bound $\|\mathbf{w}_* - \widehat{\mathbf{w}}\|_2$ by

$$\max_{\mathbf{x} \in \text{span}(\mathbf{u}_{r+1}, \dots, \mathbf{u}_d), \|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top \widehat{\mathbf{w}} = \max_{\|\mathbf{a}\|_2 \leq 1} \frac{1}{\mathbf{m}} \mathbf{a}^\top \mathbf{U}_\perp^\top \mathbf{R} \mathbf{R}^\top \mathbf{U} \widetilde{\gamma} = \frac{1}{\mathbf{m}} \|\Lambda \mathbf{A}^\top \widetilde{\gamma}\|_2. \quad (26)$$

It is easy to verify that \mathbf{A} and Λ are two independent Gaussian random matrices. Therefore, we can fix the vector $\mathbf{A}^\top \widetilde{\gamma}$ and estimate how the random matrix Λ affect the norm of vector $\Lambda \mathbf{A}^\top \widetilde{\gamma}$. According to Theorem 9 (i.e., Johnson-Lindenstrauss theorem), for a fixed vector $\mathbf{A}^\top \widetilde{\gamma}$, with a probability at least $1 - \exp(-(\mathbf{d}-\mathbf{r})/32)$

$$\frac{1}{\sqrt{\mathbf{d}-\mathbf{r}}} \|\Lambda \mathbf{A}^\top \widetilde{\gamma}\|_2 \geq \frac{1}{\sqrt{2}} \|\mathbf{A}^\top \widetilde{\gamma}\|_2. \quad (27)$$

We now bound $\|\mathbf{A}^\top \widetilde{\gamma}\|_2$. Note that we cannot directly apply Theorem 9 to bound the norm of $\mathbf{A}^\top \widetilde{\gamma}$ because $\widetilde{\gamma}$ is a random variable depending on the random matrix \mathbf{A} . To decouple the dependence between \mathbf{A} and $\widetilde{\gamma}$, we expand $\|\mathbf{A}^\top \widetilde{\gamma}\|_2$ as

$$\|\mathbf{A}^\top \widetilde{\gamma}\|_2 \geq \|\mathbf{A}^\top \gamma_*\|_2 - \|\mathbf{A}^\top (\gamma_* - \widetilde{\gamma})\|_2, \quad (28)$$

where γ_* is defined in (17). We bound the two terms on the right side of the inequality in (28) separately. Using Theorem 9, with a probability at least $1 - \exp(-\mathbf{m}/32)$, we bound $\|\mathbf{A}^\top \gamma_*\|_2$ by

$$\frac{1}{\sqrt{\mathbf{m}}} \|\mathbf{A}^\top \gamma_*\|_2 \geq \frac{1}{\sqrt{2}} \|\gamma_*\|_2 = \frac{1}{\sqrt{2}} \|\mathbf{w}_*\|_2. \quad (29)$$

To bound the second term $\|\mathbf{A}^\top(\gamma_* - \tilde{\gamma})\|$, with a probability at least $1 - \delta$, we have

$$\frac{1}{\sqrt{\mathbf{m}}} \|\mathbf{A}^\top(\gamma_* - \tilde{\gamma})\|_2 \leq \sqrt{\max(\mathbf{A}\mathbf{A}^\top/\mathbf{m})} \|\gamma_* - \tilde{\gamma}\|_2 \leq \sqrt{1 + \delta} \|\mathbf{w}_* - \tilde{\mathbf{w}}\|_2,$$

where we use the result in Corollary 7. According to Theorem 2, we have

$$\|\mathbf{w}_* - \tilde{\mathbf{w}}\|_2 \leq \frac{\delta}{1 - \delta} \|\mathbf{w}_*\|_2.$$

As a result, with a probability at least $1 - \delta$, we have

$$\frac{1}{\sqrt{\mathbf{m}}} \|\mathbf{A}^\top(\gamma_* - \tilde{\gamma})\|_2 \leq \sqrt{1 + \delta} \frac{\delta}{1 - \delta} \|\mathbf{w}_*\|_2. \quad (30)$$

We complete the proof by putting together (26), (27), (28), (29), and (30).

Appendix E. Proof of Theorem 4

It is straightforward to check that

$$\hat{\mathbf{w}} = \frac{\mathbf{R}\mathbf{R}^\top}{\mathbf{m}} \tilde{\mathbf{w}}.$$

Therefore,

$$\begin{aligned} & \max_{\mathbf{x} \in \text{span}(X), \|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top (\hat{\mathbf{w}} - \mathbf{w}_*) \\ & \leq \|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2 + \max_{\|\mathbf{x}\|_2 \leq 1, \mathbf{x} \in \text{span}(X)} \mathbf{x}^\top (\hat{\mathbf{w}} - \tilde{\mathbf{w}}) \\ & = \|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2 + \max_{\|\mathbf{a}\|_2 \leq 1} \mathbf{a}^\top \left(\frac{1}{\mathbf{m}} \mathbf{U}^\top \mathbf{R}\mathbf{R}^\top \mathbf{U} - \mathbf{I} \right) \tilde{\gamma} \\ & \leq \|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2 + \max \left(\frac{1}{\mathbf{m}} \mathbf{U}^\top \mathbf{R}\mathbf{R}^\top \mathbf{U} - \mathbf{I} \right) \|\tilde{\mathbf{w}}\|_2 \\ & \leq \|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2 + \max \left(\frac{1}{\mathbf{m}} \mathbf{A}\mathbf{A}^\top - \mathbf{I} \right) \|\mathbf{w}_*\|, \end{aligned}$$

where in the fourth line we use the fact $\|\tilde{\mathbf{w}}\|_2 = \|\tilde{\gamma}\|_2$. Using Corollary 7, we have, with a probability at least $1 - \delta$

$$\max \left(\frac{1}{\mathbf{m}} \mathbf{A}\mathbf{A}^\top - \mathbf{I} \right) \leq \delta.$$

We complete the proof by using the bound for $\|\tilde{\mathbf{w}} - \mathbf{w}_*\|_2$ stated in Theorem 2.

Appendix F. Proof of Theorem 6

Define $\mathbf{L}(\alpha)$ and $\hat{\mathbf{L}}(\alpha)$ as

$$\mathbf{L}(\alpha) = - \sum_{i=1}^n \mathbf{y}_i \mathbf{x}_i^\top - \frac{1}{2} \alpha^\top \mathbf{G} \alpha, \text{ and } \hat{\mathbf{L}}(\alpha) = - \sum_{i=1}^n \hat{\mathbf{y}}_i \hat{\mathbf{x}}_i^\top - \frac{1}{2} \alpha^\top \hat{\mathbf{G}} \alpha.$$

Since $\psi(\cdot)$ is γ -smooth, and thus $\psi_*(\cdot)$ is $\frac{1}{\gamma}$ -strongly convex. Define

$$\mathbf{g}_*(\mathbf{x}) = \psi_*(\mathbf{x}) - \frac{1}{2}\|\mathbf{x}\|^2, \quad \mathbf{H} = \mathbf{G} + \gamma\mathbf{I}, \quad \text{and} \quad \hat{\mathbf{H}} = \hat{\mathbf{G}} + \gamma\mathbf{I}.$$

Evidently, $\mathbf{g}_*(\cdot)$ is still a convex function. We write $\mathcal{L}(\boldsymbol{\alpha})$ and $\hat{\mathcal{L}}(\boldsymbol{\alpha})$ as

$$\mathcal{L}(\boldsymbol{\alpha}) = -\sum_{i=1}^n \mathbf{g}_*(\mathbf{x}_i) - \frac{1}{2}\boldsymbol{\alpha}^\top \mathbf{H} \boldsymbol{\alpha}, \quad \text{and} \quad \hat{\mathcal{L}}(\boldsymbol{\alpha}) = -\sum_{i=1}^n \mathbf{g}_*(\mathbf{x}_i) - \frac{1}{2}\boldsymbol{\alpha}^\top \hat{\mathbf{H}} \boldsymbol{\alpha}.$$

Since $\hat{\boldsymbol{\alpha}}_*$ maximizes $\hat{\mathcal{L}}(\boldsymbol{\alpha})$ over the domain Ω^n , we have

$$\hat{\mathcal{L}}(\hat{\boldsymbol{\alpha}}_*) \geq \hat{\mathcal{L}}(\boldsymbol{\alpha}_*) + \frac{1}{2}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \hat{\mathbf{H}}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*). \quad (31)$$

Using the concaveness of $\hat{\mathcal{L}}(\boldsymbol{\alpha})$, we have

$$\begin{aligned} & \hat{\mathcal{L}}(\hat{\boldsymbol{\alpha}}_*) + \frac{1}{2}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \hat{\mathbf{H}}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*) \\ & \leq \hat{\mathcal{L}}(\boldsymbol{\alpha}_*) + (\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \left(\nabla \hat{\mathcal{L}}(\boldsymbol{\alpha}_*) - \nabla \mathcal{L}(\boldsymbol{\alpha}_*) + \nabla \mathcal{L}(\boldsymbol{\alpha}_*) \right) \\ & \leq \hat{\mathcal{L}}(\boldsymbol{\alpha}_*) + \frac{1}{2}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top (\mathbf{H} - \hat{\mathbf{H}})(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*), \end{aligned} \quad (32)$$

where the last inequality follows from the fact that $(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \nabla \mathcal{L}(\boldsymbol{\alpha}_*) \leq 0$ since $\boldsymbol{\alpha}_*$ maximizes $\mathcal{L}(\boldsymbol{\alpha})$ over the domain Ω^n . Combining the inequalities in (31) and (32), we have

$$\frac{1}{2}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top (\mathbf{H} - \hat{\mathbf{H}})(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*) \geq \frac{1}{2}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \hat{\mathbf{H}}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*). \quad (33)$$

Define $\mathbf{K} = \mathbf{H}^{-1/2} \hat{\mathbf{H}} \mathbf{H}^{-1/2}$. We rewrite the bound in (33) as

$$(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \mathbf{H}^{1/2} (\mathbf{I} - \mathbf{K}) \mathbf{H}^{1/2} \boldsymbol{\alpha}_* \geq (\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)^\top \mathbf{H}^{1/2} \mathbf{K} \mathbf{H}^{1/2} (\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*).$$

To bound the spectral norm of \mathbf{K} , we have the following lemma.

Lemma 3 *With a probability at least $1 -$*

and the eigen decomposition of $\mathbf{G} = \mathbf{D}(\mathbf{y})\mathbf{X}^\top\mathbf{X}\mathbf{D}(\mathbf{y})$ is

$$\mathbf{G} = \sum_{i=1}^d \frac{2}{i} [\mathbf{D}(\mathbf{y})\mathbf{v}_i][\mathbf{D}(\mathbf{y})\mathbf{v}_i]^\top.$$

Following the Corollary 11 in Appendix G, we obtain Lemma 3. \blacksquare

From Lemma 3, we have, with a probability at least $1 - \delta$,

$$\|\mathbf{I} - \mathbf{K}\|_2 \leq \delta,$$

and therefore

$$\|\mathbf{H}^{1/2}\boldsymbol{\alpha}_*\|_2 \geq (1 - \delta) \|\mathbf{H}^{1/2}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)\|_2 \geq (1 - \delta) \|\mathbf{G}^{1/2}(\hat{\boldsymbol{\alpha}}_* - \boldsymbol{\alpha}_*)\|_2. \quad (34)$$

Since

$$\mathbf{w}_* = -\frac{1}{n}\mathbf{X}\mathbf{D}(\mathbf{y})\boldsymbol{\alpha}_* = -\frac{1}{n} \sum_{i=1}^d \frac{2}{i} ([\mathbf{D}(\mathbf{y})\mathbf{v}_i]^\top \boldsymbol{\alpha}_*) \mathbf{u}_i,$$

the assumption that \mathbf{w}_* lies in the subspace spanned by $\{\mathbf{u}_1, \dots, \mathbf{u}_k\}$ implies $\boldsymbol{\alpha}_*$ lies in the subspace spanned by $\{\mathbf{D}(\mathbf{y})\mathbf{v}_1, \dots, \mathbf{D}(\mathbf{y})\mathbf{v}_k\}$. Then, we have

$$\boldsymbol{\alpha}_*^\top \mathbf{G} \boldsymbol{\alpha}_* = \boldsymbol{\alpha}_*^\top \mathbf{G}_k \boldsymbol{\alpha}_* \geq \frac{2}{k} \boldsymbol{\alpha}_*^T \boldsymbol{\alpha}_*,$$

where \mathbf{G}_k is the rank- k best approximation of \mathbf{G} and $\frac{2}{k}$ is the k -th singular value of \mathbf{X} . As a result, we conclude

$$\|\mathbf{H}^{1/2}\boldsymbol{\alpha}_*\|_2^2 = \boldsymbol{\alpha}_*^\top \left(\mathbf{G} + -\mathbf{I} \right) \boldsymbol{\alpha}_* \leq \left(1 + \frac{1}{k} \right) \boldsymbol{\alpha}_*^\top \mathbf{G} \boldsymbol{\alpha}_* = \left(1 + \frac{1}{k} \right) \|\mathbf{G}^{1/2}\boldsymbol{\alpha}_*\|_2^2. \quad (35)$$

Combining (34) and (35), we have, with a probability at least $1 - \delta$

$$\sqrt{1 + \frac{1}{k}} \|\mathbf{G}^{1/2}\boldsymbol{\alpha}_*\|_2 \geq (1 - \delta) \|\mathbf{G}^{1/2}(\boldsymbol{\alpha}_* - \hat{\boldsymbol{\alpha}}_*)\|_2.$$

We complete the proof by using the relationship between \mathbf{w}_* , $\tilde{\mathbf{w}}$ and $\boldsymbol{\alpha}_*$, $\hat{\boldsymbol{\alpha}}_*$.

Appendix G. A matrix concentration inequality

Theorem 10 *Let $\mathbf{C} = \text{diag}(\mathbf{c}_1, \dots, \mathbf{c}_p)$ and $\mathbf{S} = \text{diag}(\mathbf{s}_1, \dots, \mathbf{s}_p)$ be $p \times p$ diagonal matrices, where $\mathbf{c}_i \neq 0$ and $\mathbf{c}_i^2 + \mathbf{s}_i^2 = 1$ for all i . Let \mathbf{R} be a Gaussian random matrix of size $p \times n$. Let $\mathbf{M} = \mathbf{C}^2 + \frac{1}{n}\mathbf{S}\mathbf{R}\mathbf{R}^\top\mathbf{S}$ and $\mathbf{r} = \sum_i \mathbf{s}_i^2$.*

$$\Pr(\lambda_{\min}(\mathbf{M}) \geq 1 + \mathbf{t}) \leq \mathbf{q} \cdot \exp\left(-\frac{\mathbf{c}n\mathbf{t}^2}{\max_i(\mathbf{s}_i^2)\mathbf{r}}\right),$$

$$\Pr(\lambda_{\max}(\mathbf{M}) \leq 1 - \mathbf{t}) \leq \mathbf{q} \cdot \exp\left(-\frac{\mathbf{c}n\mathbf{t}^2}{\max_i(\mathbf{s}_i^2)\mathbf{r}}\right),$$

where the constant \mathbf{c} is at least $1/32$, and \mathbf{q} is the rank of \mathbf{S} .

Proof [Theorem 10] The proof is similar to Theorems 5.3 and 7.1 of (Gittens and Tropp, 2011), except for adding a bias matrix. Let $\mathbf{g}(\theta) = \frac{\theta^2}{2(1-\theta)}$. We have

$$\begin{aligned}
 & \Pr \left\{ \mathbf{1}(\mathbf{C}^2 + \frac{1}{\mathbf{n}} \mathbf{SRR}^\top \mathbf{S}) \geq 1 + \mathbf{t} \right\} \\
 & \leq \inf_{\theta > 0} \text{tr exp} \left\{ \left(\mathbf{C}^2 + \frac{1}{\mathbf{n}} \mathbb{E} [\mathbf{SRR}^\top \mathbf{S}] - (1 + \mathbf{t}) \mathbf{I} \right) + \frac{1}{\mathbf{n}} \mathbf{g}(\theta) \mathbb{E} [(\mathbf{SRR}^\top \mathbf{S})^2] \right\} \\
 & \leq \inf_{\theta > 0} \text{tr exp} \left\{ -\mathbf{t} + 8\mathbf{g}(\theta) \text{tr}(\mathbf{S}^2) \mathbf{S}^2 \right\} \\
 & \leq \inf_{\theta > 0} \mathbf{q} \exp \left\{ -\mathbf{t} + 8r\mathbf{g}(\theta) \right\} \leq \mathbf{q} \exp \left(-\frac{\mathbf{n}\mathbf{t}^2}{32r} \right), \\
 & \Pr \left\{ \mathbf{q}(\mathbf{C}^2 + \frac{1}{\mathbf{n}} \mathbf{SRR}^\top \mathbf{S}) \leq 1 - \mathbf{t} \right\} \\
 & \leq \inf_{\theta > 0} \text{tr exp} \left\{ -\left(\mathbf{C}^2 + \frac{1}{\mathbf{n}} \mathbb{E} [\mathbf{SRR}^\top \mathbf{S}] - (1 - \mathbf{t}) \mathbf{I} \right) + \frac{1}{\mathbf{n}} \mathbf{g}(\theta) \mathbb{E} [(\mathbf{SRR}^\top \mathbf{S})^2] \right\} \\
 & \leq \inf_{\theta > 0} \mathbf{q} \exp \left\{ -\mathbf{t} + 8r\mathbf{g}(\theta) \right\} \leq \mathbf{q} \exp \left(-\frac{\mathbf{n}\mathbf{t}^2}{32r} \right).
 \end{aligned}$$

■

Corollary 11 *Let \mathbf{A} be a given matrix of size $\mathbf{m} \times \mathbf{p}$, and \mathbf{R} be a Gaussian random matrix of size $\mathbf{p} \times \mathbf{n}$. Let δ be a positive constant, $\frac{2}{\delta} = \max_i (\mathbf{A}^\top \mathbf{A})_i$, and $r = \sum_i \frac{\sigma_i^2}{\lambda + \sigma_i^2}$. Let $\mathbf{K} = \mathbf{I}_m + \mathbf{A}\mathbf{A}^\top$, $\tilde{\mathbf{K}} = \mathbf{I}_m + \frac{1}{n} \mathbf{A}\mathbf{R}\mathbf{R}^\top \mathbf{A}^\top$, and $\tilde{\mathbf{I}} = \mathbf{K}^{-1/2} \tilde{\mathbf{K}} \mathbf{K}^{-1/2}$. If $\mathbf{n} \geq \frac{r\sigma_1^2}{c\delta^2(\lambda + \sigma_1^2)} \log \frac{2p}{\delta}$, then with probability at least $1 - \delta$,*

$$(1 - \mathbf{t}) \mathbf{I}_m \preceq \tilde{\mathbf{I}} \preceq (1 + \mathbf{t}) \mathbf{I}_m. \quad (36)$$

Proof [Corollary 11] Let $\mathbf{s}_i^2 = \frac{\sigma_i^2}{\lambda + \sigma_i^2}$, and $\mathbf{c}_i^2 = 1 - \mathbf{s}_i^2$. By SVD and Theorem 10, we have $\Pr(1 - \mathbf{t} \leq \min_i (\tilde{\mathbf{I}}) \leq 1 + \mathbf{t}) \geq 1 - \delta$. ■