Proceedings of Machine Learning Research vol 145:1–19, 2021 2nd Annual Conference on Mathematical and Scientific Machine Learning

Average-Case Integrality Gap for Non-Negative Principal Component Analysis

Afonso S. Bandeira BANDEIRA@MATH.ETHZ.CH
Department of Mathematics, ETH Zürich

Dmitriy Kunisky KUNISKY@CIMS.NYU.EDU
Department of Mathematics, Courant Institute of Mathematical Sciences, NYU

Alexander S. Wein AWEIN@CIMS.NYU.EDU
Department of Mathematics, Courant Institute of Mathematical Sciences, NYU

Editors: Joan Bruna, Jan S Hesthaven, Lenka Zdeborova

Abstract

Montanari and Richard (2015) asked whether a natural semidefinite programming (SDP) relaxation can effectively optimize $\boldsymbol{x}^\top \boldsymbol{W}\boldsymbol{x}$ over $\|\boldsymbol{x}\| = 1$ with $x_i \geq 0$ for all coordinates i, where $\boldsymbol{W} \in \mathbb{R}^{n\times n}$ is drawn from the Gaussian orthogonal ensemble (GOE) or a spiked matrix model. In small numerical experiments, this SDP appears to be *tight* for the GOE, producing a rank-one optimal matrix solution aligned with the optimal vector $\boldsymbol{x}$. We prove, however, that as $n \to \infty$ the SDP is not tight, and certifies an upper bound asymptotically no better than the simple spectral bound $\lambda_{\max}(\boldsymbol{W})$ on this objective function. We also provide evidence, using tools from recent literature on hypothesis testing with low-degree polynomials, that no subexponential-time certification algorithm can improve on this behavior. Finally, we present further numerical experiments estimating how large n would need to be before this limiting behavior becomes evident, providing a cautionary example against extrapolating asymptotics of SDPs in high dimension from their efficacy in small "laptop scale" computations.

Keywords: semidefinite programming, principal component analysis, certification algorithms, average-case computational complexity

1. Introduction

Recovering the most significant directions or *principal components* of a matrix from noisy observations is a fundamental problem in both mathematical statistics and applications (Rao, 1964; Johnstone, 2001; Ringnér, 2008; Abdi and Williams, 2010). The asymptotics of this task have been studied at length by analyzing idealized *spiked matrix models*. These models present simplified settings where algorithms for extracting principal components having various structures may be theoretically evaluated. One important algorithmic strategy that has been evaluated in this way is *convex relaxation* of an objective function associated to extracting principal components, in particular using *semidefinite programs (SDPs)*.

The theory of convex relaxations for such problems presents two diverging phenomena: on the one hand, for some large random problems, SDPs or other relaxations are *tight*, recovering exact solutions. This behavior has supported their application to problems such as compressed sensing (Candes and Tao, 2005), matrix completion (Candès and Tao, 2010), community detection (Abbe et al., 2015; Hajek et al., 2016), geometric processing of point clouds (Bandeira et al., 2014b; Ozye-

sil, 2014; Chaudhury et al., 2015), multireference alignment (Bandeira et al., 2014a; Bendory et al., 2017), and quadratic assignment (Ferreira et al., 2018). On the other hand, for other large random problems, a recent line of work suggests that SDPs are far from tight, and do not even achieve an objective value closely approximating that of the problem they were designed to estimate. Moreover, these results suggest that any algorithm following the relaxation strategy of producing a bound on these optimization problems (perhaps as a first step towards estimating the solution by rounding the relaxation) will fail unless they have access to massive computational resources of size nearly-exponential in the problem dimension. These problems include finding large cuts, multisections, and colorings in random graphs (Bandeira et al., 2020a) and numerous random constrained principal component analysis (PCA) problems (Bandeira et al., 2020b).

Given this stark contrast between the power of convex relaxation in some settings and its severe limitations in other settings, it is of both practical and theoretical interest to understand what problems fall into which category. The purpose of this paper is to analyze a model of *non-negative PCA*, where we search for principal components with non-negative entries, from this perspective. Non-negative PCA has found applications, for example, in genomics and cell biology (Lazzeroni and Owen, 2002; Dordek et al., 2016), and is a mathematically-elegant instance of constrained PCA. This particular variant of PCA is especially interesting in the context of the performance of convex relaxations because the previous work of Montanari and Richard (2015) described numerical results concerning the performance of a natural SDP. We will show that, surprisingly, for sufficiently large problems, the natural conjecture from those numerics—that the SDP is typically successful—is incorrect, and on the contrary both the SDP and any algorithm following a similar relaxation strategy must suffer from the same limitations as the latter class of problems mentioned above.

We first review spiked matrix models and the associated "ordinary" PCA problems without extra constraints. While the first such models, proposed by Johnstone (2001), concerned Gaussian observations from a covariance matrix deformed by adding a rank-one "spike," the following simpler *Wigner spiked matrix model* captures much of the same phenomenology. We consider the following two probability distributions $\mathbb{P}$ and $\mathbb{Q}$ over $n \times n$ symmetric matrices.

- Under $\mathbb{Q} = \mathbb{Q}_n$, observe $\boldsymbol{W} \in \mathbb{R}^{n \times n}_{\mathsf{sym}}$ drawn from the *Gaussian orthogonal ensemble (GOE)*, meaning $W_{ii} \sim \mathcal{N}(0, \frac{1}{n})$ and $W_{ij} = W_{ji} \sim \mathcal{N}(0, \frac{2}{n})$ for $i < j$ with all $\frac{n(n+1)}{2}$ of these entries distributed independently. We also write $\mathbb{Q} = \mathsf{GOE}(n)$ for this distribution.
- Under $\mathbb{P} = \mathbb{P}_n$, first draw $\boldsymbol{u} \sim \mathsf{Unif}(\mathbb{S}^{n-1})$ for $\mathbb{S}^{n-1} \subset \mathbb{R}^n$ the sphere of unit radius, and then observe $\boldsymbol{W} = \boldsymbol{W}_0 + \beta \boldsymbol{u}\boldsymbol{u}^\top$ for $\boldsymbol{W}_0 \sim \mathsf{GOE}(n)$ and some fixed $\beta > 0$, held constant as $n \to \infty$.

Two natural statistical questions arise: (1) *detection* or *testing*, where we observe $\boldsymbol{W}$ drawn from $\mathbb{P}$ or $\mathbb{Q}$ and must decide which distribution $\boldsymbol{W}$ was drawn from, and (2) *recovery* or *estimation*, where we observe $\boldsymbol{W} \sim \mathbb{P}$ and seek to produce a good estimate of $\boldsymbol{u}$. In either case, the associated optimization problem of computing the largest eigenvalue is natural to consider:

$$\lambda_{\max}(\boldsymbol{W}) := \max_{\|\boldsymbol{x}\|_2 = 1} \boldsymbol{x}^\top \boldsymbol{W} \boldsymbol{x}. \tag{1}$$

For recovery, computing the maximizer $\boldsymbol{x}^\star$ (the top eigenvector of $\boldsymbol{W}$) performs maximum likelihood estimation of $\boldsymbol{u}$. For detection, computing and thresholding $\lambda_{\max}(\boldsymbol{W})$ itself is a natural and often effective strategy. The optimal value and optimizer of $\lambda_{\max}(\boldsymbol{W})$ can be approximated (to

arbitrary accuracy) in time $\mathsf{poly}(n)$, so these correspond to efficient algorithms for recovery and detection, respectively. Moreover, these algorithms are essentially *optimal*: almost whenever[1] it is possible to distinguish $\mathbb{P}$ from $\mathbb{Q}$ with high probability,[2] thresholding $\lambda_{\max}(\boldsymbol{W})$ achieves this; whenever it is possible to estimate $\boldsymbol{u}$ non-trivially then the optimizer $\boldsymbol{x}^\star$ achieves this.

Proposition 1 *Let $\boldsymbol{W} \sim \mathbb{P}$ with parameter $\beta \geq 0$ (note that taking $\beta = 0$ gives $\mathbb{P} = \mathbb{Q}$), and let $\boldsymbol{u}$ be the spike vector. Let $\boldsymbol{x}^\star(\boldsymbol{W})$ be the optimizer of $\lambda_{\max}(\boldsymbol{W})$, the top eigenvector of $\boldsymbol{W}$ scaled to have unit norm. Then, we have almost surely*

$$\lim_{n\to\infty} \lambda_{\max}(\boldsymbol{W}) = \begin{cases} 2 & \text{if } 0 \leq \beta \leq 1, \\ \beta + \beta^{-1} > 2 & \text{if } \beta > 1, \end{cases} \tag{2}$$

$$\lim_{n\to\infty} |\langle \boldsymbol{x}^\star(\boldsymbol{W}), \boldsymbol{u}\rangle| = \begin{cases} 0 & \text{if } 0 \leq \beta \leq 1, \\ \sqrt{1-\beta^{-2}} > 0 & \text{if } \beta > 1. \end{cases} \tag{3}$$

Moreover, if $0 \leq \beta < 1$ then there is no function of $\boldsymbol{W}$ that with high probability selects correctly whether $\boldsymbol{W}$ is drawn from $\mathbb{P}$ or $\mathbb{Q}$, and if $0 \leq \beta \leq 1$ then there there is no unit vector-valued function of $\boldsymbol{W}$ that has inner product with $\boldsymbol{x}$ asymptotically bounded away from zero with high probability when $\boldsymbol{W} \sim \mathbb{P}$.

More specifically, the behavior of $\lambda_{\max}(\boldsymbol{W})$ is established by Féral and Péché (2007), while the behavior of $|\langle \boldsymbol{x}^\star(\boldsymbol{W}), \boldsymbol{u}\rangle|$ is determined by Capitaine et al. (2009) (both building on the seminal results of Baik et al. (2005)). The impossibility of "detection" or of selecting whether $\boldsymbol{W}$ is drawn from $\mathbb{P}$ or $\mathbb{Q}$ with high probability is shown by Montanari et al. (2015) by establishing *contiguity* of these two sequences of probability measures. Finally, Banks et al. (2018) show that this contiguity implies the impossibility of estimating $\boldsymbol{u}$ with positive correlation.

More refined models follow from choosing more structured distributions of $\boldsymbol{u}$. This corresponds to extracting principal components under some prior knowledge of their structure. One natural example was studied by Montanari and Richard (2015), where $\boldsymbol{u}$ is chosen uniformly from the positive orthant of $\mathbb{S}^{n-1}$ instead of the entire sphere, which yields the aforementioned problem of non-negative PCA. Here, the null model $\mathbb{Q}$ remains as above, while $\mathbb{P}$ is replaced with $\mathbb{P}^+$ defined as follows:

- Under $\mathbb{P}^+$, first draw $\boldsymbol{v} \sim \mathsf{Unif}(\mathbb{S}^{n-1})$, let $\boldsymbol{u}$ have entries $u_i = |v_i|$, and then observe $\boldsymbol{W} = \boldsymbol{W}_0 + \beta \boldsymbol{u}\boldsymbol{u}^\top$ for $\boldsymbol{W}_0 \sim \mathsf{GOE}(n)$ and some fixed $\beta > 0$.

Following the case of classical PCA, we might hope to attack detection and recovery by solving the optimization problem

$$\lambda^+(\boldsymbol{W}) := \max_{\substack{\|\boldsymbol{x}\|_2 = 1 \\ \boldsymbol{x} \geq 0}} \boldsymbol{x}^\top \boldsymbol{W} \boldsymbol{x}, \tag{4}$$

where $\boldsymbol{x} \geq 0$ means $x_i \geq 0$ for each $i \in [n]$. Here, however, a crucial difference between classical PCA and non-negative PCA arises: unlike $\lambda_{\max}(\boldsymbol{W})$ and the associated optimizer, it is NP-hard to compute $\lambda^+(\boldsymbol{W})$ for general $\boldsymbol{W}$ (De Klerk and Pasechnik, 2002). Therefore, non-negative PCA poses a more substantial algorithmic challenge.

1. With the exception of the critical case $\beta = 1$, where thresholding $\lambda_{\max}(\boldsymbol{W})$ does not distinguish $\mathbb{P}$ from $\mathbb{Q}$, but it is possible to do so by considering more sophisticated statistics (Johnstone and Onatski, 2020).
2. We say that a sequence of events A_n occurs *with high probability* under a sequence of probability measures $\mathbb{P}_n$ if $\lim_{n\to\infty} \mathbb{P}_n[A_n] = 1$.

Nonetheless, using an approximate message-passing (AMP) algorithm Montanari and Richard (2015) showed that it is possible to solve this problem essentially to optimality for random inputs from $\mathbb{Q}$ or $\mathbb{P}^+$. We focus now just on the "null" case $\boldsymbol{W} \sim \mathbb{Q} = \mathsf{GOE}(n)$.

Proposition 2 (Montanari and Richard (2015)) *Almost surely for $\boldsymbol{W} \sim \mathsf{GOE}(n)$,*

$$\lim_{n\to\infty} \lambda^+(\boldsymbol{W}) = \sqrt{2}. \tag{5}$$

Moreover, for any $\varepsilon > 0$, there exists an algorithm that runs in time $\mathsf{poly}(n)$ (with runtime also depending on ε) and computes $\boldsymbol{x} \in \mathbb{S}^{n-1}$ with $\boldsymbol{x} \geq 0$ that has $\boldsymbol{x}^\top \boldsymbol{W} \boldsymbol{x} \geq \sqrt{2} - \varepsilon$ with high probability.

This same algorithm is also effective for detection between $\mathbb{Q}$ and $\mathbb{P}^+$ and recovery under $\mathbb{P}^+$.

Remark 3 *One may check that various simpler algorithms do not produce a solution of the same quality. For example, if $\boldsymbol{v}$ is the top eigenvector of $\boldsymbol{W}$, then one simple algorithm is to take $\boldsymbol{v}^+$ having entries $v_i^+ = \max(0, v_i)$ and return $\boldsymbol{x} = \boldsymbol{v}^+ / \|\boldsymbol{v}^+\|$. However, computing heuristically, we have*

$$\boldsymbol{x}^\top \boldsymbol{W} \boldsymbol{x} \approx \frac{1}{\|\boldsymbol{v}^+\|^2} \cdot 2\langle \boldsymbol{v}, \boldsymbol{v}^+ \rangle^2 \approx 2 \cdot 2 \cdot \left(\frac{1}{2}\right)^2 = 1, \tag{6}$$

whereby this choice of $\boldsymbol{x}$ is inferior to that produced by AMP.

In this paper, we study an alternative to AMP, also suggested in Montanari and Richard (2015), where we substitute for the intractable optimization problem $\lambda^+(\boldsymbol{W})$ the following tractable convex relaxation, a natural *semidefinite program (SDP)*:

$$\mathsf{SDP}(\boldsymbol{W}) := \max_{\substack{\boldsymbol{X} \succeq 0 \\ \boldsymbol{X} \geq 0 \\ \mathsf{Tr}(\boldsymbol{X}) = 1}} \langle \boldsymbol{X}, \boldsymbol{W} \rangle \geq \lambda^+(\boldsymbol{W}). \tag{7}$$

The inequality above follows since $\boldsymbol{x}^\top \boldsymbol{W} \boldsymbol{x} = \langle \boldsymbol{W}, \boldsymbol{x}\boldsymbol{x}^\top \rangle$, and $\boldsymbol{x}\boldsymbol{x}^\top$ is feasible for the SDP whenever $\boldsymbol{x}$ is feasible for $\lambda^+(\boldsymbol{W})$. Moreover, it is reasonable to hope that this relaxation procedure is "not too lax," and that the optimizer $\boldsymbol{X}^\star$ has rank close to one, with its dominating eigenvector or top eigenspace close to the optimal $\boldsymbol{x}^\star$ of $\lambda^+(\boldsymbol{W})$. Thus $\mathsf{SDP}(\boldsymbol{W})$ gives an algorithmic strategy for approximating $\lambda^+(\boldsymbol{W})$.

In Montanari and Richard (2015), numerical experiments are presented that suggest that this SDP is effective in recovering $\boldsymbol{x}$ under $\mathbb{P}^+$, the restriction of the top eigenvector of the optimizer $\boldsymbol{X}^\star$ to only positive entries indeed giving a comparable estimate of $\boldsymbol{u}$ to the AMP algorithm (see their Section 5.3). In Section 4, we present analogous experiments for $\boldsymbol{W} \sim \mathsf{GOE}(n)$, and note that in this case the SDP is often *tight*, the optimizer $\boldsymbol{X}^\star$ being rank one within numerical tolerances. While the SDP is much slower than AMP, its apparent efficacy is nevertheless tantalizing, suggesting that the algorithmic tractability of non-negative PCA might be unified with other situations where SDP relaxations of maximum likelihood estimation are tight (Bandeira et al., 2014b). Furthermore, the SDP offers some advantages over AMP: it algorithmically *proves* (or *certifies*) an upper bound on the value of $\lambda^+(\boldsymbol{W})$, and it may also exhibit robustness properties that SDPs have been shown to enjoy in other settings (Feige and Kilian, 2001; Moitra et al., 2016; Ricci-Tersenghi et al., 2016).

The problem of determining the asymptotic behavior of the SDP as $n \to \infty$ under $\boldsymbol{W} \sim \mathsf{GOE}(n)$ was also posed explicitly as Open Problem 9.5 in the lecture notes of Bandeira (2015).

We therefore take up the following two questions concerning this SDP and related algorithmic approaches when $\boldsymbol{W} \sim \mathsf{GOE}(n)$. The first concerns this specific semidefinite program:

1. When $\boldsymbol{W} \sim \mathsf{GOE}(n)$, does $\mathsf{SDP}(\boldsymbol{W}) \to \sqrt{2}$ in probability as $n \to \infty$?

As we will see, the answer in the limit $n \to \infty$, in surprising contrast to the experiments of Montanari and Richard (2015) for small n, is *no*. In fact, we instead have $\mathsf{SDP}(\boldsymbol{W}) \to \lambda_{\max}(\boldsymbol{W}) \approx 2$ as $n \to \infty$. We then ask whether any remotely efficient algorithm that *certifies* upper bounds on $\lambda^+(\boldsymbol{W})$ can improve upon this.

2. Does there exist an algorithm that runs in time $\exp(O(n^{1-\eta}))$ for some fixed $\eta > 0$ and computes $c : \mathbb{R}^{n \times n}_{\mathsf{sym}} \to \mathbb{R}$ with the following two properties?
   - For *all* $\boldsymbol{W} \in \mathbb{R}^{n \times n}_{\mathsf{sym}}$, we have $c(\boldsymbol{W}) \geq \lambda^+(\boldsymbol{W})$.
   - When $\boldsymbol{W} \sim \mathsf{GOE}(n)$, we have $c(\boldsymbol{W}) \leq 2 - \varepsilon$ with high probability for some fixed $\varepsilon > 0$.

We provide rigorous evidence, based on the *low-degree polynomial method*, that even the answer to this much broader question again is *no*.

1.1. Organization

The remainder of the paper is organized as follows. In Section 2, we state and prove a lower bound on $\mathsf{SDP}(\boldsymbol{W})$ when $\boldsymbol{W} \sim \mathsf{GOE}(n)$. In Section 3, we state and prove a *reduction* from a certain hypothesis testing problem to the problem of certifying bounds on $\lambda^+(\boldsymbol{W})$, and review evidence from prior work that this hypothesis testing problem is computationally hard. Finally, in Section 4, we present the results of larger numerical experiments that capture the departure from the "tight regime" where the optimizer of SDP has rank one, a striking example of the difference between theoretical asymptotics and computations tractable at "laptop scale" for semidefinite programming.

2. Lower Bound on Semidefinite Programming

In this section we will prove the following result, which gives the asymptotic value of $\mathsf{SDP}(\boldsymbol{W})$.

Theorem 4 *For any $\varepsilon > 0$, $\lim_{n\to\infty} \mathbb{P}[2 - \epsilon \leq \mathsf{SDP}(\boldsymbol{W}) \leq 2 + \epsilon] = 1$ where $\boldsymbol{W} \sim \mathsf{GOE}(n)$.*

The main technical tool required will be the following concentration inequality for the entries of a random projection matrix.

Proposition 5 *Let $\delta \in (0, 1)$. Let $\boldsymbol{P} \in \mathbb{R}^{n \times n}$ be the orthogonal projection matrix to a Haar-distributed subspace of $\mathbb{R}^n$ having dimension $r := \delta n$. Then, for any $K > 0$, there exist constants $C_{\delta,K}, C'_{\delta,K} > 0$ such*

$$\mathbb{P}\left[\max_{i,j\in[n]} \begin{cases} |P_{ii} - \delta| & \text{if } i = j \\ |P_{ij}| & \text{if } i \neq j \end{cases} \leq C_{\delta,K}\sqrt{\frac{\log n}{n}}\right] \geq 1 - \frac{C'_{\delta,K}}{n^K}. \tag{8}$$

See, e.g., Kunisky and Bandeira (2020) for a careful proof.

Proof (of Theorem 4) For the upper bound, note that $\langle \boldsymbol{X}, \boldsymbol{W} \rangle \leq \lambda_{\max}(\boldsymbol{W})$ for any $\boldsymbol{X}$ feasible for the SDP. The bound then follows from standard bounds on the spectrum of $\boldsymbol{W}$ (Anderson et al., 2010).

For the lower bound, fix $\alpha, \delta \in (0, 1)$. Let $r = \delta n$, assuming for the sake of simplicity that this is an integer. Let $\boldsymbol{P}$ be the orthogonal projector to the span of the r eigenvectors of $\boldsymbol{W}$ having the largest eigenvalues. Then, define

$$\boldsymbol{X} = \boldsymbol{X}^{(\alpha,\delta)} := (1-\alpha)\frac{1}{r}\boldsymbol{P} + \alpha\frac{1}{n}\mathbf{1}_n\mathbf{1}_n^\top, \tag{9}$$

where $\mathbf{1}_n \in \mathbb{R}^n$ is the vector with all entries equal to 1. Denote by F the event that $\boldsymbol{X}$ is feasible for the SDP, i.e., the event that $\mathsf{Tr}(\boldsymbol{X}) = 1$, $\boldsymbol{X} \succeq \mathbf{0}$, and $\boldsymbol{X} \geq \mathbf{0}$.

We first show that, for any fixed $\alpha, \delta \in (0, 1)$, $\mathbb{P}[F] \to 1$. Since $\mathsf{Tr}(\boldsymbol{P}) = r$ we have $\mathsf{Tr}(\boldsymbol{X}) = 1$, and $\boldsymbol{X} \succeq \mathbf{0}$ since $\boldsymbol{X}$ is a convex combination of two positive semidefinite matrices. In particular, $X_{ii} \geq 0$ for all $i \in [n]$. For the off-diagonal entries, we observe that the range of $\boldsymbol{P}$ is a Haar-distributed r-dimensional subspace of $\mathbb{R}^n$. Thus by Proposition 5, with high probability, for all $i, j \in [n]$ with $i \neq j$,

$$X_{ij} \geq \frac{\alpha}{n} - \frac{1-\alpha}{r}|P_{ij}| \geq \frac{\alpha}{n} - C_\delta(1-\alpha)\frac{1}{n}\sqrt{\frac{\log n}{n}} \tag{10}$$

for some constant $C_\delta > 0$ depending only on δ. In particular, for α, δ fixed as $n \to \infty$, this is non-negative for all sufficiently large n, thus $X_{ij} \geq 0$ for all $i, j \in [n]$ with high probability. Combining these observations, we find that F occurs with high probability.

On the event F, we have

$$\mathsf{SDP}(\boldsymbol{W}) \geq \langle \boldsymbol{W}, \boldsymbol{X} \rangle = (1-\alpha) \cdot \frac{1}{r}\sum_{i=1}^{r} \lambda_i(\boldsymbol{W}) + \alpha\frac{1}{n}\mathbf{1}_n^\top \boldsymbol{W}\mathbf{1}_n. \tag{11}$$

The second term is distributed as $\mathcal{N}(0, 2\alpha^2/n)$, and the first term is with high probability bounded below by $(1-\alpha)(2 - f(\delta))$ for some $f(\delta)$ with $\lim_{\delta\to 0} f(\delta) = 0$, by the convergence of the law of the empirical spectrum of $\boldsymbol{W}$ to the semicircle distribution (Anderson et al., 2010). In particular, for any $\varepsilon > 0$, we may choose $\alpha, \delta \in (0, 1)$ sufficiently small that the above argument shows $\mathsf{SDP}(\boldsymbol{W}) \geq 2 - \varepsilon$ with high probability. ∎

We remark that from (10) and a suitable generalization of Proposition 5 we may also extract the following more quantitative claim: even for δ depending upon n, for any given $\delta = \delta(n) \gg \frac{\log n}{n}$, the minimum $\alpha = \alpha(n)$ such that $\boldsymbol{X}^{(\alpha,\delta)}$ is feasible for the SDP with high probability scales as $\alpha = O((\delta n)^{-1/2})$ (neglecting logarithmic factors). By integrating the edge of the semicircle distribution we also have $f(\delta) \sim \delta^{-1}\int_{2-\delta}^{2}\sqrt{4-t^2}dt \sim \delta^{1/2}$ for small δ. For a given scaling of $\delta(n)$ we then expect to have $\mathsf{SDP}(\boldsymbol{W}) \geq (1 - O((\delta n)^{-1/2}))(2 - \delta^{1/2})$, which converges fastest to 2 when $\delta \sim n^{-1/2}$. Thus we expect a rate of convergence of $\mathsf{SDP}(\boldsymbol{W}) \geq 2 - O(n^{-1/4})$.

It is tempting to try to more directly address our numerical experiments and those of Montanari and Richard (2015) by establishing rigorously the dependence of the probability of tightness (i.e., of the optimizer $\boldsymbol{X}^\star$ having rank one) on n, but this seems to be a challenging question. Indeed, previous work on random semidefinite programs has only established tightness when it occurs with

high probability, and then by explicitly constructing dual certificates. Investigating such constructions for small n without relying on asymptotic effects would be informative, but likely far beyond the scope of our techniques.

We also note that the general proof technique of "nudging" an initial construction that is not feasible for a convex program towards a deterministic feasible point has been used before for relaxations of the cut polytope where the latter point is the identity matrix (Avis and Umemoto, 2003; Kunisky and Bandeira, 2020; Mohanty et al., 2020). Our proof adapts this to our different SDP constraints by using the all-ones matrix for this purpose instead.

It would be interesting to extend this result to lower bounds on higher-degree sum-of-squares relaxations of $\lambda^+(\boldsymbol{W})$; based on our results in Section 3, it is natural to conjecture that no relaxation of constant degree certifies a bound strictly smaller than 2 on $\lambda^+(\boldsymbol{W})$ as $n \to \infty$ (since this would refute Conjecture 7). We remark that, in working with inequality constraints, there are a number of reasonable ways to formulate a sum-of-squares relaxation of given degree; see, e.g., Laurent (2009); O'Donnell and Zhou (2013) for some discussion of these details. To the best of our knowledge, lower bounds for sum-of-squares relaxations with inequality constraints have not been studied for high-dimensional random problems, so this problem would be a convenient testing ground to see whether these nuances play an important technical role.

3. Evidence for General Hardness of Certification

We first formalize the notion of a certification algorithm.

Definition 6 (Certification algorithm) *Suppose an algorithm takes as input $\boldsymbol{W} \in \mathbb{R}^{n \times n}_{\mathsf{sym}}$ and outputs a number $c(\boldsymbol{W}) \in \mathbb{R}$ such that $c(\boldsymbol{W}) \geq \lambda^+(\boldsymbol{W})$ for all $\boldsymbol{W} \in \mathbb{R}^{n \times n}_{\mathsf{sym}}$. If when $\boldsymbol{W} \sim \mathsf{GOE}(n)$ then $c(\boldsymbol{W}) \leq K$ with high probability as $n \to \infty$, then we say that this algorithm* certifies *the bound $\lambda^+(\boldsymbol{W}) \leq K$.*

The key property of a certification algorithm is that it must give a valid upper bound on $\lambda^+(\boldsymbol{W})$ no matter what input matrix $\boldsymbol{W}$ is supplied; in particular, it must even do so for $\boldsymbol{W}$ that are atypical under the distribution $\mathsf{GOE}(n)$. However, this upper bound only needs to be a "good" bound for typical $\boldsymbol{W} \sim \mathsf{GOE}(n)$.

Note that $\lambda_{\max}(\boldsymbol{W})$ certifies the bound $\lambda^+(\boldsymbol{W}) \leq 2 + o(1)$. Generally speaking, one notable class of certification algorithms are convex relaxations, including the SDP given in (7). We have established above that this SDP does not improve upon the bound achieved by the simple "spectral certificate" $\lambda_{\max}(\boldsymbol{W})$. It is then natural to ask: is the issue merely that the SDP is not a sufficiently sophisticated algorithm to certify strong bounds, and better convex relaxations will improve on its performance, or is there a fundamental barrier to certifying bounds on $\lambda^+(\boldsymbol{W})$? The goal of this section is to provide formal evidence for the following conjecture, which states that the simple spectral certificate cannot be improved except by a fully exponential-time brute force search.

Conjecture 7 *For any fixed $\varepsilon > 0$ and $\eta > 0$, there is no algorithm of runtime $\exp(O(n^{1-\eta}))$ that certifies the bound $\lambda^+(\boldsymbol{W}) \leq 2 - \varepsilon$ (in the sense of Definition 6).*

We note that, while we work under the simple average-case model $\boldsymbol{W} \sim \mathsf{GOE}(n)$ as precise results are known in this setting for the true value of $\lambda^+(\boldsymbol{W})$, this conjecture casts doubt on the efficacy of convex relaxations for certifying bounds on $\lambda^+(\boldsymbol{W})$ for more general $\boldsymbol{W}$. Specifically, our proof

techniques will suggest that a similar "spectral barrier" to certifying strong bounds likely holds so long as the top eigenspaces of $\boldsymbol{W}$ are in sufficiently "general position." Establishing that this is actually the case without the technical conveniences of working under Gaussian models is an intriguing open problem for future research.

We will argue that the certification problem is hard by reduction from a particular hypothesis testing problem, which we define next.

Definition 8 (Centered Bernoulli distribution) *For a constant $\rho \in (0,1)$, let $\mathcal{X}_\rho$ be the distribution over $\mathbb{R}^n$ where $\boldsymbol{u} \sim \mathcal{X}_\rho$ is drawn by drawing each coordinate u_i independently as*

$$u_i = \begin{cases} \sqrt{\frac{1-\rho}{\rho n}} & \text{with probability } \rho, \\ -\sqrt{\frac{\rho}{(1-\rho)n}} & \text{with probability } 1-\rho. \end{cases} \tag{12}$$

This is scaled so that $\mathbb{E}[u_i] = 0$ and $\|\boldsymbol{u}\| \to 1$ in probability.

Definition 9 *Given constants $\gamma > 0$ and $\beta > -1$, the* spiked Wishart model *with spike prior $\mathcal{X}_\rho$ consists of the following pair of probability distributions. Let $N = N(n) \in \mathbb{N}$ such that $n/N \to \gamma$ as $n \to \infty$.*

- *Under $\mathbb{Q}$, draw $\boldsymbol{y}_1, \dots, \boldsymbol{y}_N \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_n)$ independently.*
- *Under $\mathbb{P}$, first draw $\boldsymbol{u} \sim \mathcal{X}_\rho$. If $\beta\|\boldsymbol{u}\|^2 \le -1$, draw $\boldsymbol{y}_1 = \cdots = \boldsymbol{y}_N = 0$. Otherwise, draw $\boldsymbol{y}_1, \dots, \boldsymbol{y}_N \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_n + \beta \boldsymbol{u}\boldsymbol{u}^\top)$ independently (noting that the covariance matrix is positive definite).*

If $\beta < 0$, we call such a model a negatively-spiked Wishart model.

We will consider the *strong detection* problem where the goal is to give a test $f : \mathbb{R}^{n \times N} \to \{\mathrm{p}, \mathrm{q}\}$ that takes input $\boldsymbol{y} = (\boldsymbol{y}_1, \dots, \boldsymbol{y}_N)$ and distinguishes between $\mathbb{P}$ and $\mathbb{Q}$ with error probability $o(1)$, i.e.,

$$\lim_{n\to\infty} \mathbb{P}[f(\boldsymbol{y}) = \mathrm{p}] = \lim_{n\to\infty} \mathbb{Q}[f(\boldsymbol{y}) = \mathrm{q}] = 1. \tag{13}$$

When $\beta^2 > \gamma$ (the "BBP transition"), it is well-known that strong detection is possible in polynomial time via the maximum (if $\beta > 0$) or minimum (if $\beta < 0$) eigenvalue of the sample covariance matrix (Baik et al., 2005; Baik and Silverstein, 2006). While strong detection is sometimes (depending on ρ) possible when $\beta^2 < \gamma$ via brute force search, this is conjectured to be impossible in subexponential time.

Conjecture 10 (Conjecture 3.1 and Corollary 3.3 of Bandeira et al. (2020b)) *For any constants $\gamma > 0$, $\beta > -1$, $\rho \in (0,1)$, $\eta > 0$ such that $\beta^2 < \gamma$, there is no algorithm of runtime $\exp(O(n^{1-\eta}))$ that achieves strong detection in the spiked Wishart model with parameters γ, β and spike prior $\mathcal{X}_\rho$.*

This conjecture is justified in Bandeira et al. (2020b) by formal evidence based on the *low-degree polynomial method*, a framework based on Barak et al. (2019); Hopkins and Steurer (2017); Hopkins et al. (2017); Hopkins (2018) that has been successful in predicting and explaining computational hardness in a wide variety of tasks in high-dimensional statistics; see Kunisky et al. (2019) for a survey. More precisely, it is shown (Theorem 3.2 of Bandeira et al. (2020b)) that when $\beta^2 < \gamma$, no

multivariate polynomial $f : \mathbb{R}^{n\times N} \to \mathbb{R}$ of degree $D = o(n/\log n)$ can distinguish $\mathbb{P}$ and $\mathbb{Q}$ in the sense of $\mathbb{E}_{\mathbb{P}}[f(\boldsymbol{y})] \to \infty$ while $\mathbb{E}_{\mathbb{Q}}[f(\boldsymbol{y})^2] = 1$. (This is true not only for the centered Bernoulli prior but more generally for any spike prior where the u_i are distributed i.i.d. as $\frac{1}{\sqrt{n}}\pi$ for a fixed distribution π on $\mathbb{R}$ that is subgaussian with $\mathbb{E}[\pi] = 0$ and $\mathbb{E}[\pi^2] = 1$.) Degree-D polynomial tests of the above form are believed to be as powerful as any $\exp(\tilde{\Omega}(D))$-time algorithm (where $\tilde{\Omega}$ hides factors of $\log n$) for a broad class of high-dimensional testing problems; see Hopkins (2018); Kunisky et al. (2019); Ding et al. (2019).

We are now prepared to state the main result of this section, which shows that if it is possible to certify a bound on $\lambda^+(\boldsymbol{W})$ below 2, then it is possible to produce a test between $\mathbb{P}$ and $\mathbb{Q}$ in a particular negatively-spiked Wishart model whose parameters lie in the "hard" regime $\beta^2 < \gamma$. An immediate consequence is that Conjecture 10 implies Conjecture 7.

Theorem 11 (Reduction from detection to certification) *Suppose there exists a constant $\varepsilon > 0$ and a $t(n)$-time certification algorithm $c : \mathbb{R}^{n\times n}_{\mathsf{sym}} \to \mathbb{R}$ for λ^+ such that, with high probability as $n \to \infty$, $c(\boldsymbol{W}) \leq 2 - \varepsilon$ when $\boldsymbol{W} \sim \mathsf{GOE}(n)$. Then there exist constants $\gamma > 1$, $\beta \in (-1, 0)$, and $\rho \in (0, 1)$ (depending on ε) such that there is a $(t(n) + \mathsf{poly}(n))$-time algorithm computing $f : \mathbb{R}^{n\times N} \to \{p, q\}$ that achieves* strong detection *(in the sense of (13)) in the negatively-spiked Wishart model with parameters γ, β and spike prior $\mathcal{X}_\rho$.*

The proof is similar to that of Theorem 3.8 in the prior work Bandeira et al. (2020b) (which gives the analogous result when the constraint set is $\{\pm 1/\sqrt{n}\}^n$ instead of the positive orthant) with one key difference. As in Bandeira et al. (2020b), the idea of the reduction is to create a GOE matrix whose top eigenspace has been "rotated" to align with the orthogonal complement of the span of the given Wishart samples. If the samples come from $\mathbb{P}$ (with β slightly greater than -1 and γ slightly greater than 1) then the Wishart samples are nearly orthogonal to the planted vector $\boldsymbol{u}$, so this has the effect of planting $\boldsymbol{u}$ in the top eigenspace of the matrix. We would like to plant a non-negative vector in the top eigenspace so that any certifier is forced to output a bound larger than $2 - \varepsilon$. However, we cannot take $\boldsymbol{u}$ to be non-negative because it is important for Wishart hardness (Conjecture 10) that $\boldsymbol{u}$ have mean zero. The key idea is to instead choose $\boldsymbol{u}$ to be a mean-zero random vector that is highly correlated with a certain non-negative vector $\widehat{\boldsymbol{z}}$; this is the purpose of introducing the centered Bernoulli prior $\mathcal{X}_\rho$.

Proof (of Theorem 11) Suppose a certification algorithm as stated exists. We will use this to design a test f achieving strong detection in the negatively-spiked Wishart model.

Call $\boldsymbol{y}_1, \dots, \boldsymbol{y}_N$ the samples from the Wishart model. Draw $\widetilde{\boldsymbol{W}} \sim \mathsf{GOE}(n)$ and let $\lambda_1 \leq \cdots \leq \lambda_n$ be its eigenvalues. Let $\boldsymbol{v}_1, \dots, \boldsymbol{v}_N$ be a uniformly random orthonormal basis for $V := \mathsf{span}(\{\boldsymbol{y}_1, \dots, \boldsymbol{y}_N\})$ and let $\boldsymbol{v}_{N+1}, \dots, \boldsymbol{v}_n$ be a uniformly random orthonormal basis for the orthogonal complement $V^\perp$. Let $\boldsymbol{W} := \sum_{i=1}^n \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^\top$.

Then, using the certification algorithm computing $c(\boldsymbol{W})$ as a subroutine, we compute the test f as

$$f(\boldsymbol{W}) := \begin{cases} \mathtt{q} & \text{if } c(\boldsymbol{W}) \leq 2 - \varepsilon, \\ \mathtt{p} & \text{otherwise.} \end{cases} \tag{14}$$

When $(\boldsymbol{y}_1, \dots, \boldsymbol{y}_N) \sim \mathbb{Q}$, then $\boldsymbol{W}$ has the law $\mathsf{GOE}(n)$, so $c(\boldsymbol{W}) \leq 2 - \varepsilon$ with high probability, and thus $f(\boldsymbol{W}) = \mathtt{q}$ with high probability. Thus to complete the proof it suffices to show that, when $(\boldsymbol{y}_1, \dots, \boldsymbol{y}_N) \sim \mathbb{P}$, then $c(\boldsymbol{W}) > 2 - \varepsilon$ with high probability, whereby we will have $f(\boldsymbol{W}) = \mathtt{p}$ with high probability.

To this end, suppose $(\boldsymbol{y}_1, \dots, \boldsymbol{y}_N) \sim \mathbb{P}$ with $\boldsymbol{u}$ the spike vector. Let $\boldsymbol{z} \geq 0$ be the vector

$$z_i = \begin{cases} 1/\sqrt{\rho n} & \text{if } u_i > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{15}$$

and note that $\|\boldsymbol{z}\| \to 1$ and $\langle \boldsymbol{z}, \boldsymbol{u} \rangle \to \sqrt{1-\rho}$ in probability. Let $\widehat{\boldsymbol{z}} := \boldsymbol{z}/\|\boldsymbol{z}\|$. We then have

$$\begin{aligned} c(\boldsymbol{W}) &\geq \lambda^+(\boldsymbol{W}) \\ &\geq \widehat{\boldsymbol{z}}^\top \boldsymbol{W} \widehat{\boldsymbol{z}} \\ &= \sum_{i=1}^n \lambda_i \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2 \end{aligned}$$

and partitioning the spectrum at λ_{N+1},

$$\begin{aligned} &\geq \lambda_1 \sum_{i=1}^N \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2 + \lambda_{N+1} \sum_{i=N+1}^n \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2 \\ &\geq \lambda_1 \sum_{i=1}^N \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2 + \lambda_{N+1} \left(1 - \sum_{i=1}^N \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2 \right) \end{aligned}$$

and, since $\{\boldsymbol{v}_1, \dots, \boldsymbol{v}_n\}$ is an orthonormal basis and $\|\widehat{\boldsymbol{z}}\| = 1$,

$$= \lambda_{N+1} - (\lambda_{N+1} - \lambda_1) \sum_{i=1}^N \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2. \tag{16}$$

Recalling that $\{\boldsymbol{v}_i\}_{i=1}^N$ is an orthonormal basis for $\mathsf{span}(\{\boldsymbol{y}_1, \dots, \boldsymbol{y}_N\})$, we have $\sum_{i=1}^N \boldsymbol{v}_i \boldsymbol{v}_i^\top \preceq \frac{1}{\mu} \boldsymbol{Y}$ where $\boldsymbol{Y} = \frac{1}{N} \sum_{i=1}^N \boldsymbol{y}_i \boldsymbol{y}_i^\top$ and μ is the smallest nonzero eigenvalue of $\boldsymbol{Y}$.

$$\sum_{i=1}^N \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2 \leq \frac{1}{\mu N} \sum_{i=1}^N \langle \widehat{\boldsymbol{z}}, \boldsymbol{y}_i \rangle^2$$

We note that, viewing $\boldsymbol{Y}$ as a sample covariance matrix under a spiked matrix model, we have $\mu \to (\sqrt{\gamma} - 1)^2 > 0$ in probability by Theorem 1.2 of Baik and Silverstein (2006). To control the second factor above, let $\boldsymbol{g}_1, \dots, \boldsymbol{g}_N \sim \mathcal{N}(0, \boldsymbol{I}_n)$ independently of all other random variables in the proof. Then, we may view $\boldsymbol{y}_i = (\boldsymbol{I}_n + \beta \boldsymbol{u}\boldsymbol{u}^\top)^{1/2} \boldsymbol{g}_i$. Thus,

$$= \frac{1}{\mu} \widehat{\boldsymbol{z}}^\top (\boldsymbol{I}_n + \beta \boldsymbol{u}\boldsymbol{u}^\top)^{1/2} \left(\frac{1}{N} \sum_{i=1}^N \boldsymbol{g}_i \boldsymbol{g}_i^\top \right) (\boldsymbol{I}_n + \beta \boldsymbol{u}\boldsymbol{u}^\top)^{1/2} \widehat{\boldsymbol{z}}$$

Let ν be the largest eigenvalue of $\frac{1}{N} \sum_{i=1}^N \boldsymbol{g}_i \boldsymbol{g}_i^\top$, noting that, by the same result cited above, $\nu \to (\sqrt{\gamma} + 1)^2$ in probability. We then have

$$\begin{aligned} &\leq \frac{\nu}{\mu} \|(\boldsymbol{I}_n + \beta \boldsymbol{u}\boldsymbol{u}^\top)^{1/2} \widehat{\boldsymbol{z}}\|^2 \\ &= \frac{\nu}{\mu} \widehat{\boldsymbol{z}}^\top (\boldsymbol{I}_n + \beta \boldsymbol{u}\boldsymbol{u}^\top) \widehat{\boldsymbol{z}} \\ &= \frac{\nu}{\mu} (1 + \beta \langle \widehat{\boldsymbol{z}}, \boldsymbol{u} \rangle^2) \end{aligned}$$

Finally, since $\|\boldsymbol{z}\| \to 1$ and $\langle \boldsymbol{z}, \boldsymbol{u} \rangle^2 \to 1 - \rho$ in probability, substituting in limiting values for all convergent quantities, we have that, in probability,

$$\to \left(\frac{\sqrt{\gamma} + 1}{\sqrt{\gamma} - 1} \right)^2 (1 + \beta(1 - \rho)). \tag{17}$$

By the convergence of $\{\lambda_i\}$ to the semicircle law on $[-2, 2]$ (Anderson et al., 2010), we have $\lambda_{N+1} - \lambda_1 \leq 5$ with high probability. Also, by choosing $\gamma > 1$ sufficiently close to 1, we can ensure $\lambda_{N+1} \geq 2 - \varepsilon/2$ with high probability and $\sqrt{\gamma} + 1 \leq 3$.

Putting it all together, the value of $c(\boldsymbol{W})$ under $(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N) \sim \mathbb{P}$ with high probability satisfies

$$c(\boldsymbol{W}) \geq \lambda_{N+1} - (\lambda_{N+1} - \lambda_1) \sum_{i=1}^{N} \langle \widehat{\boldsymbol{z}}, \boldsymbol{v}_i \rangle^2 \geq 2 - \varepsilon/2 - \frac{45}{(\sqrt{\gamma} - 1)^2} \big(1 + \beta(1 - \rho)\big),$$

which exceeds $2 - \varepsilon$ provided we choose $\beta > -1$ close enough to -1 and $\rho > 0$ small enough. ■

4. Deceptive Finite-Size Effects

Experiments for small n computing $\mathsf{SDP}(\boldsymbol{W})$ suggest, contrary to our results, that $\mathsf{SDP}(\boldsymbol{W})$ is in fact very effective in bounding $\lambda^+(\boldsymbol{W})$ for $\boldsymbol{W} \sim \mathsf{GOE}(n)$. Not only do we observe for small n the value $\mathsf{SDP}(\boldsymbol{W}) \approx \sqrt{2}$, but we also find that $\mathsf{SDP}(\boldsymbol{W})$ appears to be *tight*, the primal optimizer $\boldsymbol{X}^\star$ often having rank one to numerical tolerances. This is analogous to the efficacy of $\mathsf{SDP}(\boldsymbol{W})$ for recovering the spike for $\boldsymbol{W}$ under the spiked non-negative PCA model discussed in Montanari and Richard (2015). We describe these results here, as well as further experiments suggesting what size of n is required for these finite size effects to give way to the correct asymptotics.

The first experiment we consider solves $\mathsf{SDP}(\boldsymbol{W})$ for many random choices of $\boldsymbol{W}$, obtains the optimizer $\boldsymbol{X}^\star$, and considers the *numerical rank* of $\boldsymbol{X}^\star$. We plot the results of 50 trials of this experiment with $n = 150$ in Figure 1, and observe that most trials have the second-largest eigenvalue of $\boldsymbol{X}^\star$ of order at most 10^{-4} compared to the trace of 1, whereby $\boldsymbol{X}^\star$ is nearly rank-one and the SDP is nearly tight.

The next experiment solves the following different SDP, which is dual to $\mathsf{SDP}(\boldsymbol{W})$, and which by a standard strong duality argument has the same value as $\mathsf{SDP}(\boldsymbol{W})$:

$$\mathsf{SDP}^*(\boldsymbol{W}) := \min_{\boldsymbol{Y} \geq 0} \lambda_{\max}(\boldsymbol{W} + \boldsymbol{Y}) = \mathsf{SDP}(\boldsymbol{W}). \tag{18}$$

Having $\mathsf{SDP}(\boldsymbol{W}) \leq 2 - \varepsilon$ therefore has the elegant interpretation of it being possible to "compress" the spectrum of $\boldsymbol{W} \sim \mathsf{GOE}(n)$ below 2 by only increasing each entry. We plot the results of 50 trials of this experiment with $n = 150$ in Figure 2. (These semidefinite programs are solved using version 9.2 of the `Mosek` solver on a laptop computer with 32GB RAM and an Intel i7-1065G7 processor; the average time to solve an instance is 18.1 minutes.) From these results, this compression indeed appears possible; moreover, the compressed spectrum appears to have an interesting "wall shape" not unlike that of the GOE conditioned on its largest eigenvalue being small; see, e.g., Figure 4 of Majumdar and Schehr (2014).

The final experiment seeks to identify how large we would need to make n in order to observe that the above are all illusory finite-size effects. To do this, we consider our primal witness from the proof of Theorem 4,

$$\boldsymbol{X}^{(\alpha,\delta)} = \frac{1-\alpha}{\delta n}\boldsymbol{P} + \frac{\alpha}{n}\mathbf{1}_n\mathbf{1}_n^\top \tag{19}$$

for $\boldsymbol{P}$ the projection matrix to the top δn eigenvectors of $\boldsymbol{W} \sim \mathsf{GOE}(n)$ and α large enough that $\boldsymbol{X}^{(\alpha,\delta)} \geq 0$ entrywise. In Figure 3, we fix $\delta = 1/25 = 0.04$, and plot both the smallest α making $\boldsymbol{X}^{(\alpha,\delta)}$ feasible for the SDP and the corresponding lower bound $\langle \boldsymbol{X}^{(\alpha,\delta)}, \boldsymbol{W}\rangle$ on the SDP. We see that the smallest α only decays to zero very slowly, as $\widetilde{O}(n^{-1/2})$ per our argument. Accordingly, $\langle \boldsymbol{X}^{(\alpha,\delta)}, \boldsymbol{W}\rangle$ also only very slowly approaches its limiting value. Moreover, even to have $\langle \boldsymbol{X}^{(\alpha,\delta)}, \boldsymbol{W}\rangle > \sqrt{2}$ requires $n \sim 10^4$, suggesting that this is roughly the size of n required to observe that $\mathsf{SDP}(\boldsymbol{W})$ in fact is not typically tight (in stark contrast to $n \sim 10^2$ that is tractable to solve on commodity hardware).

Taken together, these experiments present a striking caution against extrapolating asymptotic behavior for an SDP from experimental results tractably computable in reasonable time in practice. In our case, the correct asymptotic behavior "kicks in" only for problems two orders of magnitude larger than the largest tractable with off-the-shelf software on a personal computer.

5. Conclusion

We have shown that the natural semidefinite program (SDP) for solving a principal component analysis (PCA) problem constrained to entrywise non-negative vectors fails to reach a tight objective value for an input matrix drawn from the Gaussian orthogonal ensemble. Moreover, any algorithm that always produces a valid upper bound on the value of this optimization problem (including this SDP and more sophisticated convex relaxations) must implicitly solve a hypothesis testing problem in the Wishart spiked matrix model with a negative spike. We have produced evidence that this testing problem is hard by appealing to prior work analyzing low-degree polynomial algorithms; more precisely, our evidence suggests that an algorithm solving such a problem requires nearly-exponential time in the input dimension. These results complement the numerical study of Montanari and Richard (2015), which suggested that the natural SDP was effective in small problems. To reconcile this difference, we have conducted further experiments, concluding that problem sizes intractable on commodity hardware are needed to observe the correct asymptotics.

These results extend those of Bandeira et al. (2020b,a) which showed hardness for other constrained PCA problems. The broader picture emerging from this line of work suggests that, in fact, such problems are quite generically resilient to convex relaxation approaches, or more generally to algorithms that certify bounds on the objective value of all feasible points rather than merely searching for good feasible points. While SDPs often appear effective in practice for such problems, this may only be thanks to the kind of finite-size effect that we investigate here.

Our work suggests several questions for future investigation. First, it would be of great practical interest to attempt to understand the range of problem sizes for which the finite-size effects causing strong performance persist for SDPs on non-negative PCA and related problems. Indeed, if for some application the problem sizes that arise in practice fall within this range, then it is perfectly sensible to use SDPs; our results and methods here, with their heavy dependence on asymptotics in high-dimensional probability, do not give strong guidance of this kind. Second, our strategy for proving that SDPs and certification algorithms more generally do not succeed is quite specific to

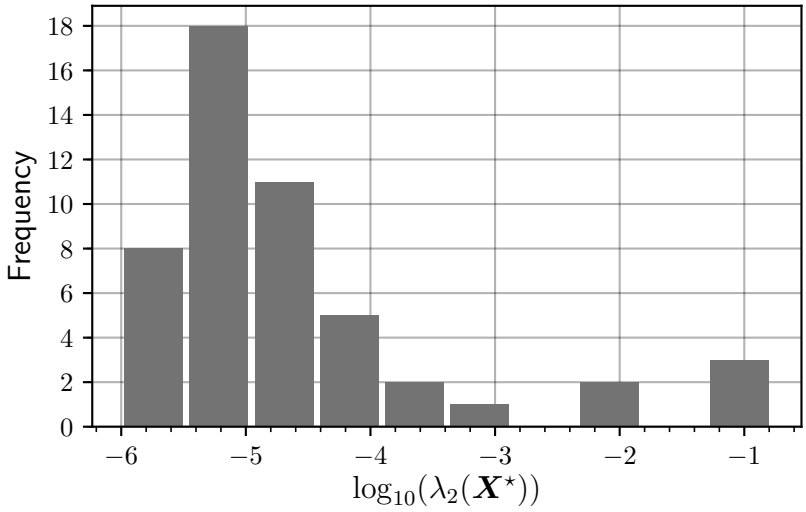

Figure 1: **SDP primal spectrum for small n.** We plot a histogram of the second-largest eigenvalue of $\boldsymbol{X}^\star$ the optimizer of $\mathsf{SDP}(\boldsymbol{W})$, over 50 trials with $n = 150$.

quadratic problems that encode some form of "searching in the spectrum" of an input matrix for certain structured directions or subspaces. Can similar strategies apply to problems such as tensor-valued PCA, which involve higher-degree polynomial optimization and thus do not have a natural spectral interpretation? Finally, Montanari and Richard (2015) showed that approximate message-passing succeeds in solving non-negative PCA problems of this kind. That is, searching for a single good solution to such problems is easy, while certifying bounds on the quality of all solutions is hard. An important question towards a unified picture of the performance guarantees of different types of algorithms is to characterize the problems where such a gap appears.

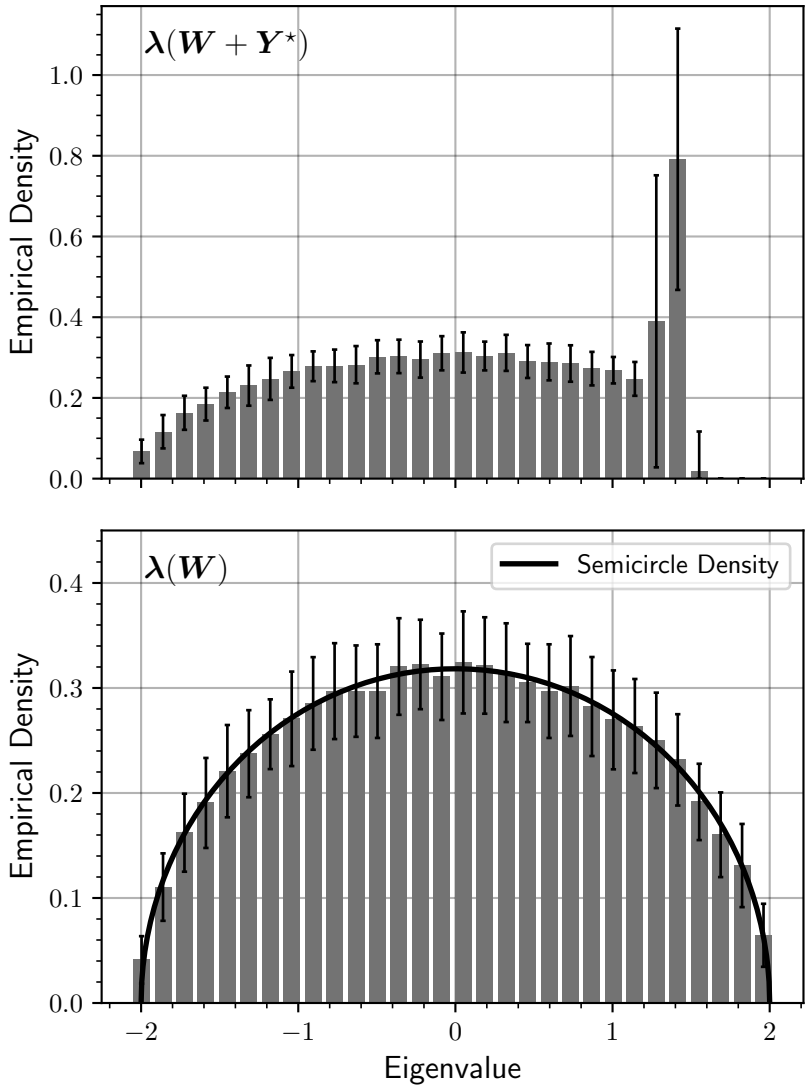

Figure 2: **SDP dual spectrum for small n.** We plot the means of histograms, with error bars of one standard deviation per bin, for the spectra of $\boldsymbol{W}$ and $\boldsymbol{W} + \boldsymbol{Y}^\star$ for $\boldsymbol{Y}^\star$ the optimizer of $\mathsf{SDP}^*(\boldsymbol{W})$, over 50 trials with $n = 150$.

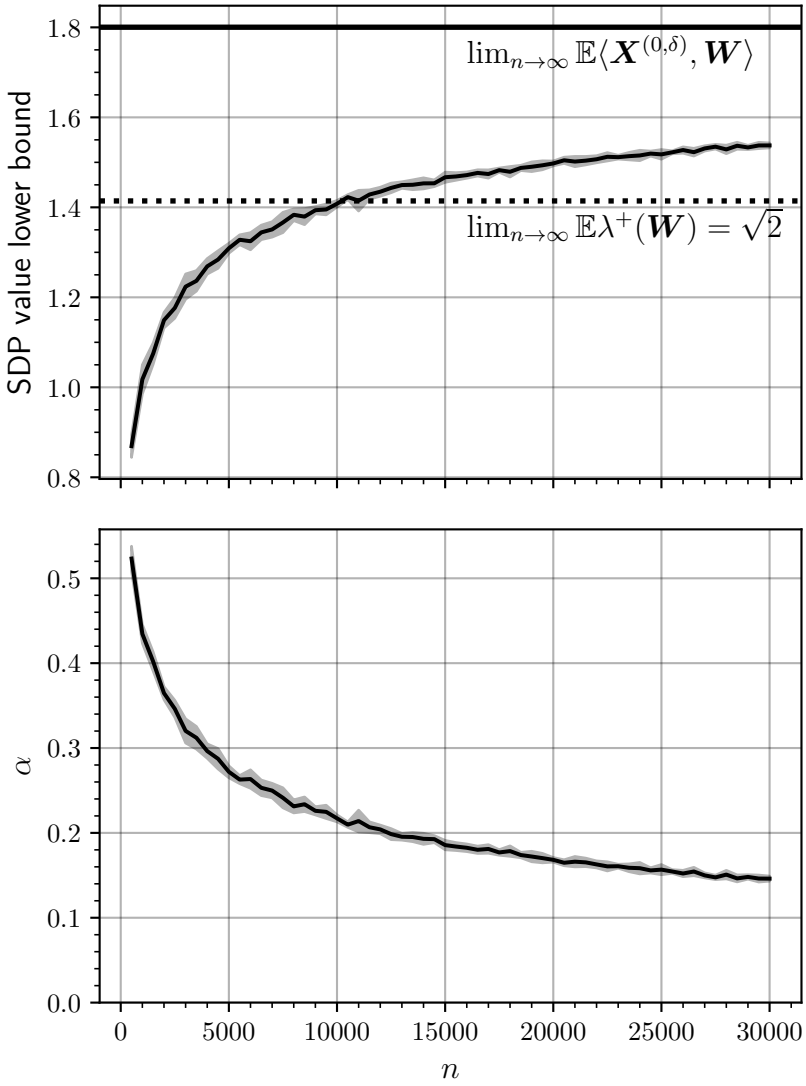

Figure 3: **SDP lower bound convergence.** We fix $\delta = 1/25 = 0.04$, and given $\boldsymbol{W}$ compute the smallest α for such that $\boldsymbol{X}^{(\alpha,\delta)}$, as defined in (19), is feasible for $\mathsf{SDP}(\boldsymbol{W})$. In the upper graph, for a range of values of n, we plot the mean and an error interval of one standard deviation of 10 values of $\langle \boldsymbol{X}^{(\alpha,\delta)}, \boldsymbol{W}\rangle$, a lower bound on $\mathsf{SDP}(\boldsymbol{W})$. We note that this clearly exceeds $\lim_{n\to\infty} \mathbb{E}\lambda^+(\boldsymbol{W}) \approx \sqrt{2}$ once $n \gtrsim 10^4$; however, it only very slowly approaches its expected limiting value. In the bottom graph, we likewise plot the mean and standard deviation of the minimum valid value of α, again observing that it only very slowly approaches its limiting value of zero (which may be verified to match the rate our theoretical calculation predicts of $\alpha = O(n^{-1/2})$ up to logarithmic factors).

Acknowledgments

We thank Yash Deshpande, Sam Hopkins, and Andrea Montanari for helpful discussions. Some of this work was done while Afonso S. Bandeira was with the Department of Mathematics at the Courant Institute of Mathematical Sciences, and the Center for Data Science, at New York University; and partially supported by NSF grants DMS-1712730 and DMS-1719545, and by a grant from the Sloan Foundation. Dmitriy Kunisky was partially supported by NSF grants DMS-1712730 and DMS-1719545. Alexander S. Wein was partially supported by NSF grant DMS-1712730 and by the Simons Collaboration on Algorithms and Geometry; part of this work was done while visiting the Simons Institute for the Theory of Computing.

References

Emmanuel Abbe, Afonso S Bandeira, and Georgina Hall. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory*, 62(1):471–487, 2015.

Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4):433–459, 2010.

Greg W Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*. Cambridge University Press, 2010.

David Avis and Jun Umemoto. Stronger linear programming relaxations of max-cut. *Mathematical Programming*, 97(3):451–469, 2003.

Jinho Baik and Jack W Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, 97(6):1382–1408, 2006.

Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.

Afonso S Bandeira. Ten lectures and forty-two open problems in the mathematics of data science. 2015.

Afonso S Bandeira, Moses Charikar, Amit Singer, and Andy Zhu. Multireference alignment using semidefinite programming. In *5th Conference on Innovations in Theoretical Computer Science (ITCS 2014)*, pages 459–470, 2014a.

Afonso S Bandeira, Yuehaw Khoo, and Amit Singer. Open problem: Tightness of maximum likelihood semidefinite relaxations. *arXiv preprint arXiv:1404.2655*, 2014b.

Afonso S Bandeira, Jess Banks, Dmitriy Kunisky, Cristopher Moore, and Alexander S Wein. Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. *arXiv preprint arXiv:2008.12237*, 2020a.

Afonso S Bandeira, Dmitriy Kunisky, and Alexander S Wein. Computational hardness of certifying bounds on constrained PCA problems. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, pages 78:1–78:29, 2020b.

Jess Banks, Cristopher Moore, Roman Vershynin, Nicolas Verzelen, and Jiaming Xu. Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *IEEE Transactions on Information Theory*, 64(7):4872–4894, 2018.

Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019.

Tamir Bendory, Nicolas Boumal, Chao Ma, Zhizhen Zhao, and Amit Singer. Bispectrum inversion with application to multireference alignment. *IEEE Transactions on Signal Processing*, 66(4): 1037–1050, 2017.

Emmanuel J Candes and Terence Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.

Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.

Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large Wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1–47, 2009.

Kunal N Chaudhury, Yuehaw Khoo, and Amit Singer. Global registration of multiple point clouds using semidefinite programming. *SIAM Journal on Optimization*, 25(1):468–501, 2015.

Etienne De Klerk and Dmitrii V Pasechnik. Approximation of the stability number of a graph via copositive programming. *SIAM Journal on Optimization*, 12(4):875–892, 2002.

Yunzi Ding, Dmitriy Kunisky, Alexander S Wein, and Afonso S Bandeira. Subexponential-time algorithms for sparse PCA. *arXiv preprint arXiv:1907.11635*, 2019.

Yedidyah Dordek, Daniel Soudry, Ron Meir, and Dori Derdikman. Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *Elife*, 5:e10094, 2016.

Uriel Feige and Joe Kilian. Heuristics for semirandom graph problems. *Journal of Computer and System Sciences*, 63(4):639–671, 2001.

Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in Mathematical Physics*, 272(1):185–228, 2007.

José FS Bravo Ferreira, Yuehaw Khoo, and Amit Singer. Semidefinite programming approach for the quadratic assignment problem with a sparse graph. *Computational Optimization and Applications*, 69(3):677–712, 2018.

Bruce Hajek, Yihong Wu, and Jiaming Xu. Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Transactions on Information Theory*, 62(5):2788–2797, 2016.

Samuel Hopkins. *Statistical inference and the sum of squares method*. PhD thesis, Cornell University, 2018.

Samuel B Hopkins and David Steurer. Efficient Bayesian estimation from few samples: community detection and related problems. In *58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*, pages 379–390. IEEE, 2017.

Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. In *58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*, pages 720–731. IEEE, 2017.

Iain M Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, pages 295–327, 2001.

Iain M Johnstone and Alexei Onatski. Testing in high-dimensional spiked models. *Annals of Statistics*, 48(3):1231–1254, 2020.

Dmitriy Kunisky and Afonso S Bandeira. A tight degree 4 sum-of-squares lower bound for the Sherrington-Kirkpatrick Hamiltonian. *Mathematical Programming*, 2020.

Dmitriy Kunisky, Alexander S Wein, and Afonso S Bandeira. Notes on computational hardness of hypothesis testing: predictions using the low-degree likelihood ratio. *arXiv preprint arXiv:1907.11636*, 2019.

Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, pages 157–270. Springer, 2009.

Laura Lazzeroni and Art Owen. Plaid models for gene expression data. *Statistica Sinica*, pages 61–86, 2002.

Satya N Majumdar and Grégory Schehr. Top eigenvalue of a random matrix: large deviations and third order phase transition. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(1): P01012, 2014.

Sidhanth Mohanty, Prasad Raghavendra, and Jeff Xu. Lifting sum-of-squares lower bounds: degree-2 to degree-4. In *52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2020)*, pages 840–853, 2020.

Ankur Moitra, William Perry, and Alexander S Wein. How robust are reconstruction thresholds for community detection? In *48th Annual ACM Symposium on Theory of Computing (STOC 2016)*, pages 828–841. ACM, 2016.

Andrea Montanari and Emile Richard. Non-negative principal component analysis: message passing algorithms and sharp asymptotics. *IEEE Transactions on Information Theory*, 62(3):1458–1484, 2015.

Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. In *Advances in Neural Information Processing Systems*, pages 217–225, 2015.

Ryan O'Donnell and Yuan Zhou. Approximability and proof complexity. In *24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2013)*, pages 1537–1556. Society for Industrial and Applied Mathematics, 2013.

Onur Ozyesil. Camera motion estimation by convex programming. *arXiv preprint arXiv:1312.5047*, 2014.

C Radhakrishna Rao. The use and interpretation of principal component analysis in applied research. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 329–358, 1964.

Federico Ricci-Tersenghi, Adel Javanmard, and Andrea Montanari. Performance of a community detection algorithm based on semidefinite programming. In *Journal of Physics: Conference Series*, volume 699, page 012015, 2016.

Markus Ringnér. What is principal component analysis? *Nature Biotechnology*, 26(3):303–304, 2008.