

OPTIMAL DETECTION OF SPARSE PRINCIPAL COMPONENTS IN HIGH DIMENSION

BY QUENTIN BERTHET¹ AND PHILIPPE RIGOLLET²

Princeton University

We perform a finite sample analysis of the detection levels for sparse principal components of a high-dimensional covariance matrix. Our minimax optimal test is based on a sparse eigenvalue statistic. Alas, computing this test is known to be NP-complete in general, and we describe a computationally efficient alternative test using convex relaxations. Our relaxation is also proved to detect sparse principal components at near optimal detection levels, and it performs well on simulated datasets. Moreover, using polynomial time reductions from theoretical computer science, we bring significant evidence that our results cannot be improved, thus revealing an inherent trade off between statistical and computational performance.

1. Introduction. The sparsity assumption has become preponderant in modern, high-dimensional statistics. In the high dimension, low sample size setting, where consistency seems to be hopeless, sparsity turns out to be the statistician's salvation. It formalizes the a priori belief that only a few parameters, among a large number of them, are significant for the statistical task at hand. This paper explores a specific high-dimensional problem, namely Principal Component Analysis (PCA). Indeed, without further assumptions, classical PCA is known to produce inconsistent estimators of the directions that explain the most variance [Johnstone and Lu (2009), Nadler (2008), Paul (2007)]. For PCA, the *spiked covariance model* introduced by Johnstone (2001) provides a natural setting for statistical problems. Namely, this model relies on the assumption that there exists a small number of directions that explain most of the variance. In this work, we assume that observations are drawn from a multivariate Gaussian distribution with mean zero and covariance matrix given by $I + \theta vv^\top$, where I is the identity matrix, v is a unit norm sparse vector and $\theta > 0$. Akin to other models, the sparsity assumption drives both methods and analysis in a wide variety of applications ranging from signal processing to biology; see Alon et al. (1999), Chen (2011), Jenatton, Obozinski and Bach (2010), Wright et al. (2011) for a few examples. Most contributions to this problem have focused on consistent estimation of the sparse principal component

Received December 2012; revised April 2013.

¹Supported in part by a Gordon S. Wu fellowship.

²Supported in part by NSF Grants DMS-09-06424 and CAREER-DMS-1053987.

MSC2010 subject classifications. Primary 62H25; secondary 62F04, 90C22.

Key words and phrases. High-dimensional detection, sparse principal component analysis, spiked covariance model, semidefinite relaxation, minimax lower bounds, planted clique.

v under various performance measures; see, for example, [Amini and Wainwright \(2009\)](#), [Birbaum et al. \(2013\)](#), [Cai, Ma and Wu \(2012\)](#), [Ma \(2013\)](#), [Shen, Shen and Marron \(2013\)](#), [Vu and Lei \(2012\)](#) and the above references.

What if there is no sparse component? In other words, what if $\theta = 0$? From a detection standpoint, one may ask the following question: How much variance should a sparse principal component explain in order to be detectable by a statistical procedure? Answering this question consists of (i) constructing a test that can detect this sparse principal component when the associated variance is above a certain level and (ii) proving that no test can detect such a principal component below a certain level.

Optimal detection levels in a high-dimensional setup have recently received a lot of attention. [Arias-Castro, Candès and Durand \(2011\)](#), [Arias-Castro, Candès and Plan \(2011\)](#), [Donoho and Jin \(2004\)](#), [Ingster, Tsybakov and Verzelen \(2010\)](#) have studied the detection of a sparse vector corrupted by noise under various sparsity assumptions. More recently, this problem has been extended from vectors to matrices by [Butucea and Ingster \(2013\)](#), [Sun and Nobel \(2008, 2013\)](#) who propose to detect a shifted sub-matrix planted in a Gaussian or binary random matrix. While the notion of sub-matrix encodes a certain sparsity structure, these two papers focus on the elementwise properties of random matrices, unlike the blooming random matrix theory that focuses on spectral aspects. [Arias-Castro, Bubeck and Lugosi \(2012\)](#) studied a problem related to sparse PCA detection, but closer to the shifted sub-matrix problem. Their goal is to detect a shifted off-diagonal sub-matrix planted in a covariance matrix. Their methods are not spectral either.

We extend the current work on detection in two directions. First, we analyze detection in the framework of sparse PCA, and more precisely, in the spiked covariance model. Second, we derive a finite sample analysis of minimax optimality in this problem with results that hold with high-probability, unlike most of the literature on detection where an asymptotic framework is usually preferred. A notable exception is the paper of [Addario-Berry et al. \(2010\)](#) where results are of the same flavor as ours. Unlike the asymptotic analysis pioneered by [Ingster \(1982\)](#) and recently extended to sparse linear regression in [Donoho and Jin \(2004\)](#), [Ingster, Tsybakov and Verzelen \(2010\)](#), this finite sample analysis is not refined enough to exhibit a qualitative difference between testing and estimation. Nevertheless, such results shed light on the delicate interplay between the important parameters of the problem: ambient dimension, sample size and sparsity.

The minimax optimal test statistic for our testing problem relies on the so-called k -sparse largest eigenvalue of the empirical covariance matrix. It captures the largest amount of empirical variance explained by any k of the original variables. It turns out that although this statistic can be used to construct an optimal test, it raises computational difficulties and can even be proved to be NP-complete in general. As a result, a large body of the optimization literature on this topic consists of numerical methods to overcome this issue; see, for example, [d'Aspremont, Bach and El Ghaoui \(2008\)](#), [d'Aspremont et al. \(2007\)](#), [Journée et al. \(2010\)](#), [Lu](#)

and Zhang (2012), Ma (2011) and references therein. Nevertheless, while these numerical methods do produce a solution, their statistical properties are rarely addressed for the estimation problem and never for the detection problem. One of the approaches introduced by d'Aspremont et al. (2007) uses a convexification technique called *semidefinite programming* (SDP). A major drawback of this technique is that it may not output a sparse direction \hat{v} . Indeed, semidefinite programs output matrices that are not rank-one in general, and an ad hoc post-processing step is often required to turn this matrix back into a unit vector. However, in the context of detection, our goal is not to estimate the eigenvector v but rather its associated eigenvalue. This notable difference allows us to bypass SDP optimization altogether, which is known to scale poorly in high dimension. Inspired by the dual SDP formulation, we propose a simple test procedure based on the *minimum dual perturbation* (MDP) that is easy to compute and for which we can derive near optimal performance bounds for the detection problem. More importantly, we bring supporting evidence to the tightness of the performance bounds that we prove. Interestingly, this evidence builds upon a conjecture from theoretical computer science. Indeed, a reduction to the *planted clique* problem shows that a better performance would contradict a widely believed conjecture on the average-case complexity of this problem.

Most of our analysis is performed in the model of sparse rank one perturbation for the covariance matrix of Gaussian random vectors. Nevertheless, our results are robust to variations around this model, and we devote Section 7 to discussing various weaker assumptions under which our results still hold. In particular, our results are more generally valid for sub-Gaussian observations and weaker notions of sparsity. We also study the case where the distance between the estimated and true covariance matrices is only controlled in sup-norm, with high probability. This setup encompasses biased estimators or adversarial noise.

The rest of the paper is organized as follows. In Section 2, we introduce the detection problem for sparse PCA. In Section 3, we discuss various links with probabilistic results on random matrix theory and more precisely, the asymptotic effect of a principal component on the spectrum of a Wishart matrix. Minimax detection levels are derived in Section 4, where in particular, we introduce a test based on spectral methods and derive the level at which it achieves detection of sparse principal components with high probability. This level is proved to be optimal in a minimax sense in Section 5. Unfortunately, this test cannot be computed efficiently, and several relaxations are proposed in Section 6. For these convex methods, we derive suboptimal levels that also hold under various weaker assumptions for which they sometimes become optimal (Section 7). Moreover, using arguments from computational complexity, we argue in Section 8 that even under the strongest assumptions of this paper, these suboptimal levels are likely to be the best achievable by the efficient relaxations. Specifically, we show that proving better bounds for these methods would lead to a contradiction of the hidden clique

conjecture, which is widely believed to be true. The numerical performance of our test and in particular its suboptimality, is illustrated in Section 9.

NOTATION. The space of $d \times d$ symmetric real matrices is denoted by \mathbf{S}_d . We write $Z \succeq 0$ whenever Z is semidefinite positive.

The elements of a vector $v \in \mathbf{R}^d$ are denoted by v_1, \dots, v_d and similarly, a matrix Z has element Z_{ij} on its i th row and j th column. For any $q > 0$, $|v|_q$ denotes the ℓ_q “norm” of a vector v and is defined by $|v|_q = (\sum_j |v_j|^q)^{1/q}$. Moreover, we denote by $|v|_0$ its so-called ℓ_0 “norm,” that is, its number of nonzero elements. Furthermore, by extension, for $Z \in \mathbf{S}_d$, we denote by $|Z|_q$ the ℓ_q norm of the vector formed by the entries of Z . We also define for $q \in [0, 2)$ the set $\mathcal{B}_q(R)$ of unit vectors within the ℓ_q -ball of radius $R > 0$

$$\mathcal{B}_q(R) = \{v \in \mathbf{R}^p : |v|_2 = 1, |v|_q \leq R\}.$$

The trace and rank functionals are denoted by \mathbf{Tr} and \mathbf{rank} , respectively, and have their usual definition. The identity matrix in \mathbf{R}^d is denoted by I_d . For a finite set S , we denote by $|S|$ its cardinality. We also write A_S for the $|S| \times |S|$ submatrix with elements $(A_{ij})_{i,j \in S}$, and v_S for the vector of $\mathbf{R}^{|S|}$ with elements v_i for $i \in S$. Finally, for two real numbers a and b , we write $a \wedge b = \min(a, b)$, $a \vee b = \max(a, b)$ and $a_+ = a \vee 0$.

2. Statement of the hypothesis testing problem. Let X_1, \dots, X_n be n i.i.d. copies of a random variable X in \mathbf{R}^p . Our objective is to perform the following test:

$$\begin{aligned} H_0 : X &\sim \mathcal{N}(0, I_p), \\ H_1 : X &\sim \mathcal{N}(0, I_p + \theta vv^\top), \quad v \in \mathcal{B}_0(k), \end{aligned}$$

where $\theta > 0$, and we remind the reader that $\mathcal{B}_0(k)$ is the set of k -sparse unit vectors. Note that the model under H_1 is an adaptation of the spiked covariance model since it only allows v to be k -sparse on the unit Euclidean sphere. This is precisely the model of sparse PCA introduced in [Johnstone and Lu \(2009\)](#). In particular, the distribution of X under H_1 is invariant under rotation of the k relevant variables. We use this simplified model for reasons of clarity: to highlight the importance of relative variance, only one direction v is used for signal, and only one parameter θ is used to express the signal-to-noise ratio. Note that our upper and lower bounds for optimal testing are valid for the general hypotheses

$$\begin{aligned} H_0 : X &\sim \mathcal{N}(0, \Sigma_0), \quad \lambda_{\max}^k(\Sigma_0) \leq 1, \\ H_1 : X &\sim \mathcal{N}(0, \Sigma_1), \quad \lambda_{\max}^k(\Sigma_1) \geq 1 + \theta, \end{aligned}$$

where λ_{\max}^k , is the k -sparse eigenvalue defined in (4.1) below. In particular, the model under H_1 encompasses that of [Amini and Wainwright \(2009\)](#).

Let $\Sigma = \mathbb{E}[XX^\top]$ denote the covariance matrix of the centered random vector X , and denote by $\hat{\Sigma}$ the empirical covariance matrix defined by

$$(2.1) \quad \hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n X_i X_i^\top.$$

We say that a test *discriminates* between H_0 and H_1 with probability $1 - \delta$ if both type I and type II errors have a probability smaller than δ . Our goal is therefore to find a statistic $\varphi(\hat{\Sigma})$ and quantiles $\tau_0 < \tau_1$, depending on (p, n, k, δ) such that

$$\mathbf{P}_{H_0}(\varphi(\hat{\Sigma}) > \tau_0) \leq \delta, \quad \mathbf{P}_{H_1}(\varphi(\hat{\Sigma}) < \tau_1) \leq \delta.$$

For $\tau \in [\tau_0, \tau_1]$ define the test

$$\psi(\hat{\Sigma}) = \mathbf{1}\{\varphi(\hat{\Sigma}) > \tau\},$$

where $\mathbf{1}\{\cdot\}$ denotes the indicator function. As desired, this test discriminates between the two hypotheses with probability $1 - \delta$. We assume that the user is testing for a specific sparsity k . Nevertheless, using a Bonferroni correction, this test can be performed for various values of k if needed.

Note that throughout the paper, we assume that all of the parameters (k, n, p) are known so that τ_0 and τ_1 are easily determined.

3. Link with random matrix theory.

3.1. *Spectral methods.* It is not hard to see that, under H_1 , for any $\theta > 0$, v is an eigenvector associated to the largest eigenvalue of the population covariance matrix Σ . Moreover, if $\hat{\Sigma}$ is close to Σ in spectral norm, then its largest eigenvector should be a good candidate to approximate v . It is therefore natural to consider spectral methods for the spiked covariance model. Understanding the behavior of our test statistic under both the null and the alternative is key in proving that it discriminates between the hypotheses.

Spectral convergence of the empirical covariance matrix to the true covariance matrix has received some attention recently [see, e.g., [Bickel and Levina \(2008\)](#), [Cai, Zhang and Zhou \(2010\)](#), [El Karoui \(2008\)](#)] under various elementwise sparsity assumptions and using thresholding methods. However, since our assumption allows for relevant variables to produce arbitrarily small entries under the alternative hypothesis, we cannot use such results. A natural statistic to discriminate between the null and the alternative would be, for example, the largest eigenvalue of the covariance matrix.

Spectral properties of random matrices have received a lot of attention from both a statistical and probabilistic perspective. We devote the rest of this section to reviewing some of the classical results from random matrix theory, and we argue that even in moderate dimension, the largest eigenvalue cannot discriminate between the null and alternative hypotheses.

It is easily seen that for any unit vector v ,

$$(3.1) \quad \lambda_{\max}(I_p) = 1 \quad \text{and} \quad \lambda_{\max}(I_p + \theta vv^\top) = 1 + \theta.$$

If we could allow, for a fixed p , to let n go to infinity, the consistency of the estimator $\hat{\Sigma}$ (for fixed p , entry by entry) and the continuity of the largest eigenvalue as a function of the entries of a matrix would imply that the largest eigenvalue can be used to discriminate between the two alternatives, at least asymptotically.

However, in a high-dimension setting, where p is typically much larger than n , the behavior of $\lambda_{\max}(\hat{\Sigma})$ under the null hypothesis is quite different. If $p/n \rightarrow \alpha > 0$, [Geman \(1980\)](#) showed that, in accordance with the Marcenko–Pastur distribution, we have

$$\lambda_{\max}(\hat{\Sigma}) \rightarrow (1 + \sqrt{\alpha})^2 > 1,$$

where the convergence holds almost surely see also [Bai \(1999\)](#), [Johnstone \(2001\)](#) and references therein. Moreover, [Yin, Bai and Krishnaiah \(1988\)](#) established that finite fourth moment is a necessary and sufficient condition for this almost sure convergence to hold. Furthermore, since $\hat{\Sigma} \geq 0$, its number of positive eigenvalues is equal to its rank (which is at most n), and we have

$$\lambda_{\max}(\hat{\Sigma}) \geq \frac{1}{\text{rank}(\hat{\Sigma})} \sum_{i=1}^p \lambda_i(\hat{\Sigma}) \geq \frac{1}{n} \text{Tr}(\hat{\Sigma}) = \frac{p}{n} \frac{\sum_{i=1}^n |X_i|_2^2}{np}.$$

Note that under H_0 , it holds $\sum_{i=1}^n |X_i|_2^2 \sim \chi_{np}^2$. Hence almost surely, for $p/n \rightarrow \infty$, we have $\lambda_{\max}(\hat{\Sigma}) \rightarrow \infty$.

These two results hint at an intrinsic limitation of the largest eigenvalue statistic: its fluctuations are too large to discriminate between the two hypotheses in a “large p /small n ” scenario unless the signal strength θ is very strong.

In the next subsection, we show that the above argument can be made formal using spectral results in random matrix theory.

3.2. Finite rank perturbations of covariance matrices. In a moderate-dimensional regime, where $p/n \rightarrow \alpha \in (0, 1)$, [Baik, Ben Arous and P ech e \(2005\)](#) describe a phase transition for the spectral behavior of the sample covariance matrix $\hat{\Sigma}$ of complex Gaussian vectors between two different regimes. This phenomenon is now widely known as the BBP transition for the name of the authors. The same phenomenon, for real random variables, was subsequently established in [Baik and Silverstein \(2006\)](#).

Qualitatively, there exists a critical value θ^* such that if $\theta > \theta^*$, the spectrum of $\hat{\Sigma}$ exhibits an isolated eigenvalue significantly larger than the others, and such that if $\theta < \theta^*$, the spectrum has a similar behavior under the two hypotheses. More precisely, Theorem 1.1 of [Baik and Silverstein \(2006\)](#) implies that under H_1 , the largest eigenvalue will either exhibit an important concentration around a deterministic value strictly larger than $1 + \theta$ if the perturbation is strong enough, or

around the upper edge of the Marcenko–Pastur distribution, as if the perturbation was nonexistent, when it is too weak. The critical level is $\theta^* = \sqrt{\alpha}$, and suggests a minimum signal level of order $\sqrt{p/n}$ which is high already when p is of the order of n .

These results are even proved to hold for weakened assumptions on the distribution of the vectors, in Féral and Pécché (2009). On the statistical side, these are coherent with the detection levels shown in Onatski, Moreira and Hallin (2013) for testing of the sphericity hypothesis with no assumption on the alternative, by spectral methods.

4. Sparse principal component detection. In sparse principal component detection, we are testing the existence of a sparse direction v with a significantly higher explained variance $v^\top \Sigma v$ than any other direction. To exploit the sparsity assumption, we use the fact that only a small submatrix of the covariance is affected by the perturbation. Let A be a $p \times p$ matrix and fix $k < p$. We define the k -sparse largest³ eigenvalue by

$$(4.1) \quad \lambda_{\max}^k(A) = \max_{|S|=k} \lambda_{\max}(A_S).$$

It can be defined equivalently to (4.1) by

$$(4.2) \quad \lambda_{\max}^k(A) = \max_{x \in \mathcal{B}_0(k)} x^\top Ax.$$

Therefore, we study the behavior of the test statistic $\varphi(\hat{\Sigma}) = \lambda_{\max}^k(\hat{\Sigma})$ under both hypotheses.

4.1. *Deviation bounds for the k -sparse eigenvalue.* Optimal detection levels are governed by the deviations of the test statistic $\lambda_{\max}^k(\hat{\Sigma})$ both under the null and the alternative hypotheses. We begin with the following proposition, which guarantees that our test statistic remains large enough under the alternative hypothesis.

PROPOSITION 4.1. *Under H_1 , we have with probability $1 - \delta$,*

$$\lambda_{\max}^k(\hat{\Sigma}) \geq 1 + \theta - 2(1 + \theta) \sqrt{\frac{\log(1/\delta)}{n}}.$$

PROOF. Under H_1 , there exists a unit vector v with sparsity k , such that $X \sim \mathcal{N}(0, I_p + \theta vv^\top)$. Therefore, we have

$$\lambda_{\max}^k(\hat{\Sigma}) \geq v^\top \hat{\Sigma} v = \frac{1}{n} \sum_{i=1}^n (X_i^\top v)^2$$

³In the rest of the paper, we drop the qualification “largest” since we only refer to this one.

by definition of $\hat{\Sigma}$. Since $X \sim \mathcal{N}(0, I_p + \theta vv^\top)$, we have $X^\top v \sim \mathcal{N}(0, 1 + \theta)$.

Define the random variable

$$Y = \frac{1}{n} \sum_{i=1}^n \left(\frac{(X_i^\top v)^2}{1 + \theta} - 1 \right).$$

Using Lemma A.1, we get for any $t > 0$, that

$$\mathbf{P}(Y \leq -2\sqrt{t/n}) \leq e^{-t}.$$

Hence, taking $t = \log(1/\delta)$ yields the desired inequality. \square

Note that our proof relies only on the existence of a sparse vector v associated to the eigenvalue $(1 + \theta)$ of the population covariance matrix Σ . In particular, the result of Proposition 4.1 extends to more general alternative hypotheses, as long as they satisfy this condition.

Note that much more than detection can actually be achieved under this model. Indeed, Amini and Wainwright (2009) prove optimal rates of support recovery when θ is known and large enough, and for v taking only values in $\{0, \pm 1/\sqrt{k}\}$.

We now study the behavior of the k -sparse eigenvalue under the null hypothesis, that is, for a Wishart matrix with mean I_p . We adapt a technique from Vershynin (2012) to obtain the desired deviation bounds.

PROPOSITION 4.2. *Under H_0 , with probability $1 - \delta$*

$$\lambda_{\max}^k(\hat{\Sigma}) \leq 1 + 4\sqrt{\frac{k \log(9ep/k) + \log(1/\delta)}{n}} + 4\frac{k \log(9ep/k) + \log(1/\delta)}{n}.$$

PROOF. Using a $1/4$ -net over the unit sphere of \mathbf{R}^k , it can be easily shown [see, e.g., Vershynin (2012)] that there exists a subset \mathcal{N}_k of the unit sphere of \mathbf{R}^k , with cardinality smaller than 9^k , such that for any $A \geq 0$

$$(4.3) \quad \lambda_{\max}(A) \leq 2 \max_{x \in \mathcal{N}_k} x^\top Ax.$$

Under H_0 , since $\hat{\Sigma}$ is positive semidefinite, we have

$$\lambda_{\max}^k(\hat{\Sigma}) = 1 + \max_{|S|=k} \{\lambda_{\max}(\hat{\Sigma}_S) - 1\}.$$

For all $u \in \mathbf{R}^k$, $|u|_2 = 1$ and $S \subset \{1, \dots, p\}$ such that $|S| = k$, let $\tilde{u} \in \mathbf{R}^p$ be the vector with support in S such that $\tilde{u}_S = u$. We have

$$u^\top \hat{\Sigma}_S u - 1 = \tilde{u}^\top \hat{\Sigma} \tilde{u} - 1 = \frac{1}{n} \sum_{i=1}^n [(\tilde{u}^\top X_i)^2 - 1].$$

Since $|\tilde{u}|_2 = |u|_2 = 1$, Lemma A.1 yields for any $t > 0$,

$$(4.4) \quad \mathbf{P}\left(\frac{1}{n} \sum_{i=1}^n [(\tilde{u}^\top X_i)^2 - 1] \geq 2\sqrt{\frac{t}{n}} + 2\frac{t}{n}\right) \leq e^{-t}.$$

For any $S \subset \{1, \dots, p\}$, define \mathbf{R}^S to be the subset of \mathbf{R}^p defined such that $x \in \mathbf{R}^S$ if and only if $x_j = 0, \forall j \notin S$. Let $\mathcal{N}_k(S)$ be a subset of the unit sphere of \mathbf{R}^S , with cardinality smaller than 9^k such that for any $A \geq 0$, inequality (4.3) holds with $\mathcal{N}_k = \mathcal{N}_k(S)$. Fix $t > 0$ and define the event \mathcal{A}_S by

$$\mathcal{A}_S = \left\{ \lambda_{\max}(\hat{\Sigma}_S) - 1 \geq 4\sqrt{\frac{t}{n}} + 4\frac{t}{n} \right\}.$$

Observe that a union bound over the elements of $\mathcal{N}_k(S)$ together with (4.4) yields that for any $t > 0$,

$$\mathbf{P}(\mathcal{A}_S) \leq \mathbf{P}\left(\max_{v \in \mathcal{N}_k(S)} \frac{1}{n} \sum_{i=1}^n (v^\top X_i)^2 - 1 \geq 2\sqrt{\frac{t}{n}} + 2\frac{t}{n}\right) \leq 9^k e^{-t}.$$

Let now \mathcal{A} be the event defined by

$$\mathcal{A} = \bigcup_{|S|=k} \mathcal{A}_S = \left\{ \max_{|S|=k} \{ \lambda_{\max}(\hat{\Sigma}_S) - 1 \} \geq 4\sqrt{\frac{t}{n}} + 4\frac{t}{n} \right\}.$$

Therefore, by a union bound on the $\binom{p}{k}$ subsets S of $\{1, \dots, p\}$ that have cardinality k , we get

$$\mathbf{P}\left(\lambda_{\max}^k(\hat{\Sigma}) \geq 1 + 4\sqrt{\frac{t}{n}} + 4\frac{t}{n}\right) = \mathbf{P}(\mathcal{A}) \leq \binom{p}{k} 9^k e^{-t}.$$

To complete our proof, it is sufficient to use the standard inequality $\binom{p}{k} \leq \left(\frac{ep}{k}\right)^k$ and to take $t = k \log(9ep/k) + \log(1/\delta)$. \square

4.2. *Hypothesis testing with λ_{\max}^k .* Using these results, we have, with the notation from Section 2,

$$\mathbf{P}_{H_0}(\lambda_{\max}^k(\hat{\Sigma}) > \tau_0) \leq \delta, \quad \mathbf{P}_{H_1}(\lambda_{\max}^k(\hat{\Sigma}) < \tau_1) \leq \delta,$$

where τ_0 and τ_1 are given by

$$\tau_0 = 1 + 4\sqrt{\frac{k \log(9ep/k) + \log(1/\delta)}{n}} + 4\frac{k \log(9ep/k) + \log(1/\delta)}{n},$$

$$\tau_1 = 1 + \theta - 2(1 + \theta)\sqrt{\frac{\log(1/\delta)}{n}}.$$

Whenever $\tau_1 > \tau_0$, we take $\tau \in [\tau_0, \tau_1]$ and define the following test:

$$\psi(\hat{\Sigma}) = \mathbf{1}\{\lambda_{\max}^k(\hat{\Sigma}) > \tau\}.$$

It follows from the previous subsection that the test discriminates between H_1 and H_0 with probability $1 - \delta$. It remains to find for which values of θ the condition $\tau_1 > \tau_0$ holds. It corresponds to our minimum detection level.

THEOREM 4.1. *Assume that k, p, n and δ are such that $\bar{\theta} \leq 1$, where*

$$(4.5) \quad \bar{\theta} := 4\sqrt{\frac{k \log(9ep/k) + \log(1/\delta)}{n}} + 4\frac{k \log(9ep/k) + \log(1/\delta)}{n} + 4\sqrt{\frac{\log(1/\delta)}{n}}.$$

Then, for any $\theta > \bar{\theta}$ and for any $\tau \in [\tau_0, \tau_1]$, the test $\psi(\hat{\Sigma}) = \mathbf{1}\{\lambda_{\max}^k(\hat{\Sigma}) > \tau\}$ discriminates between H_0 and H_1 with probability $1 - \delta$.

If we consider high-dimensional asymptotic regimes, for large p, n, k , taking $\delta = p^{-\beta}$ with $\beta > 0$, provides a sequence of tests ψ_n that discriminate between H_0 and H_1 with probability converging to 1, for any fixed $\theta > 0$, as soon as $k \log(p)/n \rightarrow 0$.

5. Minimax lower bounds for detection. The goal of this section is to prove that for any $\nu > 0$, there exists $\underline{\theta}_\nu$ (5.1) of the same order as $\bar{\theta}$ (up to logarithmic terms), and such that if $\theta < \underline{\theta}_\nu$, then no test can discriminate between H_0 and H_1 with probability greater than $\frac{1}{2} + \nu$. Recall that \mathbf{P}^n denotes the joint distribution of n i.i.d. random variables with distribution \mathbf{P} .

THEOREM 5.1. *Fix $\nu > 0$. There exists a constant $C_\nu > 0$ defined in (5.5) such that if*

$$(5.1) \quad \theta < \underline{\theta}_\nu := \sqrt{\frac{k \log(C_\nu p/k^2 + 1)}{n}} \wedge \frac{1}{\sqrt{2}},$$

it holds

$$\inf_{\psi} \left\{ \mathbf{P}_0^n(\psi = 1) \vee \max_{v \in \mathcal{B}_0(k)} \mathbf{P}_v^n(\psi = 0) \right\} \geq \frac{1}{2} - \nu,$$

where the infimum is taken over all possible tests, that is, measurable functions of the n observations, that take values in $\{0, 1\}$.

In order to find lower bounds for the probability of error, we study the χ^2 distance between probability measures; see, for example, [Tsybakov \(2009\)](#), Chapter 2. For any $v \in \mathbf{R}^p$ such that $|v|_2 = 1$, define the matrix $\Sigma_v = I_p + \theta vv^\top$, and

let \mathbf{P}_v denote the distribution of a Gaussian random variable $X \sim \mathcal{N}(0, \Sigma_v)$. Moreover, let $\mathcal{S} = \{S \subset \{1, \dots, p\} : |S| = k\}$, and for any $S \in \mathcal{S}$, define $u(S) \in \mathbf{R}^p$ to be the unit vector with j th coordinate equal to $1/\sqrt{k}$ if $j \in S$ and 0 otherwise. Finally, define the Gaussian mixture $\mathbf{P}_\mathcal{S}$ by

$$\mathbf{P}_\mathcal{S} = \frac{1}{|\mathcal{S}|} \sum_{S \in \mathcal{S}} \mathbf{P}_{u(S)}.$$

We write for simplicity $\mathbf{P}_S = \mathbf{P}_{u(S)}$ when this leads to no confusion. Our proof relies on the following lemma.

LEMMA 5.1. *For any $S, T \in \mathcal{S}$ and any $\theta < 1$, it holds*

$$\mathbb{E}_{\mathbf{P}_0} \left(\frac{d\mathbf{P}_S}{d\mathbf{P}_0} \frac{d\mathbf{P}_T}{d\mathbf{P}_0} \right) = (1 - \theta^2 (u(S)^\top u(T))^2)^{-1/2}.$$

PROOF. Fix $S \in \mathcal{S}$, and observe that

$$\frac{d\mathbf{P}_S}{d\mathbf{P}_0}(X) = \frac{\det(I_p)^{1/2} \exp(-X^\top \Sigma_{u(S)}^{-1} X/2)}{\det(\Sigma_{u(S)})^{1/2} \exp(-X^\top I_p^{-1} X/2)}.$$

Furthermore, since $\det(I_p) = 1$ and $|u(S)|_2 = 1$, we get by Sylvester’s determinant theorem that

$$\det(\Sigma_{u(S)}) = \det(I_p + \theta u(S)u(S)^\top) = \det(I_1 + \theta u(S)^\top u(S)) = 1 + \theta.$$

Moreover, the Sherman–Morrison formula yields

$$\Sigma_{u(S)}^{-1} = (I_p + \theta u(S)u(S)^\top)^{-1} = I_p - \frac{\theta u(S)u(S)^\top}{1 + \theta}.$$

By substitution, the above three displays yield

$$\frac{d\mathbf{P}_S}{d\mathbf{P}_0}(X) = \frac{1}{\sqrt{1 + \theta}} \exp\left(\frac{1}{2} \frac{\theta}{1 + \theta} (X^\top u(S))^2\right)$$

and

$$(5.2) \quad \frac{d\mathbf{P}_S}{d\mathbf{P}_0} \frac{d\mathbf{P}_T}{d\mathbf{P}_0}(X) = \frac{1}{1 + \theta} \exp(X^\top M X),$$

where M is defined by

$$M := \frac{1}{2} \frac{\theta}{1 + \theta} (u(S)u(S)^\top + u(T)u(T)^\top).$$

Note that M has at most two nonzero eigenvalues given by

$$\lambda_1 = \frac{1}{2} \frac{\theta}{1 + \theta} (1 + u(S)^\top u(T)) < \frac{1}{2} \quad \text{and} \quad \lambda_2 = \frac{1}{2} \frac{\theta}{1 + \theta} (1 - u(S)^\top u(T)) < \frac{1}{2},$$

and let Λ denote the diagonal matrix with elements $(\lambda_1, \lambda_2, 0, \dots, 0) \in \mathbf{R}^p$. Together with (5.2), it yields

$$\begin{aligned} \mathbb{E}_{\mathbf{P}_0} \left(\frac{d\mathbf{P}_S}{d\mathbf{P}_0} \frac{d\mathbf{P}_T}{d\mathbf{P}_0} \right) &= \frac{1}{1 + \theta} \mathbb{E}_{\mathbf{P}_0} [\exp(X^\top M X)] \\ &= \frac{1}{1 + \theta} \mathbb{E}_{\mathbf{P}_0} [\exp(X^\top \Lambda X)] \\ &= \frac{1}{1 + \theta} \mathbb{E}_{\mathbf{P}_0} [\exp(\lambda_1 X_1^2)] \mathbb{E}_{\mathbf{P}_0} [\exp(\lambda_2 X_2^2)] \\ &= \frac{1}{1 + \theta} [(1 - 2\lambda_1)(1 - 2\lambda_2)]^{-1/2}, \end{aligned}$$

where, in the second equality, the substitution of M by Λ is valid by rotational invariance of the distribution of X under \mathbf{P}_0 . The last equation yields the desired result. \square

We now turn to the proof of Theorem 5.1.

PROOF OF THEOREM 5.1. Observe now that

$$\chi^2(\mathbf{P}_S, \mathbf{P}_0) = \mathbb{E}_{\mathbf{P}_0} \left[\left(\frac{d\mathbf{P}_S}{d\mathbf{P}_0} - 1 \right)^2 \right] = \frac{1}{|S|^2} \sum_{S, T \in \mathcal{S}} \mathbb{E}_{\mathbf{P}_0} \left(\frac{d\mathbf{P}_S}{d\mathbf{P}_0} \frac{d\mathbf{P}_T}{d\mathbf{P}_0} \right) - 1.$$

Lemma 5.1 together with the fact $u(S)^\top u(T) = |S \cap T|/k$ yield

$$\chi^2(\mathbf{P}_S, \mathbf{P}_0) = \sum_{r=0}^k \left\{ \frac{\mathcal{C}(\mathcal{S}, r)}{|S|^2} \left(1 - \theta^2 \frac{r^2}{k^2} \right)^{-1/2} \right\} - 1,$$

where $\mathcal{C}(\mathcal{S}, r)$ denotes the number of subsets $S, T \in \mathcal{S}$ such that $|S \cap T| = r$. Let S, T be chosen uniformly at random in \mathcal{S} , and observe that $\mathbf{P}(|S \cap T| = r) = \mathbf{P}(R = r)$, where $R = |S \cap \{1, \dots, k\}|$. Jensen’s inequality yields

$$\begin{aligned} \chi^2(\mathbf{P}_S^n, \mathbf{P}_0^n) &= \prod_{i=1}^n (1 + \chi^2(\mathbf{P}_S, \mathbf{P}_0)) - 1 \\ &\leq \mathbb{E}_{S, T} \left\{ \left[1 - \theta^2 \frac{|S \cap T|^2}{k^2} \right]^{-n/2} \right\} - 1 \\ &= \mathbb{E}_R \left\{ \left[1 - \theta^2 \frac{R^2}{k^2} \right]^{-n/2} \right\} - 1, \end{aligned}$$

where $\mathbb{E}_{S, T}$ denotes the expectation with respect to the random subsets S, T and \mathbb{E}_R the expectation with respect to R .

Using now the convexity inequality $(1 - t)^{-n/2} \leq e^{nt/(2(1-t))} \leq e^{nt}$ valid for $1 - t \geq 1/2$, and noticing that $R \leq k$, the above display leads to

$$(5.3) \quad \chi^2(\mathbf{P}_S^n, \mathbf{P}_0^n) \leq \mathbb{E}_R \left[\exp\left(\frac{n\theta^2 R}{k}\right) \right] - 1.$$

Define $\mu^2 = n\theta^2/k$. We have, as in [Addario-Berry et al. \(2010\)](#), [Arias-Castro, Bubeck and Lugosi \(2012\)](#), that

$$\begin{aligned} \mathbb{E}_R[e^{\mu^2 R}] &= \mathbb{E}_S \left[\prod_{i=1}^k \exp(\mu^2 \mathbf{1}\{i \in S\}) \right] \\ &\leq \prod_{i=1}^k \mathbb{E}_S[\exp(\mu^2 \mathbf{1}\{i \in S\})] \leq \left((e^{\mu^2} - 1) \frac{k}{p} + 1 \right)^k. \end{aligned}$$

The first inequality holds by the negative association [see, e.g., [Addario-Berry et al. \(2010\)](#), Section 3] of negatively correlated dependent random variables. Assume now that $\theta < \theta_v$. It yields

$$\left((e^{\mu^2} - 1) \frac{k}{p} + 1 \right)^k \leq \left(\left(\frac{C_v p}{k^2} \right) \frac{k}{p} + 1 \right)^k \leq \left(1 + \frac{C_v}{k} \right)^k \leq e^{C_v}.$$

Together with (5.3), the previous two displays yield

$$(5.4) \quad \chi^2(\mathbf{P}_S^n, \mathbf{P}_0^n) \leq e^{C_v} - 1.$$

We are now in a position to apply standard results from minimax theory. Define

$$(5.5) \quad C_v := \log \left[(1 + 8v^2) \wedge \log \left(\frac{e}{2 - 4v} \right) \right],$$

and note that for all measurable tests ψ , we have

$$\begin{aligned} \mathbf{P}_0^n(\psi = 1) \vee \max_{v \in \mathcal{B}_0(k)} \mathbf{P}_v^n(\psi = 0) &\geq \mathbf{P}_0^n(\psi = 1) \vee \max_{S \in \mathcal{S}} \mathbf{P}_{u(S)}^n(\psi = 0) \\ &\geq \mathbf{P}_0^n(\psi = 1) \vee \mathbf{P}_S^n(\psi = 0) \\ &\geq \frac{e^{1-e^{C_v}}}{4} \vee \frac{1 - \sqrt{(e^{C_v} - 1)/2}}{2} = \frac{1}{2} - v, \end{aligned}$$

where the last inequality is a direct consequence of (5.4) and [Tsybakov \(2009\)](#), Theorem 2.2, case (iii). \square

We observe a gap between our upper and lower bounds, with a term in $\log(p/k)$ in the upper bound, and one in $\log(p/k^2)$ in the lower bound. This gap has been observed in the detection literature before [see, e.g., [Baraud \(2002\)](#), [Verzelen \(2012\)](#), for an explicit remark] and, to our knowledge, has never been addressed. However, if $p \geq k^{2+\varepsilon}$, $\varepsilon > 0$, upper and lower bounds match up to constants, and the detection rate for the sparse eigenvalue is optimal in a minimax sense. Under this assumption, detection becomes impossible if $\theta < C \sqrt{(k/n) \log(p/k)}$ for a small enough constant $C > 0$.

6. Efficient methods for sparse principal component testing. Computing the largest k -sparse eigenvalue λ_{\max}^k of a symmetric matrix A is, in general, a hard computational problem. To see this, consider the particular case where A is a $p \times p$ symmetric matrix with values in $\{0, 1\}$ and $A_{ii} = 1$ for all diagonal entries, so that A corresponds to the adjacency matrix of an undirected graph. It is not hard to see that $\lambda_{\max}^k(A) \leq k$, with equality if and only if the graph of A contains a clique of size k . It is a well-known fact of computational complexity [Karp (1972)] that the decision problem associated to finding whether a graph contains a clique of size k is NP-complete.

6.1. *Semidefinite relaxation for λ_{\max}^k .* Semidefinite programming (SDP) is the matrix equivalent of linear programming. Define the scalar product in \mathbf{S}_d by $\langle A, B \rangle = \mathbf{Tr}(AB)$. A semidefinite program can be written in the canonical form:

$$(6.1) \quad \begin{aligned} \text{SDP} = \max. \quad & \mathbf{Tr}(CX) \\ \text{subject to} \quad & \mathbf{Tr}(A_i X) \leq b_i \quad \forall i \in \{1, \dots, m\}, \\ & X \geq 0. \end{aligned}$$

As convex problems, they are computationally efficient and can be solved using interior point or first order methods; see, for example, Boyd and Vandenberghe (2004), Nesterov and Nemirovskii (1987). Using SDP relaxations of problems with nonconvex constraints such as integer programs is a common method to find approximate solutions. Approximation bounds, up to a constant, can sometimes be proved as in the celebrated result of Goemans and Williamson (1995) for the MAXCUT problem. A major breakthrough for sparse PCA was achieved by d'Aspremont et al. (2007), who introduced a SDP relaxation for λ_{\max}^k , but tightness of this relaxation is, to this day, unknown. Our task is not as difficult though. Indeed, we only need to prove that the SDP objective criterion has significantly different behavior under H_0 and H_1 .

Making the change of variables $Z = xx^\top$ in (4.2) yields

$$\begin{aligned} \lambda_{\max}^k(A) = \max. \quad & \mathbf{Tr}(AZ) \\ \text{subject to} \quad & \mathbf{Tr}(Z) = 1, |Z|_0 \leq k^2, \\ & Z \geq 0, \mathbf{rank}(Z) = 1. \end{aligned}$$

Note that this problem contains two sources of nonconvexity: the ℓ_0 norm constraint and the rank constraint. We make two relaxations in order to have a convex feasible set. First, for a semidefinite matrix Z , with trace 1, and sparsity k^2 , the Cauchy–Schwarz inequality yields $|Z|_1 \leq k$, which is substituted to the cardinality constraint in this relaxation. Simply dropping the rank constraint leads to the

following relaxation of our original problem:

$$(6.2) \quad \begin{aligned} \text{SDP}_k(A) = \max. \quad & \text{Tr}(AZ) \\ \text{subject to} \quad & \text{Tr}(Z) = 1, |Z|_1 \leq k, \\ & Z \geq 0. \end{aligned}$$

Note that this optimization problem is convex since it consists in minimizing a linear objective over a convex set. Moreover, it is a standard exercise to show that it can be expressed in the canonical form (6.1). As such, it can be solved efficiently using any of the aforementioned algorithms. This natural relaxation was originally developed in d’Aspremont et al. (2007). Note that building on an earlier version of this paper, d’Aspremont, Bach and Ghaoui (2012) proposed a new SDP relaxation to the same problem and derive somewhat larger detection levels, at least for the interesting case where k is small compared to p .

Let us now study the behavior of the objective value $\text{SDP}_k(\hat{\Sigma})$ under H_1 and H_0 , respectively. First, as a relaxation of the original problem, for any $A \geq 0$, it holds

$$(6.3) \quad \lambda_{\max}^k(A) \leq \text{SDP}_k(A).$$

Since we have proved in Section 4 that $\lambda_{\max}^k(\hat{\Sigma})$ takes large values under H_1 , this inequality tells us that so does $\text{SDP}_k(\hat{\Sigma})$. It remains to show that it stays small under H_0 . This can be achieved by using the dual formulation of the SDP.

LEMMA 6.1 [Bach, Ahipasaoglu and d’Aspremont (2010)]. *For a given $A \geq 0$, we have by duality*

$$\text{SDP}_k(A) = \min_{U \in \mathbf{S}_p} \{ \lambda_{\max}(A + U) + k|U|_{\infty} \}.$$

Together with (6.3), Lemma 6.1 implies that for any $z \geq 0$ and any matrix $U \in \mathbf{S}_p$ such that $|U|_{\infty} \leq z$, it holds

$$(6.4) \quad \lambda_{\max}^k(A) \leq \text{SDP}_k(A) \leq \lambda_{\max}(A + U) + kz.$$

A direct consequence of (6.4) is that the functional λ_{\max}^k is robust to small perturbations in $|\cdot|_{\infty}$ -norm. Let $A \geq 0$ be such that its largest eigenvector is k sparse. Then, for any matrix N , (6.4) yields

$$\lambda_{\max}^k(A + N) \leq \lambda_{\max}((A + N) - N) + k|N|_{\infty} = \lambda_{\max}^k(A) + k|N|_{\infty}.$$

6.2. *High probability bounds for convex relaxation.* We now study the properties of $\text{SDP}_k(\hat{\Sigma})$ and other computationally efficient variants as test statistics for our detection problem. In view of (6.3), the following proposition follows directly from Proposition 4.1.

PROPOSITION 6.1. Under H_1 , we have, with probability $1 - \delta$

$$\text{SDP}_k(\hat{\Sigma}) \geq 1 + \theta - 2(1 + \theta)\sqrt{\frac{\log(1/\delta)}{n}}.$$

We now turn to the upper deviations under H_0 .

PROPOSITION 6.2. Under H_0 , we have, with probability $1 - \delta$,

$$\begin{aligned} \text{SDP}_k(\hat{\Sigma}) \leq & 1 + 2\sqrt{\frac{k^2 \log(4p^2/\delta)}{n}} + 2\frac{k \log(4p^2/\delta)}{n} \\ & + 2\sqrt{\frac{\log(2p/\delta)}{n}} + 2\frac{\log(2p/\delta)}{n}. \end{aligned}$$

PROOF. Let $st_z(A)$ be the soft-threshold of A , with threshold z , defined by $(st_z(A))_{ij} = \text{sign}(A_{ij})(|A_{ij}| - z)_+$. It follows from (6.4) that for any $A \geq 0$,

$$(6.5) \quad \text{SDP}_k(A) \leq \lambda_{\max}(st_z(A)) + kz.$$

Let $\hat{\Delta} = \text{diag}(\hat{\Sigma})$ be the diagonal matrix with the same diagonal entries as $\hat{\Sigma}$, and $\hat{\Psi} = \hat{\Sigma} - \hat{\Delta}$ the matrix of its off-diagonal entries, so that $\hat{\Sigma} = \hat{\Delta} + \hat{\Psi}$. Since $\hat{\Psi}$ and $\hat{\Delta}$ have disjoint supports, it follows that

$$(6.6) \quad st_z(\hat{\Sigma}) = st_z(\hat{\Delta}) + st_z(\hat{\Psi}).$$

We first control the largest off-diagonal element of $\hat{\Sigma}$ by bounding $|\hat{\Psi}|_\infty$ with high probability. For every i, j , we have

$$\hat{\Psi}_{ij} = \frac{1}{2} \left[\frac{1}{n} \sum_{k=1}^n \left[\frac{1}{2} (X_{ki} + X_{kj})^2 - 1 \right] - \frac{1}{n} \sum_{k=1}^n \left[\frac{1}{2} (X_{ki} - X_{kj})^2 - 1 \right] \right].$$

Under H_0 , we have $X \sim \mathcal{N}(0, I_p)$, so by Lemma A.1, it holds for $t > 0$ that

$$\mathbf{P} \left(|\hat{\Psi}_{ij}| \geq 2\sqrt{\frac{t}{n}} + 2\frac{t}{n} \right) \leq 4e^{-t}.$$

Hence, by union bound on the off-diagonal terms, we get

$$\mathbf{P} \left(\max_{i < j} |\hat{\Psi}_{ij}| \geq 2\sqrt{\frac{t}{n}} + 2\frac{t}{n} \right) \leq 2p^2 e^{-t}.$$

Taking $t = \log(4p^2/\delta)$ yields that $|\hat{\Psi}|_\infty \leq z$, with probability $1 - \delta/2$, where

$$(6.7) \quad z = 2\sqrt{\frac{\log(4p^2/\delta)}{n}} + 2\frac{\log(4p^2/\delta)}{n}.$$

Note now that if we take z as in (6.7), then $st_z(\hat{\Psi}) = 0$ on an event \mathcal{E} of probability $1 - \delta/2$. Furthermore, since $\hat{\Delta}$ is a nonnegative diagonal matrix, then (6.6) yields that on the event \mathcal{E} , it holds

$$(6.8) \quad \lambda_{\max}(st_z(\hat{\Sigma})) = \lambda_{\max}(st_z(\hat{\Delta})) \leq \lambda_{\max}(\hat{\Delta}) = \max_{1 \leq i \leq p} \hat{\Delta}_{ii}.$$

Next, we control the largest diagonal element of $\hat{\Sigma}$ as follows. We have by definition of $\hat{\Delta}$, for every $i = 1, \dots, p$

$$\hat{\Delta}_{ii} = \frac{1}{n} \sum_{j=1}^n X_{ji}^2.$$

Applying Lemma A.1 and a union bound over the p diagonal terms, we get

$$\mathbf{P}\left(\max_{1 \leq i \leq p} \hat{\Delta}_{ii} \geq 1 + 2\sqrt{\frac{t}{n}} + 2\frac{t}{n}\right) \leq pe^{-t}.$$

Taking $t = \log(2p/\delta)$ yields with probability $1 - \delta/2$,

$$(6.9) \quad \max_{1 \leq i \leq p} \hat{\Delta}_{ii} \leq 1 + 2\sqrt{\frac{\log(2p/\delta)}{n}} + 2\frac{\log(2p/\delta)}{n}.$$

To conclude the proof of Proposition 6.2, observe that (6.5) implies that for all $z \geq 0$, we have

$$\text{SDP}_k(\hat{\Sigma}) \leq \lambda_{\max}(st_z(\hat{\Sigma})) + kz \leq \lambda_{\max}(st_z(\hat{\Delta})) + \lambda_{\max}(st_z(\hat{\Psi})) + kz,$$

where we used (6.6) and the triangle inequality for the operator norm.

Putting together (6.8) and (6.9) completes the proof. \square

6.3. *Hypothesis testing with convex methods.* Using the notation from Section 2, the results of the previous subsection can be written as

$$\mathbf{P}_{H_0}(\text{SDP}_k(\hat{\Sigma}) > \tilde{\tau}_0) \leq \delta, \quad \mathbf{P}_{H_1}(\text{SDP}_k(\hat{\Sigma}) < \tilde{\tau}_1) \leq \delta,$$

where $\tilde{\tau}_0$ and $\tilde{\tau}_1$ are given by

$$\begin{aligned} \tilde{\tau}_0 &= 1 + 2\sqrt{\frac{k^2 \log(4p^2/\delta)}{n}} + 2\frac{k \log(4p^2/\delta)}{n} + 2\sqrt{\frac{\log(2p/\delta)}{n}} + 2\frac{\log(2p/\delta)}{n}, \\ \tilde{\tau}_1 &= 1 + \theta - 2(1 + \theta)\sqrt{\frac{\log(1/\delta)}{n}}. \end{aligned}$$

Whenever $\tilde{\tau}_1 > \tilde{\tau}_0$, we take $\tau \in [\tilde{\tau}_0, \tilde{\tau}_1]$ and define the following computationally efficient test $\tilde{\psi}(\hat{\Sigma}) = \mathbf{1}\{\text{SDP}_k(\hat{\Sigma}) > \tau\}$. It discriminates between H_1 and H_0 with probability $1 - \delta$.

It remains to find for which values of θ the condition $\tilde{\tau}_1 > \tilde{\tau}_0$ holds. It corresponds to our minimum detection level.

THEOREM 6.1. Assume that p, n, k and δ are such that $\tilde{\theta} \leq 1$, where

$$(6.10) \quad \begin{aligned} \tilde{\theta} := & 2\sqrt{\frac{k^2 \log(4p^2/\delta)}{n}} + 2\frac{k \log(4p^2/\delta)}{n} + 2\sqrt{\frac{\log(2p/\delta)}{n}} \\ & + 2\frac{\log(2p/\delta)}{n} + 4\sqrt{\frac{\log(1/\delta)}{n}}. \end{aligned}$$

Then, for any $\theta > \tilde{\theta}$, any $\tau \in [\tilde{\tau}_0, \tilde{\tau}_1]$, the test $\tilde{\psi}(\hat{\Sigma}) = \mathbf{1}\{\text{SDP}_k(\hat{\Sigma}) > \tau\}$ discriminates between H_0 and H_1 with probability $1 - \delta$.

If we consider asymptotic regimes, for large p, n, k , taking $\delta = p^{-\beta}$ with $\beta > 0$, provides a sequence of tests $\tilde{\psi}_n$ that discriminate between H_0 and H_1 with probability converging to 1, for any fixed $\theta > 0$, if $k^2 \log(p)/n \rightarrow 0$.

Note that, compared to Theorem 4.1, the price to pay for using this convex relaxation is to multiply the minimum detection level by a factor \sqrt{k} . Such a gap is observed for these techniques in Amini and Wainwright (2009). Nevertheless, in most examples, k remains small and so is this price. As we will see in Section 8, there is strong evidence that $\tilde{\tau}_0$, which dominates the detection rate, cannot be made smaller and that therefore, our proof is tight.

6.4. *Simple methods.* While the SDP relaxation proposed in the previous subsection is provably computationally efficient, it is also known to scale poorly on large problems. Simple heuristics such as the diagonal method of Johnstone and Lu (2009) become more attractive for larger problems. A careful inspection of the proofs in the previous subsection is quite informative. It indicates that our results not only hold for the test $\tilde{\psi}(\hat{\Sigma})$ but for a test based on a simpler statistic arising from the dual formulation (6.5). Indeed, to control the behavior of $\text{SDP}_k(\hat{\Sigma})$ under H_0 , we showed that it was no larger than the *minimum dual perturbation* $\text{MDP}_k(\hat{\Sigma})$ defined by

$$(6.11) \quad \text{MDP}_k(\hat{\Sigma}) = \min_{z \geq 0} \{\lambda_{\max}(st_z(\hat{\Sigma})) + kz\}.$$

Clearly $\text{MDP}_k(\hat{\Sigma}) \geq \text{SDP}_k(\hat{\Sigma}) \geq \lambda_{\max}^k(\hat{\Sigma})$ so that both Propositions 6.1 and 6.2 still hold for $\text{SDP}_k(\hat{\Sigma})$ replaced by $\text{MDP}_k(\hat{\Sigma})$. As a result, for any $\theta > \tilde{\theta}$ the test $\hat{\psi}(\hat{\Sigma}) = \mathbf{1}\{\text{MDP}_k(\hat{\Sigma}) > \tau\}$ discriminates between H_0 and H_1 with probability $1 - \delta$.

Actually, a detection level of the same order as $\tilde{\theta}$ holds already for an even simpler test statistic: the largest diagonal element of $\hat{\Sigma}$. This method called *Johnstone's diagonal method* was first proposed by Johnstone and Lu (2009) and later studied by Amini and Wainwright (2009). For the problem of detection considered here, it dictates one to employ the test statistic $D(\hat{\Sigma}) = \max_{1 \leq i \leq p} \hat{\Sigma}_{ii}$. Using even simpler techniques than in Propositions 6.1 and 6.2, it is not hard to show that

$$\mathbf{P}_{H_0}(D(\hat{\Sigma}) > \tau_0^d) \leq \delta, \quad \mathbf{P}_{H_1}(D(\hat{\Sigma}) < \tau_1^d) \leq \delta$$

for quantiles τ_0^d and τ_1^d given by

$$\tau_0^d = 1 + \frac{1}{k}\theta - 2\left(1 + \frac{1}{k}\theta\right)\sqrt{\frac{\log(1/\delta)}{n}},$$

$$\tau_1^d = 1 + 2\sqrt{\frac{\log(p/\delta)}{n}} + 2\frac{\log(p/\delta)}{n}.$$

However, as we shall see in Section 9, on simulated datasets, MDP_k behaves much better than D in practice. It was proved by Amini and Wainwright (2009) that if the SDP (6.2) has a solution of rank one, then it is strictly better than Johnstone’s diagonal method. While they study a support recovery problem different from the detection problem considered here, it seems to indicate that the two methods are qualitatively different. However, the assumption that the SDP (6.2) has a solution of rank one is strong and unnecessary in our problem. Indeed, our results from Section 8 indicate that, if detecting a planted clique in a random graph is computationally hard, then for large (p, n, k) , the SDP method does not achieve better rates than the ones we prove. In particular, this result is a good indication that with high probability, the solution of the SDP is not rank-one for parameters in a range around the minimax detection level.

7. Generalization with weakened assumptions. In this section we investigate several extensions of our original problem. For simplicity, we denote by $\ast\text{DP}_k$ any of the two functionals MDP_k or SDP_k .

7.1. *Sparsity in terms of ℓ_q norm.* Fix $q \in (0, 2)$, and recall that $\mathcal{B}_q(R)$ is the set of unit vectors that are in an ℓ_q ball of radius $R > 0$. This relaxed notion of sparsity allows for vectors $v \in \mathbf{R}^p$ to have ordered coordinates that decay fast enough but never take value zero. Note that $q = 2$ corresponds to no sparsity and requires different techniques. It is therefore excluded from this section. Consider the following hypothesis testing problem:

$$H_0 : X \sim \mathcal{N}(0, I_p),$$

$$\tilde{H}_1^q : X \sim \mathcal{N}(0, I_p + \theta vv^\top), \quad v \in \mathcal{B}_q(k^{1/q-1/2}).$$

The radius $k^{1/q-1/2}$ is the smallest $R > 0$ such that $\mathcal{B}_0(k) \subset \mathcal{B}_q(R)$, making it the most natural relaxation of the notion of k -sparse vectors. Below, we show that it yields the same detection levels as for $q = 0$.

THEOREM 7.1. Fix $\nu > 0$. There exists a constant $C_\nu > 0$ such that if

$$(7.1) \quad \theta < \underline{\theta}_\nu := \sqrt{\frac{k \log(C_\nu p/k^2 + 1)}{n}} \wedge \frac{1}{\sqrt{2}},$$

it holds, for $q \in (0, 2)$

$$(7.2) \quad \inf_{\psi} \left\{ \mathbf{P}_0^n(\psi = 1) \vee \max_{v \in \mathcal{B}_q(k^{1/q-1/2})} \mathbf{P}_v^n(\psi = 0) \right\} \geq \frac{1}{2} - \nu,$$

where the infimum is taken over all possible tests.

PROOF. Let $v \in \mathbf{R}^p$ be a unit vector with sparsity k . It follows from Hölder’s inequality that $|v|_q \leq k^{1/q-1/2}$. Therefore, for any test ψ , we have

$$\max_{v \in \mathcal{B}_q(k^{1/q-1/2})} \mathbf{P}_v^n(\psi = 0) \geq \max_{v \in \mathcal{B}_0(k)} \mathbf{P}_v^n(\psi = 0),$$

and the result follows as a direct consequence of Theorem 5.1. \square

To show a matching upper bound, we use the following lemma.

LEMMA 7.1. *Let $v \in \mathbf{R}^p$ be a unit vector, $|v|_2 = 1$. Then, for any $r \geq 1$, there exists a r -sparse unit vector $x \in \mathcal{B}_0(r)$ such that*

$$1 - r^{1-2/q} |v|_q^2 \leq (x^\top v)^2 \leq 1.$$

PROOF. Assume without loss of generality that $|v_1| \geq \dots \geq |v_p|$. Define $\tilde{x}_j = v_j$ if $j \leq r$, $\tilde{x}_j = 0$ otherwise, and $x = \tilde{x}/|\tilde{x}|_2$. We have $(x^\top v)^2 = \tilde{x}^\top v = \sum_{j=1}^r |v_j|^2 = 1 - \sum_{j=r+1}^p |v_j|^2$. Moreover, since $|v_r| \leq r^{-1/q} |v|_q$,

$$\begin{aligned} \sum_{j=r+1}^p |v_j|^2 &\leq \sum_{j=r+1}^p |v_r|^{2-q} |v_j|^q \leq |v|_q^{2-q} r^{1-2/q} \sum_{j=r+1}^p |v_j|^q \\ &\leq r^{1-2/q} |v|_q^2. \end{aligned} \quad \square$$

Vectors in $\mathcal{B}_q(k^{1/q-1/2})$ can therefore be approximated by sparse unit vectors. This property can be leveraged to show that for an appropriate choice of k_q , a test based on $\lambda_{\max}^{k_q}(\hat{\Sigma})$ is optimal.

PROPOSITION 7.1. *Under \tilde{H}_1^q , let $\varepsilon \in (0, 1)$, and define k_q to be the smallest integer such that $k_q \geq k\varepsilon^{1/(1-2/q)}$. Then with probability $1 - \delta$,*

$$\lambda_{\max}^{k_q}(\hat{\Sigma}) \geq 1 + (1 - \varepsilon)\theta - 2(1 + \theta) \sqrt{\frac{\log(2/\delta)}{n}}.$$

PROOF. Let $x \in \mathbf{R}^p$ be the k_q -sparse unit norm approximation of v from Lemma 7.1. It follows from the proof of Proposition 4.1 that

$$\lambda_{\max}^{k_q}(\hat{\Sigma}) \geq 1 + \theta(v^\top x)^2 - 2(1 + \theta(v^\top x)^2) \sqrt{\frac{\log(2/\delta)}{n}}.$$

Lemma 7.1 with $r = k_q \geq k\varepsilon^{1/(1-2/q)}$ yields $1 - \varepsilon \leq (x^\top v)^2 \leq 1$. \square

Moreover, it follows from Proposition 4.2 that for any $\varepsilon \in (0, 1)$ and integer k_q , with probability $1 - \delta$, it holds

$$\lambda_{\max}^{k_q}(\hat{\Sigma}) \leq 1 + 2 \left(\frac{k_q \log(9ep/k_q) + \log(2/\delta)}{n} + \sqrt{\frac{k_q \log(9ep/k_q) + \log(2/\delta)}{n}} \right).$$

Since k_q is only a constant factor away from k for all $q \in (0, 2)$ and $\varepsilon \in (0, 1)$, the statistic $\lambda_{\max}^{k_q}(\hat{\Sigma})$ achieves optimal rates of detection.

In an estimation context, Vu and Lei (2012) [see also Birnbaum et al. (2013), Paul and Johnstone (2012) for related results using a different method] have examined the ℓ_q sparsity assumption for $q \in (0, 1]$. Their estimation method consists of maximizing the quadratic form $x \mapsto x^\top \hat{\Sigma} x$ over $\mathcal{B}_q(R)$ for some given $R > 0$. We argue that in light of Lemma 7.1, the estimation problem of Vu and Lei (2012) can be solved by maximizing the quadratic form over $\mathcal{B}_0(R')$ for some appropriate choice of R' that depends on k and q and extended to $q \in [0, 2)$. In particular, an algorithm for ℓ_0 -sparse PCA can be used for ℓ_q -sparse PCA.

Similar results hold for our convex relaxations. Following the same steps as in the proof of Theorem 6.1, we find that there exists a constant $C_q > 0$ such that tests based on $*\text{DP}_{k_q}$ discriminate between H_0 and \tilde{H}_1^q with probability $1 - \delta$, for any $\theta > C_q \tilde{\theta}$, where $\tilde{\theta}$ is defined in (6.10). In particular, a gap of size \sqrt{k} is observed between these methods and the optimal ones.

7.2. Sub-Gaussian random variables. Our results can be extended to more general assumptions, where the variables $X_1, \dots, X_n \in \mathbf{R}^p$ are sub-Gaussian in the following sense.

DEFINITION 7.1. A real-valued random variable G is said to be standard sub-Gaussian if $\mathbb{E}[\exp(t(G - \mathbb{E}[G]))] \leq \exp(t^2/2)$ for all $t \in \mathbf{R}$.

Let $Z_1, \dots, Z_n \in \mathbf{R}^p$ be i.i.d. vectors with i.i.d. standard sub-Gaussian coefficients, such that for all $i = 1, \dots, n$ it holds $\mathbb{E}[Z_i] = 0$, $\mathbb{E}[Z_i Z_i^\top] = I_p$.

Given a scatter matrix $\Sigma \succeq 0$, for any $i = 1, \dots, n$, define $X_i = \Sigma^{1/2} Z_i$. Sub-Gaussian random vectors were generated in the same way by Vu and Lei (2012). Under this condition, we define the new detection problem with hypotheses H'_0 and H'_1 , for $\theta > 0$ by

$$\begin{aligned} H'_0 : \Sigma &= I_p, \\ H'_1 : \Sigma &= I_p + \theta v v^\top, \quad v \in \mathcal{B}_0(k). \end{aligned}$$

Replacing Lemma A.1 by Lemma A.2 in the proofs of Propositions 4.2 and 6.2, we get, respectively, the two following results.

PROPOSITION 7.2. Under H'_1 , for $\theta \leq 1$, it holds with probability $1 - \delta$

$$\lambda_{\max}^k(\hat{\Sigma}) \geq 1 + \theta - 6 \left(64 \frac{\log(2/\delta)}{n} + 32 \sqrt{\frac{\log(2/\delta)}{n}} \right).$$

Moreover, under H'_0 , it holds with probability $1 - \delta$,

$$\lambda_{\max}^k(\hat{\Sigma}) \leq 1 + 352 \left(2 \frac{k \log(9ep/k) + \log(2/\delta)}{n} + \sqrt{\frac{k \log(9ep/k) + \log(2/\delta)}{n}} \right).$$

Similarly, for the $*DP_k$ statistic, we obtain the following bound.

PROPOSITION 7.3. Under H_0 , we have, with probability $1 - \delta$,

$$SDP_k(\hat{\Sigma}) \leq MDP_k(\hat{\Sigma}) \leq 1 + 6 \left(64 \sqrt{\frac{k^2 \log(4p^2/\delta)}{n}} + 128 \frac{k \log(4p^2/\delta)}{n} \right).$$

As a consequence, all the results from Sections 4 and 6 can be extended to the present sub-Gaussian case. In particular, the same gap between the detection levels of the two procedures is observed.

7.3. *Adversarial noise.* While our previous results rely heavily on the fact that the X_i are sub-Gaussian random vectors, we can find much weaker assumptions under which the results for detection using the $*DP$ statistics are still valid. We also describe an adversarial noise setting in which the detection level attained by $*DP_k$ is actually optimal. Assume that

$$(7.3) \quad \hat{\Sigma} = \Sigma + N.$$

Here the only assumption on N is that $|N|_{\infty} \leq \sqrt{\log(p/\delta)/n}$ with probability $1 - \delta$. Up to constant factor, this is a generalization of our initial setting, and can describe a situation where the data is censored, akin to the setting of [Loh and Wainwright \(2012\)](#), but where the censored entries are not necessarily chosen at random.

PROPOSITION 7.4. Under H_1 , we have with probability $1 - \delta$

$$*DP_k(\hat{\Sigma}) \geq \lambda_{\max}^k(\hat{\Sigma}) \geq 1 + \theta - k \sqrt{\frac{\log(p/\delta)}{n}}.$$

PROOF. Recall that for any v such that $|v|_0 \leq k$, we have

$$\begin{aligned} *DP_k(\hat{\Sigma}) &\geq \lambda_{\max}^k(\hat{\Sigma}) \geq v^{\top} \hat{\Sigma} v \geq v^{\top} (I_p + \theta v v^{\top}) v + v^{\top} N v \\ &\geq 1 + \theta - |N|_{\infty} |v|_1^2 \geq 1 + \theta - k |N|_{\infty}, \end{aligned}$$

which yields the desired result. \square

PROPOSITION 7.5. *Under H_0 , we have with probability $1 - \delta$*

$$\lambda_{\max}^k(\hat{\Sigma}) \leq *DP_k(\hat{\Sigma}) \leq 1 + k\sqrt{\frac{\log(p/\delta)}{n}}.$$

PROOF. It follows from (6.4) that $\lambda_{\max}^k(\hat{\Sigma}) \leq *DP_k(\hat{\Sigma}) \leq \lambda_{\max}(I_p) + k|N|_{\infty}$, which yields the desired result. \square

The following theorem follows from Propositions 7.4 and 7.5. We omit its proof.

THEOREM 7.2. *Let ψ^{adv} be the test defined by*

$$\psi^{\text{adv}}(\hat{\Sigma}) = \mathbf{1}\left\{ *DP_k(\hat{\Sigma}) > 1 + k\sqrt{\frac{\log(p/\delta)}{n}} \right\}.$$

Then ψ^{adv} discriminates between H_0 and H_1 with probability $1 - \delta$ if $\theta > 2k\sqrt{\log(p/\delta)/n}$.

We now prove the corresponding lower bound. Let $v = (v_1, \dots, v_p)^\top \in \mathbf{R}^p$ be such that $v_j = 1/\sqrt{k}$ if $j \leq k$ and $v_j = 0$ otherwise. Define the random matrix N that takes values $\pm \frac{\theta}{2}vv^\top$, each with probability $1/2$.

THEOREM 7.3. *There exists an adversarial model of the form (7.3) where $|N|_{\infty} \leq \sqrt{\log(p)/n}$, such that if $\theta \leq 2k\sqrt{\log(p)/n}$, then for any test $\psi(\hat{\Sigma}) \in \{0, 1\}$ it holds*

$$\mathbf{P}_{H_1}(\psi(\hat{\Sigma}) = 0) \vee \mathbf{P}_{H_0}(\psi(\hat{\Sigma}) = 1) \geq 1/2.$$

PROOF. Note first that $|N|_{\infty} = \theta/(2k) \leq \sqrt{(\log p)/n}$ so that

$$\mathbf{P}_{H_0}\left(\hat{\Sigma} = I_p + \frac{\theta}{2}vv^\top\right) = \frac{1}{2}, \quad \mathbf{P}_{H_1}\left(\hat{\Sigma} = I_p + \frac{\theta}{2}vv^\top\right) = \frac{1}{2}.$$

Therefore, if $\psi(I_p + \frac{\theta}{2}vv^\top) = 1$, then $\mathbf{P}_{H_0}(\psi(\hat{\Sigma}) = 1) \geq 1/2$ and if $\psi(I_p + \frac{\theta}{2}vv^\top) = 0$, then $\mathbf{P}_{H_1}(\psi(\hat{\Sigma}) = 0) \geq 1/2$. \square

Note that the lower bound in Theorem 7.3 below is not minimax since there exists one model under which all tests cannot discriminate between H_0 and H_1 with probability less than $1/2$. It implies that tests based on either $*DP_k$ and λ_{\max}^k are optimal.

8. Complexity theoretic lower bounds. The difference between the detection rates proved for the testing statistic λ_{\max}^k and the convex optimization based statistics SDP_k and MDP_k suggests a statistical cost for computational efficiency. Such phenomena are hinted at by [Chandrasekaran and Jordan \(2013\)](#). While it is not hard to see that our bounds are tight for the diagonal method, it is legitimate to wonder if the observed gap for SDP_k and MDP_k comes from a proof artifact, or an intrinsic limitation of the problem. The computational hardness of the related RIP certification has recently attracted a lot of interest. By reductions to problems with known complexity theoretic limitations, [Bandeira et al. \(2012\)](#) and [Koiran and Zouzias \(2012\)](#) prove that it is in general impossible to approximate in polynomial time the λ_{\max}^k statistic up to an arbitrarily small constant. Clearly, a constant factor approximation of λ_{\max}^k would suffice to achieve optimal detection rates. However, such results are not sufficient for two reasons. First they do not rule out the existence of a polynomial time algorithm that approximates λ_{\max}^k within a large enough constant. Second, such results are in nature *worst case*, meaning that the input matrix can be arbitrarily difficult for an algorithm. Rather, in our problem, the entry matrix is an empirical covariance matrix constructed from i.i.d. random vectors with Gaussian distribution. Hereafter, we develop a *polynomial time reduction* from another problem which is believed to be hard in average: the planted clique problem.

8.1. *Reduction to the planted clique problem.* A careful inspection of the proof of [Corollary 6.1](#) and the results of [Section 6.4](#) reveals that the only way to obtain better detection levels for the SDP and MDP statistics is to prove a better control of the statistics under the null hypothesis. We argue below that this is unlikely.

Let $X_1, \dots, X_n \in \mathbf{R}^p$ be i.i.d. Gaussian vectors with distribution $\mathcal{N}(0, I_p)$ and for any $\alpha \in [1, 2]$, consider the following hypothetical bound:

$$(B_\alpha) \quad \text{SDP}_k(\hat{\Sigma}) \leq 1 + C_\alpha \sqrt{\frac{k^\alpha \log(p/\delta)}{n}} \quad \text{with probability } 1 - \delta,$$

where $C_\alpha > 0$. Our prior results hinge on proving that B_2 holds. However, to achieve minimax optimal detection rates, one would need to prove B_1 . Reasoning by contradiction, we examine the consequences of B_α with $\alpha \in (1, 2)$. In particular, such bounds would yield polynomial time algorithms to detect small planted cliques in random graphs. Hereafter, we argue that the existence of such algorithms is unlikely.

Fix an integer $k \geq 0$, and let $\mathcal{G}(n, 1/2, k)$ be the distribution over the set of graphs on n vertices generated as follows. Pick k vertices at random, and place a clique⁴ between them; then connect every other pair of vertices by an edge independently with probability $1/2$. Note that for $k = 0$, $\mathcal{G}(n, 1/2, 0) = \mathcal{G}(n, 1/2)$ is

⁴A clique is a subset of fully connected vertices.

simply the distribution of an Erdős–Rényi random graph. In the decision version of the planted clique problem, called **Planted Clique**, one is given a graph G on n vertices and the goal is to test

$$H_0^{\text{PC}} : G \sim \mathcal{G}(n, 1/2),$$

$$H_1^{\text{PC}} : G \sim \mathcal{G}(n, 1/2, k)$$

for some given $k \geq 2$ with probability of error at most $\delta > 0$. The search version of this problem consists of finding the clique planted under H_1 . The search problem was introduced by [Jerrum \(1992\)](#) and [Kuřera \(1995\)](#) while the decision version is traditionally attributed to Saks; see [Krivelevich and Vu \(2002\)](#), Section 5. It is known [see, e.g., [Spencer \(1994\)](#)] that if $k > 2 \log_2 n$, the planted clique under H_1 is the only clique of size k in the graph, asymptotically almost surely. We consider only such values of k hereafter.

For $k = o(\sqrt{n})$ there is no known polynomial time algorithm that solves this problem. The first polynomial time algorithm for the case $k = C\sqrt{n}$ was proposed in [Alon, Krivelevich and Sudakov \(1998\)](#) and is based on spectral techniques. Subsequent algorithms with similar performance appeared in [Ames and Vavasis \(2011\)](#), [Dekel, Gurel-Gurevich and Peres \(2011\)](#), [Feige and Krauthgamer \(2000\)](#), [Feige and Ron \(2010\)](#). It is widely believed that there is no polynomial time algorithm that solves **Planted Clique** for any k of order n^c for some fixed positive $c < 1/2$, and it can even be proved that certain algorithmic techniques such as the Metropolis process [[Jerrum \(1992\)](#)] and the Lovász–Schrijver hierarchy of relaxations [[Feige and Krauthgamer \(2003\)](#)] fail at this task. Moreover, **Planted Clique** is provably hard in certain computational models, as seen in [Rossman \(2010\)](#), [Feldman et al. \(2013\)](#) which brings more evidence toward its hardness. Note that recent results by [Brubaker and Vempala \(2009\)](#), [Frieze and Kannan \(2008\)](#) based on r -dimensional tensors, suggest an algorithmic approach capable of finding a planted clique of size $O(n^{1/r})$, but currently this tensor-based approach is not known to yield a polynomial time algorithm for $r > 2$. The confidence in the difficulty of this problem has led researchers to prove hardness results assuming that the planted clique problem is indeed hard. Examples include cryptographic applications [[Juels and Peinado \(2000\)](#)], testing for k -wise dependence [[Alon et al. \(2007\)](#)], approximating Nash equilibria [[Hazan and Krauthgamer \(2011\)](#)] and approximating solutions to the densest k -subgraph problem [[Alon et al. \(2011\)](#)].

Consider the following polynomial-time reduction from a graph instance to random vectors, valid for the case $p = n$. Let A be the $n \times n$ adjacency matrix of a random graph G , and let U be the $n \times n$ matrix defined for any $1 \leq i \leq j$ by

$$U_{ij} = \begin{cases} 2A_{ij} - 1, & \text{if } i < j, \\ \varepsilon_{ij}, & \text{otherwise,} \end{cases}$$

where $\{\varepsilon_{ij}\}_{i,j}$ is a sequence of i.i.d. Rademacher ± 1 random variables. Moreover, let $Z^{(1)}, \dots, Z^{(n)} \in \mathbf{R}^n$ be n i.i.d. $\mathcal{N}(0, I_n)$ random vectors, and define

$X_{ij} = |Z_j^{(i)}|U_{ij}$. Finally define the $n \times n$ empirical covariance matrix $\hat{\Sigma}$ associated to the vectors $X_i = (X_{i1}, \dots, X_{in})^\top \in \mathbf{R}^n$ as in (2.1). This construction clearly takes polynomial time.

If $G \sim \mathcal{G}(n, 1/2)$, by construction, $X_1, \dots, X_n \in \mathbf{R}^n$ are i.i.d. centered standard Gaussian vectors, where all the coefficients are independent. If $G \sim \mathcal{G}(n, 1/2, k)$, it is no longer the case, and the $\ast\text{DP}_k$ statistic behaves in a qualitatively different manner.

8.2. *Computational theoretic lower bounds for SDP and MDP.* In this subsection, we illustrate the intrinsic limitations of the SDP and MDP methods in the detection problem using arguments borrowed from computational complexity. We begin by showing that both statistics take large values on the problem reduced from a graph with a planted clique.

LEMMA 8.1. *Let $G \sim \mathcal{G}(n, 1/2, k)$, $k \geq 14$ even, and $X_1, \dots, X_n \in \mathbf{R}^n$ be constructed as above. It holds, with probability $1 - \delta$,*

$$\text{MDP}_k(\hat{\Sigma}) \geq \text{SDP}_k(\hat{\Sigma}) \geq 1 + \frac{k^2}{4\pi n} - 3\sqrt{\frac{k \log(2/\delta)}{n}}.$$

PROOF. Let $S \subset \{1, \dots, n\}$ be the random subset of k vertices on which the clique has been planted. By construction, there are subsets S_1 and S_2 of S , of cardinality $k/2$, such that the random variables $X_{ij}, i \in S_1, j \in S_2$ are all positive almost surely. Assume without loss of generality that $S = \{1, \dots, k\}$, $S_1 = \{k/2 + 1, \dots, k\}$ and $S_2 = \{1, \dots, k/2\}$. Let $v = v(S_2)$ be the unit vector with j th coordinate equal to $2/\sqrt{k}$ if $j \in S_2$ and 0 otherwise. It yields

$$\begin{aligned} \text{MDP}_k(\hat{\Sigma}) &\geq \text{SDP}_k(\hat{\Sigma}) \geq v^\top \hat{\Sigma} v = \frac{1}{n} \sum_{i=1}^n (v^\top X_i)^2 \\ &\geq \frac{1}{n} \sum_{i \in S_1} (v^\top X_i)^2 + \frac{1}{n} \sum_{i \notin S} (v^\top X_i)^2 \\ (8.1) \quad &= \frac{1}{n} \sum_{i \in S_1} \frac{1}{|S_2|} \left(\sum_{j \in S_2} |Z_j^{(i)}| \right)^2 + \frac{n-k}{n} \frac{1}{n-k} \sum_{i \notin S} (v^\top X_i)^2. \end{aligned}$$

We begin by controlling the first term on the right-hand side of (8.1). For all $i \in S_1$, define the centered sub-Gaussian random variable

$$Y_i = \sum_{j \in S_2} \{ |Z_j^{(i)}| - \sqrt{2/\pi} \}$$

and observe that

$$\begin{aligned} \frac{1}{n} \sum_{i \in S_1} \frac{1}{|S_2|} \left(\sum_{j \in S_2} |Z_j^{(i)}| \right)^2 &= \frac{1}{n} \sum_{i \in S_1} \frac{1}{|S_2|} (Y_i + |S_2| \sqrt{2/\pi})^2 \\ &\geq \frac{2}{\pi} \frac{|S_1| \cdot |S_2|}{n} + 2\sqrt{\frac{2}{\pi}} \frac{1}{n} \sum_{i \in S_1} Y_i. \end{aligned}$$

It follows from Lemma A.3 that with probability $1 - \delta/2$, we have

$$\sum_{i \in S_1} Y_i \geq -\sqrt{2|S_1| \cdot |S_2| \log(2/\delta)}.$$

Together, the previous two displays yield

$$(8.2) \quad \frac{1}{n} \sum_{i \in S_1} \frac{1}{|S_2|} \left(\sum_{j \in S_2} |Z_j^{(i)}| \right)^2 \geq \frac{k^2}{2\pi n} - \frac{2k}{\sqrt{\pi n}} \sqrt{\log(2/\delta)}.$$

To control the second term on the right-hand side of (8.1), we use Lemma A.1. It holds with probability $1 - \delta/2$ that

$$(8.3) \quad \frac{1}{n-k} \sum_{i \notin S} (v^\top X_i)^2 = 1 + \frac{1}{n-k} \sum_{i \notin S} [(v^\top X_i)^2 - 1] \geq 1 - 2\sqrt{\frac{\log(2/\delta)}{n-k}}.$$

Therefore, with probability $1 - \delta$, we get from (8.1), (8.2) and (8.3) that

$$\begin{aligned} \text{SDP}_k(\hat{\Sigma}) &\geq \frac{k^2}{2\pi n} - \frac{2k}{\sqrt{\pi n}} \sqrt{\log(2/\delta)} + \frac{n-k}{n} \left(1 - 2\sqrt{\frac{\log(2/\delta)}{n-k}} \right) \\ &\geq 1 + \frac{k^2}{4\pi n} - 3\sqrt{\frac{k \log(2/\delta)}{n}}, \end{aligned}$$

where the last inequality holds for $k \geq 14$. This yields the desired result. \square

Next, we prove that improving substantially the bound of Proposition 6.2 (i.e., if B_α were to hold for some $\alpha \in [1, 2)$) would allow us to detect the presence of cliques of size n^c for some $c < 1/2$.

THEOREM 8.1. *Let $X_1, \dots, X_n \in \mathbf{R}^n$ be i.i.d. $\mathcal{N}(0, I_n)$ random vector and let $\hat{\Sigma}$ be their corresponding empirical covariance matrix as defined in (2.1). If for any $\alpha \in [1, 2]$, (B_α) is valid for $p = n$, that is,*

$$(B_\alpha) \quad \text{SDP}_k(\hat{\Sigma}) \leq 1 + C_\alpha \sqrt{\frac{k^\alpha \log(n/\delta)}{n}} \quad \text{with probability } 1 - \delta,$$

where $C_\alpha > 0$, then there exists a polynomial time algorithm that discriminates between $\mathcal{G}(n, 1/2)$ and $\mathcal{G}(n, 1/2, k)$ with probability $1 - \delta$, as soon as

$k \geq [Cn \log(n/\delta)]^{1/(4-\alpha)}$ for some constant $C > 0$ that depends only on C_α . The same holds if $\text{SDP}_k(\hat{\Sigma})$ is replaced by $\text{MDP}_k(\hat{\Sigma})$. In particular, for any fixed $\alpha < 2$ and $\delta > 0$, it allows one to detect the presence of cliques of size n^c for some $c < 1/2$ with probability $1 - \delta$.

PROOF. Note first that since $\text{SDP}_k(\hat{\Sigma}) \leq \text{MDP}_k(\hat{\Sigma})$, it suffices to prove the result for $\text{SDP}_k(\hat{\Sigma})$.

Let G be a random graph from H_0^{PC} or H_1^{PC} . Our goal is to construct a test ϕ that discriminates between the two hypotheses. Let $X_1, \dots, X_n \in \mathbf{R}^n$ be n random vectors obtained by the polynomial time reduction described in the previous subsection, and denote by $\hat{\Sigma}$ their associated empirical covariance matrix. We propose the following test:

$$\phi = \phi(\hat{\Sigma}) = \mathbf{1} \left\{ \text{SDP}_k^{(\varepsilon)}(\hat{\Sigma}) > 1 + C_\alpha \sqrt{\frac{k^\alpha \log(n/\delta)}{n}} \right\},$$

where $\text{SDP}_k^{(\varepsilon)}(\hat{\Sigma}) \geq \text{SDP}_k(\hat{\Sigma}) - \varepsilon$ is an approximation of the SDP solution with tolerance $\varepsilon \leq 1/\sqrt{n}$. In particular, $\text{SDP}_k^{(\varepsilon)}(\hat{\Sigma})$ and thus ϕ can be computed in polynomial time.

Recall that under H_0 (no planted clique), the X_i 's are i.i.d. $\mathcal{N}(0, I_n)$ so that $\phi = 0$ with probability $1 - \delta$, which controls the type I error appropriately. Moreover, we know from Lemma 8.1 that under H_1 , we have $\phi = 1$ with probability $1 - \delta$ as soon as

$$1 + C_\alpha \sqrt{\frac{k^\alpha \log(n/\delta)}{n}} \leq 1 + \frac{k^2}{4\pi n} - 3\sqrt{\frac{k \log(2/\delta)}{n}}.$$

Solving for k yields that it is sufficient to have $k \geq [Cn \log(n/\delta)]^{1/(4-\alpha)}$, for some constant $C > 0$ that depends only on C_α . As a result, our test allows us to detect the presence of cliques of size $[Cn \log(n/\delta)]^{1/(4-\alpha)}$. \square

The consequences of Theorem 8.1 can be taken two ways. If one believes that detecting planted cliques of size at most $O(n^c)$, $c < 1/2$ is hard, then suboptimality by a factor \sqrt{k} is intrinsic to the SDP relaxation. Otherwise, the SDP_k statistic allows to reach new detection levels for Planted Clique.

To conclude, observe that the above results apply to the specific tests based on MDP and SDP only. An interesting question is to find whether this limitation is intrinsic to *all* polynomial time computable tests. Currently, the main limitation of the above proof is that $\text{SDP}_k(\hat{\Sigma})$ is well controlled under H_1 , but it may no longer be the case for any other statistic.

9. Numerical experiments. Computational cost is a crucial element in this study. In [Bach, Ahipasaoglu and d’Aspremont \(2010\)](#), the SDP relaxation with accuracy ε is shown to have a total complexity of $\mathcal{O}(kp^3\sqrt{\log(p)}/\varepsilon)$. This is achieved by minimizing a smooth approximation of the dual function, using first order methods from [Nesterov \(2003\)](#). However, this polynomial cost is already prohibitive in a high-dimensional setting, and we study only tests based on the MDP_k statistic. The latter is the solution of a one-dimensional minimization problem, and is approximately solved by taking a uniform grid on the variable z . The purpose of this section is to illustrate the empirical behavior of tests based on MDP_k and to compare it with the diagonal method.

9.1. Comparison of simple methods. We simulate $N = 1000$ samples of n independent random vectors $X_1^0, \dots, X_n^0 \sim \mathcal{N}(0, I_p)$ and $X_1^1, \dots, X_n^1 \sim \mathcal{N}(0, I_p + \theta vv^\top)$, for random unit vectors v supported on $S = \{1, \dots, k\}$. The vector v_S is distributed uniformly on the unit sphere of dimension k .

It yields N empirical covariance matrices $\hat{\Sigma}_1^0, \dots, \hat{\Sigma}_N^0$ under H_0 and N of them, $\hat{\Sigma}_1^1, \dots, \hat{\Sigma}_N^1$ under H_1 . We compute the D and MDP_k statistics for these samples and compare their densities. We take $\theta = 4$ and observe that the D statistic yields two distributions under H_0 and H_1 that are hard to distinguish ([Figure 1](#), left). In particular, it is clear that the statistic D cannot discriminate between H_0 and H_1 for $\theta = 4$, with this set of parameters. However, the distributions of $\text{MDP}_k(\hat{\Sigma})$ under H_0 and H_1 have almost disjoint support so that it can discriminate between the two hypotheses with probability close to one.

9.2. Tightness of error bounds. In [Section 6](#), we prove that both the D and MDP_k statistics discriminate between H_0 and H_1 with high probability as long as




FIG. 1. For $p = 500$, $n = 200$, $k = 30$, $N = 1000$, estimated densities for the two statistics, under H_0 (whole line) and under H_1 (dashed line).

$\theta \geq Ck\sqrt{\log(p/k)/n}$. The previous subsection indicates that MDP_k actually performs better than D and it is pertinent to wonder if detection levels of order smaller than $\theta \geq Ck\sqrt{\log(p/k)/n}$ can be achieved. In this subsection, we bring numerical evidence that it is not the case and thus corroborate evidence from Section 8.

For MDP_k to be considered a tight (up to constant factor) approximation of λ_{\max}^k , it needs to discriminate between H_0 and H_1 with high probability as soon as θ is of the order $\sqrt{k\log(p/k)/n}$, which is the minimax optimal detection level that is also achieved by λ_{\max}^k . This behavior can be illustrated by showing a phase transition for the probability of error in the testing problem, as a function of θ , for different choices of (p, n, k) . More precisely, if MDP_k were a tight approximation of λ_{\max}^k , there should exist a critical value θ_{crit} and a constant C_{crit} , such that $\theta > \theta_{\text{crit}} = C_{\text{crit}}\sqrt{k\log(p/k)/n}$, the probability of type II error is close to 0. Moreover, C_{crit} should not depend on (p, n, k) . Our numerical results show that this is not the case. Instead, as predicted by the analysis of Section 6, our experiments point to θ_{crit} of order $k\sqrt{\log(p/k)/n}$.

In order to substantiate such effects, we use a reciprocal setting. For fixed $\theta = 1$, fixed probability of type I error (test level) and several choices of parameters (p, k) , we exhibit a phase transition for the probability of type II error $P_{\text{II}}(\cdot)$ as a function of the optimal and suboptimal scalings, defined respectively by

$$\eta^* = \frac{k}{n} \log\left(\frac{p}{k}\right) \quad \text{and} \quad \eta^\circ = \frac{k^2}{n} \log\left(\frac{p}{k}\right).$$

If $\eta \in \{\eta^*, \eta^\circ\}$ is the correct scaling, there should exist a critical value η_{crit} , independent of (p, n, k) , such that one of the following two scenarios hold. On the one hand, if MDP_k actually exhibits optimal rates, that is, if η^* is the correct scaling, then $\eta^* \mapsto P_{\text{II}}(\eta^*)$ should have a sharp transition from 0 to 1 around η_{crit} for all choices of parameters (p, k) . On the other hand, if the correct scaling for MDP_k is η° , then it is the function $\eta^\circ \mapsto P_{\text{II}}(\eta^\circ)$ that has a sharp transition around η_{crit} for all choices of parameters (p, k) .

We simulate $N = 1400$ samples of n independent random variables $X_1^0, \dots, X_n^0 \sim \mathcal{N}(0, I_p)$. It yields $\hat{\Sigma}_1^0, \dots, \hat{\Sigma}_N^0$ that are drawn under H_0 , and used to estimate the quantiles $q_{0.01}, q_{0.05}$ at 1% and 5% for the MDP_k statistic. The same process is repeated under H_1 to estimate the probability of type II error $\mathbf{P}_{H_1}(\text{MDP}_k(\hat{\Sigma}) > q_\alpha)$. To that end, we simulate $X_1^1, \dots, X_n^1 \sim \mathcal{N}(0, I_p + \theta vv^\top)$, for random unit vectors v supported on $S = \{1, \dots, k\}$. The restriction of v to S is distributed uniformly on the unit sphere of dimension k . To display a one-dimensional dependence, k is chosen equal to the integer part of \sqrt{p} .

Figure 2 compares the behavior of the functions $\eta^* \mapsto P_{\text{II}}(\eta^*)$ and $\eta^\circ \mapsto P_{\text{II}}(\eta^\circ)$. It clearly demonstrates the presence of a critical level $\eta_{\text{crit}} \simeq 0.1$ independent of (p, n, k) . The concomitance of the right curves for different choices of (p, n, k) indicates that η° rather than η^* is the correct scaling factor for the MDP_k statistic. This confirms the results of Section 8 and the existence of a statistical price to pay for computational efficiency.




FIG. 2. Probability of type II error $P_{\text{II}}(\eta)$ with scalings $\eta = \eta^*$ (left) and $\eta = \eta^\circ$ (right) for $p = \{50, 100, 200, 500\}$, $k = \lfloor \sqrt{p} \rfloor$, $N = 1400$. Test levels are $\alpha = 5\%$ (top) and $\alpha = 1\%$ (bottom).

APPENDIX: TECHNICAL LEMMAS

We gather in this Appendix various useful concentration inequalities. The first Lemma is due to Laurent and Massart.

LEMMA A.1 [Laurent and Massart (2000), Lemma 1]. *Let $Z_1, \dots, Z_n \sim \mathcal{N}(0, 1)$ be i.i.d. ce random variables, and define $Y = \frac{1}{n} \sum_{i=1}^n Z_i^2 - 1$. Then the two following tail bounds hold for any $t > 0$:*

$$\mathbf{P}\left(Y \leq -2\sqrt{\frac{t}{n}}\right) \leq e^{-t}, \quad \mathbf{P}\left(Y \geq 2\sqrt{\frac{t}{n}} + 2\frac{t}{n}\right) \leq e^{-t}.$$

This second lemma generalizes the previous one to sums of squares of sub-Gaussian random variables.

LEMMA A.2. *Let G_1, \dots, G_n be i.i.d. standard sub-Gaussian centered random variables. It holds*

$$\mathbf{P}\left(\left|\frac{1}{n} \sum_{i=1}^n (G_i^2 - \mathbb{E}[G_i^2])\right| > 2e\left(64\frac{t}{n} + 32\sqrt{\frac{t}{n}}\right)\right) \leq 2e^{-t}.$$

PROOF. Using a Chernoff bound and integrating the tails yields that $\mathbb{E}[|G|^p]^{1/p} \leq 2\sqrt{p}$, for any integer $p \geq 0$. It follows from these bounds, by a series expansion, that

$$(A.1) \quad \mathbb{E}[e^{t(G^2 - \mathbb{E}[G^2])}] \leq \exp(512e^2 t^2) \quad \text{for } 0 < t < 1/(32e).$$

For any $u \in \mathbf{R}^n$, define

$$S_n = \sum_{i=1}^n u_i (G_i^2 - \mathbb{E}[G_i^2]).$$

By a Chernoff bound, using equation (A.1), it holds for all $t > 0$,

$$\mathbf{P}(S_n \geq t) \leq \exp\left(-\min\left(\frac{t^2}{2048e^2|u|_2^2}, \frac{t}{64e|u|_\infty}\right)\right).$$

This implies our final result. \square

LEMMA A.3. *Let Z_1, \dots, Z_n be i.i.d. $\mathcal{N}(0, 1)$ random variables and define $Y = \sum_{i=1}^n |Z_i|$. Then, for any $t > 0$, it holds*

$$\mathbf{P}(Y - \mathbb{E}Y < -t) \leq e^{-t^2/(2n)}.$$

PROOF. Using a Chernoff bound, observe first that for any $s > 0$, we have

$$\begin{aligned} \mathbf{P}(Y - \mathbb{E}Y < -t) &= \mathbf{P}(\mathbb{E}Y - Y > t) \leq e^{-st} \mathbb{E}[e^{s(\mathbb{E}Y - Y)}] \\ &= e^{-st} \prod_{i=1}^n \mathbb{E}[e^{-s(\mathbb{E}|Z_i| - |Z_i|)}]. \end{aligned}$$

Moreover,

$$\mathbb{E}[e^{-s(\mathbb{E}|Z_i| - |Z_i|)}] \leq 2e^{-s\mathbb{E}|Z_i|} \mathbb{E}[e^{sZ_i}] = 2e^{-s\sqrt{\pi/2}} e^{s^2/2}.$$

The above two displays yield

$$\begin{aligned} \mathbf{P}(Y - \mathbb{E}Y < -t) &\leq 2^n \inf_{s>0} \exp\left(-st - ns\sqrt{\pi/2} + n\frac{s^2}{2}\right) \\ &= 2^n \exp\left(-\frac{(t + n\sqrt{\pi/2})^2}{2n}\right) \leq e^{-t^2/(2n)}. \quad \square \end{aligned}$$

REFERENCES

- ADDARIO-BERRY, L., BROUTIN, N., DEVROYE, L. and LUGOSI, G. (2010). On combinatorial testing problems. *Ann. Statist.* **38** 3063–3092. [MR2722464](#)
- ALON, N., KRIVELEVICH, M. and SUDAKOV, B. (1998). Finding a large hidden clique in a random graph. *Random Structures Algorithms* **13** 457–466. [MR1662795](#)

- ALON, U., BARKAI, N., NOTTERMAN, D. A., GISH, K., YBARRA, S., MACK, D. and LEVINE, A. J. (1999). Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. *Proc. Natl. Acad. Sci. USA* **96** 6745–6750.
- ALON, N., ANDONI, A., KAUFMAN, T., MATULEF, K., RUBINFELD, R. and XIE, N. (2007). Testing k -wise and almost k -wise independence. In *STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing* 496–505. ACM, New York. [MR2402475](#)
- ALON, N., ARORA, S., MANOKARAN, R., MOSHKOVITZ, D. and WEINSTEIN, O. (2011). On the inapproximability of the densest κ -subgraph problem. Unpublished manuscript.
- AMES, B. P. W. and VAVASIS, S. A. (2011). Nuclear norm minimization for the planted clique and biclique problems. *Math. Program.* **129** 69–89. [MR2831403](#)
- AMINI, A. A. and WAINWRIGHT, M. J. (2009). High-dimensional analysis of semidefinite relaxations for sparse principal components. *Ann. Statist.* **37** 2877–2921. [MR2541450](#)
- ARIAS-CASTRO, E., BUBECK, S. and LUGOSI, G. (2012). Detection of correlations. *Ann. Statist.* **40** 412–435. [MR3014312](#)
- ARIAS-CASTRO, E., CANDÈS, E. J. and DURAND, A. (2011). Detection of an anomalous cluster in a network. *Ann. Statist.* **39** 278–304. [MR2797847](#)
- ARIAS-CASTRO, E., CANDÈS, E. J. and PLAN, Y. (2011). Global testing under sparse alternatives: ANOVA, multiple comparisons and the higher criticism. *Ann. Statist.* **39** 2533–2556. [MR2906877](#)
- BACH, F., AHIPASAOGLU, S. D. and D'ASPREMONT, A. (2010). Convex relaxations for subset selection. Preprint. Available at [arXiv:1006.3601v1](#).
- BAI, Z. D. (1999). Methodologies in spectral analysis of large-dimensional random matrices, a review. *Statist. Sinica* **9** 611–677. [MR1711663](#)
- BAIK, J., BEN AROUS, G. and PÉCHÉ, S. (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Ann. Probab.* **33** 1643–1697. [MR2165575](#)
- BAIK, J. and SILVERSTEIN, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *J. Multivariate Anal.* **97** 1382–1408. [MR2279680](#)
- BANDEIRA, A. S., DOBRIBAN, E., MIXON, D. G. and SAWIN, W. F. (2012). Certifying the restricted isometry property is hard. Preprint. Available at [arXiv:1204.1580](#).
- BARAUD, Y. (2002). Non-asymptotic minimax rates of testing in signal detection. *Bernoulli* **8** 577–606. [MR1935648](#)
- BICKEL, P. J. and LEVINA, E. (2008). Covariance regularization by thresholding. *Ann. Statist.* **36** 2577–2604. [MR2485008](#)
- BIRNBAUM, A., JOHNSTONE, I. M., NADLER, B. and PAUL, D. (2013). Minimax bounds for sparse PCA with noisy high-dimensional data. *Ann. Statist.* **41** 1055–1084.
- BOYD, S. and VANDENBERGHE, L. (2004). *Convex Optimization*. Cambridge Univ. Press, Cambridge. [MR2061575](#)
- BRUBAKER, S. C. and VEMPALA, S. S. (2009). Random tensors and planted cliques. In *Approximation, Randomization, and Combinatorial Optimization. Lecture Notes in Computer Science* **5687** 406–419. Springer, Berlin. [MR2551018](#)
- BUTUCEA, C. and INGSTER, Y. I. (2013). Detection of a sparse submatrix of a high-dimensional noisy matrix. *Bernoulli*. To appear. Available at [arXiv:1109.0898v1](#).
- CAI, T. T., MA, Z. and WU, Y. (2012). Sparse PCA: Optimal rates and adaptive estimation. Preprint. Available at [arXiv:1211.1309](#).
- CAI, T. T., ZHANG, C.-H. and ZHOU, H. H. (2010). Optimal rates of convergence for covariance matrix estimation. *Ann. Statist.* **38** 2118–2144. [MR2676885](#)
- CHANDRASEKARAN, V. and JORDAN, M. I. (2013). Computational and statistical tradeoffs via convex relaxation. *Proc. Natl. Acad. Sci.* **110** E1181–E1190. [MR3047651](#)
- CHEN, X. (2011). Adaptive elastic-net sparse principal component analysis for pathway association testing. *Stat. Appl. Genet. Mol. Biol.* **10** Art. 48, 23. [MR2851290](#)

- D'ASPREMONT, A., BACH, F. and EL GHAOUI, L. (2008). Optimal solutions for sparse principal component analysis. *J. Mach. Learn. Res.* **9** 1269–1294. [MR2426043](#)
- D'ASPREMONT, A., BACH, F. and GHAOUI, L. E. (2012). Approximation bounds for sparse principal component analysis. Preprint. Available at [arXiv:1205.0121](#).
- D'ASPREMONT, A., EL GHAOUI, L., JORDAN, M. I. and LANCKRIET, G. R. G. (2007). A direct formulation for sparse PCA using semidefinite programming. *SIAM Rev.* **49** 434–448 (electronic). [MR2353806](#)
- DEKEL, Y., GUREL-GUREVICH, O. and PERES, Y. (2011). Finding hidden cliques in linear time with high probability. In *ANALCO11—Workshop on Analytic Algorithmics and Combinatorics* 67–75. SIAM, Philadelphia, PA. [MR2815485](#)
- DONOHO, D. and JIN, J. (2004). Higher criticism for detecting sparse heterogeneous mixtures. *Ann. Statist.* **32** 962–994. [MR2065195](#)
- EL KAROUI, N. (2008). Operator norm consistent estimation of large-dimensional sparse covariance matrices. *Ann. Statist.* **36** 2717–2756. [MR2485011](#)
- FEIGE, U. and KRAUTHGAMER, R. (2000). Finding and certifying a large hidden clique in a semi-random graph. *Random Structures Algorithms* **16** 195–208. [MR1742351](#)
- FEIGE, U. and KRAUTHGAMER, R. (2003). The probable value of the Lovász–Schrijver relaxations for maximum independent set. *SIAM J. Comput.* **32** 345–370 (electronic). [MR1969394](#)
- FEIGE, U. and RON, D. (2010). Finding hidden cliques in linear time. In *21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (AofA'10)* 189–203. Assoc. Discrete Math. Theor. Comput. Sci., Nancy. [MR2735341](#)
- FELDMAN, V., GRIGORESCU, E., REYZIN, L., VEMPALA, S. and XIAO, Y. (2013). Statistical algorithms and a lower bound for planted clique. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing, STOC'13* 655–664. ACM, New York.
- FÉRAL, D. and PÉCHÉ, S. (2009). The largest eigenvalues of sample covariance matrices for a spiked population: Diagonal case. *J. Math. Phys.* **50** 073302, 33. [MR2548630](#)
- FRIEZE, A. and KANNAN, R. (2008). A new approach to the planted clique problem. In *FSTTCS 2008: IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. LIPIcs. Leibniz Int. Proc. Inform.* **2** 187–198. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern. [MR2874057](#)
- GEMAN, S. (1980). A limit theorem for the norm of random matrices. *Ann. Probab.* **8** 252–261. [MR0566592](#)
- GOEMANS, M. X. and WILLIAMSON, D. P. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.* **42** 1115–1145. [MR1412228](#)
- HAZAN, E. and KRAUTHGAMER, R. (2011). How hard is it to approximate the best Nash equilibrium? *SIAM J. Comput.* **40** 79–91. [MR2765712](#)
- INGSTER, Y. I. (1982). The asymptotic efficiency of tests for a simple hypothesis against a composite alternative. *Teor. Veroyatn. Primen.* **27** 587–592. [MR0673934](#)
- INGSTER, Y. I., TSYBAKOV, A. B. and VERZELEN, N. (2010). Detection boundary in sparse regression. *Electron. J. Stat.* **4** 1476–1526. [MR2747131](#)
- JENATTON, R., OBOZINSKI, G. and BACH, F. (2010). Structured sparse principal component analysis. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- JERRUM, M. (1992). Large cliques elude the Metropolis process. *Random Structures Algorithms* **3** 347–359. [MR1179827](#)
- JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* **29** 295–327. [MR1863961](#)
- JOHNSTONE, I. M. and LU, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *J. Amer. Statist. Assoc.* **104** 682–693. [MR2751448](#)

- JOURNÉE, M., NESTEROV, Y., RICHTÁRIK, P. and SEPULCHRE, R. (2010). Generalized power method for sparse principal component analysis. *J. Mach. Learn. Res.* **11** 517–553. [MR2600619](#)
- JUELS, A. and PEINADO, M. (2000). Hiding cliques for cryptographic security. *Des. Codes Cryptogr.* **20** 269–280. [MR1779310](#)
- KARP, R. M. (1972). Reducibility among combinatorial problems. In *Complexity of Computer Computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972)* 85–103. Plenum, New York. [MR0378476](#)
- KOIRAN, P. and ZOUZIAS, A. (2012). Hidden cliques and the certification of the restricted isometry property. Preprint. Available at [arXiv:1211.0665](#).
- KRIVELEVICH, M. and VU, V. H. (2002). Approximating the independence number and the chromatic number in expected polynomial time. *J. Comb. Optim.* **6** 143–155. [MR1885488](#)
- KUČERA, L. (1995). Expected complexity of graph partitioning problems. *Discrete Appl. Math.* **57** 193–212. [MR1327775](#)
- LAURENT, B. and MASSART, P. (2000). Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.* **28** 1302–1338. [MR1805785](#)
- LOH, P.-L. and WAINWRIGHT, M. J. (2012). High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity. *Ann. Statist.* **40** 1637–1664. [MR3015038](#)
- LU, Z. and ZHANG, Y. (2012). An augmented Lagrangian approach for sparse principal component analysis. *Math. Program.* **135** 149–193. [MR2968253](#)
- MA, S. (2011). Alternating direction method of multipliers for sparse principal component analysis. Preprint. Available at [arXiv:1111.6703v1](#).
- MA, Z. (2013). Sparse principal component analysis and iterative thresholding. *Ann. Statist.* **41** 772–801.
- NADLER, B. (2008). Finite sample approximation results for principal component analysis: A matrix perturbation approach. *Ann. Statist.* **36** 2791–2817. [MR2485013](#)
- NESTEROV, Y. (2003). *Introductory Lectures on Convex Optimization*. Springer, New York.
- NESTEROV, Y. and NEMIROVSKII, A. (1987). *Interior-Point Polynomial Algorithms in Convex Programming*. Society for Industrial Mathematics **13**. SIAM, Philadelphia, PA.
- ONATSKI, A., MOREIRA, M. J. and HALLIN, M. (2013). Asymptotic power of sphericity tests for high-dimensional data. *Ann. Statist.* **41** 1204–1231.
- PAUL, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statist. Sinica* **17** 1617–1642. [MR2399865](#)
- PAUL, D. and JOHNSTONE, I. M. (2012). Augmented sparse principal component analysis for high dimensional data. Preprint. Available at [arXiv:1202.1242v1](#).
- ROSSMAN, B. (2010). Average-case complexity of detecting cliques. Ph.D. thesis, Massachusetts Institute of Technology. [MR2873600](#)
- SHEN, D., SHEN, H. and MARRON, J. S. (2013). Consistency of sparse PCA in high dimension, low sample size contexts. *J. Multivariate Anal.* **115** 317–333. [MR3004561](#)
- SPENCER, J. (1994). *Ten Lectures on the Probabilistic Method*, 2nd ed. *CBMS-NSF Regional Conference Series in Applied Mathematics* **64**. SIAM, Philadelphia, PA. [MR1249485](#)
- SUN, X. and NOBEL, A. B. (2008). On the size and recovery of submatrices of ones in a random binary matrix. *J. Mach. Learn. Res.* **9** 2431–2453. [MR2460888](#)
- SUN, X. and NOBEL, A. B. (2013). On the maximal size of large-average and ANOVA-fit submatrices in a Gaussian random matrix. *Bernoulli* **19** 275–294.
- TSYBAKOV, A. B. (2009). *Introduction to Nonparametric Estimation*. Springer, New York. Revised and extended from the 2004 French original, translated by Vladimir Zaiats. [MR2724359](#)
- VERSHYNIN, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing* 210–268. Cambridge Univ. Press, Cambridge. [MR2963170](#)
- VERZELEN, N. (2012). Minimax risks for sparse regressions: Ultra-high dimensional phenomenons. *Electron. J. Stat.* **6** 38–90. [MR2879672](#)

- VU, V. and LEI, J. (2012). Minimax rates of estimation for sparse pca in high dimensions. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics April 21–23, 2012, La Palma, Canary Islands, Vol. 22 of JMLR W&CP* 1278–1286.
- WRIGHT, J., GANESH, A., YANG, A., ZHOU, Z. and MA, Y. (2011). Sparsity and robustness in face recognition. Preprint. Available at [arXiv:1111.1014](https://arxiv.org/abs/1111.1014).
- YIN, Y. Q., BAI, Z. D. and KRISHNAIAH, P. R. (1988). On the limit of the largest eigenvalue of the large-dimensional sample covariance matrix. *Probab. Theory Related Fields* **78** 509–521. [MR0950344](https://doi.org/10.1007/BF01233444)

DEPARTMENT OF OPERATIONS RESEARCH
AND FINANCIAL ENGINEERING
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: qberthet@princeton.edu
rigollet@princeton.edu