# APPROXIMATION BOUNDS FOR SPARSE PRINCIPAL COMPONENT ANALYSIS 

ALEXANDRE D'ASPREMONT, FRANCIS BACH, AND LAURENT EL GHAOUI


#### Abstract

We produce approximation bounds on a semidefinite programming relaxation for sparse principal component analysis. These bounds control approximation ratios for tractable statistics in hypothesis testing problems where data points are sampled from Gaussian models with a single sparse leading component.


We study approximation bounds for a semidefinite relaxation of the sparse eigenvalue problem, written here in penalized form

$$
\max _{\|x\|_{2}=1} x^{T} \Sigma x-\rho \operatorname{Card}(x)
$$

in the variable $x \in \mathbb{R}^{n}$, where $\Sigma \in \mathbf{S}_{n}$ and $\rho \geq 0$. Sparse eigenvalues appear in many applications in statistics and machine learning. Sparse eigenvectors are often used, for example, to improve the interpretability of principal component analysis, while sparse eigenvalues control recovery thresholds in compressed sensing [Candes and Tao, 2007]. Several convex relaxations and greedy algorithms have been developed to find approximate solutions (see d'Aspremont et al. [2007, 2008], Journée et al. [2008], Journée et al. [2008] among others), but except in simple scenarios where $\rho$ is small and the two leading eigenvalues of $\Sigma$ are separated, very little is known about the tightness of these approximation methods.

Here, using randomization techniques based on [Ben-Tal and Nemirovski, 2002], we derive simple approximation bounds for the semidefinite relaxation derived in [d'Aspremont, Bach, and El Ghaoui, 2008]. We do not produce a constant approximation ratio and our bounds depend on the optimum value of the semidefinite relaxation: the higher this value, the better the approximation. A similar behavior was observed by Zwick [1999] for the semidefinite relaxation to MAXCUT, who showed that the classical approximation ratio of Goemans and Williamson [1995] can be improved when the value of the cut is high enough.

We then show that, in some applications, it is possible to bound a priori the optimum value of the semidefinite relaxation, hence produce a lower bound on the approximation ratio. In particular, following recent works by [Amini and Wainwright, 2009, Berthet and Rigollet, 2012], we focus on the problem of detecting the presence of a (significant) sparse principal component in a Gaussian model, hence test the significance of eigenvalues isolated by sparse principal component analysis. More precisely, we apply our approximation results to the problem of discriminating between the two Gaussian models

$$
\mathcal{N}\left(0, \mathbf{I}_{n}\right) \quad \text { and } \quad \mathcal{N}\left(0, \mathbf{I}_{n}+\theta v v^{T}\right)
$$

where $v \in \mathbb{R}^{n}$ is a sparse vector with unit Euclidean norm and cardinality $k$. We use a convex relaxation for the sparse eigenvalue problem to produce a tractable statistic for this hypothesis testing problem and show that in a high-dimensional setting where the dimension $n$, the number of samples $m$ and the cardinality $k$ grow towards infinity proportionally, the detection threshold on $\theta$ remains finite.

More broadly speaking, in the spirit of smoothed analysis [Spielman and Teng, 2001], this shows that analyzing the performance of semidefinite relaxations on random problem instances is sometimes easier and provides a somewhat more realistic description of typical approximation ratios. Another classical example of this phenomenon is a MAXCUT-like problem arising in statistical physics, for which explicit (asymptotic) formulas can be derived for certain random instances, e.g. the Parisi formula [Mezard et al., 1987, Mezard and Montanari, 2009, Talagrand, 2010] for computing the ground state of spin glasses in the SherringtonKirkpatrick model. It thus seems that comparing the performance of convex relaxations on random problem

[^0]instances (e.g. in detection problems) often yields a more nuanced understanding of their performance in cases where uniform approximation ratios are hard to derive or analyze.

The paper is organized as follows. The next section recalls a few definitions on sparse eigenvalue problems. Section 2 recalls the construction of the semidefinite relaxation in [d'Aspremont et al., 2008]. Section 3 derives approximation bounds on the solution of this relaxation. Section 4 studies the performance of these approximation bounds on the sparse eigenvector detection problem. Section 5 presents some algorithms for solving the semidefinite relaxation used as a test statistic in the detection problem. Finally, we present some numerical results in Section 6.

## 1. Sparse eigenvalues

We begin by formally defining sparse eigenvalues. Let $\Sigma \in \mathbf{S}_{n}$ be a symmetric matrix. We define the sparse maximum eigenvalues of the matrix $\Sigma$ as

$$
\begin{array}{rll}
\lambda_{\max }^{k}(\Sigma) \triangleq & \text { max. } & x^{T} \Sigma x \\
\text { s.t. } & \operatorname{Card}(x) \leq k  \tag{1}\\
& \|x\|_{2}=1,
\end{array}
$$

in the variable $x \in \mathbb{R}^{n}$ where the parameter $k>0$ controls the sparsity of the solution. We can similarly define sparse minimum eigenvalues as

$$
\begin{array}{rll}
\lambda_{\min }^{k}(\Sigma) \triangleq & \min . & x^{T} \Sigma x \\
\text { s.t. } & \operatorname{Card}(x) \leq k  \tag{2}\\
& \|x\|_{2}=1,
\end{array}
$$

in the variable $x \in \mathbb{R}^{n}$. Because $\lambda_{\max }^{k}(\Sigma+\alpha \mathbf{I})$ is affine in $\alpha$, we have

$$
\lambda_{\min }^{k}(\Sigma)=\lambda_{\max }(\Sigma)-\lambda_{\max }^{k}\left(\lambda_{\max }(\Sigma) \mathbf{I}-\Sigma\right)
$$

and the following sections will be focused on approximating $\lambda_{\max }^{k}(\Sigma)$.

## 2. SEmidefinite relaxation

Here, we first recall the semidefinite relaxation for (1) derived in [d'Aspremont, Bach, and El Ghaoui, 2008]. We assume without loss of generality that $\Sigma \in \mathbf{S}_{n}$ is positive semidefinite (we can always add a multiple of the identity) and that the $n$ variables are ordered by decreasing marginal variances, i.e. that $\Sigma_{11} \geq \ldots \geq \Sigma_{n n}$. We also assume that we are given a square root $A$ of the matrix $\Sigma$ with

$$
\Sigma=A^{T} A,
$$

where $A \in \mathbb{R}^{n \times n}$ and we denote by $a_{1}, \ldots, a_{n} \in \mathbb{R}^{n}$ the columns of $A$. Note that the problem and our algorithms are invariant by permutations of $\Sigma$ and by the choice of square root $A$. In practice, we are very often given the data matrix $A$ instead of the covariance $\Sigma$. As we will see below, we can directly exclude variables for which $\Sigma_{i i}<\rho$, hence we can assume w.l.o.g. that

$$
0<\rho<\min _{i \in[1, n]} \Sigma_{i i} .
$$

If this condition is not satisfied, the variable $i$ will never be part of the optimal support and we can focus on the reduced problem.
2.1. Relaxation bounds on sparse eigenvalues. We can rewrite the maximum eigenvalue problem in terms of the data matrix $A$. We start by writing

$$
\begin{aligned}
\lambda_{\max }^{k}(\Sigma) & =\max _{\substack{\operatorname{Card}(x) \leq k \\
\|x\|_{2}=1}}^{\operatorname{lon}} x^{T} \Sigma x \\
& =\max _{\substack{u \in\{0,1\}^{n} \\
\mathbf{1}^{T} u=k}} \lambda_{\max }\left(\sum_{i=1}^{n} u_{i} a_{i} a_{i}^{T}\right) \\
& =\max _{\substack{u \in\{0,1\}^{n} \\
\mathbf{1}^{T} u=k}} \max _{\|x\|_{2}=1} \sum_{i=1}^{n} u_{i}\left(a_{i}^{T} x\right)^{2} \\
& =\max _{\|x\|_{2}=1} \max _{u \in\{0,1\}^{n}} \sum_{i=1}^{n} u_{i}\left(a_{i}^{T} x\right)^{2},
\end{aligned}
$$

and use the fact that

$$
\max _{\substack{u \in\{0,1\}^{n} \\ \mathbf{1}^{T} u=k}} \sum_{i=1}^{n} u_{i} b_{i}=\min _{\rho \geq 0}\left\{\sum_{i=1}^{n}\left(b_{i}-\rho\right)_{+}+\rho k\right\}
$$

for any $b \in \mathbb{R}^{n}$, to write

$$
\begin{aligned}
& \max _{\|x\|_{2}=1} \max _{\substack{u \in\{0,1\}^{n} \\
\mathbf{1}^{T} u=k}} \sum_{i=1}^{n} u_{i}\left(a_{i}^{T} x\right)^{2} \\
= & \max _{\|x\|_{2}=1} \min _{\rho \geq 0}\left\{\sum_{i=1}^{n}\left(\left(a_{i}^{T} x\right)^{2}-\rho\right)_{+}+\rho k\right\} \\
= & \max _{\substack{\operatorname{Rank}(X)=1 \\
X \succeq 0, \operatorname{Tr}(X)=1}} \min _{\rho \geq 0}\left\{\sum_{i=1}^{n}\left(X^{1 / 2}\left(a_{i} a_{i}^{T}-\rho \mathbf{I}\right) X^{1 / 2}\right)_{+}+\rho k\right\},
\end{aligned}
$$

where the last equality follows from the fact, when $\operatorname{Rank}(X)=1$ the only potentially nonnegative eigenvalue of $\left(X^{1 / 2}\left(a_{i} a_{i}^{T}-\rho \mathbf{I}\right) X^{1 / 2}\right)$ is $\left(a_{i}^{T} x\right)^{2}-\rho$, and $X^{1 / 2}=X=x x^{T}$ when $\|x\|_{2}=1$. We then produce a semidefinite relaxation for (1) by simply dropping the rank constraint to get the following bound

$$
\begin{equation*}
\lambda_{\max }^{k}(\Sigma) \leq \min _{\rho \geq 0}\left\{\max _{\substack{\operatorname{Tr}(X)=1 \\ X \succeq 0}} \sum_{i=1}^{n}\left(X^{1 / 2}\left(a_{i} a_{i}^{T}-\rho \mathbf{I}\right) X^{1 / 2}\right)_{+}+\rho k\right\} \tag{3}
\end{equation*}
$$

which is equivalent to a semidefinite program. Note that because $\operatorname{Rank}(X)=1$ defines a non-convex set, we cannot simply switch the min and the max and this last inequality is potentially strict. i.e. the semidefinite relaxation only produces an upper bound on $\lambda_{\max }^{k}(\Sigma)$.
2.2. Penalized problem. We now focus on the inner optimization problem in (3). Starting from a penalized version of problem (1), written

$$
\begin{equation*}
\phi(\rho) \triangleq \max _{\|x\|_{2}=1} x^{T} \Sigma x-\rho \operatorname{Card}(x) \tag{4}
\end{equation*}
$$

it was shown in d'Aspremont et al. [2008] that

$$
\begin{aligned}
\phi(\rho) & =\max _{\|x\|_{2}=1} \sum_{i=1}^{n}\left(\left(a_{i}^{T} x\right)^{2}-\rho\right)_{+} \\
& =\max _{\substack{\operatorname{Rak}(X)=1 \\
X \succeq 0, \operatorname{Tr}(X)=1}} \sum_{i=1}^{n} \operatorname{Tr}\left(X^{1 / 2}\left(a_{i} a_{i}^{T}-\rho \mathbf{I}\right) X^{1 / 2}\right)_{+}
\end{aligned}
$$

and we write $\psi(\rho)$ the semidefinite relaxation of this last problem

$$
\begin{array}{rll}
\psi(\rho) \triangleq & \max . & \sum_{i=1}^{n} \operatorname{Tr}\left(X^{1 / 2} a_{i} a_{i}^{T} X^{1 / 2}-\rho X\right)_{+}  \tag{5}\\
\text {s.t. } & \operatorname{Tr}(X)=1, X \succeq 0
\end{array}
$$

which is equivalent to a semidefinite program [d'Aspremont et al., 2008] and is the inner problem in (3). In the next section, we use this quantity as a test statistic for detecting significant sparse eigenvectors.

## 3. Approximation Bounds

Using the randomization argument detailed in [Ben-Tal and Nemirovski, 2002, El Ghaoui, 2006], we can derive an explicit bound on the quality of the semidefinite relaxation (5).
Proposition 3.1. Let us call $X$ the optimal solution to problem (5) and let $r=\operatorname{Rank}(X)$, we have

$$
\begin{equation*}
n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right) \leq \phi(\rho) \leq \psi(\rho) \tag{6}
\end{equation*}
$$

where

$$
\begin{equation*}
\vartheta_{r}(x) \triangleq \mathbf{E}\left[\left(x \xi_{1}^{2}-\frac{1}{r-1} \sum_{j=2}^{r} \xi_{j}^{2}\right)_{+}\right] \tag{7}
\end{equation*}
$$

controls the approximation ratio.
Proof. We can assume w.l.o.g. that $0<\rho<\min _{i \in[1, n]} \Sigma_{i i}$. This means that $B_{i}(X)=X^{1 / 2}\left(a_{i} a_{i}^{T}-\right.$ $\rho \mathbf{I}) X^{1 / 2}$ has rank $r$ and exactly one positive eigenvalue $\alpha_{i}$, with $\alpha_{i}=\operatorname{Tr} B_{i}(X)_{+}$for $i=1, \ldots, n$. We also denote by $-\beta_{j}^{i}$ for $j=2, \ldots, k$, the $(k-1)$ negative eigenvalues of $B_{i}(X)$. We follow the randomization procedure in [Ben-Tal and Nemirovski, 2002, El Ghaoui, 2006] and let $\xi$ denote normally distributed variables on $\mathbb{R}^{n}$, we have, using the rotational invariance of the normal distribution

$$
\mathbf{E}\left[\left(\xi^{T} B_{i}(X) \xi\right)_{+}\right]=\mathbf{E}\left[\left(\alpha_{i} \xi_{1}^{2}-\sum_{j=2}^{r} \beta_{j}^{i} \xi_{j}^{2}\right)_{+}\right], \quad \text { for } i=1, \ldots, n
$$

We then get

$$
\sum_{j=2}^{r} \beta_{j}^{i}=\operatorname{Tr}(B(X))_{+}-\operatorname{Tr}(B(X))=\alpha_{i}-\left(a_{i}^{T} X a_{i}-\rho\right) \leq \rho
$$

because $\lambda_{\max }\left(B_{i}(X)\right) \leq a_{i}^{T} X a_{i}$, hence

$$
\begin{aligned}
\mathbf{E}\left[\left(\xi^{T} B_{i}(X) \xi\right)_{+}\right] & \geq \min _{\beta}\left\{\mathbf{E}\left[\left(\alpha_{i} \xi_{1}^{2}-\sum_{j=2}^{r} \beta_{j}^{i} \xi_{j}^{2}\right)_{+}\right]: \sum_{j=2}^{r} \beta_{j}^{i} \leq \rho, \beta_{j}^{i} \geq 0\right\} \\
& =\mathbf{E}\left[\left(\alpha_{i} \xi_{1}^{2}-\frac{\rho}{r-1} \sum_{j=2}^{r} \xi_{j}^{2}\right)_{+}\right]
\end{aligned}
$$

by convexity and symmetry. By homogeneity and convexity, with $\psi(\rho)=\sum_{i=1}^{n} \alpha_{i}$, we then get

$$
\begin{aligned}
\mathbf{E}\left[\sum_{i=1}^{n}\left(\xi^{T} B_{i}(X) \xi\right)_{+}\right] & \geq \sum_{i=1}^{n} \mathbf{E}\left[\left(\alpha_{i} \xi_{1}^{2}-\frac{\rho}{r-1} \sum_{j=2}^{r} \xi_{j}^{2}\right)_{+}\right] \\
& \geq \mathbf{E}\left[\left(\psi(\rho) \xi_{1}^{2}-\frac{n \rho}{r-1} \sum_{j=2}^{r} \xi_{j}^{2}\right)_{+}\right]
\end{aligned}
$$

and we define $\vartheta_{r}(x)$ as $\operatorname{in}(7)$ above. Having shown

$$
\mathbf{E}\left[\sum_{i=1}^{n}\left(\xi^{T} B_{i}(X) \xi\right)_{+}\right] \geq n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right)
$$

and using $\mathbf{E}\left[\xi^{T} X \xi\right]=\operatorname{Tr}(X)=1$, we get

$$
\mathbf{E}\left[\sum_{i=1}^{n}\left(\xi^{T} B_{i}(X) \xi\right)_{+}\right] \geq n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right) \mathbf{E}\left[\xi^{T} X \xi\right]
$$

and this bound implies that there exists a nonzero $\xi$ such that

$$
\sum_{i=1}^{n}\left(\xi^{T} B_{i}(X) \xi\right)_{+} \geq n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right)\left(\xi^{T} X \xi\right)
$$

Suppose we set

$$
v_{i}= \begin{cases}1 & \text { if }\left(\xi^{T} B_{i}(X) \xi\right)>0 \\ 0 & \text { otherwise }\end{cases}
$$

we now have

$$
\xi^{T}\left(\sum_{i=1}^{n} v_{i} B_{i}(X)\right) \xi \geq n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right)\left(\xi^{T} X \xi\right)
$$

which is also, with $z=X^{1 / 2} \xi$ and $B_{i}(X)=X B_{i} X$

$$
z^{T}\left(\sum_{i=1}^{n} v_{i} B_{i}\right) z \geq n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right) z^{T} z
$$

This finally means that for our choice of $v$, with $B_{i}=a_{i} a_{i}^{T}-\rho$

$$
\phi(\rho)=\max _{v \in\{0,1\}^{n}} \lambda_{\max }\left(\sum_{i=1}^{n} u_{i} a_{i} a_{i}^{T}\right)-\rho \mathbf{C a r d}(u) \geq \lambda_{\max }\left(\sum_{i=1}^{n} v_{i} B_{i}\right) \geq n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right)
$$

hence our lower bound $n \rho \vartheta_{r}(\psi(\rho) / n \rho)<\phi(\rho)$ (which holds whenever $X$ is feasible in (5), i.e. whenever $X \succeq 0$ with $\operatorname{Tr} X=1$ ). Furthermore, if $X$ is an optimal solution of the relaxation in (5), we also get an upper bound on $\phi(\rho)$, with

$$
n \rho \vartheta_{r}\left(\frac{\psi(\rho)}{n \rho}\right) \leq \phi(\rho) \leq \psi(\rho)
$$

which is the desired result.

An explicit formula involving trigonometric integrals was derived in [El Ghaoui, 2006] (note that our definition for the function $\vartheta(x)$ is slightly different here). When $r$ is large, we can approximate $\vartheta_{r}(\cdot)$ by the function

$$
\begin{equation*}
\vartheta(x) \triangleq \mathbf{E}\left[\left(x \xi^{2}-1\right)_{+}\right] \tag{8}
\end{equation*}
$$

where $\xi \sim \mathcal{N}(0,1)$. Indeed, the central limit theorem shows that

$$
\frac{\sqrt{r}}{r-1} \sum_{j=2}^{r}\left(\xi_{i}^{2}-1\right) \stackrel{\mathcal{L}}{\rightarrow} \mathcal{N}(0,1)
$$

when $r$ grows to infinity. By convexity, we also have $\vartheta(x) \leq \vartheta_{r}(x)$. The function $\vartheta(\cdot)$ can be computed explicitly, with

$$
\begin{aligned}
\vartheta(x) & =\mathbf{E}\left[\left(x \xi^{2}-1\right)_{+}\right] \\
& =2 \int_{x^{-\frac{1}{2}}}^{\infty} x u^{2} \frac{e^{-u^{2} / 2}}{\sqrt{2 \pi}} d u-2 \mathcal{N}\left(-x^{-\frac{1}{2}}\right) \\
& =\frac{2 e^{-1 / 2 x}}{\sqrt{2 \pi x}}+2(x-1) \mathcal{N}\left(-x^{-\frac{1}{2}}\right)
\end{aligned}
$$

where $\mathcal{N}(\cdot)$ is the Gaussian cumulative density function. As with all results based on the central limit theorem, the approximation starts to be very good at relatively low values of $r$. The fact that $\vartheta^{\prime}\left(0^{+}\right)=0$ means that we cannot obtain a constant approximation ratio (à la MAXCUT). However, because $\vartheta(x)$ is
convex and increasing, we can still derive meaningful lower bounds in (6) if we can bound $\psi(\rho) / n \rho$ from below by a given $c>0$, with

$$
\begin{equation*}
\frac{\vartheta(c)}{c} \psi(\rho) \leq n \rho \vartheta\left(\frac{\psi(\rho)}{n \rho}\right) \leq \phi(\rho), \quad \text { when } c \leq \frac{\psi(\rho)}{n \rho} \tag{9}
\end{equation*}
$$

We also observe that $\psi\left(0^{+}\right)=\lambda_{\max }(\Sigma)>0$, and $\lim _{x \rightarrow \infty} \vartheta(x) / x=1$ mean that, when $n$ is fixed, we have

$$
\frac{\vartheta(\psi(\rho) / n \rho)}{\psi(\rho) / n \rho} \underset{\rho \rightarrow 0}{\longrightarrow} 1
$$

i.e. the approximation ratio converges to one as $\rho$ goes to zero (and solutions get less sparse). In fact

$$
\vartheta(x) / x=1-\sqrt{\frac{2}{\pi}} x^{-1 / 2}-x^{-1}+O\left(x^{-3 / 2}\right), \quad \text { as } x \rightarrow \infty
$$

We illustrate these last points by plotting $\vartheta(x)$ and $\vartheta(x) / x$ on the interval $[0,2]$ in Figure 1 , together with the function $\vartheta_{r}(x)$ for $r=5$.


Figure 1. Left: The functions $\vartheta(x)$ (solid black line), and $\vartheta_{5}(x)$ (dotted blue line). Right:
The functions $\vartheta(x) / x$ (solid black line), and $\vartheta_{5}(x) / x$ (dotted blue line).

Remark also that a naive lower bound on $\psi(\rho) / n \rho$, hence on the approximation ratio, can be obtained by plugging $X=\mathbf{I} / n$ in problem (5), which yields

$$
\begin{aligned}
\frac{\psi(\rho)}{n \rho} & \geq \frac{\sum_{i=1}^{n}\left(\Sigma_{i i}-\rho\right)_{+}}{n^{2} \rho} \\
& =\frac{\operatorname{Tr}(\Sigma)-\rho}{n^{2} \rho}
\end{aligned}
$$

because we have assumed that $\rho \leq \min _{i} \Sigma_{i i}$. This shows for example that when

$$
\rho \leq \frac{\operatorname{Tr}(\Sigma)}{n^{2}+1}
$$

then the approximation ration is greater than $\vartheta(1)$. We will see below that this range of values for $\rho$ is actually quite typical in detection problems where the matrix $\Sigma$ is Wishart.

## 4. DETECTION PROBLEMS

In this section, we focus on the problem of detecting the presence of a sparse leading component in a Gaussian model. It was shown in [Berthet and Rigollet, 2012] that the sparse eigenvalue statistic is minimax optimal in this setting. Computing sparse maximum eigenvalues is NP-Hard, but we show here that the relaxation detailed in the previous section achieve detection rates that are a multiple of the minimax optimum, in a high-dimensional setting where the ambient dimension $n$, the number of samples $m$ and the sparsity level $k$ all grow towards infinity proportionally. More specifically, we focus on the following hypothesis testing problem, where

$$
\begin{cases}\mathcal{H}_{0}: & x \sim \mathcal{N}\left(0, \mathbf{I}_{n}\right)  \tag{10}\\ \mathcal{H}_{1}: & x \sim \mathcal{N}\left(0, \mathbf{I}_{n}+\theta v v^{T}\right)\end{cases}
$$

where $\theta>0$ and $v \in \mathbb{R}^{n}$ is a sparse vector satisfying $\operatorname{Card}(v) \leq k^{*}$ and $\|v\|_{2}=1$. Given $m$ sample variables $x_{i} \in \mathbb{R}^{n}$, we let $\hat{\Sigma} \in \mathbf{S}_{n}$ be the sample covariance matrix, with

$$
\hat{\Sigma}=\frac{1}{m} \sum_{i=1}^{m} x_{i} x_{i}^{T}
$$

We will now seek to bound the value of the statistics $\phi(\rho)$ and $\psi(\rho)$ defined in (4) and (5) respectively, under the two hypotheses above.
4.1. The optimal statistic $\phi(\rho)$. We start by the easy part, namely bounding $\phi(\rho)$ from below under $\mathcal{H}_{1}$.

Proposition 4.1. Given $\hat{\Sigma} \in \mathbf{S}_{n}$ under $\mathcal{H}_{1}$, we have

$$
\begin{equation*}
\phi(\rho) \geq 1+\theta-\rho k^{*}-2(1+\theta) \sqrt{\frac{\log (1 / \delta)}{m}} \tag{11}
\end{equation*}
$$

with probability $1-\delta$.
Proof. By construction, we have $\phi(\rho) \geq \lambda_{\max }^{k^{*}}(\hat{\Sigma})-\rho k^{*}$, and [Berthet and Rigollet, 2012, Prop. 4.1] then yields the desired result.

Using again the results in [Berthet and Rigollet, 2012], we now show an upper bound on the value of the statistic $\phi(\rho)$ under $\mathcal{H}_{0}$.

Proposition 4.2. Given $\hat{\Sigma} \in \mathbf{S}_{n}$ under $\mathcal{H}_{0}$, and assuming $\rho \geq \Delta / m$, where

$$
\Delta=4 \log \left(9 e n / k^{*}\right)+4 \log (1 / \delta)
$$

we have

$$
\begin{equation*}
\phi(\rho) \leq 1+\frac{4 k^{*} \rho}{e \Delta}+\frac{1}{\rho m / \Delta-1} \tag{12}
\end{equation*}
$$

with probability $1-2 \delta$ when $\delta$ is small enough.
Proof. Under $\mathcal{H}_{0}$, [Berthet and Rigollet, 2012, Prop. 4.2] shows

$$
\operatorname{Prob}\left[\lambda_{\max }^{k}(\hat{\Sigma})-\rho k \geq 1-\rho k+4 \sqrt{t / m}+4 t / m\right] \leq\binom{ n}{k} 9^{k} e^{-t}
$$

and writing $1+u=1-\rho k+4 v / \sqrt{m}+4 v^{2} / m$, with $v=\sqrt{t}$, yields

$$
v=\frac{-\sqrt{m}+\sqrt{m(1+u+\rho k)}}{2}
$$

hence

$$
\operatorname{Prob}\left[\lambda_{\max }^{k}(\hat{\Sigma})-\rho k \geq 1+u\right] \leq\binom{ n}{k} 9^{k} e^{-v^{2}}
$$

Using the fact that $\phi(\rho)=\max _{k} \lambda_{\max }^{k}(\hat{\Sigma})-\rho k$ and $\binom{n}{k} \leq\left(\frac{e n}{k}\right)^{k}$ we then get, using union bounds

$$
\operatorname{Prob}[\phi(\rho) \geq 1+u] \leq \sum_{k=1}^{n} \exp \left(k \log \frac{9 e n}{k}-\frac{m}{4}(\sqrt{(1+u+\rho k)}-1)^{2}\right)
$$

We write

$$
k \log \frac{9 e n}{k}-\frac{m}{4}(\sqrt{(1+u+\rho k)}-1)^{2}=k \log \frac{9 e n}{k^{*}}-\frac{m}{4}(\sqrt{(1+u+\rho k)}-1)^{2}+\left(\log k^{*}-\log k\right) k .
$$

When $\alpha<1$, the function

$$
(\sqrt{1+x}-1)^{2}-\alpha x
$$

is convex and reaches its minimum when $x=1 /(1-\alpha)^{2}-1$, with value $-\alpha^{2} /(1-\alpha)$. A similar argument shows that $\left(\log k^{*}-\log k\right) k \leq k^{*} / e$. Setting $x=u+\rho k$ we get

$$
\alpha u+\alpha \rho k-(\sqrt{1+u+\rho k}-1)^{2} \leq \alpha^{2} /(1-\alpha)
$$

Setting $\alpha=\Delta / \rho m$ above, imposing $\rho>\Delta / m$, and

$$
\alpha u=\alpha^{2} /(1-\alpha)+\frac{4 k^{*}}{e m}
$$

we can ensure

$$
\frac{m}{4}\left(\frac{\Delta}{m} k+\left(\log k^{*}-\log k\right) \frac{4 k}{m}-(\sqrt{(1+u+\rho k)}-1)^{2}\right) \leq 0
$$

for all $k \geq 1$, hence

$$
k \log \frac{9 e p}{k^{*}}-\frac{m}{4}(\sqrt{(1+u+\rho k)}-1)^{2}+\left(\log k^{*}-\log k\right) k \leq-k \log (1 / \delta)
$$

and

$$
\operatorname{Prob}[\phi(\rho) \geq 1+u] \leq \frac{\delta}{1-\delta}
$$

which yields the desired result.
We now use these last two results to determine the minimum signal level $\theta$ which can be detected using the statistic $\phi(\rho)$. We define the following levels

$$
\left\{\begin{array}{l}
\tau_{0}=1+\sqrt{\frac{k^{*}(\Delta+4 / e)}{m}}+\frac{4 k^{*}}{e m} \tag{13}
\end{array}+\frac{4}{e \Delta} \sqrt{\frac{k^{*} \Delta}{m(1+4 /(e \Delta))}} .\right.
$$

for some $\gamma>0$. Given $\hat{\Sigma} \in \mathbf{S}_{n}$ and $\tau_{\phi} \in\left[\tau_{0}, \tau_{1}\right]$, the corresponding test is given by

$$
\begin{equation*}
\mathbf{1}_{\left\{\phi(\rho)>\tau_{\phi}\right\}} \tag{14}
\end{equation*}
$$

The following proposition shows that if $\theta$ is high enough, then this test discriminates between $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$ with probability $1-3 \delta$.

Proposition 4.3. Suppose we set

$$
\begin{equation*}
\Delta=4 \log \left(9 e n / k^{*}\right)+4 \log (1 / \delta) \quad \text { and } \quad \rho=\frac{\Delta}{m}+\frac{\Delta}{\sqrt{k^{*} m(\Delta+4 / e)}} \tag{15}
\end{equation*}
$$

and define $\theta_{\phi}$ such that

$$
\begin{equation*}
\theta_{\phi}=\left(2 \sqrt{\frac{k^{*}(\Delta+4 / e)}{m}}+\frac{k^{*}(\Delta+4 / e)}{m}+2 \sqrt{\frac{\log (1 / \delta)}{m}}\right)\left(1-2 \sqrt{\frac{\log (1 / \delta)}{m}}\right)^{-1} \tag{16}
\end{equation*}
$$

then if $\theta>\theta_{\phi}$ in the Gaussian model (10), the test statistic (14) based on $\phi(\rho)$ discriminates between $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$ with probability $1-3 \delta$.

Proof. If $\theta_{\phi}$ is set as in (16), setting $\rho$ as in (15) means $\Delta / m \rho<1$, we have $\tau_{0} \leq \tau_{1}$ and propositions 4.2 and 4.1 show that (14) discriminates between $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$.

This detection level was shown to be minimax optimal in [Berthet and Rigollet, 2012]. This is not surprising, since the statistic $\phi(\rho)$ is simply a penalized formulation of $\lambda_{\max }^{k}(\cdot)$ which was shown to reach a similar detection level in [Berthet and Rigollet, 2012]. Both $\phi(\rho)$ and $\lambda_{\max }^{k}(\cdot)$ are intractable however, and we will now focus on an efficiently computable statistic based on $\psi(\rho)$.
4.2. The tractable statistic $\psi(\rho)$. We can directly use proposition 4.1 to produce a lower bound on $\psi(\rho)$ and on the approximation ratio under both $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$.

Corollary 4.4. Setting $\Delta$ as in (15), under both $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$, we have

$$
\psi(\rho) \geq 1-\rho k^{*}-2 \sqrt{\frac{\log (1 / \delta)}{m}}
$$

with probability $1-\delta$.
Proof. We simply set $\theta=0$ in proposition (4.1) and use the fact that $\phi(\rho) \leq \psi(\rho)$ by construction.
We are now ready to prove the main result of this section, showing that in a high dimensional setting, the tractable statistic $\psi(\rho)$ discriminates between $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$ when $\theta \geq \theta_{\psi}$, where $\theta_{\psi}$ is comparable to $\theta_{\phi}$, and $\theta_{\psi}$ is independent of $n$. The approximation ratio in (9) is controlled by $n \rho / \psi(\rho)$ which depends on $\rho$ so we cannot explicitly minimize the detection level in $\rho$ as we did above. Instead, we will control the quality of the approximation of $\phi(\rho)$ by $\psi(\rho)$ for the value of $\rho$ used in computing $\theta_{\phi}$. We suppose $n=\mu m$ and $k^{*}=\kappa n$, where $\mu>0$ and $\kappa \in(0,1)$. Setting $\rho$ as in (15), we get

$$
n \rho=\mu \Delta+\frac{\mu \Delta}{\sqrt{\kappa(\Delta+4 / e)}}
$$

with Corollary 4.4 implying

$$
\psi(\rho) \geq 1-\mu \Delta \kappa-\frac{\sqrt{\mu \kappa}}{\sqrt{(\Delta+4 / e)}}-2 \sqrt{\frac{\log (1 / \delta)}{m}}
$$

This means that the approximation ratio in (9) is bounded below by $\beta(\mu, \kappa)$, with

$$
\begin{equation*}
\beta(\mu, \kappa)=\frac{\vartheta(c)}{c} \quad \text { where } \quad c=\frac{1-\mu \Delta \kappa-\frac{\sqrt{\mu \kappa}}{\sqrt{(\Delta+4 / e)}}-2 \sqrt{\frac{\log (1 / \delta)}{m}}}{\mu \Delta+\frac{\mu \Delta}{\sqrt{\kappa(\Delta+4 / e)}}} \tag{17}
\end{equation*}
$$

Given $\hat{\Sigma} \in \mathbf{S}_{n}$ and $\tau_{\psi} \in\left[\beta(\mu, \kappa)^{-1} \tau_{0}, \tau_{1}\right]$, where $\tau_{0}$ and $\tau_{1}$ are defined in (13), the corresponding test is then

$$
\begin{equation*}
\mathbf{1}_{\left\{\psi(\rho)>\tau_{\psi}\right\}} \tag{18}
\end{equation*}
$$

with $\rho$ set as in (15). The following proposition shows that if $\theta$ is high enough, then this test discriminates between $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$ with probability $1-3 \delta$.
Theorem 4.5. suppose $n=\mu m$ and $k^{*}=\kappa n$, where $\mu>0$ and $\kappa \in(0,1)$ are fixed and $n$ is large. Define the detection threshold $\theta_{\psi}$ such that

$$
\begin{equation*}
\theta_{\psi} \geq \beta(\mu, \kappa)^{-1} \theta_{\phi} \tag{19}
\end{equation*}
$$

where $\beta(\mu, \kappa)$ is defined in (17) and $\theta_{\phi}$ is defined in (16), then if $\theta>\theta_{\psi}$ in the Gaussian model (10) the test statistic (18) based on $\psi(\rho)$ discriminates between $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$ with probability $1-3 \delta$.

Proof. Having bounded the approximation ratio $\beta(\mu, \kappa)$ defined in (17), the result follows from (9) and Proposition 4.3.

In Figure 2, we plot the level sets of $\beta(\mu, \kappa)$ for $\Delta=5$. Observe that whenever $\mu$ is small enough, $\beta(\mu, \kappa)>0$ for all values of $\kappa \in(0,1)$ and the approximation ratio converges to one as $\mu$ goes to zero.

This means that the detection threshold $\theta$ of the statistic $\psi(\rho)$ remains finite when $n$ goes to infinity in the proportional regime. By contrast, the detection threshold of the MDP statistic in [Berthet and Rigollet, 2012] blows up to infinity as soon when $k$ goes to infinity in this scenario.


Figure 2. Level sets of $\beta(\mu, \kappa)$ for $\Delta=5$.
4.3. Detection thresholds for $\psi(\rho)$ and $\lambda_{\max }(\cdot)$. The fact that $\psi(0)=\lambda_{\max }(\Sigma)$ means the statistic $\psi(\rho)$ should always perform better than $\lambda_{\max }(\cdot)$ for at least some values of $\rho$. However, we notice in Figure 2 that $\beta(\mu, \kappa)$ goes to zero as $\kappa$ goes to zero, which is a direct consequence of our choice of $\rho$ in (15). Our choice of $\rho$ is optimal for the statistic $\phi(\rho)$ but not for $\psi(\rho)$ and the main issue here is that we cannot explicitly maximize the difference $\tau_{1}-\beta(\mu, \kappa)^{-1} \tau_{0}$ as a function of $\rho$. On the other hand, it is easy to show that a better guess for $\rho$, when both $\kappa$ and $\mu$ are small, is to pick

$$
\begin{equation*}
\rho=\frac{1}{n}, \tag{20}
\end{equation*}
$$

in which case, one can show that the detection threshold for $\theta_{\phi}$ becomes

$$
\begin{equation*}
\theta_{\phi}=\left(\left(1+\frac{4}{e \Delta}\right) \kappa+\frac{\mu \Delta}{1-\mu \Delta}+2 \sqrt{\frac{\log (1 / \delta)}{m}}\right)\left(1-2 \sqrt{\frac{\log (1 / \delta)}{m}}\right)^{-1} \tag{21}
\end{equation*}
$$

while the approximation ratio $\beta(\mu, \kappa)$ is given by $\vartheta(\psi(1 / n))$ which is of order one. This means that the detection threshold for $\psi(\rho)$ is controlled by

$$
\left(1+\frac{4}{e \Delta}\right) \kappa+\frac{\mu \Delta}{1-\mu \Delta} \simeq\left(1+\frac{4}{e \Delta}\right) \kappa+\mu \Delta
$$

when both $\kappa$ and $\mu \Delta$ are small. On the other hand, [Benaych-Georges et al., 2011] show that, in our regime, the statistic $\lambda_{\max }(\cdot)$ can only distinguish between $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$ when $\theta$ is larger than

$$
\sqrt{\mu}+\mu,
$$

which means that even for our suboptimal choice of $\rho$, the statistic $\psi(\rho)$ outperforms $\lambda(\cdot)$ by a factor $\Delta \sqrt{\mu}$.

## 5. Algorithms

The approximation performance we studied in the previous section comes at a price. While the semidefinite program (5) is tractable, its complexity is significantly higher than that of the simpler relaxations derived in [d'Aspremont et al., 2007], and very significantly higher than the MDP statistic in [Berthet and Rigollet, 2012] for example. d'Aspremont et al. [2008] derived greedy algorithms to produce good upper bounds on $\psi(\rho)$ from approximate dual solutions to problem (5), but there are of course no guarantees on the
quality of their output. In this section, we describe a simple algorithm to compute $\psi(\rho)$, with comparatively low storage and iteration costs.

Recall from [d'Aspremont et al., 2008] that the dual to problem (5) is written

$$
\begin{array}{ll}
\operatorname{minimize} & \lambda_{\max }\left(\sum_{i=1}^{n} Y_{i}\right) \\
\text { subject to } & Y_{i} \succeq a_{i} a_{i}^{T}-\rho \mathbf{I}  \tag{22}\\
& Y_{i} \succeq 0, \quad i=1, \ldots, n
\end{array}
$$

in the variables $Y_{i} \in \mathbf{S}_{n}$, where $\rho>0$ and $a_{i} \in \mathbb{R}^{n}$ are defined in Section 2 . We first show how to regularize this problem, then discuss how to solve the regularized instance using a Frank-Wolfe type algorithm.
5.1. Smoothing. Problem (22) is not smooth but we can write

$$
\lambda_{\max }(Y)=\max _{\operatorname{Tr}(X)=1, X \succeq 0} \operatorname{Tr}(Y X)
$$

and as in [Ben-Tal and Nemirovski, 2005], a natural way to regularize $\lambda_{\max }(\cdot)$ is to add a (strongly convex) matrix entropy penalty to this variational formulation. We also add an explicit lower bound on the eigenvalues of $X$ to ensure that the gradient matrix $X$ is invertible and well conditioned. We summarize this in the next lemma.

Lemma 5.1. Let $\epsilon>0$, the function

$$
\begin{equation*}
f(Y) \triangleq \max _{\substack{\operatorname{Tr}(X)=1, X \succeq(\epsilon / n) \mathbf{I}}} \operatorname{Tr}(Y X)-\frac{\epsilon}{\log n}(\operatorname{Tr}(X \log (X))+\log n) \tag{23}
\end{equation*}
$$

has a Lipschitz continuous gradient with constant

$$
L_{f} \leq \frac{\log n}{\epsilon}
$$

with respect to the trace norm and satisfies

$$
(1-\epsilon) \lambda_{\max }(Y)-\epsilon \leq f(Y) \leq \lambda_{\max }(Y)
$$

for all $Y \in \mathbf{S}_{n}$.
Proof. In the spectahedron setting, we know from [Ben-Tal and Nemirovski, 2005] that the matrix entropy function

$$
d(X)=\operatorname{Tr}(X \log (X))+\log n
$$

is strongly convex with parameter $1 / 2$ with respect to the trace norm (the dual of the spectral norm) and satisfies

$$
\max _{\operatorname{Tr}(X)=1, X \succeq 0} d(X) \leq \log n
$$

Then, [Nesterov, 2005, Thm. 1] shows that $\nabla f(Y)$ is Lipschitz continuous with constant $\log n / \epsilon$ with respect to the trace norm. By construction, we have $f(Y) \leq \lambda_{\max }(Y)$ and

$$
(1-\epsilon) \lambda_{\max }(Y)-\epsilon \leq \max _{\substack{\operatorname{Tr}(X)=1, X \succeq(\epsilon / n) \mathbf{I}}} \operatorname{Tr}(Y X)-\epsilon \leq f(Y)
$$

hence the desired result.

We can thus form a smooth approximation of problem (22), written

$$
\begin{array}{ll}
\operatorname{minimize} & f\left(\sum_{i=1}^{n} Y_{i}\right) \\
\text { subject to } & Y_{i} \succeq a_{i} a_{i}^{T}-\rho \mathbf{I}  \tag{25}\\
& Y_{i} \succeq 0, \quad i=1, \ldots, n
\end{array}
$$

in the variables $Y_{i} \in \mathbf{S}_{n}$, where $f$ is the smooth approximation of the function $\lambda_{\max }(\cdot)$ defined in (23) and solve this problem using Algorithm 1.
5.2. Complexity. We first show how to efficiently compute both $f(Z)$ and $\nabla f(Z)$.

Lemma 5.2. Assume $Z=V \operatorname{diag}(y) V^{T}$ and suppose $\lambda \in \mathbb{R}$ solves

$$
\begin{equation*}
\min _{\lambda} \sum_{i=1}^{n} \min \left\{\frac{\epsilon\left(y_{i}+\lambda\right)}{n}-\frac{\beta \epsilon}{n} \log \frac{\epsilon}{n}, \beta e^{\frac{y_{i}+\lambda}{\beta}-1}\right\} \tag{26}
\end{equation*}
$$

then $Y=V \operatorname{diag}(x) V^{T}$ solves the maximization problem in (23), where

$$
x_{i}=\max \left\{\frac{\epsilon}{n}, e^{\frac{y_{i}+\lambda}{\beta}-1}\right\}, \quad i=1, \ldots, n
$$

and we have $\nabla f(Z)=Y$.
Proof. The function $f(Z)$ is spectral so solving (23) is equivalent to solving

$$
\max _{\substack{1^{T} x=1, x \geq(\epsilon / n)}} y^{T} x-\frac{\epsilon}{\log n} \sum_{i=1}^{n} x_{i} \log x_{i}
$$

whose dual is

$$
\min _{\lambda} \max _{x \geq \epsilon / n}(y-\lambda \mathbf{1})^{T} x-\frac{\epsilon}{\log n} \sum_{i=1}^{n} x_{i} \log x_{i}-\lambda
$$

which is equivalent to (26). Now, the fact that $f(Z)$ is a maximum of affine functions of $Z$ shows that $\nabla f(Z)=Y$.

Besides computing the gradient, the main cost at each iteration is the problem of solving the $n$ subproblems (24). We will see that when $\nabla f(Z)$ is positive definite, then this can be done in closed form, with complexity $O\left(n^{2} \log n\right)$. Furthermore, the matrices $Y_{i}$ do not need to be stored, only their sum is required.
Lemma 5.3. Given $X \in \mathbf{S}_{n}$ such that $X \succ 0$, together with $B_{i}=a_{i} a_{i}^{T}-\rho \mathbf{I}$ for some $\rho>0$, we have

$$
\min _{\substack{Y \succeq B_{i} \\ Y \succeq 0}} \operatorname{Tr}(X Y)=\operatorname{Tr}\left(X^{1 / 2} B_{i} X^{1 / 2}\right)_{+}
$$

and the optimal solution has rank one and is given by

$$
\begin{cases}Y=X^{-1 / 2} v v^{T} X^{1 / 2} B_{i}, & \text { if } \rho<\left\|a_{i}\right\|_{2}^{2} \\ 0 & \text { otherwise }\end{cases}
$$

where $v$ is the leading eigenvector of the matrix $X^{1 / 2} B_{i} X^{1 / 2}$.

```
Algorithm 1 Frank-Wolfe algorithm for computing \(\psi(\rho)\).
Input: \(\rho>0\) and a feasible starting point \(Z_{0}\).
    : for \(k=1\) to \(N_{\max }\) do
        Compute \(X=\nabla f(Z)\), together with \(X^{-1}\) and \(X^{1 / 2}\).
    : Solve the \(n\) subproblems
\[
\begin{array}{ll}
\operatorname{minimize} & \operatorname{Tr}\left(Y_{i} X\right) \\
\text { subject to } & Y_{i} \succeq a_{i} a_{i}^{T}-\rho \mathbf{I}  \tag{24}\\
& Y_{i} \succeq 0,
\end{array}
\]
```

in the variables $Y_{i} \in \mathbf{S}_{n}$ for $i=1, \ldots, n$.
Compute $W=\sum_{i=1}^{n} Y_{i}$.
5: Update the current point, with

$$
Z_{k}=\left(1-\frac{2}{k+2}\right) Z_{k-1}+\frac{2}{k+2} W
$$

end for
Output: A matrix $Z \in \mathbf{S}_{n}$.

Proof. Recall from d'Aspremont et al. [2008] that

$$
\begin{aligned}
\operatorname{Tr}\left(X^{1 / 2} B_{i} X^{1 / 2}\right)_{+} & =\max _{\{0 \preceq P \preceq X\}} \operatorname{Tr}\left(P B_{i}\right) \\
& =\min _{\{Y \succeq B, Y \succeq 0\}} \operatorname{Tr}(Y X)
\end{aligned}
$$

and a solution to the dual of problem (24) can be obtained from the solution to

$$
\max _{\{0 \preceq Q \preceq \mathbf{I}\}} \operatorname{Tr}\left(Q X^{1 / 2} B_{i} X^{1 / 2}\right) .
$$

By Sylvester's theorem, when $\rho<\left\|a_{i}\right\|_{2}^{2}$, the matrix $X^{1 / 2} B_{i} X^{1 / 2}$ has exactly one nonnegative eigenvalue, so the optimal solution to this last problem is $Q=v v^{T}$ where $v$ is the leading eigenvector of the matrix $X^{1 / 2} B_{i} X^{1 / 2}$. This means that the optimal dual solution is $P=X^{1 / 2} v v^{T} X^{1 / 2}$. Finally, the KKT optimality conditions impose $X Y=P B_{i}$, which together with $X \succ 0$ yields the desired result.

The last lemma shows that solving problem (24) requires the following steps. Assume $X^{-1}$ and $X^{1 / 2}$ have been precomputed, we first form the matrix $X^{1 / 2} B_{i} X^{1 / 2}$ at cost $O\left(n^{2}\right)$. We then compute its leading eigenvector at cost $O\left(n^{2} \log n\right)$ and form the rank one matrix $P=X^{1 / 2} v v^{T} X^{1 / 2}$ at cost $O\left(n^{2}\right)$. Because $P$ is rank one, computing $X^{-1} P B_{i}$ also costs $O\left(n^{2}\right)$. This means that the total cost of solving problem (24) is bounded by $O\left(n^{2} \log n\right)$. Furthermore, without loss of generality, we can restrict the matrices $Y_{i}$ to have norm less than

$$
B=\left\|\sum_{i=1}^{n} B_{i}^{+}\right\|_{2}
$$

which means we can assume the feasible set of problem (25) is compact.
Proposition 5.4. Assume $A \in \mathbb{R}^{n \times n}$ and $\rho>0$, Algorithm 1 will produce and $\epsilon$ solution to problem (25) in

$$
O\left(\frac{D^{2} \log n}{\epsilon^{2}}\right)
$$

iterations, where $D$ is the trace norm diameter of the feasible set

$$
D=\operatorname{diam}\left\{\sum_{i=1}^{n} Y_{i}: Y_{i} \succeq B, Y_{i} \succeq 0,\left\|Y_{i}\right\|_{2} \leq B\right\}
$$

Each iteration has complexity $O\left(n^{3} \log n\right)$ and storage cost $O\left(n^{2}\right)$.
Proof. We use the complexity bounds in [Frank and Wolfe, 1956, Clarkson, 2010, Jaggi, 2011] for example and Lemma (5.1) to control the curvature of $f$.

## 6. NUMERICAL EXPERIMENTS

We test the detection procedure based on $\psi(\rho)$ described in (18). We generate 3000 experiments, where $m$ points $x_{i} \in \mathbb{R}^{n}$ are sampled under both hypotheses, with

$$
\begin{cases}\mathcal{H}_{0}: & x \sim \mathcal{N}\left(0, \mathbf{I}_{n}\right) \\ \mathcal{H}_{1}: & x \sim \mathcal{N}\left(0, \mathbf{I}_{n}+\theta v v^{T}\right)\end{cases}
$$

as in (10). In each experiment, we pick the leading dimension $n=100$, the number of samples $m=50$ and the cardinality $k=20$. We set $\theta=3, v_{i}=1 / \sqrt{k}$ when $i \in[1, k]$ and zero otherwise. In Figure 3 we plot the distributions of the test statistic $\psi(\rho)$ defined in (18), the MDP statistic in [Berthet and Rigollet, 2012], the $\lambda_{\max }(\cdot)$ statistic and the diagonal statistic in [Amini and Wainwright, 2009]. As in [Berthet and Rigollet, 2012], we observe that all the test statistics perform very similarly, except for the diagonal test.


Figure 3. Distribution of test statistic $\psi(\rho)$ (top left), the $M D P_{k}$ statistic in [Berthet and Rigollet, 2012] (top right), the $\lambda_{\max }(\cdot)$ statistic (bottom left) and the diagonal statistic from [Amini and Wainwright, 2009] (bottom right) under both $\mathcal{H}_{0}$ and $\mathcal{H}_{1}$. All experiments are performed on random Gaussian matrices with ambient dimension $n=100, m=50$ samples and cardinality for $v$ under $\mathcal{H}_{1}$ set to $k=20$.

## Acknowledgments

AA would like to acknowledge partial support from NSF grants SES-0835550 (CDI), CMMI-0844795 (CAREER), CMMI-0968842, a starting grant from the European Research Council (project SIPA), a Peek junior faculty fellowship, a Howard B. Wentz Jr. award and a gift from Google. The authors would like to thank Philippe Rigollet and Quentin Berthet for very constructive discussions and introducing them to the detection problem in Section 4.

## References

A.A. Amini and M. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. The Annals of Statistics, 37(5B):2877-2921, 2009.
A. Ben-Tal and A. Nemirovski. On tractable approximations of uncertain linear matrix inequalities affected by interval uncertainty. SIAM Journal on Optimization, 12(3):811-833, 2002.
A. Ben-Tal and A. Nemirovski. Non-Euclidean restricted memory level method for large-scale convex optimization. Mathematical Programming, 102(3):407-456, 2005.
F. Benaych-Georges, A. Guionnet, and M. Maida. Fluctuations of the extreme eigenvalues of finite rank deformations of random matrices. Electron. J. Probab., 16:no. 60, 1621-1662, 2011. ISSN 1083-6489. doi: 10.1214/EJP. v16-929. URL http://dx.doi.org/10.1214/EJP.v16-929.
Q. Berthet and P. Rigollet. Optimal detection of sparse principal components in high dimension. Arxiv preprint arXiv:1202.5070, 2012.
E. Candes and T. Tao. The Dantzig selector: statistical estimation when p is much larger than n . Annals of Statistics, 35(6):2313-2351, 2007.
K.L. Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. ACM Transactions on Algorithms (TALG), 6(4):63, 2010.
A. d'Aspremont, L. El Ghaoui, M.I. Jordan, and G. R. G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. SIAM Review, 49(3):434-448, 2007.
A. d'Aspremont, F. Bach, and L. El Ghaoui. Optimal solutions for sparse principal component analysis. Journal of Machine Learning Research, 9:1269-1294, 2008.
L. El Ghaoui. On the quality of a semidefinite programming bound for sparse principal component analysis. ArXiV math.OC/0601448, 2006.
M. Frank and P. Wolfe. An algorithm for quadratic programming. Naval research logistics quarterly, 3(1-2):95-110, 1956.
M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. J. ACM, 42:1115-1145, 1995.
M. Jaggi. Convex optimization without projection steps. Arxiv preprint arXiv:1108.1170, 2011.
M. Journée, F. Bach, P.A. Absil, and R. Sepulchre. Low-rank optimization for semidefinite convex problems. Arxiv preprint arXiv:0807.4423, 2008.
M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. arXiv:0811.4724, 2008.
M. Mezard and A. Montanari. Information, physics, and computation. Oxford University Press, USA, 2009.
M. Mezard, G. Parisi, and M.A. Virasoro. Spin glass theory and beyond, volume 9. World scientific Singapore, 1987.
Y. Nesterov. Smooth minimization of non-smooth functions. Mathematical Programming, 103(1):127-152, 2005.
D. Spielman and S.H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. In Proceedings of the thirty-third annual ACM symposium on Theory of computing, pages 296-305. ACM, 2001.
M. Talagrand. Mean Field Models for Spin Glasses: Basic examples, volume 1. Springer Verlag, 2010.
U. Zwick. Outward rotations: a tool for rounding solutions of semidefinite programming relaxations, with applications to max cut and other problems. In Proceedings of the thirty-first annual ACM symposium on Theory of computing, pages 679-687. ACM, 1999.

In alphabetical order.
CMAP, Ecole Polytechnique, UMR CNRS 7641, Palaiseau, France.
E-mail address: alexandre.daspremont@m4x.org
INRIA, SIERRA Project-Team,
Laboratoire d'Informatique de l'École Normale Supérieure, Paris, France.
E-mail address: francis.bach@ens.fr
EECS Department, U.C. Berkeley, Berkeley, CA 94720, USA.
E-mail address: elghaoui@eecs.berkeley.edu


[^0]:    Date: June 19, 2012.
    2010 Mathematics Subject Classification. 62H25, 90C22, 90C27.
    Key words and phrases. Sparse PCA, convex relaxation, semidefinite programming, approximation bounds, detection.

