
Minimax Rates of Estimation for Sparse PCA in High Dimensions

Vincent Q. Vu

Department of Statistics
Carnegie Mellon University
vqv@stat.cmu.edu

Jing Lei

Department of Statistics
Carnegie Mellon University
jinglei@andrew.cmu.edu

Abstract

We study sparse principal components analysis in the high-dimensional setting, where p (the number of variables) can be much larger than n (the number of observations). We prove optimal, non-asymptotic lower and upper bounds on the minimax estimation error for the leading eigenvector when it belongs to an ℓ_q ball for $q \in [0, 1]$. Our bounds are sharp in p and n for all $q \in [0, 1]$ over a wide class of distributions. The upper bound is obtained by analyzing the performance of ℓ_q -constrained PCA. In particular, our results provide convergence rates for ℓ_1 -constrained PCA.

1 Introduction

High-dimensional data problems, where the number of variables p exceeds the number of observations n , are pervasive in modern applications of statistical inference and machine learning. Such problems have increased the necessity of dimensionality reduction for both statistical and computational reasons. In some applications, dimensionality reduction is the end goal, while in others it is just an intermediate step in the analysis stream. In either case, dimensionality reduction is usually data-dependent and so the limited sample size and noise may have an adverse affect. Principal components analysis (PCA) is perhaps one of the most well known and widely used techniques for unsupervised dimensionality reduction. However, in the high-dimensional situation, where p/n does not tend to 0 as $n \rightarrow \infty$, PCA may not give consistent estimates of eigenvalues and eigenvectors of the population covariance matrix [12]. To remedy this situation,

Appearing in Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS) 2012, La Palma, Canary Islands. Volume 22 of JMLR: W&CP 22. Copyright 2012 by the authors.

sparsity constraints on estimates of the leading eigenvectors have been proposed and shown to perform well in various applications. In this paper we prove optimal minimax error bounds for sparse PCA when the leading eigenvector is sparse.

1.1 Subspace Estimation

Suppose we observe i.i.d. random vectors $X_i \in \mathbb{R}^p$, $i = 1, \dots, n$ and we wish to reduce the dimension of the data from p down to k . PCA looks for k uncorrelated, linear combinations of the p variables that have maximal variance. This is equivalent to finding a k -dimensional linear subspace whose orthogonal projection A minimizes the mean squared error

$$\text{mse}(A) = \mathbb{E} \|(X_i - \mathbb{E}X_i) - A(X_i - \mathbb{E}X_i)\|_2^2 \quad (1)$$

[see 10, Chapter 7.2.3 for example]. The optimal subspace is determined by spectral decomposition of the population covariance matrix

$$\Sigma = \mathbb{E}X_iX_i^T - (\mathbb{E}X_i)(\mathbb{E}X_i)^T = \sum_{j=1}^p \lambda_j \theta_j \theta_j^T, \quad (2)$$

where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 0$ are the eigenvalues and $\theta_1, \dots, \theta_p \in \mathbb{R}^p$, orthonormal, are eigenvectors of Σ . If $\lambda_k > \lambda_{k+1}$, then the optimal k -dimensional linear subspace is the span of $\Theta = (\theta_1, \dots, \theta_k)$ and its projection is given by $\Pi = \Theta\Theta^T$. Thus, if we know Σ then we may optimally (in the sense of eq. (1)) reduce the dimension of the data from p to k by the mapping $x \mapsto \Theta\Theta^T x$.

In practice, Σ is not known and so Θ must be estimated from the data. In that case we replace Θ by an estimate $\hat{\Theta}$ and reduce the dimension of the data by the mapping $x \mapsto \hat{\Pi}x$, where $\hat{\Pi} = \hat{\Theta}\hat{\Theta}^T$. PCA uses the spectral decomposition of the sample covariance matrix

$$S = \frac{1}{n} \sum_{i=1}^n X_iX_i^T - \bar{X}\bar{X}^T = \sum_{j=1}^{p \wedge n} l_j u_j u_j^T,$$

where \bar{X} is the sample mean, and l_j and u_j are eigenvalues and eigenvectors of S defined analogously to

eq. (2). It reduces the dimension of the data to k by the mapping $x \mapsto UU^T x$, where $U = (u_1, \dots, u_k)$.

In the classical regime where p is fixed and $n \rightarrow \infty$, PCA is a consistent estimator of the population eigenvectors. However, this scaling is not appropriate for modern applications where p is comparable to or larger than n . In that case, it has been observed [18, 17, 12] that if $p, n \rightarrow \infty$ and $p/n \rightarrow c > 0$, then PCA can be an inconsistent estimator in the sense that the angle between u_1 and θ_1 can remain bounded away from 0 even as $n \rightarrow \infty$.

1.2 Sparsity Constraints

Estimation in high-dimensions may be beyond hope without additional structural constraints. In addition to making estimation feasible, these structural constraints may also enhance interpretability of the estimators. One important example of this is sparsity. The notion of sparsity is that a few variables have large effects, while most others are negligible. This type of assumption is often reasonable in applications and is now widespread in high-dimensional statistical inference.

Many researchers have proposed sparsity constrained versions of PCA along with practical algorithms, and research in this direction continues to be very active [e.g., 13, 27, 6, 21, 25]. Some of these works are based on the idea of adding an ℓ_1 constraint to the estimation scheme. For instance, Jolliffe, Trendafilov, and Uddin [13] proposed adding an ℓ_1 constraint to the variance maximization formulation of PCA. Others have proposed convex relaxations of the “hard” ℓ_0 -constrained form of PCA [6]. Nearly all of these proposals are based on an iterative approach where the eigenvectors are estimated in a one-at-a-time fashion with some sort of deflation step in between [14]. For this reason, we consider the basic problem of estimating the leading population eigenvector θ_1 .

The ℓ_q balls for $q \in [0, 1]$ provide an appealing way to make the notion of sparsity concrete. These sets are defined by

$$\mathbb{B}_q^p(R_q) = \{\theta \in \mathbb{R}^p : \sum_{j=1}^p |\theta_j|^q \leq R_q\}$$

and

$$\mathbb{B}_0^p(R_0) = \{\theta \in \mathbb{R}^p : \sum_{j=1}^p 1_{\{\theta_j \neq 0\}} \leq R_0\}.$$

The case $q = 0$ corresponds to “hard” sparsity where R_0 is the number of nonzero entries of the vectors. For $q > 0$ the ℓ_q balls capture “soft” sparsity where a few of the entries of θ are large, while most are small. The soft sparsity case may be more realistic for applications where the effects of many variables may be very small, but still nonzero.

1.3 Minimax Framework and High-Dimensional Scaling

In this paper, we use the statistical minimax framework to elucidate the difficulty/feasibility of estimation when the leading eigenvector θ_1 is assumed to belong to $\mathbb{B}_q^p(R_q)$ for $q \in [0, 1]$. The framework can make clear the fundamental limitations of statistical inference that *any estimator* $\hat{\theta}_1$ must satisfy. Thus, it can reveal gaps between optimal estimators and computationally tractable ones, and also indicate when practical algorithms achieve the fundamental limits.

Parameter space There are two main ingredients in the minimax framework. The first is the class of probability distributions under consideration. These are usually associated with some parameter space corresponding to the structural constraints. Formally, suppose that $\lambda_1 > \lambda_2$. Then we may write eq. (2) as

$$\Sigma = \lambda_1 \theta_1 \theta_1^T + \lambda_2 \Sigma_0, \quad (3)$$

where $\lambda_1 > \lambda_2 \geq 0$, $\theta_1 \in \mathbb{S}_2^{p-1}$ (the unit sphere of ℓ_2), $\Sigma_0 \succeq 0$, $\Sigma_0 \theta = 0$, and $\|\Sigma_0\|_2 = 1$ (the spectral norm of Σ_0). In model (3), the covariance matrix Σ has a unique largest eigenvalue λ_1 . Throughout this paper, for $q \in [0, 1]$, we consider the class

$$\mathcal{M}_q(\lambda_1, \lambda_2, \bar{R}_q, \alpha, \kappa)$$

that consists of all probability distributions on $X_i \in \mathbb{R}^p$, $i = 1, \dots, n$ satisfying model (3) with $\theta_1 \in \mathbb{B}_q^p(\bar{R}_q + 1)$, and Assumption 2.1 (below) with α and κ depending on q only.

Loss function The second ingredient in the minimax framework is the loss function. In the case of subspace estimation, an obvious criterion for evaluating the quality of an estimator $\hat{\Theta}$ is the squared distance between $\hat{\Theta}$ and Θ . However, it is not appropriate because Θ is not unique— Θ and ΘV span the same subspace for any $k \times k$ orthogonal matrix V . On the other hand, the orthogonal projections $\Pi = \Theta \Theta^T$ and $\hat{\Pi} = \hat{\Theta} \hat{\Theta}^T$ are unique. So we consider the loss function defined by the Frobenius norm of their difference:

$$\|\hat{\Pi} - \Pi\|_F.$$

In the case where $k = 1$, the only possible non-uniqueness in the leading eigenvector is its sign ambiguity. Still, we prefer to use the above loss function in the form

$$\|\hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T\|_F$$

because it generalizes to the case $k > 1$. Moreover, when $k = 1$, it turns out to be equivalent to both the Euclidean distance between $\theta_1, \hat{\theta}_1$ (when they belong to the same half-space) and the magnitude of the

sine of the angle between θ_1 , $\hat{\theta}_1$. (See Lemmas A.1.1 and A.1.2 in the Appendix.)

Scaling Our goal in this work is to provide non-asymptotic bounds on the minimax error

$$\min_{\hat{\theta}_1} \max_{P \in \mathcal{M}_q(\lambda_1, \lambda_2, \bar{R}_q, \alpha, \kappa)} \mathbb{E}_P \|\hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T\|_F,$$

where the minimum is taken over all estimators that depend only on X_1, \dots, X_n , that explicitly track the dependence of the minimax error on the vector $(p, n, \lambda_1, \lambda_2, \bar{R}_q)$. As we stated early, the classical p fixed, $n \rightarrow \infty$ scaling completely misses the effect of high-dimensionality; we, on the other hand, want to highlight the role that sparsity constraints play in high-dimensional estimation. Our lower bounds on the minimax error use an information theoretic technique based on Fano's Inequality. The upper bounds are obtained by constructing an ℓ_q -constrained estimator that nearly achieves the lower bound.

1.4 ℓ_q -Constrained Eigenvector Estimation

Consider the constrained maximization problem

$$\begin{aligned} & \text{maximize} && b^T S b, \\ & \text{subject to} && b \in \mathbb{S}_2^{p-1} \cap \mathbb{B}_q^p(\rho_q) \end{aligned} \quad (4)$$

and the estimator defined to be the solution of the optimization problem. The feasible set is non-empty when $\rho_q \geq 1$, and the ℓ_q constraint is active only when $\rho_q \leq p^{1-\frac{1}{q}}$. The ℓ_q -constrained estimator corresponds to ordinary PCA when $q = 2$ and $\rho_q = 1$. When $q \in [0, 1]$, the ℓ_q constraint promotes sparsity in the estimate. Since the criterion is a convex function of b , the convexity of the constraint set is inconsequential—it may be replaced by its convex hull without changing the optimum.

The case $q = 1$ is the most interesting from a practical point of view, because it corresponds to the well-known Lasso estimator for linear regression. In this case, eq. (4) coincides with the method proposed by Jolliffe, Trendafilov, and Uddin [13], though (4) remains a difficult *convex maximization* problem. Subsequent authors [21, 25] have proposed efficient algorithms that can approximately solve eq. (4). Our results below are (to our knowledge) the first convergence rate results available for this ℓ_1 -constrained PCA estimator.

1.5 Related Work

Amini and Wainwright [1] analyzed the performance of a semidefinite programming (SDP) formulation of sparse PCA for a generalized spiked covariance model [11]. Their model assumes that the nonzero entries of

the eigenvector all have the same magnitude, and that the covariance matrix corresponding to the nonzero entries is of the form $\beta \theta_1 \theta_1^T + I$. They derived upper and lower bounds on the success probability for model selection under the constraint that $\theta_1 \in \mathbb{B}_0^p(R_0)$. Their upper bound is conditional on the SDP based estimate being rank 1. Model selection accuracy and estimation accuracy are different notions of accuracy. One does not imply the other. In comparison, our results below apply to a wider class of covariance matrices and in the case of ℓ_0 we provide sharp bounds for the estimation error.

Operator norm consistent estimates of the covariance matrix automatically imply consistent estimates of eigenspaces. This follows from matrix perturbation theory [see, e.g., 22]. There has been much work on finding operator norm consistent covariance estimators in high-dimensions under assumptions on the sparsity or bandability of the entries of Σ or Σ^{-1} [see, e.g., 3, 2, 7]. Minimax results have been established in that setting by Cai, Zhang, and Zhou [5]. However, sparsity in the covariance matrix and sparsity in the leading eigenvector are different conditions. There is some overlap (e.g. the spiked covariance model), but in general, one does not imply the other.

Raskutti, Wainwright, and Yu [20] studied the related problem of minimax estimation for linear regression over ℓ_q balls. Remarkably, the rates that we derive for PCA are nearly identical to those for the Gaussian sequence model and regression. The work of Raskutti, Wainwright, and Yu [20] is close to ours in that they inspired us to use some similar techniques for the upper bounds.

While writing this paper we became aware of an unpublished manuscript by Paul and Johnstone [19]. They also study PCA under ℓ_q constraints with a slightly different but equivalent loss function. Their work provides asymptotic lower bounds for the minimax rate of convergence over ℓ_q balls for $q \in (0, 2]$. They also analyze the performance of an estimator based on a multistage thresholding procedure and show that asymptotically it nearly attains the optimal rate of convergence. Their analysis used spiked covariance matrices (corresponding to $\lambda_2 \Sigma_0 = (I_p - \theta_1 \theta_1^T)$ in eq. (3) when $k = 1$), while we allow a more general class of covariance matrices. We note that our work provides *non-asymptotic* bounds that are optimal over (p, n, \bar{R}_q) when $q \in \{0, 1\}$ and optimal over (p, n) when $q \in (0, 1)$.

In next section, we present our main results along with some additional conditions to guarantee that estimation over \mathcal{M}_q remains non-trivial. The main steps of the proofs are in Section 3. In the proofs we state

some auxiliary lemmas. They are mainly technical, so we defer their proofs to the Appendix. Section 4 concludes the paper with some comments on extensions of this work.

2 Main Results

Our minimax results are formulated in terms of non-asymptotic bounds that depend explicitly on $(n, p, R_q, \lambda_1, \lambda_2)$. To facilitate presentation, we introduce the notations

$$\bar{R}_q = R_q - 1 \text{ and } \sigma^2 = \frac{\lambda_1 \lambda_2}{(\lambda_1 - \lambda_2)^2}.$$

\bar{R}_q appears naturally in our lower bounds because the eigenvector θ_1 belongs to the sphere of dimension $p-1$ due to the constraint that $\|\theta_1\|_2 = 1$. Intuitively, σ^2 plays the role of the effective noise-to-signal ratio. When comparing with minimax results for linear regression over ℓ_q balls, σ^2 is exactly analogous to the noise variance in the linear model. Throughout the paper, there are absolute constants c, C, c_1, \dots that may take different values in different expressions.

The following assumption on R_q , the size of the ℓ_q ball, is to ensure that the eigenvector is not too dense.

Assumption 2.1. *There exists $\alpha \in (0, 1]$, depending only on q , such that*

$$\bar{R}_q \leq \kappa^q (p-1)^{1-\alpha} \bar{R}_q^{\frac{2\alpha}{2-q}} \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{\frac{2}{2-q}}} \right]^{\frac{q}{2}}, \quad (5)$$

where $\kappa \leq c\alpha/16$ is a constant depending only on q , and

$$1 \leq \bar{R}_q \leq e^{-1} (p-1)^{1-q/2}. \quad (6)$$

Assumption 2.1 also ensures that the effective noise σ^2 is not too small—this may happen if the spectral gap $\lambda_1 - \lambda_2$ is relatively large or if λ_2 is relatively close to 0. In either case, the distribution of $X_i/\lambda_1^{1/2}$ would concentrate on a 1-dimensional subspace and the problem would effectively degrade into a low-dimensional one. If R_q is relatively large, then $\mathbb{S}_2^{p-1} \cap \mathbb{B}_q^p(R_q)$ is not much smaller than \mathbb{S}_2^{p-1} and the parameter space will include many non-sparse vectors. In the case $q = 0$, Assumption 2.1 simplifies because we may take $\alpha = 1$ and only require that

$$1 \leq \bar{R}_0 \leq e^{-1} (p-1).$$

In the high-dimensional case that we are interested, where $p > n$, the condition that

$$1 \leq \bar{R}_q \leq e^{-1} \kappa^q \sigma^q p^{(1-\alpha)/2},$$

for some $\alpha' \in [0, 1]$, is sufficient to ensure that (5) holds for $q \in (0, 1]$. Alternatively, if we let $\alpha = 1 - q/2$ then (5) is satisfied for $q \in (0, 1]$ if

$$1 \leq \kappa^2 \sigma^2 ((p-1)/n) \log((p-1)/\bar{R}_q^{\frac{2}{2-q}}).$$

The relationship between n, p, R_q and σ^2 described in Assumption 2.1 indicates a regime in which the inference is neither impossible nor trivially easy. We can now state our first main result.

Theorem 2.1 (Lower Bound for Sparse PCA). *Let $q \in [0, 1]$. If Assumption 2.1 holds, then there exists a universal constant $c > 0$ depending only on q , such that every estimator $\hat{\theta}_1$ satisfies*

$$\begin{aligned} & \max_{P \in \mathcal{M}_q(\lambda_1, \lambda_2, \bar{R}_q, \alpha, \kappa)} \mathbb{E}_P \|\hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T\|_F \\ & \geq c \min \left\{ 1, \bar{R}_q^{\frac{1}{2}} \left[\frac{\sigma^2}{n} \log \left((p-1)/\bar{R}_q^{\frac{2}{2-q}} \right) \right]^{\frac{1}{2} - \frac{q}{4}} \right\}. \end{aligned}$$

Our proof of Theorem 2.1 is given in Section 3.1. It follows the usual nonparametric lower bound framework. The main challenge is to construct a rich packing set in $\mathbb{S}_2^{p-1} \cap \mathbb{B}_q^p(R_q)$. (See Lemma 3.1.2.) We note that a similar construction has been independently developed and applied in similar a context by Paul and Johnstone [19].

Our upper bound result is based on analyzing the solution to the ℓ_q -constrained maximization problem (4), which is a special case of empirical risk minimization. In order to bound the empirical process, we assume the data vector has sub-Gaussian tails, which is nicely described by the Orlicz ψ_α -norm.

Definition 2.1. *For a random variable $Y \in \mathbb{R}$, the Orlicz ψ_α -norm is defined for $\alpha \geq 1$ as*

$$\|Y\|_{\psi_\alpha} = \inf\{c > 0 : \mathbb{E} \exp(|Y|/c)^\alpha \leq 2\}.$$

Random variables with finite ψ_α -norm correspond to those whose tails are bounded by $\exp(-Cx^\alpha)$.

The case $\alpha = 2$ is important because it corresponds to random variables with sub-Gaussian tails. For example, if $Y \sim \mathcal{N}(0, \sigma^2)$ then $\|Y\|_{\psi_2} \leq C\sigma$ for some positive constant C . See [24, Chapter 2.2] for a complete introduction.

Assumption 2.2. *There exist i.i.d. random vectors $Z_1, \dots, Z_n \in \mathbb{R}^p$ such that $\mathbb{E} Z_i = 0$, $\mathbb{E} Z_i Z_i^T = I_p$,*

$$X_i = \mu + \Sigma^{1/2} Z_i \text{ and } \sup_{x \in \mathbb{S}_2^{p-1}} \|\langle Z_i, x \rangle\|_{\psi_2} \leq K,$$

where $\mu \in \mathbb{R}^p$ and $K > 0$ is a constant.

Assumption 2.2 holds for a variety of distributions, including the multivariate Gaussian (with $K^2 = 8/3$) and those of bounded random vectors. Under this assumption, we have the following theorem.

Theorem 2.2 (Upper Bound for Sparse PCA). *Let $\hat{\theta}_1$ be the ℓ_q constrained PCA estimate in eq. (4) with $\rho_q = R_q$, and let*

$$\epsilon = \|\hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T\|_F \text{ and } \tilde{\sigma} = \lambda_1 / (\lambda_1 - \lambda_2).$$

If the distribution of (X_1, \dots, X_n) belongs to $\mathcal{M}_q(\lambda_1, \lambda_2, \bar{R}_q, \alpha, \kappa)$ and satisfies Assumptions 2.1 and 2.2, then there exists a constant $c > 0$ depending only on K such that the following hold:

1. *If $q \in (0, 1)$, then*

$$\mathbb{E}\epsilon^2 \leq c \min \left\{ 1, R_q^2 \left[\frac{\tilde{\sigma}^2}{n} \log p \right]^{1-\frac{2}{q}} \right\}.$$

2. *If $q = 1$, then*

$$\mathbb{E}\epsilon^2 \leq c \min \left\{ 1, R_1 \left[\frac{\tilde{\sigma}^2}{n} \log(p/R_1^2) \right]^{\frac{1}{2}} \right\}.$$

3. *If $q = 0$, then*

$$[\mathbb{E}\epsilon]^2 \leq c \min \left\{ 1, R_0 \frac{\tilde{\sigma}^2}{n} \log(p/R_0) \right\}.$$

The proof of Theorem 2.2 is given in Section 3.2. The different bounds for $q = 0$, $q = 1$, and $q \in (0, 1)$ are due to the different tools available for controlling empirical processes in ℓ_q balls. Comparing with Theorem 2.1, when $q = 0$, the lower and upper bounds agree up to a factor $\sqrt{\lambda_2/\lambda_1}$. In the cases of $p = 1$ and $p \in (0, 1)$, a lower bound in the squared error can be obtained by using the fact $\mathbb{E}Y^2 \geq (\mathbb{E}Y)^2$. Therefore, over the class of distributions in $\mathcal{M}_q(\lambda_1, \lambda_2, \bar{R}_q, \alpha, \kappa)$ satisfying Assumptions 2.1 and 2.2, the upper and lower bound agree in terms of (p, n) for all $q \in (0, 1)$, and are sharp in (p, n, R_q) for $q \in \{0, 1\}$.

3 Proofs of Main Results

We use the following notation in the proofs. For matrices A and B whose dimensions are compatible, we define $\langle A, B \rangle = \text{Tr}(A^T B)$. Then the Frobenius norm is $\|A\|_F^2 = \langle A, A \rangle$. The Kullback-Leibler (KL) divergence between two probability measures $\mathbb{P}_1, \mathbb{P}_2$ is denoted by $D(\mathbb{P}_1 \|\| \mathbb{P}_2)$.

3.1 Proof of the Lower Bound (Theorem 2.1)

Our main tool for proving the minimax lower bound is the generalized Fano Method [9]. The following version is from [26, Lemma 3].

Lemma 3.1.1 (Generalized Fano method). *Let $N \geq 1$ be an integer and $\theta_1, \dots, \theta_N \subset \Theta$ index a collection of probability measures \mathbb{P}_{θ_i} on a measurable space $(\mathcal{X}, \mathcal{A})$. Let d be a pseudometric on Θ and suppose that for all $i \neq j$*

$$d(\theta_i, \theta_j) \geq \alpha_N$$

and

$$D(\mathbb{P}_{\theta_i} \|\| \mathbb{P}_{\theta_j}) \leq \beta_N.$$

Then every \mathcal{A} -measurable estimator $\hat{\theta}$ satisfies

$$\max_i \mathbb{E}_{\theta_i} d(\hat{\theta}, \theta_i) \geq \frac{\alpha_N}{2} \left(1 - \frac{\beta_N + \log 2}{\log N} \right).$$

The method works by converting the problem from estimation to testing by discretizing the parameter space, and then applying Fano's Inequality to the testing problem. (The β_N term that appears above is an upper bound on the mutual information.)

To be successful, we must find a sufficiently large finite subset of the parameter space such that the points in the subset are α_N -separated under the loss, yet nearly indistinguishable under the KL divergence of the corresponding probability measures. We will use the subset given by the following lemma.

Lemma 3.1.2 (Local packing set). *Let $\bar{R}_q = R_q - 1 \geq 1$ and $p \geq 5$. There exists a finite subset $\Theta_\epsilon \subset \mathbb{S}_2^{p-1} \cap \mathbb{B}_q^p(R_q)$ and an absolute constant $c > 0$ such that every distinct pair $\theta_1, \theta_2 \in \Theta_\epsilon$ satisfies*

$$\epsilon/\sqrt{2} < \|\theta_1 - \theta_2\|_2 \leq \sqrt{2}\epsilon,$$

and

$$\log|\Theta_\epsilon| \geq c \left(\frac{\bar{R}_q}{\epsilon^q} \right)^{\frac{2}{2-q}} \left[\log(p-1) - \log \left(\frac{\bar{R}_q}{\epsilon^q} \right)^{\frac{2}{2-q}} \right]$$

for all $q \in [0, 1]$ and $\epsilon \in (0, 1]$.

Fix $\epsilon \in (0, 1]$ and let Θ_ϵ denote the set given by Lemma 3.1.2. With Lemma A.1.2 we have

$$\epsilon^2/2 \leq \|\theta_1 \theta_1^T - \theta_2 \theta_2^T\|_F^2 \leq 4\epsilon^2 \quad (7)$$

for all distinct pairs $\theta_1, \theta_2 \in \Theta_\epsilon$. For each $\theta \in \Theta_\epsilon$, let

$$\Sigma_\theta = (\lambda_1 - \lambda_2)\theta\theta^T + \lambda_2 I_p.$$

Clearly, Σ_θ has eigenvalues $\lambda_1 > \lambda_2 = \dots = \lambda_p$. Then Σ_θ satisfies eq. (3). Let \mathbb{P}_θ denote the n -fold product of the $\mathcal{N}(0, \Sigma_\theta)$ probability measure. We use the following lemma to help bound the KL divergence.

Lemma 3.1.3. For $i = 1, 2$, let $x_i \in \mathbb{S}_2^{p-1}$, $\lambda_1 > \lambda_2 > 0$,

$$\Sigma_i = (\lambda_1 - \lambda_2)x_i x_i^T + \lambda_2 I_p,$$

and \mathbb{P}_i be the n -fold product of the $\mathcal{N}(0, \Sigma_i)$ probability measure. Then

$$D(\mathbb{P}_1 \| \mathbb{P}_2) = \frac{n}{2\sigma^2} \|x_1 x_1^T - x_2 x_2^T\|_F^2,$$

where $\sigma^2 = \lambda_1 \lambda_2 / (\lambda_1 - \lambda_2)^2$.

Applying this lemma with eq. (7) gives

$$D(\mathbb{P}_{\theta_1} \| \mathbb{P}_{\theta_2}) = \frac{n}{2\sigma^2} \|\theta_1 \theta_1^T - \theta_2 \theta_2^T\|_F^2 \leq \frac{2n\epsilon^2}{\sigma^2}.$$

Thus, we have found a subset of the parameter space that conforms to the requirements of Lemma 3.1.1, and so

$$\max_{\theta \in \Theta_\epsilon} \mathbb{E}_\theta \|\hat{\theta} \hat{\theta}^T - \theta \theta^T\|_F \geq \frac{\epsilon}{2\sqrt{2}} \left(1 - \frac{2n\epsilon^2/\sigma^2 + \log 2}{\log|\Theta_\epsilon|}\right)$$

for all $\epsilon \in (0, 1]$. The final step is to choose ϵ of the correct order. If we can find ϵ so that

$$\frac{2n\epsilon^2/\sigma^2}{\log|\Theta_\epsilon|} \leq \frac{1}{4} \quad (8)$$

and

$$\log|\Theta_\epsilon| \geq 4 \log 2, \quad (9)$$

then we may conclude that

$$\max_{\theta \in \Theta_\epsilon} \mathbb{E}_\theta \|\hat{\theta} \hat{\theta}^T - \theta \theta^T\|_F \geq \frac{\epsilon}{4\sqrt{2}}.$$

For a constant $C \in (0, 1)$ to be chosen later, let

$$\epsilon^2 = \min \left\{ 1, C^{2-q} \bar{R}_q \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{2-q}} \right]^{1-\frac{q}{2}} \right\}. \quad (10)$$

We consider each of the two cases in the above $\min\{\dots\}$ separately.

Case 1: Suppose that

$$1 \leq C^{2-q} \bar{R}_q \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{2-q}} \right]^{1-\frac{q}{2}}. \quad (11)$$

Then $\epsilon^2 = 1$ and by rearranging (11)

$$\frac{n}{C^2 \sigma^2} \leq \bar{R}_q^{2-q} \left[\log(p-1) - \log \bar{R}_q^{2-q} \right].$$

So by Lemma 3.1.2,

$$\log|\Theta_\epsilon| \geq c \bar{R}_q^{2-q} \left[\log(p-1) - \log \bar{R}_q^{2-q} \right] \geq \frac{cn}{C^2 \sigma^2}.$$

If we choose $C^2 \leq c/16$, then

$$\frac{2n\epsilon^2/\sigma^2}{\log|\Theta_\epsilon|} \leq \frac{4C^2}{c} \leq \frac{1}{4}.$$

To lower bound

$$\log|\Theta_\epsilon| \geq c \bar{R}_q^{2-q} \left[\log(p-1) - \log \bar{R}_q^{2-q} \right],$$

observe that the function $x \mapsto x \log[(p-1)/x]$ is increasing on $[1, (p-1)/e]$, and, by Assumption 2.1, this interval contains $\bar{R}_q^{2/(2-q)}$. If p is large enough so that $p-1 \geq \exp\{(4/c) \log 2\}$, then

$$\log|\Theta_\epsilon| \geq c \log(p-1) \geq 4 \log 2.$$

Thus, eqs. (8) and (9) are satisfied, and we conclude that

$$\max_{\theta \in \Theta_\epsilon} \mathbb{E}_\theta \|\hat{\theta} \hat{\theta}^T - \theta \theta^T\|_F \geq \frac{\epsilon}{4\sqrt{2}},$$

as long as $C^2 \leq c/16$ and $p-1 \geq \exp\{(4/c) \log 2\}$.

Case 2: Now let us suppose that

$$1 > C^{2-q} \bar{R}_q \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{2-q}} \right]^{1-\frac{q}{2}}. \quad (12)$$

Then

$$\left(\frac{\bar{R}_q}{\epsilon^q} \right)^{\frac{2}{2-q}} = \frac{\bar{R}_q}{C^q} \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{2-q}} \right]^{-\frac{q}{2}}, \quad (13)$$

and it is straightforward to check that Assumption 2.1 implies that if $C^q \geq \kappa^q$, then there is $\alpha \in (0, 1]$, depending only on q , such that

$$\left(\frac{1}{\epsilon^q} \right)^{\frac{2}{2-q}} \leq \left(\frac{p-1}{\bar{R}_q^{2-q}} \right)^{1-\alpha}. \quad (14)$$

So by Lemma 3.1.2,

$$\begin{aligned} \log|\Theta_\epsilon| &\geq c \left(\frac{\bar{R}_q}{\epsilon^q} \right)^{\frac{2}{2-q}} \left[\log \frac{p-1}{\bar{R}_q^{2-q}} - \log \left(\frac{1}{\epsilon^q} \right)^{\frac{2}{2-q}} \right] \\ &\geq c \alpha \frac{\bar{R}_q}{C^q} \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{2-q}} \right]^{-\frac{q}{2}} \left[\log \frac{p-1}{\bar{R}_q^{2-q}} \right], \end{aligned} \quad (15)$$

where the last inequality is obtained by plugging in (13) and (14).

If we choose $C^2 \leq c\alpha/16$, then combining (10) and (15), we have

$$\frac{4n\epsilon^2}{\sigma^2} \leq \frac{4C^2}{c\alpha} \leq \frac{1}{4} \quad (16)$$

and eq. (8) is satisfied. On the other hand, by (12) and the fact that $\bar{R}_q \geq 1$, we have

$$C^{-q} \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{\frac{2}{2-q}}} \right]^{-\frac{q}{2}} \geq 1,$$

and hence (15) becomes

$$\log|\Theta_\epsilon| \geq c\alpha \bar{R}_q \left[\log \frac{p-1}{\bar{R}_q^{\frac{2}{2-q}}} \right]. \quad (17)$$

The function $x \mapsto x \log[(p-1)/x^{2/(2-q)}]$ is increasing on $[1, (p-1)^{1-q/2}/e]$ and, by Assumption 2.1, $1 \leq \bar{R}_q \leq (p-1)^{1-q/2}/e$. If $p-1 \geq \exp\{[4/(\alpha)] \log 2\}$, then

$$\log|\Theta_\epsilon| \geq c\alpha \log(p-1) \geq 4 \log 2$$

and eq. (9) is satisfied. So we can conclude that

$$\max_{\theta \in \Theta_\epsilon} \mathbb{E}_\theta \|\hat{\theta}\hat{\theta}^T - \theta\theta\|_F \geq \frac{\epsilon}{4\sqrt{2}},$$

as long as $C^2 \leq c\alpha/16$ and $p-1 \geq \exp\{[4/(\alpha)] \log 2\}$.

Cases 1 and 2 together: Looking back at cases 1 and 2, we see that because $\alpha \leq 1$, the conditions that $\kappa^2 \leq C^2 \leq c\alpha/16$ and $p-1 \geq \exp\{[4/(\alpha)] \log 2\}$ are sufficient to ensure that

$$\begin{aligned} & \max_{\theta \in \Theta_\epsilon} \mathbb{E}_\theta \|\hat{\theta}\hat{\theta}^T - \theta\theta\|_F \\ & \geq c' \min \left\{ 1, \bar{R}_q^{\frac{1}{2}} \left[\frac{\sigma^2}{n} \log \frac{p-1}{\bar{R}_q^{\frac{2}{2-q}}} \right]^{\frac{1}{2}-\frac{q}{4}} \right\}, \end{aligned}$$

for a constant $c' > 0$ depending only on q . \blacksquare

3.2 Proof of the Upper Bound (Theorem 2.2)

We begin with a lemma that bounds the curvature of the matrix functional $\langle \Sigma, bb^T \rangle$.

Lemma 3.2.1. *Let $\theta \in \mathbb{S}_2^{p-1}$. If $\Sigma \succeq 0$ has a unique largest eigenvalue λ_1 with corresponding eigenvector θ_1 , then*

$$\frac{1}{2}(\lambda_1 - \lambda_2) \|\theta\theta^T - \theta_1\theta_1^T\|_F^2 \leq \langle \Sigma, \theta_1\theta_1^T - \theta\theta^T \rangle.$$

Now consider $\hat{\theta}_1$, the ℓ_q -constrained sparse PCA estimator of θ_1 . Let $\epsilon = \|\hat{\theta}_1\hat{\theta}_1^T - \theta_1\theta_1^T\|_F$. Since $\theta_1 \in \mathbb{S}_2^{p-1}$, it follows from Lemma 3.2.1 that

$$\begin{aligned} (\lambda_1 - \lambda_2)\epsilon^2/2 & \leq \langle \Sigma, \theta_1\theta_1^T - \hat{\theta}_1\hat{\theta}_1^T \rangle \\ & = \langle S, \theta_1\theta_1^T \rangle - \langle \Sigma, \hat{\theta}_1\hat{\theta}_1^T \rangle - \langle S - \Sigma, \theta_1\theta_1^T \rangle \\ & \leq \langle S - \Sigma, \hat{\theta}_1\hat{\theta}_1^T \rangle - \langle S - \Sigma, \theta_1\theta_1^T \rangle \\ & = \langle S - \Sigma, \hat{\theta}_1\hat{\theta}_1^T - \theta_1\theta_1^T \rangle. \end{aligned} \quad (18)$$

We consider the cases $q \in (0, 1)$, $q = 1$, and $q = 0$ separately.

3.2.1 Case 1: $q \in (0, 1)$

By applying Hölder's Inequality to the right side of eq. (18) and rearranging, we have

$$\epsilon^2/2 \leq \frac{\|\text{vec}(S - \Sigma)\|_\infty}{\lambda_1 - \lambda_2} \|\text{vec}(\theta_1\theta_1^T - \hat{\theta}_1\hat{\theta}_1^T)\|_1, \quad (19)$$

where $\text{vec}(A)$ denotes the $1 \times p^2$ matrix obtained by stacking the columns of a $p \times p$ matrix A . Since θ_1 and $\hat{\theta}_1$ both belong to $\mathbb{B}_q^p(R_q)$,

$$\begin{aligned} \|\text{vec}(\theta_1\theta_1^T - \hat{\theta}_1\hat{\theta}_1^T)\|_q^q & \leq \|\text{vec}(\theta_1\theta_1^T)\|_q^q + \|\text{vec}(\hat{\theta}_1\hat{\theta}_1^T)\|_q^q \\ & \leq 2R_q^2. \end{aligned}$$

Let $t > 0$. We can use a standard truncation argument [see, e.g., 20, Lemma 5] to show that

$$\begin{aligned} & \|\text{vec}(\theta_1\theta_1^T - \hat{\theta}_1\hat{\theta}_1^T)\|_1 \\ & \leq \sqrt{2}R_q \|\text{vec}(\theta_1\theta_1^T - \hat{\theta}_1\hat{\theta}_1^T)\|_2 t^{-q/2} + 2R_q^2 t^{1-q} \\ & = \sqrt{2}R_q \|\theta_1\theta_1^T - \hat{\theta}_1\hat{\theta}_1^T\|_F t^{-q/2} + 2R_q^2 t^{1-q} \\ & = \sqrt{2}R_q \epsilon t^{-q/2} + 2R_q^2 t^{1-q}. \end{aligned}$$

Letting $t = \|\text{vec}(S - \Sigma)\|_\infty / (\lambda_1 - \lambda_2)$ and joining with eq. (19) gives us

$$\epsilon^2/2 \leq \sqrt{2} t^{1-q/2} R_q \epsilon + 2t^{2-q} R_q^2.$$

If we define m implicitly so that $\epsilon = m\sqrt{2}t^{1-q/2}R_q$, then the preceding inequality reduces to $m^2/2 \leq m+1$. If $m \geq 3$, then this is violated. So we must have $m < 3$ and hence

$$\epsilon \leq 3\sqrt{2}t^{1-q/2}R_q = 3\sqrt{2}R_q \left(\frac{\|\text{vec}(S - \Sigma)\|_\infty}{\lambda_1 - \lambda_2} \right)^{1-q/2}. \quad (20)$$

Combining the above discussion with the sub-Gaussian assumption, the next lemma allows us to bound $\|\text{vec}(S - \Sigma)\|_\infty$.

Lemma 3.2.2. *If Assumption 2.2 holds and Σ satisfies (2), then there is an absolute constant $c > 0$ such that*

$$\|\|\text{vec}(S - \Sigma)\|_\infty\|_{\psi_1} \leq cK^2 \lambda_1 \max \left\{ \sqrt{\frac{\log p}{n}}, \frac{\log p}{n} \right\}.$$

Applying Lemma 3.2.2 to eq. (20) gives

$$\begin{aligned} & \|\epsilon^{2/(2-q)}\|_{\psi_1} \\ & \leq cR_q^{2/(2-q)} \frac{\|\|\text{vec}(S - \Sigma)\|_\infty\|_{\psi_1}}{\lambda_1 - \lambda_2} \\ & \leq cK^2 R_q^{2/(2-q)} \tilde{\sigma} \max \left\{ \sqrt{\frac{\log p}{n}}, \frac{\log p}{n} \right\}. \end{aligned}$$

The fact that $\mathbb{E}|X|^m \leq (m!)^m \|X\|_{\psi_1}^m$ for $m \geq 1$ [see 24, Chapter 2.2] implies the following bound:

$$\mathbb{E}\epsilon^2 \leq c_K R_q^2 \tilde{\sigma}^{2-q} \max \left\{ \sqrt{\frac{\log p}{n}}, \frac{\log p}{n} \right\}^{2-q} =: \mathbf{M},$$

Combining this with the trivial bound $\epsilon \leq 2$, yields

$$\mathbb{E}\epsilon^2 \leq \min(2, \mathbf{M}). \quad (21)$$

If $\log p > n$, then $\mathbb{E}\epsilon^2 \leq 2$. Otherwise, we need only consider the square root term inside $\max\{\cdot\}$ in the definition of \mathbf{M} . Thus,

$$\mathbb{E}\epsilon^2 \leq c \min \left\{ 1, R_q^2 \left[\frac{\tilde{\sigma}^2}{n} \log p \right]^{1-\frac{q}{2}} \right\}.$$

for an appropriate constant $c > 0$, depending only on K . This completes the proof for the case $q \in (0, 1)$.

3.2.2 Case 2: $q = 1$

θ_1 and $\hat{\theta}_1$ both belong to $\mathbb{B}_1^p(R_1)$. So applying the triangle inequality to the right side of eq. (18) yields

$$\begin{aligned} (\lambda_1 - \lambda_2)\epsilon^2/2 &\leq \langle S - \Sigma, \hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T \rangle \\ &\leq |\hat{\theta}_1^T (S - \Sigma) \hat{\theta}_1| + |\theta_1^T (S - \Sigma) \theta_1| \\ &\leq 2 \sup_{b \in \mathbb{S}_2^{p-1} \cap \mathbb{B}_1^p(R_1)} |b^T (S - \Sigma) b|. \end{aligned}$$

The next lemma provides a bound for the supremum.

Lemma 3.2.3. *If Assumption 2.2 holds and Σ satisfies (2), then there is an absolute constant $c > 0$ such that*

$$\begin{aligned} \mathbb{E} \sup_{b \in \mathbb{S}_2^{p-1} \cap \mathbb{B}_1^p(R_1)} |b^T (S - \Sigma) b| \\ \leq c \lambda_1 K^2 \max \left\{ R_1 \sqrt{\frac{\log(p/R_1^2)}{n}}, R_1^2 \frac{\log(p/R_1^2)}{n} \right\} \end{aligned}$$

for all $R_1^2 \in [1, p/e]$.

Assumption 2.1 guarantees that $R_1^2 \in [1, p/e]$. Thus, we can apply Lemma 3.2.3 and an argument similar to that used with (21) to complete the proof for the case $q = 1$.

3.2.3 Case 3: $q = 0$

We continue from eq. (18). Since $\hat{\theta}_1$ and θ_1 belong to $\mathbb{B}_0^p(R_0)$, their difference belongs to $\mathbb{B}_0^p(2R_0)$. Let Π denote the diagonal matrix whose diagonal entries are 1 wherever $\hat{\theta}_1$ or θ_1 are nonzero, and 0 elsewhere. Then Π has at most $2R_0$ nonzero diagonal entries, and

$\Pi \hat{\theta}_1 = \hat{\theta}_1$ and $\Pi \theta_1 = \theta_1$. So by the Von Neumann trace inequality and Lemma A.1.1,

$$\begin{aligned} (\lambda_1 - \lambda_2)\epsilon^2/2 &\leq |\langle S - \Sigma, \Pi(\hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T) \Pi \rangle| \\ &= |\langle \Pi(S - \Sigma) \Pi, \hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T \rangle| \\ &\leq \|\Pi(S - \Sigma) \Pi\|_2 \|\hat{\theta}_1 \hat{\theta}_1^T - \theta_1 \theta_1^T\|_{S_1} \\ &= \|\Pi(S - \Sigma) \Pi\|_2 \sqrt{2}\epsilon \\ &\leq \sup_{b \in \mathbb{S}_2^{p-1} \cap \mathbb{B}_0^p(2R_0)} |b^T (S - \Sigma) b| \sqrt{2}\epsilon, \end{aligned}$$

where $\|\cdot\|_{S_1}$ denotes the sum of the singular values. Divide both sides by ϵ , rearrange terms, and then take the expectation to get

$$\mathbb{E}\epsilon \leq \frac{c}{\lambda_1 - \lambda_2} \mathbb{E} \sup_{b \in \mathbb{S}_2^{p-1} \cap \mathbb{B}_0^p(2R_0)} |b^T (S - \Sigma) b|.$$

Lemma 3.2.4. *If Assumption 2.2 holds and Σ satisfies (2), then there is an absolute constant $c > 0$ such that*

$$\begin{aligned} \mathbb{E} \sup_{b \in \mathbb{S}_2^{p-1} \cap \mathbb{B}_0^p(d)} |b^T (S - \Sigma) b| \\ \leq c K^2 \lambda_1 \max \left\{ \sqrt{(d/n) \log(p/d)}, (d/n) \log(p/d) \right\} \end{aligned}$$

for all integers $d \in [1, p/2]$.

Taking $d = 2R_0$ and applying an argument similar to that used with (21) completes the proof of the $q = 0$ case. \blacksquare

4 Conclusion and Further Extensions

We have presented upper and lower bounds on the minimax estimation error for sparse PCA over ℓ_q balls. The bounds are sharp in (p, n) , and they show that ℓ_q constraints on the leading eigenvector make estimation possible in high-dimensions even when the number of variables greatly exceeds the sample size. Although we have specialized to the case $k = 1$ (for the leading eigenvector), our methods and arguments can be extended to the multi-dimensional subspace case ($k > 1$). One nuance in that case is that there are different ways to generalize the notion of ℓ_q sparsity to multiple eigenvectors. A potential difficulty there is that if there is multiplicity in the eigenvalues or if eigenvalues coalesce, then the eigenvectors need not be unique (up to sign). So care must be taken to handle this possibility.

Acknowledgements

V. Q. Vu was supported by a NSF Mathematical Sciences Postdoctoral Fellowship (DMS-0903120). J. Lei was supported by NSF Grant BCS0941518. We thank the anonymous reviewers for their helpful comments.

References

- [1] A. A. Amini and M. J. Wainwright. “High-dimensional analysis of semidefinite relaxations for sparse principal components”. In: *Annals of Statistics* 37.5B (2009), pp. 2877–2921.
- [2] P. J. Bickel and E. Levina. “Covariance Regularization by Thresholding”. In: *Annals of Statistics* 36.6 (2008), pp. 2577–2604.
- [3] P. J. Bickel and E. Levina. “Regularized estimation of large covariance matrices”. In: *Annals of Statistics* 36.1 (2008), pp. 199–227.
- [4] L. Birgé and P. Massart. “Minimum contrast estimators on sieves: exponential bounds and rates of convergence”. In: *Bernoulli* 4.3 (1998), pp. 329–375.
- [5] T. T. Cai, C.-H. Zhang, and H. H. Zhou. “Optimal rates of convergence for covariance matrix estimation”. In: *Annals of Statistics* 38.4 (2010), pp. 2118–2144.
- [6] A. d’Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. G. Lanckriet. “A direct formulation for sparse PCA using semidefinite programming”. In: *SIAM Review* 49.3 (2007), pp. 434–448.
- [7] N. El Karoui. “Operator norm consistent estimation of large-dimensional sparse covariance matrices”. In: *Annals of Statistics* 36.6 (2008), pp. 2717–2756.
- [8] Y. Gordon, A. Litvak, S. Mendelson, and A. Pajor. “Gaussian averages of interpolated bodies and applications to approximate reconstruction”. In: *Journal of Approximation Theory* 149 (2007), pp. 59–73.
- [9] T. S. Han and S. Verdú. “Generalizing the Fano inequality”. In: *IEEE Transactions on Information Theory* 40 (1994), pp. 1247–1251.
- [10] A. J. Izenman. *Modern Multivariate Statistical Techniques: Regression, Classification, and Manifold Learning*. New York, NY: Springer, 2008.
- [11] I. M. Johnstone. “On the distribution of the largest eigenvalue in principal components analysis”. In: *Annals of Statistics* 29.2 (2001), pp. 295–327.
- [12] I. M. Johnstone and A. Y. Lu. “On consistency and sparsity for principal components analysis in high dimensions”. In: *Journal of the American Statistical Association* 104.486 (2009), pp. 682–693.
- [13] I. T. Jolliffe, N. T. Trendafilov, and M. Uddin. “A modified principal component technique based on the Lasso”. In: *Journal of Computational and Graphical Statistics* 12 (2003), pp. 531–547.
- [14] L. Mackey. “Deflation methods for sparse PCA”. In: *Advances in Neural Information Processing Systems 21*. Ed. by D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou. 2009, pp. 1017–1024.
- [15] P. Massart. *Concentration Inequalities and Model Selection*. Ed. by J. Picard. Springer-Verlag, 2007.
- [16] S. Mendelson. “Empirical processes with a bounded ψ_1 diameter”. In: *Geometric and Functional Analysis* 20.4 (2010), pp. 988–1027.
- [17] B. Nadler. “Finite sample approximation results for principal component analysis: A matrix perturbation approach”. In: *Annals of Statistics* 36.6 (2008), pp. 2791–2817.
- [18] D. Paul. “Asymptotics of sample eigenstructure for a large dimensional spiked covariance model”. In: *Statistica Sinica* 17 (2007), pp. 1617–1642.
- [19] D. Paul and I. Johnstone. “Augmented sparse principal component analysis for high dimensional Data”. manuscript. 2007.
- [20] G. Raskutti, M. J. Wainwright, and B. Yu. “Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls”. In: *IEEE Transactions on Information Theory* (2011). to appear.
- [21] H. Shen and J. Z. Huang. “Sparse principal component analysis via regularized low rank matrix approximation”. In: *Journal of Multivariate Analysis* 99 (2008), pp. 1015–1034.
- [22] G. W. Stewart and J. Sun. *Matrix Perturbation Theory*. Academic Press, 1990.
- [23] M. Talagrand. *The Generic Chaining*. Springer-Verlag, 2005.
- [24] A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag, 1996.
- [25] D. M. Witten, R. Tibshirani, and T. Hastie. “A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis”. In: *Biostatistics* 10 (2009), pp. 515–534.
- [26] B. Yu. “Assouad, Fano, and Le Cam”. In: *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*. 1997, pp. 423–435.
- [27] H. Zou, T. Hastie, and R. Tibshirani. “Sparse principal component analysis”. In: *Journal of Computational and Graphical Statistics* 15.2 (2006), pp. 265–286.