Submitted to the Annals of Statistics

OPTIMAL DIFFERENTIALLY PRIVATE PCA AND ESTIMATION FOR SPIKED COVARIANCE MATRICES

BY T. TONY CAI[1,a], DONG XIA[2,b] AND MENGYUE ZHA[2,c]

[1]*Department of Statistics and Data Science, University of Pennsylvania ,* [a]*tcai@wharton.upenn.edu*

[2]*Department of Mathematics, Hong Kong University of Science and Technology ,* [b]*madxia@ust.hk;* [c]*mzha@connect.ust.hk*

Estimating a covariance matrix and its associated principal components is a fundamental problem in contemporary statistics. While optimal estimation procedures have been developed with well-understood properties, the increasing demand for privacy preservation introduces new complexities to this classical problem. In this paper, we study optimal differentially private Principal Component Analysis (PCA) and covariance estimation within the spiked covariance model.

We precisely characterize the sensitivity of eigenvalues and eigenvectors under this model and establish the minimax rates of convergence for estimating both the principal components and covariance matrix. These rates hold up to logarithmic factors and encompass general Schatten norms, including spectral norm, Frobenius norm, and nuclear norm as special cases.

We propose computationally efficient differentially private estimators and prove their minimax optimality for sub-Gaussian distributions, up to logarithmic factors. Additionally, matching minimax lower bounds are established. Notably, compared to the existing literature, our results accommodate a diverging rank, a broader range of signal strengths, and remain valid even when the sample size is much smaller than the dimension, provided the signal strength is sufficiently strong.

1. Introduction.

The covariance structure plays a fundamental role in multivariate analysis, and Principal Component Analysis (PCA) is a widely recognized technique known for its efficacy in dimension reduction and feature extraction [4]. PCA is particularly adept in settings where the data is high-dimensional but the underlying signal displays a low-dimensional structure. The estimation of covariance matrices and principal components finds applications across a diverse spectrum, encompassing tasks such as image recognition, data compression, clustering, risk management, portfolio allocation, mean tests, independence tests, and correlation analysis. Methodologies and theoretical advancements, including minimax optimality, for covariance matrix estimation and PCA, have been well-established in both low-dimensional and high-dimensional settings. See, for example, [7, 12, 13, 18, 31, 37, 52, 55, 58, 65]. For a survey on optimal estimation of high-dimensional covariance structures, see [14].

Amidst the increasing availability of large datasets containing sensitive personal information, privacy concerns in statistical data analysis have gained heightened prominence. The utilization of personal information in statistical analyses raises apprehensions about the potential compromise of individual privacy. Consequently, there is a growing emphasis on developing methodologies and techniques that offer robust privacy guarantees while still facilitating accurate statistical insights. This motivates a comprehensive exploration of the optimal tradeoff between privacy and accuracy in fundamental statistical problems, including PCA and covariance matrix estimation.

MSC2020 subject classifications: Primary 62H25; secondary 62C20.

Keywords and phrases: Cost of privacy, Covariance matrix estimation, differential privacy, minimax optimality, PCA, spiked covariance model.

Differential privacy (DP), a concept introduced by [26], provides a framework for safeguarding individual privacy in statistical analysis. DP has become a commonly accepted standard in both industrial and governmental applications [1, 2, 22, 29, 56]. The goal of the present paper is to develop methods and optimality results for PCA and covariance matrix estimation within the framework of the spiked covariance model under DP constraints.

1.1. *Problem formulation.* We begin by formally introducing the spiked covariance model and general formulation of the privacy constrained estimation problems.

The spiked covariance structure [31, 32] naturally arises from factor models with homoscedastic noise and has found diverse applications in signal processing, chemometrics, econometrics, population genetics, and various other fields. See, for example, [30, 40, 49, 51]. The spiked covariance model assumes that the population covariance matrix can be decomposed as

$$\Sigma = U\Lambda U^\top + \sigma^2 I_p, \tag{1}$$

where $U \in \mathbb{O}_{p,r}$ and $\Lambda = \operatorname{diag}(\lambda_1, \cdots, \lambda_r)$ represent the leading eigenvectors and eigenvalues (excluding σ^2), respectively. Here, $\mathbb{O}_{p,r}$ denotes the set of $p \times r$ matrices satisfying $U^\top U = I_r$. The spiked covariance model is convenient for studying the distribution of sample eigenvalues and eigenvectors, which play a critical role in the statistical inference of Σ and its eigenvectors. For instance, [24] studied the optimal shrinkage of sample eigenvalues in the spiked covariance model. In particular, [13] and [65] established the minimax optimal rates

$$\begin{aligned}
&\inf_{\widehat{U}} \sup_{\Sigma \in \Theta(\lambda, \sigma^2)} \mathbb{E}\|\widehat{U}\widehat{U}^\top - UU^\top\| \asymp \left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\sqrt{\frac{p}{n}};\\
&\inf_{\widehat{\Sigma}} \sup_{\Sigma \in \Theta(\lambda, \sigma^2)} \mathbb{E}\|\widehat{\Sigma} - \Sigma\| \asymp \lambda\sqrt{\frac{r}{n}} + \sqrt{\sigma^2(\lambda + \sigma^2)}\sqrt{\frac{p}{n}},
\end{aligned} \tag{2}$$

where the infimum is taken over all possible estimators based on the data $X = (X_1, \cdots, X_n)$ consisting of n observations independently sampled from the spiked covariance model (1), the parameter set $\Theta(\lambda, \sigma^2)$ is a collection of covariance matrices in the form (1) with all spiked eigenvalues have magnitudes of order λ (see formal definition in (4)), and $\|\cdot\|$ denotes the matrix spectral norm.

The concept of differential privacy was first introduced in [26]. For a given dataset X and any $\varepsilon > 0$ and $\delta \in [0, 1)$, a randomized algorithm A that maps X into $\mathbb{R}^{d_1 \times d_2}$ is called (ε, δ)-differentially private ((ε, δ)-DP) over the dataset X if

$$\mathbb{P}\big(A(X) \in \mathcal{Q}\big) \le e^{\varepsilon}\mathbb{P}\big(A(X') \in \mathcal{Q}\big) + \delta,$$

for all measurable subset $\mathcal{Q} \subset \mathbb{R}^{d_1 \times d_2}$ and all neighboring data set X'. In the standard definition, a dataset X' is a neighbor of X if they differ by only one datum, i.e., one observation in X is replaced by some other, possibly arbitrary, datum. In the context of PCA and covariance matrix estimation, as observations in X are independently sampled from a common distribution, a neighboring dataset X' is obtained by replacing one datum in X with an independent copy. This facilitates exploration of the statistical properties of the sample data.

Under the (ε, δ)-DP constraint, our goal is to investigate the cost of privacy in PCA and covariance matrix estimation. This includes designing minimax optimal (ε, δ)-DP estimators of the principal components and covariance matrix and establishing the privacy-constrained minimax lower bounds.

1.2. *Main contribution.* In this paper, we establish the minimax optimal rates for PCA and covariance matrix estimation in the spiked model under DP constraints. Over the collection of sub-Gaussian distributions, these rates, up to logarithmic factors, are given by:

$$\begin{aligned}
&\inf_{\widehat{U}\in\mathcal{U}_{\varepsilon,\delta}}\ \sup_{\Sigma\in\Theta(\lambda,\sigma^2)} \frac{\mathbb{E}\|\widehat{U}\widehat{U}^\top - UU^\top\|_q}{r^{1/q}} \asymp \left(\frac{\sigma^2}{\lambda}+\sqrt{\frac{\sigma^2}{\lambda}}\right)\left(\sqrt{\frac{p}{n}}+\frac{p\sqrt{r}}{n\epsilon}\right)\bigwedge 1;\\
&\inf_{\widehat{\Sigma}\in\mathcal{M}_{\varepsilon,\delta}}\ \sup_{\Sigma\in\Theta(\lambda,\sigma^2)} \frac{\mathbb{E}\|\widehat{\Sigma}-\Sigma\|_q}{r^{1/q}} \asymp \left(\lambda\left(\sqrt{\frac{r}{n}}+\frac{r^{3/2}}{n\varepsilon}\right)+\sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p}{n}}+\frac{\sqrt{r}p}{n\varepsilon}\right)\right)\bigwedge \lambda,
\end{aligned}\tag{3}$$

where the infimum is taken over all possible (ε,δ)-DP algorithms denoted by $\mathcal{U}_{\varepsilon,\delta}$ for principal components and $\mathcal{M}_{\varepsilon,\delta}$ for the covariance matrix. The expectation is taken with respect to the randomness of both the data and the differentially private algorithm. These rates hold in Schatten-q norms for all $q\in[1,\infty]$, including spectral norm ($q=\infty$), Frobenius norm ($q=2$), and nuclear norm ($q=1$) as special cases. The rank r can grow with respect to p as long as $r\le p/2$, and the sample size can be much smaller than p as long as the signal-to-noise ratio (SNR) satisfies $\lambda/\sigma^2\ge C_1(\sqrt{p/n}+p/n)$. This condition is minimal since no consistent estimation is possible when this condition does not hold. To our knowledge, this represents the first comprehensive presentation of minimax optimal rates for PCA and covariance matrix estimation under DP constraints. For technical convenience and theoretical clarity, we focus on sub-Gaussian distributions in this paper. However, we believe that our results can be extended beyond sub-Gaussian distributions. For further details, see the discussion in Section 5.

Our contributions are multifod. Methodologically, we introduce (ε,δ)-DP estimators for PCA and covariance matrices that are computationally efficient. Specifically, we employ the Gaussian mechanism for the sample spectral projector in differentially private PCA. Notably, our DP estimator for the covariance matrix incorporates a novel design to handle unknown orthogonal rotations. These estimators are shown to achieve minimax optimality, up to logarithmic factors. Theoretically, we provide a comprehensive understanding of the minimax optimal rates for PCA and covariance estimation under privacy constraints, valid across all Schatten norms. The derivation of minimax lower bounds employs Fano's lemma with a differential privacy constraint and the construction of well-separated spectral projectors based on the packing complexity of Grassmannians [38, 64].

Differentially private PCA and covariance estimation are challenging because it is difficult to characterize a sharp sensitivity bound for the eigenvectors. Our main technical contribution lies in a precise characterization of the sensitivity of the sample spectral projector $\widehat{U}\widehat{U}^\top$, quantifying its deviation when one datum X_i is replaced by an independent copy X_i'. A key technical tool is an explicit spectral representation formula for $\widehat{U}\widehat{U}^\top$ adapted from [61]. We derive a similar formula specifically for the spiked covariance model, which is of independent interest. Based on this sharp sensitivity analysis, we apply the Gaussian mechanism to achieve the upper bounds in (3), up to logarithmic terms.

1.3. *Related work.* Minimax optimal rates under (ϵ,δ)-DP guarantees have been established for several statistical problems, such as mean estimation, linear regression, pairwise comparisons, matrix completion, factorization, generalized linear models (GLMs), and sparse GLMs [11, 15, 16, 20, 60]. Additionally, optimality results have also been developed under local privacy constraints. For example, [25] established minimax rates for mean estimation, GLMs, and nonparametric density estimation, while [53] developed minimax theory for estimating linear functionals under local privacy. It is worth noting that local privacy is a stronger notion of privacy compared to (ε, δ)-DP, and it may not be compatible with high-dimensional problems [25]. A refined fingerprint lower bound method was introduced by [46], allowing for a broader range of δ and establishing a minimax lower bound for covariance matrix estimation (see also [23] and [43]). Both studies focused on general covariance matrix estimation, but their results become suboptimal in the case of spiked covariance matrices. The Johnson-Lindenstrauss mechanism was examined by [47], providing optimal sample complexity for differentially private covariance estimation of a bounded high-dimensional distribution. While these privacy-preserving methods are centered on covariance estimation, their applicability and performance for PCA remain largely unclear. Additionally, although enforcing boundedness can guarantee worst-case privacy protection, it may result in a pessimistic estimator in certain settings. See Remarks 1 and 3 for a detailed comparison with existing literature.

Differentially private PCA algorithms were proposed in [9, 19, 28] based on the perturbation mechanism, treating each datum X_i as a fixed vector and investigating the sensitivity of sample eigenvectors. However, their deterministic sensitivity analysis disregards the statistical properties of sample data, resulting in suboptimal error rates when X_i's are i.i.d. sampled from a common distribution, such as the spiked covariance model. Differentially private methods that explore statistical properties have been studied in [10, 33] and related works. However, optimal differentially private PCA has received much less attention, and existing results for private covariance estimation are generally suboptimal under the spiked covariance model. Recently, [42] introduced an online PCA algorithm with DP, providing a much sharper upper bound for differentially private PCA under the spiked covariance model. The online Oja's algorithm in [42] consumes one datum at a time, allowing for an explicit representation formula in the updated estimate of eigenvectors and enabling a study of their sensitivity. However, their bound is valid only for the rank-one case ($r = 1$) and is minimax optimal only when $\lambda \leq \sigma^2$. The optimality of their algorithm for general rank r or $\lambda \gg \sigma^2$ remains unclear. Moreover, the minimax optimal rates for estimating Σ under privacy constraints are still unknown under the spiked covariance model.

1.4. *Organization of the paper.* The rest of the paper is organized as follows. In Section 2, we introduce the Gaussian mechanism and study the sensitivity of the empirical spectral projector under the spiked covariance model. We present a DP algorithm for estimating the spectral projector and spiked covariance matrix in the same section. The upper bounds for our proposed DP algorithms are proven in Section 3, where an explicit spectral representation formula under the spiked covariance model is also developed. Section 4 establishes a differentially private Fano's lemma and minimax lower bounds. Extensions to the settings with diverging conditioning number and sub-Gaussian distributions are discussed in Section 5. Some of the key technical lemmas are presented in Section 8. All the proofs as well as additional simulation results are given in the Supplementary Materials [17].

2. Methodology: Gaussian Mechanism and Sensitivity.

Our differentially private PCA and covariance estimation method relies on a precise characterization of the sensitivity for both eigenvectors and eigenvalues under the spiked covariance model. For technical convenience, we first focus on Gaussian PCA and provide a broader extension to general sub-Gaussian PCA in Section 5.

For brevity, let $X := (X_1, \cdots, X_n)$ represent the $p \times n$ matrix collecting all i.i.d. observations X_i sampled from a centered normal distribution $\mathcal{N}(0, \Sigma)$. The sensitivity of eigenvectors and eigenvalues denotes their perturbation if an observation X_i is replaced by an independent copy X_i' expressed briefly as $X^{(i)} := (X_1, \cdots, X_{i-1}, X_i', X_{i+1}, \cdots, X_n)$. Here, X and $X^{(i)}$ form a pair of neighboring datasets [26]. Notably, the sensitivity is contingent on the covariance matrix Σ.

Through out this paper, we consider the spiked covariance matrix model where Σ is from the following parameter space

(4)

$$\Theta(p,r,\lambda,\sigma^2)=\Big\{\Sigma=U\Lambda U^\top+\sigma^2 I_p: \\ U\in\mathbb{O}_{p,r},\Lambda=\operatorname{diag}(\lambda_1,\cdots,\lambda_r),c_0\lambda\le\lambda_r\le\lambda_1\le C_0\lambda\Big\},$$

where I_p is the identity matrix and and $\mathbb{O}_{p,r}$ refers to the set of matrices with orthonormal columns, i.e., matrices satisfying $U^\top U=I_r$. Thus, our focus is on spiked covariance matrices with a bounded condition number, a common assumption in existing literature [12, 19, 42]. However, our methodology remains valid, and the theoretical framework can be extended to the case of an unbounded condition number, as discussed in Section 5. For simplicity, we use $\Theta(\lambda,\sigma^2)$ without explicitly stating the dimensions p and rank r. Let $\mathcal{P}$ denote the family of normal distributions $\mathcal{N}(0,\Sigma)$ with the population covariance matrix $\Sigma\in\Theta(\lambda,\sigma^2)$. Without loss of generality, we assume that σ^2 is known.

Formally, the sensitivity and Gaussian mechanism are described as follows without proofs. See, for example, [26, Proposition 1] and [27, Theorem A.1] for more details. Here, $\|\cdot\|_{\mathrm{F}}$ stands for the matrix Frobenius norm.

LEMMA 2.1 (sensitivity and Gaussian mechanism). *Let X be a given data set and X' be any neighboring data set of X, i.e., X and X' differs by at most one observation. The sensitivity of a function f that maps X into $\mathbb{R}^{d_1\times d_2}$ is defined by*

$$\omega_f:=\sup_{neighboring(X,X')}\|f(X)-f(X')\|_{\mathrm{F}}. \tag{5}$$

Then, for any $\varepsilon>0$ and $\delta\in[0,1)$, the randomized algorithm A defined by $A(X)=f(X)+Z$ where Z has i.i.d. $\mathcal{N}\big(0,2\omega_f^2\varepsilon^{-2}\log(1.25/\delta)\big)$ entries is (ε,δ)-DP over the dataset X.

The definition of sensitivity in Lemma 2.1 relies on the pair of neighboring data sets. Here, X is simply the data matrix where each column represents one observation. While X and X' differ only by one observation, the sensitivity can still be unbounded if no restriction is posed on the difference, e.g., by replacing one observation of X by infinite. Since X consists of i.i.d. columns under the spiked covariance model, we assume that a neighboring data set X' is obtained by replacing some column of X by its i.i.d. copy throughout this paper.

2.1. *Differentially private estimation by Gaussian mechanism.* Our DP-estimators of principal components and covariance matrix are built on Gaussian mechanism. Here, we assume that the rank r and nuisance variance σ^2 are known for simplicity. Let $\widehat{U}$ be the top-r eigenvectors of the sample covariance matrix $\widehat{\Sigma}:=n^{-1}\sum_{i=1}^n X_iX_i^\top$ and denote $\widehat{U}\widehat{U}^\top$ the sample spectral projector. By Lemma 2.1, differentially private PCA can be obtained by adding Gaussian noise Z to $\widehat{U}\widehat{U}^\top$ provided that the entrywise variance of Z dominates the sensitivity of $\widehat{U}\widehat{U}^\top$. While publishing $\widehat{U}\widehat{U}^\top+Z$ protects privacy, it is certainly not a preferable estimator of principal components as it generally lacks validity as a spectral projector. We therefore take the eigenvectors of $\widehat{U}\widehat{U}^\top+Z$ as the ultimate estimator. This choice maintains differential privacy, as the post-processing of a differentially private algorithm retains differential privacy according to well-established results, as discussed in [26].

Our proposed differentially private PCA and covariance estimation procedures are given in Algorithm 1. The proper choice of sensitivities Δ_1 and Δ_2 is determined by Lemma 2.3 and Lemma 2.4 in Section 2.2, respectively. However, $\widetilde{U}$ and $\widehat{U}$ are close up to an orthogonal rotation. As a result, our algorithm chooses to add Gaussian noise to $\widetilde{U}^\top\widehat{\Sigma}\widetilde{U}$ instead of the

Algorithm 1 Differentially private PCA and covariance estimation

```
**Input**: data matrix $X = (X_1, \cdots, X_n) \in \mathbb{R}^{n\times p}$; eigenvectors and eigenvalues sensitivity $\Delta_1$ and $\Delta_2 > 0$; rank $r$; nuisance variance $\sigma^2$; privacy budget $\varepsilon > 0, \delta \in (0,1)$.
**Output**: $(\varepsilon,\delta)$-DP estimate of $U$ and $\Sigma$.
Compute the sample covariance matrix and top-$r$ eigenvectors:
    $\widehat{\Sigma} \longleftarrow \frac{1}{n}\sum_{i=1}^{n} X_i X_i^\top$   and   $\widehat{U} \longleftarrow \mathrm{SVD}_r(\widehat{\Sigma})$;
Compute $(\varepsilon/2,\delta/2)$-DP PCA by adding artificial Gaussian noise:
    $\widetilde{U} \longleftarrow \mathrm{SVD}_r\Big(\widehat{U}\widehat{U}^\top + Z\Big)$   where   $Z_{ij} = Z_{ji} \overset{\text{i.i.d.}}{\sim} \mathcal{N}\Big(0, \frac{8\Delta_1^2}{\varepsilon^2}\log\frac{2.5}{\delta}\Big)$,   $\forall 1 \le i \le j \le p$;
Compute $(\varepsilon/2,\delta/2)$-DP estimates of eigenvalues up to rotations:
    $\widetilde{\Lambda} \quad \longleftarrow \widetilde{U}^\top(\widehat{\Sigma} - \sigma^2 I_p)\widetilde{U} + E$   where   $E_{ij} = E_{ji} \overset{i.i.d.}{\sim} \mathcal{N}\Big(0, \frac{8\Delta_2^2}{\varepsilon^2}\log\frac{2.5}{\delta}\Big)$,   $\forall 1 \le i \le j \le r$;
Compute $(\varepsilon,\delta)$-DP covariance estimate by :
    $\widetilde{\Sigma} \longleftarrow \widetilde{U}\widetilde{\Lambda}\widetilde{U}^\top + \sigma^2 I_p.$
**Return**: $\widetilde{U}$ and $\widetilde{\Sigma}$
```

empirical eigenvalues $\widehat{\Lambda} := (\widehat{\lambda}_1, \cdots, \widehat{\lambda}_r)^\top$. The added noise level depends on the sensitivity of $\widetilde{U}^\top\widehat{\Sigma}\widetilde{U}$, within which $\widetilde{U}$ is already differentially private. It thus suffices to study the upper bound of $\|\widetilde{U}^\top(\widehat{\Sigma} - \widehat{\Sigma}^{(i)})\widetilde{U}\|_{\mathrm{F}} \le \|\widehat{\Sigma} - \widehat{\Sigma}^{(i)}\|_{\mathrm{F}}$, which will be established in Lemma 2.4.

Our approach to differentially privately estimating the main covariance term involves separately privatizing the eigenvectors and eigenvalues. This separation is driven by the observation that the relative sensitivity of eigenvalues is significantly larger than that of eigenvectors. Note that a natural estimator of $U(\Lambda + \sigma^2 I_r)U^\top$ is $\widehat{U}\widehat{U}^\top\widehat{\Sigma}\widehat{U}\widehat{U}^\top$. It is possible to characterize the sensitivity of this estimator by directly studying the bound $\|\widehat{U}\widehat{U}^\top\widehat{\Sigma}\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)}\widehat{U}^{(i)\top}\widehat{\Sigma}^{(i)}\widehat{U}^{(i)}\widehat{U}^{(i)\top}\|_{\mathrm{F}}$. However, the sensitivity of eigenvalues will be the dominating factor and force us to add unnecessarily large noise to a $p \times p$ matrix. This delivers a statistically sub-optimal estimator of the spiked covariance matrix.

The estimated eigenvectors $\widetilde{U}$ is $(\varepsilon/2,\delta/2)$-DP and eigenvalues $\widetilde{\Lambda}$ is $(\varepsilon/2,\delta/2)$-DP with high probability. By the composition property of differentially private algorithm, the estimator $\widetilde{U}\widetilde{\Lambda}\widetilde{U}^\top$ is (ε,δ)-DP. The conclusion is formally stated in the following lemma. Recall that $\tilde{r} = (r\lambda + p\sigma^2)/(\lambda + \sigma^2)$ is the effective rank of Σ. Here, λ is regarded as the signal strength.

LEMMA 2.2. *Let the data matrix $X = (X_1, \cdots, X_n)$ consists of i.i.d. columns sampled from $\mathcal{N}(0,\Sigma)$ with $\Sigma \in \Theta(\lambda,\sigma^2)$, $\varepsilon > 0, \delta \in (0,1)$, and assume $n \ge C_1(r\log n + \log^2 n), 2r \le p$, and $\lambda/\sigma^2 \ge C_1(p/n + \sqrt{p/n})$ for some large absolute constant $C_1 > 0$. If we choose*

$$\Delta_1 := C_2\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\frac{\sqrt{p(r+\log n)}}{n} \quad \text{and} \quad \Delta_2 := C_3\frac{\lambda(r+\log n) + \sigma^2(p+\log n)}{n},$$

for some large absolute constants $C_2, C_3 > 0$, then Algorithm 1 is (ε,δ)-DP with probability at least $1 - 4n^{-99} - e^{-c_1(n\wedge p)}$ for some absolute constant $c_1 > 0$.

REMARK 1 (Worst-case and high-probability privacy guarantee). Compared to existing literature [19, 23, 34, 42, 47], our algorithm does not truncate the observations, allowing $\|X_i\|$ to remain unbounded. As a result, our algorithm is differentially private with high probability.

The DP-Oja algorithm proposed by [42] ensures worst-case privacy guarantees due to its online nature. However, it is limited to the rank-one case, performs poorly in both simulation and real data experiments (see Section 6), and the established error rate is much larger than ours under spiked covariance model when signal strength $\lambda \gg \sigma^2$ (see Remark 2). The DP-Gauss method [28, 43] ensures worst-case privacy by applying a global scaling, limiting each observation to at most unit norm. While we could apply global scaling to our method to ensure worst-case privacy, as discussed in Remark 3, this approach would result in an overly pessimistic estimator with a significantly larger error rate under the spiked covariance model. Therefore, we do not pursue worst-case privacy guarantees in this paper. Moreover, note that the probability terms n^{-99} in Lemma 2.2 can be replaced by n^{-C_5} with any absolute constant $C_5 > 0$ (by adjusting the constants C_2, C_3 in the definitions of Δ_1 and Δ_2 accordingly). The failure probability decreases polynomially fast with respect to sample size n.

The sensitivities Δ_1 and Δ_2 play a critical role in guaranteeing the differential privacy of Algorithm 1, which shall be developed in next section. The conditions $r\log n + \log^2 n = O(n)$ and $2r \le p$ are mild. The SNR condition $\lambda/\sigma^2 \ge C_1(p/n + \sqrt{p/n})$ is typical in the existing literature of spiked covariance matrix model. See, e.g., [45, 65] and references therein.

2.2. *Sensitivity analysis.* In this section, we analyze the sensitivities of sample eigenvectors and eigenvalues under the spiked covariance model. The data matrix $X = (X_1, \cdots, X_n) \sim \mathcal{N}(0, \Sigma)^{\otimes n}$ for some $\Sigma \in \Theta(\lambda, \sigma^2)$. Similarly, its neighboring data matrix $X^{(i)} = (X_1, \cdots, X_i', \cdots, X_n) \sim \mathcal{N}(0, \Sigma)^{\otimes n}$. Define the sample covariance matrices by

$$\widehat{\Sigma} := \frac{1}{n}\sum_{i=1}^{n} X_i X_i^\top \quad \text{and} \quad \widehat{\Sigma}^{(i)} := \frac{1}{n}\Big(X_i' X_i'^\top + \sum_{j\neq i} X_j X_j^\top\Big),$$

Denote $\widehat{U} \in \mathbb{O}_{p,r}$ and $\widehat{U}^{(i)} \in \mathbb{O}_{p,r}$ the top-r left eigenvectors of $\widehat{\Sigma}$ and $\widehat{\Sigma}^{(i)}$, respectively. The sensitivity of sample eigenvectors characterizes the deviation between $\widehat{U}$ and $\widehat{U}^{(i)}$ caused by replacing the i-th observation by its i.i.d. copy. Since eigenvectors are determined up to an orthogonal rotation (note that we allow the eigengap $|\lambda_i - \lambda_j|$ to be zero), a commonly used metric for measuring the distance between eigenvectors is the projection distance defined by $\|\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)}\widehat{U}^{(i)\top}\|_{\mathrm{F}}$.

The primary challenge in differentially private PCA lies in characterizing a precise upper bound for $\|\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)\top}\widehat{U}^{(i)}\|_{\mathrm{F}}$. In most existing works [9, 19, 28], the data matrix X is assumed to be fixed, and its columns are all bounded, denoted as $\|X_i\| \le \gamma$, where we slightly abuse the notation by letting $\|\cdot\|$ denote the ℓ_2-norm for vectors and γ is a deterministic value. This immediately implies an upper bound $\|\widehat{\Sigma} - \widehat{\Sigma}^{(i)}\| \le 2\gamma^2/n$ and the sensitivity of $\widehat{U}\widehat{U}^\top$ is guaranteed by the Davis-Kahan theorem.

However, this approach becomes invalid when observations are unbounded and suboptimal when observations are randomly sampled from a common distribution. A more recent work [42] aimed to exploit the statistical properties of i.i.d. samples to achieve a sharper bound for differentially private PCA. This work focused on the rank-one case ($r = 1$) and the Oja's algorithm, well-known for online PCA, which iteratively updates the estimation with one additional observation. The online fashion of Oja's algorithm in the rank-one case allows for an explicit representation of the eigenvector estimator, enabling a sharp upper bound of the sensitivity to be derived. Consequently, nearly optimal differentially private PCA for the case $r = 1$ was achieved. However, it remains unclear how this approach can be extended to the rank-r case and what the minimax optimal convergence rates are.

We take a fundamentally different approach by directly focusing on $\|\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)}\widehat{U}^{(i)\top}\|_{\mathrm{F}}$. This task presents two challenges: the spectral projector $\widehat{U}\widehat{U}^\top$ involves a complicated function of the data matrix X, and a sharp perturbation analysis is required for a set of r empirical eigenvectors. Fortunately, we leverage an explicit spectral representation formula adapted from [61] and successfully establish a precise upper bound for $\|\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)}\widehat{U}^{(i)\top}\|_{\mathrm{F}}$.

LEMMA 2.3. *Suppose the conditions in Lemma 2.2 hold and assume* $n \geq C_1(r\log n + \log^2 n)$ *and* $2r \leq p$. *There exist absolute constants* $c_1, C_2 > 0$ *such that with probability at least* $1 - 3n^{-99} - e^{-c_1(n\wedge p)}$,

$$\max_{i\in[n]} \|\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)}\widehat{U}^{(i)\top}\|_{\mathrm{F}} \leq C_2\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\frac{\sqrt{p(r+\log n)}}{n}. \tag{6}$$

The $\log n$ term in upper bound (6) is due to the maximization over n. Nevertheless, the bound is much smaller than that achieved by the deterministic analysis in [9, 19, 28]. Indeed, a direct application of Davis-Kahan theorem yields an upper bound $O\big(\|\widehat{\Sigma} - \widehat{\Sigma}^{(i)}\|\sqrt{r}/\lambda\big)$, which is at least in the order $O\big((r\lambda + p\sigma^2)\sqrt{r}/(n\lambda)\big)$, with high probability. The significant improvement is due to a sharp spectral characterization showing that the difference $\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)}\widehat{U}^{(i)\top}$ is mainly contributed by the term $\|U^\top(X_iX_i^\top - X_i'X_i'^\top)U_\perp\|_{\mathrm{F}}/(n\lambda)$. Here, $U_\perp \in \mathbb{O}_{p,p-r}$ denotes the orthogonal complement of U such that $(U, U_\perp)$ is an orthogonal matrix. The proof of Lemma 2.3 is technically involved and deferred to Section C.2. It is worth noting that the original spectral representation formula developed in [61] is inapplicable here because Σ is not exactly rank-r. Interestingly, we establish a similar spectral representation formula exclusively for spiked covariance matrix, which may be of independent interest. See Lemma 3.1 in Section 3.1.

The sensitivity of eigenvalues is also necessary for constructing differentially private covariance estimation. Let $\lambda_k(\widehat{\Sigma})$ and $\lambda_k(\widehat{\Sigma}^{(i)})$ denote the k-th largest eigenvalue of $\widehat{\Sigma}$ and $\widehat{\Sigma}^{(i)}$, respectively. Compared to the eigenvectors, the sensitivity of eigenvalues can be easily characterized by Hoffman-Weilandt's inequality. The proof of Lemma 2.4 is deferred to Section C.3.

LEMMA 2.4. *Suppose the conditions in Lemma 2.2 hold. There exists an absolute constant* $C_2 > 0$ *such that with probability at least* $1 - n^{-100}$,

$$\sum_{k=1}^{p}\left|\lambda_k(\widehat{\Sigma}) - \lambda_k(\widehat{\Sigma}^{(i)})\right|^2 \leq C_2\left(\frac{\lambda(r+\log n) + \sigma^2(p+\log n)}{n}\right)^2, \tag{7}$$

for all $i \in [n]$.

We can regard $\Big(\sum_{k=1}^p \big(\lambda_k(\widehat{\Sigma}) - \lambda_k(\widehat{\Sigma}^{(i)})\big)^2\Big)^{1/2}/\lambda$ and $\|\widehat{U}\widehat{U}^\top - \widehat{U}^{(i)}\widehat{U}^{(i)\top}\|_{\mathrm{F}}/\sqrt{r}$ as the relative sensitivity of eigenvectors and eigenvalues, respectively. Lemmas 2.3 and 2.4 show that the relative sensitivity of eigenvalues can be considerably larger than that of eigenvectors. This insight implies that, when designing a differentially private optimal estimation procedure for the population covariance matrix, it is advisable to privatize the eigenvalues and eigenvectors separately, as elaborated in Algorithm 1.

3. Upper Bounds with Differential Privacy.

3.1. *Spectral representation formula.* Our key technical tool is the following spectral representation formula. Recall that $\widehat{U}$ and U denote the top-r eigenvectors of $\widehat{\Sigma}$ and Σ, respectively. Denote the deviation matrix by $\widehat{\Delta} := \widehat{\Sigma} - \Sigma$ so that $\widehat{\Sigma} = \Sigma + \widehat{\Delta}$ is viewed as a perturbation of the "signal" matrix Σ. The spectral representation formula was first introduced in [61], which, however, requires the "signal" matrix to be exactly rank-r. This is certainly not the case here since Σ is full-rank. Here, we develop the spectral representation formula exclusively for the perturbation of a spiked covariance matrix.

The spectral representation formula is actually deterministic. Let the symmetric matrix $\Delta \in \mathbb{R}^{p\times p}$ be an arbitrary perturbation. Denote $\widehat{U}$ the top-r eigenvectors of $\Sigma + \Delta$ where $\Sigma = U\Lambda U^\top + \sigma^2 I_p$ with $\Lambda = \mathrm{diag}(\lambda_1, \cdots, \lambda_r)$. We are interested in developing an explicit representation formula for the spectral projector $\widehat{U}\widehat{U}^\top$ in terms of Δ. Let $Q^\perp := U_\perp U_\perp^\top = I_p - UU^\top$ denotes the orthogonal projection. For all $t \geq 1$, we define $Q^{-t} := U\Lambda^{-t}U^\top$. We slightly abuse the notation and denote $Q^0 := Q^\perp = U_\perp U_\perp^\top$.

LEMMA 3.1. *Suppose that Σ is a spiked covariance matrix as in (1) and $2\|\Delta\| \leq \lambda_r$, then*

$$\widehat{U}\widehat{U}^\top - UU^\top = \sum_{k\geq 1} \mathcal{S}_{\Sigma,k}(\Delta),$$

where the k-th order term $\mathcal{S}_{\Sigma,k}(\Delta)$ is a summation of $\binom{2k}{k}$ terms defined by

$$\mathcal{S}_{\Sigma,k}(\Delta) = \sum_{\mathbf{s}: s_1+\ldots+s_{k+1}=k} (-1)^{1+\tau(\mathbf{s})} \cdot Q^{-s_1}\Delta Q^{-s_2}\ldots\Delta Q^{-s_{k+1}},$$

where $\mathbf{s} = (s_1, \ldots, s_{k+1})$ contains non-negative indices and $\tau(\mathbf{s}) = \sum_{j=1}^{k+1}\mathbb{I}(s_j > 0)$. A simple upper bound of the k-th order term is

$$\left\|\mathcal{S}_{\Sigma,k}(\Delta)\right\| \leq \binom{2k}{k}\left(\frac{\|\Delta\|}{\lambda_r}\right)^k.$$

Based on Lemma 3.1, the leading term, i.e., the 1st-order term, of $\widehat{U}\widehat{U}^\top - UU^\top$ is contributed by $\Lambda^{-1}U^\top\Delta U_\perp$ and $U_\perp^\top\Delta U\Lambda^{-1}$. The latter terms can be sharply controlled by exploiting the statistical properties of Δ if observations are i.i.d. sampled.

3.2. *Upper bounds.* In this section, we present the upper bounds of our (ε, δ)-DP estimator $\widetilde{U}\widetilde{U}^\top$ and $\widetilde{\Sigma}$. In this section, we focus on the Gaussian setting, with an extension to the sub-Gaussian case provided in Section 5.2. Cases beyond sub-Gaussian distributions are discussed in Section 7.3. Let $\|\cdot\|_q$ denotes the matrix Schatten-q norm for any $q \in [1, \infty]$, e.g., the spectral norm $\|\cdot\|$ if $q = \infty$, the Frobenius norm $\|\cdot\|_{\mathrm{F}}$ if $q = 2$, and the nuclear norm $\|\cdot\|_*$ if $q = 1$. A straightforward application of the triangle inequality

$$\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_q \leq \|\widetilde{U}\widetilde{U}^\top - \widehat{U}\widehat{U}^\top\|_q + \|\widehat{U}\widehat{U}^\top - UU^\top\|_q,$$

leads to the following theorem.

THEOREM 3.2. *Suppose that $X_1, \cdots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \Sigma)$, $n \geq C_1(r\log n + \log^2 n)$, $2r \leq p$, and $\lambda/\sigma^2 \geq C_1(p/n + \sqrt{p/n})$ for some large absolute constant $C_1 > 0$. If we choose*

$$\Delta_1 := C_2\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\frac{\sqrt{p(r+\log n)}}{n},$$

then, there exist absolute constants $c_1, C_4 > 0$ such that, for any $\varepsilon > 0, \delta \in (0,1)$, Algorithm 1 outputs an (ε,δ)-DP estimator $\widetilde{U}\widetilde{U}^\top$ satisfying

$$\frac{\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_q}{r^{1/q}} \le C_4\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{r+\log n}}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}}\right),$$

with probability at least $1 - e^{-c_1(n\wedge p)}$. Moreover, if $\lambda/\sigma^2 \le (p/n)e^{c_2(p\wedge n)}$ for some small absolute constant $c_2 > 0$, then

$$\frac{\mathbb{E}\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_q}{r^{1/q}} \le C_4\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{r+\log n}}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}}\right).$$

Here, q can be any number in $[1,\infty]$.

Basically, the upper bounds consist of two parts: the first one represent the statistical error rate and the second one is the cost of privacy constraint. It is well-known that the first term is minimax optimal [13, 37, 45]. The second term decays at the rate $O\big(p/(n\varepsilon)\log^{1/2}\delta^{-1}\big)$ with respect to the sample size, dimension and privacy-related parameters, which is typical in differentially private algorithms [16, 42]. In Section 4, we shall develop matching minimax lower bounds showing that the rates in Theorem 3.2 are minimax optimal up the $\log n$ and $\log(2.5/\delta)$ terms.

It worth to mention that the $\log n$ term appearing in the privacy-related rate is due to the requirement of differential privacy that applies to each of the n observations. This $\log n$ term seems to be present in the upper bounds of most differentially private algorithms. See, e.g., [15, 16, 28] and references therein. A slight difference here is that the $\log n$ term appears not as an additional factor, but as an additive term. If $r \ge \log n$, the logarithmic factor can be ignored and the rate becomes minimax optimal except for the $\log\delta^{-1}$ factor.

REMARK 2 (Comparison with Oja's algorithm [42]). The DP-Oja algorithm introduced in [42, Corollary 5.2] delivered a rank-one ($r = 1$) PCA estimator achieving the following error bound, with probability 0.99:

$$\|\widehat{u}_{\mathsf{oja}}\widehat{u}_{\mathsf{oja}}^\top - uu^\top\| = \widetilde{O}\left(\left(1 + \frac{\sigma^2}{\lambda}\right)\cdot\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{\log 1/\delta}}{\varepsilon n}\right)\right),$$

where $\widetilde{O}(\cdot)$ hides logarithmic factors in n and p. Their established upper bound is much larger than ours when the signal strength $\lambda \gg \sigma^2$, and their failure probability is a constant while ours decay polynomially fast as sample size n increases.

We now present the performance bound for the differentially private estimator $\widetilde{\Sigma}$.

THEOREM 3.3. *Suppose that $X_1, \cdots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,\Sigma)$, $n \ge C_1(r\log n + \log^2 n), 2r \le p$, and $\lambda/\sigma^2 \ge C_1(p/n + \sqrt{p/n})$ for some large absolute constant $C_1 > 0$. If we choose*

$$\Delta_1 := C_2\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\frac{\sqrt{p(r+\log n)}}{n} \quad \text{and} \quad \Delta_2 := C_2\frac{\lambda(r+\log n) + \sigma^2(p+\log n)}{n},$$

then, there exist absolute constants $c_1, C_4 > 0$ such that, for any $\varepsilon > 0, \delta \in (0,1)$, Algorithm 1 outputs an (ε,δ)-DP estimator $\widetilde{\Sigma}$ satisfying

$$\frac{\|\widetilde{\Sigma} - \Sigma\|_q}{r^{1/q}}$$

$$\le C_4\left(\lambda\left(\sqrt{\frac{r}{n}} + \frac{\sqrt{r}(r+\log n)}{n\varepsilon}\cdot\sqrt{\log\frac{2.5}{\delta}}\right) + \sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{(r+\log n)}}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}}\right)\right),$$

with probability at least $1 - 3n^{-99} - e^{-c_1(n \wedge p)}$. *Moreover, if* $\lambda/\sigma^2 \le (p/n)e^{c_2(p\wedge n)}$ *for some small absolute constant* $c_2 > 0$, *then*

$$\frac{\mathbb{E}\|\widetilde{\Sigma} - \Sigma\|_q}{r^{1/q}}$$
$$\le C_4\left(\lambda\left(\sqrt{\frac{r}{n}} + \frac{\sqrt{r}(r+\log n)}{n\varepsilon}\cdot\sqrt{\log\frac{2.5}{\delta}}\right) + \sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{(r+\log n)}}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}}\right)\right).$$

By Theorem 3.3, the privacy-irrelevant error rate

$$\lambda\sqrt{\frac{r}{n}} + \sqrt{\sigma^2(\lambda+\sigma^2)}\sqrt{\frac{p}{n}},$$

matches the minimax optimal rate of spiked covariance estimation in the existing literature [13, 18]. For ease of discussion, let us focus on the error rate in spectral norm. There are two terms related to the cost of privacy:

$$\lambda\cdot\frac{\sqrt{r}(r+\log n)}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}} \quad \text{and} \quad \sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{(r+\log n)}}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}}\right),$$

where the second term is approximately of order $\lambda\|\widetilde{U}\widetilde{U}^\top - UU^\top\|$, contributed by the cost of estimating the eigenvectors. The first term grows at the rate $O(r^{3/2})$ with respect to the rank, which is contributed by the cost of estimating the eigenvalues. Due to the unknown orthogonal rotation measuring the alignment between $\widehat{U}$ and $\widetilde{U}$, privacy cost is also paid for the $r \times r$ unknown rotation matrix. Minimax lower bounds are developed in Section 4 demonstrating the optimality of these bound up to the $\log n$ and $\log(2.5/\delta)$ related terms.

REMARK 3 (Comparison with [28] and [43]). The DP-Gauss method is a privacy-preserving low-rank approximation method originally proposed by [28] and later improved by [43]. The method applies the Gaussian mechanism to find the rank-r approximation of the sample covariance matrix $\widehat{\Sigma}$, denoted by $\widehat{\Sigma}_r$ hereafter. Under the spiked covariance model, the DP-Gauss method provides an (ϵ, δ)-DP estimator, denoted as $\widetilde{\Sigma}_r$, achieving the rate (Corollary 2.3 in [43])

$$\left\|\widetilde{\Sigma}_r - \widehat{\Sigma}_r\right\|_{\mathrm{F}} \le C_3 \max_{i\in[n]}\|X_i\|^2\left(1+\frac{\sigma^2}{\lambda}\right)\frac{\sqrt{rp}}{n\varepsilon}\log^{1/2}(1/\delta), \tag{8}$$

which can be viewed as the cost of privacy in their method. Note that the term $\max_{i\in[n]}\|X_i\|_2$ appears here because [28] and [43] require that each observation has at most unit norm. Under the spiked model, we have $\max_i\|X_i\|^2 \asymp r\lambda + p\sigma^2$ up to $\log n$ factors. Plugging this into Equation (8), we can conclude that the bound attained by DP-Gauss in [43] is much larger than ours under the spiked covariance model.

Private covariance estimation for Gaussian distributions was studied by [34] using the Gaussian mechanism. Their rate is optimal in the case $\lambda \asymp \sigma^2$ and $r = p$, but becomes sub-optimal otherwise. In contrast, our rate is optimal, allowing a much more relaxed condition on λ and σ^2. Moreover, their method cannot be applied to differentially private PCA.

REMARK 4 (High-dimensional data). Our methods work as long as the signal-to-noise ratio satisfies $\lambda/\sigma^2 \gtrsim p/n + \sqrt{p/n}$, meaning that a strong signal is required when the dimension p is much larger than the sample size n. However, we emphasize that such a signal strength condition is necessary for a non-trivial estimate of the population eigenvectors, even in the conventional non-private setting. See, e.g. , [37]. Moreover, simulation results in Section 6 and Appendix A demonstrate that our method is much more robust than other methods (DP-Oja [42] and DP-Gauss [28, 43]) when dimension p is relatively large compared to the sample size n.

4. Minimax Lower Bounds with Differential Privacy. In this section, we establish the minimax lower bound of PCA and covariance matrix estimation under the constraint of differential privacy. Our main technical tool is a version of Fano's lemma with privacy constraint.

4.1. *DP-constrained Fano's Lemma.* Several techniques have been developed to establish minimax lower bounds under the constraint of differential privacy. Notable examples include the fingerprint method [33], Le Cam's method under differential privacy [6], differentially private Fano's lemma [3], and the recently introduced Score Attack method [16]. Le Cam's method and Fano's lemma construct a multitude of hypotheses that are difficult to distinguish, while the fingerprint method and Score Attack design a test statistic with a prior distribution.

For our convenience, we employ the differentially private Fano's lemma, as detailed in Lemma 4.1, whose proof is provided in Section C.5 of the Supplementary Materials [17]. Here, $\mathrm{KL}(\cdot,\cdot)$ and $\mathrm{TV}(\cdot,\cdot)$ denote the Kullback-Leibler divergence and total variation distance between two distributions.

LEMMA 4.1. *Let $\mathcal{P} := \{P : P = \mu^{(1)} \times \cdots \times \mu^{(n)}\}$ be a family of product measures indexed by a parameter from a pseudo-metric space (Θ, ρ). Denote $\theta(P) \in \Theta$ the parameter associated with the distribution P. Let $\mathcal{Q} = \{P_1, \cdots, P_N\} \subset \mathcal{P}$ contain N probability measures and there exist constants $\rho_0, l_0, t_0 > 0$ such that for all $i \neq i' \in [N]$,*

$$\rho\left(\theta(P_i), \theta(P_{i'})\right) \geqslant \rho_0, \quad \mathrm{KL}\left(P_i \| P_{i'}\right) \leq l_0,$$

and

$$\sum_{k\in[n]} \mathrm{TV}\left(\mu_i^{(k)}, \mu_{i'}^{(k)}\right) \leq t_0,$$

where $P_i = \mu_i^{(1)} \times \cdots \times \mu_i^{(n)}$ and $P_{i'} = \mu_{i'}^{(1)} \times \cdots \times \mu_{i'}^{(n)}$. Then,

$$\inf_{A\in\mathcal{A}_{\varepsilon,\delta}(\mathcal{P})} \sup_{P\in\mathcal{P}} \mathbb{E}_A\, \rho(A, \theta(P)) \geqslant \max\left\{\frac{\rho_0}{2}\left(1 - \frac{l_0 + \log 2}{\log N}\right), \frac{\rho_0}{4}\left(1 \bigwedge \frac{N-1}{\exp\left(4\varepsilon t_0\right)}\right)\left(1 - \frac{2\delta e^{4\varepsilon t_0}}{e^{\varepsilon} - 1}\right)\right\}, \tag{9}$$

where the infimum is taken over all the (ε, δ)-DP randomized algorithm defined by $\mathcal{A}_{\varepsilon,\delta}(\mathcal{P}) := \{A : X \mapsto \Theta \text{ and } A \text{ is } (\varepsilon,\delta)\text{-differentially private for all } X \sim P \in \mathcal{P}\}$.

Lemma 4.1 provides a powerful tool for developing a minimax lower bound in estimation problems under the constraint of differential privacy. Basically, if one can construct a sufficiently large set of distributions which are pairwise close in both Kullback-Leibler divergence and total variation distance, then a minimax lower bound can be derived if the underlying parameters are well-separated. The first term in the RHS of (9) is derived from the classic Fano's Lemma without privacy constraint and serves as a lower bound for the statistical error rate. This term is a well-established outcome in information theory by the framework of hypothesis testing and has been extensively employed in the statistics literature. The second term in the RHS of (9) characterizes the price one needs to pay for differential privacy. It is noteworthy that the cost of privacy is determined by t_0, which is the summation of marginal total variances. Intuitively, if the marginal total variance distances between $P_i = \mu_i^{(1)} \times \cdots \times \mu_i^{(n)}$ and $P_{i'} = \mu_{i'}^{(1)} \times \cdots \times \mu_{i'}^{(n)}$ are small , it becomes challenging to identify the distribution from which the dataset is drawn. Therefore, the cost of privacy is expected to be low when t_0 is small. Moreover, if we assume that $X = (X_1, \cdots, X_n) \sim P_i$, then the cost of privacy resulting from replacing $X_k \sim \mu_i^{(k)}$ by $X_k' \sim \mu_{i'}^{(k)}$ should be upper bounded in terms of $\mathrm{TV}(\mu_i^{(k)}, \mu_{i'}^{(k)})$. We remark that the bound given in (9) is meaningful only when $\delta \leq (e^{\varepsilon} - 1)e^{-4\varepsilon t_0}$.

4.2. *Minimax lower bounds.* In this section, we apply Lemma 4.1 to establish the minimax lower bounds for differentially private PCA and covariance estimation under the spiked covariance model. Denote the family of normal distribution with a spiked covariance matrix by

$$\mathcal{P}(\lambda, \sigma^2) := \Big\{ \mathcal{N}(0, \Sigma) : \Sigma = U\Lambda U^\top + \sigma^2 I_p \in \Theta(\lambda, \sigma^2) \Big\}.$$

By definition, each distribution $P \in \mathcal{P}(\lambda, \sigma^2)$ is indexed by the pair of eigenvalues Λ and eigenvectors $U \in \mathbb{O}_{p,r}$. We first focus on the minimax lower bounds for estimating the spectral projector $UU^\top$. Similarly, the minimax lower bounds are established in all Schatten-q norms for $q \in [1, \infty]$.

THEOREM 4.2. *Let the $p \times n$ data matrix X have i.i.d. columns sampled from a distribution $P = \mathcal{N}(0, U^\top \Lambda U^\top + \sigma^2 I_p) \in \mathcal{P}(\lambda, \sigma^2)$. Suppose $\delta \leq c_0' \exp\big\{2\varepsilon - c_0\big(\varepsilon\sqrt{npr} + pr\big)\big\}$ for some small constants $c_0, c_0' > 0$. Then, there exists an absolute constant $c_1 > 0$ such that*

$$\inf_{\widetilde{U} \in \mathcal{U}_{\varepsilon,\delta}} \sup_{P \in \mathcal{P}(\lambda,\sigma^2)} \frac{\mathbb{E}\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_q}{r^{1/q}} \geq c_1 \left(\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}} \right) \left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{r}}{n\epsilon} \right) \right) \bigwedge 1,$$

where the infimum is taken over all the possible (ε, δ)-DP algorithms, denoted by $\mathcal{U}_{\varepsilon,\delta}$, and the expectation is taken with respect to both $\widetilde{U}$ and P.

Theorem 4.2 imposes a strong restriction on the parameter δ. For most interesting cases, δ needs to be near zero. Therefore, the minimax lower bounds hold primarily for the pure differential privacy case, i.e., $\delta = 0$. It is worth noting the two terms in the minimax lower bound of spectral norm ($q = \infty$):

$$\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}} \right) \sqrt{\frac{p}{n}} \quad \text{and} \quad \left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}} \right) \frac{p\sqrt{r}}{n\varepsilon}. \tag{10}$$

The first term concerns the statistical error of PCA without privacy constraint. The error bound is free of the rank r, which is very typical in spectral norm error rate and the rate matches the existing minimax optimal rate of PCA for spiked covariance model. See, e.g., [13, 63, 64]. The second term is the price paid for differential privacy. Interestingly, the second term is dependent on the rank r even though spectral norm is considered here. The technical explanation is that the sensitivity of empirical spectral projector increases as the number of PC's grows. Comparing the two terms in (10), we observe that if $\varepsilon \geq (rp/n)^{1/2}$, the cost of privacy is dominated by the statistical error.

A minimax lower bound for *rank-one* PCA has been established in [42, Theorem 5.3]. Their developed rate in spectral norm also have two terms:

$$\sqrt{\frac{\sigma^2}{\lambda + \sigma^2}} \cdot \sqrt{\frac{p}{n}} \quad \text{and} \quad \sqrt{\frac{\sigma^2}{\lambda + \sigma^2}} \cdot \frac{p}{n\varepsilon}.$$

Their rate matches ours when $r=1$ and $\lambda \geq \sigma^2$. On the other hand, if $\lambda \ll \sigma^2$, our minimax lower bound is much stronger. Moreover, our minimax lower bounds hold for a diverging rank as long as $2r \leq p$.

We now shift our focus to the minimax lower bound of differentially private estimation of the spiked covariance matrix. Here, we assume σ^2 is known and it suffices to estimate the signal part $U\Lambda U^\top$. As a result, the minimax lower bound is essentially determined jointly by the lower bounds in estimating eigenvalues and eigenvectors.

THEOREM 4.3. *Let the $p \times n$ data matrix X have i.i.d. columns sampled from a distribution $P = \mathcal{N}(0, U^\top \Lambda U^\top + \sigma^2 I_p) \in \mathcal{P}(\lambda, \sigma^2)$. Suppose $\delta \leq c_0' \exp\left\{2\varepsilon - c_0\big(\varepsilon\sqrt{npr} + pr\big)\right\}$ for some small constants $c_0, c_0' > 0$. Then, there exists an absolute constant $c_1 > 0$ such that*

$$\inf_{\widetilde{\Sigma} \in \mathcal{M}_{\varepsilon,\delta}} \sup_{P \in \mathcal{P}(\lambda,\sigma^2)} \frac{\mathbb{E}\big\|\widetilde{\Sigma} - \Sigma\big\|_q}{r^{1/q}} \geq c_1 \left(\lambda\left(\sqrt{\frac{r}{n}} + \frac{r^{3/2}}{n\varepsilon}\right) + \sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p}{n}} + \frac{\sqrt{rp}}{n\varepsilon}\right)\right) \bigwedge \lambda,$$

where the infimum is taken over all the possible (ε,δ)-DP algorithms, denoted by $\mathcal{M}_{\varepsilon,\delta}$, and the expectation is taken with respect to both $\widetilde{\Sigma}$ and P. Here, q can be any number in $[1,\infty]$.

Without loss of generality, let us discuss the two terms in the spectral norm distance

$$\lambda\left(\sqrt{\frac{r}{n}} + \frac{r^{3/2}}{n\varepsilon}\right) \quad \text{and} \quad \sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p}{n}} + \frac{\sqrt{rp}}{n\varepsilon}\right). \tag{11}$$

The second term is contributed by the differentially private estimation error of PCA in the form of $\lambda\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_{\mathrm{F}}^2$. The first term dominates if the signal strength is exceedingly large, or more precisely, when $\lambda/\sigma^2 \gg p/r$. In this case, we can simply regard $\sigma = 0$ and the stochastic error mainly comes from the randomness of a low-dimensional distribution. Basically, it suffices to consider the minimax optimal estimation under a smaller family of normal distributions $\{\mathcal{N}(0, \lambda UU^\top + \lambda I_r) : \ U \in \mathbb{O}_{r,r/4}\}$. By replacing $\sigma \leftarrow \lambda$, $r \leftarrow r/4$, and $p \leftarrow r$, the second term reduces to the first term in (11). Without the privacy constraint, the first term also matches the existing optimal rate in covariance estimation under spiked covariance model [13, 18].

5. Extensions. For the sake of clarity, we have assumed uniformity in the order of spiked eigenvalues and Gaussian distributions. In this section, we extend our analysis to provide upper bounds for differentially private PCA and covariance estimation without requiring these specific conditions.

5.1. *Diverging condition number.* Suppose that $X_1, \cdots, X_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,\Sigma)$ where $\Sigma = U\Lambda U^\top + \sigma^2 I_p$ with spiked eigenvalues $\Lambda = \mathrm{diag}(\lambda_1, \cdots, \lambda_r)$. Denote $\kappa_0 := \lambda_1/\lambda_r$, the ratio of the largest and smallest spiked eigenvalues. The proof of Corollary 5.1 is almost identical to that of Theorems 3.2 and 3.3, and thus omitted. We only present the upper bounds of the expected error in Schatten norms, but high probability bounds hold similarly.

COROLLARY 5.1. *Suppose that $n \geq C_1(\kappa_0^2 r \log n + \log^2 n)$, $2r \leq p$, $\lambda_1/\sigma^2 \leq (p/n)e^{(p\wedge n)/C_1}$, and $\lambda_r/\sigma^2 \geq C_1(\kappa_0 p/n + \sqrt{p/n})$ for some large absolute constant $C_1 > 0$. If we choose*

$$\Delta_1 := C_2\left(\frac{\sigma^2}{\lambda_r} + \sqrt{\frac{\kappa_0\sigma^2}{\lambda_r}}\right)\frac{\sqrt{p(r+\log n)}}{n} \quad \text{and} \quad \Delta_2 := C_2\frac{\lambda_1(r+\log n) + \sigma^2(p+\log n)}{n},$$

then, there exist absolute constants $C_4 > 0$ such that, for any $\varepsilon > 0, \delta \in (0,1)$, Algorithm 1 outputs an (ε,δ)-DP estimators $\widetilde{U}\widetilde{U}^\top$ and $\widetilde{\Sigma}$ satisfying

$$\frac{\mathbb{E}\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_q}{r^{1/q}} \le C_4\left(\frac{\sigma^2}{\lambda_r} + \sqrt{\frac{\kappa_0\sigma^2}{\lambda_r}}\right)\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{r+\log n}}{n\varepsilon}\log^{1/2}\left(\frac{2.5}{\delta}\right)\right).$$

and

$$\frac{\mathbb{E}\|\widetilde{\Sigma} - \Sigma\|_q}{r^{1/q}}$$
$$\le C_4\left(\lambda_1\left(\sqrt{\frac{r}{n}} + \frac{\sqrt{r}(r+\log n)}{n\varepsilon}\cdot\sqrt{\log\frac{2.5}{\delta}}\right) + \sqrt{\sigma^2(\lambda_1+\sigma^2)}\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{(r+\log n)}}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}}\right)\right).$$

for all $q \in [1,\infty]$.

5.2. *Sub-Gaussian.* Suppose that X follows a sub-Gaussian distribution satisfying that, for any $u \in \mathbb{R}^p$, the following bound holds

$$\mathbb{E}\exp\left\{\frac{\langle X,u\rangle^2}{u^\top\Sigma u}\right\} \le 2,$$

where $\Sigma \in \Theta(\lambda,\sigma^2)$. For ease of exposition, we focus on the case of bounded condition number. Interestingly, the sensitivity of eigenvectors and eigenvalues is actually identical to that under Gaussian distributions.

COROLLARY 5.2. *Suppose that $n \ge C_1\big(r\log(p+n)\log^2 r + \log^2 n\big), 2r \le p$, $\lambda/\sigma^2 \le (p/n)e^{(p\wedge n)/C_1}$, and $\lambda/\sigma^2 \ge C_1(p/n + \sqrt{p/n})\log(p+n)$ for some large absolute constant $C_1 > 0$. If we choose*

$$\Delta_1 := C_2\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\frac{\sqrt{p(r+\log n)}}{n} \quad \text{and} \quad \Delta_2 := C_2\frac{\lambda(r+\log n) + \sigma^2(p+\log n)}{n},$$

then, there exist absolute constant $C_4 > 0$ such that, for any $\varepsilon > 0, \delta \in (0,1)$, Algorithm 1 outputs an (ε,δ)-DP estimators $\widetilde{U}\widetilde{U}^\top$ and $\widetilde{\Sigma}$ satisfying

$$\frac{\mathbb{E}\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_q}{r^{1/q}} \le C_4\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\left(\sqrt{\frac{p\log p}{n}} + \frac{p\sqrt{r+\log n}}{n\varepsilon}\log^{1/2}\left(\frac{2.5}{\delta}\right)\right),$$

and

$$\frac{\mathbb{E}\|\widetilde{\Sigma} - \Sigma\|_q}{r^{1/q}}$$
$$\le C_4\left(\lambda\left(\sqrt{\frac{r}{n}} + \frac{\sqrt{r}(r+\log n)}{n\varepsilon}\cdot\sqrt{\log\frac{2.5}{\delta}}\right) + \sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p\log p}{n}} + \frac{p\sqrt{r+\log n}}{n\varepsilon}\sqrt{\log\frac{2.5}{\delta}}\right)\right).$$

for all $q \in [1,\infty]$.

As shown by Corollary 5.2, the upper bounds of differentially private sub-Gaussian PCA and covariance estimation are almost the same as those for Gaussian distributions, implying that these bounds are minimax optimal. However, some additional logarithmic factors appear in the upper bound and signal-to-noise ratio condition when controlling the higher-order terms in spectral perturbation.

5.3. *Private estimation of nuisance variance.* In this section, we provide a differentially private estimator of σ^2, demonstrating that minimax optimal estimation of the spiked covariance matrix is still achievable even if σ^2 is unknown.

The estimation of σ^2 in spiked covariance models has been studied by [13, 24, 54]. These methods exploit the eigenvalues of the sample covariance matrix or the properties of the empirical spectral distribution. We utilize the robust estimator of σ^2, originally proposed by [35] for detecting the number of spikes. The basic idea is to average several bulk eigenvalues that are separated from the spike eigenvalues. We begin by reviewing the well-known Marchenko-Pastur (MP) law [5, 44].

Let Z be a $p \times n$ matrix whose entries are independent, identically distributed random variables with mean 0 and variance σ^2. Let $Y_n := ZZ^\top/n$ be the sample covariance matrix, and let $\lambda_1(Y_n) \geq \cdots \geq \lambda_{p\wedge n}(Y_n)$ be its non-zero eigenvalues. Define the empirical spectral distribution (ESD) of Y_n by $\mu_n(A) := (p \wedge n)^{-1} \sum_{j=1}^{p\wedge n} \mathbb{1}\big(\lambda_j(Y_n) \in A\big)$, $\forall A \subset \mathbb{R}$. Assume that $p/n \to \gamma$ as $n \to \infty$ for some $\gamma > 0$. It is known that the ESD converges in distribution to the MP distribution $\mu(\cdot)$, whose density function is given as follows.

DEFINITION 1. Given $\gamma > 0$, the zero-excluded MP distribution is defined by the density

$$f_{\gamma,\sigma^2}(x) = \frac{1}{2\pi\sigma^2} \cdot \frac{1}{x(1\wedge\gamma)} \sqrt{(x - \sigma^2\gamma_-)(\sigma^2\gamma_+ - x)} \cdot \mathbb{1}\big(x \in [\sigma^2\gamma_-,\, \sigma^2\gamma_+]\big),$$

where $\gamma_\pm := (1 \pm \sqrt{\gamma})^2$.

Since $\mathbb{E}Y_n = \sigma^2 I_p$, a natural estimator of σ^2 is by taking the average of several bulk eigenvalues of Y_n. Recall the sample covariance matrix $\widehat{\Sigma}$ under the spiked model (1), and let $\lambda_1(\widehat{\Sigma}) \geq \cdots \geq \lambda_{p\wedge n}(\widehat{\Sigma})$ denote the non-zero eigenvalues of $\widehat{\Sigma}$. The eigenvalue sticking property [8, Theorem 2.7] tells that $\lambda_{j+r}(\widehat{\Sigma}) \approx \lambda_j(Y_n)$ with high probability for all $(p\wedge n)/4 \leq j \leq 3(p\wedge n)/4$ if $r \ll (p \wedge n)$. Denote q_k the $k/(p\wedge n)$-upper quantile of the MP distribution with $\gamma_n := p/n$ and $\sigma^2 = 1$, i.e., $\int_{q_k}^{(1+\sqrt{\gamma_n})^2} f_{\gamma_n,1}(x)dx = k/(p\wedge n)$.

We define the non-private estimator of σ^2 by

$$\widehat{\sigma}^2 := \frac{\sum_{(p\wedge n)/4 \leq k \leq 3(p\wedge n)/4} q_k \lambda_k(\widehat{\Sigma})}{\sum_{(p\wedge n)/4 \leq k \leq 3(p\wedge n)/4} q_k^2}.$$

The convergence rate of $\widehat{\sigma}^2$ was established by Theorem 1 of [35]. Its sensitivity is characterized in Theorem 5.3. Let $\widehat{\sigma}^{(i)2}$ be defined as $\widehat{\sigma}^2$ using $\widehat{\Sigma}^{(i)}$ instead of $\widehat{\Sigma}$.

THEOREM 5.3. *Suppose the conditions in Theorem 3.3 hold, $r \leq C_3$ for any large constant C_3, and $p/n \to \gamma$ for a constant $\gamma > 0$. Then, there exists an absolute constant $C_4 > 0$ such that, with probability at least $1 - n^{-100}$,*

$$\big|\widehat{\sigma}^2 - \widehat{\sigma}^{(i)2}\big| \leq \Delta_3 := \frac{C_4}{\sqrt{p\wedge n}} \cdot \frac{\lambda(r + \log n) + \sigma^2(p + \log n)}{n}.$$

Consequently, the estimator $\widetilde{\sigma}^2 := \big|\widehat{\sigma}^2 + \mathcal{N}\big(0, 18(\Delta_3/\varepsilon)^2 \log(3.75/\delta)\big)\big|$ is an $(\varepsilon/3, \delta/3)$-DP estimate of σ^2 with probability at least $1 - n^{-100}$.

We now study the DP PCA and covariance estimator using the private estimate $\widetilde{\sigma}^2$.

THEOREM 5.4. *Suppose the conditions in Theorem 5.3, $p/n \to \gamma$ for a constant $\gamma > 0$, and there exists a small constant $c_1 > 0$ such that*

$$\frac{\lambda}{\sigma^2} \le c_1 \frac{n\sqrt{(p \wedge n)/\log n}}{(r+\log n)} \cdot \frac{\varepsilon}{\sqrt{\log(4/\delta)}} \quad \text{and} \quad \frac{p}{n}\sqrt{\frac{\log n}{p \wedge n}} \cdot \frac{\varepsilon}{\sqrt{\log(4/\delta)}} \le c_1 \tag{12}$$

Let $\widetilde{U}, \widetilde{\Lambda}$, and $\widetilde{\Sigma}$ be defined as Algorithm 1 with replacing $(\varepsilon, \delta, \sigma^2)$ by $(\varepsilon/3, \delta/3, \widetilde{\sigma}^2)$, respectively. Then, with probability at least $1 - e^{-c_2(n \wedge p)} - 4n^{-99}$, $\widetilde{U}$ and $\widetilde{\Lambda}$ are $(\varepsilon/3, \delta/3)$-DP, $\widetilde{\Sigma}$ is (ε, δ)-DP, and the following bounds hold

$$\|\widetilde{U}\widetilde{U}^\top - UU^\top\|_{\mathrm{F}} \le C_4\left(\frac{\sigma^2}{\lambda} + \sqrt{\frac{\sigma^2}{\lambda}}\right)\left(\sqrt{\frac{pr}{n}} + \frac{p\sqrt{r(r+\log n)}}{n\varepsilon}\sqrt{\log\frac{4}{\delta}}\right),$$

and

$$\max\left\{\|\widetilde{\Sigma} - \Sigma\|, \|\widetilde{\Sigma} - \Sigma\|_{\mathrm{F}}\right\}$$

$$\le C_{5,\gamma}\left(\lambda\left(\sqrt{\frac{r}{n}} + \frac{(r+\log n)^{3/2}}{n\varepsilon} \cdot \sqrt{\log\frac{4}{\delta}}\right) + \sqrt{\sigma^2(\lambda+\sigma^2)}\left(\sqrt{\frac{p}{n}} + \frac{p\sqrt{(r+\log n)}}{n\varepsilon}\sqrt{\log\frac{4}{\delta}}\right)\right),$$

where $c_2, C_{5,\gamma} > 0$ are constants.

These rates are nearly identical to those in Theorems 3.2 and 3.3, making them minimax optimal up to logarithmic factors. Condition (12) is imposed to ensure that the outputs of Algorithm 1 remain differentially private when using the private estimate $\widetilde{\sigma}^2$ instead of the true σ^2. Essentially, this condition guarantees that $\widetilde{\sigma}^2 \asymp \sigma^2$ with high probability, ensuring that the artificial noise added in Algorithm 1 is sufficiently strong to maintain the privacy guarantee.

6. Numerical Experiments.

6.1. *Simulations.* We present simulation results comparing the performance of our differentially private (DP) algorithms with existing methods in the literature. The DP-Oja algorithm, proposed by [42], estimates the first principal component under privacy constraints by extending Oja's algorithm, originally introduced by [48] for online PCA. We also compare it with the DP-Gauss algorithm proposed by [28], which uses the Gaussian mechanism for privacy-preserving PCA. Both our method and DP-Gauss use the Gaussian mechanism to ensure privacy. However, our method is specifically motivated by the spiked covariance model, while DP-Gauss is designed for deterministic data, assuming each observation has at most unit norm. This distinction is also reflected in the artificial noise levels introduced by the two methods. For the spiked covariance model, we implement the DP-Gauss method after applying a universal scaling to ensure each observation has at most unit norm. The DP-Oja method, as discussed by [42], is directly applicable to the spiked covariance model for Gaussian distributions.

Under the spiked covariance model (1), we carefully and fairly choose the constants involved in the artificial noise level for the DP-Oja, DP-Gauss, and our methods. For simplicity, we set the nuisance variance parameter $\sigma^2 = 1$ in all experiments. In our method, the constant C_2 in the definition of Δ_1 in Lemma 2.2 arises from the concentration of the sub-exponential random variable $\max_{i\in[n]} \|U^\top X_i X_i^\top U_\perp\|$ (see Lemma 8.3). In DP-Gauss, a scaling by $\left(\max_{i\in[n]} \|X_i\|^2\right)^{-1}$ is applied to the sample covariance matrix to enforce the unit norm upper bound assumption. By Lemma 8.2, $\max_{i\in[n]} \|X_i\|^2 \le (r + C_5 \log n)\lambda + p\sigma^2$.

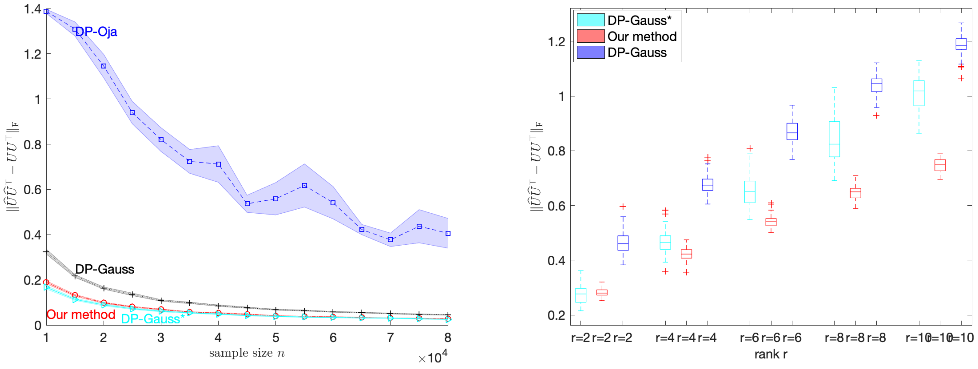

(a) Fixed rank $r = 1$.

(b) Fixed sample size $n = 10^4$.

Fig 1: Comparison of our method, DP-Oja [42], and DP-Gauss, DP-Gauss* [28] in differentially private PCA with varying n and r. The dimension $p = 50$, $\lambda = 10, \sigma^2 = 1$, and privacy constraints $\varepsilon = 1, \delta = 0.1$.

The scaling factor is set as $(r + C_5 \log n)\lambda + p\sigma^2$ in DP-Gauss. For a fair comparison, we set both the C_2 in our method and the C_5 in the scaling factor in DP-Gauss to 4. We can also set the scaling factor in DP-Gauss exactly as $\left(\max_{i\in[n]} \|X_i\|^2\right)^{-1}$. The resultant algorithm is denoted as DP-Gauss*. The artificial noise level in DP-Oja involves several unsettled constant factors, with the dominating one determined by the concentration property of $\max_{i\in[n]} \|X_i X_i^\top \hat{u}_{i-1}\|$ (see Lemma 3.2 in [42]). For a fair comparison, we should set the constant $C = 2$ in their Algorithm 2. However, since DP-Oja is an online PCA algorithm with random initialization, it is unsurprising that this method significantly underperforms our method and DP-Gauss. For clear illustration, we set $C = 0.2$ in their Algorithm 2. Note that DP-Oja only works for rank-one PCA, and its stepsize is set as $0.5/n$ after fine tuning for the best performance.

In the first simulation setting, we set $p = 50$, $\Sigma = \lambda UU^\top + I_p$ with $\lambda = 10$, and U contains the left singular vectors of a $p \times r$ random matrix with i.i.d. $N(0, 1)$ entries. The privacy budget is set as $\varepsilon = 1$ and $\delta = 0.1$. We first set $r = 1$ and compare the utility performances of the three methods in terms of the error $\|\widehat{U}\widehat{U}^\top - UU^\top\|_\mathrm{F}$ as the sample size n varies. For each n, the simulation is repeated for 40 times, and the average error and standard deviation are recorded. The results are displayed in the left panel of Figure 1. It shows that DP-Oja significantly underperforms compared to other methods, while our method slightly outperforms DP-Gauss. Our method and DP-Gauss* achieve similar performance. Moreover, we observe that DP-Oja runs much slower than the others. We then compare our method and DP-Gauss while varying the rank r. The boxplots shown in the right panel of Figure 1 are based on 40 independent simulations, demonstrating that our method consistently outperforms DP-Gauss for different ranks. Interestingly, as the rank r increases, our method becomes better than DP-Gauss*.

The second simulation setting compares these methods with respect to the varying privacy parameter ε and signal strength λ. The results, presented in Figure 2, show that the error rates of all methods increase rapidly as ε decreases, which aligns with our theoretical predictions. When the signal strength λ is small, DP-Gauss* performs the best. However, as λ increases, our method outperforms the others.

The third simulation setting aims to test the performance of our method and others in the high-dimensional case where $p \geq n$. We demonstrate that our method is applicable as long as the signal strength condition $\lambda/\sigma^2 \geq C(p/n + \sqrt{p/n})$ hold. The dimension, rank,

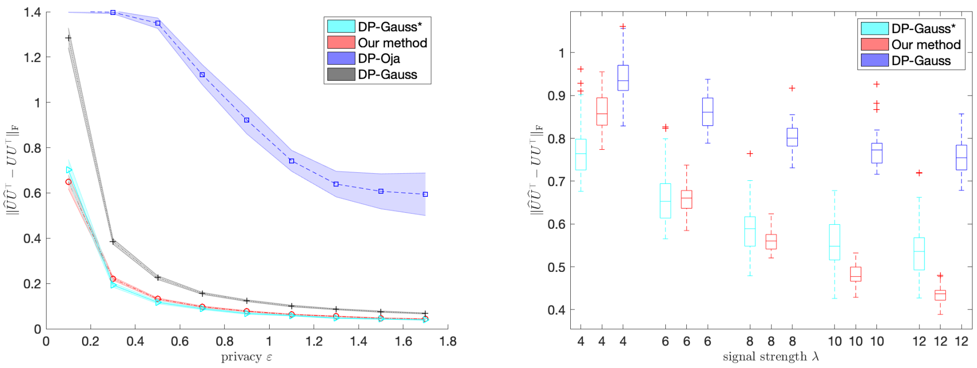

(a) Fixed rank $r = 1$, $\lambda = 10$, and sample size $n = 3 \times 10^4$.

(b) Fixed rank $r = 5$, $\varepsilon = 1$, and sample size $n = 10^4$.

Fig 2: Comparison of our method, DP-Oja [42], and DP-Gauss, DP-Gauss* [28] in differentially private PCA with varying ε and λ. The dimension $p = 50$, $\sigma^2 = 1$, and privacy constraint $\delta = 0.1$.

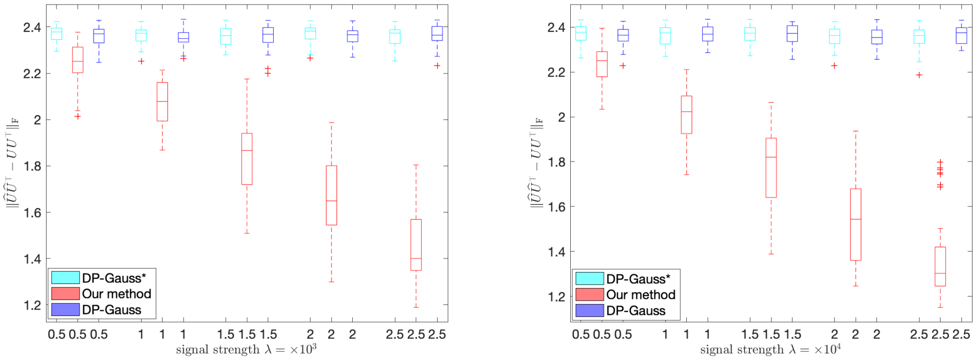

(a) Privacy constraint $\varepsilon = 3$.

(b) Privacy constraint $\varepsilon = 1$.

Fig 3: Comparison of our method, DP-Gauss, and DP-Gauss* [28] in differentially private PCA when $p \geq n$ and the signal strength λ changes. The dimension $p = 50$, $n = 30, r = 3, \sigma^2 = 1$, and privacy constraints $\delta = 0.1$.

and sample size are set to $p = 50$, $r = 3$, and $n = 30$, respectively. The boxplots of error $\|\widehat{U}\widehat{U}^\top - UU^\top\|_F$ based on 40 simulations are displayed in Figure 3. They show that our method significantly outperforms DP-Gauss and DP-Gauss* in this setting. The error rates of DP-Gauss and DP-Gauss* hardly improve as the signal strength increases. One possible reason is that the universal scaling procedure used in DP-Gauss and DP-Gauss* leads to significant information loss, especially when sample size n is small.

In the fourth simulation settings, we compare the performance of our method, DP-Gauss, and DP-Gauss* for covariance estimation. Due to space constraints, the results are provided in Appendix A.

6.2. *MNIST dataset.* We implemented our method, DP-Gauss, and DP-Gauss* for differentially private PCA on the MNIST dataset, which contains grayscale images of handwritten digits from 0 to 9. The dataset includes 500 samples for each digit. For a clear illustration,

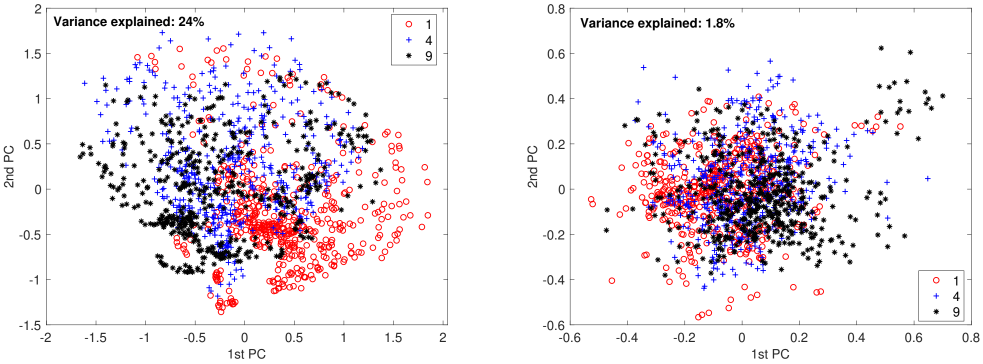

(a) Our method

(b) DP-Gauss*

Fig 4: Comparison of our method and DP-Gauss* [28] in differentially private PCA on MNIST dataset. The privacy constraints are $\varepsilon = 2$ and $\delta = 0.1$. The total sample size is $n = 1500$. All images are downscaled to a size 14×14.

we used only the images corresponding to digits 1, 4, and 9, creating a sample of $n = 1,500$ images. Each image is an observation of length $p = 28 \times 28 = 784$. The dimension is relatively large compared to the sample size. We downscaled the original images to 14×14. The rank was set to $r = 3$. We estimated λ by averaging the first three eigenvectors of the sample covariance matrix and σ^2 by the mean the 50-th to 140-th sample eigenvalues. The privacy constraints were set as $\varepsilon = 2$ and $\delta = 0.1$. After obtaining the DP estimates of eigenvectors, we applied dimension reduction to each observation, reported the variance explained, and visualized the scores corresponding to first and second components. The results are shown as in Figure 4. As expected, the DP-Gauss* method performed poorly due to the relatively small sample size. These methods add too much noise, resulting in principal components that can explain only 2% of the total variance.

7. Discussion.

In this paper, we establish optimal rates of convergence, up to logarithmic factors, for differentially private estimation of both the principal components and the covariance matrix under the spiked covariance model. We propose computationally efficient algorithms, and our results accommodate a diverging rank and a wider range of signal strengths.

7.1. *Private estimation of unknown rank.*

In the present paper, we assume that the number of components, r, in the spiked covariance model is known. However, in practice, r is typically unknown. The consistent estimation of the rank r in such models has been extensively studied in the conventional setting (see, e.g., [41], [13], [35] and references therein). For instance, we can use the eigen-ratio estimator $\widehat{r} := \arg\max_{1 \le k \le R} \big(\lambda_k(\widehat{\Sigma}) + Z_k\big)\big(\lambda_{k+1}(\widehat{\Sigma}) + Z_{k+1}\big)^{-1}$, where $Z_1, \cdots, Z_R$ are i.i.d centered $\mathcal{N}\big(0, 8\Delta_2^2\varepsilon^{-2}\log(2.5/\delta)\big)$ noise, and $R \le (n \wedge p)$ is a postulated upper bound for r. Lemma 2.4 ensures that $\widehat{r}$ is $(\varepsilon/2, \delta/2)$-DP with probability $1 - 4n^{-99}$. By the concentration property of edge eigenvalues and the sticking property of bulk eigenvalues of the sample covariance matrix [8], we can show that $\widehat{r}$ is a consistent estimate of r as long as the signal-to-noise ratio satisfies $\lambda/\sigma^2 \ge C_0 p\sqrt{\log(R)\log(2.5/\delta)}/(n\varepsilon)$ for some large enough constant $C_0 > 0$, in addition to the conditions required in Theorem 3.3.

7.2. *Private estimation of unknown eigenvalue.*

We assume the signal strength λ is known for simplicity. It can be estimated by averaging the first r eigenvalues of the sample covariance matrix. This approach works well in our numerical experiment on the MNIST dataset. The properties of the sample eigenvalues under the spiked covariance model are well-understood (see, e.g., [8] for a precise characterization). To protect privacy, one can also perturb sample eigenvalues with Gaussian noise, as discussed above using sensitivity Δ_2 developed in Lemma 2.4. The resultant estimator $\widetilde{\lambda}$ is consistent in the sense that $C_0^{-1} \le |\widetilde{\lambda}/\lambda| \le C_0$ with high probability under the conditions of Theorem 3.3 and if $n\varepsilon \ge C_1(r+\log n)$ and $\lambda/\sigma^2 \ge C_1 p\sqrt{\log(2.5/\delta)}/(n\varepsilon)$, where $C_0 > 0$ is some absolute constant and $C_1 > 0$ is a large constant depending only on C_0.

7.3. *Beyond sub-Gaussian distribution.* Our differentially private PCA method can be extended to distributions beyond the sub-Gaussian case. The primary technical challenge lies in analyzing the sensitivity of the sample spectral projector. The leading term of sensitivity is primarily determined by the quantity $\Delta_1 \approx \max_{i\in[n]} \|U^\top X_i\|\|U_\perp^\top X_i\|/(n\lambda)$, where $(\lambda+\sigma^2)I_r \preccurlyeq \operatorname{cov}\big(U^\top X_i\big) \preccurlyeq (\lambda+\sigma^2)I_r$ and $\sigma^2 I_{p-r} \preccurlyeq \operatorname{cov}(U_\perp^\top X_i) \preccurlyeq \sigma^2 I_{p-r}$. Assuming that $U^\top X_i/\lambda$ (also $U_\perp^\top X_i/\sigma^2$) has independent entries with a finite fourth moment, we can show that $\|U^\top X_i\|\|U_\perp^\top X_i\|/(n\lambda) \asymp (\sigma^2/\lambda + \sqrt{\sigma^2/\lambda})\sqrt{p(r+\log n)}/n$ (same as Lemma 2.3) up to additional logarithmic factors with probability $1 - O\big(\log^{-1}(n)\big)$. Therefore, we believe that our established bounds for differentially private PCA and covariance matrix estimation are optimal for a wide range of distributions. However, completing the proof still requires establishing sharp upper bounds for the following terms: $\|\widehat{\Sigma}-\Sigma\|$, $\|U^\top(\widehat{\Sigma}-\Sigma)U_\perp\|$, $\|X_i\|$, $\|U_\perp^\top X_i\|$, $\|U^\top X_i\|$, $\big\|h_1\big(\sum_{j\ne i} X_j X_j^\top\big)U_\perp^\top X_i\big\|$, and, $\big\|h_2\big(\sum_{j\ne i} X_j X_j^\top\big)U^\top X_i\big\|$, etc., which should hold for all $i\in[n]$, where $h_1(\cdot)$ and $h_2(\cdot)$ are some deterministic functions. It is possible to establish these bounds by imposing some moment condition on the distribution of X, such as requiring $\mathbb{E}|\langle X, a\rangle|^m \le L^m$ for some sufficiently large $m \ge 4$ and constant $L > 0$. See, e.g., [5] and [8]. Developing these bounds would significantly complicate the current proof and presentation of theoretical results, introduce additional logarithmic factors, and weaken the privacy guarantee. We leave these extensions for future work.

8. Technical lemmas.

In this section, we provide some technical lemmas that will be frequently used in the subsequent proofs. Due to page constraints, all proofs are given in the supplement [17].

Lemma 8.1 is a well-known *dimension-free* concentration inequality of sample covariance matrix developed by [37]. Here, $\|\cdot\|$ denote the spectral norm of a matrix and ℓ_2-norm of a vector.

LEMMA 8.1 ([37]). *Suppose $X_1,\cdots,X_n$ are i.i.d. sampled from $\mathcal{N}(0,\Sigma)$ and $\widehat{\Sigma} := \sum_{i=1}^n X_iX_i^\top/n$. Then,*

$$\mathbb{E}\|\widehat{\Sigma}-\Sigma\| \asymp \left(\sqrt{\frac{\operatorname{tr}(\Sigma)\|\Sigma\|}{n}} \bigvee \frac{\operatorname{tr}(\Sigma)}{n}\right).$$

Moreover, there exists an absolute constant $C_1 > 0$ such that, for all $t \ge 1$, with probability at least $1-e^{-t}$,

$$\Big|\|\widehat{\Sigma}-\Sigma\| - \mathbb{E}\|\widehat{\Sigma}-\Sigma\|\Big| \le C_1\|\Sigma\|\left(\frac{t}{n} + \sqrt{\frac{t}{n}}\left(1+\sqrt{\frac{\operatorname{tr}(\Sigma)/\|\Sigma\|}{n}}\right)\right).$$

The following lemma characterizes the concentration of the norm of a Gaussian random vector. Recall that X_i' is an independent copy of X_i.

LEMMA 8.2. *Let $X \sim \mathcal{N}(0, \Sigma)$ and the eigenvalues of Σ are $\lambda_1 \geq \cdots \geq \lambda_p \geq 0$. Then, there exist absolute constants $C_1, C_2, c_1 > 0$ such that*

$$\mathbb{P}\Big(\Big|\|X\|^2 - \mathrm{tr}(\Sigma)\Big| \leq C_1\Big(u\sum_{i=1}^{p}\lambda_i^2\Big)^{1/2} + C_2\lambda_1 u\Big) \geq 1 - e^{-c_1 u},$$

for any $u > 0$. Under the spiked covariance model $\Sigma \in \Theta(\lambda, \sigma^2)$ and the condition that $p \geq C_6 \log n$ for some absolute constant $C_6 > 0$, we have

$$\begin{aligned}\mathbb{P}\Big(\Big\{&\max_{i\in[n]}\|X_i\|^2 + \|X_i'\|^2 \leq C_3(r\lambda + p\sigma^2) + C_4\sqrt{(r\lambda^2 + p\sigma^4)\log n} + C_5(\lambda + \sigma^2)\log n\Big\}\Big)\\ &\bigcap\Big\{\max_{i\in[n]}\|U^\top X_i\|^2 + \max_{i\in[n]}\|U^\top X_i'\|^2 \leq C_3 r(\lambda + \sigma^2) + C_4\sqrt{r(\lambda^2 + \sigma^4)\log n} + C_5(\lambda + \sigma^2)\log n\Big\}\\ &\bigcap\Big\{\max_{i\in[n]}\|U_\perp^\top X_i\|^2 + \max_{i\in[n]}\|U_\perp^\top X_i'\|^2 \leq C_3 p\sigma^2\Big\} \geq 1 - n^{-100},\end{aligned}$$

where $C_3, C_4, C_5 > 0$ are some absolute constants. Let $\mathcal{E}_0$ denote the above event. Moreover,

$$\mathbb{E}\|X_i\|^2 \leq C_3(r\lambda + p\sigma^2), \quad \mathbb{E}\|U^\top X_i\|^2 \leq C_3 r(\lambda + \sigma^2) \quad \text{and} \quad \mathbb{E}\|U_\perp^\top X_i\|^2 \leq C_3 p\sigma^2.$$

Denote $\Delta := \widehat{\Sigma} - \Sigma$ and $\Delta^{(i)} := \widehat{\Sigma}^{(i)} - \Sigma$. We shall frequently use several concentration bounds related to Δ and $\Delta^{(i)}$ throughout the proof. For reader's convenience, these concentration bounds are collected in the following lemma.

LEMMA 8.3. *Suppose that $\Sigma \in \Theta(\lambda, \sigma^2)$, $n \geq C_1(r\log n + \log^2 n)$, $2r \leq p$, and $\lambda/\sigma^2 \geq C_1 p/n$ for some absolute constant $C_1 > 0$. There exist absolute constants $c_0 > 0$ and $C_2 > 0$ such that the event*

$$\begin{aligned}\mathcal{E}_\Delta := &\left\{\|\Delta\| + \max_{i\in[n]}\|\Delta^{(i)}\| \leq C_2\sqrt{\frac{(\lambda + \sigma^2)(r\lambda + p\sigma^2)}{n}} + \frac{\lambda}{10}\right\}\\ &\bigcap\left\{\big\|U^\top\Delta U_\perp\big\| + \max_{i\in[n]}\|U^\top\Delta^{(i)}U_\perp\| \leq C_3\sqrt{\frac{\sigma^2(\lambda + \sigma^2)p}{n}}\right\}\end{aligned}\tag{13}$$

holds with probability $\mathbb{P}(\mathcal{E}_\Delta) \geq 1 - e^{-c_0(n\wedge p)}$. Meanwhile, we have

$$\mathbb{E}\|\Delta\| \leq C_2\sqrt{\frac{(\lambda + \sigma^2)(r\lambda + p\sigma^2)}{n}} \quad \text{and} \quad \mathbb{E}\big\|U^\top\Delta U_\perp\big\| \leq C_3\sqrt{\frac{\sigma^2(\lambda + \sigma^2)p}{n}}.$$

There exist absolute constants $c_1 > 0$ and $C_3 > 0$ such that the event

$$\begin{aligned}\mathcal{E}_1 := &\left\{\max_{i\in[n]}\big\|U^\top(X_iX_i^\top/n)U_\perp\big\| + \big\|U^\top(X_i'X_i'^\top/n)U_\perp\big\| \leq C_3\frac{\sqrt{\sigma^2(\lambda + \sigma^2)p(r + \log n)}}{n}\right\}\\ &\bigcap\left\{\max_{i\in[n]}\left\|U^\top\Big(\frac{1}{n}\sum_{j\neq i}X_jX_j^\top\Big)U_\perp U_\perp^\top X_i\right\| \leq C_3\sigma\cdot\sqrt{\frac{\sigma^2(\lambda + \sigma^2)p(r + \log n)}{n}}\right\}\end{aligned}\tag{14}$$

holds with probability $\mathbb{P}(\mathcal{E}_1) \geq 1 - e^{-c_1(n\wedge p)} - 2n^{-99}$.

The following perturbation bound of principal subspace will be useful.

LEMMA 8.4. *Suppose that $\lambda_r \geq (4 + \delta)\|\Delta\|$ for some $\delta > 0$, then*

$$\big\|\widehat{U}\widehat{U}^\top - UU^\top\big\| \leq 2\left\|\Lambda^{-1}U^\top\Delta U_\perp\right\| + 6(4 + \delta)\frac{\|\Delta\|\,\big\|U^\top\Delta U_\perp\big\|}{\delta\lambda_r^2}.$$

Funding. The research of Tony Cai was supported in part by NIH grants R01-GM123056 and R01-GM129781 and NSF grant DMS-2413106. Dong Xia's research was partially supported by Hong Kong RGC Grant GRF 16302323.

SUPPLEMENTARY MATERIAL

Supplement to "Optimal Differentially Private PCA and Estimation for Spiked Covariance Matrices"
In this supplement, we present the proofs of Lemmas 2.2 through 2.4, additional technical lemmas, and all the theorems. Additionally, we provide simulation results for differentially private covariance matrix estimation.

REFERENCES

[1] ABOWD, J. M. (2016). The challenge of scientific reproducibility and privacy protection for statistical agencies. *Census Scientific Advisory Committee*.

[2] ABOWD, J. M., RODRIGUEZ, I. M., SEXTON, W. N., SINGER, P. E. and VILHUBER, L. (2020). The modernization of statistical disclosure limitation at the US Census Bureau. *US Census Bureau*.

[3] ACHARYA, J., SUN, Z. and ZHANG, H. (2021). Differentially Private Assouad, Fano, and Le Cam. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory* (V. FELDMAN, K. LIGETT and S. SABATO, eds.). *Proceedings of Machine Learning Research* **132** 48–78. PMLR.

[4] ANDERSON, T. W. (2003). *An Introduction to Multivariate Statistical Analysis*, 3rd ed. Wiley.

[5] BAI, Z. and SILVERSTEIN, J. W. (2010). *Spectral analysis of large dimensional random matrices* **20**. Springer.

[6] BARBER, R. F. and DUCHI, J. C. (2014). Privacy and Statistical Risk: Formalisms and Minimax Bounds. *CoRR* **abs/1412.4451**.

[7] BICKEL, P. J. and LEVINA, E. (2008). Regularized estimation of large covariance matrices. *The Annals of Statistics* **36** 199 – 227. https://doi.org/10.1214/009053607000000758

[8] BLOEMENDAL, A., KNOWLES, A., YAU, H.-T. and YIN, J. (2016). On the principal components of sample covariance matrices. *Probability theory and related fields* **164** 459–552.

[9] BLUM, A., DWORK, C., MCSHERRY, F. and NISSIM, K. (2005). Practical privacy: The SulQ framework. *Proceedings of the ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* 128-138. https://doi.org/10.1145/1065167.1065184

[10] BROWN, G., GABOARDI, M., SMITH, A., ULLMAN, J. and ZAKYNTHINOU, L. (2021). Covariance-aware private mean estimation without private covariance estimation. *Advances in neural information processing systems* **34** 7950–7964.

[11] CAI, T. T., CHAKRABORTY, A. and WANG, Y. (2023). Optimal differentially private ranking from pairwise comparisons. *Technical Report*.

[12] CAI, T. T., MA, Z. and WU, Y. (2013). Sparse PCA: Optimal rates and adaptive estimation. *The Annals of Statistics* **41** 3074-3110. https://doi.org/10.1214/13-AOS1178

[13] CAI, T. T., MA, Z. and WU, Y. (2015). Optimal estimation and rank detection for sparse spiked covariance matrices. *Probability Theory and Related Fields* **161** 781-815.

[14] CAI, T. T., REN, Z. and ZHOU, H. H. (2016). Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation. *Electronic Journal of Statistics* **10** 1–59. https://doi.org/10.1214/15-EJS1081

[15] CAI, T. T., WANG, Y. and ZHANG, L. (2021). The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics* **49** 2825-2850.

[16] CAI, T. T., WANG, Y. and ZHANG, L. (2023). Score attack: A lower bound technique for optimally differentially private learning. *arXiv preprint arXiv:2303.07152*.

[17] CAI, T. T., XIA, D. and ZHA, M. (2024). Supplement to "Optimal Differentially Private PCA and Estimation for Spiked Covariance Matrices".

[18] CAI, T. T., ZHANG, C.-H. and ZHOU, H. H. (2010). Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics* **38** 2118 – 2144. https://doi.org/10.1214/09-AOS752

[19] CHAUDHURI, K., MONTELEONI, C. and SARWATE, A. D. (2011). Differentially private empirical risk minimization. *Journal of Machine Learning Research* **12**.

[20] CHIEN, S., JAIN, P., KRICHENE, W., RENDLE, S., SONG, S., THAKURTA, A. and ZHANG, L. (2021). Private alternating least squares: Practical private matrix completion with tighter rates. In *International Conference on Machine Learning* 1877–1887. PMLR.

[21] Devroye, L. (1987). *A Course in Density Estimation*. Birkhauser Boston Inc.

[22] Ding, B., Kulkarni, J. and Yekhanin, S. (2017). Collecting telemetry data privately. *Advances in Neural Information Processing Systems* **30**.

[23] Dong, W., Liang, Y. and Yi, K. (2022). Differentially private covariance revisited. *Advances in Neural Information Processing Systems* **35** 850–861.

[24] Donoho, D. L., Gavish, M. and Johnstone, I. M. (2018). Optimal shrinkage of eigenvalues in the spiked covariance model. *The Annals of Statistics* **46** 1742.

[25] Duchi, J. C., Jordan, M. I. and Wainwright, M. J. (2018). Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association* **113** 182–201.

[26] Dwork, C., McSherry, F., Nissim, K. and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. *Theory of Cryptography Conference* 265-284.

[27] Dwork, C., Roth, A. et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* **9** 211–407.

[28] Dwork, C., Talwar, K., Thakurta, A. and Zhang, L. (2014). Analyze gauss: optimal bounds for privacy-preserving principal component analysis. *Proceedings of the forty-sixth annual ACM symposium on Theory of computing* 11–20.

[29] Erlingsson, U., Pihur, V. and Korolova, A. (2014). RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. *CCS '14* 1054–1067. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/2660267.2660348

[30] Fan, J., Fan, Y. and Lv, J. (2008). High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics* **147** 186–197.

[31] Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics* **29** 295 – 327. https://doi.org/10.1214/aos/1009210544

[32] Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* **104** 682–693.

[33] Kamath, G., Li, J., Singhal, V. and Ullman, J. (2019). Privately Learning High-Dimensional Distributions. In *Proceedings of the Thirty-Second Conference on Learning Theory* (A. Beygelzimer and D. Hsu, eds.). *Proceedings of Machine Learning Research* **99** 1853–1902. PMLR.

[34] Kamath, G., Li, J., Singhal, V. and Ullman, J. (2019). Privately learning high-dimensional distributions. In *Conference on Learning Theory* 1853–1902. PMLR.

[35] Ke, Z. T., Ma, Y. and Lin, X. (2023). Estimation of the number of spiked eigenvalues in a covariance matrix by bulk eigenvalue matching analysis. *Journal of the American Statistical Association* **118** 374–392.

[36] Koltchinskii, V. (2011). Von Neumann entropy penalization and low-rank matrix estimation. *The Annals of Statistics* **39** 2936 – 2973. https://doi.org/10.1214/11-AOS926

[37] Koltchinskii, V. and Lounici, K. (2017). Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli* 110–133.

[38] Koltchinskii, V. and Xia, D. (2015). Optimal Estimation of Low Rank Density Matrices. *Journal of Machine Learning Research* **16** 1757–1792.

[39] Koltchinskii, V. and Xia, D. (2016). Perturbation of linear forms of singular vectors under Gaussian noise. In *High Dimensional Probability VII: The Cargèse Volume* 397–423. Springer.

[40] Kritchman, S. and Nadler, B. (2008). Determining the number of components in a factor model from limited noisy data. *Chemometrics and Intelligent Laboratory Systems* **94** 19–32.

[41] Lam, C. and Yao, Q. (2012). Factor modeling for high-dimensional time series: inference for the number of factors. *The Annals of Statistics* 694–726.

[42] Liu, X., Kong, W., Jain, P. and Oh, S. (2022). DP-PCA: Statistically Optimal and Differentially Private PCA. *Advances in Neural Information Processing Systems* **35** 29929–29943.

[43] Mangoubi, O. and Vishnoi, N. (2022). Re-analyze Gauss: Bounds for private matrix approximation via Dyson Brownian motion. *Advances in Neural Information Processing Systems* **35** 38585–38599.

[44] Marchenko, V. A. and Pastur, L. A. (1967). Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik* **114** 507–536.

[45] Nadler, B. (2008). Finite sample approximation results for principal component analysis: a matrix perturbation approach. *The Annals of Statistics* **36** 2791–2817.

[46] Narayanan, S. (2023). Better and simpler lower bounds for differentially private statistical estimation. *arXiv preprint arXiv:2310.06289*.

[47] Nikolov, A. (2023). Private query release via the johnson-lindenstrauss transform. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* 4982–5002. SIAM.

[48] Oja, E. (1982). Simplified neuron model as a principal component analyzer. *Journal of mathematical biology* **15** 267–273.

[49] ONATSKI, A. (2012). Asymptotics of the principal components estimator of large factor models with weak factors. *J. Econometrics* **168** 244-258.

[50] PAJOR, A. (1998). Metric entropy of the Grassmann manifold. *Convex Geometric Analysis* **34** 0942–46013.

[51] PATTERSON, N., PRICE, A. L. and REICH, D. (2006). Population structure and eigenanalysis. *PLoS Genet.* **2** e190. https://doi.org/10.1371/journal.pgen.0020190

[52] RAVIKUMAR, P., WAINWRIGHT, M. J., RASKUTTI, G. and YU, B. (2011). High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence. *Electronic Journal of Statistics* **5** 935 – 980. https://doi.org/10.1214/11-EJS631

[53] ROHDE, A. and STEINBERGER, L. (2020). Geometrizing rates of convergence under local differential privacy constraints. *The Annals of Statistics* **48** 2646 – 2670. https://doi.org/10.1214/19-AOS1901

[54] SHABALIN, A. A. and NOBEL, A. B. (2013). Reconstruction of a low-rank matrix in the presence of Gaussian noise. *Journal of Multivariate Analysis* **118** 67–76.

[55] SRIVASTAVA, N. and VERSHYNIN, R. (2013). Covariance estimation for distributions with $2+\varepsilon$ moments. *The Annals of Probability* **41** 3081 – 3111. https://doi.org/10.1214/12-AOP760

[56] APPLE DIFFERENTIAL PRIVACY TEAM (2017). Learning with Privacy at Scale.

[57] TSYBAKOV, A. B. (2008). *Introduction to Nonparametric Estimation*. Springer.

[58] VERSHYNIN, R. (2012). How close is the sample covariance matrix to the actual covariance matrix? *Journal of Theoretical Probability* **25** 655–686.

[59] VERSHYNIN, R. (2020). High-dimensional probability. *University of California, Irvine*.

[60] WANG, L., ZHAO, B. and KOLAR, M. (2023). Differentially Private Matrix Completion through Low-rank Matrix Factorization. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics* (F. RUIZ, J. DY and J.-W. VAN DE MEENT, eds.). *Proceedings of Machine Learning Research* **206** 5731–5748. PMLR.

[61] XIA, D. (2021). Normal approximation and confidence region of singular subspaces. *Electronic Journal of Statistics* **15** 3798–3851.

[62] YU, B. (1997). Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam: research papers in probability and statistics* 423–435. Springer.

[63] YU, Y., WANG, T. and SAMWORTH, R. J. (2015). A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika* **102** 315–323.

[64] ZHANG, A. and XIA, D. (2018). Tensor SVD: Statistical and computational limits. *IEEE Transactions on Information Theory* **64** 7311–7338.

[65] ZHANG, A. R., CAI, T. T. and WU, Y. (2022). Heteroskedastic PCA: Algorithm, optimality, and applications. *The Annals of Statistics* **50** 53–80.