# Efficient Non-oblivious Randomized Reduction for Risk Minimization with Improved Excess Risk Guarantee

**Yi Xu**[1], **Haiqin Yang**[2], **Lijun Zhang**[3], **Tianbao Yang**[1]
[1]Department of Computer Science, The University of Iowa, Iowa City, IA 52242, USA
[2]Department of Computing, Hang Seng Management College, Hong Kong
[3]National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China
{yi-xu, tianbao-yang}@uiowa.edu, hqyang@ieee.org, zhanglj@lamda.nju.edu.cn

## Abstract

In this paper, we address learning problems for high dimensional data. Previously, oblivious random projection based approaches that project high dimensional features onto a random subspace have been used in practice for tackling high-dimensionality challenge in machine learning. Recently, various non-oblivious randomized reduction methods have been developed and deployed for solving many numerical problems such as matrix product approximation, low-rank matrix approximation, etc. However, they are less explored for the machine learning tasks, e.g., classification. More seriously, the theoretical analysis of excess risk bounds for risk minimization, an important measure of generalization performance, has not been established for non-oblivious randomized reduction methods. It therefore remains an open problem what is the benefit of using them over previous oblivious random projection based approaches. To tackle these challenges, we propose an algorithmic framework for employing non-oblivious randomized reduction method for general empirical risk minimizing in machine learning tasks, where the original high-dimensional features are projected onto a random subspace that is derived from the data with a small matrix approximation error. We then derive the first excess risk bound for the proposed non-oblivious randomized reduction approach without requiring strong assumptions on the training data. The established excess risk bound exhibits that the proposed approach provides much better generalization performance and it also sheds more insights about different randomized reduction approaches. Finally, we conduct extensive experiments on both synthetic and real-world benchmark datasets, whose dimension scales to $O(10^7)$, to demonstrate the efficacy of our proposed approach.

## Introduction

Recently, the scale and dimensionality of data associated with machine learning and data mining applications have seen unprecedented growth, spurring the BIG DATA research and development. Learning from large-scale ultrahigh-dimensional data remains a computationally challenging problem. The big size of data not only increases the memory footprint but also increases the computational costs pertaining to optimization. A popular approach for addressing the high-dimensionality challenge is to perform dimensionality reduction. Nowadays, randomized reduction methods are emerging to be attractive for dimensionality reduction. Compared with traditional dimensionality reduction methods (e.g., PCA and LDA), randomized reduction methods (i) can lead to simpler algorithms that are easier to analyze (Mahoney 2011); (ii) can often be organized to exploit modern computational architectures better than classical dimensional reduction methods (Halko, Martinsson, and Tropp 2011); (iii) can be more efficient without loss in efficacy (Paul et al. 2013).

Generally, randomized reduction methods can be cast into two types: the first type of methods reduces a set of high-dimensional vectors into a low dimensional space independent of each other. These methods usually sample a random matrix independent of the data and then use it to reduce the dimensionality of the data. The second type of methods projects a set of vectors (in the form of a matrix) onto a subspace such that the original matrix can be well reconstructed from the projected matrix and the subspace. Therefore, the subspace to which the data is projected depends on the original data. These methods have been deployed for solving many numerical problems related to matrices, e.g., matrix product approximation, low-rank matrix approximation, approximate singular value decomposition (Boutsidis and Gittens 2013; Halko, Martinsson, and Tropp 2011). To differentiate these two types of randomized reduction methods, we refer to the first type as oblivious randomized reduction, and refer to the second type as non-oblivious randomized reduction. We note that in literature oblivious and non-oblivious are used interchangeably with data-independent and data-dependent. Here, we use the terminology commonly appearing in matrix analysis and numerical linear algebra due to that the general excess risk bound depends on the matrix approximation error.

However, we have not seen any comprehensive study on the statistical property (in particular the excess risk bound) of these randomized reduction methods applied to risk minimization in machine learning. The excess risk bound measures the generalization performance of a learned model compared to the optimal model from a class that has the best generalization performance. The excess risk bounds facilitate a better understanding of different learning algorithms and have the potential to guide us to design better algorithms (Kukliansky and Shamir 2015). It is worth

noting that several studies have been devoted to understanding the theoretical properties of oblivious randomized reduction methods applied to classification and regression problems. For example, (Blum 2005; Shi et al. 2012; Paul et al. 2013) analyzed the preservation of the margin of SVM based classification methods with randomized dimension reduction. (Zhang et al. 2014; Yang et al. 2015; Pilanci and Wainwright 2015) studied the problem from the perspective of optimization. Nonetheless, these results are limited in the sense that (i) they focus on only oblivious randomized reduction where the data is projected onto a random subspace independent of the data; (ii) they depend heavily on strong assumptions of the training data or the problem, e.g., low-rank of the data matrix, linear separability of training examples, or the sparsity of optimal solution, and (iii) some of these results do not directly carry over to the excess risk bounds.

To tackle the above challenges, we propose an algorithmic framework for employing non-oblivious randomized reduction (NOR) method to project the original high-dimensional features onto a random subspace that is derived from the original data. We study and establish the excess risk bound of the presented randomized algorithms for risk minimization. Different from previous results for oblivious randomized reduction methods, our theoretical analysis does not require assumptions of the training data or the problem, such as low-rank of the data matrix, linear separability of training examples, and the sparsity of optimal solution. When the data matrix is of low-rank or has a fast spectral decay, the excess risk bound of NOR is much better than that of oblivious randomized reduction based methods. Empirical studies on synthetic and real data sets corroborate the theoretical results and demonstrate the effectiveness of the proposed methods.

## Related Work

In literature, tremendous studies are devoted to non-oblivious randomized reduction in matrix applications. The focus of these studies is to establish matrix approximation error or the recovery error of the solution (e.g., in least-squares regression). Few studies have examined their properties for risk minimization in machine learning. For oblivious randomized reduction methods, there exist some theoretical work trying to understand their impact on prediction performance (Blum 2005; Shi et al. 2012; Paul et al. 2013). This work differentiates from these studies in that we focus on the statistical property (the generalization property) of non-oblivious randomized reduction for expected risk minimization.

We employ tools in statistical learning theory to study the excess risk bounds of randomized reduction and results from randomized matrix theory to understand the order of the excess risk bound. A popular method for expected risk minimization is regularized empirical risk minimization (Vapnik 1998). The excess risk bounds of regularized empirical risk minimization have been well understood. In general, given a sample of size $n$ it can achieve a risk bound of $O(1/\sqrt{n})$. Under some special conditions (e.g., low noise condition) this bound can be further improved (Bousquet, Boucheron, and Lugosi 2003). However, it is still not entirely clear what is the order of excess risk for learning from randomized dimensionality reduced data. The recovery result from (Zhang et al. 2014; Yang et al. 2015) could end up with an order of $O(1/\sqrt{m})$ excess risk for oblivious randomized reduction, where $m$ is the reduced dimensionality. However, it relies on strong assumptions of the data. (Durrant and Kaban 2013) proved the generalization error of the linear classifier trained on randomly projected data by oblivious randomized reduction, which is upper bounded by the training error of the classifier learned in the original feature space by empirical risk minimization plus the VC-complexity in the projection space (proportional to $O(1/\sqrt{m})$ and plus terms depending on the average flipping probabilities on the training points defined as $(1/n)\sum_{i=1}^{n}\Pr(sign(\mathbf{w}_n^T A^\intercal A\mathbf{x}_i) \neq sign(\mathbf{w}_n^\intercal \mathbf{x}_i))$, where $\mathbf{w}_n$ is a model learned from the original data by empirical risk minimization. However, the order of the average flipping probabilities is generally unknown.

Random sampling (in particular uniform sampling) has been used in the Nyström method for approximating a big kernel matrix. There are some related work focusing on the statistical properties of the Nyström based kernel method (Yang et al. 2012; Bach 2013; Jin et al. 2013; Alaoui and Mahoney 2015). We note that the presented empirical risk minimization with non-oblivious randomized reduction using random sampling is similar to using the Nyström approximation on the linear kernel. However, in the present work besides random sampling, we also study other efficient randomized reduction methods using different random matrices. By leveraging recent results of these randomized reduction methods we are able to obtain better performance than using random sampling.

## Preliminaries

Let $(\mathbf{x}, y)$ denote a feature vector and a label that follow a distribution $\mathcal{P} = \mathcal{P}(\mathbf{x}, y)$, where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$ and $y \in \mathcal{Y}$. In the sequel, we will focus on $\mathcal{Y} = \{+1, -1\}$ and $\mathcal{Y} = \mathbb{R}$. However, we emphasize that the results are applicable to other problems (e.g., multi-class and multi-label classification). We denote by $\ell(z, y)$ a non-negative loss function that measures the inconsistency between a prediction $z$ and the label $y$. Let $\mathbf{w} \in \mathbb{R}^d$, then by assuming a linear model $z = \mathbf{w}^\intercal \mathbf{x}$ for prediction, the risk minimization problem in machine learning is to solve following problem:

$$\mathbf{w}_* = \arg\min_{\mathbf{w}\in\mathbb{R}^d} \mathrm{E}_{\mathcal{P}}[\ell(\mathbf{w}^\intercal \mathbf{x}, y)] \tag{1}$$

where $\mathrm{E}_{\mathcal{P}}[\cdot]$ denotes the expectation over $(\mathbf{x}, y) \sim \mathcal{P}$.

Let $\mathcal{A}$ be an algorithm that learns an approximate solution $\mathbf{w}_n$ from a sample of size $n$, i.e., $\{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_n, y_n)\}$. The excess risk of $\mathbf{w}_n$ is defined as the difference between the expected risk of the solution $\mathbf{w}_n$ and that of the optimal solution $\mathbf{w}_*$:

$$\mathrm{ER}(\mathbf{w}_n, \mathbf{w}_*) = \mathrm{E}_{\mathcal{P}}[\ell(\mathbf{w}_n^\intercal \mathbf{x}, y)] - \mathrm{E}_{\mathcal{P}}[\ell(\mathbf{w}_*^\intercal \mathbf{x}, y)] \tag{2}$$

A popular method for learning an approximate solution $\mathbf{w}_n$ is based on regularized empirical risk minimization (ERM),

i.e.,

$$\mathbf{w}_n = \arg\min_{\mathbf{w}\in\mathbb{R}^d} \frac{1}{n}\sum_{i=1}^n \ell(\mathbf{w}^\intercal \mathbf{x}_i, y_i) + \frac{\lambda}{2}\|\mathbf{w}\|_2^2 \quad (3)$$

The ERM problem is sometimes solved by solving its dual problem:

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

where $\alpha$ is usually called dual variable, $\ell_i^*(\alpha) = \max_z \alpha z - \ell(z, y_i)$ is the conjugate dual of the loss function, and $X = (\mathbf{x}_1, \ldots, \mathbf{x}_n) \in \mathbb{R}^{d\times n}$ is the data matrix. With $\alpha_*$, we have $\mathbf{w}_n = -\frac{1}{\lambda n} X\alpha_*$.

## Oblivious Randomized Reduction

In this section, we present an excess risk bound of oblivious randomized reduction building on previous theoretical results to facilitate the comparison with our result of non-oblivious randomized reduction. The idea of oblivious randomized reduction is to reduce a high-dimensional feature vector $\mathbf{x} \in \mathbb{R}^d$ to a low dimensional vector by $\widehat{\mathbf{x}} = A\mathbf{x} \in \mathbb{R}^m$, where $A \in \mathbb{R}^{m\times d}$ is a random matrix that is independent of the data. A traditional approach is to use a Gaussian matrix with each entry independently sampled from a normal distribution with mean zero and variance $1/m$ (Dasgupta and Gupta 2003). Recently, many other types of random matrix $A$ are proposed that lead to much more efficient computation of reduction, including subsampled randomized Hadamard transform (SRHT) (Boutsidis and Gittens 2013) and random hashing (RH) (Kane and Nelson 2014). The key property of $A$ that plays an important role in the analysis is that it should preserve the Euclidean length of a high-dimensional vector with a high probability, which is stated formally in Johnson-Lindenstrauss (JL) lemma below.

**Lemma 1 (JL Lemma)** *For any $0 < \epsilon$, $\delta < 1/2$, there exists a probability distribution on matrices $A \in \mathbb{R}^{m\times d}$ such that there exists a small universal constant $c > 0$ and for any fixed $\mathbf{x} \in \mathbb{R}^d$, with a probability at least $1-\delta$, we have*

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

The key consequence of the JL lemma is that we can reduce a set of $d$-dimensional vectors into a low dimensional space with a reduced dimensionality independent of $d$ such that the pairwise distance between any two points can be well preserved.

Given the JL transform $A \in \mathbb{R}^{m\times d}$, the problem can be imposed as,

$$\min_{\mathbf{v}\in\mathbb{R}^m} \mathrm{E}_{\mathcal{P}}[\ell(\mathbf{v}^\intercal A\mathbf{x}, y)] \quad (5)$$

Previous studies have focused on using the ERM of the above problem

$$\min_{\mathbf{v}\in\mathbb{R}^m} \frac{1}{n}\sum_{i=1}^n \ell(\mathbf{v}^\intercal \widehat{\mathbf{x}}_i, y_i) + \frac{\lambda}{2}\|\mathbf{v}\|_2^2 \quad (6)$$

to learn a model in the reduced feature space or using its dual solution to recover a model in the original high-dimensional space:

$$\widehat{\alpha} = \arg\max_{\alpha\in\mathbb{R}^n} -\frac{1}{n}\sum_{i=1}^n \ell_i^*(\alpha_i) - \frac{1}{2\lambda n^2}\alpha^\intercal \widehat{X}^\intercal \widehat{X}\alpha \quad (7)$$

where $\widehat{X} = (\widehat{\mathbf{x}}_1, \ldots, \widehat{\mathbf{x}}_n) \in \mathbb{R}^{m\times n}$. For example, (Zhang et al. 2014) proposed a dual recovery approach to recover a model in the original high-dimensional space that is close to the optimal solution $\mathbf{w}_n$ in (3). The dual recovery approach consists of two steps (i) the first step obtains an approximate dual solution $\widehat{\alpha} \in \mathbb{R}^n$ by solving the dual problem in (7), and (ii) the second step recovers a high-dimensional model by $\widehat{\mathbf{w}}_n = -\frac{1}{\lambda n} X\widehat{\alpha}$. By making a low-rank assumption of the data matrix, they established a recovery error $\|\mathbf{w}_n - \widehat{\mathbf{w}}_n\|_2$ in the order of $O(\sqrt{r/m}\|\mathbf{w}_n\|_2)$, where $r$ represents the rank of the data matrix. The theory has been generalized to full rank data matrix but with an additional assumption that the optimal primal solution $\mathbf{w}_n$ or the optimal dual solution $\alpha_*$ is sparse (Zhang et al. 2014; Yang et al. 2015). A similar order $O(\sqrt{r/m}\|\mathbf{w}_n\|_2)$ of recovery error was established, where $r$ represents the number of non-zero elements in the optimal solution. The proposition below exhibits the excess risk bound building on the recovery error.

**Proposition 1** *Suppose $\ell(z, y)$ is Lipschitz continuous and $A$ is a JL transform. Let $\widehat{\mathbf{w}}_n$ denote a recovered model by an ERM approach such that $\|\mathbf{w}_n - \widehat{\mathbf{w}}_n\|_2 \leq O(\sqrt{r/m}\|\mathbf{w}_n\|_2)$. Then*

$$\begin{aligned} ER(\widehat{\mathbf{w}}_n, \mathbf{w}_*) &\triangleq \mathrm{E}_{\mathcal{P}}[\ell(\widehat{\mathbf{w}}_n^\intercal \mathbf{x}, y)] - \mathrm{E}_{\mathcal{P}}[\ell(\mathbf{w}_*^\intercal \mathbf{x}, y)] \\ &\leq O(\sqrt{r/m}\|\mathbf{w}_n\|_2 + 1/\sqrt{n}) \end{aligned}$$

**Remark:** In the above bound, we omit dependence on upper bound of the data norm $\|\mathbf{x}\| \leq R$. Although the synthesis of the proposition and previous recovery error analysis can give us guarantee on the excess risk bound, it relies on certain assumptions of the data, which may not hold in practice.

## Non-Oblivious Randomized Reduction

The key idea of non-oblivious randomized reduction is to compute a subspace $\widehat{U} \in \mathbb{R}^{d\times m}$ from the data matrix $X \in \mathbb{R}^{d\times n}$ such that the projection of the data matrix to the subspace is close to the data matrix. To compute the subspace $\widehat{U}$, we first sample a random matrix $\Omega \in \mathbb{R}^{n\times m}$ and compute $Y = X\Omega \in \mathbb{R}^{d\times m}$. Then let $\widehat{U}$ be the left singular vector matrix of $Y$. This technique has been used in low-rank matrix approximation, matrix product approximation and approximate singular value decomposition (SVD) of a large matrix (Halko, Martinsson, and Tropp 2011). Various random matrices $\Omega$ can be used as long as the matrix approximation error defined below can be well bounded.

$$\|X - \widehat{U}\widehat{U}^\intercal X\|_2 = \|X - P_Y X\|_2 \quad (8)$$

where $P_Y = \widehat{U}\widehat{U}^\intercal$ denotes the projection to the subspace $\widehat{U}$. We defer more discussions on different random matrices and their impact on the excess risk bound to subsection "Matrix Approximation Error".

**Algorithm 1** ERM with Non-Oblivious Randomized Reduction (NOR)

1: Compute $Y = X\Omega \in \mathbb{R}^{d\times m}$, where $\Omega \in \mathbb{R}^{n\times m}$ is a random subspace embedding matrix
2: Compute SVD of $Y = \widehat{U}\widehat{\Sigma}\widehat{V}^\intercal$, where $\widehat{U} \in \mathbb{R}^{d\times m}$
3: Compute the reduced data by $\widehat{X} = \widehat{U}^\intercal X \in \mathbb{R}^{m\times n}$
4: Solve the reduced problem in Eqn. (9)
5: Output $\widehat{\mathbf{v}}_n$

Next, we focus on the non-oblivious reduction defined by $\widehat{U}$ for risk minimization. Let $\widehat{\mathbf{x}}_i = \widehat{U}^\intercal \mathbf{x}_i$ denote the reduced feature vector and $\widehat{X} = \widehat{U}^\intercal X \in \mathbb{R}^{m\times n}$ be the reduced data matrix. We propose to solve the following ERM problem:

$$\widehat{\mathbf{v}}_n = \arg\min_{\mathbf{v}\in\mathbb{R}^m} \frac{1}{n}\sum_{i=1}^n \ell(\mathbf{v}^\intercal \widehat{\mathbf{x}}_i, y_i) + \frac{\lambda}{2}\|\mathbf{v}\|_2^2. \quad (9)$$

To understand the non-oblivious randomized reduction for ERM, we first see that the problem above is equivalent to

$$\min_{\mathbf{w}=\widehat{U}\mathbf{v},\mathbf{v}\in\mathbb{R}^m} \frac{1}{n}\sum_{i=1}^n \ell(\mathbf{w}^\intercal \mathbf{x}_i, y_i) + \frac{\lambda}{2}\|\mathbf{w}\|_2^2.$$

due to that $\|\widehat{U}\mathbf{v}\|_2 = \|\mathbf{v}\|_2$. Compared to (3), we can see that the ERM with non-oblivious randomized reduction is restricting the model $\mathbf{w}$ to be $\widehat{U}\mathbf{v}$. Since $\widehat{U}$ can capture the top column space of $X$ due to the way it is constructed, and therefore the resulting model $\widehat{\mathbf{w}}_n = \widehat{U}\widehat{\mathbf{v}}_n$ is close to the top column space of $X$. Thus, we expect $\widehat{\mathbf{w}}_n$ to be close to $\mathbf{w}_n$. The procedure is described in details in Algorithm 1. We note that the SVD in step 2 can be computed efficiently for sparse data. We defer the details into the supplement.

## Excess Risk Bound

Here, we show an excess risk bound of the proposed ERM with non-oblivious randomized reduction. The logic of the analysis is to first derive the optimization error of the approximate model $\widehat{\mathbf{w}}_n = \widehat{U}\widehat{\mathbf{v}}_n$ and then explore the statistical learning theory to bound the excess risk. In particular, we will show that the optimization error and consequentially the excess risk is bounded by the matrix approximation error in (8). To simplify the presentation, we introduce some notations:

$$F(\mathbf{w}) = \frac{1}{n}\sum_{i=1}^n \ell(\mathbf{w}^\intercal \mathbf{x}_i, y_i) + \frac{\lambda}{2}\|\mathbf{w}\|^2 \quad (10)$$

$$\bar{F}(\mathbf{w}) = \mathrm{E}_{\mathcal{P}}[\ell(\mathbf{w}^\intercal \mathbf{x}, y)] + \frac{\lambda}{2}\|\mathbf{w}\|_2^2 \quad (11)$$

Next, we derive the optimization error of $\widehat{\mathbf{w}}_n = \widehat{U}\widehat{\mathbf{v}}_n$.

**Lemma 2** *Suppose the loss function is G-Lipschitz continuous. Let $\widehat{\mathbf{w}}_n = \widehat{U}\widehat{\mathbf{v}}_n$. We have*

$$F(\widehat{\mathbf{w}}_n) \leq F(\mathbf{w}_n) + \frac{G^2}{2\lambda n}\|X - P_Y X\|_2^2$$

The lemma below bounds the excess risk by the optimization error.

**Lemma 3 (Theorem 1 (Sridharan et al. 2008))** *Assume the loss function is G-Lipschitz continuous and $\|\mathbf{x}\|_2 \leq R$. Then, for any $\delta > 0$ and any $a > 0$, with probability at least $1-\delta$, we have that for any $\mathbf{w}_* \in \mathbb{R}^d$*

$$\bar{F}(\widehat{\mathbf{w}}_n) - \bar{F}(\mathbf{w}_*) \leq (1+a)(F(\widehat{\mathbf{w}}_n) - F(\mathbf{w}_n)) + \frac{8(1+1/a)G^2R^2(32+\log(1/\delta))}{\lambda n}$$

Using the lemma above and the the result in Lemma 2, we have the following theorem about the excess risk bound.

**Theorem 1** *Assume the loss function is G-Lipschitz continuous and $\|\mathbf{x}\|_2 \leq R$. Then, for any $\delta > 0$ and any $a > 0$, with probability at least $1-\delta$, we have that for any $\mathbf{w}_*$ such that $\|\mathbf{w}_*\|_2 \leq B$*

$$ER(\widehat{\mathbf{w}}_n, \mathbf{w}_*) \leq \frac{\lambda B^2}{2} + \frac{G^2(1+a)}{2\lambda n}\|X - P_Y X\|_2^2 + \frac{8(1+1/a)G^2R^2(32+\log(1/\delta))}{\lambda n}$$

*In particular, if we optimize $\lambda$ over the R.H.S., we obtain*

$$ER(\widehat{\mathbf{w}}_n, \mathbf{w}_*) \leq \frac{GB\sqrt{(1+a)}}{\sqrt{n}}\|X - P_Y X\|_2 + \frac{4GRB\sqrt{(1+1/a)(32+\log(1/\delta))}}{\sqrt{n}}$$

**Remark:** Note that the above theorem bounds the excess risk by the matrix approximation error. Thus, we can leverage state-of-the-art results on the matrix approximation to study the excess risk bound. Importantly, future results about matrix approximation can be directly plugged into the excess risk bound. When the data matrix is of low rank $r$, then if $m \geq \Omega(r\log r)$ the matrix approximation error can be made zero (see below). As a result, the excess risk bound of $\widehat{\mathbf{w}}_n$ is $O(1/\sqrt{n})$, the same to that of $\mathbf{w}_n$. In contrast, the excess risk bound in Proposition 1 of oblivious randomized reduction for ERM is $O(\sqrt{r/m})$ for the dual recovery approach under the low rank assumption.

## Matrix Approximation Error

In this subsection, we will present some recent results on the matrix approximation error of four commonly used randomized reduction operators $\Omega \in \mathbb{R}^{n\times m}$, i.e., random sampling (RS), random Gaussian (RG), subsampled randomized Hadamard transform (SRHT), and random hashing (RH), and discuss their impact on the excess risk bound. More details of these four randomized reduction operators can be found in (Yang et al. 2015). We first introduce some notations used in matrix approximation analysis. Let $r \leq \min(n,d)$ denote the rank of $X$ and $k \in \mathbb{N}^+$ such that $1 \leq k \leq r$. We write the SVD of $X \in \mathbb{R}^{d\times n}$ as $X = U_1\Sigma_1 V_1^\intercal + U_2\Sigma_2 V_2^\intercal$, where $\Sigma_1 \in \mathbb{R}^{k\times k}$, $\Sigma_2 \in \mathbb{R}^{(r-k)\times(r-k)}$, $U_1 \in \mathbb{R}^{d\times k}$, $U_2 \in \mathbb{R}^{d\times(r-k)}$, $V_1 \in \mathbb{R}^{n\times k}$ and $V_2 \in \mathbb{R}^{n\times(r-k)}$. We use $\sigma_1, \sigma_2, \ldots, \sigma_r$ to denote the singular values of $X$ in the descending order. Let $\mu_k$ denote the coherence measure of $V_1$ defined as $\mu_k = \frac{n}{k}\max_{1\leq i\leq n}\sum_{j=1}^k [V_1]_{ij}^2$.

**Theorem 2 (RS (Gittens 2011))** *Let $\Omega \in \mathbb{R}^{n\times m}$ be a random sampling matrix corresponding to sampling the columns of $X$ uniformly at random with or without replacement. If for any $\epsilon \in (0,1)$ and $\delta > 0$, $m$ satisfies $m \geq \frac{2\mu_k}{(1-\epsilon)^2} k \log \frac{k}{\delta}$, then with a probability at least $1-\delta$,*

$$\|X - P_Y X\|_2 \leq \sqrt{1 + \frac{n}{\epsilon m}}\sigma_{k+1}$$

**Remark:** The matrix approximation error using RS implies the excess risk bound of RS for risk minimization is dominated by $O(\frac{\sigma_{k+1}}{\sqrt{m}})$ provided $m \geq \Omega(\mu_k k \log k)$, which is in the same order in terms of $m$ to that in Proposition 1 of oblivious randomized reduction. However, random sampling is not guaranteed to work in oblivious randomized reduction since it does not satisfy the JL lemma in general (Yang et al. 2015) as required in Proposition 1. Moreover, if the data matrix is low rank such that $m \geq \Omega(\mu_r r \log r)$, then the excess risk bound of NOR with RS is $O(1/\sqrt{n})$, the same order to that of $\mathbf{w}_n$ learned from the original high-dimensional features.

**Theorem 3 (RG (Gittens and Mahoney 2013))** *Let $\Omega \in \mathbb{R}^{n\times m}$ be a random Gaussian matrix. If for any $\epsilon \in (0,1)$ and $k > 4$, $m$ satisfies $m \geq 2\epsilon^{-2} k \log k$, then with a probability at least $1 - 2k^{-1} - 4k^{-k/\epsilon^2}$,*

$$\|X - P_Y X\|_2 \leq O(\sigma_{k+1}) + O\left(\frac{\epsilon}{\sqrt{k \log k}}\right)\sqrt{\sum\nolimits_{j>k}\sigma_j^2}$$

**Remark:** We are interested in comparing the error bound of RG with that of RS. In the worse case, when the tail singular values are flat, then $\|X - P_Y X\|_2 \leq O(\sqrt{\frac{n}{m}}\sigma_{k+1})$, which is in the same order to that of RS. However, if the tail eigenvalues decay fast such that $\sqrt{\sum_{j>k}\sigma_j^2} \ll \sqrt{n}\sigma_{k+1}$, then the matrix approximation error could be much better than $O(\sqrt{\frac{n}{m}}\sigma_{k+1})$, and consequentially the excess risk bound could be much better than $O(\sigma_{k+1}/\sqrt{m})$ that is suffered by RS.

**Theorem 4 (SRHT (Boutsidis and Gittens 2013))** *Let $\Omega = \sqrt{\frac{n}{m}} DHP \in \mathbb{R}^{n\times m}$ be a SRHT with $P \in \mathbb{R}^{n\times m}$ being a random sampling matrix, $D \in \mathbb{R}^{n\times n}$ is a diagonal matrix with each entry sampled from $\{1,-1\}$ with equal probabilities and $H \in \mathbb{R}^{n\times n}$ is a normalized Hadamard transform. If for any $0 < \epsilon < 1/3$, $2 \leq k \leq r$ and $\delta \in (0,1)$, $m$ satisfies*

$$6C^2\epsilon^{-1}[\sqrt{k} + \sqrt{8\log(n/\delta)}]^2 \log(k/\delta) \leq m \leq n,$$

*then with a probability at least $1 - 5\delta$,*

$$\|X - P_Y X\|_2 \leq \left(4 + \sqrt{\frac{3\log(n/\delta)\log(r/\delta)}{m}}\right)\sigma_{k+1} + \sqrt{\frac{3\log(r/\delta)}{m}}\sqrt{\sum\nolimits_{j>k}\sigma_j^2}$$

*where $C$ is a universal constant.*

**Remark:** The order of the matrix approximation error of SRHT is similar to that of RG up to a logarithmic factor.

Finally, we summarize the matrix approximation error of the RH matrix $\Omega$. This has been studied in (Cohen, Nelson, and Woodruff 2015), in which RH is also referred to as sparse subspace embedding. We first describe the construction of random hashing matrix $\Omega$. Let $h_k(i) : [n] \to [m/s], k = 1,\ldots,s$ denote $s$ independent random hashing functions and let $\Omega = ((H_1 D_1)^\intercal, (H_2 D_2)^\intercal, \ldots, (H_s D_s)^\intercal)^\intercal \in \mathbb{R}^{m\times n}$ be a random matrix with a block of $s$ random hashing matrices, where $D_k \in \mathbb{R}^{n\times n}$ is a diagonal matrix with each entry sampled from $\{-1,+1\}$ with equal probabilities, and $H_k \in \mathbb{R}^{m/s,n}$ with $[H_k]_{j,i} = \delta_{j,h_k(i)}$. The following theorem below summarizes the matrix approximation error using such a random matrix $\Omega$.

**Theorem 5 (RH (Cohen, Nelson, and Woodruff 2015))** *For any $\delta \in (0,1)$ and $\epsilon \in (0,1)$. If $s = 1$ and $m = O(k/(\epsilon\delta))$ or $s = O(\log^3(k/\delta)/\sqrt{\epsilon})$ and $m = O(k\log^6(k/\delta)/\epsilon)$, then with a probability $1-\delta$*

$$\|X - P_Y X\|_2 \leq (1+\sqrt{\epsilon})\sigma_{k+1} + \sqrt{\frac{\epsilon}{k}}\sqrt{\sum\nolimits_{j>k}\sigma_j^2}$$

**Remark:** With the second choice of $s$ and $m$, the order of the matrix approximation error of RH is similar to that of SRHT up to a logarithmic factor.

To conclude this section, we can see that the excess risk bound of the ERM with non-oblivious randomized reduction is dominated by $O(1/\sqrt{m})$ in the worst case, and could be much better than RG, SRHT and RH if the tail singular values decay fast.

## Experiments

In this section, we provide empirical evaluations in support of the proposed algorithms and the theoretical analysis. We implement and compare the following algorithms: (i) NOR: ERM with non-oblivious randomized reduction; (ii) previous ERM approaches with oblivious randomized reduction, including two dual recovery approaches, namely random projection with dual recovery (RPDR) (Zhang et al. 2014), and dual-sparse regularized randomized (DSRR) approach (Yang et al. 2015), and the pure random projection (RP) (Paul et al. 2013). We also implement and compare three randomized reduction operators for these different approaches, i.e., RH, RG and RS [1]. For RH, we use only one block of random hashing matrix (i.e., $s = 1$). A similar result to Therorem 5 can be established for one-block of random hashing but with a constant success probability (Nelson and Nguyen 2012). The loss function for the binary classification problem is the hinge loss and for the multi-class classification problem is the softmax loss.

Experiments are conducted on four real-world datasets and three synthetic datasets. The four real-world datasets are described in Table 1. To generate synthetic data, we first draw a random standard Gaussian matrix $M \in \mathbb{R}^{d\times n}$ and

[1] We do not report the performance of SRHT because it has similar performance to RH but it is less efficient than RH.

Table 1: Statistics of real datasets

| Name | #Training | #Testing | #Features | #Classes |
|---|---|---|---|---|
| RCV1.b | 677,399 | 20,242 | 47,236 | 2 |
| Splice | 1,000,000 | 4,627,840 | 12,495,340 | 2 |
| RCV1.m | 15,564 | 518,571 | 47,236 | 53 |
| News20 | 15,935 | 3,993 | 62,061 | 20 |

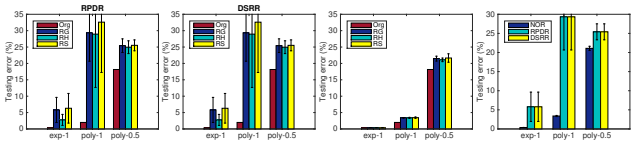


(a) RPDR and DSRR (b) NOR (c) RH

Figure 1: (a) RPDR and DSRR with different randomized reduction operators. (b) NOR with different randomized reduction operators. (c) Different approaches with RH.

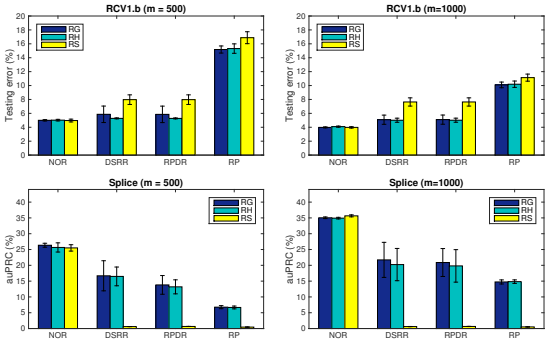


Figure 2: Testing performance of different approaches on RCV1.b and Splice datasets.

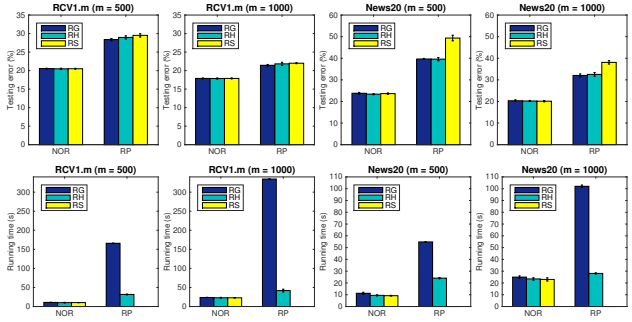


Figure 3: Testing performance and running time of different approaches on RCV1.m and News20 datasets.

compute its SVD $M = USV^\intercal$. Then we construct singular values following three different decay: an exponential decay (exp-$\tau$) with $\sigma_i = e^{-i\tau}, (\tau = 1)$ and polynomial decay (poly-$\tau$) with $\sigma_i = i^{-\tau}, (\tau = 0.5, 1)$. This will generate 3 synthetic datasets. We compute a base data matrix by $X_b = \sqrt{n}U\Sigma V^\intercal$, where $\Sigma = \text{diag}\{\sigma_1, \ldots, \sigma_d\}$. Then the binary labels are computed by $\mathbf{y} = sign(X_b^\intercal \mathbf{w})$, where $\mathbf{w} \in \mathbb{R}^d$ is a standard Gaussian random vector. To increase the difficulty of the problem, we add some Gaussian random features to each data in $X_b$ and form a full data matrix $X \in \mathbb{R}^{(d+t)\times n}$. We use the first $90\%$ examples as training data and the remaining $10\%$ examples as testing data. In particular, we generate the synthetic datasets with $d = 1000, n = 10^5, t = 10$. We note that the synthetic data is not high-dimensional and it is solely for verifying the proposed approach and analysis. We perform data reduction to reduce features to the dimensionality of $m = 100$.

We first compare the performance of NOR, RPDR and DSRR with different randomized reduction operators on the synthetic datasets in order to verify the excess risk bounds established in Proposition 1 and subsection "Matrix Approximation Error" (MAE). The results are shown in Figure 1, where we also include the performance of SVM on the original features (denoted by Org). From the results, we can observe that (i) when the singular values follow an exponential decay (which yields almost low-rank data matrices), NOR performs almost the same to SVM on the original data; however the two recovery approaches RPDR and DSRR perform much worse than SVM, verifying our theoretical analysis in Proposition 1 and subsection "MAE"; (ii) The performance of NOR decreases gradually as the decay of singular values becomes slower, which is consistent with the theoretical results in subsection "MAE"; (iii) for NOR, RS is comparable to RH and RG when the decay of singular-values is fast, but is slightly worse when the decay of singular values becomes slower. This is also expected according to the discussions in subsection "MAE"; (iv) NOR always performs better than RPDR and DSRR. One reason that RPDR and DSRR do not perform well on these synthetic data sets is that the recovered dual solution is not accurate because that the data has noise. In contrast, NOR is much more robust to noise.

Secondly, we present some experimental results on two binary classification real datasets, namely Reuters Text Categorization both for binary version (RCV1.b) and Splice Site Recognition (Splice), which are tested in previous studies (Sonnenburg and Franc 2010). For Splice dataset, we evaluate different algorithms by computing the same measure, namely area under precision recall curve (auPRC), as in (Sonnenburg and Franc 2010). We compare NOR with RPDR, DSRR and RP. The results are shown in Figure 2. We can see that when $m$ increases, the testing error/auPRC is monotonically decreasing/increasing. Comparing with other three algorithms, NOR has the best performance. In addition, RS does not work well for oblivious randomized reduction approaches (RP, RPDR and DSRR), but performs similarly to other randomized operators for NOR, which is consistent with our analysis (i.e., RS is not a JL transform as required in Proposition 1 for RPDR and DSRR; however, RS provides guarantee on the matrix approximation error that renders NOR work).

Finally, we compare NOR with RP on RCV.m and News20 datasets for multi-class classification. The results are shown in Figure 3 (upper panel). We can see that NOR

clearly outperforms RP. In addition, running time results [2] are reported in Figure 3 (lower panel). The running time consists of the reducation time and the optimization time in the reduced feture space. The results show that (i) NOR is more efficient than RP; (ii) RH and/or RS are much more efficient than RG for a certain approach. It is interesting to note that the total running time of NOR is less than RP. The reason is that the optimization of NOR is more efficient than RP [3] due to that the new data of NOR is better suited for classification (higher prediction performance), making the optimization easier, though NOR has slightly higher data reduction time than RP.

## Conclusions

In this paper, we have established the excess risk bound of non-oblivious randomized reduction method for risk minimization problems. More importantly, the new excess risk bound does not require stringent assumptions of the data and the loss functions, which is nontrivial and significant theoretical results. The empirical studies on synthetic datasets and real datasets validate our theoretical analysis and also demonstrate the effectiveness of the proposed non-oblivious randomized reduction approach.

## Acknowlegements

We thank the anonymous reviewers for their helpful comments. Y. Xu and T. Yang are partially supported by National Science Foundation (IIS-1463988, IIS-1545995). L. Zhang is partially supported by NSFC (61603177) and JiangsuSF (BK20160658).

## References

Alaoui, A., and Mahoney, M. W. 2015. Fast randomized kernel ridge regression with statistical guarantees. In *NIPS*, 775–783.

Bach, F. 2013. Sharp analysis of low-rank kernel matrix approximations. In *COLT*, 185–209.

Blum, A. 2005. Random projection, margins, kernels, and feature-selection. In *Proceedings of the 2005 international conference on Subspace, Latent Structure and Feature Selection*, 52–68. Springer-Verlag.

Bousquet, O.; Boucheron, S.; and Lugosi, G. 2003. Introduction to statistical learning theory. In *Advanced Lectures on Machine Learning*, volume 3176, 169–207. Springer.

Boutsidis, C., and Gittens, A. 2013. Improved matrix algorithms via the subsampled randomized hadamard transform. *SIAM J. Matrix Analysis Applications* 34(3):1301–1340.

Cohen, M. B.; Nelson, J.; and Woodruff, D. P. 2015. Optimal approximate matrix product in terms of stable rank. *CoRR* abs/1507.02268.

Dasgupta, S., and Gupta, A. 2003. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures & Algorithms* 22(1):60–65.

Durrant, R. J., and Kaban, A. 2013. Sharp generalization error bounds for randomly-projected classifiers. In *ICML*, 693–701.

Gittens, A., and Mahoney, M. 2013. Revisiting the nystrom method for improved large-scale machine learning. In *ICML*, 567–575.

Gittens, A. 2011. The spectral norm error of the naive nystrom extension. *CoRR* abs/1110.5305.

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.

Jin, R.; Yang, T.; Mahdavi, M.; Li, Y.-F.; and Zhou, Z.-H. 2013. Improved bounds for the nyström method with application to kernel classification. *IEEE Transactions on Information Theory* 59(10):6939–6949.

Kane, D. M., and Nelson, J. 2014. Sparser johnson-lindenstrauss transforms. *Journal of the ACM* 61:4:1–4:23.

Kukliansky, D., and Shamir, O. 2015. Attribute efficient linear regression with distribution-dependent sampling. In *ICML*, 153–161.

Mahoney, M. W. 2011. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning* 3(2):123–224.

Nelson, J., and Nguyen, H. L. 2012. OSNAP: faster numerical linear algebra algorithms via sparser subspace embeddings. *CoRR* abs/1211.1002.

Paul, S.; Boutsidis, C.; Magdon-Ismail, M.; and Drineas, P. 2013. Random projections for support vector machines. In *AISTATS*, 498–506.

Pilanci, M., and Wainwright, M. J. 2015. Randomized sketches of convex programs with sharp guarantees. *IEEE Transactions on Information Theory* 61(9):5096–5115.

Shi, Q.; Shen, C.; Hill, R.; and Hengel, A. 2012. Is margin preserved after random projection? In *ICML*, 591–598.

Sonnenburg, S., and Franc, V. 2010. Coffin: A computational framework for linear svms. In *ICML*, 999–1006.

Sridharan, K.; Shalev-Shwartz, S.; ; and Srebro, N. 2008. Fast rates for regularized objectives. In *NIPS*, 1545–1552.

Vapnik, V. N. 1998. *Statistical learning theory*. Wiley, 1 edition.

Yang, T.; Li, Y.; Mahdavi, M.; Jin, R.; and Zhou, Z. 2012. Nyström method vs random fourier features: A theoretical and empirical comparison. In *NIPS*, 485–493.

Yang, T.; Zhang, L.; Jin, R.; and Zhu, S. 2015. Theory of dual-sparse regularized randomized reduction. In *ICML*, 305–314.

Zhang, L.; Mahdavi, M.; Jin, R.; Yang, T.; and Zhu, S. 2014. Random projections for classification: A recovery approach. *IEEE Transactions on Information Theory* 60(11):7300–7316.

[2] We do not report the running time for RP using RS as randomized reduction operator since it has the worst performance

[3] We terminate the optimization for both methods by the same criterion, i.e., the duality gap is less than $10^{-3}$.