

Wishart Mechanism for Differentially Private Principal Components Analysis

Wuxuan Jiang, Cong Xie and Zhihua Zhang

Department of Computer Science and Engineering
Shanghai Jiao Tong University, Shanghai, China
jiangwuxuan@gmail.com, {xcgoner, zhihua}@sjtu.edu.cn

Abstract

We propose a new input perturbation mechanism for publishing a covariance matrix to achieve $(\epsilon, 0)$ -differential privacy. Our mechanism uses a Wishart distribution to generate matrix noise. In particular, we apply this mechanism to principal component analysis (PCA). Our mechanism is able to keep the positive semi-definiteness of the published covariance matrix. Thus, our approach gives rise to a general publishing framework for input perturbation of a symmetric positive semidefinite matrix. Moreover, compared with the classic Laplace mechanism, our method has better utility guarantee. To the best of our knowledge, the Wishart mechanism is the best input perturbation approach for $(\epsilon, 0)$ -differentially private PCA. We also compare our work with previous exponential mechanism algorithms in the literature and provide near optimal bound while having more flexibility and less computational intractability.

1 Introduction

Plenty of machine learning tasks deal with sensitive information such as financial and medical data. A common concern regarding data security arises on account of the rapid development of data mining techniques. Several data privacy definitions are proposed in the literature. Among them differential privacy (DP) has been widely used (Dwork et al. 2006). Differential privacy controls the fundamental quantity of information that can be revealed with changing one individual. Beyond a concept in database security, differential privacy has been used by many researchers to develop privacy-preserving learning algorithms (Chaudhuri and Monteleoni 2009; Chaudhuri, Monteleoni, and Sarwate 2011; Bojarski et al. 2014). Indeed, this class of algorithms is applied to a large number of machine learning models including logistic regression (Chaudhuri and Monteleoni 2009), support vector machine (Chaudhuri, Monteleoni, and Sarwate 2011), random decision tree (Bojarski et al. 2014), etc. Accordingly, these methods can protect the raw data even though the output and algorithm itself are published.

Differential privacy (DP) aims to hide the individual information while keeping basic statistics of the whole dataset.

A simple idea to achieve this purpose is to add some special noise to the original model. After that, the attacker, who has two outputs generated by slightly different inputs, cannot distinguish whether the output change comes from the artificial noise or input difference. However, the noise might influence the quality of regular performance of the model. Thus, we should carefully trade off between privacy and utility.

No matter what the procedure is, a query, a learning algorithm, a game strategy or something else, we are able to define differential privacy if this procedure takes a dataset as input and returns the corresponding output. In this paper, we study the problem of designing differential private principal component analysis (PCA). PCA reduces the data dimension while keeping the optimal variance. More specifically, it finds a projection matrix by computing a low rank approximation to the sample covariance matrix of the given data points.

Privacy-preserving PCA is a well-studied problem in the literature (Dwork et al. 2014; Hardt and Roth 2012; 2013; Hardt and Price 2014; Blum et al. 2005; Chaudhuri, Sarwate, and Sinha 2012; Kapralov and Talwar 2013). It outputs a noisy projection matrix for dimension reduction while preserving the privacy of any single data point. The extant privacy-preserving PCA algorithms have been devised based on two major features: the notion of differential privacy and the stage of randomization. Accordingly, the privacy-preserving PCA algorithms can be divided into distinct categories.

The notion of differential privacy has two types: $(\epsilon, 0)$ -DP (also called pure DP) and (ϵ, δ) -DP (also called approximate DP). (ϵ, δ) -DP is a weaker version of $(\epsilon, 0)$ -DP as the former allows the privacy guarantee to be broken with tiny probability (more precisely, δ). In the seminal work on privacy-preserving PCA (Dwork et al. 2014; Hardt and Roth 2012; 2013; Hardt and Price 2014; Blum et al. 2005), the authors used the notion of (ϵ, δ) -DP. In contrast, there is only a few work (Chaudhuri, Sarwate, and Sinha 2012; Kapralov and Talwar 2013), which is based on $(\epsilon, 0)$ -DP.

In terms of the stage of randomization, there are two mainstream classes of approaches. The first is randomly computing the eigenspace (Hardt and Roth 2013; Hardt and Price 2014; Chaudhuri, Sarwate, and Sinha 2012; Kapralov and Talwar 2013). The noise is added in the computing pro-

cedure. An alternative way is directly adding noise to the covariance matrix. Then one runs the non-private eigenspace computing algorithm to produce the output. This class of approaches is called input perturbation (Blum et al. 2005; Dwork et al. 2014). The input perturbation algorithms publish a noisy sample covariance matrix before computing the eigenspace. Thus, any further operation on the noisy covariance matrix does not violate privacy guarantee. So far as the flexibility is concerned, the input perturbation has better performance because it is not limited only to computing eigenspace. Besides, the input perturbation approach is efficient because it merely takes extra efforts on generating the noise. In view of these advantages, our mechanism for privacy-preserving PCA is also based on input perturbation.

Related Work

Blum et al. (2005) proposed an early input perturbation framework (named SULQ), and the parameters of noise are refined by Dwork et al. (2006). Dwork et al. (2014) proved the state-of-the-art utility bounds for (ϵ, δ) -DP. Hardt and Roth (2012) provided a better bound under the coherence assumption. In (Hardt and Roth 2013; Hardt and Price 2014), the authors used a noisy power method to produce the principal eigenvector iteratively with removing the previous generated ones. Hardt and Price (2014) provided a special case for $(\epsilon, 0)$ -DP as well.

Chaudhuri, Sarwate, and Sinha (2012) proposed the first useful privacy-preserving PCA algorithm for $(\epsilon, 0)$ -DP based on an exponential mechanism (McSherry and Talwar 2007). Kapralov and Talwar (2013) argued that the algorithm in (Chaudhuri, Sarwate, and Sinha 2012) lacks convergence time guarantee and used heuristic tests to check convergence of the chain, which may affect the privacy guarantee. They also devised a mixed algorithm for low rank matrix approximation. However, their algorithm is quite complicated to implement and takes $O(d^6/\epsilon)$ running time. Here d is the dimension of the data point.

Our work is mainly inspired by Dwork et al. (2014). Since they provided the algorithms for (ϵ, δ) -DP, we seek the similar approach for $(\epsilon, 0)$ -DP with a different noise matrix design. As input perturbation methods, Blum et al. (2005) and Dwork et al. (2014) both used the Gaussian symmetric noise matrix for privately publishing a noisy covariance matrix. A reasonable worry is that the published matrix might be no longer positive semidefinite, a normal attribute for a covariance matrix.

Contribution and Organization

In this paper we propose a new mechanism for privacy-preserving PCA that we call *Wishart mechanism*. The key idea is to add a Wishart noise matrix to the original sample covariance matrix. A Wishart matrix is always positive semidefinite, which in turn makes the perturbed covariance matrix positive semidefinite. Additionally, Wishart matrix can be regarded as the scatter matrix of some random Gaussian vectors (Gupta and Nagar 2000). Consequently, our Wishart mechanism equivalently adds Gaussian noise to the original data points.

Setting appropriate parameters of Wishart distribution, we derive the $(\epsilon, 0)$ -privacy guarantee (Theorem 4). Compared to the present Laplace mechanism, our Wishart mechanism adds less noise (Section 4), which implies our mechanism always has better utility bound. We also provide a general framework for choosing Laplace or Wishart input perturbation for $(\epsilon, 0)$ -DP in Section 4.

Not only using the Laplace mechanism as a baseline, we also conduct theoretical analysis to compare our work with other privacy-preserving PCA algorithms based on the $(\epsilon, 0)$ -DP. With respect to the different criteria, we provide sample complexity bound (Theorem 7) for comparison with Chaudhuri, Sarwate, and Sinha (2012) and derive the low rank approximation closeness when comparing to Kapralov and Talwar (2013). Other than the principal eigenvector guarantee in (Chaudhuri, Sarwate, and Sinha 2012), we have the guarantee for rank- k subspace closeness (Theorem 6). With using a stronger definition of adjacent matrices, we achieve a k -free utility bound (Theorem 9). Converting the lower bound construction in (Chaudhuri, Sarwate, and Sinha 2012; Kapralov and Talwar 2013) into our case, we can see the Wishart mechanism is near-optimal.

The remainder of the paper is organized as follows. Section 2 gives the notation and definitions used in our paper. Section 3 lists the baseline and our designed algorithms. Section 4 provides the thorough analysis on privacy and utility guarantee of our mechanism together with comparison to several highly-related work. Finally, we conclude the work in Section 5. Note that we put some proofs and more explanation into the full version.

2 Preliminaries

We first give some notation that will be used in this paper. Let I_m denote the $m \times m$ identity matrix. Given an $m \times n$ real matrix $Z = [Z_{ij}]$, let its full singular value decomposition (SVD) as $Z = U\Sigma V^T$, where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal (i.e., $U^T U = I_m$ and $V^T V = I_n$), and $\Sigma \in \mathbb{R}^{m \times n}$ is a diagonal matrix with the i th diagonal entry σ_i being the i th largest singular value of Z . Assume that the rank of Z is $\rho \leq \min(m, n)$. This implies that Z has ρ nonzero singular values. Let U_k and V_k be the first k ($< \rho$) columns of U and V , respectively, and Σ_k be the $k \times k$ top sub-block of Σ . Then the $m \times n$ matrix $Z_k = U_k \Sigma_k V_k^T$ is the best rank- k approximation to Z .

The Frobenius norm of Z is defined as $\|Z\|_F = \sqrt{\sum_{i,j} Z_{ij}^2} = \sqrt{\sum_{i=1}^{\rho} \sigma_i^2}$, the spectral norm is defined as $\|Z\|_2 = \max_{x \neq 0} \frac{\|Zx\|_2}{\|x\|_2} = \sigma_1$, the nuclear norm is defined as $\|Z\|_* = \sum_{i=1}^{\rho} \sigma_i$, and the $\ell_{1,1}$ norm is defined as $\|Z\|_{1,1} = \sum_{i,j} |Z_{ij}|$.

Given a set of n raw data points $X = [x_1, \dots, x_n]$ where $x_i \in \mathbb{R}^d$, we consider the problem of publishing a noisy empirical sample covariance matrix for doing PCA. Following previous work on privacy-preserving PCA, we also assume

$\|x_i\|_2 \leq 1$. The standard PCA computes the sample covariance matrix of the raw data $A = \frac{1}{n} X X^T = \frac{1}{n} \sum_{i=1}^n x_i x_i^T$. Since A is a $d \times d$ symmetric positive semidefinite matrix, its SVD is equivalent to the spectral decomposition. That is, $A = V \Sigma V^T$. PCA uses V_k as projection matrix to compute the low-dimensional representation of raw data: $Y \triangleq V_k^T X$.

In this work we use Laplace and Wishart distributions, which are defined as follows.

Definition 1. A random variable z is said to have a Laplace distribution $z \sim \text{Lap}(\mu, b)$, if its probability density function is

$$p(z) = \frac{1}{2b} \exp\left(-\frac{|z - \mu|}{b}\right).$$

Definition 2 ((Gupta and Nagar 2000)). A $d \times d$ random symmetric positive definite matrix W is said to have a Wishart distribution $W \sim W_d(m, C)$, if its probability density function is

$$p(W) = \frac{|W|^{\frac{m-d-1}{2}}}{2^{\frac{md}{2}} |C|^{\frac{m}{2}} \Gamma_d(\frac{m}{2})} \exp\left(-\frac{1}{2} \text{tr}(C^{-1}W)\right),$$

where $m > d - 1$ and C is a $d \times d$ positive definite matrix.

Now we introduce the formal definition of differential privacy.

Definition 3. A randomized mechanism M takes a dataset D as input and outputs a structure $s \in R$, where R is the range of M . For any two adjacent datasets D and \hat{D} (with only one distinct entry), M is said to be $(\epsilon, 0)$ -differential private if for all $S \subseteq R$ we have

$$\Pr\{M(D) \in S\} \leq e^\epsilon \Pr\{M(\hat{D}) \in S\},$$

where $\epsilon > 0$ is a small parameter controlling the strength of privacy requirement.

This definition actually sets limitation on the similarity of output probability distributions for the given similar inputs. Here the adjacent datasets can have several different interpretations. In the scenario of privacy-preserving PCA, our definition is as follows. Two datasets X and \hat{X} are adjacent provided $X = [x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n]$ and $\hat{X} = [x_1, \dots, x_{i-1}, \hat{x}_i, x_{i+1}, \dots, x_n]$ for $x_i \neq \hat{x}_i$. It should be pointed out that our definition of adjacent datasets is slightly different from (Kapralov and Talwar 2013), which leads to significant difference on utility bounds. We will give more specifically discussions in Section 4.

We also give the definition of (ϵ, δ) -differential privacy. This notion requires less privacy protection so that it often brings better utility guarantee.

Definition 4. A randomized mechanism M takes a dataset as input and outputs a structure $s \in R$, where R is the range of M . For any two adjacent datasets D and \hat{D} (with only one distinct entry), M is said to be (ϵ, δ) -differential private if for all $S \subseteq R$ we have

$$\Pr[M(D) \in S] \leq e^\epsilon \Pr[M(\hat{D}) \in S] + \delta.$$

Sensitivity analysis is a general approach to achieving differential privacy. The following definitions show the two typical kinds of sensitivity.

Definition 5. The ℓ_1 sensitivity is defined as

$$s_1(M) = \max_{d(D, \hat{D})=1} \|M(D) - M(\hat{D})\|_1.$$

The ℓ_2 sensitivity is defined as

$$s_2(M) = \max_{d(D, \hat{D})=1} \|M(D) - M(\hat{D})\|_2.$$

The sensitivity describes the possible largest change as a result of individual data entry replacement. The ℓ_1 sensitivity is used in Laplace Mechanism for $(\epsilon, 0)$ -differential privacy, while the ℓ_2 sensitivity is used in Gaussian Mechanism for (ϵ, δ) -differential privacy. We list the two mechanisms for comparison.

Theorem 1 (Laplace Mechanism). Let $\lambda > s_1(M)/\epsilon$. Add Laplace noise $\text{Lap}(0, \lambda)$ to each dimension of $M(D)$. This mechanism provides $(\epsilon, 0)$ -differential privacy.

Theorem 2 (Gaussian Mechanism). For $c^2 > 2 \ln(1.25/\delta)$, let $\sigma > cs_2(M)/\epsilon$. Add Gaussian noise $N(0, \sigma^2)$ to each dimension of $M(D)$. This mechanism provides (ϵ, δ) -differential privacy.

The above mechanisms are all perturbation methods. Another widely used method is exponential mechanism (McSherry and Talwar 2007) which is based on sampling techniques.

3 Algorithms

First we look at a general framework of privacy-preserving PCA. According to the definition of differential privacy, a privacy-preserving PCA takes the raw data matrix X as input and then calculates the sample covariance matrix $A = \frac{1}{n} X X^T$. Finally, it computes the top- k subspace of A as the output.

The traditional approach adds noise to the computing procedure. For example, Chaudhuri, Sarwate, and Sinha (2012) and Kapralov and Talwar (2013) used a sampling based mechanism during computing eigenvectors to obtain approximate results. Our mechanism adds noise in the first stage, publishing A in a differential private manner. Thus, our mechanism takes X as input and outputs A . Afterwards we follow the standard PCA to compute the top- k subspace. This can be seen as a differential private preprocessing procedure.

Our baseline is the Laplace mechanism (Algorithm 1 and Theorem 1). To the best of our knowledge, Laplace mechanism is the only input perturbation method for $(\epsilon, 0)$ -DP PCA. Since this private procedure ends before computing the subspace, this shows $M(D) = \frac{1}{n} D D^T$ in sensitivity definition.

Note that to make \hat{A} be symmetric, we use a symmetric matrix-variate Laplace distribution in Algorithm 1. However, this mechanism cannot guarantee the positive semidefiniteness of \hat{A} , a desirable attribute for a covariance matrix. This motivates us to use a Wishart noise alternatively, giving rise to the Wishart mechanism in Algorithm 2.

Algorithm 1 Laplace input perturbation

Input: Raw data matrix $X \in \mathbb{R}^{d \times n}$; Privacy parameter ϵ ; Number of data n ;

1: Draw $\frac{d^2+d}{2}$ i.i.d. samples from $Lap(0, \frac{2d}{n\epsilon})$, then form a symmetric matrix L . These samples are put in the upper triangle part. Each entry in lower triangle part is copied from the opposite position.

2: Compute $A = \frac{1}{n}XX^T$;

3: Add noise $\hat{A} = A + L$;

Output: \hat{A} ;

Algorithm 2 Wishart input perturbation

Input: Raw data matrix $X \in \mathbb{R}^{d \times n}$; Privacy parameter ϵ ; Number of data n ;

1: Draw a sample W from $W_d(d+1, C)$, where C has d same eigenvalues equal to $\frac{3}{2n\epsilon}$;

2: Compute $A = \frac{1}{n}XX^T$;

3: Add noise $\hat{A} = A + W$;

Output: \hat{A} ;

4 Analysis

In this section, we are going to conduct theoretical analysis of Algorithms 1 and 2 under the framework of differential private matrix publishing. The theoretical support has two parts: privacy and utility guarantee. The former is the essential requirement for privacy-preserving algorithms and the latter tells how well the algorithm works against a non-private version. Chiefly, we list the valuable theorems and analysis. Due to space constraints, all the technical proofs omitted are presented in the full version.

Privacy guarantee

We first show that both algorithms satisfy privacy guarantee. Suppose there are two adjacent datasets $X = [x_1, \dots, v, \dots, x_n] \in \mathbb{R}^{d \times n}$ and $\hat{X} = [x_1, \dots, \hat{v}, \dots, x_n] \in \mathbb{R}^{d \times n}$ where $v \neq \hat{v}$ (i.e., only v and \hat{v} are distinct). Without loss of generality, we further assume that each data vector has the ℓ_2 norm at most 1.

Theorem 3. *Algorithm 1 provides $(\epsilon, 0)$ -differential privacy.*

This theorem can be quickly proved by some simple derivations so we omit the proof.

Theorem 4. *Algorithm 2 provides $(\epsilon, 0)$ -differential privacy.*

Proof. Assume the outputs for the adjacent inputs X and \hat{X} are identical (denoted $A + W_0$). Here $A = \frac{1}{n}XX^T$ and $\hat{A} = \frac{1}{n}\hat{X}\hat{X}^T$. We define the difference matrix $\Delta \triangleq A - \hat{A} = \frac{1}{n}(vv^T - \hat{v}\hat{v}^T)$. Actually the privacy guarantee is to bound

the following term:

$$\begin{aligned} \frac{p(A + W = A + W_0)}{p(\hat{A} + W = A + W_0)} &= \frac{p(W = W_0)}{p(W = A + W_0 - \hat{A})} \\ &= \frac{p(W = W_0)}{p(W = W_0 + \Delta)} \end{aligned}$$

As $W \sim W_d(d+1, C)$, we have that

$$\begin{aligned} \frac{p(W = W_0)}{p(W = W_0 + \Delta)} &= \frac{\exp[-\frac{1}{2}\text{tr}(C^{-1}W_0)]}{\exp[-\frac{1}{2}\text{tr}(C^{-1}(W_0 + \Delta))]} \\ &= \exp[\frac{1}{2}\text{tr}(C^{-1}(W_0 + \Delta)) - \text{tr}(C^{-1}W_0)] \\ &= \exp[\frac{1}{2}\text{tr}(C^{-1}\Delta)]. \end{aligned}$$

Then apply Von Neumann's trace inequality: For matrices $A, B \in \mathbb{R}^{d \times d}$, denote their i th-largest singular value as $\sigma_i(\cdot)$. Then $|\text{tr}(AB)| \leq \sum_{i=1}^d \sigma_i(A)\sigma_i(B)$. So that

$$\begin{aligned} \exp[\frac{1}{2}\text{tr}(C^{-1}\Delta)] &\leq \exp[\frac{1}{2}\sum_{i=1}^d \sigma_i(C^{-1})\sigma_i(\Delta)] \\ &\leq \exp[\frac{1}{2}\|C^{-1}\|_2\|\Delta\|_*]. \end{aligned} \tag{1}$$

Since $\Delta = A - \hat{A} = \frac{1}{n}(vv^T - \hat{v}\hat{v}^T)$ has rank at most 2, and by singular value inequality $\sigma_{i+j-1}(A+B) \leq \sigma_i(A) + \sigma_j(B)$, we can bound $\|\Delta\|_*$:

$$\begin{aligned} n\|\Delta\|_* &\leq \sigma_1(vv^T) + \sigma_1(-\hat{v}\hat{v}^T) + \max\{\sigma_1(vv^T) \\ &\quad + \sigma_2(-\hat{v}\hat{v}^T), \sigma_2(vv^T) + \sigma_1(-\hat{v}\hat{v}^T)\} \\ &= \sigma_1(vv^T) + \sigma_1(\hat{v}\hat{v}^T) + \max\{\sigma_1(vv^T), \sigma_1(\hat{v}\hat{v}^T)\} \\ &\leq 3\max\sigma_1(vv^T) = 3\max\|vv^T\|_2 \\ &= 3\max\|v\|_2^2 \leq 3. \end{aligned}$$

In Algorithm 2, the scale matrix C in Wishart distribution has d same eigenvalues equal to $\frac{3}{2n\epsilon}$, which implies $\|C^{-1}\|_2 = \frac{2n\epsilon}{3}$. Substituting these terms in Eq. (1) yields

$$\begin{aligned} \frac{p(A + W = A + W_0)}{p(\hat{A} + W = A + W_0)} &\leq \exp[\frac{1}{2}\|C^{-1}\|_2\|\Delta\|_*] \\ &\leq \exp[\frac{1}{2} \cdot \frac{2n\epsilon}{3} \cdot \frac{3}{n}] = e^\epsilon. \end{aligned}$$

□

Utility guarantee

Then we give bounds about how far the noisy results are from optimal. Since the Laplace and Wishart mechanisms are both input perturbation methods, their analyses are similar.

In order to ensure privacy guarantee, we add a noise matrix to the input data. Such noise may have effects on the property of the original matrix. For input perturbation methods, the *magnitude* of the noise matrix directly determines how large the effects are. For example, if the *magnitude* of the noise matrix is even larger than data, the matrix after

perturbation is surely covered by noise. Better utility bound means less noise added. We choose the spectral norm of the noise matrix to measure its *magnitude*. Since we are investigating the privacy-preserving PCA problem, the usefulness of the subspace of the top- k singular vectors is mainly cared.

The noise matrix in the Laplace mechanism is constructed with $\frac{d^2+d}{2}$ i.i.d random variables of $Lap(2d/n\epsilon)$. Using the tail bound for an ensemble matrix in (Tao 2012), we have that the spectral norm of the noise matrix in Algorithm 1 satisfies $\|L\|_2 = O(2d\sqrt{d}/n\epsilon)$ with high probability.

Then we turn to analyze the Wishart mechanism. We use the tail bound of the Wishart distribution in (Zhu 2012):

Lemma 1 (Tail Bound of Wishart Distribution). *Let $W \sim W_d(m, C)$. Then for $\theta \geq 0$, with probability at most $d \exp(-\theta)$,*

$$\lambda_1(W) \geq (m + \sqrt{2m\theta(r+2)} + 2\theta r)\lambda_1(C)$$

where $r = \text{tr}(C)/\|C\|_2$.

In our settings that $r = d$ and $m = d + 1$, we thus have that with probability at most $d \exp(-\theta)$,

$$\lambda_1(W) \geq (d + 1 + \sqrt{2(d+1)(d+2)\theta} + 2\theta d)\lambda_1(C).$$

Let $\theta = c \log d$ ($c > 1$). Then $d \exp(-\theta) = d^{1-c}$. So we can say with high probability

$$\lambda_1(W) = O([d + 1 + \sqrt{2(d+1)(d+2)\theta} + 2\theta d]\lambda_1(C)).$$

For convenience, we write

$$\lambda_1(W) = O(d \log d \lambda_1(C)) = O(3d \log d / 2n\epsilon).$$

We can see that the spectral norm of noise matrix generated by the Wishart mechanism is $O(d \log d / n\epsilon)$ while the Laplace mechanism requires $O(d\sqrt{d}/n\epsilon)$. This implies that the Wishart mechanism adds less noise to obtain privacy guarantee. We list the present four input perturbation approaches for comparison. Compared to the state-of-the-art results about (ϵ, δ) case (Dwork et al. 2014), our noise magnitude of $O(\frac{d \log d}{n\epsilon})$ is obviously worse than their $O(\frac{\sqrt{d}}{n\epsilon})$. It can be seen as the utility gap between (ϵ, δ) -DP and $(\epsilon, 0)$ -DP.

Table 1: Spectral norm of noise matrix in input perturbation.

Approach	Noise magnitude	Privacy
Laplace	$O(d\sqrt{d}/n\epsilon)$	$(\epsilon, 0)$
(Blum et al. 2005)	$O(d\sqrt{d \log d}/n\epsilon)$	(ϵ, δ)
Wishart	$O(d \log d / n\epsilon)$	$(\epsilon, 0)$
(Dwork et al. 2014)	$O(\sqrt{d}/n\epsilon)$	(ϵ, δ)

General framework We are talking about the intrinsic difference between the Laplace and Wishart mechanisms. The key element is the difference matrix Δ of two adjacent matrices. Laplace mechanism adds a noise matrix according to the ℓ_1 sensitivity, which equals to $\max \|\Delta\|_{1,1}$. Thus, the spectral norm of noise matrix is $O(\max \|\Delta\|_{1,1} \sqrt{d}/n\epsilon)$. When it comes to the Wishart

mechanism, the magnitude of noise is determined by $\|C\|_2$. For purpose of satisfying privacy guarantee, we take $\|C\|_2 = \omega(\max \|\Delta\|_*/n\epsilon)$. Then the spectral norm of noise matrix is $O(\max \|\Delta\|_* d \log d / n\epsilon)$. Consequently, we obtain the following theorem.

Theorem 5. *M is a $d \times d$ symmetric matrix generated by some input. For two arbitrary adjacent inputs, the generated matrices are M and \hat{M} . Let $\Delta = M - \hat{M}$. Using the Wishart mechanism to publish M in differential private manner works better if*

$$\frac{\max \|\Delta\|_{1,1}}{\max \|\Delta\|_*} = \omega(\sqrt{d \log d});$$

otherwise the Laplace mechanism works better.

Top- k subspace closeness We now conduct comparison between our mechanism and the algorithm in (Chaudhuri, Sarwate, and Sinha 2012). Chaudhuri, Sarwate, and Sinha (2012) proposed an exponential-mechanism-based method, which outputs the top- k subspace by drawing a sample from the matrix Bingham-von Mises-Fisher distribution. Wang, Wu, and Wu (2013) applied this algorithm to private spectral analysis on graph and showed that it outperforms the Laplace mechanism for output perturbation. Because of the scoring function defined, it is hard to directly sample from the original Bingham-von Mises-Fisher distribution. Instead, Chaudhuri, Sarwate, and Sinha (2012) used Gibbs sampling techniques to reach an approximate solution. However, there is no guarantee for convergence. They check the convergence heuristically, which may affect the basic privacy guarantee.

First we provide our result on the top- k subspace closeness:

Theorem 6. *Let \hat{V}_k be the top- k subspace of $A+W$ in Algorithm 2. Denote the non-noisy subspace as V_k corresponding to A . Assume $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_d$ are singular values of A . If $\sigma_k - \sigma_{k+1} \geq 2\|W\|_2$, then with high probability*

$$\|V_k V_k^T - \hat{V}_k \hat{V}_k^T\|_F \leq \frac{2\sqrt{k}\|W\|_2}{\sigma_k - \sigma_{k+1}}.$$

We apply the well-known Davis-Kahan $\sin \theta$ theorem (Davis 1963) to obtain this result. This theorem characterizes the usefulness of our noisy top- k subspace. Nevertheless, Chaudhuri, Sarwate, and Sinha (2012) only provided the utility guarantee on the principal eigenvector. So we can only compare the top-1 subspace closeness, correspondingly.

Before the comparison, we introduce the measure in (Chaudhuri, Sarwate, and Sinha 2012).

Definition 6. *A randomized algorithm $\mathcal{A}(\cdot)$ is an (ρ, η) -close approximation to the top eigenvector if for all data sets \mathcal{D} of n points, output a vector \hat{v}_1 such that*

$$\mathbb{P}(\langle \hat{v}_1, v_1 \rangle \geq \rho) \geq 1 - \eta.$$

Under this measure, we derive the sample complexity of the Wishart mechanism.

Theorem 7. If $n \geq \frac{3(d+1) + \sqrt{2(d+1)(d+2) \log \frac{d}{\eta} + 2d \log \frac{d}{\eta}}}{2\epsilon(1-\rho^2)(\lambda_1-\lambda_2)}$ and $\rho \geq \frac{\sqrt{2}}{2}$, then the Wishart mechanism is a (ρ, η) -close approximation to PCA.

Because a useful algorithm should output an eigenvector making ρ close to 1, our condition of $\rho \geq \frac{\sqrt{2}}{2}$ is quite weak. Comparing to the sample complexity bound of the algorithm in (Chaudhuri, Sarwate, and Sinha 2012):

Theorem 8. If $n \geq \frac{d}{\epsilon(1-\rho)(\lambda_1-\lambda_2)} \left(\frac{\log \frac{1}{\eta}}{d} + \log \frac{4\lambda_1}{(1-\rho^2)(\lambda_1-\lambda_2)} \right)$, then the algorithm in (Chaudhuri, Sarwate, and Sinha 2012) is a (ρ, η) -close approximation to PCA.

Our result has a factor up to $\log d$ with dropping the term $\log \frac{\lambda_1}{\lambda_1-\lambda_2}$. Actually, the relationship between d and $\frac{\lambda_1}{\lambda_1-\lambda_2}$ heavily depends on the data. Thus, as a special case of top- k subspace closeness, our bound for the top-1 subspace is comparable to Chaudhuri, Sarwate, and Sinha's (2012).

Low rank approximation Here we discuss the comparison between the Wishart mechanism and privacy-preserving rank- k approximation algorithm proposed in (Kapralov and Talwar 2013; Hardt and Price 2014). PCA can be seen as a special case of low rank approximation problems. Kapralov and Talwar (2013) combined the exponential and Laplace mechanisms to design a low rank approximation algorithm for a symmetric matrix, providing strict guarantee on convergence. However, the implementation of the algorithm contains too many approximation techniques and it takes $O(d^6/\epsilon)$ time complexity while our algorithm takes $O(kd^2)$ running time. Hardt and Price (2014) proposed an efficient meta algorithm, which can be applied to (ϵ, δ) -differentially private PCA. Additionally, they provided a $(\epsilon, 0)$ -differentially private version.

We need to point out that the definition of adjacent matrix in privacy-preserving low rank approximation is different from ours (our definition is the same as (Dwork et al. 2014; Chaudhuri, Sarwate, and Sinha 2012)). In the definition (Kapralov and Talwar 2013; Hardt and Price 2014), two matrices A and B are called adjacent if $\|A - B\|_2 \leq 1$, while we restrict the difference to a certain form $vv^T - \hat{v}\hat{v}^T$. In fact, we make a stronger assumption so that we are dealing with a case of less sensitivity. This difference impacts the lower bound provided in (Kapralov and Talwar 2013).

For the consistence of comparison, we remove the term $\frac{1}{n}$ in Algorithm 2, which means we use the XX^T for PCA instead of $\frac{1}{n}XX^T$. This is also used by Dwork et al. (2014).

Applying Lemma 1 in (Achlioptas and McSherry 2001), we can immediately have the following theorem:

Theorem 9. Suppose the original matrix is $A = XX^T$ and \hat{A}_k is the rank- k approximation of output by the Wishart mechanism. Denote the k -th largest eigenvalue of A as λ_k . Then

$$\|A - \hat{A}_k\|_2 \leq \lambda_{k+1} + O\left(\frac{d \log d}{\epsilon}\right).$$

Kapralov and Talwar (2013) provided a bound of $O\left(\frac{k^3 d}{\epsilon}\right)$ and Hardt and Price (2014) provided $O\left(\frac{k^{\frac{3}{2}} d \log^2 d}{\epsilon}\right)$ for the

same scenario. If k^3 is larger than $\log d$, our algorithm will work better. Moreover, our mechanism has better bounds than that of Hardt and Price (2014) while both algorithms are computationally efficient. Kapralov and Talwar (2013) established a lower bound of $O\left(\frac{kd}{\epsilon}\right)$ according to their definition of adjacent matrix. If replaced with our definition, the lower bound will become $O\left(\frac{d}{\epsilon}\right)$. The details will be given in the full version. So our mechanism is near-optimal.

5 Concluding Remarks

We have studied the problem of privately publishing a symmetric matrix and provided an approach for choosing Laplace or Wishart noise properly. In the scenario of PCA, our Wishart mechanism adds less noise than the Laplace, which leads to better utility guarantee. Compared with the privacy-preserving PCA algorithm in (Chaudhuri, Sarwate, and Sinha 2012), our mechanism has reliable rank- k utility guarantee while the former (Chaudhuri, Sarwate, and Sinha 2012) only has rank-1. For rank-1 approximation we have the comparable performance on sample complexity. Compared with the low rank approximation algorithm in (Kapralov and Talwar 2013), the bound of our mechanism does not depend on k . Moreover, our method is more tractable computationally. Compared with the tractable algorithm in (Hardt and Price 2014), our utility bound is better.

Since input perturbation only publishes the matrix for PCA, any other procedure can take the noisy matrix as input. Thus, our approach has more flexibility. While other entry-wise input perturbation techniques make the covariance not be positive semidefinite, in our case the noisy covariance matrix still preserves this property.

Acknowledgments

We thank Luo Luo for the meaningful technical discussion. We also thank Yujun Li, Tianfan Fu for support on the early stage of the work. This work is supported by the National Natural Science Foundation of China (No. 61572017) and the Natural Science Foundation of Shanghai City (No. 15ZR1424200).

References

- Achlioptas, D., and McSherry, F. 2001. Fast computation of low rank matrix approximations. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, 611–618. ACM.
- Blum, A.; Dwork, C.; McSherry, F.; and Nissim, K. 2005. Practical privacy: the sulq framework. In *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 128–138. ACM.
- Bojarski, M.; Choromanska, A.; Choromanski, K.; and LeCun, Y. 2014. Differentially-and non-differentially-private random decision trees. *arXiv preprint arXiv:1410.6973*.
- Chaudhuri, K., and Monteleoni, C. 2009. Privacy-preserving logistic regression. In *Advances in Neural Information Processing Systems*, 289–296.

- Chaudhuri, K.; Monteleoni, C.; and Sarwate, A. D. 2011. Differentially private empirical risk minimization. *The Journal of Machine Learning Research* 12:1069–1109.
- Chaudhuri, K.; Sarwate, A.; and Sinha, K. 2012. Near-optimal differentially private principal components. In *Advances in Neural Information Processing Systems*, 989–997.
- Davis, C. 1963. The rotation of eigenvectors by a perturbation. *Journal of Mathematical Analysis and Applications* 6(2):159–173.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography*. Springer. 265–284.
- Dwork, C.; Talwar, K.; Thakurta, A.; and Zhang, L. 2014. Analyze gauss: optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, 11–20. ACM.
- Gupta, A. K., and Nagar, D. K. 2000. *Matrix Variate Distributions*. Chapman & Hall/CRC.
- Hardt, M., and Price, E. 2014. The noisy power method: A meta algorithm with applications. In *Advances in Neural Information Processing Systems*, 2861–2869.
- Hardt, M., and Roth, A. 2012. Beating randomized response on incoherent matrices. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 1255–1268. ACM.
- Hardt, M., and Roth, A. 2013. Beyond worst-case analysis in private singular vector computation. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 331–340. ACM.
- Kapralov, M., and Talwar, K. 2013. On differentially private low rank approximation. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1395–1414. SIAM.
- McSherry, F., and Talwar, K. 2007. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, 94–103. IEEE.
- Tao, T. 2012. *Topics in random matrix theory*, volume 132. American Mathematical Soc.
- Wang, Y.; Wu, X.; and Wu, L. 2013. Differential privacy preserving spectral graph analysis. In *Advances in Knowledge Discovery and Data Mining*. Springer. 329–340.
- Zhu, S. 2012. A short note on the tail bound of wishart distribution. *arXiv preprint arXiv:1212.5860*.