

PRECISE ERROR RATES FOR COMPUTATIONALLY EFFICIENT TESTING

BY ANKUR MOITRA^{1,a}, AND ALEXANDER S. WEIN^{2,b}¹*Department of Mathematics and CSAIL, MIT, ^amoitra@mit.edu*²*Department of Mathematics, UC Davis, ^baswein@ucdavis.edu*

We revisit the fundamental question of simple-versus-simple hypothesis testing with an eye towards computational complexity, as the statistically optimal likelihood ratio test is often computationally intractable in high-dimensional settings. In the classical spiked Wigner model with a general i.i.d. spike prior we show (conditional on a conjecture) that an existing test based on linear spectral statistics achieves the best possible tradeoff curve between type I and type II error rates among all computationally efficient tests, even though there are exponential-time tests that do better. This result is conditional on an appropriate complexity-theoretic conjecture, namely a natural strengthening of the well-established low-degree conjecture. Our result shows that the spectrum is a sufficient statistic for computationally bounded tests (but not for all tests).

To our knowledge, our approach gives the first tool for reasoning about the precise asymptotic testing error achievable with efficient computation. The main ingredients required for our hardness result are a sharp bound on the norm of the low-degree likelihood ratio along with (counterintuitively) a positive result on achievability of testing. This strategy appears to be new even in the setting of unbounded computation, in which case it gives an alternate way to analyze the fundamental statistical limits of testing.

1. Introduction. A fundamental question in mathematical statistics is that of testing between two simple hypotheses, that is, the task of deciding which of two known distributions produced a given sample. The celebrated Neyman–Pearson lemma states that tests based on thresholding the likelihood ratio are optimal in the sense that they sweep out the best possible tradeoff between type I and type II error rates as the threshold is varied. However, in high-dimensional settings, the likelihood ratio is often intractable to compute because it involves summing over an exponential number of possible values for a latent variable. In this work we revisit the classical problem of simple-versus-simple testing with an eye towards computational complexity: we aim to precisely characterize the best possible testing error achievable by a *computationally efficient algorithm*, which may differ from the statistical limit.

1.1. *Spiked Wigner model.* While our methods have the potential to be applied more broadly, we focus on one canonical high-dimensional testing problem: the *spiked Wigner model*.

DEFINITION 1.1 (Spiked Wigner testing problem). *For a positive integer n , a signal-to-noise ratio (SNR) $\lambda \geq 0$, and a spike prior π which is a distribution on \mathbb{R} with mean 0 and variance 1, let $\mathbb{P}_{\lambda, \pi, n}$ denote the distribution over $n \times n$ symmetric matrices Y generated as*

$$(1) \quad Y = \lambda \cdot \frac{xx^\top}{\|x\|^2} + W$$

MSC2020 subject classifications: Primary 62F03; secondary 68Q17.

Keywords and phrases: Hypothesis testing, spiked Wigner model, low-degree likelihood ratio.

where $x \in \mathbb{R}^n$ has entries drawn i.i.d. from π , and $W \in \mathbb{R}^{n \times n}$ is drawn (independent from x) according to the Gaussian Orthogonal Ensemble (GOE): W is symmetric with $\{W_{ij} : i \leq j\}$ independent, $W_{ij} \sim \mathcal{N}(0, 1/n)$ for $i \neq j$, and $W_{ii} \sim \mathcal{N}(0, 2/n)$. We adopt the convention $xx^\top/\|x\|^2 = 0$ in the (unlikely) event $\|x\| = 0$. For ease of notation, we will suppress the dependence on π, n and write $\mathbb{P}_\lambda = \mathbb{P}_{\lambda, \pi, n}$. Our focus is on the asymptotic regime $n \rightarrow \infty$ with λ, π held fixed. We consider testing the null hypothesis $Y \sim \mathbb{P}_0$ (the case $\lambda = 0$) against a specific alternative $Y \sim \mathbb{P}_\lambda$, where λ, π are known.

This testing problem amounts to detecting the presence of a rank-1 “signal” buried in random “noise.” We emphasize that only a *single* sample $Y \in \mathbb{R}^{n \times n}$, drawn either from \mathbb{P}_0 or \mathbb{P}_λ , is observed.

Spiked (or *deformed*) random matrix models such as (1) and the related *spiked Wishart model* have been extensively studied from the perspective of random matrix theory [40, 6, 9, 60, 59, 30, 51, 5, 21, 56, 17, 62]. Notably, the model (1) undergoes a spectral transition akin to that of Baik, Ben Arous, and Péché [6] at the threshold $\lambda = 1$: when $\lambda \leq 1$, the empirical distribution of eigenvalues converges to the Wigner semicircle law (supported on $[-2, 2]$) and the maximum eigenvalue converges to 2; when $\lambda > 1$, the maximum eigenvalue converges to $\lambda + 1/\lambda > 2$ due to a single “signal” eigenvalue that exits the semicircular bulk [30, 51]. As an immediate corollary, thresholding the maximum eigenvalue gives a test achieving *strong detection* when $\lambda > 1$, that is, testing \mathbb{P}_0 versus \mathbb{P}_λ with both type I and type II error probabilities tending to 0 as $n \rightarrow \infty$.

More recently, spiked models have been studied from a statistical perspective with the aim of identifying the best possible test, which may not be based on the maximum eigenvalue [57, 58, 53, 48, 14, 61, 28, 22, 42]. For the moment, we restrict ourselves to *spectral tests*, that is, tests that only use the spectrum (multiset of eigenvalues) of Y . In the Wigner model (Definition 1.1), no spectral test can achieve strong detection when $\lambda < 1$ [53] but *weak detection*, i.e., testing with a non-vanishing advantage over random guessing, is possible for any $\lambda > 0$ by thresholding the trace of Y . More precisely, the best possible asymptotic trade-off curve between type I and type II error rates for spectral tests is known for each $\lambda < 1$ [7, 8] (see [22]), and is furthermore achieved by a computationally efficient (polynomial-time) test based on *linear spectral statistics* (LSS) [22]. Specifically, [22] considers a family of tests, henceforth *LSS*, based on thresholding the value

$$(2) \quad \sum_{i=1}^n h_\lambda(\mu_i) \quad \text{with} \quad h_\lambda(\mu) := -\log(1 - \lambda\mu + \lambda^2),$$

where $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$ are the eigenvalues of Y , and the optimal function h_λ has been carefully chosen. The LSS paradigm has been studied in a long line of work, including [3, 2, 4, 7, 13, 22, 41, 42, 66, 49, 25].

However, the spectrum of Y only contains information about the norm λ of the rank-1 signal term, not its direction $x/\|x\|$. Conceivably a better test can be constructed by exploiting the eigenvectors of Y in addition to the eigenvalues. Indeed, for some (but not all) spike priors π , there are (non-spectral) tests that achieve strong detection for some $\lambda < 1$ [14, 61], beating the spectral threshold. For instance, if π places enough mass on 0 so that x is sufficiently sparse, strong detection is possible below the spectral threshold by exhaustively enumerating all possible sparsity patterns. For any prior π with bounded support, [28] resolves the optimal statistical performance among all possible tests, identifying the exact threshold $\lambda^* = \lambda^*(\pi) \leq 1$ above which strong detection is information-theoretically possible (confirming a conjecture of [48]), and showing that the error tradeoff (between types I and II) of LSS is information-theoretically optimal when $\lambda < \lambda^*$.

While the above results resolve the fundamental statistical limits of the spiked Wigner testing problem, the computational complexity remains unclear: for all π and all $\lambda < 1$, the best known *computationally efficient* test is LSS. In other words, the best known computationally efficient tests use only the spectrum, suggesting that the spectrum may be a “*computationally sufficient statistic*.” In this work we aim to confirm this by answering (affirmatively) the following question:

Do linear spectral statistics (LSS) achieve the best possible tradeoff between type I and type II error rates among all polynomial-time tests for all π and all $\lambda < 1$?

The regime of interest for this question is $\lambda^* < \lambda < 1$, where strong detection is information-theoretically possible but all known algorithms achieving this require exponential time. For instance, optimal statistical performance is of course achieved by thresholding the likelihood ratio

$$(3) \quad \frac{d\mathbb{P}_\lambda}{d\mathbb{P}_0}(Y) = \mathbb{E}_{\substack{x \sim \pi \\ x \text{ iid}}} \exp \left(-\frac{n}{4} \left\| Y - \lambda \frac{xx^\top}{\|x\|^2} \right\|_F^2 + \frac{n}{4} \|Y\|_F^2 \right),$$

but naive evaluation of this quantity requires a sum over $\exp(\Omega(n))$ possible values for x (or an integral of comparable complexity, if π is not discrete).

1.2. Average-case complexity. Our main question above lies in the realm of *average-case computational complexity*, as we are interested in inherent limitations on both the runtime and statistical properties of algorithms over a particular distribution of random inputs. While such considerations are prevalent in high-dimensional statistics, we currently lack the tools to prove a lower bound on the runtime of *all* algorithms in these settings, even under a standard complexity assumption such as $P \neq NP$. Instead, a rich landscape of frameworks have emerged for giving various forms of “rigorous evidence” for computational hardness of statistical tasks. Most approaches either establish hardness conditional on the assumed hardness of a “standard” statistical problem such as *planted clique*, or establish unconditional failure of some restricted class of methods such as statistical query (SQ) algorithms or the sum-of-squares (SoS) hierarchy. Some pioneering works in this area include [39, 23, 29, 18, 34, 15], and we also refer the reader to [63, 46, 20, 33] for further exposition.

In this work we explore inherent computational barriers from the perspective of the *low-degree polynomial framework*, which first arose from the work of [15, 38, 37, 36] (see also the survey [46]) and was later refined and extended in various directions, e.g. [65, 32, 64, 44]. This approach amounts to studying the power and limitations of algorithms that can be represented as low-degree polynomial functions of the input variables (in our case, the entries of Y). The polynomial degree is thought of as a measure of the algorithm’s complexity and a proxy for runtime, with degree D corresponding to runtime roughly n^D (up to $\log n$ factors in the exponent), which is the number of terms in such a polynomial. We refer to the *low-degree conjecture* as the informal belief that the class of degree- D polynomials is as powerful as *all* algorithms of the corresponding runtime, for a particular style of high-dimensional testing problems. This heuristic is by now well-established as a reliable method for predicting and rigorously vindicating suspected computational barriers in a wide array of statistical problems. For instance, in the spiked Wigner testing problem that we consider in this work, with any prior π of bounded support, the following phase transition is known for low-degree polynomials: if $\lambda < 1$ then *any* degree- $o(n/\log n)$ polynomial *fails* to achieve strong detection (in an appropriate sense), whereas degree $\omega(1)$ suffices when $\lambda > 1$ [46]. We consider this as rigorous evidence that strong detection requires exponential runtime $\exp(n^{1-o(1)})$ when $\lambda < 1$, at least for a broad class of known approaches. In contrast, recall that strong detection is possible in polynomial time when $\lambda > 1$.

1.3. *Our contributions.* While a myriad of prior work has shown hardness in the low-degree framework for strong (or weak) detection in various models, in this work we seek to answer an even more precise question, as we are concerned with the *exact error probability* in a regime where the type I and type II error rates will both converge to nontrivial constants. To our knowledge, no existing tools allow us to approach this question — not in the low-degree framework nor in any other framework for average-case complexity. Our main contribution is a new way to argue, based on low-degree polynomials, about the exact error probability achievable by computationally efficient algorithms. We illustrate this with the following result in the spiked Wigner model.

THEOREM 1.2 (Main result, informal). *Consider the spiked Wigner testing problem (Definition 1.1) with any prior π of bounded support and any $\lambda < 1$. Assuming a natural strengthening of the low-degree conjecture (Conjecture 2.15), any test with error tradeoff (between types I and II) asymptotically better than LSS requires runtime $\exp(n^{1-o(1)})$.*

The precise statement of this result is presented in Section 2.3, particularly Corollary 2.21.

One consequence of our result is a certain *computational universality* with respect to the prior. While the optimal statistical performance (namely the threshold λ^* for strong detection) depends on π , our result shows that the best *computationally efficient* test only uses the spectrum and thus its performance depends only on λ (not π). Put another way, the spectrum is a *sufficient statistic* for computationally-efficient testing. It was known previously that the low-degree threshold for strong detection does not depend on the prior [46], and our result extends this to weak detection.

The computational universality for testing is in contrast to the related task of *estimating* the rank-1 spike. For the estimation problem, the best known computationally efficient algorithm uses *approximate message passing (AMP)* [27, 16], which achieves nontrivial mean squared error whenever $\lambda > 1$ and this error depends on both λ and π [31, 54]. For more on the estimation problem, we refer the reader to [52] and references therein.

The argument we use to prove Theorem 1.2 can be specialized to the case where there is no restriction on runtime, in which case it yields a new approach for establishing the statistically optimal error tradeoff. To our knowledge, this method has not appeared before in the literature. We illustrate this by giving an alternate re-proof of the following result of [1, 28].

THEOREM 1.3 (Special case of [28], informal). *Consider the spiked Wigner testing problem (Definition 1.1) where π is the Rademacher prior (uniform on $\{-1, +1\}$) and $\lambda < 1$. No test (regardless of runtime) has error tradeoff asymptotically better than LSS.*

The precise statement of this result is presented in Section 2.2, particularly Theorem 2.5. The proof of [28] characterizes the limiting distribution of the likelihood ratio by leveraging some powerful machinery of Guerra and Talagrand from spin-glass theory. The same result can be deduced from the earlier work [1, Proposition 2.2], which is specific to the Rademacher prior and takes a combinatorial approach called *cluster expansion*. Our proof is quite different, and arguably more elementary. We require two ingredients: first, a sharp bound on the second moment of the likelihood ratio; and second, a positive result showing achievability of some tradeoff curve (between errors of types I and II) that “saturates” the first bound. In the case of Theorem 1.3, the positive result is the analysis of LSS from [22]. In a nutshell, the standard approach requires a direct analysis of the likelihood ratio, whereas our approach requires the analysis of *any* test combined with a matching second moment bound.

Notation. By default, asymptotic notation refers to the limit $n \rightarrow \infty$ with all other parameters held fixed (aside from those that are explicitly allowed to scale with n , such as $D = D_n$). In other words, notation such as $O(\cdot)$, $\Omega(\cdot)$, $o(\cdot)$, $\omega(\cdot)$ may hide factors depending on constants such as λ, π . We use $\text{poly}(n)$ as shorthand for $n^{O(1)}$, and $\text{polylog}(n)$ as shorthand for $(\log n)^{O(1)}$. The term “polynomial time” refers to an algorithm of runtime $\text{poly}(n)$.

2. Main Results.

2.1. Background.

2.1.1. *ROC curve.* We recall some basic notions from the theory of hypothesis testing, referring the reader to [47] for a standard reference. Given two distributions \mathcal{P}, \mathcal{Q} on a set Ω , we consider testing the null hypothesis $Y \sim \mathcal{Q}$ against the (simple) alternative $Y \sim \mathcal{P}$. (Some authors use the opposite meaning of \mathcal{P} and \mathcal{Q} but we use the mnemonic \mathcal{P} = “planted.”) A *test* is a (possibly random) function $t : \Omega \rightarrow \{\mathbf{p}, \mathbf{q}\}$, where the symbols \mathbf{p}, \mathbf{q} encode the assertion that Y was drawn from \mathcal{P} or \mathcal{Q} respectively. The *size* $\alpha \in [0, 1]$ of a test (also called the *type I error rate* or *false positive rate*) is defined as

$$(4) \quad \alpha = \mathcal{Q}(t(Y) = \mathbf{p}),$$

and the *power* $\beta \in [0, 1]$ (or *true positive rate*) is defined as

$$(5) \quad \beta = \mathcal{P}(t(Y) = \mathbf{p}).$$

(Here we follow the convention of [47] but note that some authors use β for the type II error rate, which is $1 - \beta$ in our notation.) Constrained to a given value of α , it is desirable to find a test maximizing β .

For a class of tests \mathcal{C} , let $N = N_{\mathcal{C}} \subseteq [0, 1]^2$ denote the set of (α, β) pairs for which there exists a test $t \in \mathcal{C}$ satisfying (4) and (5). There are trivial tests achieving the points $(0, 0)$ and $(1, 1)$, so these are contained in N (assuming the class \mathcal{C} contains these trivial tests). Also, by considering probabilistic mixtures of two tests (i.e., run t_1 with probability p and t_2 with probability $1 - p$), it is clear that N is a convex set (assuming \mathcal{C} is closed under probabilistic mixtures). Finally, by flipping the output of a test, N is symmetric with respect to the point $(1/2, 1/2)$ in the sense that $(\alpha, \beta) \in N$ if and only if $(1 - \alpha, 1 - \beta) \in N$ (again assuming \mathcal{C} is closed under flipping the output). We therefore restrict our attention to (α, β) pairs in the upper triangle

$$\Delta := \{(\alpha, \beta) \in [0, 1]^2 : \alpha \leq \beta\}.$$

The curve $\phi : [0, 1] \rightarrow [0, 1]$ that bounds the upper edge of the region N , namely

$$\phi(\alpha) = \sup\{\beta : (\alpha, \beta) \in N\},$$

is called the *ROC (receiver operating characteristic) curve* for the class \mathcal{C} . The properties of N discussed above imply that ϕ is increasing and concave, with $\phi(\alpha) \geq \alpha$. This curve describes the possible tradeoffs between type I and type II errors achievable by \mathcal{C} . At one extreme, \mathcal{C} could be the class of tests that thresholds a particular statistic at varying thresholds. At the other extreme, \mathcal{C} could be the class of *all* tests, in which case the corresponding *optimal* ROC curve describes the fundamental limits for testing error.

REMARK 2.1 (Terminology). *The ROC curve is a popular notion in machine learning. A related notion from classical statistics is the power function of a test, although typically this describes the power as a function of the alternative hypothesis (e.g., the parameter λ in the spiked Wigner model), rather than a function of α . Similarly, the optimal ROC curve (when \mathcal{C} is the class of all tests) is related to the notion of power envelope (or envelope power function).*

We will be interested in an asymptotic setting where $\mathcal{P} = \mathcal{P}_n$ and $\mathcal{Q} = \mathcal{Q}_n$ are sequences of distributions, namely the spiked Wigner distributions $\mathcal{Q}_n = \mathbb{P}_{0,\pi,n}$ and $\mathcal{P}_n = \mathbb{P}_{\lambda,\pi,n}$ for a fixed choice of λ, π . Further, we will often restrict ourselves to tests $t = t_n$ with a given runtime \mathcal{T}_n , where \mathcal{T}_n may stand for a class of asymptotic runtimes such as $O(n)$ or $\text{poly}(n)$.

DEFINITION 2.2. *Fix sequences \mathcal{P}_n (alternative) and \mathcal{Q}_n (null), and a runtime bound \mathcal{T}_n . A point $(\alpha, \beta) \in \Delta$ is asymptotically achievable in time \mathcal{T}_n if there is a sequence of tests $t = t_n$ computable in time \mathcal{T}_n such that*

$$\mathcal{Q}_n(t_n(Y) = \mathbf{p}) \leq \alpha + o(1) \quad \text{and} \quad \mathcal{P}_n(t_n(Y) = \mathbf{p}) \geq \beta - o(1)$$

as $n \rightarrow \infty$.

We emphasize that α, β do not depend on n . For given sequences $\mathcal{P}_n, \mathcal{Q}_n$ and a given runtime bound \mathcal{T}_n , one can define the set of asymptotically achievable (α, β) pairs analogous to N , as well as the associated asymptotic ROC curve.

REMARK 2.3 (Model of computation). *To rigorously define the notion of runtime, we must specify a model of computation for real-valued inputs. The arguments we use in this paper are not particularly sensitive to the model of computation, and for our purposes we need a model that is powerful enough to implement the LSS test in polynomial time (Theorem 2.10) but weak enough that we believe the strong low-degree conjecture (Conjecture 2.15). The PhD thesis of Hopkins [36], which introduced the original low-degree conjecture, proposes (in a footnote on pg. 80) to use the real RAM model. For concreteness, one can take the real RAM model as the model of computation throughout this paper, specifically the variant defined by Blum–Shub–Smale [19]. The proof of Theorem 2.10 will include an explanation of how to implement the LSS test in this model. The real RAM model is an abstract model of computation where each memory cell can hold a real number, and the operations of addition, subtraction, multiplication, division, and comparison can be performed in a single step. This is arguably an unrealistic model for real-world digital computers because issues of numerical precision are ignored, but it is a convenient and popular theoretical abstraction. Plus, we are ultimately proving a negative result, which only becomes stronger when using a stronger model of computation. As a final note, we will allow algorithms to depend on the known parameters $\lambda, \pi, \alpha, \beta$ (rather than, say, taking these as input), which means that any real-valued constants depending on these parameters may be “hard-coded” into the algorithm.*

2.1.2. Likelihood ratio and second moment. Again consider sequences of distributions $\mathcal{P}_n, \mathcal{Q}_n$ on a set Ω_n , and further assume \mathcal{P}_n is absolutely continuous with respect to \mathcal{Q}_n for each n . The *likelihood ratio* $L_n = L_n(Y)$ is defined to be the Radon–Nikodym derivative $d\mathcal{P}_n/d\mathcal{Q}_n$. Working in the function space $\mathcal{L}^2(\mathcal{Q}_n)$ with inner product $\langle f, g \rangle := \mathbb{E}_{Y \sim \mathcal{Q}_n}[f(Y) \cdot g(Y)]$ and norm $\|f\| := \sqrt{\langle f, f \rangle}$, an important quantity is the *squared norm* (or second moment) of the likelihood ratio:

$$\|L_n\|^2 = \mathbb{E}_{Y \sim \mathcal{Q}_n}[L_n(Y)^2] = \mathbb{E}_{Y \sim \mathcal{P}_n}[L_n(Y)].$$

It is a standard fact that asymptotic bounds on $\|L_n\|$ as $n \rightarrow \infty$ have implications for statistical indistinguishability of \mathcal{P}_n and \mathcal{Q}_n (see e.g. [53, Lemma 2]), namely:

- If $\|L_n\| = O(1)$ then strong detection is impossible, or in other words, $(\alpha, \beta) = (0, 1)$ is not asymptotically achievable (by any test, regardless of runtime).
- If $\|L_n\| = 1 + o(1)$ then weak detection is impossible, or in other words, no $(\alpha, \beta) \in \Delta$ with $\alpha < \beta$ is asymptotically achievable.

(We always have $\|L_n\| \geq 1$, using Jensen's inequality and the fact $\mathbb{E}_{\mathcal{Q}_n}[L_n] = 1$.) This gives a powerful tool for ruling out strong or weak detection via a relatively tractable second moment calculation, which is carried out for the spiked Wigner model in [53, 14, 61].

However, to exactly pin down the optimal asymptotic ROC curve in a regime where weak (but not strong) detection is possible, a more refined strategy is needed. The standard approach is the following. The Neyman–Pearson lemma implies that the optimal ROC curve is swept out by tests that threshold $L_n(Y)$, as the threshold is varied. It therefore suffices to determine the limiting distribution of L_n under both \mathcal{P}_n and \mathcal{Q}_n . This is often shown directly for one of the two hypotheses (\mathcal{P}_n or \mathcal{Q}_n), and then the distribution of L_n under the other can be deduced immediately via Le Cam's third lemma. This approach is carried out for the spiked Wigner model in [28].

2.1.3. Low-degree testing. Luckily, some parts of the above theory have natural analogues in the setting where we restrict our attention to computationally efficient tests — or rather, to low-degree polynomial tests as a proxy for this [38, 37, 36] (see also the survey [46]). We consider the same asymptotic setting as above and additionally assume the domain Ω_n is a subset of \mathbb{R}^M for some $M = M_n$ so that we may speak of (multivariate) polynomial functions $f : \Omega_n \rightarrow \mathbb{R}$. The analogue of the likelihood ratio L (suppressing n -dependence for ease of notation) is the *low-degree likelihood ratio* $L^{\leq D}$, which is the orthogonal projection in $\mathcal{L}^2(\mathcal{Q})$ of L onto the subspace of degree- D polynomial functions. The norm of $L^{\leq D}$ plays a similar role as its statistical analogue (see e.g. [11, Proposition 6.2]), namely:

- If $\|L^{\leq D}\| = O(1)$ for some $D = D_n$ then no degree- D polynomial $f = f_n$ achieves *strong separation* between \mathcal{P} and \mathcal{Q} , defined as

$$\sqrt{\max \left\{ \text{Var}_{\mathcal{P}}[f], \text{Var}_{\mathcal{Q}}[f] \right\}} = o \left(\left| \mathbb{E}_{\mathcal{P}}[f] - \mathbb{E}_{\mathcal{Q}}[f] \right| \right)$$

as $n \rightarrow \infty$. Note that strong separation is a natural sufficient condition for strong detection by thresholding $f(Y)$.

- If $\|L^{\leq D}\| = 1 + o(1)$ for some $D = D_n$ then no degree- D polynomial $f = f_n$ achieves *weak separation* between \mathcal{P} and \mathcal{Q} . Weak separation has the same definition as strong separation but with $O(\dots)$ in place of $o(\dots)$, and is a natural sufficient condition for weak detection using the value of $f(Y)$ [11, Proposition 6.1].

Results of this form are considered “evidence” that the associated strong/weak detection problem is inherently hard for algorithms of the corresponding runtime.

CONJECTURE 2.4 (Low-degree conjecture, informal). *If degree- D polynomials fail to solve a testing problem (in the sense of strong/weak separation) for some $D = \omega(\log n)$, then there is no polynomial-time algorithm for the associated strong/weak (respectively) detection task. In general, if degree- D polynomials fail then there is no algorithm of runtime $\exp(\tilde{\Omega}(D))$ where $\tilde{\Omega}(\cdot)$ hides a polylog(n) factor.*

This informal conjecture is inspired by [36, Hypothesis 2.1.5 & Conjecture 2.2.4], along with the fact that low-degree polynomials capture the best known algorithms for a wide variety of statistical tasks. The conjecture appears to hold up for distributions \mathcal{P}, \mathcal{Q} of a particular style that often arises in high-dimensional statistics (including the spiked Wigner and Wishart models [12, 46, 26, 10]), but it does not hold for *all* distributions \mathcal{P}, \mathcal{Q} , and we refer the reader to [36, 46, 35, 45, 43, 67, 24] for further discussion on which distributions are appropriate. Given the current state of average-case complexity theory, we do not realistically hope to

prove any variant of the low-degree conjecture. Rather, its purpose is to serve as a guide for making principled conjectures about average-case complexity in settings where we would otherwise have no way to make progress.

The above framework provides a useful tool for probing the computational feasibility of strong or weak detection, but our goal in this work is to ask a more refined question: we seek the best possible ROC curve achievable by algorithms of a given runtime, in a regime where weak (but not strong) detection is possible by this class of algorithms. As discussed previously, the statistical analogue of this question is traditionally attacked via the Neyman–Pearson lemma combined with Le Cam’s third lemma. This approach does not seem viable in the computationally bounded setting because there is no analogue of Neyman–Pearson, that is, there is no guarantee that thresholding $L^{\leq D}$ achieves the best possible ROC curve among low-degree tests. We will overcome this by taking a different approach, explained in the following sections.

2.2. Our Approach: Statistical Limits. To explain our approach, we first consider the purely statistical question of determining the optimal ROC curve with no constraints on runtime. This section serves as a “warm-up” to the more general framework presented in the next section, which accounts for computational complexity. The ideas in this section may also be of independent interest, as they provide a way to indirectly characterize the optimal ROC curve without analyzing the distribution of the likelihood ratio.

We will illustrate our approach by re-proving the following known result in the spiked Wigner model. For $\lambda > 0$ and $\alpha \in [0, 1]$, define

$$(6) \quad \phi_\lambda(\alpha) := 1 - \Phi \left[\Phi^{-1}(1 - \alpha) - \sqrt{\frac{1}{2} \log \left(\frac{1}{1 - \lambda^2} \right)} \right]$$

where Φ denotes the standard normal CDF function, and we take the conventions $\phi_\lambda(0) = 0$, $\phi_\lambda(1) = 1$. This is the ROC curve achieved by linear spectral statistics (LSS) [22]. We will show this is statistically optimal for the Rademacher spike prior.

THEOREM 2.5 (Special case of [28]). *Consider the spiked Wigner testing problem (Definition 1.1) where π is uniform on $\{-1, +1\}$ and $\lambda \in (0, 1)$. Any $(\alpha, \beta) \in [0, 1]^2$ with $\beta > \phi_\lambda(\alpha)$ is not asymptotically achievable (in the sense of Definition 2.2), regardless of runtime.*

Since we are considering the Rademacher prior, the model lacks spherical symmetry, which precludes approaches such as [57, 7, 8]. The proof of [28] uses machinery developed by Guerra and Talagrand in their study of the Sherrington–Kirkpatrick spin-glass model, and the same result also follows from [1, Proposition 2.2], which takes a combinatorial approach specific to the Rademacher prior. Our alternate proof instead uses the analysis of LSS from [22], combined with a relatively simple second moment calculation from [61].

We begin with the second moment calculation.

THEOREM 2.6 (Special case of [61], Theorem 3.10). *Consider the spiked Wigner testing problem (Definition 1.1) where π is uniform on $\{-1, +1\}$ and $\lambda \in (0, 1)$. The likelihood ratio $L_\lambda := d\mathbb{P}_\lambda/d\mathbb{P}_0$ has $\mathcal{L}^2(\mathbb{P}_0)$ -norm (as defined in Section 2.1)*

$$(7) \quad \lim_{n \rightarrow \infty} \|L_\lambda\| = (1 - \lambda^2)^{-1/4}.$$

The utility of this result stems from the following well-known operational definition for the norm of the likelihood ratio, which is a special case of (13) below:

$$(8) \quad \|L_\lambda\| = \sup_f \frac{\mathbb{E}_{Y \sim \mathbb{P}_\lambda}[f(Y)]}{\sqrt{\mathbb{E}_{Y \sim \mathbb{P}_0}[f(Y)^2]}},$$

where the supremum is over all real-valued functions $f \in \mathcal{L}^2(\mathbb{P}_0)$. The fact (8) on its own places some constraints on what the optimal ROC curve can look like. For instance, the relation

$$(9) \quad \frac{\beta^2}{\alpha} + \frac{(1-\beta)^2}{1-\alpha} \leq \|L_\lambda\|^2$$

must be satisfied by every achievable pair $(\alpha, \beta) \in (0, 1)^2$ [61, Proposition 2.5]. This is a weaker condition than the desired one, $\beta \leq \phi_\lambda(\alpha)$. However, the information we get from (8) is more subtle than (9) alone. At this point there remain many viable candidates for the optimal ROC curve — one of which is ϕ_λ — that are incomparable to each other. To illustrate, each (α, β) pair satisfying (9) with equality is achieved by one such viable ROC curve, namely a straight line from $(0, 0)$ to (α, β) followed by a straight line from (α, β) to $(1, 1)$. However, it is *not* possible that *all* the (α, β) pairs satisfying (9) are simultaneously achieved by the true ROC curve, or else we could construct a function f that yields a too-good-to-be-true value for the ratio in (8). We will use the positive result on the achievability of ϕ_λ to disambiguate between these candidate ROC curves and conclude that ϕ_λ is in fact the true ROC curve.

Now in more detail: as we will show (see Section 2.5), the fact that ϕ_λ is achievable allows us to construct a function f that makes the ratio in (8) asymptotically equal to

$$\text{val}(\phi_\lambda) := \sqrt{\int_0^1 (\phi'_\lambda(\alpha))^2 d\alpha} = (1 - \lambda^2)^{-1/4},$$

saturating (7). Any hypothetical improvement to ϕ_λ , even at a single point, would allow us to construct an even better function f , leading to an even larger value for the ratio, contradicting (7). This is formalized in Proposition 2.7 below and illustrated in Figure 1.

PROPOSITION 2.7 (Special case of Proposition 2.19). *Given sequences of distributions $\mathcal{P} = \mathcal{P}_n$ and $\mathcal{Q} = \mathcal{Q}_n$ on $\Omega = \Omega_n$, consider testing the null hypothesis $Y \sim \mathcal{Q}$ against the alternative $Y \sim \mathcal{P}$. For some $\lambda \in (0, 1)$, suppose we have the following positive result: for every $\alpha, \beta \in [0, 1]$ with $\alpha \leq \beta < \phi_\lambda(\alpha)$, there is a test that asymptotically achieves (α, β) , in the sense of Definition 2.2. Suppose further there exists some $(\alpha^*, \beta^*) \in [0, 1]^2$ with $\beta^* > \phi_\lambda(\alpha^*)$ that is also asymptotically achievable. Then there exists a function $f = f_n : \Omega_n \rightarrow \mathbb{R}$ such that*

$$(10) \quad \liminf_{n \rightarrow \infty} \frac{\mathbb{E}_{Y \sim \mathcal{P}}[f(Y)]}{\sqrt{\mathbb{E}_{Y \sim \mathcal{Q}}[f(Y)^2]}} > \text{val}(\phi_\lambda).$$

Our desired conclusion now follows by combining the ingredients above.

PROOF OF THEOREM 2.5. We will apply Proposition 2.7 with $\mathcal{P} = \mathbb{P}_\lambda$ and $\mathcal{Q} = \mathbb{P}_0$. The required positive result follows from the LSS analysis of [22], which is stated formally as Theorem 2.10 in the next section. Assume on the contrary that some (α^*, β^*) with $\beta^* > \phi_\lambda(\alpha^*)$ is asymptotically achievable. By Proposition 2.7 we obtain f such that

$$\liminf_{n \rightarrow \infty} \frac{\mathbb{E}_{Y \sim \mathbb{P}_\lambda}[f(Y)]}{\sqrt{\mathbb{E}_{Y \sim \mathbb{P}_0}[f(Y)^2]}} > \text{val}(\phi_\lambda),$$

and the computation $\text{val}(\phi_\lambda) = (1 - \lambda^2)^{-1/4}$ is Lemma 2.12 in the next section. This now contradicts (7),(8). \square

The next section contains a more general form of this argument that also considers computational complexity. The general form of Proposition 2.7 (i.e., Proposition 2.19) applies to a wide class of curves ϕ rather than simply ϕ_λ . We note, however, that ROC curves of the form ϕ_λ are “common” in that they arise from thresholding a statistic whose limiting distributions under null and alternative are two different Gaussians with the same variance.

For simplicity, we have focused here on recovering the result of [28] for one particular spike prior — Rademacher. The main obstacle to recovering the full result of [28] for other priors π is establishing $\|L_\lambda\| = O(1)$ for all λ below the critical threshold $\lambda^*(\pi)$. For some priors this is simply not true and would need to be replaced with a conditional second moment calculation. See [14, 61] for some results of this form, albeit only reaching the sharp threshold $\lambda^*(\pi)$ for certain priors.

REMARK 2.8. *Theorem 2.5 can be generalized to a modified spiked Wigner testing problem where the noise matrix W has diagonal entries $W_{ii} \sim \mathcal{N}(0, \sigma^2/n)$ instead of $\mathcal{N}(0, 2/n)$ for a constant $\sigma > 0$, and everything else is unchanged. This case is covered by [28], and our method also extends to this case in a straightforward way. As in the proof of Theorem 2.10, the results of [22] imply that LSS asymptotically achieves the ROC curve $\phi_{\lambda,\sigma}(\alpha) := 1 - \Phi \left[\Phi^{-1}(1 - \alpha) - \sqrt{\mu/2} \right]$ with $\mu := -\log(1 - \lambda^2) + \lambda^2(2 - \sigma^2)/\sigma^2$. As in the proof of Lemma 2.12, $\text{val}(\phi_{\lambda,\sigma}) = \exp(\mu/4) = (1 - \lambda^2)^{-1/4} \exp(\lambda^2(2 - \sigma^2)/(4\sigma^2))$. For the analysis of the likelihood ratio $L_{\lambda,\sigma}$, we have the standard formula (see Lemma 1 of [14])*

$$\|L_{\lambda,\sigma}\|^2 = \mathbb{E}_{x,x'} \exp \left(\frac{\lambda^2 n}{2} \frac{\langle x, x' \rangle^2}{\|x\|^2 \|x'\|^2} + \lambda^2 n \left(\frac{1}{\sigma^2} - \frac{1}{2} \right) \sum_{i=1}^n \frac{x_i^2 (x'_i)^2}{\|x\|^2 \|x'\|^2} \right)$$

where x is i.i.d. Rademacher and x' is an independent copy of x . This differs from the original second moment $\|L_\lambda\|^2$ due to the additional second term in $\exp(\dots)$, which (due to the Rademacher prior) is equal to the constant $\lambda^2(2 - \sigma^2)/(2\sigma^2)$, i.e., $\|L_{\lambda,\sigma}\|^2 = \|L_\lambda\|^2 \exp(\lambda^2(2 - \sigma^2)/(2\sigma^2)) \rightarrow (1 - \lambda^2)^{-1/2} \exp(\lambda^2(2 - \sigma^2)/(2\sigma^2))$. Therefore, the limiting value of $\|L_{\lambda,\sigma}\|$ matches $\text{val}(\phi_{\lambda,\sigma})$. The conclusion for the modified model is identical to Theorem 2.5 except the optimal ROC curve is $\phi_{\lambda,\sigma}$ in place of ϕ_λ .

To what extent is our new proof of Theorem 2.5 simpler or more broadly applicable than the existing proofs of [1, 28]? Certainly the second moment bound (Theorem 2.6) is more elementary than the interpolation methods used in [28], and it also seems more broadly applicable than the combinatorial approach of [1] (see [61]). On the other hand, our argument also relies on the LSS analysis of [22], which is rather involved (although quite different from [1, 28]). At the very least, our approach gives the flexibility to trade off one argument for another, and if one is planning to do the LSS analysis anyway then the matching impossibility result is obtained essentially for free. Perhaps the greatest strength of our approach is its ability to generalize to the computationally bounded setting, which we discuss next.

2.3. Our Approach: Computational Limits. Building on ideas from the previous section, we now outline our approach for arguing that no (computationally) efficient algorithm can beat the ROC curve of LSS in the spiked Wigner model. As this is a statement about average-case complexity, there will inevitably need to be a conjecture involved (Conjecture 2.15). Our final conclusion, conditional on the conjecture, is Corollary 2.21. The proofs of the results in this section are deferred to Section 3.

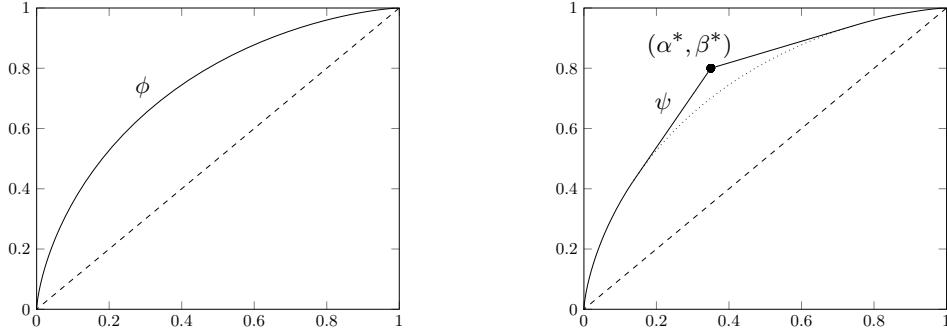


Fig 1: Illustration of the proof of Theorem 2.5. **Left:** The curve $\phi = \phi_\lambda$ for $\lambda = 0.9$. All points (α, β) below the curve ϕ (and above the “trivial” dashed line) are achievable. **Right:** If hypothetically there were an achievable point (α^*, β^*) above the curve ϕ then every point below ψ — the upper concave envelope of ϕ and (α^*, β^*) — would also be achievable. The improved curve ψ would allow us to construct a function f that makes the ratio in (8) equal to $\text{val}(\psi)$, which is strictly greater than $\text{val}(\phi_\lambda) = (1 - \lambda^2)^{-1/4}$, contradicting (7).

Throughout this section, we consider the spiked Wigner testing problem described in Definition 1.1: we are testing the null hypothesis $Y \sim \mathbb{P}_0$ against a specific alternative $Y \sim \mathbb{P}_\lambda$, in the regime $n \rightarrow \infty$ with λ, π fixed. We place the following mild assumptions on the spike prior π .

ASSUMPTION 2.9. *The spike prior π is a probability distribution on \mathbb{R} such that:*

- *(Centered)* $\mathbb{E}[\pi] = 0$,
- *(Unit variance)* $\mathbb{E}[\pi^2] = 1$,
- *(Subgaussian)* there exists a constant $c > 0$ such that $\mathbb{E}[\exp(t\pi)] \leq \exp(ct^2)$ for all $t \in \mathbb{R}$.
(This is satisfied by any distribution with bounded support, by Hoeffding’s lemma.)

While we focus on the spiked Wigner model here, it will be clear that our approach can potentially be applied more generally.

2.3.1. Step 1: Guess the ROC curve and prove a positive result. The first step in our approach is to “guess” the ROC curve that we aim to show is optimal among efficient algorithms. While our final goal is to show a hardness result, it will (perhaps counterintuitively) be important to prove the corresponding positive result, that is, our ROC curve is achievable by an efficient algorithm. In our case, we choose the ROC curve (6) that is known to be achievable by LSS, which we repeat here for convenience:

$$(11) \quad \phi_\lambda(\alpha) := 1 - \Phi \left[\Phi^{-1}(1 - \alpha) - \sqrt{\frac{1}{2} \log \left(\frac{1}{1 - \lambda^2} \right)} \right]$$

where Φ denotes the standard normal CDF function, and we take the conventions $\phi_\lambda(0) = 0$, $\phi_\lambda(1) = 1$. We now state the corresponding positive result, which will be extracted from [22] in a straightforward way; the details of the proof are deferred to the Supplement.

THEOREM 2.10. *Fix any $\lambda \in (0, 1)$, any spike prior π , and any $\alpha, \beta \in [0, 1]$ with $\alpha \leq \beta \leq \phi_\lambda(\alpha)$. There is a polynomial-time algorithm that asymptotically achieves (α, β) for the spiked Wigner testing problem (in the sense of Definition 2.2).*

We have defined spike priors to have mean 0 and variance 1 throughout the paper, but Theorem 2.10 holds with no assumptions on π (as long as $x \neq 0$ with probability $1 - o(1)$). Due to GOE rotational invariance, the LSS test is not affected by the direction of the spike.

2.3.2. Step 2: Matching bound on $\|L^{\leq D}\|$. A key step in our approach is to establish a sharp bound on the norm of the low-degree likelihood ratio (defined in Section 2.1) that matches the “value” of our ROC curve, defined as follows.

DEFINITION 2.11. *For a function $\phi : [0, 1] \rightarrow [0, 1]$ whose derivative ϕ' exists on $(0, 1)$ except at finitely many points, define*

$$(12) \quad \text{val}(\phi) = \sqrt{\int_0^1 (\phi'(\alpha))^2 d\alpha},$$

provided the (improper Riemann) integral exists.

The intuition for this value is the following fact, which will be implicit in the proof of Proposition 2.19 and is discussed further in Section 2.5: if some (concave) ϕ is achievable then we can construct a function f that makes a particular ratio — namely $R_\lambda(f)$ defined in (14) below — equal to $\text{val}(\phi)$.

LEMMA 2.12. *For the function ϕ_λ defined in (11),*

$$\text{val}(\phi_\lambda) = (1 - \lambda^2)^{-1/4}.$$

The details of the above calculation are deferred to the Supplement.

THEOREM 2.13. *Fix any $\lambda \in (0, 1)$ and any π satisfying Assumption 2.9. Let $L_\lambda = L_{\lambda, \pi, n}$ denote the spiked Wigner likelihood ratio $d\mathbb{P}_\lambda/d\mathbb{P}_0$. Let $D = D_n$ be any sequence satisfying $D = \omega(1)$ and $D = o(n/\log n)$. The $\mathcal{L}^2(\mathbb{P}_0)$ -norm of the degree- D likelihood ratio (as defined in Section 2.1) satisfies*

$$\lim_{n \rightarrow \infty} \|(L_\lambda)^{\leq D}\| = (1 - \lambda^2)^{-1/4}.$$

Crucially, the value on the right-hand side above matches $\text{val}(\phi_\lambda)$. The assumption $D = \omega(1)$ is only needed for the lower bound, i.e., the upper bound $\|(L_\lambda)^{\leq D}\| \leq (1 - \lambda^2)^{-1/4} + o(1)$ still holds without this assumption. The proof of Theorem 2.13 builds on some existing tools for similar second moment calculations [61, 12, 46, 10, 11].

Theorem 2.13 gives a surprisingly sharp phase transition: for a fixed $\lambda \in (0, 1)$, the limiting value of $\|(L_\lambda)^{\leq D}\|$ has no dependence on D when D ranges between $\omega(1)$ and $o(n/\log n)$. Once D reaches order n , the value of $\|(L_\lambda)^{\leq D}\|$ can (for some priors) suddenly jump to size $\exp(\Omega(n))$ as discussed in Remark 2.14 below. This behavior is consistent with the fact that we do not know any algorithm of subexponential time $\exp(O(n^\delta))$ for a constant $\delta < 1$ that beats the ROC curve of LSS, but some priors do admit such an algorithm of exponential time $\exp(O(n))$.

REMARK 2.14. *Regarding the condition $D = o(n/\log n)$ in Theorem 2.13: while the log factor might be an artifact of the proof, it is necessary to assume $D = o(n)$, or else $\|(L_\lambda)^{\leq D}\|$ can diverge to infinity (for some choices of $\lambda \in (0, 1)$ and π). To see this, take π to be the sparse Rademacher prior which takes values $\pm 1/\sqrt{\rho}$ each with probability $\rho/2$ and value*

0 with probability $1 - \rho$, for a small constant $\rho > 0$ to be chosen later. With $x \in \mathbb{R}^n$ having entries i.i.d. from π and x' an independent copy of x , a standard formula (20) gives

$$\|(L_\lambda)^{\leq D}\|^2 = \sum_{d=0}^D \frac{1}{d!} \mathbb{E}_{x,x'} \left[\left(\frac{\lambda^2 n \langle x, x' \rangle^2}{2\|x\|^2\|x'\|^2} \right)^d \right] \geq \frac{1}{D!} \Pr\{x = x' \neq 0\} \left(\frac{\lambda^2 n}{2} \right)^D.$$

Using the bounds $\Pr\{x = x' \neq 0\} \geq \Pr\{x = x' = e_1/\sqrt{\rho}\} = (\rho/2)^2(1 - \rho)^{2(n-1)}$ and $D! \leq D^D$, the above is at least $\exp[D \log(\lambda^2 n/(2D)) + 2n \log(1 - \rho) + 2 \log(\rho/2)]$. Set $D = \lfloor \epsilon n \rfloor$ for a constant $\epsilon > 0$ small enough so that $\lambda^2/(2\epsilon) > 1$. Now choosing ρ small enough (depending on λ, ϵ), we have $\|(L_\lambda)^{\leq D}\|^2 = \exp(\Omega(n))$.

2.3.3. Step 3: Strengthening of the low-degree conjecture. As discussed in Section 2.1, the low-degree conjecture posits that (for certain testing problems) low-degree polynomials are at least as powerful as all algorithms of the corresponding runtime (where the correspondence is described in Conjecture 2.4). For instance, in the spiked Wigner model with $\lambda < 1$, we have from Theorem 2.13 that $\|(L_\lambda)^{\leq D}\| = O(1)$ for any $D = o(n/\log n)$, and we consider this to be evidence that strong detection requires runtime $\exp(n^{1-o(1)})$. It is well known (see [36, 46]) that the norm of the low-degree likelihood ratio admits the operational definition

$$(13) \quad \left\| \left(\frac{d\mathcal{P}}{d\mathcal{Q}} \right)^{\leq D} \right\| = \sup_{f: \deg(f) \leq D} \frac{\mathbb{E}_{Y \sim \mathcal{P}}[f(Y)]}{\sqrt{\mathbb{E}_{Y \sim \mathcal{Q}}[f(Y)^2]}},$$

where the supremum is over polynomials $f : \Omega \rightarrow \mathbb{R}$ of degree at most D . In fact, the supremum is achieved and the optimizer is $f = (d\mathcal{P}/d\mathcal{Q})^{\leq D}$. For our purposes, we introduce a natural refinement of the low-degree conjecture: we posit that low-degree polynomials perform at least as well as all algorithms of the corresponding runtime *in terms of the value of the ratio* on the right-hand side of (13). In the spiked Wigner case, this ratio is

$$(14) \quad R_\lambda(f) := \frac{\mathbb{E}_{Y \sim \mathbb{P}_\lambda}[f(Y)]}{\sqrt{\mathbb{E}_{Y \sim \mathbb{P}_0}[f(Y)^2]}}.$$

CONJECTURE 2.15. *Consider the spiked Wigner testing problem (Definition 1.1). Fix any $0 < \delta_1 < \delta_2$, any $\lambda \in (0, 1)$, and any π satisfying Assumption 2.9. Any sequence of functions $f = f_n$ computable in time $\exp(O(n^{\delta_1}))$ must satisfy*

$$(15) \quad \limsup_{n \rightarrow \infty} R_\lambda(f) \leq \limsup_{n \rightarrow \infty} \sup_{g: \deg(g) \leq n^{\delta_2}} R_\lambda(g).$$

Recalling (13) and Theorem 2.13, the specific case of interest for us will be the following: for any fixed $\epsilon > 0$, any $f = f_n$ computable in time $\exp(O(n^{1-\epsilon}))$ must satisfy

$$(16) \quad \limsup_{n \rightarrow \infty} R_\lambda(f) \leq (1 - \lambda^2)^{-1/4}.$$

To be conservative, we have stated Conjecture 2.15 only for the spiked Wigner model, but one can imagine making similar conjectures in other settings. As with the standard low-degree conjecture we cannot realistically hope to prove Conjecture 2.15, but we can continue to amass evidence in its favor.

REMARK 2.16. *To justify Conjecture 2.15, one must believe that low-degree polynomials capture the most powerful approaches in our algorithmic toolkit. To this end, it may be instructive to see how the best known efficient algorithm for maximizing $R_\lambda(f)$ —*

which uses LSS — can be implemented as a low-degree polynomial. Recall the LSS statistic $H(Y) := \sum_{i=1}^n h_\lambda(\mu_i)$ from (2). After appropriate shifting and scaling, this is shown in [22] to converge to a Gaussian limit: $aH(Y) + b \Rightarrow \mathcal{N}(\pm c, 1)$ where the plus sign holds under \mathbb{P}_λ and the minus sign holds under \mathbb{P}_0 , and $c := \sqrt{-\log(1 - \lambda^2)/8}$. Now let $f(Y) = p(aH(Y) + b)$ and choose the optimal function p to maximize $R_\lambda(f)$, which turns out to be the likelihood ratio between $\mathcal{N}(c, 1)$ and $\mathcal{N}(-c, 1)$, i.e., $p(z) = \exp(2cz)$. This yields $R_\lambda(f) \rightarrow (1 - \lambda^2)^{-1/4}$, matching (16). To approximate this with a polynomial, consider polynomial approximations \hat{h}_λ and \hat{p} for h_λ and p , of degrees D_1 and D_2 respectively. Now $\hat{f}(Y) := \hat{p}(a \sum_i \hat{h}_\lambda(\mu_i) + b) = \hat{p}(a \cdot \text{Tr}(\hat{h}_\lambda(Y)) + b)$ is a polynomial of degree $D = D_1 D_2$ that approximates f . We expect that any slowly growing degree $D = \omega(1)$ should suffice to construct such an \hat{f} achieving $R_\lambda(\hat{f}) \rightarrow (1 - \lambda^2)^{-1/4}$. While we have not attempted to rigorously analyze this polynomial approximation, we do know from Theorem 2.13 that some degree- $\omega(1)$ polynomial achieves $R_\lambda \rightarrow (1 - \lambda^2)^{-1/4}$.

2.3.4. Steps 1,2,3 are sufficient. We now argue that the above steps are sufficient to deduce our desired conclusion. The idea is the following, which is similar to Figure 1. Our achievable ROC curve ϕ_λ allows us to construct an efficiently computable function f that achieves $R_\lambda(f) \approx \text{val}(\phi_\lambda)$. If hypothetically there were an efficient algorithm achieving some point (α, β) above the curve ϕ_λ , this could be combined with f to produce an efficiently computable g that achieves $R_\lambda(g) > \text{val}(\phi_\lambda)$, contradicting our refined low-degree conjecture.

We will state this part of the argument (Proposition 2.19) in high generality so that it can potentially be used for other problems beyond spiked Wigner in the future. We will require some conditions on the ROC curve ϕ (which in our case is ϕ_λ). As discussed in Section 2.1, an ROC curve should be increasing and concave with $\phi(1) = 1$. We also impose some additional technical conditions (some of which are likely removable, but they allow for a cleaner proof).

ASSUMPTION 2.17. *Assume $\phi : [0, 1] \rightarrow [0, 1]$ has the following properties:*

- $\phi(0) = 0$ and $\phi(1) = 1$,
- ϕ is continuous on $[0, 1]$ and differentiable on $(0, 1)$,
- ϕ' is continuous, strictly positive, and decreasing on $(0, 1)$,
- $\lim_{\alpha \rightarrow 0^+} \phi'(\alpha) = +\infty$ and $\lim_{\alpha \rightarrow 1^-} \phi'(\alpha) = 0$,
- the integral (12) defining $\text{val}(\phi)$ is finite.

Some immediate consequences are that ϕ is concave and strictly increasing, and $\phi(\alpha) \geq \alpha$ for all $\alpha \in [0, 1]$.

LEMMA 2.18. *For any $\lambda \in (0, 1)$, the function ϕ_λ defined in (11) satisfies Assumption 2.17.*

The proof of Lemma 2.18 is deferred to the Supplement.

The result below refers to a runtime bound \mathcal{T} , which may stand for a class of asymptotic runtimes such as $O(n)$ or $\text{poly}(n)$, or it may be ∞ (no bound on runtime). We will elaborate on this further in Remark 2.20 below.

PROPOSITION 2.19. *Given sequences of distributions $\mathcal{P} = \mathcal{P}_n$ and $\mathcal{Q} = \mathcal{Q}_n$ on $\Omega = \Omega_n$, consider testing the null hypothesis $Y \sim \mathcal{Q}$ against the alternative $Y \sim \mathcal{P}$. Suppose we have a function ϕ satisfying Assumption 2.17 and a runtime bound $\mathcal{T} = \mathcal{T}_n$ along with the corresponding positive result: for every $\alpha, \beta \in [0, 1]$ with $\alpha \leq \beta < \phi(\alpha)$, there is an algorithm that asymptotically achieves (α, β) in time \mathcal{T} , in the sense of Definition 2.2. Suppose further*

there exists some $(\alpha^*, \beta^*) \in [0, 1]^2$ with $\beta^* > \phi(\alpha^*)$ that is also asymptotically achievable in time \mathcal{T} . Then there is a function $f = f_n : \Omega_n \rightarrow \mathbb{R}$ computable in time $O(\mathcal{T})$ such that

$$(17) \quad \liminf_{n \rightarrow \infty} \frac{\mathbb{E}_{Y \sim \mathcal{P}}[f(Y)]}{\sqrt{\mathbb{E}_{Y \sim \mathcal{Q}}[f(Y)^2]}} > \text{val}(\phi).$$

REMARK 2.20 (Runtime). *The runtime bound \mathcal{T} may stand for a class of asymptotic runtimes such as $O(n)$ or $\text{poly}(n)$, or it may be ∞ (no bound on runtime). The runtime for the positive result need not be uniform in α, β , e.g., if $\mathcal{T} = \text{poly}(n)$ it would be fine to have a different algorithm for each α, β with runtime $O(n^{C(\alpha, \beta)})$.*

The precise meaning of the runtime $O(\mathcal{T})$ for f is as follows. A finite list of (α, β) pairs, depending only on ϕ, α^, β^* (not n), is chosen in advance and hard-coded into the algorithm. The algorithm runs the corresponding tests that asymptotically achieve these points, to get a vector of responses $s \in \{\mathfrak{p}, \mathfrak{q}\}^r$ where $r = r(\phi, \alpha^*, \beta^*)$. Finally, a hard-coded lookup table maps each of the $2^r = O(1)$ possible values for s to a corresponding output $f(Y)$.*

We now state our conclusion specialized to the case of spiked Wigner: ϕ_λ is the best possible ROC curve for computationally efficient algorithms.

COROLLARY 2.21. *Consider the spiked Wigner testing problem and assume Conjecture 2.15. Fix any $\epsilon > 0$, any $\lambda \in (0, 1)$, and any π satisfying Assumption 2.9. Any $(\alpha, \beta) \in [0, 1]^2$ with $\beta > \phi_\lambda(\alpha)$ is not asymptotically achievable in time $\exp(O(n^{1-\epsilon}))$.*

In general, the runtime bound $\exp(O(n^{1-\epsilon}))$ cannot be made any larger: for some λ, π pairs, strong detection is possible in time $\exp(O(n))$ by directly evaluating the likelihood ratio (3).

PROOF. We will apply Proposition 2.19 with $\mathcal{P} = \mathbb{P}_\lambda$, $\mathcal{Q} = \mathbb{P}_0$, $\phi = \phi_\lambda$ (which satisfies Assumption 2.17 by Lemma 2.18), and $\mathcal{T} = \exp(O(n^{1-\epsilon}))$. Theorem 2.10 gives the required positive result. Assume on the contrary that some (α^*, β^*) with $\beta^* > \phi_\lambda(\alpha^*)$ is asymptotically achievable in time $\exp(O(n^{1-\epsilon}))$. By Proposition 2.19 we obtain f computable in time $\exp(O(n^{1-\epsilon}))$ such that

$$(18) \quad \liminf_{n \rightarrow \infty} R_\lambda(f) > \text{val}(\phi_\lambda) = (1 - \lambda^2)^{-1/4},$$

where the final equality is Lemma 2.12. On the other hand, Conjecture 2.15 implies that any f computable in time $\exp(O(n^{1-\epsilon}))$ must satisfy

$$(19) \quad \limsup_{n \rightarrow \infty} R_\lambda(f) \leq \limsup_{n \rightarrow \infty} \sup_{g : \deg(g) \leq n^{1-\epsilon/2}} R_\lambda(g) = (1 - \lambda^2)^{-1/4},$$

where the final equality is Theorem 2.13 combined with (13). Together, (18) and (19) give a contradiction. \square

2.4. *Discussion.* In this work we have presented a general-purpose framework for arguing about the precise testing error achievable by computationally efficient algorithms — the first such framework, to our knowledge. Our method involves a two-stage argument: we first conjecture that low-degree polynomials are optimal for the related task of optimizing the ratio $R_\lambda(f)$, and we then provably characterize the best possible ROC curve for efficient algorithms conditional on this conjecture.

An alternative approach, and a potential direction for future work, would be to prove unconditionally that a particular class of tests (say, those based on thresholding low-degree polynomials) cannot surpass a particular ROC curve. At this point, even the simpler task

of showing that polynomial threshold functions fail to achieve *strong detection* is not fully understood, as the existing results along these lines have technical conditions that are likely unnecessary (see Section 4.1 of [46] and Section 2.3 of [50]).

While we view unconditional results as an important future direction, we also emphasize the merits of our current approach. Notably, the only low-degree calculation required is to bound $\|L^{\leq D}\|$, which tends to be relatively simple. This has allowed us to prove a very sharp result that captures the precise threshold at $\lambda = 1$, the precise asymptotic ROC curve ϕ_λ , and the nearly optimal degree $D = o(n/\log n)$. In contrast, other low-degree arguments not based on $\|L^{\leq D}\|$ — such as [65, 55], which deal with *estimation* rather than testing — tend to be more difficult, and the known results are less sharp. For this reason, the approach we have presented in this work appears to be an especially user-friendly tool that may be useful for other problems.

We now discuss some specific extensions of our result that appear to be feasible directions for future work. First, we have considered the spiked Wigner model with GOE noise, but one can imagine relaxing the assumptions on the noise matrix W . For instance, one non-GOE (but still Gaussian) noise model takes the diagonal entries of W to be $\mathcal{N}(0, \sigma^2/n)$ instead of $\mathcal{N}(0, 2/n)$, while the off-diagonal entries remain $\mathcal{N}(0, 1/n)$. The analysis of LSS has been extended to this case [22] (with a different ROC curve that depends on σ), and we have shown in Remark 2.8 how our method can prove *statistical* optimality of LSS when the spike prior is Rademacher. We expect it should also be possible to prove *computational* optimality of LSS for more general priors, generalizing our Corollary 2.21. This would require a probabilistic analysis of the additional second term discussed in Remark 2.8, which is no longer deterministic for general priors.

A more significant generalization of the Wigner model would allow a known non-Gaussian distribution for the entries of W . While LSS has also been analyzed here, it is not the optimal test: a better test involves first applying an entrywise transformation followed by LSS [22]. We expect this modified test to be optimal among computationally efficient tests for a broad class of spike priors, and this can potentially be approached using our method. This would require a precise analysis of the low-degree likelihood ratio in this more general setting.

Throughout, we have assumed the spike prior π is subgaussian, and it is plausible that this could be relaxed. Doing so may require a more complicated argument such as a conditional second moment computation. Certainly it is necessary to make some assumption on the growth of π 's moments or else LSS may cease to be the optimal test: if π is so heavy-tailed that the maximum of n independent copies exceeds $\tilde{O}(n^{1/4})$ then the maximum entry of the observed matrix Y will achieve strong detection for all $\lambda > 0$. We also note that our assumption that π have mean 0 is needed, or else LSS will be sub-optimal compared to a simple test based on the sum of all entries of Y .

We also hope that our approach may become useful beyond the spiked Wigner setting. A key ingredient in our approach is a precise bound on the norm of the low-degree likelihood ratio. The main step in deriving this precise limit is to show boundedness of the norm, as this allows us to deduce the desired convergence in expectation from the corresponding convergence in distribution. There are many settings where this boundedness has already been established, including spiked Wishart matrices [12] and the stochastic block model [38, 10], just to name a few. However, the precise positive results for testing do not yet seem to be known in these settings.

2.5. Proof Overview. The bulk of our technical work goes into proving Proposition 2.19. A key fact underlying the proof is the following: given an achievable ROC curve ϕ , one can construct a function f with $\mathbb{E}_P[f]/\sqrt{\mathbb{E}_Q[f^2]} = \text{val}(\phi)$. We will give an overview of how this is done and where the formula for $\text{val}(\cdot)$ comes from.

For intuition it will help to consider a simplified setting: imagine our ROC curve ϕ is achieved by thresholding a particular test statistic that only takes r different values. There are $r+1$ choices for the threshold, yielding $r+1$ different tests. Say test i achieves $(\alpha, \beta) = (a_i, b_i)$ for $i = 0, 1, \dots, r$, where $0 = a_0 < a_1 < \dots < a_r = 1$ and $0 = b_0 < b_1 < \dots < b_r = 1$. Also assume that no (a_i, b_i) is dominated by the convex hull of other (a_j, b_j) pairs, or else test i is redundant and can be removed. This means ϕ is the concave piecewise linear function formed by connecting each (a_i, b_i) to (a_{i+1}, b_{i+1}) with a line segment. (These points are achieved by taking appropriate probabilistic combinations between tests i and $i+1$.) Since we're assuming all our tests are thresholding the same underlying statistic, there will be, for any input Y , a unique $i^* = i^*(Y) \in \{0, 1, \dots, r-1\}$ such that all tests $i \leq i^*$ output \mathbf{q} and all tests $i > i^*$ output \mathbf{p} . By definition of a_i, b_i , the probability that $i^* = i$ is $a_{i+1} - a_i$ under \mathcal{Q} and $b_{i+1} - b_i$ under \mathcal{P} . Letting Q^* and P^* denote the distributions of i^* under \mathcal{Q} and \mathcal{P} respectively, it is a standard fact that the function f maximizing $\mathbb{E}_{\mathcal{P}}[f]/\sqrt{\mathbb{E}_{\mathcal{Q}}[f^2]}$ (given access to i^*) is the likelihood ratio (see e.g., [46, Proposition 1.9])

$$f(Y) = \frac{dP^*}{dQ^*}(i^*) = \frac{b_{i^*+1} - b_{i^*}}{a_{i^*+1} - a_{i^*}},$$

which is the slope of the line connecting (a_{i^*}, b_{i^*}) and (a_{i^*+1}, b_{i^*+1}) . For this choice of f , both $\mathbb{E}_{\mathcal{P}}[f]$ and $\mathbb{E}_{\mathcal{Q}}[f^2]$ are equal to

$$\sum_{i=0}^{r-1} \frac{(b_{i+1} - b_i)^2}{a_{i+1} - a_i} = \sum_{i=0}^{r-1} \left(\frac{b_{i+1} - b_i}{a_{i+1} - a_i} \right)^2 (a_{i+1} - a_i) = \text{val}(\phi)^2,$$

and so $\mathbb{E}_{\mathcal{P}}[f]/\sqrt{\mathbb{E}_{\mathcal{Q}}[f^2]} = \text{val}(\phi)$ as desired.

The full proof of Proposition 2.19 is more involved for a number of reasons. We do not assume the ROC curve is piecewise linear, but the proof will involve an appropriate discretization. We also do not assume the ROC curve is achieved by thresholding a single statistic at varying thresholds: there may be no relation between the tests that achieve two different (α, β) pairs. Finally, we must show that given a hypothetical achievable point (α^*, β^*) *above* the graph of ϕ , we can achieve an improved ROC curve whose value is *strictly* larger than $\text{val}(\phi)$ (see Figure 1 for illustration). We note that the value $\text{val}(\phi)$ is only meaningful for concave functions ϕ .

3. Proofs.

3.1. *Proof of Theorem 2.13.* A standard formula for $\|(L_\lambda)^{\leq D}\|$ gives (see Theorem 2.6 of [46])

$$(20) \quad \|(L_\lambda)^{\leq D}\|^2 = \mathbb{E}_{x, x'} \exp^{\leq D} \left(\frac{\lambda^2 n \langle x, x' \rangle^2}{2\|x\|^2\|x'\|^2} \right)$$

where x' is an independent copy of x (with entries i.i.d. from π), and $\exp^{\leq D}(z) := \sum_{d=0}^D z^d/d!$ (the degree- D Taylor polynomial for $\exp(z)$). Define the random variable $A = A_n \geq 0$ by

$$(21) \quad A := \frac{\lambda^2 n \langle x, x' \rangle^2}{2\|x\|^2\|x'\|^2},$$

so that

$$(22) \quad \|(L_\lambda)^{\leq D}\|^2 = \mathbb{E} \exp^{\leq D}(A) = \mathbb{E}[\mathbb{1}_{A \leq t} \cdot \exp^{\leq D}(A)] + \mathbb{E}[\mathbb{1}_{A > t} \cdot \exp^{\leq D}(A)]$$

where $t = t_n \geq 0$ is a threshold to be chosen later. The first term on the right-hand side is bounded by

$$(23) \quad \mathbb{E}[\mathbb{1}_{A \leq t} \cdot \exp(A)] - \sum_{d=D+1}^{\infty} \frac{t^d}{d!} \leq \mathbb{E}[\mathbb{1}_{A \leq t} \cdot \exp^{\leq D}(A)] \leq \mathbb{E}[\mathbb{1}_{A \leq t} \cdot \exp(A)],$$

where the second inequality relies on $A \geq 0$ so that every term in the Taylor series for \exp is non-negative. We will proceed to bound the various terms in (22) and (23) separately.

LEMMA 3.1. *If $t = o(D)$ then*

$$\sum_{d=D+1}^{\infty} \frac{t^d}{d!} = o(1).$$

PROOF. Using the standard bound $d! \geq (d/e)^d$,

$$\sum_{d=D+1}^{\infty} \frac{t^d}{d!} \leq \sum_{d=D+1}^{\infty} \left(\frac{et}{d}\right)^d \leq \sum_{d=D+1}^{\infty} \left(\frac{et}{D}\right)^d \leq \sum_{d=1}^{\infty} \left(\frac{et}{D}\right)^d = o(1),$$

since $t = o(D)$. \square

LEMMA 3.2. *If $t = \omega(1)$ and $t \leq \rho n$ for a particular constant $\rho = \rho(\lambda, \pi) > 0$, then*

$$\lim_{n \rightarrow \infty} \mathbb{E}[\mathbb{1}_{A \leq t} \cdot \exp(A)] = (1 - \lambda^2)^{-1/2}.$$

PROOF. By the central limit theorem, we have convergence in distribution $\langle x, x' \rangle / \sqrt{n} \xrightarrow{d} \mathcal{N}(0, 1)$. Also, $\|x\|^2/n \xrightarrow{d} 1$ and the same for x' . As a result, $A \xrightarrow{d} \lambda^2 \chi_1^2 / 2$ where χ_1^2 is a chi-squared random variable with 1 degree of freedom. Since $t = \omega(1)$, we have $\mathbb{1}_{A \leq t} \cdot \exp(A) \xrightarrow{d} \exp(\lambda^2 \chi_1^2 / 2)$. We aim to conclude the corresponding convergence of the expectation $\mathbb{E}[\mathbb{1}_{A \leq t} \cdot \exp(A)] \rightarrow \mathbb{E}[\exp(\lambda^2 \chi_1^2 / 2)]$. This completes the proof because, using the chi-squared moment-generating function, $\mathbb{E}[\exp(\lambda^2 \chi_1^2 / 2)] = (1 - \lambda^2)^{-1/2}$.

Convergence in distribution plus uniform integrability implies convergence of the expectation, so it remains to verify that the sequence $X_n := \mathbb{1}_{A_n \leq t_n} \cdot \exp(A_n)$ is uniformly integrable, i.e., for every $\epsilon > 0$ there exists $K \geq 0$ such that for all n we have $\mathbb{E}[\mathbb{1}_{|X_n| \geq K} \cdot |X_n|] \leq \epsilon$. Lemma 3.3 below establishes for all n that $\mathbb{E}[|X_n|^{1+\gamma}] \leq C$ for some positive constants γ, C (depending on λ, π). This implies uniform integrability because

$$\mathbb{E}[\mathbb{1}_{|X_n| \geq K} \cdot |X_n|] = \mathbb{E}[\mathbb{1}_{|X_n| \geq K} \cdot |X_n|^{1+\gamma} |X_n|^{-\gamma}] \leq \mathbb{E}[\mathbb{1}_{|X_n| \geq K} \cdot |X_n|^{1+\gamma} K^{-\gamma}] \leq CK^{-\gamma},$$

which can be made smaller than ϵ by choosing K sufficiently large. \square

We note that the rest of the argument is similar to [46, Theorem 3.9], with the main difference being that we have normalized the spike (a convention we have adopted for consistency with [22]). We will defer many of the details to the Supplement and state only the main claims here. The following lemma, proved in the Supplement, was used to establish uniform integrability above.

LEMMA 3.3. *Fix $\lambda \in (0, 1)$ and choose $\gamma > 0$ small enough so that $(1 + \gamma)\lambda^2 < 1$. If $t \leq \rho n$ for a particular constant $\rho = \rho(\lambda, \pi, \gamma) > 0$ then*

$$\mathbb{E}[\mathbb{1}_{A \leq t} \cdot \exp((1 + \gamma)A)] = O(1).$$

We now return to the final term in (22).

LEMMA 3.4. *For any sequences $D = D_n$ and $t = t_n$ that scale as $D = o(n/\log n)$ and $t = \Omega(n)$,*

$$\mathbb{E}[\mathbb{1}_{A>t} \cdot \exp^{\leq D}(A)] = o(1).$$

The proof of Lemma 3.4 is deferred to the Supplement. The proof of Theorem 2.13 now follows by combining the ingredients above; the full details of this are also deferred to the Supplement.

3.2. *Proof of Proposition 2.19.* We will first argue that $\alpha^* \in (0, 1)$ without loss of generality. First, since $\phi(1) = 1$ and $\phi(\alpha^*) < \beta^* \leq 1$, we must have $\alpha^* \neq 1$. Second, if $\alpha^* = 0$, we will find a *different* asymptotically achievable point (α, β) with $\beta > \phi(\alpha)$ and $\alpha > 0$, to use in place of (α^*, β^*) . Since ϕ is continuous, this can be done via the following test: with probability p for a sufficiently small constant $p > 0$, output \mathbf{p} ; otherwise, apply the original decision rule that achieved (α^*, β^*) . In the sequel we therefore assume $\alpha^* > 0$.

Define $\psi : [0, 1] \rightarrow [0, 1]$ to be the upper concave envelope of ϕ and (α^*, β^*) , as illustrated in Figure 1. By this we mean the minimal concave curve such that $\psi(\alpha) \geq \phi(\alpha)$ for all $\alpha \in [0, 1]$ and $\psi(\alpha^*) = \beta^*$. Note that ψ has the following structure for certain constants $0 < A_1 < \alpha^* < A_2 < 1$.

- For $\alpha \in [0, A_1] \cup [A_2, 1]$, $\psi(\alpha) = \phi(\alpha)$.
- For $\alpha \in [A_1, \alpha^*]$, the graph of ψ is a straight line connecting $(A_1, \phi(A_1))$ and (α^*, β^*) , and this line has slope $\phi'(A_1)$.
- For $\alpha \in [\alpha^*, A_2]$, the graph of ψ is a straight line connecting (α^*, β^*) and $(A_2, \phi(A_2))$, and this line has slope $\phi'(A_2)$.
- Finally, $\phi'(A_1) > \phi'(A_2)$, since ϕ is concave and (α^*, β^*) lies strictly above the graph of ϕ .

Here we have used some basic properties of ϕ from Assumption 2.17, namely the fact that ϕ' is continuous and decreasing, as well as the limits of ϕ' near 0 and 1.

LEMMA 3.5. *$\text{val}(\psi)$ exists and $\text{val}(\psi) > \text{val}(\phi)$.*

PROOF. The existence of $\text{val}(\psi)$ is immediate from the existence of $\text{val}(\phi)$ and the piecewise structure of ψ discussed above. Since ϕ and ψ coincide outside the interval $[A_1, A_2]$, it suffices to show

$$(24) \quad \int_{A_1}^{A_2} (\psi'(\alpha))^2 d\alpha > \int_{A_1}^{A_2} (\phi'(\alpha))^2 d\alpha.$$

Introduce the shorthand $\Delta_1 := \alpha^* - A_1$ and $\Delta_2 := A_2 - \alpha^*$. Since ψ is piecewise linear on $[A_1, A_2]$, the left-hand side can be computed directly:

$$(25) \quad \int_{A_1}^{A_2} (\psi'(\alpha))^2 d\alpha = (\phi'(A_1))^2 \Delta_1 + (\phi'(A_2))^2 \Delta_2.$$

Since $\phi(A_1) = \psi(A_1)$, $\phi(\alpha^*) < \psi(\alpha^*)$, and $\phi(A_2) = \psi(A_2)$, we have for some $\eta > 0$,

$$(26) \quad \int_{A_1}^{\alpha^*} \phi'(\alpha) d\alpha = \int_{A_1}^{\alpha^*} \psi'(\alpha) d\alpha - \eta = \phi'(A_1) \Delta_1 - \eta$$

and

$$(27) \quad \int_{\alpha^*}^{A_2} \phi'(\alpha) d\alpha = \int_{\alpha^*}^{A_2} \psi'(\alpha) d\alpha + \eta = \phi'(A_2) \Delta_2 + \eta.$$

Also note that an integrable function f on $[a, b]$ with $0 \leq m \leq f(x) \leq M$ for all $x \in [a, b]$ must satisfy

$$\begin{aligned} \int_a^b f(x)^2 dx &= \int_a^b (f(x) - m)(f(x) + m) dx + m^2(b - a) \\ &\leq (M + m) \int_a^b (f(x) - m) dx + m^2(b - a) \\ (28) \quad &= (M + m) \int_a^b f(x) dx - Mm(b - a). \end{aligned}$$

Recalling that ϕ' is positive and decreasing, combine (26), (27), (28) to bound the right-hand side of (24):

$$\begin{aligned} \int_{A_1}^{A_2} (\phi'(\alpha))^2 d\alpha &= \int_{A_1}^{\alpha^*} (\phi'(\alpha))^2 d\alpha + \int_{\alpha^*}^{A_2} (\phi'(\alpha))^2 d\alpha \\ &\leq [\phi'(A_1) + \phi'(\alpha^*)][\phi'(A_1)\Delta_1 - \eta] - \phi'(A_1)\phi'(\alpha^*)\Delta_1 \\ &\quad + [\phi'(\alpha^*) + \phi'(A_2)][\phi'(A_2)\Delta_2 + \eta] - \phi'(\alpha^*)\phi'(A_2)\Delta_2 \\ &= (\phi'(A_1))^2\Delta_1 + (\phi'(A_2))^2\Delta_2 - \eta[\phi'(A_1) - \phi'(A_2)]. \end{aligned}$$

Compare this with (25) and recall $\phi'(A_1) > \phi'(A_2)$ to complete the proof. □

Our next step will be to approximate ψ using a finite collection of points $u = (u_0, \dots, u_r)$ where $u_i = (a_i, \psi(a_i))$ for some choice of $0 = a_0 < a_1 < \dots < a_r = 1$. By virtue of lying on the graph of ψ , these points will be in “concave position,” which we define as follows.

DEFINITION 3.6 (Concave position and conc). *Let $u = (u_0, \dots, u_r)$ be a sequence of points in $[0, 1]^2$, sorted by strictly ascending first coordinate, with $u_0 = (0, 0)$ and $u_r = (1, 1)$. We say u is in concave position if the slope of the line through u_i and u_{i+1} is strictly positive and strictly decreasing as a function of i . In this case, we let $\text{conc}(u)$ denote the upper convex envelope of these points, i.e., the piecewise linear function $[0, 1] \rightarrow [0, 1]$ whose graph on $[a_i, a_{i+1}]$ is a straight line connecting u_i and u_{i+1} .*

LEMMA 3.7. *There exist values $r \geq 1$ and $0 = a_0 < a_1 < \dots < a_r = 1$ such that the points $u = (u_0, \dots, u_r)$ where $u_i = (a_i, \psi(a_i))$ are in concave position with*

$$\text{val}(\text{conc}(u)) > \text{val}(\phi).$$

PROOF. In light of Lemma 3.5 we can write $\text{val}(\psi)^2 = \text{val}(\phi)^2 + \epsilon$ for some $\epsilon > 0$. Since $\text{val}(\psi)$ exists, it is possible to choose a_1, a_{r-1} with $0 < a_1 < A_1$ and $A_2 < a_{r-1} < 1$ so that $\int_0^{a_1} (\psi'(\alpha))^2 d\alpha \leq \epsilon/6$ and $\int_{a_{r-1}}^1 (\psi'(\alpha))^2 d\alpha \leq \epsilon/6$. We will also include A_1, α^*, A_2 in the list of a_i ’s. It remains to partition the intervals $[a_1, A_1]$ and $[A_2, a_{r-1}]$. In each case, choose a fine enough partition so that each sub-interval $[a_i, a_{i+1}]$ satisfies $\psi'(a_i) \leq (1 + \gamma)\psi'(a_{i+1})$ for a constant $\gamma > 0$ to be chosen later. This is possible because ϕ' is continuous and strictly

positive, and $\psi' = \phi'$ on $[a_1, A_1] \cup [A_2, a_{r-1}]$. Let I denote the set of i that index the sub-intervals $[a_i, a_{i+1}]$ of $[a_1, A_1] \cup [A_2, a_{r-1}]$. For each $i \in I$, the mean value theorem implies that the slope m_i of the line connecting u_i and u_{i+1} satisfies

$$\psi'(a_{i+1}) \leq m_i \leq \psi'(a_i) \leq (1 + \gamma)\psi'(a_{i+1}).$$

Note that the interval $[A_1, A_2]$ has the same contribution to both $\text{val}(\text{conc}(u))$ and $\text{val}(\psi)$, and so

$$\begin{aligned} \text{val}(\psi)^2 - \text{val}(\text{conc}(u))^2 &\leq \sum_{i \in I} \left[\int_{a_i}^{a_{i+1}} (\psi'(\alpha))^2 d\alpha - m_i^2(a_{i+1} - a_i) \right] \\ &\quad + \int_{[0, a_1] \cup [a_{r-1}, 1]} (\psi'(\alpha))^2 d\alpha \\ &\leq \sum_{i \in I} [(\psi'(a_i))^2(a_{i+1} - a_i) - m_i^2(a_{i+1} - a_i)] + \epsilon/3 \\ &\leq \sum_{i \in I} [(1 + \gamma)^2 m_i^2(a_{i+1} - a_i) - m_i^2(a_{i+1} - a_i)] + \epsilon/3 \\ &= (2\gamma + \gamma^2) \sum_{i \in I} m_i^2(a_{i+1} - a_i) + \epsilon/3 \leq (2\gamma + \gamma^2)(\psi'(a_1))^2 + \epsilon/3, \end{aligned}$$

implying $\text{val}(\psi)^2 - \text{val}(\text{conc}(u))^2 \leq 2\epsilon/3$ for a sufficiently small choice of $\gamma > 0$. Recalling $\text{val}(\psi)^2 = \text{val}(\phi)^2 + \epsilon$, this completes the proof that $\text{val}(\text{conc}(u)) > \text{val}(\phi)$. Finally, remove any redundant u_i 's for which $m_{i-1} = m_i$, noting that this has no effect on $\text{val}(\text{conc}(u))$. The remaining points are in concave position because they lie on the graph of ψ , which is concave and strictly increasing. \square

Next we modify the points u slightly to create a new sequence of points $v = (v_0, \dots, v_r)$ where $v_i = (a_i, b_i)$. The α -coordinates are inherited from the u 's and the β -coordinates will be decreased slightly to ensure these points are asymptotically achievable. This step is only necessarily because we have only assumed asymptotic achievability for points strictly below the graph of ϕ , not on it. The proof is deferred to the Supplement.

LEMMA 3.8. *Let r and $\{a_i\}$ be as in Lemma 3.7. There exist values $0 = b_0 < b_1 < \dots < b_r = 1$ such that the points $v = (v_0, \dots, v_r)$ where $v_i = (a_i, b_i)$ are in concave position with*

$$\text{val}(\text{conc}(v)) > \text{val}(\phi),$$

and furthermore, $b_i < \psi(a_i)$ for each $i \notin \{0, r\}$.

We now have a finite number of points v_0, \dots, v_r such that each is asymptotically achievable in time \mathcal{T} . These points depend only on ϕ, α^*, β^* but not n . To complete the proof, we construct an algorithm to compute a function f achieving (17). The points v_0, \dots, v_r , as well as the algorithms for asymptotically achieving each of them, will be hard-coded into our algorithm for f .

PROOF OF PROPOSITION 2.19. Given input Y , we describe the algorithm to compute $f(Y)$. Let t_0, \dots, t_r denote the tests that asymptotically achieve the points v_0, \dots, v_r from Lemma 3.8, and let $s = (s_0, \dots, s_r)$ denote the sequence of outputs $s_i = t_i(Y)$. Since $v_0 = (0, 0)$ and $v_r = (1, 1)$, we always have $s_0 = \mathbf{q}$ and $s_r = \mathbf{p}$. Our goal is to choose $f(Y)$ so that

$$(29) \quad \frac{\mathbb{E}_{Y \sim \mathcal{P}}[f(Y)]}{\sqrt{\mathbb{E}_{Y \sim \mathcal{Q}}[f(Y)^2]}} \geq \text{val}(\text{conc}(v)) - o(1),$$

which implies (17) since $\text{val}(\text{conc}(v)) > \text{val}(\phi)$ by Lemma 3.8.

For intuition, it is natural to expect the rejection regions of the t_i 's to be nested in the sense that $t_{i+1}(Y) = \mathfrak{p}$ whenever $t_i(Y) = \mathfrak{p}$. If this is the case then the sequence s is always equal to one of the “monotone” sequences

$$S^{(0)} = (\mathfrak{q}, \mathfrak{p}, \mathfrak{p}, \dots, \mathfrak{p}), \quad S^{(1)} = (\mathfrak{q}, \mathfrak{q}, \mathfrak{p}, \dots, \mathfrak{p}), \quad \dots, \quad S^{(r-1)} = (\mathfrak{q}, \mathfrak{q}, \dots, \mathfrak{q}, \mathfrak{p}).$$

If $s = S^{(j)}$ for some j , let $f(Y) = \ell_j$, which, recall, is the slope of the line connecting v_i and v_{i+1} . In the case where the rejection regions of the t_i 's are nested, this rule achieves (29). However, there is no guarantee that the rejection regions will be nested, and so we need to also decide on $f(Y)$ for non-monotone s .

The general construction for $f(Y)$ is as follows. For $i = 0, \dots, r-1$ define

$$\sigma_i(s) = \begin{cases} +1 & \text{if } s_i = \mathfrak{q}, s_{i+1} = \mathfrak{p}, \\ -1 & \text{if } s_i = \mathfrak{p}, s_{i+1} = \mathfrak{q}, \\ 0 & \text{otherwise.} \end{cases}$$

Then choose any $f(Y) \geq 0$ satisfying

$$\sum_{i=0}^{r-1} \sigma_i(s) \ell_i \leq f(Y) \leq \sqrt{\sum_{i=0}^{r-1} \sigma_i(s) \ell_i^2},$$

for example, $f(Y)$ can be defined to simply equal the left-hand side above. The interval for possible $f(Y)$ values is nonempty due to Lemma 3.9 below, using the following facts: ℓ_i is positive and decreasing in i , by definition of concave position; and the nonzero σ_i 's alternate in sign, starting with positive. Since there are only a finite number of possible values for s , a lookup table with the corresponding $f(Y)$ values can be hard-coded into the algorithm. Note that if s is a monotone sequence $S^{(j)}$, we recover the rule $f(Y) = \ell_j$ from above.

Since $f(Y)$ depends only on s , we write $f(s) = f(Y)$. Now compute

$$\begin{aligned} \mathbb{E}_{\mathcal{P}}[f] &= \sum_s \mathcal{P}(s) f(s) \geq \sum_s \mathcal{P}(s) \sum_{i=0}^{r-1} \sigma_i(s) \ell_i = \sum_{i=0}^{r-1} \ell_i \sum_s \sigma_i(s) \mathcal{P}(s) \\ &= \sum_{i=0}^{r-1} \ell_i [\mathcal{P}(s_{i+1} = \mathfrak{p}) - \mathcal{P}(s_i = \mathfrak{p})] \\ &= \sum_{i=1}^r \mathcal{P}(s_i = \mathfrak{p}) (\ell_{i-1} - \ell_i) \quad \text{where } \ell_r := 0 \text{ and using } \mathcal{P}(s_0 = \mathfrak{p}) = 0 \\ &\geq \sum_{i=1}^r b_i (\ell_{i-1} - \ell_i) - o(1) = \sum_{i=0}^{r-1} \ell_i (b_{i+1} - b_i) - o(1) = \text{val}(\text{conc}(v))^2 - o(1). \end{aligned}$$

Similarly,

$$\begin{aligned} \mathbb{E}_{\mathcal{Q}}[f^2] &= \sum_s \mathcal{Q}(s) f(s)^2 \leq \sum_s \mathcal{Q}(s) \sum_{i=0}^{r-1} \sigma_i(s) \ell_i^2 = \sum_{i=0}^{r-1} \ell_i^2 \sum_s \sigma_i(s) \mathcal{Q}(s) \\ &= \sum_{i=0}^{r-1} \ell_i^2 [\mathcal{Q}(s_{i+1} = \mathfrak{p}) - \mathcal{Q}(s_i = \mathfrak{p})] = \sum_{i=1}^r \mathcal{Q}(s_i = \mathfrak{p}) (\ell_{i-1}^2 - \ell_i^2) \\ &\leq \sum_{i=1}^r a_i (\ell_{i-1}^2 - \ell_i^2) - o(1) = \sum_{i=0}^{r-1} \ell_i^2 (a_{i+1} - a_i) - o(1) = \text{val}(\text{conc}(v))^2 - o(1). \end{aligned}$$

Together, the above calculations imply (29) as desired. \square

Finally, the following fact was used above. Its proof is deferred to the Supplement.

LEMMA 3.9. *For an integer $m \geq 1$, suppose $h_0 \geq h_1 \geq \dots \geq h_{m-1} \geq 0$. Then*

$$0 \leq \sum_{i=0}^{m-1} (-1)^i h_i \leq \sqrt{\sum_{i=0}^{m-1} (-1)^i h_i^2}.$$

Acknowledgments. The authors thank Ahmed El Alaoui for helpful discussions.

Funding. The first author was partially supported by a Microsoft Trustworthy AI Grant, an ONR grant and a David and Lucile Packard Fellowship. The second author was partially supported by an Alfred P. Sloan Research Fellowship and NSF CAREER Award CCF-2338091.

SUPPLEMENTARY MATERIAL

Supplement to “Precise Error Rates for Computationally Efficient Testing”
Contains some additional proofs that were omitted from the main text.

REFERENCES

- [1] AIZENMAN, M., LEBOWITZ, J. L. and RUELLE, D. (1987). Some rigorous results on the Sherrington–Kirkpatrick spin glass model. *Communications in mathematical physics* **112** 3–20.
- [2] ANDERSON, G. W. and ZEITOUNI, O. (2006). A CLT for a band matrix model. *Probability Theory and Related Fields* **134** 283–338.
- [3] BAI, Z. and SILVERSTEIN, J. W. (2004). CLT for linear spectral statistics of large-dimensional sample covariance matrices. *The Annals of Probability* **32** 553–605.
- [4] BAI, Z., WANG, X. and ZHOU, W. (2009). CLT for linear spectral statistics of Wigner matrices. *Electronic Journal of Probability* **14** 2391–2417.
- [5] BAI, Z. and YAO, J.-F. (2008). Central limit theorems for eigenvalues in a spiked population model. In *Annales de l’IHP Probabilités et statistiques* **44** 447–474.
- [6] BAIK, J., BEN AROUS, G. and PÉCHÉ, S. (2005). Phase Transition of the Largest Eigenvalue for Nonnull Complex Sample Covariance Matrices. *Annals of Probability* 1643–1697.
- [7] BAIK, J. and LEE, J. O. (2016). Fluctuations of the free energy of the spherical Sherrington–Kirkpatrick model. *Journal of Statistical Physics* **165** 185–224.
- [8] BAIK, J. and LEE, J. O. (2017). Fluctuations of the free energy of the spherical Sherrington–Kirkpatrick model with ferromagnetic interaction. In *Annales Henri Poincaré* **18** 1867–1917. Springer.
- [9] BAIK, J. and SILVERSTEIN, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *Journal of multivariate analysis* **97** 1382–1408.
- [10] BANDEIRA, A. S., BANKS, J., KUNISKY, D., MOORE, C. and WEIN, A. S. (2021). Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. In *Conference on Learning Theory* 410–473. PMLR.
- [11] BANDEIRA, A. S., EL ALAOUI, A., HOPKINS, S., SCHRAMM, T., WEIN, A. S. and ZADIK, I. (2022). The Franz-Parisi criterion and computational trade-offs in high dimensional statistics. *Advances in Neural Information Processing Systems* **35** 33831–33844.
- [12] BANDEIRA, A. S., KUNISKY, D. and WEIN, A. S. (2020). Computational Hardness of Certifying Bounds on Constrained PCA Problems. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [13] BANERJEE, D. and MA, Z. (2017). Optimal hypothesis testing for stochastic block models with growing degrees. *arXiv preprint arXiv:1705.05305*.
- [14] BANKS, J., MOORE, C., VERSHYNIN, R., VERZELEN, N. and XU, J. (2018). Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *IEEE Transactions on Information Theory* **64** 4872–4894.

- [15] BARAK, B., HOPKINS, S., KELNER, J., KOTHARI, P. K., MOITRA, A. and POTECHIN, A. (2019). A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing* **48** 687–735.
- [16] BAYATI, M. and MONTANARI, A. (2011). The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory* **57** 764–785.
- [17] BENAYCH-GEORGES, F. and NADAKUDITI, R. R. (2011). The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics* **227** 494–521.
- [18] BERTHET, Q. and RIGOLLET, P. (2013). Complexity theoretic lower bounds for sparse principal component detection. In *Conference on learning theory* 1046–1066. PMLR.
- [19] BLUM, L., SHUB, M. and SMALE, S. (1989). On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin of the American Mathematical Society* **21** 1–46.
- [20] BRENNAN, M. and BRESLER, G. (2020). Reducibility and statistical-computational gaps from secret leakage. In *Conference on Learning Theory* 648–847. PMLR.
- [21] CAPITAIN, M., DONATI-MARTIN, C. and FÉRAL, D. (2009). The largest eigenvalues of finite rank deformation of large Wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability* **37** 1–47.
- [22] CHUNG, H. W. and LEE, J. O. (2022). Weak detection in the spiked Wigner model. *IEEE Transactions on Information Theory* **68** 7427–7453.
- [23] DECELLE, A., KRZAKALA, F., MOORE, C. and ZDEBOROVÁ, L. (2011). Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E* **84** 066106.
- [24] DIAKONIKOLAS, I. and KANE, D. (2022). Non-gaussian component analysis via lattice basis reduction. In *Conference on Learning Theory* 4535–4547. PMLR.
- [25] DING, X. and WANG, Z. (2023). Global and local CLTs for linear spectral statistics of general sample covariance matrices when the dimension is much larger than the sample size with applications. *arXiv preprint arXiv:2308.08646*.
- [26] DING, Y., KUNISKY, D., WEIN, A. S. and BANDEIRA, A. S. (2023). Subexponential-time algorithms for sparse PCA. *Foundations of Computational Mathematics* 1–50.
- [27] DONOHO, D. L., MALEKI, A. and MONTANARI, A. (2009). Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences* **106** 18914–18919.
- [28] EL ALAOUI, A., KRZAKALA, F. and JORDAN, M. I. (2020). Fundamental limits of detection in the spiked Wigner model. *The Annals of Statistics* **48** 863–885.
- [29] FELDMAN, V., GRIGORESCU, E., REYZIN, L., VEMPALA, S. S. and XIAO, Y. (2017). Statistical algorithms and a lower bound for detecting planted cliques. *Journal of the ACM* **64** 1–37.
- [30] FÉRAL, D. and PÉCHÉ, S. (2007). The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in mathematical physics* **272** 185–228.
- [31] FLETCHER, A. K. and RANGAN, S. (2018). Iterative reconstruction of rank-one matrices in noise. *Information and Inference: A Journal of the IMA* **7** 531–562.
- [32] GAMARNIK, D., JAGANNATH, A. and WEIN, A. S. (2020). Low-degree hardness of random optimization problems. In *61st Annual Symposium on Foundations of Computer Science (FOCS)* 131–140. IEEE.
- [33] GAMARNIK, D., MOORE, C. and ZDEBOROVÁ, L. (2022). Disordered systems insights on computational hardness. *Journal of Statistical Mechanics: Theory and Experiment* **2022** 114015.
- [34] GAMARNIK, D. and SUDAN, M. (2017). Limits of Local Algorithms over Sparse Random Graphs. *The Annals of Probability* 2353–2376.
- [35] HOLMGREN, J. and WEIN, A. S. (2021). Counterexamples to the Low-Degree Conjecture. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)* **185**.
- [36] HOPKINS, S. (2018). Statistical Inference and the Sum of Squares Method, PhD thesis, Cornell University.
- [37] HOPKINS, S. B., KOTHARI, P. K., POTECHIN, A., RAGHAVENDRA, P., SCHRAMM, T. and STEURER, D. (2017). The power of sum-of-squares for detecting hidden structures. In *58th Annual Symposium on Foundations of Computer Science (FOCS)* 720–731. IEEE.
- [38] HOPKINS, S. B. and STEURER, D. (2017). Efficient Bayesian estimation from few samples: community detection and related problems. In *58th Annual Symposium on Foundations of Computer Science (FOCS)* 379–390. IEEE.
- [39] JERRUM, M. (1992). Large cliques elude the Metropolis process. *Random Structures & Algorithms* **3** 347–359.
- [40] JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *The Annals of statistics* **29** 295–327.
- [41] JUNG, J. H., CHUNG, H. W. and LEE, J. O. (2020). Weak detection in the spiked Wigner model with general rank. *arXiv preprint arXiv:2001.05676*.

- [42] JUNG, J. H., CHUNG, H. W. and LEE, J. O. (2021). Detection of signal in the spiked rectangular models. In *International Conference on Machine Learning* 5158–5167. PMLR.
- [43] KOEHLER, F. and MOSSEL, E. (2021). Reconstruction on Trees and Low-Degree Polynomials. *arXiv preprint arXiv:2109.06915*.
- [44] KOTHARI, P., VEMPALA, S. S., WEIN, A. S. and XU, J. (2023). Is Planted Coloring Easier than Planted Clique? In *The Thirty Sixth Annual Conference on Learning Theory* 5343–5372. PMLR.
- [45] KUNISKY, D. (2021). Hypothesis testing with low-degree polynomials in the Morris class of exponential families. In *Conference on Learning Theory* 2822–2848. PMLR.
- [46] KUNISKY, D., WEIN, A. S. and BANDEIRA, A. S. (2019). Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio. In *ISAAC Congress (International Society for Analysis, its Applications and Computation)* 1–50. Springer.
- [47] LEHMANN, E. and ROMANO, J. P. (2005). Testing statistical hypotheses. *Springer texts in statistics*.
- [48] LESIEUR, T., KRZAKALA, F. and ZDEBOROVÁ, L. (2015). Phase transitions in sparse PCA. In *International Symposium on Information Theory (ISIT)* 1635–1639. IEEE.
- [49] LIU, Z., HU, J., BAI, Z. and SONG, H. (2022). A CLT for the LSS of large dimensional sample covariance matrices with diverging spikes. *arXiv preprint arXiv:2212.05896*.
- [50] LÖFFLER, M., WEIN, A. S. and BANDEIRA, A. S. (2022). Computationally efficient sparse clustering. *Information and Inference: A Journal of the IMA* **11** 1255–1286.
- [51] MAÏDA, M. (2007). Large deviations for the largest eigenvalue of rank one deformations of Gaussian ensembles. *Electronic Journal of Probability* **12** 1131–1150.
- [52] MIOLANE, L. (2018). Phase transitions in spiked matrix estimation: information-theoretic analysis. *arXiv preprint arXiv:1806.04343*.
- [53] MONTANARI, A., REICHMAN, D. and ZEITOUNI, O. (2015). On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. *Advances in Neural Information Processing Systems* **28**.
- [54] MONTANARI, A. and VENKATARAMAN, R. (2021). Estimation of Low-Rank Matrices via Approximate Message Passing. *The Annals of Statistics* **49** 321–345.
- [55] MONTANARI, A. and WEIN, A. S. (2022). Equivalence of Approximate Message Passing and Low-Degree Polynomials in Rank-One Matrix Estimation. *arXiv preprint arXiv:2212.06996*.
- [56] NADLER, B. (2008). Finite sample approximation results for principal component analysis: a matrix perturbation approach. *The Annals of Statistics* **36** 2791–2817.
- [57] ONATSKI, A., MOREIRA, M. J. and HALLIN, M. (2013). Asymptotic power of sphericity tests for high-dimensional data. *The Annals of Statistics* **41** 1204–1231.
- [58] ONATSKI, A., MOREIRA, M. J. and HALLIN, M. (2014). Signal detection in high dimension: The multi-spiked case. *The Annals of Statistics* 225–254.
- [59] PAUL, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica* 1617–1642.
- [60] PÉCHÉ, S. (2006). The largest eigenvalue of small rank perturbations of Hermitian random matrices. *Probability Theory and Related Fields* **134** 127–173.
- [61] PERRY, A., WEIN, A. S., BANDEIRA, A. S. and MOITRA, A. (2018). Optimality and sub-optimality of PCA I: Spiked random matrix models. *The Annals of Statistics* **46** 2416–2451.
- [62] PIZZO, A., RENFREW, D. and SOSHNIKOV, A. (2013). On finite rank deformations of Wigner matrices. In *Annales de l'IHP Probabilités et Statistiques* **49** 64–94.
- [63] RAGHAVENDRA, P., SCHRAMM, T. and STEURER, D. (2018). High dimensional estimation via sum-of-squares proofs. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018* 3389–3423. World Scientific.
- [64] RUSH, C., SKERMAN, F., WEIN, A. S. and YANG, D. (2023). Is It Easier to Count Communities Than Find Them? In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [65] SCHRAMM, T. and WEIN, A. S. (2022). Computational barriers to estimation from low-degree polynomials. *The Annals of Statistics* **50** 1833–1858.
- [66] WANG, Z. and YAO, J. (2021). Central limit theorem for linear spectral statistics of block-Wigner-type matrices. *arXiv preprint arXiv:2110.12171*.
- [67] ZADIK, I., SONG, M. J., WEIN, A. S. and BRUNA, J. (2022). Lattice-based methods surpass sum-of-squares in clustering. In *Conference on Learning Theory* 1247–1248. PMLR.