

Improved Subsampled Randomized Hadamard Transform for Linear SVM

Zijian Lei

Department of Computer Science,
Hong Kong Baptist University,
Hong Kong SAR, China
cszjlei@comp.hkbu.edu.hk

Liang Lan

Department of Computer Science,
Hong Kong Baptist University,
Hong Kong SAR, China
lanliang@comp.hkbu.edu.hk

Abstract

Subsampled Randomized Hadamard Transform (SRHT), a popular random projection method that can efficiently project a d -dimensional data into r -dimensional space ($r \ll d$) in $O(d \log(d))$ time, has been widely used to address the challenge of high-dimensionality in machine learning. SRHT works by rotating the input data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ by Randomized Walsh-Hadamard Transform followed with a subsequent uniform column sampling on the rotated matrix. Despite the advantages of SRHT, one limitation of SRHT is that it generates the new low-dimensional embedding without considering any specific properties of a given dataset. Therefore, this data-independent random projection method may result in inferior and unstable performance when used for a particular machine learning task, e.g., classification. To overcome this limitation, we analyze the effect of using SRHT for random projection in the context of linear SVM classification. Based on our analysis, we propose importance sampling and deterministic top- r sampling to produce effective low-dimensional embedding instead of uniform sampling SRHT. In addition, we also proposed a new supervised non-uniform sampling method. Our experimental results have demonstrated that our proposed methods can achieve higher classification accuracies than SRHT and other random projection methods on six real-life datasets.

Introduction

Classification is one of the fundamental machine learning tasks. In the era of big data, huge amounts of high-dimensional data become common in a wide variety of applications. As the dimensionality of the data for real-life classification tasks has increased dramatically in recent years, it is essential to develop accurate and efficient classification algorithms for high-dimensional data.

One of the most popular methods to address the high-dimensionality challenge is dimensionality reduction. It allows the classification problems to be efficiently solved in a lower dimensional space. Popular dimensionality reduction methods includes Principal Component Analysis (PCA) (Jolliffe 2011), Linear Discriminant Analysis (LDA) (Mika

et al. 1999) etc. However, these traditional dimensionality reduction methods are computationally expensive which makes them not suitable for big data. Recently, random projection (Bingham and Mannila 2001) that projects a high-dimensional feature vector $\mathbf{x} \in \mathbb{R}^d$ into a low-dimensional feature $\tilde{\mathbf{x}} \in \mathbb{R}^r$ ($r \ll d$) using a random orthonormal matrix $\mathbf{R} \in \mathbb{R}^{d \times r}$. Namely, $\tilde{\mathbf{x}} = \mathbf{R}^T \mathbf{x}$. Since random projection methods are (1) computationally efficient while maintaining accurate results; (2) simple to implement in practice; and (3) easy to analyze (Mahoney and others 2011; Xu et al. 2017), they have attracted significant research attentions in recent years.

The theoretical foundation of random projection is the Johnson-Lindenstrauss Lemma (JL lemma) (Johnson, Lindenstrauss, and Schechtman 1986). JL lemma proves that in Euclidean space, high-dimensional data can be randomly projected into lower dimensional space while the pairwise distances between all data points are preserved with a high probability. Based on the JL lemma, several different ways are proposed for constructing the random matrix \mathbf{R} : (1) Gaussian Matrix (Dasgupta and Gupta 1999): each entry in \mathbf{R} is generated from a Gaussian distribution with mean equals to 0 and variance equals to $\frac{1}{d}$; (2) Achlioptas matrix (Achlioptas 2003): each entry in \mathbf{R} is generated from $\{-1, 0, 1\}$ from a discrete distribution. This method generates a sparse random matrix \mathbf{R} . Therefore it requires less memory and computation cost; (3) sparse embedding matrix (or called count sketch matrix) (Clarkson and Woodruff 2017) and its similar variant feature hashing (Weinberger et al. 2009; Freksen, Kamma, and Larsen 2018): for each row \mathbf{R} , only one column is randomly selected and assign either 1 or -1 with probability 0.5 to this entry; and (4) Subsampled Randomized Hadamard Transform (SRHT) (Tropp 2011): It uses a highly structured matrix \mathbf{R} for random projection. The procedure of SRHT can be summarized as follows: It first rotates the input matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ (n is the number of samples, d is the dimensionality of the data) by Randomized Walsh-Hadamard Transform and then uniformly sampling r columns from the rotated matrix. More details of SRHT will be discussed in next section.

Although random projection methods have attracted a great deal of research attention in recent years (Sarlos

2006; Choromanski, Rowland, and Weller 2017) and have been studied for regression (Lu et al. 2013), classification (Paul et al. 2013; Zhang et al. 2013; Paul, Magdon-Ismail, and Drineas 2015) and clustering (Boutsidis, Zouzias, and Drineas 2010; Liu, Shen, and Tsang 2017), most of the existing studies focused on data-independent projection. In other words, the random matrix \mathbf{R} was constructed without considering any specific properties of a given dataset.

In this paper, we focus on SRHT for random projection and study how it affects the classification performance of linear Support Vector Machines (SVM). SRHT achieve dimensionality reduction by uniformly sampling r columns from a rotated version of the input data matrix \mathbf{X} . Even though Randomized Walsh-Hadamard transform in SRHT tends to equalize column norms (Boutsidis and Gittens 2013), we argue that the importance of different columns in the rotated data matrix are not equal. Therefore, the uniformly random sampling procedure in SRHT would result in low accuracy when used as a dimensionality reduction method in high-dimensional data classification.

To overcome the limitation of SRHT, we propose to produce effective low-dimensional embedding by using non-uniform sampling instead of uniform sampling. To achieve this goal, we first analyze the effect of using SRHT for random projection in linear SVM classification. Based on our analysis, we have proposed importance sampling and deterministic top- r sampling methods to improve SRHT. Furthermore, we also propose a new sampling method by incorporating label information. It samples the columns which can achieve optimal inter-class separability along with the intra-class compactness of the data samples.

Finally, we performed experiments to evaluate our proposed methods on six real-life datasets. Our experimental results clearly demonstrate that our proposed method can obtain higher classification accuracies than SRHT and other three popular random projection methods while only slightly increasing the running time.

Preliminaries

Random Projection

Given a data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, random projection methods reduce the dimensionality of \mathbf{X} by multiplying it by a random orthonormal matrix $\mathbf{R} \in \mathbb{R}^{d \times r}$ with parameter $r \ll d$. The projected data in the low-dimensional space is

$$\tilde{\mathbf{X}} = \mathbf{X}\mathbf{R} \in \mathbb{R}^{n \times r}. \quad (1)$$

Note that the matrix \mathbf{R} is randomly generated and is independent of the input data \mathbf{X} . The theoretical foundation of random projection is the Johnson-Lindenstrauss Lemma (JL lemma) as shown in following,

Lemma 1 (Johnson-Lindenstrauss Lemma (JL lemma) (Johnson, Lindenstrauss, and Schechtman 1986)). *For any $0 < \epsilon < 1$ and any integer n , let $r = O(\log n/\epsilon^2)$ and $\mathbf{R} \in \mathbb{R}^{d \times r}$ be a random orthonormal matrix. Then for any set \mathbf{X} of n points in \mathbb{R}^d , the following inequality about pairwise distance between any two data points \mathbf{x}_i and \mathbf{x}_j in \mathbf{X} holds true with high probability:*

$$(1-\epsilon)\|\mathbf{x}_i - \mathbf{x}_j\|_2 \leq \|\mathbf{R}^T \mathbf{x}_i - \mathbf{R}^T \mathbf{x}_j\|_2 \leq (1+\epsilon)\|\mathbf{x}_i - \mathbf{x}_j\|_2.$$

JL lemma proves that in Euclidean space, high-dimensional data can be randomly projected into lower dimensional space while the pairwise distances between all data points are well preserved with high probability. There are several different ways to construct the random projection matrix \mathbf{R} . In this paper, we focus on SRHT which uses a structured orthonormal matrix for random projection. The details about SRHT are as follows.

Subsampled Randomized Hadamard Transform (SRHT)

For $d = 2^q$ where q is any positive integer¹, SRHT defines a $d \times r$ matrix as:

$$\mathbf{R} = \sqrt{\frac{d}{r}} \mathbf{DHS} \quad (2)$$

where

- $\mathbf{D} \in \mathbb{R}^{d \times d}$ is a diagonal matrix whose elements are independent random signs $\{1, -1\}$;
- $\mathbf{H} \in \mathbb{R}^{d \times d}$ is a normalized Walsh-Hadamard matrix. The Walsh-Hadamard matrix is defined recursively as:

$$\mathbf{H}_d = \begin{bmatrix} \mathbf{H}_{d/2} & \mathbf{H}_{d/2} \\ \mathbf{H}_{d/2} & -\mathbf{H}_{d/2} \end{bmatrix}$$
 with $\mathbf{H}_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$, and

$$\mathbf{H} = \frac{1}{\sqrt{d}} \mathbf{H}_d \in \mathbb{R}^{d \times d}$$
;
- $\mathbf{S} \in \mathbb{R}^{d \times r}$ is a subset of randomly sampling r columns from the $d \times d$ identity matrix. The purpose of multiplying \mathbf{S} is to uniformly sample r columns from the rotated data matrix $\mathbf{X}_r = \mathbf{X}\mathbf{D}\mathbf{H}$.

There are several advantages of using SRHT for random projection. Firstly, we do not need to explicitly represent \mathbf{H} in memory and only need to store a diagonal random sign matrix \mathbf{D} and a sampling matrix \mathbf{S} with r non-zero values. Therefore, the memory cost of storing \mathbf{R} is only $O(d+r)$. Secondly, due to the recursive structure of Hadamard matrix \mathbf{H} , matrix multiplication $\mathbf{X}\mathbf{R}$ only takes $O(nd \log(d))$ time by using Fast Fourier Transform (FFT) for matrix multiplication (Tropp 2011). Ailon and Liberty (Ailon and Liberty 2009) further improve the time complexity of SRHT to $O(nd \log(r))$ if only r columns in \mathbf{X}_r are needed.

Methodology

Limitation of Uniform Sampling in SRHT

Despite the advantages of SRHT as mentioned in the previous section, one limitation of SRHT is that the random matrix \mathbf{R} is constructed without considering any specific properties of input data \mathbf{X} . Based on the definition in (2), SRHT works by first rotating the input data matrix \mathbf{X} by multiplying it with $\mathbf{D}\mathbf{H}$ and then uniformly random sampling r columns from the rotated matrix $\mathbf{X}_r = \mathbf{X}\mathbf{D}\mathbf{H}$. Since this uniform column sampling method does not consider the underlying data distribution of \mathbf{X}_r , it would result in low and unstable accuracy when used as a dimensionality reduction method in a specific classification problem. To illustrate the limitation of uniformly random sampling step in SRHT, we generate a simple two-dimensional synthetic data

¹We can ensure this by padding zeros to original data




Figure 1: Limitation of uniform sampling

as shown in Figure 1(a). The first class (i.e., class label ‘+’) was generated from $N(\mu_1 = [3, 3], \Sigma = \begin{bmatrix} 1 & 1.75 \\ 1.75 & 1 \end{bmatrix})$. The second class (i.e., class label ‘•’) was generated from $N(\mu_2 = [-3, -3], \Sigma = \begin{bmatrix} 1 & 1.75 \\ 1.75 & 1 \end{bmatrix})$. After multiplying with matrix DH , the original data was rotated into new feature space as shown in Figure 1(b). In this simple illustration example, SRHT might choose the feature that can not distinguish these two classes (i.e., the vertical feature in Figure 1(b)) and then leads to low classification accuracy. In Figure 1(c), we show the classification accuracy of using SRHT on a real-life dataset for 15 different runs. It clearly shows that this data-independent random projection method could result in low and unstable performance on this synthetic dataset. In the following sections, we will describe our proposed methods to construct the data-dependent sampling matrix S in (2) by exploiting the underlying data properties from both unsupervised and supervised perspectives.

Improved SRHT by unsupervised non-uniform sampling

First, we present our proposed non-uniform sampling methods from unsupervised perspective. In this paper, we study SRHT in the context of using linear SVM for classification. Assume we are given a training data set $D = \{\mathbf{x}_i, y_i\}_{i=1}^n$, where \mathbf{x}_i is a d -dimensional input feature vector, y_i is its corresponding class label. The dual problem of SVM can be written in matrix format as shown in following,

$$\begin{aligned} \max \quad & \mathbf{1}^T \boldsymbol{\alpha} - \frac{1}{2} \boldsymbol{\alpha}^T \mathbf{Y} \mathbf{X} \mathbf{X}^T \mathbf{Y} \boldsymbol{\alpha} \\ \text{s.t.} \quad & \mathbf{1}^T \mathbf{Y} \boldsymbol{\alpha} = 0 \\ & \mathbf{0} \leq \boldsymbol{\alpha} \leq \mathbf{C}, \end{aligned} \quad (3)$$

where $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the input feature matrix, and \mathbf{Y} is a diagonal matrix where the elements in the diagonal are class labels, $\mathbf{Y}_{ii} = y_i$, \mathbf{C} is the regularization parameter.

By applying SRHT to project original \mathbf{X} from d -dimensional space into r -dimensional space, i.e., $\tilde{\mathbf{X}} = \sqrt{\frac{d}{r}} \mathbf{X} \mathbf{D} \mathbf{H} \mathbf{S}$, the SVM optimization problem on projected

data $\tilde{\mathbf{X}}$ will be

$$\begin{aligned} \max \quad & \mathbf{1}^T \tilde{\boldsymbol{\alpha}} - \frac{1}{2} \tilde{\boldsymbol{\alpha}}^T \mathbf{Y} \tilde{\mathbf{X}} \tilde{\mathbf{X}}^T \mathbf{Y} \tilde{\boldsymbol{\alpha}} \\ \text{s.t.} \quad & \mathbf{1}^T \mathbf{Y} \tilde{\boldsymbol{\alpha}} = 0 \\ & \mathbf{0} \leq \tilde{\boldsymbol{\alpha}} \leq \mathbf{C}. \end{aligned} \quad (4)$$

Comparing the dual problem (3) in the original space and the dual problem (4) in the projected low-dimensional space, the only difference is that $\mathbf{X} \mathbf{X}^T$ in (3) is replaced by $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T$ in (4). The constraints on optimal solutions (i.e., $\boldsymbol{\alpha}^*$ of (3) and $\tilde{\boldsymbol{\alpha}}^*$ of (4)) do not depend on input data matrix. Therefore, $\tilde{\boldsymbol{\alpha}}^*$ is expected to be close to $\boldsymbol{\alpha}^*$ when $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T$ is close to $\mathbf{X} \mathbf{X}^T$.

Let us rewrite $\tilde{\mathbf{X}} = \mathbf{X}_r \mathbf{S} \sqrt{\frac{d}{r}}$ where $\mathbf{X}_r = \mathbf{X} \mathbf{D} \mathbf{H}$, we can view the effect of $\mathbf{X}_r \mathbf{S} \sqrt{\frac{d}{r}}$ is to uniformly sample r columns from \mathbf{X}_r and then re-scale each selected column by a scaling factor $\sqrt{\frac{d}{r}}$. It corresponds to use uniform column sampling for approximating matrix multiplication (Mahoney and others 2011). It has been proved in (Drineas, Kannan, and Mahoney 2006) that (i, j) -th element in $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T$ equals to the (i, j) -th element of exact product $\mathbf{X}_r \mathbf{X}_r^T$ in expectation regardless of the sampling probabilities. However, the variance of the (i, j) -th element in $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T$ depends on the choice of sampling probabilities. Therefore, simply using uniform sampling would result in large variance of each element in $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T$ and then cause a large approximation error between $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T$ and $\mathbf{X}_r \mathbf{X}_r^T$.

Motivated by this observation, we first propose to improve SRHT by employing importance sampling for randomized matrix multiplication. Specifically, we would like to design a method to construct $\tilde{\mathbf{S}}$ based on specific data properties of \mathbf{X}_r instead of using random sampling matrix \mathbf{S} in original SRHT. The idea is based on importance sampling for approximating matrix multiplication: (1) we assign each column i in \mathbf{X}_r a probability p_i such that $p_i \geq 0$ and $\sum_{i=1}^d p_i = 1$; (2) we then select r columns from \mathbf{X}_r based on the probability distribution $\{p_i\}_{i=1}^d$ to form $\tilde{\mathbf{X}}$; (3) a column in $\tilde{\mathbf{X}}$ is formed as $\sqrt{\frac{1}{r p_i}} \mathbf{X}_{r(:, i)}$ if this column is chosen

from the i -th column in \mathbf{X}_r where $\sqrt{\frac{1}{rp_i}}$ is the re-scaling factor. In other word, the matrix $\tilde{\mathbf{S}} \in \mathbb{R}^{d \times r}$ is constructed as: for every column j in $\tilde{\mathbf{S}}$ only one entry i (i is a random number from $\{1, \dots, d\}$) is selected by following the probability distribution $\{p_i\}_{i=1}^d$ and a non-zero value is assigned to this entry. The non-zero value is set as

$$\tilde{\mathbf{S}}_{ij} = \sqrt{\frac{1}{rp_i}}. \quad (5)$$

In the uniform sampling setting, $p_i = \frac{1}{d}$, therefore $\tilde{\mathbf{S}} = \sqrt{\frac{d}{r}}\mathbf{S}$ which is consistent with the setting in original SRHT as shown in (2).

Now, the question is what choice of probability $\{p_i\}_{i=1}^d$ can optimally reduce the expected approximation error between $\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T$ and $\mathbf{X}\mathbf{X}^T$. We use Frobenius norm to measure the approximation error between $\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T$ and $\mathbf{X}\mathbf{X}^T$.

Let us use $\mathbf{X}_{r(:,j)}$ to denote the j -th column of \mathbf{X}_r , as shown in the Proposition 1, the optimal sampling probability to minimize $\mathbf{E}[\|\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T - \mathbf{X}\mathbf{X}^T\|_F]$ is

$$p_j = \frac{\|\mathbf{X}_{r(:,j)}\|_2^2}{\sum_{j=1}^d \|\mathbf{X}_{r(:,j)}\|_2^2}. \quad (6)$$

Proposition 1. Suppose $\mathbf{X} \in \mathbb{R}^{n \times d}$, $\mathbf{X}_r = \mathbf{X}\mathbf{D}\mathbf{H} \in \mathbb{R}^{n \times d}$ and $\tilde{\mathbf{X}} = \mathbf{X}_r\tilde{\mathbf{S}}_{ij}$ where $\tilde{\mathbf{S}}_{ij} \in \mathbb{R}^{d \times r}$ is defined as in (5). Then, the optimal sampling probability to minimize the expectation of the Frobenius norm of the approximation error between $\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T$ and $\mathbf{X}\mathbf{X}^T$, i.e., $\mathbf{E}[\|\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T - \mathbf{X}\mathbf{X}^T\|]$, is

$$p_j = \frac{\|\mathbf{X}_{r(:,j)}\|_2^2}{\sum_{j=1}^d \|\mathbf{X}_{r(:,j)}\|_2^2}.$$

To prove Proposition 1, we first use the fact $\mathbf{X}\mathbf{X}^T = \mathbf{X}_r\mathbf{X}_r^T$ since $\mathbf{D}\mathbf{H}\mathbf{H}^T\mathbf{D}^T = \mathbf{I}$. Then, we apply the Lemma 4 from (Drineas, Kannan, and Mahoney 2006) to obtain the optimal sampling probability.

In addition, we also propose a deterministic top- r sampling, which select the r columns with largest Euclidean norms from \mathbf{X}_r . As shown in the Proposition 2, this deterministic top- r sampling method directly optimizes an upper bound of $\|\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T - \mathbf{X}\mathbf{X}^T\|_F$ without using re-scaling factor.

Proposition 2. Suppose $\mathbf{X} \in \mathbb{R}^{n \times d}$, $\mathbf{X}_r = \mathbf{X}\mathbf{D}\mathbf{H} \in \mathbb{R}^{n \times d}$ and $\tilde{\mathbf{X}} = \mathbf{X}_r\mathbf{P}$ where $\mathbf{P} \in \mathbb{R}^{d \times d}$ is a diagonal matrix whose elements are 1 or 0. $\mathbf{P}_{ii} = 1$ denotes the i -th column in \mathbf{X}_r is selected. Then, $\sum_{\mathbf{P}_{jj}=0} \|\mathbf{X}_{r(:,j)}\|_2^2$ is an upper bound of the approximation error $\|\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T - \mathbf{X}\mathbf{X}^T\|_F$.

Proof. Since $\tilde{\mathbf{X}} = \mathbf{X}_r\mathbf{P}$ and $\mathbf{X}\mathbf{X}^T = \mathbf{X}_r\mathbf{X}_r^T$, $\|\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T - \mathbf{X}\mathbf{X}^T\|_F = \|\mathbf{X}_r\mathbf{P}\mathbf{P}^T\mathbf{X}_r^T - \mathbf{X}_r\mathbf{X}_r^T\|_F$. Let us rewrite the matrix product $\mathbf{X}_r\mathbf{X}_r^T$ as the sum of d rank one matrices $\sum_{j=1}^d \mathbf{X}_{r(:,j)}\mathbf{X}_{r(:,j)}^T$. Similarly, $\mathbf{X}_r\mathbf{P}\mathbf{P}^T\mathbf{X}_r^T$ can be rewritten as $\sum_{j=1}^d \mathbf{P}_{jj}\mathbf{X}_{r(:,j)}\mathbf{X}_{r(:,j)}^T$. Therefore, the approximation error $\|\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T - \mathbf{X}\mathbf{X}^T\|_F$ can be upper bounded

as follows,

$$\begin{aligned} \|\tilde{\mathbf{X}}\tilde{\mathbf{X}}^T - \mathbf{X}\mathbf{X}^T\|_F &= \left\| \sum_{\mathbf{P}_{jj}=0} \mathbf{X}_{r(:,j)}\mathbf{X}_{r(:,j)}^T \right\|_F \\ &\leq \sum_{\mathbf{P}_{jj}=0} \|\mathbf{X}_{r(:,j)}\mathbf{X}_{r(:,j)}^T\|_F = \sum_{\mathbf{P}_{jj}=0} \|\mathbf{X}_{r(:,j)}\|_2^2 \end{aligned} \quad (7)$$

In our experimental section, we have shown that this deterministic top- r sampling obtains better classification accuracy than importance sampling method. Adelman and Silberstein(Adelman and Silberstein 2018) also observed that deterministic top- r sampling gets better results when they use randomized matrix multiplication approximation to speedup the training of deep neural networks.

Improved SRHT by Supervised Non-uniform Sampling

In this section, we describe our proposed non-uniform sampling method from supervised perspective. We propose to construct the sampling matrix $\tilde{\mathbf{S}}$ by incorporating label information. Our proposed method is based on the idea of metric learning (Xing et al. 2003). The idea is to select r columns from \mathbf{X}_r for improving the inter-class separability along with the intra-class compactness as used in (Lan et al. 2019).

Let \mathbf{z}_i be new feature representation of the i -th data sample in \mathbf{X}_r . Then, we measure the interclass separability by the sum of squared pair-wise distances between \mathbf{z}_i 's from different classes as $\sum_{y_i \neq y_j} \|\mathbf{z}_i - \mathbf{z}_j\|^2$. Similarly, we measure the intra-class tightness by the sum of squared pair-wise distances between \mathbf{z}_i 's from the same class as $\sum_{y_i = y_j} \|\mathbf{z}_i - \mathbf{z}_j\|^2$. We formulate our objective as a combination of these two terms,

$$\begin{aligned} \sum_{i,j \in D} \|\mathbf{z}_i - \mathbf{z}_j\|^2 \mathbf{A}_{ij} &= \text{trace}((\mathbf{X}_r\tilde{\mathbf{S}})^T \mathbf{L}(\mathbf{X}_r\tilde{\mathbf{S}})) \\ &= \text{trace}(\mathbf{X}_r^T \mathbf{L} \mathbf{X}_r \tilde{\mathbf{S}} \tilde{\mathbf{S}}^T), \end{aligned} \quad (8)$$

where $\mathbf{A}_{ij} = 1$ if $y_i = y_j$, and $\mathbf{A}_{ij} = -a$ if $y_i \neq y_j$, parameter a is used to balance the tradeoff between inter-class separability and intra-class tightness, \mathbf{L} is the Laplacian matrix of \mathbf{A} , defined as $\mathbf{L} = \mathbf{D} - \mathbf{A}$, and \mathbf{D} is a diagonal degree matrix of \mathbf{A} such that $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$. $\tilde{\mathbf{S}}$ is the sampling matrix whose elements are 0 or 1. We want to learn $\tilde{\mathbf{S}}$ by minimizing (8).

Without losing any information, let us rewrite our variable $\tilde{\mathbf{S}}$ as a $d \times d$ diagonal matrix \mathbf{P} where the diagonal element $\mathbf{P}_{ii} \in \{0, 1\}$. $\mathbf{P}_{ii} = 1$ means the i -th column in \mathbf{X}_r is selected and $\mathbf{P}_{ii} = 0$ means the i -th column is not selected. Then $\tilde{\mathbf{S}}$ can be viewed as a submatrix of \mathbf{P} that selects the columns with $\mathbf{P}_{ii} = 1$ from \mathbf{P} .

It is straight-forward to verify that $\tilde{\mathbf{S}}\tilde{\mathbf{S}}^T$ in (8) equals to \mathbf{P} . Therefore, our optimization is formulated as

$$\begin{aligned} \min_{\mathbf{P}} \text{trace}(\mathbf{X}_r^T \mathbf{L} \mathbf{X}_r \mathbf{P}) \\ \text{s.t. } \mathbf{P}_{ii} \in \{0, 1\}, \\ \sum_{i=1}^d \mathbf{P}_{ii} = r. \end{aligned} \quad (9)$$

Algorithm 1 Improved Subsampled Randomized Hadamard Transform (ISRHT) for linear SVM classification

Training

Input: training set \mathbf{X}_{train} , reduced dimension r , regularization parameter c ;

Output: model $\mathbf{w} \in \mathbb{R}^{r \times 1}$, matrix $\mathbf{D} \in \mathbb{R}^{d \times d}$ and $\tilde{\mathbf{S}} \in \mathbb{R}^{d \times r}$;

- 1: Generate a diagonal random sign matrix \mathbf{D} as in (2)
- 2: Compute $\mathbf{X}_r = \mathbf{X}_{train} \mathbf{D} \mathbf{H} \in \mathbb{R}^{n \times d}$ by FFT
- 3: Obtain the sampling matrix $\tilde{\mathbf{S}}$
- 4: Compute $\tilde{\mathbf{X}}_{train} = \mathbf{X}_r \tilde{\mathbf{S}} \in \mathbb{R}^{n \times r}$
- 5: Train a linear SVM on $\tilde{\mathbf{X}}_{train}$ and obtain the model \mathbf{w}

Prediction

Input: test data \mathbf{X}_{test} , model \mathbf{w} , matrix \mathbf{D} , $\tilde{\mathbf{S}}$;

Output: predicted labels $\hat{\mathbf{y}}_{test}$;

- 1: compute $\tilde{\mathbf{X}}_{test} = \mathbf{X}_{test} \mathbf{D} \mathbf{H} \tilde{\mathbf{S}} \in \mathbb{R}^{n \times r}$
 - 2: predict by $\hat{\mathbf{y}}_{test} = \tilde{\mathbf{X}}_{test} \mathbf{w}$;
-

Let vector \mathbf{b} denote the diagonal of the matrix $\mathbf{X}_r^\top \mathbf{L} \mathbf{X}_r$, i.e. $b_i = (\mathbf{X}_r^\top \mathbf{L} \mathbf{X}_r)_{ii}$, vector \mathbf{p} denote the diagonal of \mathbf{P} . Then, the objective (9) is simplified to

$$\begin{aligned} \min_{\mathbf{p}} \quad & \mathbf{b}^\top \mathbf{p} \\ \text{s.t.} \quad & p_i \in \{0, 1\}, \\ & \sum_{i=1}^d p_i = r. \end{aligned} \quad (10)$$

The optimal solution of (10) is just to select the largest r elements in \mathbf{b}_i and set the corresponding entries in \mathbf{p}_i to 1. Then $\tilde{\mathbf{S}}$ can easily be obtained based on non-zero values in \mathbf{p} .

Algorithm Implementation and Analysis

We summarize our proposed algorithm in algorithm 1 and name it as Improved Subsampled Randomized Hadamard Transform (ISRHT). With respect to training time complexity, step 2 needs $O(nd \log(d))$ time by using FFT. In step 3, for unsupervised column selection, we need $O(nd)$ time to compute the Euclidean norm for each column. For supervised column selection, b_i in (10) are computed as $b_i = (\mathbf{X}_r^\top \mathbf{L} \mathbf{X}_r)_{ii}$ and the Laplacian matrix \mathbf{L} is equal to $\mathbf{D} - \mathbf{y} \mathbf{y}^\top$, therefore, the required time is still $O(nd)$. Step 4 needs $O(r)$ time and step 5 needs $O(tnr)$ where t is the number of iterations needed for training linear SVM. Therefore, the overall time complexity of our proposed algorithm is $O(nd \log(d) + nd + tnr)$. With respect to prediction time complexity, our proposed method needs $O(d \log(d) + r)$ time to make a prediction for a new test sample. With respect to space complexity, \mathbf{D} is a diagonal matrix, matrix $\tilde{\mathbf{S}}$ is a sampling matrix with r non-zero entries. Therefore, the required space for storing the projection matrix \mathbf{R} of our proposed methods is $O(d + r)$. Our proposed methods achieve very efficient time and space complexities for model deployment. We compare the time and space complexity of different random projection methods in Table 1.

Table 1: Comparison of different random projection methods

Algorithms	Space for \mathbf{R}	Time for projection
Gaussian	$O(dr)$	$O(ndr)$
Achlioptas	$O(\frac{1}{3}dr)$	$O(\frac{1}{3}ndr)$
Sparse Embedding	$O(d)$	$O(nnz(\mathbf{X}))$
SRHT	$O(d + r)$	$O(nd \log(r))$
ISRHT-unsupervised	$O(d + r)$	$O(nd \log(d) + nd)$
ISRHT-supervised	$O(d + r)$	$O(nd \log(d) + nd)$

Extension to High-dimensional Sparse Data

In this section, we describe an extension of our algorithm on high-dimensional sparse data. One disadvantage of applying Hadamard transform on sparse data is that it produces a dense matrix. As shown in the step 2 in Algorithm 1, the memory cost is increased from $O(nnz(\mathbf{X}))$ to $O(nd)$. To tackle this challenge, we use a combination of sparse embedding and SRHT as proposed in (Chen et al. 2015) for memory-efficient projection. We improve their work by replacing the original SRHT by our proposed ISRHT methods. We expect that our proposed data-dependent projection methods can obtain higher accuracies than data-independent projection methods on high-dimensional sparse data.

Specifically, we first use sparse embedding matrix $\mathbf{R}_{sparse} \in \mathbb{R}^{d \times r'}$ to project \mathbf{X} into r' -dimensional space and then use our improved SRHT methods to project into r -dimensional space. The sparse embedding matrix \mathbf{R}_{sparse} is generated as follows, for each row in random matrix \mathbf{R}_{sparse} , randomly selected one column and assign either 1 or -1 with probability 0.5 to this entry. All other entries are 0s. Due to the sparse structure of \mathbf{R}_{sparse} , matrix multiplication $\mathbf{X} \mathbf{R}_{sparse}$ takes $O(nnz(\mathbf{X}))$ time. By doing this, the time complexity of the proposed method on high-dimensional sparse data was reduced to $O(nnz(\mathbf{X}) + nr' \log(r') + nr')$ and the memory cost was reduced to $O(nnz(\mathbf{X}) + nr')$. In our experimental setting, we set $r' = 2r$ as suggested in (Chen et al. 2015).

Experiments

In this section, we compare our proposed methods with other four popular random projection methods on six real-world benchmark datasets. These benchmark datasets are downloaded from LIBSVM website². For datasets (mushrooms, real-sim, rcv1-binary, news20-binary), we refer to the setting in (Pourkamali-Anaraki, Becker, and Wakin 2018) and randomly select 70% of the original training data as training data and the rest 30% as test data. The detailed information about these six datasets is summarized in Table 3.

We evaluate the performance of the following seven random projection algorithms in our experiments:

- Gaussian: random projection by Gaussian Matrix (Dasgupta and Gupta 1999);

²<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

Table 2: The classification accuracy and running time (in seconds) (The best results for each dataset are in bold and italic and the best results among the seven random projection methods are in bold).

Algorithms	Datasets					
	mushrooms ($r = 16$)	usps-b ($r = 16$)	gisette ($r = 256$)	real-sim ($r = 256$)	rcv1-binary ($r = 256$)	news20-binary ($r = 256$)
All features	<i>99.85</i> 0.1s	<i>91.03</i> 2.8s	<i>97.5</i> 3.0s	<i>97.43</i> 0.1s	<i>96.43</i> 0.1s	<i>96.01</i> 0.3s
PCA	96.23 0.3s	85.85 0.6s	96.50 3.2s	93.10 51.0s	95.05 34.5s	- -
DLSS	94.66 0.2s	80.21 1.0s	93.90 5.8s	89.89 2.0s	91.43 1.6s	79.33 9.3s
Gaussian	89.15±3.85 0.2s	80.05±3.33 1.0s	89.76±0.56 5.9s	76.48±0.55 19s	78.90±0.61 1.2s	70.10±0.78 20.3s
Achlioptas	91.74±2.75 0.3s	81.50±2.28 0.9s	89.56±0.65 7.8s	76.20±0.63 6.3s	79.03±1.28 1.5s	69.29±0.99 28.6s
Sparse Embedding	89.06±4.18 0.3s	82.09±2.05 0.8s	89.65±0.86 1.5s	76.82±0.45 2.4s	79.11±1.20 0.2s	70.08±0.69 0.6s
SRHT	92.45±2.59 0.3s	81.87±1.55 0.3s	91.77±0.76 1.5s	72.80±0.23 23.3s	79.08±0.97 0.5s	70.58±0.26 0.7s
ISRHT-nps	94.30±1.15 0.4s	83.34±1.31 0.6s	91.17±0.51 1.4s	73.04±0.51 24.2s	79.28±0.63 0.6s	70.36±0.84 0.9s
ISRHT-top- r	94.23±1.58 0.4s	82.91±1.07 0.4s	93.20±0.61 1.4s	76.81±0.46 18.7s	82.81±0.89 0.6s	73.46±0.90 0.9s
ISRHT-supervised	<i>96.25±1.22</i> 0.2s	<i>86.25±0.81</i> 0.4s	<i>93.90±0.72</i> 1.4s	<i>80.04±0.52</i> 44.3s	<i>87.81±0.42</i> 1.0s	<i>78.33±0.74</i> 1.1s

Table 3: Experiment dataset

Dataset	train size	test size	dimensionality
mushrooms	6,000	2124	112
usps-b	7,291	2007	256
gisette	6,000	1000	5,000
real-sim	48,447	23862	20,958
rcv1-binary	13,562	6680	47,236
news20-binary	13,397	6599	1,355,191

- Achlioptas: random projection by Achlioptas matrix (Achlioptas 2003);
- Sparse Embedding: random projection by count sketch matrix (Clarkson and Woodruff 2017);
- SRHT: original subsampled randomized Hadamard transform with uniform sampling (Tropp 2011);
- ISRHT-nps: our proposed method for ISRHT with norm-proportional sampling as defined in (6);
- ISRHT-top- r : our proposed deterministic sampling method for ISRHT which select r columns with largest Euclidean norms;
- ISRHT-supervised: our proposed method by incorporating label information as defined in (10);

We also include the results of using all features and two other popular data-dependent dimensionality reduction methods: PCA and Deterministic Leverage Score Sampling (DLS) (Papailiopoulos, Kyrillidis, and Boutsidis 2014) in Table 2. Due to high computational complexity of Singular Value Decomposition (SVD) in PCA and leverage score

sampling, we employed randomized SVD (Halko, Martinson, and Tropp 2011) in our experiments.

Experimental setup. The feature values for all data sets are linearly scaled to $[-1, 1]$. For dense datasets (mushrooms, usps-b and gisette), we directly apply our proposed ISRHT methods on them. For the high-dimensional sparse datasets (real-sim, rcv1-binary and news20-binary), we use an extension of our proposed methods that combines of sparse embedding and ISRHT for memory efficient projection as discussed in previous section. The regularization parameter C in linear SVM is chosen from $\{2^{-5}, 2^{-4}, \dots, 2^4, 2^5\}$ by 5-fold cross validation. The tradeoff parameter a for ISRHT-supervised is fixed to 1.0. Our experiments are performed on a server with Dual 6-core Intel Xeon 2.4GHz CPU and 128 GB RAM.

Experimental Results. The accuracy and training time of all algorithms are reported in Table 2. '-' is used to denote that PCA cannot be completed because of space complexity. The number of reduced dimension r is shown in the first column of the table. We also explore the impact of parameter r in Figure 2 and will be discussed later. The reported accuracies are the averaged accuracies on test data based on 15 repetitions. As shown in Table 2, our proposed ISRHT-(top- k) method achieves significant higher accuracies than other four data-independent random projection methods (i.e., Gaussian, Achlioptas, Sparse Embedding and SRHT) on all six datasets. The ISRHT-nps gets slightly better results than SRHT. Furthermore, by incorporating the label information, our proposed ISRHT-supervised method gets the best accuracy on all six datasets among the seven random projection methods. Using all features gets the best accuracy on all six datasets. PCA gets better accuracy than

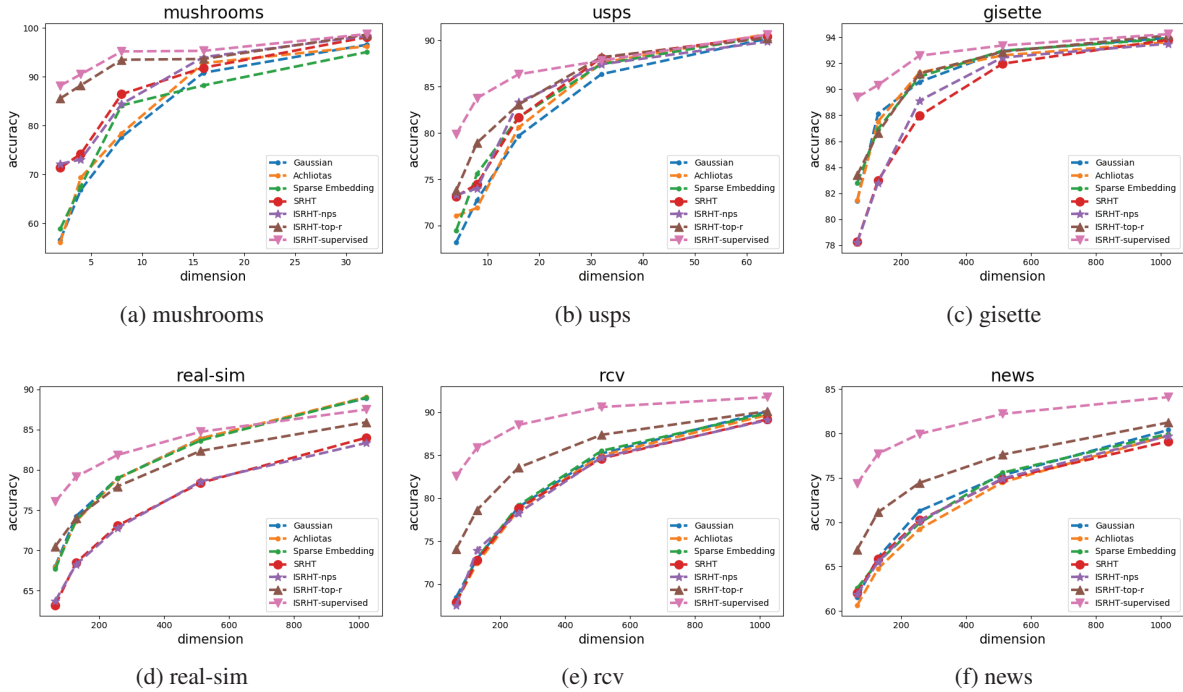


Figure 2: classification accuracies of different algorithms with different r

DLSS and other random projection methods while needs longer running time. We also observe that our proposed methods produce dense feature representation after random projection and therefore it needs longer running time than directly applying liblinear on high-dimensional sparse data.

With respect to the running time, all the random projection methods are more efficient than data-dependent dimension reduction method such as PCA and DLSS especially on large datasets. Among these random projection methods Gaussian and Achilotas are much slower than Sparse Embedding and SRHT since they need $O ndr$ time for projection. The running times of our proposed methods ISRHT-nps and ISRHT-top- r are very close to SRHT. The results demonstrate that our proposed methods only slightly increase the running time since the computing norm can be done very efficiently.

Impact of parameter r . We evaluate the impact of parameter r in our proposed algorithms. We show the accuracies of all the algorithms with respect to different r in Figure 2. We can observe that our proposed methods gets higher accuracy than other four methods. The accuracy improvement is large when the parameter r is set to relative small number. As expected, the difference will become smaller as the parameter r increases.

Stability Comparison Between ISRHT and SRHT. In this section, we would like to further investigate the stability of our proposed methods and SRHT with respect to different choices of the regularization parameter C in SVM and reduced dimension r . We evaluate all methods for different parameter combinations of C and r . Regularization param-

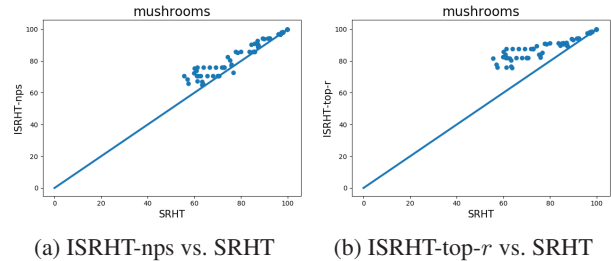


Figure 3: Classification Accuracies of SRHT (x-axis) and ISRHT (y-axis) for different parameters.

eter C is chosen from $\{2^{-5}, 2^{-4}, \dots, 2^5\}$ and reduced dimension r is chosen from $\{\frac{d}{2^6}, \frac{d}{2^5}, \dots, \frac{d}{2}\}$.

Due to the space limitation, we only show the results of comparing the prediction accuracies of ISRHT-nps and ISRHT-top- r with SRHT on mushrooms dataset in Figure 3. The accuracies plotted in Figure 3 are based on the average of 15 repetitions. For each parameter combination, we plot its corresponding accuracy of SRHT on the x -axis and the accuracy of our proposed methods on y -axis. So the points above $y = x$ line indicate an accuracy improvement of our proposed methods (y -axis) (ISRHT-nps or ISRHT-top- r) over SRHT(x -axis). Overall speaking, our proposed methods is more stable with different choices of parameter combinations.

Conclusion and Future work

In this paper, we propose to produce more effective low-dimensional embedding than original SRHT by using non-uniform sampling instead of uniform sampling in SRHT. To achieve this goal, we first analyze the effect of using SRHT for random projection in the context of linear SVM classification. Based on our analysis, we have proposed importance sampling and deterministic top- r sampling to improve the embedding. Secondly, we also propose a new sampling method to incorporate label information based on the idea of metric learning. We performed extensive experiments to evaluate our proposed non-uniform samplings methods. Our experimental results demonstrate that our proposed new methods can achieve better accuracy than original SRHT and other three popular random projection methods. Our results also demonstrate that our proposed method only slightly increase the running time but results in more effective embedding. In the future, we would like to extend our proposed ISRHT methods to nonlinear classification problems. Another interesting direction is to design data-dependent sparse embedding methods.

Acknowledgments

We would like to thank the anonymous reviewers for their insightful comments and valuable suggestions on our paper. This work was supported by NSFC 61906161.

References

- Achlioptas, D. 2003. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of computer and System Sciences* 66(4):671–687.
- Adelman, M., and Silberstein, M. 2018. Faster neural network training with approximate tensor operations. *arXiv preprint arXiv:1805.08079*.
- Ailon, N., and Liberty, E. 2009. Fast dimension reduction using rademacher series on dual bch codes. *Discrete & Computational Geometry* 42(4):615.
- Bingham, E., and Mannila, H. 2001. Random projection in dimensionality reduction: applications to image and text data. In *Proceedings of the seventh ACM SIGKDD*, 245–250. ACM.
- Boutsidis, C., and Gittens, A. 2013. Improved matrix algorithms via the subsampled randomized hadamard transform. *SIAM Journal on Matrix Analysis and Applications* 34(3):1301–1340.
- Boutsidis, C.; Zouzias, A.; and Drineas, P. 2010. Random projections for k -means clustering. In *NIPS*, 298–306.
- Chen, S.; Liu, Y.; Lyu, M. R.; King, I.; and Zhang, S. 2015. Fast relative-error approximation algorithm for ridge regression. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, 201–210. AUAI Press.
- Choromanski, K. M.; Rowland, M.; and Weller, A. 2017. The unreasonable effectiveness of structured random orthogonal embeddings. In *NeurIPS*, 219–228.
- Clarkson, K. L., and Woodruff, D. P. 2017. Low-rank approximation and regression in input sparsity time. *Journal of the ACM (JACM)* 63(6):54.
- Dasgupta, S., and Gupta, A. 1999. An elementary proof of the johnson-lindenstrauss lemma. *International Computer Science Institute, Technical Report* 22(1):1–5.
- Drineas, P.; Kannan, R.; and Mahoney, M. W. 2006. Fast monte carlo algorithms for matrices i: Approximating matrix multiplication. *SIAM Journal on Computing* 36(1):132–157.
- Freksen, C. B.; Kamma, L.; and Larsen, K. G. 2018. Fully understanding the hashing trick. In *Advances in Neural Information Processing Systems*, 5389–5399.
- Halko, N.; Martinsson, P.-G.; and Tropp, J. A. 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review* 53(2):217–288.
- Johnson, W. B.; Lindenstrauss, J.; and Schechtman, G. 1986. Extensions of lipschitz maps into banach spaces. *Israel Journal of Mathematics* 54(2):129–138.
- Jolliffe, I. 2011. *Principal component analysis*. Springer.
- Lan, L.; Wang, Z.; Zhe, S.; Cheng, W.; Wang, J.; and Zhang, K. 2019. Scaling up kernel SVM on limited resources: A low-rank linearization approach. *IEEE Trans. Neural Netw. Learning Syst.* 30(2):369–378.
- Liu, W.; Shen, X.; and Tsang, I. 2017. Sparse embedded k -means clustering. In *NIPS*, 3319–3327.
- Lu, Y.; Dhillon, P.; Foster, D. P.; and Ungar, L. 2013. Faster ridge regression via the subsampled randomized hadamard transform. In *NIPS*, 369–377.
- Mahoney, M. W., et al. 2011. Randomized algorithms for matrices and data. *Foundations and Trends® in Machine Learning* 3(2):123–224.
- Miika, S.; Ratsch, G.; Weston, J.; Scholkopf, B.; and Mullers, K.-R. 1999. Fisher discriminant analysis with kernels. In *Neural networks for signal processing IX: Proceedings of the 1999 IEEE signal processing society workshop (cat. no. 98th8468)*, 41–48. Ieee.
- Papailiopoulos, D.; Kyrillidis, A.; and Boutsidis, C. 2014. Provable deterministic leverage score sampling. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 997–1006. ACM.
- Paul, S.; Boutsidis, C.; Magdon-Ismael, M.; and Drineas, P. 2013. Random projections for support vector machines. In *Artificial intelligence and statistics*, 498–506.
- Paul, S.; Magdon-Ismael, M.; and Drineas, P. 2015. Feature selection for linear svm with provable guarantees. In *Artificial Intelligence and Statistics*, 735–743.
- Pourkamali-Anaraki, F.; Becker, S.; and Wakin, M. B. 2018. Randomized clustered nystrom for large-scale kernel machines. In *AAAI’18*.
- Sarlos, T. 2006. Improved approximation algorithms for large matrices via random projections. In *FOCS’06*, 143–152. IEEE.
- Tropp, J. A. 2011. Improved analysis of the subsampled randomized hadamard transform. *Advances in Adaptive Data Analysis* 3(01n02):115–126.
- Weinberger, K.; Dasgupta, A.; Langford, J.; Smola, A.; and Attenberg, J. 2009. Feature hashing for large scale multitask learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, 1113–1120. ACM.
- Xing, E. P.; Jordan, M. I.; Russell, S. J.; and Ng, A. Y. 2003. Distance metric learning with application to clustering with side-information. In *NIPS*, 521–528.
- Xu, Y.; Yang, H.; Zhang, L.; and Yang, T. 2017. Efficient non-oblivious randomized reduction for risk minimization with improved excess risk guarantee. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Zhang, L.; Mahdavi, M.; Jin, R.; Yang, T.; and Zhu, S. 2013. Recovering the optimal solution by dual random projection. In *Conference on Learning Theory*, 135–157.