

Statistical limits of spiked tensor models

Amelia Perry*^{†1}, Alexander S. Wein*^{‡2}, and Afonso S. Bandeira^{§2,3}

¹Department of Mathematics, Massachusetts Institute of Technology

²Department of Mathematics, Courant Institute of Mathematical Sciences, New York University

³Center for Data Science, New York University

June 24, 2020

Abstract

We study the statistical limits of both detecting and estimating a rank-one deformation of a symmetric random Gaussian tensor. We establish upper and lower bounds on the critical signal-to-noise ratio, under a variety of priors for the planted vector: (i) a uniformly sampled unit vector, (ii) i.i.d. ± 1 entries, and (iii) a sparse vector where a constant fraction ρ of entries are i.i.d. ± 1 and the rest are zero. For each of these cases, our upper and lower bounds match up to a $1+o(1)$ factor as the order d of the tensor becomes large. For sparse signals (iii), our bounds are also asymptotically tight in the sparse limit $\rho \rightarrow 0$ for any fixed d (including the $d=2$ case of sparse PCA). Our upper bounds for (i) demonstrate a phenomenon reminiscent of the work of Baik, Ben Arous and Péché: an ‘eigenvalue’ of a perturbed tensor emerges from the bulk at a strictly lower signal-to-noise ratio than when the perturbation itself exceeds the bulk; we quantify the size of this effect. We also provide some general results for larger classes of priors. In particular, the large d asymptotics of the threshold location differs between problems with discrete priors versus continuous priors. Finally, for priors (i) and (ii) we carry out the replica prediction from statistical physics, which is conjectured to give the exact information-theoretic threshold for any fixed d .

Of independent interest, we introduce a new improvement to the second moment method for continuity, on which our lower bounds are based. Our technique conditions away from rare ‘bad’ events that depend on interactions between the signal and noise. This enables us to close $\sqrt{2}$ -factor gaps present in several previous works.

1 Introduction

Among the central problems in random matrix theory is to determine the spectral properties of *spiked* or *deformed* random matrix ensembles. Introduced by Johnstone [Joh01], such matrices consist of a random matrix (e.g. Wigner or Wishart) with a low-rank perturbation. These distributions serve as models for data consisting of “signal plus noise” and thus, results on these models form the basis of our theoretical understanding of principal components analysis (PCA) throughout the sciences.

*The first two authors contributed equally.

[†]This work was supported in part by NSF CAREER Award CCF-1453261 and a grant from the MIT NEC Corporation.

[‡]Email: aewin@cims.nyu.edu. Part of this work was done while A.S.W. was with the Department of Mathematics at the Massachusetts Institute of Technology. A.S.W. received Government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. A.S.W. is also supported by NSF grant DMS-1712730 and by the Simons Collaboration on Algorithms and Geometry.

[§]Email: bandeira@cims.nyu.edu. A.S.B. was supported by NSF Grants DMS-1712730 and DMS1719545, and by a grant from the Sloan foundation.

Perhaps the most studied phenomenon in these spiked ensembles is the transition first examined by [BBP05] in the Wishart setting. We will be interested in the Wigner analogue: if W is a Gaussian Wigner matrix¹ and x is a unit vector (the ‘spike’), the spectrum of the spiked matrix $\lambda xx^\top + W$ undergoes a sharp phase transition at $\lambda = 1$ (see e.g. [Péc06, FP07, CDMF09, BGN11]). Namely, when $\lambda \leq 1$, many properties of the spectrum resemble those of a random (‘unspiked’) matrix: the empirical distribution of eigenvalues is semicircular and the top eigenvalue concentrates about 2. When $\lambda > 1$, however, the spectrum becomes indicative of the spike: a single eigenvalue exceeds 2, exiting the semicircular bulk, and the associated eigenvector is correlated with the spike (and the precise correlation is known).

We emphasize that this ‘BBP-style’ ([BBP05]) transition exhibits a *push-out* effect: the top eigenvalue of a random Wigner matrix is 2 (in the high-dimensional limit), but one only needs to add a planted signal of spectral norm 1 before it becomes visible in the spectrum [FP07]. Once $\lambda > 1$, the planted signal aligns well enough with fluctuations of the random matrix in order to create an eigenvalue greater than 2.

More recent work shows a second, statistical role of this $\lambda = 1$ threshold: not only does the top eigenvalue fail to distinguish the spiked and unspiked models for $\lambda < 1$, but in fact every hypothesis test fails with constant probability [OMH13, MRZ15, PWBM16]. Thus, this transition indicates the point at which the spiked and unspiked models become markedly different.

It is natural to ask how the phenomena above generalize to tensors of higher order. Such tensors lack a well-behaved spectral theory, and many standard tools of random matrix theory (e.g. the method of resolvents) fail to cleanly generalize. However, there remain a number of interesting probabilistic questions to ask in this setting:

- Is there a sharp transition point for λ , below which the spiked model resembles the unspiked model?

In particular, we compare the Wigner tensor model (entries of W are Gaussian, i.i.d. apart from permutation symmetry; see Definition 1.1) to a spiked analogue $T = \lambda x^{\otimes d} + W$. Previous work of [MRZ15] provides a bound on λ below which these two models are information-theoretically indistinguishable. On the other hand, [RM14] notes that once λ exceeds the injective norm (defined below) of the noise, the spiked and unspiked models can be distinguished via injective norm. There remained a $\sqrt{2}$ -factor gap between these bounds as the order $d \rightarrow \infty$. In this paper, we improve both the lower and upper bounds, saving an asymptotic $\sqrt{2}$ factor in the lower bound, and thus obtaining a $(1 + o(1))$ factor gap as $d \rightarrow \infty$; see Figure 1.

- By analogy with the top eigenvalue (spectral norm) of a random matrix, what is the *injective norm*

$$\|T\| \triangleq \max_{\|x\|=1} \langle T, x^{\otimes d} \rangle$$

of a spiked Wigner d -tensor?

For unspiked tensors, the value of the injective norm was predicted by [CS92] through non-rigorous methods from statistical physics, and was later rigorously proven [Tal06a, ABC13, Sub15]. However, the spiked question has not been studied to our knowledge. Our statistical lower bound shows in particular that the injective norm of a spiked tensor remains identical to this value for all λ below a critical value. However, starting slightly above this critical value, our work provides a strong lower bound on the injective norm of the spiked model $T = \lambda x^{\otimes d} + W$, which exceeds the injective norm of the unspiked model.

- Do tensors exhibit a BBP-style push-out effect, and if so, how large is it?

We show that the injective norm of a spiked tensor exceeds that of an unspiked tensor strictly before the injective norm of the spike exceeds that of the noise, much as the $\lambda = 1$ threshold in the matrix case strictly precedes the spectral norm of 2 in the noise. We identify the asymptotic size of this gap as the order d becomes large, up to a small constant factor.

¹ W is symmetric with off-diagonal entries $\mathcal{N}(0, 1/n)$, diagonal entries $\mathcal{N}(0, 2/n)$, and all entries independent (except for symmetry).




Figure 1: Critical SNR values λ as a function of the tensor order d . Left: the spherical prior; from top to bottom, black: the injective norm μ_d of a noise tensor (see Theorem 3.3); red: our upper bound $\Lambda_{\text{sph},d}^*$ (Corollary 3.5); dashed: replica prediction for the exact threshold (Appendix B); green: our lower bound $\lambda_{\text{sph},d}^*$ (Theorem 3.2); blue: the lower bound of [MRZ15] which is loose by a factor of $\sqrt{2}$ for large d . Right: the Rademacher prior; from top to bottom, red: our upper bound (Theorem 1.4); dashed: replica prediction, which is a rigorous upper bound (see Appendix B); green: our lower bound $\lambda_{\text{Rad},d}^*$ (Theorem 4.6).

Much as random matrix theory provides a theoretical foundation for PCA, these questions probe at the statistical limits of *tensor PCA*, the estimation of a low-rank spike from a spiked tensor. Such problems arise naturally in topic modeling [AGH⁺14], in the study of hypergraphs [DBKP11], and more generally in moment-based approaches to estimation, when it may be desirable to detect low-rank structure in higher empirical moments. As many such problems involve extra structure such as sparsity in the signal, we allow the spike to be drawn from various priors, such as a distribution of sparse vectors. For each prior, we will investigate the detection problem of distinguishing the spiked and unspiked models, as well as the recovery problem of estimating the spike.

Our techniques and prior work. We will prove information-theoretic lower bounds using the second moment method associated with the statistical notion of *contiguity*. By computing a particular second moment, one can show that the spiked and unspiked models are contiguous, implying that they cannot be reliably distinguished. This second moment method originated in the study of random graphs (see e.g. [Wor99, RW94, Jan95, MRRW97]) but has since been applied to various average-case computational problems such as the stochastic block model [MNS15, BMNN16], submatrix localization [BMV⁺17], Gaussian mixture clustering [BMV⁺17], synchronization problems over compact groups [PWBM16], and even spiked tensor PCA [MRZ15] and sparse PCA [BMV⁺17]. However, many of the previous results are not tight in particular asymptotic regimes. In fact, there are many instances where curiously, in certain regimes they are loose by precisely a factor of $\sqrt{2}$ in the signal-to-noise ratio λ [MRZ15, BMNN16, BMVX16, PWBM16].

Our main technical contribution is a modification of the second moment method that closes (at least some of) these gaps. Specifically, we close the gap for spherically-spiked² tensor PCA in the limit of large tensor order [MRZ15], and the gap³ for sparse PCA in the limit of low sparsity [BMVX16]. Our technique, which we call *noise conditioning*, is based on conditioning away from rare bad events that depend jointly on the signal and noise. We expect that this technique can also be used to close several other $\sqrt{2}$ -factor gaps of the same nature. Another application in which we have found noise conditioning to be fruitful is in contiguity results for the Rademacher-spiked Wishart model [PWBM16]. For this problem the basic second moment method struggles quite badly because, due to a certain symmetry, it gives the same results for

²Here, the spike is a uniformly random unit vector.

³After the first appearance of this paper, an update [BMV⁺17] independently closed the asymptotic $\sqrt{2}$ -factor gap from [BMVX16], using a different modification of the second moment method.




Figure 2: Upper and lower bounds for the critical λ value in sparse Rademacher PCA ($d = 2$) as a function of sparsity ρ . Note that $\lambda = 1$ corresponds to the eigenvalue transition above which efficient algorithms (such as PCA) are known to succeed. From top to bottom, red: MLE upper bound (Theorem 1.5, first due to [BMV⁺17]); dashed: replica prediction of [LKZ15] which has been rigorously shown to be correct for the weak recovery problem [KXZ16, BDM⁺16, LM16]; green: our noise conditioning lower bound $\lambda_{\text{spRad}(\rho),d}^*$ (Theorem 4.11); blue: our lower bound without noise conditioning [PWBM16]; purple: lower bound of [BMVX16] (equations (33),(34) in [BMVX16]). In the limit $\rho \rightarrow 0$, noise conditioning (green) achieves the correct asymptotic behavior (matching the upper bound) whereas the blue and purple curves are loose by a factor of $\sqrt{2}$.

the positively- and negatively-spiked regimes, even though these two regimes have very different thresholds; the noise conditioning method is able to break this symmetry, yielding tight or almost-tight results in both regimes.

Our noise conditioning method is somewhat reminiscent of other modified second moment methods that have appeared in other contexts such as community detection in random networks [ACV14, VAC15], and thresholds for random satisfiability problems (e.g. k -colorability, k -sat) [COZ12, COP12, COV13, COP13, COP16].

Our upper bound for the spherical prior, via a lower bound on the spiked injective norm, is based on a direct analysis of how vectors close to the spike can align constructively with fluctuations in the noise to produce a larger injective norm than the spike alone. In particular, we consider how such vectors align with submatrices of the given tensor, and leverage existing results from spiked matrix models. Upper bounds for structured priors are obtained through naïve union bounds on the maximum likelihood value.

A variety of other techniques originating from statistical physics have also been successful in tackling structured inference or optimization problems involving large random systems, including random matrices and tensors. For instance, random tensors are intimately connected to spin glasses with d -spin interactions [Gar85, CS92]. The so-called *replica method* gives extremely precise non-rigorous solutions to these types of problems (see [MM09] for an introduction). In some cases, such as the celebrated Parisi formula for the

ground state of a Sherrington–Kirkpatrick spin glass⁴, the replica prediction has been rigorously proven to be correct [Tal06b]; we also rigorously know the injective norm of a random tensor of any order (in the high-dimensional limit), as well as various structural properties of the critical points of the associated maximization problem [Tal06a, ABČ13, Sub15]. Furthermore, for a variety of structured spiked matrix problems such as sparse PCA (with constant-fraction sparsity), the statistically-optimal mean squared error can be exactly characterized in the high-dimensional limit for any level of sparsity and any signal-to-noise ratio [KXZ16, BDM⁺16, LM16]. Additionally, [KM09] implies bounds on the statistical threshold for tensor PCA with a $\{\pm 1\}$ -valued spike (which we discuss in Appendix B). Recent independent work confirms the replica prediction for various spiked tensor problems [LML⁺17]. Some key techniques used in the above works include Guerra interpolation [Gue03], the Kac–Rice formula (see [ABČ13]), and the approximate message passing (AMP) framework [DMM09, DMM10, BM11, JM13] (see also [DM14, DAM16, MR16, BDM⁺16]).

In comparison to these techniques from statistical physics, the second moment method for contiguity typically does not yield results that are as sharp, but it has the advantage of being quite simple and widely applicable. In particular, it can be applied to problems such as the sparse stochastic block model [MNS15, BMNN16] (constant average degree) for which various techniques in statistical physics do not seem to apply (see e.g. [DAM16, LM16]). Another advantage of the second moment method is that it addresses the detection problem instead of only the recovery problem, and furthermore implies bounds on hypothesis testing power below the detection threshold [PWBM16]. (However, unlike e.g. AMP, the second moment method only tells us about the threshold for nontrivial recovery and not the optimal recovery error at each value of λ above the threshold.)

We remark that our results on statistical indistinguishability have concrete implications for various probabilistic quantities. In particular, for any quantity that converges in probability to a constant under the spiked and unspiked models (such as the injective norm), the limiting values must agree throughout the subcritical region of the spiked model (signal strength below the detection threshold)⁵. For example, in the unspiked model, [ABČ13, Sub15] give a detailed description of the energy landscape, i.e. the number of critical points of $x \mapsto \langle T, x^{\otimes d} \rangle$ of any given value and index; our results immediately imply that the same is true for subcritical spiked tensors. An interesting problem, that we largely do not address here, is to characterize the energy landscape above the detection threshold.

It is important to note that we are studying information-theoretic limits rather than computational ones. All of the upper bounds in this paper are inefficient algorithms such as exhaustive search over all possible spikes. There is good evidence in the form of sum-of-squares lower bounds that there is a significant gap between what is possible statistically and what is possible computationally for tensor PCA and related problems [HSS15, BHK⁺16]. The $d = 2$ case of matrices is of course an exception; in this case spectral algorithms achieve the optimal detection and recovery threshold in many cases. Various efficient algorithms for spiked tensor PCA are considered by [RM14, HSS15, HSSS15], but (as is believed to be necessary) these operate only in a regime that is quite far from the information-theoretic threshold.

1.1 Preliminaries

Definition 1.1. We define a *Wigner tensor* $W \in \mathbb{R}^{n^d}$ of order d by the following sampling procedure: an asymmetric precursor $W' \in \mathbb{R}^{n^d}$ is drawn with entries sampled i.i.d. from $\mathcal{N}(0, 2/n)$, and is then averaged over all permutations of the indices to form a symmetric tensor W . Thus, the distribution of W has density proportional to $\exp(-\frac{n}{4}\langle W, W \rangle)$ on the space of symmetric tensors. We denote this distribution by $\text{WT}(d)$.

Note that a typical entry of W (with no repeated indices) is distributed as $\mathcal{N}(0, 2/(nd!))$, and for any unit vector x we have $\langle x^{\otimes d}, W \rangle \sim \mathcal{N}(0, 2/n)$. This normalization agrees with that of [MRZ15], but other conventions exist in the literature.

⁴This can be thought of as the maximum value of $x^\top W x$ over $x \in \{\pm 1\}^n$ where W is a Gaussian Wigner matrix.

⁵In fact, for tensors of order $d \geq 3$ we show that the unspiked and subcritical-spiked distributions differ by $o(1)$ in total variation distance (in the high-dimensional limit), implying that any quantity with a limit in distribution must converge to the same distribution under both models.

Let the *prior* $\mathcal{X} = \{\mathcal{X}_n\}$ be a family of distributions (one for each n) over unit vectors $x \in \mathbb{R}^n$, $\|x\| = 1$. Define the *spiked* distribution $\text{WT}(d, \lambda, \mathcal{X})$ by $T = \lambda x^{\otimes d} + W$, where x is drawn from \mathcal{X} , $\lambda \geq 0$ is a signal-to-noise parameter, and W is drawn from $\text{WT}(d)$.

The injective norm of a d -tensor T is defined as $\|T\| = \max\langle x^{\otimes d}, T \rangle$ over unit vectors x . For each $d \geq 2$ it is known that the injective norm $\|\text{WT}(d)\|$ of a random tensor converges in probability to a particular value μ_d (see Theorem 3.3) [CS92, Tal06a, ABC13, RM14, Sub15]. For $d = 2$ we have $\mu_2 = 2$, the spectral norm of a Wigner matrix. We also have, for instance, $\mu_3 \approx 2.3433$, and $\mu_d = (1 + o(1))\sqrt{2 \log d}$ as $d \rightarrow \infty$.

In the matrix case ($d = 2$) we have the following classical eigenvalue transition for spiked Wigner matrices.

Theorem 1.2 ([Péc06, FP07, CDMF09, BGN11]). *Let $T = \lambda x x^\top + W$ where x is a unit vector in \mathbb{R}^n and W is an $n \times n$ Gaussian Wigner matrix (symmetric with off-diagonal entries $\mathcal{N}(0, 1/n)$, diagonal entries $\mathcal{N}(0, 2/n)$, and all entries independent up to symmetry).*

- *If $\lambda \leq 1$, the top eigenvalue of T converges almost surely to 2 as $n \rightarrow \infty$, and the top eigenvector v (unit-norm) has trivial correlation with the spike: $\langle v, x \rangle^2 \rightarrow 0$ almost surely.*
- *If $\lambda > 1$, the top eigenvalue converges almost surely to $\lambda + 1/\lambda > 2$ and v has nontrivial correlation with the spike: $\langle v, x \rangle^2 \rightarrow 1 - 1/\lambda^2$ almost surely.*

All of our results will consider the $n \rightarrow \infty$ limit. We will be interested in the following detection and recovery problems. For the detection problems, we are given a single sample from either the spiked distribution $\text{WT}(d, \lambda, \mathcal{X})$ or the corresponding unspiked distribution $\text{WT}(d)$ (say each is chosen with probability $1/2$) and our goal is to decide which distribution the sample came from.

- **Strong detection:** Distinguish between the spiked and unspiked distributions with success probability $1 - o(1)$ as $n \rightarrow \infty$.
- **Weak detection:** Distinguish between the spiked and unspiked distributions with success probability $\frac{1}{2} + \varepsilon$ as $n \rightarrow \infty$, for some $\varepsilon > 0$ that does not depend on n .
- **Weak recovery:** Given a sample from the spiked model, output a unit vector \hat{x} with nontrivial correlation with the spike: $\mathbb{E}[\langle x, \hat{x} \rangle^d] \geq \varepsilon$ for some $\varepsilon > 0$ that does not depend on n .

Note that the exponent of d in $\langle x, \hat{x} \rangle^d$ captures the fact that when d is even we can only hope to learn the spike up to a global sign flip.

The three problems above are related as follows. Clearly strong detection implies weak detection, but there is no formal connection between detection and recovery in general (see e.g. [BMV⁺17] for a simple counterexample). Typically, however, strong detection and weak recovery tend to be equivalent in the sense that they exhibit the same threshold for λ above which they are possible. For spiked tensors with $d \geq 3$, we will see that this also tends to coincide with the weak detection threshold. However, for matrices ($d = 2$), weak detection is actually possible below the strong detection threshold; in fact it is possible for any $\lambda > 0$ simply by inspecting the trace of the matrix.

1.2 Summary of results

Our results apply to a wide range of spike priors, but here we present specific results for three priors:

- the spherical prior \mathcal{X}_{sph} , in which x is drawn uniformly from the unit sphere in \mathbb{R}^n ,
- the Rademacher prior \mathcal{X}_{Rad} , in which the entries of x are drawn i.i.d. from $\{\pm 1/\sqrt{n}\}$,
- the sparse Rademacher prior $\mathcal{X}_{\text{spRad}}(\rho)$ for $\rho \in (0, 1]$, in which a random support of exactly ρn entries is chosen uniformly, and those entries of x are drawn i.i.d. from $\{\pm 1/\sqrt{\rho n}\}$ (with all other entries 0). (We restrict to n for which ρn is an integer.)

Although we will give bounds for every d , our results are most precisely stated in the limit $d \rightarrow \infty$ in which case the lower and upper bounds match asymptotically.

Theorem 1.3 (Spherical prior, large d). *Consider the spherical prior \mathcal{X}_{sph} . There exist bounds $\lambda_{\text{sph},d}^*$ and $\Lambda_{\text{sph},d}^*$, both behaving as $\sqrt{2 \log d} + o(1)$ as $d \rightarrow \infty$, and with $\Lambda_{\text{sph},d}^* < \mu_d$ for each d , such that*

- if $\lambda < \lambda_{\text{sph},d}^*$ then weak detection and weak recovery are impossible,
- if $\lambda > \Lambda_{\text{sph},d}^*$ then strong detection and weak recovery are possible.

Recall that μ_d is the value of the injective norm of $\text{WT}(d)$, which also behaves as $\sqrt{2 \log d} + o(1)$; see Theorem 3.3 for a formal statement, including the value of μ_d . Explicit formulas for the bounds are given by Theorem 3.2 ($\lambda_{\text{sph},d}^*$) and Theorem 3.4 ($\Lambda_{\text{sph},d}^*$). We obtain more detailed asymptotic descriptions in Appendix A:

$$\begin{aligned}\mu_d^2 &= 2 \log d + 2 \log \log d + o(1), \\ (\Lambda_{\text{sph},d}^*)^2 &= 2 \log d + 2 \log \log d + o(1), \\ (\lambda_{\text{sph},d}^*)^2 &= 2 \log d + 2 \log \log d + 2 - 4 \log 2 + o(1)\end{aligned}$$

in the limit $d \rightarrow \infty$. Hence the quantities μ_d , $\Lambda_{\text{sph},d}$, and $\lambda_{\text{sph},d}$ all behave as $\sqrt{2 \log d} + o(1)$. Our lower bound closes a $\sqrt{2}$ -factor gap in [MRZ15]. A non-rigorous calculation via the replica method (Appendix B) suggests that the true statistical threshold matches the asymptotics $2 \log d + 2 \log \log d + o(1)$ of the upper bound. These various quantities are depicted in Figure 1.

It is clear that once $\lambda > \mu_d$, strong detection is possible (by thresholding the injective norm). Our upper bound implies that for each d , the spherically-spiked and unspiked tensor models can be distinguished as soon as $\lambda > \Lambda_{\text{sph},d}^*$, a threshold strictly below μ_d ; indeed, we show (Theorem 3.4) that the injective norm of $\text{WT}(d, \lambda, \mathcal{X}_{\text{sph}})$ exceeds that of $\text{WT}(d)$ for such λ . This mirrors the eigenvalue transition for $d = 2$ (see Theorem 1.2), in which an eigenvalue leaves the spectrum bulk when $\lambda > 1$, exhibiting a gap from when the spike exceeds the bulk in spectral norm at $\lambda = 2$. The results above imply that the size of this ‘‘BBP gap’’ lies between $\mu_d - \Lambda_{\text{sph},d}^*$ and $\mu_d - \lambda_{\text{sph},d}^*$ and is therefore of order $1/\sqrt{\log d}$ as $d \rightarrow \infty$.

We now present our results for discrete priors. One qualitative difference from the above is that for discrete priors, the statistical threshold tends to remain bounded as $d \rightarrow \infty$:

Theorem 1.4 (Rademacher prior, large d). *Consider the Rademacher prior \mathcal{X}_{Rad} . There exists a bound $\lambda_{\text{Rad},d}^*$, with $\lim_{d \rightarrow \infty} \lambda_{\text{Rad},d}^* = 2\sqrt{\log 2}$, such that*

- if $\lambda < \lambda_{\text{Rad},d}^*$ then weak detection and weak recovery are impossible,
- if $\lambda > 2\sqrt{\log 2}$ then strong detection and weak recovery are possible.

Theorem 1.5 (Sparse Rademacher prior, large d). *Consider the sparse Rademacher prior $\mathcal{X}_{\text{spRad}}(\rho)$ with $\rho \in (0, 1]$ fixed. There exists a bound $\lambda_{\text{spRad}(\rho),d}^*$, with $\lim_{d \rightarrow \infty} \lambda_{\text{spRad}(\rho),d}^* = 2\sqrt{H(\rho) + \rho \log 2}$, such that*

- if $\lambda < \lambda_{\text{spRad}(\rho),d}^*$ then weak detection and weak recovery are impossible,
- if $\lambda > 2\sqrt{H(\rho) + \rho \log 2}$ then strong detection and weak recovery are possible.

Here H denotes the binary entropy: $H(\rho) = -\rho \log \rho - (1 - \rho) \log(1 - \rho)$.

See Theorems 4.6 and 4.11 for explicit formulas for the lower bounds. The upper bounds follow from Proposition 4.1. We remark that the hardest case for sparse Rademacher is $\rho = 2/3$ (maximizing $H(\rho) + \rho \log 2$), where the three types of entries $(1/\sqrt{\rho n}, -1/\sqrt{\rho n}, 0)$ occur in equal proportion. In the above three cases (spherical, Rademacher, sparse Rademacher), the lower and upper bounds agree up to a $1 + o(1)$ factor as $d \rightarrow \infty$. The bounds λ^*, Λ^* for each d are described as finite-dimensional optimization problems and are easy to compute numerically. See Figure 1 for a comparison of the bounds for various values of d .

The special case of $d = 2$ has been previously understood for the spherical and Rademacher priors [DAM16, MRZ15, PWBM16]. Namely, the threshold for both strong detection and weak recovery occurs precisely at $\lambda = 1$, matching the eigenvalue transition.

Another asymptotic regime that we consider is the sparse Rademacher prior with d fixed and sparsity $\rho \rightarrow 0$.

Theorem 1.6 (Sparse Rademacher prior, $\rho \rightarrow 0$). *Fix $d \geq 2$ and consider $\mathcal{X}_{\text{spRad}}(\rho)$. There exists a bound $\lambda_{\text{spRad}(\rho),d}^*$ that behaves as $2\sqrt{-\rho \log \rho + O(\rho)}$ in the limit $\rho \rightarrow 0$, such that*

- if $\lambda < \lambda_{\text{spRad}(\rho),d}^*$ then strong detection is impossible,
- if $\lambda > 2\sqrt{H(\rho) + \rho \log 2} = 2\sqrt{-\rho \log \rho + O(\rho)}$ then strong detection and weak recovery are possible.

For $d \geq 3$, the lower bound also rules out weak detection and weak recovery.

The $d = 2$ case of the above theorem was previously considered by [BMVX16] where they give a tight upper bound and a lower bound that is loose⁶ by a factor of $\sqrt{2}$. Here we close the $\sqrt{2}$ gap by improving the lower bound. Our upper bounds for discrete priors are straightforward generalizations of their upper bound, based on exhaustive search over all possible spikes. The recovery problem for the sparse Rademacher prior (with $d = 2$) has also been studied using tools from statistical physics. In particular, the weak recovery threshold is known exactly, as well as the optimal recovery quality at each value of λ [LKZ15, KXZ16, BDM⁺16, LM16]. (These results actually consider a variant of the sparse Rademacher prior where the entries of the spike are i.i.d. but we believe this does not change the information-theoretic limits of the problem.) See Figure 2 for a comparison of bounds for $d = 2$ sparse Rademacher.

In Appendix B, we present non-rigorous calculations through the replica method, predicting the precise location of the phase transition for each d , in the spherical and Rademacher cases. For the Rademacher prior, it can be deduced from [KM09] that the replica prediction is a rigorous upper bound on the true threshold; see Appendix B. More recent independent work [LML⁺17] confirms the replica prediction for various spiked tensor problems, but only for the recovery problem (not detection).

Our upper bounds for discrete priors (Rademacher and sparse Rademacher) are obtained by straightforward analysis of the MLE (maximum likelihood estimator), i.e. the non-efficient procedure that enumerates all possible spikes and tests which one is most likely. Our upper bound for the spherical prior showing strict separation from the injective norm, is proven by harnessing known properties of the $d = 2$ case (eigenvalue transition in spiked matrices). Our lower bound techniques are discussed in the next section.

The rest of the paper is organized as follows. In Section 2 we present the main tools used in our lower bounds, including the statement of our main lower bound theorem (Theorem 2.7) along with a sketch of its proof using our noise conditioning method. In Section 3 we prove our lower and upper bounds for the spherical prior (assuming Theorem 2.7). In Section 4 we prove our lower and upper bounds for discrete priors (again assuming Theorem 2.7), including the Rademacher and sparse Rademacher priors. In Section 5 we prove Theorem 2.7. Some results are deferred to the appendix, including the replica calculations for the spherical and Rademacher priors (Appendix B).

2 Lower bound techniques

2.1 χ^2 -divergence, contiguity, and non-recovery

χ^2 -divergence and TV distance For probability distributions P and Q , with P absolutely continuous with respect to Q , the χ^2 -divergence is defined as

$$\chi^2(P \parallel Q) = \mathbb{E}_Q \left[\left(\frac{dP}{dQ} \right)^2 \right] - 1.$$

When P, Q are continuous distributions with densities p, q respectively, note that this is simply

$$\chi^2(P \parallel Q) = \mathbb{E}_{y \sim Q} \left[\left(\frac{p(y)}{q(y)} \right)^2 \right] - 1.$$

⁶After the first appearance of this paper, an update [BMV⁺17] independently closed the asymptotic $\sqrt{2}$ -factor gap from [BMVX16], using a different modification of the second moment method.

Let $P_n = \text{WT}_n(d, \lambda, \mathcal{X})$ be the spiked ensemble, and let $Q_n = \text{WT}_n(d)$ be its unspiked analogue. Our lower bounds will proceed by bounding $\chi^2(P_n \parallel Q_n)$ (or slight variants) and drawing various conclusions from the value.

For tensors of order $d \geq 3$, we will bound the TV (total variation) distance via the inequality

$$\text{TV}(P_n, Q_n) \leq 2\sqrt{\chi^2(P_n \parallel Q_n)}. \quad (1)$$

In particular, if we can establish that $\chi^2(P_n \parallel Q_n) = o(1)$ as $n \rightarrow \infty$, it follows that $\text{TV}(P_n, Q_n) = o(1)$, implying that weak detection is impossible.

χ^2 -divergence and contiguity For matrices ($d = 2$), we cannot hope to show $\text{TV}(P_n, Q_n) = o(1)$ because weak detection is possible for all $\lambda > 0$ using the trace of the matrix. We instead show lower bounds against strong detection, using the notion of *contiguity* introduced by Le Cam [LC60].

Definition 2.1 ([LC60]). Let P_n, Q_n be sequences of distributions defined on the measurable space $(\Omega_n, \mathcal{F}_n)$. We say that P_n is *contiguous* to Q_n , and write $P_n \triangleleft Q_n$, if for any sequence A_n of events with $Q_n(A_n) \rightarrow 0$ as $n \rightarrow \infty$, we also have $P_n(A_n) \rightarrow 0$ as $n \rightarrow \infty$.

Note that $P_n \triangleleft Q_n$ implies that strong detection is impossible; to see this, suppose we had a distinguisher and let A_n be the event that it outputs ‘ P_n ’ to arrive at a contradiction. (The definition of contiguity is asymmetric but contiguity in either direction implies that strong detection is impossible.)

The following second moment method connects the χ^2 -divergence with contiguity.

Lemma 2.2 (see e.g. [MRZ15, BMV⁺17]). *If $\chi^2(P_n \parallel Q_n) = O(1)$ as $n \rightarrow \infty$, then $P_n \triangleleft Q_n$.*

This is referred to as a *second moment method* due to the “second moment” $\mathbb{E}_{Q_n} \left(\frac{dP_n}{dQ_n} \right)^2$ appearing in the definition of χ^2 -divergence.

Proof. Let A_n be a sequence of events. Using Cauchy–Schwarz,

$$P_n(A_n) = \int_{A_n} \frac{dP_n}{dQ_n} dQ_n \leq \sqrt{\int_{A_n} \left(\frac{dP_n}{dQ_n} \right)^2 dQ_n} \cdot \sqrt{\int_{A_n} dQ_n}.$$

The bound on χ^2 -divergence implies that the first factor on the right-hand side is bounded. This means if $Q_n(A_n) \rightarrow 0$ then also $P_n(A_n) \rightarrow 0$. \square

To summarize, we now know that if the χ^2 -divergence is $o(1)$ then weak detection is impossible, and if the χ^2 -divergence is $O(1)$ then strong detection is impossible.

We remark that computing $\chi^2(Q_n \parallel P_n)$ seems significantly harder than computing $\chi^2(P_n \parallel Q_n)$ (where P_n is spiked and Q_n is unspiked). Establishing contiguity in the opposite direction $Q_n \triangleleft P_n$ typically requires additional work such as the small subgraph conditioning method (see e.g. [MNS15, BMNN16, Wor99]). Our methods will only require us to compute χ^2 -divergence in the ‘easy’ direction.

χ^2 -divergence and conditioning It turns out that in some cases, the χ^2 -divergence can be dominated by extremely rare ‘bad’ events, causing it to be unbounded even though $P_n \triangleleft Q_n$. Towards fixing this, a key observation is that if we replace P_n by a modified distribution \tilde{P}_n that only differs from P_n with probability $o(1)$ as $n \rightarrow \infty$ (i.e. $\text{TV}(P_n, \tilde{P}_n) = o(1)$), this does not affect whether or not the detection and recovery problems can be solved. Therefore, by choosing \tilde{P} to condition on the high-probability ‘good’ events we can hope to make $\chi^2(\tilde{P}_n \parallel Q_n)$ controlled even in some cases when $\chi^2(P_n \parallel Q_n)$ is not. If we can show $\chi^2(\tilde{P}_n \parallel Q_n) = O(1)$ it follows that $\tilde{P}_n \triangleleft Q_n$, which implies $P_n \triangleleft Q_n$ as desired. Prior work (e.g. [BMNN16, BMVX16, PWBM16]) has applied this conditioning technique to ‘good’ events depending on the spike x ; specifically, \tilde{P}_n enforces that the entries of x occur in close-to-typical proportions. As the main technical novelty of the current work, we introduce a technique that we call *noise conditioning* where

we condition on ‘good’ events that depend jointly on the spike x and the noise W . Specifically, we require $\langle W, x^{\otimes d} \rangle$ to have close-to-typical value. This method will allow us to obtain asymptotically tight lower bounds in cases where prior work [MRZ15, BMVX16] has been loose by a constant factor of $\sqrt{2}$. The noise conditioning method is discussed further in Section 2.4.

χ^2 -divergence and non-recovery In the case $d \geq 3$, we are able to show $\text{TV}(P_n, Q_n) = o(1)$ as $n \rightarrow \infty$. In Section 5.5 we deduce from this that weak recovery is impossible.

In the case $d = 2$, non-recovery is less straightforward. Various works have forged a connection between bounded χ^2 -divergence and non-recovery [BMNN16, BMVX16]. For a wide class of problems with additive Gaussian noise, Theorem 4 of [BMVX16] implies that if the χ^2 -divergence is $O(1)$ then weak recovery is impossible. This result cannot be immediately applied in our setting because, due to our noise conditioning technique, we do not exactly have an additive Gaussian noise model. We will leave the proof of non-recovery for $d = 2$ for future work⁷.

2.2 Large deviation rate functions

Our lower bounds on tensor PCA problems will depend on the prior through tail probabilities of the correlation $\langle x, x' \rangle$ of two independent spikes drawn from the prior; note that this quantity is typically of order $1/\sqrt{n}$ but may be as large as 1. We require detailed tail information, which we summarize in two objects: a *rate function*, describing the asymptotic behavior of deviations of order 1, and a *local subgaussian condition* [CCK06], bounding the deviations of size $[0, \varepsilon]$ non-asymptotically.

First we define the rate function $f_{\mathcal{X}}$ corresponding to a prior \mathcal{X} . Intuitively, $f_{\mathcal{X}} : [0, 1) \rightarrow [0, \infty)$ is roughly defined by $\Pr[|\langle x, x' \rangle| \geq t] \approx \exp(-nf_{\mathcal{X}}(t))$. More formally, one should think of $f_{\mathcal{X}}(t)$ as being equal to $-\lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr[|\langle x, x' \rangle| \geq t]$ but we will not technically require this in our definition since it will be ok to use a weaker (smaller) rate function than the ‘true’ one (although we will always use the ‘true’ one in our examples). For technical reasons, we formally define the rate function as follows.

Definition 2.3. Let $\mathcal{X} = \{\mathcal{X}_n\}$ be a prior supported on unit vectors in \mathbb{R}^n . For x, x' drawn independently from \mathcal{X}_n and $t \in [0, 1)$, let

$$f_{n, \mathcal{X}}(t) = -\frac{1}{n} \log \Pr[|\langle x, x' \rangle| \geq t].$$

Suppose we have $f_{n, \mathcal{X}}(t) \geq b_{n, \mathcal{X}}(t)$ for some sequence of functions $b_{n, \mathcal{X}}$ that converges uniformly on $[0, 1)$ to $f_{\mathcal{X}}$ as $n \rightarrow \infty$. Then we call such $f_{\mathcal{X}}$ the *rate function* of the prior \mathcal{X} .

Remark 2.4. The condition $f_{n, \mathcal{X}} \geq b_{n, \mathcal{X}} \rightarrow f_{\mathcal{X}}$ (uniformly) is satisfied if we have a tail bound of the form $\Pr[|\langle x, x' \rangle| \geq t] \leq \text{poly}(n) \exp(-nf_{\mathcal{X}}(t))$.

Remark 2.5. The above definition coincides with the standard notion of a rate function for the random variable $|\langle x, x' \rangle|$ provided that the latter rate function is monotone increasing. This is indeed the case for all priors that we consider as well as for many other natural priors (e.g. any i.i.d. prior).

Note the following basic properties of the rate function (which we can assume without loss of generality): $f_{\mathcal{X}}(0) = 0$ and $f_{\mathcal{X}}$ is monotone increasing. We will see that for continuous priors (such as the spherical prior), the rate function blows up to infinity at $t = 1$, whereas for discrete priors (such as Rademacher) it remains finite at $t = 1$.

We now present the local subgaussianity condition which gives tighter control of the small deviations.

Definition 2.6. Let \mathcal{X} be a prior supported on unit vectors in \mathbb{R}^n . We say that \mathcal{X} is *locally subgaussian with constant σ^2* if for all $\eta > 0$, there exist $C, T > 0$ such that

$$\Pr_{x, x' \sim \mathcal{X}}[|\langle x, x' \rangle| \geq t] \leq C \exp\left(-\frac{nt^2}{2\sigma^2 + \eta}\right) \quad \forall t \in [0, T].$$

⁷After the first appearance of this paper, an update [BMV⁺17] resolved this, extending the [BMVX16] result to the noise-conditioned case.

Note that if $\langle x, x' \rangle$ is (σ^2/n) -subgaussian, then \mathcal{X} is locally subgaussian with constant σ^2 . However, the converse is false: for instance, the sparse Rademacher prior with sufficiently small density ρ has a strictly better local subgaussian constant than its subgaussian constant.

2.3 Main lower bound theorem

The following theorem encapsulates our lower bound result:

Theorem 2.7. *Let $d \geq 2$, and let \mathcal{X} be a prior supported on unit vectors in \mathbb{R}^n . Suppose the following holds for some $\lambda^* > 0$.*

(i) \mathcal{X} has a rate function $f_{\mathcal{X}}$,

(ii) $f_{\mathcal{X}}(t) \geq \frac{(\lambda^*)^2}{2} \frac{t^d}{1+t^d}$ for all $t \in [0, 1)$,

(iii) \mathcal{X} is locally subgaussian with some constant σ^2 ,

(iv) if $d = 2$, $\lambda^* \leq 1/\sigma$.

Then for all $\lambda < \lambda^*$:

- $\text{WT}(d, \lambda, \mathcal{X})$ is contiguous to $\text{WT}(d)$, so no hypothesis test to distinguish them achieves full power (i.e. strong detection is impossible),
- if $d \geq 3$, $\text{TV}(\text{WT}(d, \lambda, \mathcal{X}), \text{WT}(d)) = o(1)$, so every hypothesis test has asymptotically zero power (i.e. weak detection is impossible),
- if $d \geq 3$, no estimator $\hat{x} = \hat{x}(T)$, when given a sample $T \sim \text{WT}(d, \lambda, \mathcal{X})$, achieves expected correlation $\mathbb{E}|\langle x, \hat{x} \rangle|$ (where x is the true spike) bounded above 0 as $n \rightarrow \infty$ (i.e. weak recovery is impossible).

The proof is deferred to Section 5. The essential condition is (ii) and the others exist for technical reasons. We remark that condition (iv) is simply a slight strengthening of condition (ii) in the following sense.

Proposition 2.8. *Let $d = 2$. Suppose the rate function $f_{\mathcal{X}}$ admits a local Chernoff-style tail bound of the following form: there exist $C, T > 0$ such that*

$$\Pr[|\langle x, x' \rangle| \geq t] \leq C \exp(-nf_{\mathcal{X}}(t)) \quad \forall t \in [0, T].$$

Then condition (ii) implies conditions (iii) and (iv).

Note that this is stronger than Remark 2.4 because there is no $\text{poly}(n)$ factor.

Proof. Fix $\eta > 0$. Let $\varepsilon > 0$, to be chosen later. For all $t \in [0, T']$ we have $\frac{t^2}{1+t^2} \geq \frac{t^2}{1+T'^2}$ and so there exists $T' > 0$ such that $\frac{t^2}{1+t^2} \geq (1 - \varepsilon)t^2$ for all $t \in [0, T']$. For $t \in [0, \min(T, T')]$,

$$\Pr[|\langle x, x' \rangle| \geq t] \leq C \exp(-nf_{\mathcal{X}}(t)) \leq C \exp\left(-n \frac{(\lambda^*)^2}{2} \frac{t^2}{1+t^2}\right) \leq C \exp\left(-\frac{nt^2}{2} (\lambda^*)^2 (1 - \varepsilon)\right).$$

Choose ε small enough so that this is at most $C \exp\left(-\frac{nt^2}{2/(\lambda^*)^2 + \eta}\right)$, the local subgaussian condition. \square

2.4 Proof idea: noise conditioning

The proof of our lower bounds hinge on a simple ‘‘noise conditioning’’ modification to the second moment method (χ^2 -divergence) for proving contiguity. We have found that this method yields asymptotically tight bounds in various regimes where the basic second moment method is loose by a factor of $\sqrt{2}$. (Below we will see how the number $\sqrt{2}$ arises.) We expect that the same idea could be used to close the various other $\sqrt{2}$ -factor gaps that exist for contiguity arguments in prior work: the sparse stochastic block model in the regime of a constant number of equally-sized communities with low signal-to-noise [BMNN16]; submatrix localization in the limit of a large constant number of blocks [BMVX16]; Gaussian mixture clustering with a large constant number of clusters [BMVX16]; and synchronization over a finite group of large constant size (with either truth-or-Haar noise or Gaussian noise on all frequencies) [PWBM16]. However, we leave investigation of these other problems for future work⁸.

The following is a proof sketch of our main lower bound theorem (Theorem 2.7) using the noise conditioning method. As discussed in Section 2.1, if the second moment $\mathbb{E}_{Q_n} \left(\frac{dP_n}{dQ_n} \right)^2$ is bounded as $n \rightarrow \infty$ then P_n is contiguous to Q_n and so it is impossible to reliably distinguish the two distributions. Taking P_n to be the spiked tensor model and Q_n the corresponding unspiked model, the second moment can be computed to be

$$\mathbb{E}_{Q_n} \left(\frac{dP_n}{dQ_n} \right)^2 = \mathbb{E}_{x, x' \sim \mathcal{X}} \mathbb{E}_{T \sim Q_n} \exp \left(\frac{n\lambda}{2} \langle T, x^{\otimes d} + x'^{\otimes d} \rangle - \frac{n\lambda^2}{2} \right) \quad (2)$$

where x, x' are drawn independently from the prior \mathcal{X} . The standard approach used by previous work is to next apply the Gaussian moment-generating function to compute the expectation over $T \sim Q_n$ in closed form, yielding

$$\mathbb{E}_{x, x' \sim \mathcal{X}} \exp \left(\frac{n\lambda^2}{2} \beta \right) \quad (3)$$

where $\beta = \langle x, x' \rangle^d$. Our approach is instead to use the fact that (as discussed in Section 2.1) we can change P_n to \tilde{P}_n that excludes low-probability bad events, without affecting contiguity. Specifically, we take \tilde{P}_n to condition on (informally) $\langle T, x^{\otimes d} \rangle \approx \lambda$, which is a high-probability event under P_n . Now, instead of (2), the second moment can be computed to be

$$\begin{aligned} \mathbb{E}_{Q_n} \left(\frac{d\tilde{P}_n}{dQ_n} \right)^2 &\approx \mathbb{E}_{x, x' \sim \mathcal{X}} \mathbb{E}_{T \sim Q_n} \mathbb{1}_{\langle T, x^{\otimes d} \rangle \approx \lambda} \mathbb{1}_{\langle T, x'^{\otimes d} \rangle \approx \lambda} \exp \left(\frac{n\lambda}{2} \langle T, x^{\otimes d} + x'^{\otimes d} \rangle - \frac{n\lambda^2}{2} \right) \\ &\approx \mathbb{E}_{x, x' \sim \mathcal{X}} \exp \left(\frac{n\lambda}{2} (2\lambda) - \frac{n\lambda^2}{2} \right) \Pr_{T \sim Q_n} [\langle T, x^{\otimes d} \rangle \approx \langle T, x'^{\otimes d} \rangle \approx \lambda], \end{aligned}$$

which can be computed to be

$$\mathbb{E}_{x, x' \sim \mathcal{X}} \exp \left(\frac{n\lambda^2}{2} \frac{\beta}{1 + \beta} \right) \quad (4)$$

where again $\beta = \langle x, x' \rangle^d$. Note that when the contribution from $\beta = 1$ dominates, this gives a factor of $\sqrt{2}$ advantage on λ compared to (3). This explains the gaps of $\sqrt{2}$ present in prior work, and allows us to close them.

The final step is to bound (4) using the rate function $f_{\mathcal{X}}$ of \mathcal{X} . Roughly, the rate function gives us the tail bound $\Pr[|\langle x, x' \rangle| \geq t] \approx \exp(-nf_{\mathcal{X}}(t))$. To use this, we write (4) as a tail bound integral and then apply

⁸After the first appearance of this paper, an update [BMV⁺17] closed the asymptotic gaps for sparse PCA and submatrix localization using a different modification of the second moment method; the problem for Gaussian mixtures remains open, to our knowledge.

a change of variables:

$$\begin{aligned} \mathbb{E}_{x, x' \sim \mathcal{X}} \exp \left(\frac{n\lambda^2}{2} \frac{\beta}{1+\beta} \right) &= \int_0^\infty \Pr \left[\exp \left(\frac{n\lambda^2}{2} \frac{\beta}{1+\beta} \right) \geq u \right] du \\ &\simeq \int_0^1 \Pr[|\langle x, x' \rangle| \geq t] \exp \left(\frac{n\lambda^2}{2} \frac{t^d}{1+t^d} \right) \frac{n\lambda^2}{2} \frac{dt^{d-1}}{(1+t^d)^2} dt \end{aligned}$$

where $u = \exp \left(\frac{n\lambda^2}{2} \frac{t^d}{1+t^d} \right)$,

$$\approx \int_0^1 \exp \left[n \left(-f_{\mathcal{X}}(t) + \frac{\lambda^2}{2} \frac{t^d}{1+t^d} \right) \right] dt.$$

Note that this would be bounded as $n \rightarrow \infty$ if we had $-f_{\mathcal{X}}(t) + \frac{\lambda^2}{2} \frac{t^d}{1+t^d} < 0$ for all $t \in [0, 1]$, which gives condition (ii) in Theorem 2.7. Note however that this cannot be satisfied at $t = 0$ (since $f_{\mathcal{X}}(t) = 0$), so our proof will require extra work in order to handle the $t \in [0, \varepsilon]$ interval. This is where the local subgaussianity condition will be used.

3 Results for the spherical prior

In this section we discuss lower and upper bounds for the *spherical prior* \mathcal{X}_{sph} , the uniform prior on the unit sphere.

3.1 Rate function

Lemma 3.1. *The spherical prior \mathcal{X}_{sph} has rate function $f_{\mathcal{X}_{\text{sph}}}(t) = -\frac{1}{2} \log(1-t^2)$ and is locally subgaussian (with some constant).*

Proof. Note that with $x, x' \sim \mathcal{X}_{\text{sph}}$, $\frac{1}{2} + \frac{1}{2} \langle x, x' \rangle \sim \text{Beta}(\frac{n}{2}, \frac{n}{2})$. It follows from subgaussianity of the beta distribution [Eld16] that $\langle x, x' \rangle$ is $O(1/n)$ -subgaussian, so that \mathcal{X}_{sph} is locally subgaussian (with some constant). Moreover, note that

$$-\frac{1}{n} \log \Pr[|\langle x, x' \rangle| \geq t] = -\frac{1}{n} \log \left[2I_{(1-t)/2} \left(\frac{n}{2}, \frac{n}{2} \right) \right],$$

where I denotes the incomplete beta function. By equation 8.18.9 from [DLMF], this converges uniformly to $-\frac{1}{2} \log(1-t^2)$ on $[0, 1]$. Thus we satisfy the convergence conditions in Definition 2.3, taking $b_{n, \mathcal{X}} = f_{n, \mathcal{X}}$. \square

3.2 Lower bound

We can now complete the lower bound portion of Theorem 1.3. From the discussion above together with Theorem 2.7, we have the following.

Theorem 3.2. *Consider the spherical prior \mathcal{X}_{sph} with $d \geq 3$. Let $\lambda_{\text{sph}, d}^*$ be the supremum of all values λ^* for which*

$$\frac{(\lambda^*)^2}{2} \frac{t^d}{1+t^d} \leq -\frac{1}{2} \log(1-t^2) \quad \forall t \in [0, 1].$$

Then weak detection and weak recovery are impossible for all $\lambda < \lambda_{\text{sph}, d}^$.*

Note that the $d = 2$ case is already well-understood: strong detection and weak recovery are possible when $\lambda > 1$ and impossible when $\lambda < 1$ [MRZ15, PWBM16], matching the eigenvalue transition. (Weak recovery is possible for all $\lambda > 0$ simply by inspecting the trace of the matrix.) In Appendix A we establish the following asymptotics as $d \rightarrow \infty$:

$$(\lambda_{\text{sph}, d}^*)^2 = 2 \log d + 2 \log \log d + 2 - 4 \log 2 + o(1).$$

3.3 Upper bound

In this section we prove an upper bound showing that for sufficiently large λ , detection and recovery are possible in the spherically-spiked tensor model. Indeed, our upper bound will apply to any prior supported on the unit sphere. Recall that the injective norm of a d -tensor is defined as $\|T\| = \max_{\|x\|=1} \langle x^{\otimes d}, T \rangle$. As noted by [RM14], as soon as the injective norm λ of the spike exceeds that of the noise, detection and recovery are possible. We improve on this bound, showing that the injective norm of the spiked model exceeds that of the unspiked model at a lower threshold. This is achieved by studying how perturbations away from the spike can line up with fluctuations in the noise to achieve a large injective norm.

For each d , the precise limit value (in probability, as $n \rightarrow \infty$) of $\|\text{WT}(d)\|$ has been non-rigorously computed using the replica method [CS92]; following [RM14], we refer to this limit value as μ_d (although our normalization differs from theirs). This was later proved rigorously, first only for even values of d [Tal06a, ABC13] (see [RM14] for a summary) and later for all d [Sub15].

Theorem 3.3 ([CS92, Tal06a, ABC13, RM14, Sub15]). *Fix $d \geq 3$. Define $\mu_d = x\sqrt{2/d}$ where $x \geq 2\sqrt{d-1}$ is the unique solution to*

$$\frac{2-d}{d} - \log\left(\frac{dz^2}{2}\right) + \frac{d-1}{2}z^2 - \frac{2}{d^2z^2} = 0, \quad z(x) = \frac{1}{(d-1)\sqrt{2d}} \left(x - \sqrt{x^2 - 4(d-1)}\right). \quad (5)$$

With probability $1 - o(1)$ as $n \rightarrow \infty$ we have

$$\|\text{WT}(d)\| = \mu_d + o(1).$$

Now suppose $T = \lambda x^{\otimes d} + W$ is a spiked d -tensor. Clearly if $\lambda > \mu_d$ then we will have $\|T\| > \mu_d$ with probability $1 - o(1)$ and so strong detection is possible by thresholding the injective norm. However, recall that for matrices ($d = 2$) this is not tight: $\mu_2 = 2$ yet detection is possible for all $\lambda > 1$ (see Theorem 1.2). Our result will imply that for all d it remains true that the detection threshold λ^* is strictly less than μ_d , but the gap between them shrinks as $d \rightarrow \infty$. The argument simply harnesses the $d = 2$ result in a “black box” fashion (which, perhaps surprisingly, gives quite a good upper bound; see Figure 1). We will first give a lower bound on the injective norm of a spiked tensor, stronger than the trivial lower bound of λ . This easily implies results for strong detection and weak recovery (Corollary 3.5 below).

Theorem 3.4. *Fix $d \geq 3$ and $\lambda \geq 0$. Let \mathcal{X} be any prior supported on the unit sphere in \mathbb{R}^d . Let $T = \lambda x^{\otimes d} + W$ be a spiked tensor drawn from $\text{WT}(d, \lambda, \mathcal{X})$. With probability $1 - o(1)$ as $n \rightarrow \infty$ we have*

$$\|T\| \geq L_d(\lambda) - o(1)$$

where

$$L_d(\lambda) = \max_{m \in [0,1]} m^d \left(\lambda + \sqrt{\frac{2d}{d-1}} \sqrt{M(1+M)} \right), \quad M(m) = (d-1) \frac{1-m^2}{m^2}. \quad (6)$$

Proof. By Gaussian spherical symmetry we can assume without loss of generality that the spike is $x = e_1$, the first standard basis vector; thus we write $T = \lambda e_1^{\otimes d} + W$ with $W \sim \text{WT}(d)$. Writing a general unit vector as $v = me_1 + \sqrt{1-m^2}u$, where $u \perp e_1$ is a unit vector, we expand $v^{\otimes d}$ to write

$$\langle T, v^{\otimes d} \rangle = \lambda m^d + \sum_{i=0}^d \binom{d}{i} m^{d-i} (1-m^2)^{i/2} \langle W_i, u^{\otimes i} \rangle,$$

where W_i is the sub-tensor of W for which the first i indices are not 1 and the remaining indices are 1.

Note that the W_i are independent. We will choose u to optimize the interactions with W_1 and W_2 , independently from all other W_i . It follows that for all $i \neq 1, 2$, the terms $\langle W_i, u^{\otimes i} \rangle$ are $o(1)$ as $n \rightarrow \infty$ with high probability, and we have

$$\langle T, v^{\otimes d} \rangle = \lambda m^d + dm^{d-1} \sqrt{1-m^2} \langle W_1, u \rangle + \binom{d}{2} m^{d-2} (1-m^2) u^\top W_2 u + o(1). \quad (7)$$

Specifically, let us take u to be the top eigenvector of the auxiliary spiked matrix

$$Y = \beta \frac{d}{2} W_1 W_1^\top + \sqrt{\frac{d(d-1)}{2}} W_2$$

for some parameter $\beta \geq 1$ to be optimized later. Since $\sqrt{\frac{d}{2}} W_1$ has norm converging to 1 in probability and $\sqrt{\frac{d(d-1)}{2}} W_2$ is a Gaussian Wigner matrix, the classical eigenvalue transition for spiked Wigner matrices (Theorem 1.2) implies that the top eigenvalue $u^\top Y u$ converges almost surely to $\beta + 1/\beta$ (as $n \rightarrow \infty$), and the top eigenvector u is correlated with the normalized spike as $\langle u, W_1 \sqrt{d/2} \rangle^2 \rightarrow 1 - 1/\beta^2$ almost surely [Péc06, FP07, CDMF09, BGN11]. Noting that

$$\beta + \frac{1}{\beta} + o(1) = u^\top Y u = \frac{\beta d}{2} \langle u, W_1 \rangle^2 + \sqrt{\frac{d(d-1)}{2}} u^\top W_2 u,$$

we solve for $u^\top W_2 u = \sqrt{\frac{2}{d(d-1)}} \cdot \frac{2}{\beta} + o(1)$ and plug into (7) to obtain

$$\|T\| \geq \langle v^{\otimes d}, T \rangle = \lambda m^d + m^{d-1} \sqrt{2d(1-m^2)(1-1/\beta^2)} + m^{d-2}(1-m^2) \sqrt{2d(d-1)}/\beta + o(1).$$

To obtain the strongest possible bound, we maximize this expression over all $m \in [0, 1]$ and $\beta \geq 1$. We can optimize in closed form for

$$\beta = \left(1 + \frac{m^2}{(1-m^2)(d-1)} \right)^{1/2}.$$

The result as stated now follows by simple algebra. \square

Corollary 3.5. *Fix $d \geq 3$ and $\lambda \geq 0$. Let \mathcal{X} be any prior supported on the unit sphere in \mathbb{R}^d . Let μ_d be given by Theorem 3.3 and let $L_d(\lambda)$ be defined as in Theorem 3.4. If $L_d(\lambda) > \mu_d$ then strong detection and weak recovery are possible for $\text{WT}(d, \lambda, \mathcal{X})$.*

Proof. If $T = \lambda x^{\otimes d} + W$ we have $\|T\| \geq L_d(\lambda) - o(1)$ with probability $1 - o(1)$ as $n \rightarrow \infty$. If instead $T = W$ is unspiked, we have $\|T\| \leq \mu_d + o(1)$. It follows that strong detection is possible by thresholding $\|T\|$.

To perform weak recovery given $T = \lambda x^{\otimes d} + W$, output the unit vector v maximizing $\langle v^{\otimes d}, T \rangle$. With probability $1 - o(1)$ we have

$$L_d(\lambda) - o(1) \leq \|T\| = \langle v^{\otimes d}, T \rangle = \lambda \langle v, x \rangle^d + \langle v^{\otimes d}, W \rangle \leq \lambda \langle v, x \rangle^d + \mu_d + o(1)$$

and so $\langle v, x \rangle^d \geq \frac{1}{\lambda} (L_d(\lambda) - \mu_d) - o(1)$, which is bounded above 0 as $n \rightarrow \infty$ (by assumption). Therefore weak recovery is possible. \square

To prove the upper bound portion of Theorem 1.3, we let $\Lambda_{\text{sph},d}^*$ be the point at which $\lambda > \Lambda_{\text{sph},d}^*$ implies $L_d(\lambda) > \mu_d$. Note that by considering $m = 1$ in Theorem 3.4, we obtain that the injective norm of a spiked tensor is at least the size of the spike (i.e. $L_d(\lambda) \geq \lambda$), as noted in [RM14]. However, the derivative of (6) at $m = 1$ is $-\infty$, implying a strict separation: $L_d(\lambda) > \lambda$ and so $\Lambda_{\text{sph},d}^* < \mu_d$. Therefore, for any d it is possible to distinguish the spiked and unspiked models for some λ strictly less than the injective norm of a random tensor.

In Appendix A we determine asymptotic forms of Theorem 3.3 and Corollary 3.5. Specifically, the value μ_d behaves as $d \rightarrow \infty$ as

$$\mu_d^2 = 2 \log d + 2 \log \log d + 2 + o(1),$$

while $L_d(\lambda) > \mu_d$ as soon as $\lambda > \Lambda_{\text{sph},d}^*$, with

$$(\Lambda_{\text{sph},d}^*)^2 = 2 \log d + 2 \log \log d + o(1).$$

4 Results for discrete priors

In contrast to the upper and lower bounds discussed in the previous section for the spherical prior, which diverge with d , the statistical thresholds for discrete priors tend to remain bounded as $d \rightarrow \infty$. This dichotomy is somewhat surprising given that the spherical and Rademacher thresholds agree exactly ($\lambda = 1$) for $d = 2$. In this section we present some general and some specific results confirming this phenomenon. In particular, we establish upper and lower bounds under the Rademacher and sparse Rademacher priors, which match in either of the limits $d \rightarrow \infty$ and $\rho \rightarrow 0$ (recall that ρ is the density of nonzeros). We thus establish Theorems 1.4, 1.5, and 1.6.

4.1 Upper bounds

A natural approach to hypothesis testing is the MLE (maximum likelihood estimate), the maximum of the tensor form over the support of the prior. For discrete priors of exponential cardinality (including Rademacher and sparse Rademacher), we can control this via a naïve union bound to obtain positive results for both detection and recovery (similar to [BMVX16]). Note however that this only yields non-efficient procedures for detection and recovery, since they require exhaustive search over all possible spikes.

Proposition 4.1. *Let \mathcal{X} be a discrete prior supported on the unit sphere in \mathbb{R}^n and with exponential support:*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log |\text{supp } \mathcal{X}_n| = c.$$

For any $d \geq 2$, when $\lambda > 2\sqrt{c}$, there exists a hypothesis test distinguishing $\text{WT}(d, \lambda, \mathcal{X})$ from $\text{WT}(d)$ with probability $o(1)$ of error as $n \rightarrow \infty$ (i.e. strong detection is possible) and furthermore there exists an estimator that achieves nontrivial correlation with the spike (i.e. weak recovery is possible).

The hypothesis test computes

$$m = \max_{v \in \text{supp } \mathcal{X}} \langle T, v^{\otimes d} \rangle \tag{8}$$

and reports ' P_n ' iff $m \geq \lambda - \varepsilon$ where ε is any fixed constant in the range $0 < \varepsilon < \lambda/2 - \sqrt{c}$. The estimator that achieves weak recovery is the argmax v in (8).

Proof of Proposition 4.1. We will show that thresholding m suitably yields a hypothesis test that distinguishes the spiked and unspiked models with $o(1)$ probability of error of either type.

Suppose first that $T = \lambda x^{\otimes d} + W$ is drawn from the spiked model $\text{WT}(d, \lambda, \mathcal{X})$ with spike x . Then

$$m \geq \langle T, x^{\otimes d} \rangle = \lambda + \langle W, x^{\otimes d} \rangle = \lambda + \mathcal{N}(0, 2/n),$$

which, for any fixed $\varepsilon > 0$, is greater than $\lambda - \varepsilon$ with probability $1 - o(1)$.

On the other hand, suppose that T is drawn from the unspiked model $\text{WT}(d)$. Then by taking a union bound over all $v \in \text{supp } \mathcal{X}$,

$$\begin{aligned} \Pr[m \geq \lambda - 2\varepsilon] &\leq |\text{supp } \mathcal{X}_n| \Pr[\langle W, v^{\otimes d} \rangle \geq \lambda - 2\varepsilon] \\ &= |\text{supp } \mathcal{X}_n| \Pr[\mathcal{N}(0, 2/n) \geq \lambda - 2\varepsilon] \\ &\leq \exp\left(n\left(\frac{1}{n} \log |\text{supp } \mathcal{X}_n| - \frac{1}{4}(\lambda - 2\varepsilon)^2\right)\right), \end{aligned}$$

and so long as $\lambda > 2\sqrt{c}$, we can choose $\varepsilon > 0$ such that this probability of error is $o(1)$. Hence, thresholding m at $\lambda - \varepsilon$, we obtain a hypothesis test with $o(1)$ probability of error of either type.

In order to perform weak recovery given $T = \lambda x^{\otimes d} + W$, output the v that maximizes (8). From the analysis above we have with probability $1 - o(1)$, $\langle T, v^{\otimes d} \rangle \geq \lambda - \varepsilon$ and $\langle W, v^{\otimes d} \rangle \leq \lambda - 2\varepsilon$. This means we have

$$\lambda - \varepsilon \leq \langle T, v^{\otimes d} \rangle = \lambda \langle x, v \rangle^d + \langle W, v^{\otimes d} \rangle \leq \lambda \langle x, v \rangle^d + \lambda - 2\varepsilon$$

and so $\langle x, v \rangle^d \geq \varepsilon/\lambda$, implying weak recovery. \square

The following MAP (maximum a posteriori) result is stronger for non-uniform priors. Recall that the Rényi entropy of order $\alpha \geq 0$ for a discrete random variable X (taking values x_i with probability p_i) is $H_\alpha(X) = \frac{1}{1-\alpha} \log(\sum_i p_i^\alpha)$.

Proposition 4.2. *Fix $d \geq 2$ and let \mathcal{X} be a prior on the unit sphere in \mathbb{R}^n , with bounded Rényi entropy density: there exists $\delta > 0$ with*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} H_{1-\delta}(\mathcal{X}) < \infty.$$

Suppose furthermore that the varentropy $V(\mathcal{X}) \triangleq \text{Var}_{x \sim \mathcal{X}}[-\log \Pr_{\mathcal{X}}(x)]$ is $o(n^2)$. Let $s = \limsup_{n \rightarrow \infty} \frac{1}{n} H_1(\mathcal{X})$, the Shannon entropy density. Then when $\lambda > 2\sqrt{s}$, there exists a hypothesis test distinguishing $\text{WT}(d, \lambda, \mathcal{X})$ from $\text{WT}(d)$ with probability $o(1)$ of error (i.e. strong detection is possible) and furthermore there exists an estimator that achieves nontrivial correlation with the spike (i.e. weak recovery is possible).

The proof is deferred to Appendix E but the idea is the same as above. Note that as H_α is monotonically decreasing in α , the Rényi entropy bound guarantees that the Shannon entropy density is finite.

Remark 4.3. We remark that the optimal hypothesis test for distinguishing two distributions is likelihood ratio thresholding [NP33]. However, the simpler MLE and MAP tests that we used above are much easier to analyze.

4.2 Lower bounds as $d \rightarrow \infty$

As a corollary of our main lower bound theorem (Theorem 2.7) we have the following lower bound in the $d \rightarrow \infty$ limit.

Proposition 4.4. *Suppose that \mathcal{X} has a rate function $f_{\mathcal{X}}(t)$ (Definition 2.3) and is locally subgaussian (with some constant). Suppose the limit $F = \lim_{t \rightarrow 1^-} f_{\mathcal{X}}(t)$ exists. Then for any $\lambda < 2\sqrt{F}$, it holds for sufficiently large d that weak detection and weak recovery are impossible.*

Note that (assuming $f_{\mathcal{X}}$ is continuous at $t = 1$) we have $F = f_{\mathcal{X}}(1) = -\lim_{n \rightarrow \infty} \frac{1}{n} \log(\Pr[x = x'])$ which is the collision entropy density.

Proof. By Theorem 2.7, it suffices to show that, for all $\lambda < 2\sqrt{F}$, it holds for sufficiently large d that $\frac{\lambda^2}{2} \frac{t^d}{1+t^d} \leq f_{\mathcal{X}}(t)$ for all $t \in [0, 1]$. The intuition for the proof is that when d is large, the function $\frac{t^d}{1+t^d}$ is essentially zero for most of the interval $t \in [0, 1)$ but it jumps sharply up to $\frac{1}{2}$ near $t = 1$; it follows that only the value of $f_{\mathcal{X}}$ near $t = 1$ is important.

Let $\varepsilon > 0$, to be chosen later. Provided ε is sufficiently small, local subgaussianity guarantees that $f_{\mathcal{X}}(t) \geq ct^2$ on $[0, \varepsilon]$ for some $c > 0$; thus $\frac{\lambda^2}{2} \frac{t^d}{1+t^d} \leq f_{\mathcal{X}}(t)$ on $[0, \varepsilon]$, for all $d \geq 3$.

We next establish that for d sufficiently large we have $\frac{\lambda^2}{2} \frac{t^d}{1+t^d} \leq f_{\mathcal{X}}(t)$ for all $t \in [\varepsilon, 1 - \varepsilon]$. We have $f_{\mathcal{X}}(\varepsilon) > 0$ by local subgaussianity and so (by monotonicity of $f_{\mathcal{X}}$) we have a positive lower bound on $f_{\mathcal{X}}$ on $[\varepsilon, 1 - \varepsilon]$. By contrast, $\frac{\lambda^2}{2} \frac{t^d}{1+t^d}$ converges uniformly to 0 on $[\varepsilon, 1 - \varepsilon]$ as $d \rightarrow \infty$; the claim follows.

It remains to show the bound on $[1 - \varepsilon, 1]$. Here $f_{\mathcal{X}}(t) \geq f_{\mathcal{X}}(1 - \varepsilon)$, while $\frac{\lambda^2}{2} \frac{t^d}{1+t^d} \leq \frac{\lambda^2}{4}$ (since $x \mapsto \frac{x}{1+x}$ is increasing on $[0, 1]$). It is therefore sufficient to choose ε small enough so that $f_{\mathcal{X}}(1 - \varepsilon) \geq \frac{\lambda^2}{4}$. This is possible due to the assumption $\lambda^2 < 4F$ along with the definition of F . \square

Remark 4.5. This argument goes through even if the collision entropy density F is ∞ , in which case we have that for any λ there exists d such that weak detection and weak recovery are impossible. In other words, the threshold λ^* must necessarily go to infinity as $d \rightarrow \infty$, in contrast to the behavior $\lambda^* = O(1)$ for discrete priors. For instance, recall that the spherical prior has $F = \infty$ and indeed the threshold λ^* diverges with d .

4.3 Rademacher prior

Consider the Rademacher prior \mathcal{X}_{Rad} where each entry of x is drawn i.i.d. from $\{\pm 1/\sqrt{n}\}$. In this section we prove Theorem 1.4 (large d behavior) as well as give an explicit lower bound for each d . The upper bound is immediate from Proposition 4.1.

The rate function of \mathcal{X}_{Rad} follows from Cramér's theorem, with the rate function being the convex dual of the cumulant generating function $\log \cosh(t)$, which is

$$f_{\mathcal{X}_{\text{Rad}}}(t) = \log 2 - H\left(\frac{1+t}{2}\right),$$

where H is the binary entropy: $H(p) = -p \log p - (1-p) \log(1-p)$. The necessary convergence condition is satisfied with $b_{n,\mathcal{X}_{\text{Rad}}}(t) = f_{\mathcal{X}_{\text{Rad}}}(t)$, as the Chernoff bound gives $\Pr[|\langle x, x' \rangle| \geq t] \leq 2 \exp(-n f_{\mathcal{X}_{\text{Rad}}}(t))$ and so $f_{n,\mathcal{X}_{\text{Rad}}}(t) \geq f_{\mathcal{X}_{\text{Rad}}}(t) - \frac{1}{n} \log 2$. Moreover, $\langle x, x' \rangle$ is an average of n i.i.d. Rademacher ± 1 random variables, and thus is $(1/n)$ -subgaussian by Hoeffding's inequality. Hence \mathcal{X}_{Rad} is locally subgaussian with constant 1.

The asymptotic lower bound for large d now follows from Proposition 4.4. By Theorem 2.7 we have the following explicit bound for any fixed d .

Theorem 4.6. *Consider the Rademacher prior \mathcal{X}_{Rad} with $d \geq 3$. Let $\lambda_{\text{Rad},d}^*$ be the supremum of all values λ^* for which*

$$\frac{(\lambda^*)^2}{2} \frac{t^d}{1+t^d} \leq f_{\mathcal{X}_{\text{Rad}}}(t) \quad \forall t \in [0, 1].$$

Then weak detection and weak recovery are impossible for all $\lambda < \lambda_{\text{sph},d}^$.*

As with the spherical prior, the $d = 2$ case is already well-understood: strong detection and weak recovery are possible only above the spectral threshold $\lambda = 1$ [DAM16, PWBM16], and weak detection is possible for all $\lambda > 0$ by taking the trace.

4.4 Sparse Rademacher prior

Recall that the sparse Rademacher prior $\mathcal{X}_{\text{spRad}}(\rho)$ with sparsity $\rho \in (0, 1]$ draws the spike x by choosing at random a subset of exactly ρn out of the n indices, setting these indices to $\pm 1/\sqrt{\rho n}$ at random, and setting the remaining indices to 0. (We restrict to n for which ρn is an integer.)

Remark 4.7. As in [KXZ16], one can also define a similar i.i.d. prior where each entry is drawn independently from the appropriate distribution: $\pm 1/\sqrt{\rho n}$ each with probability $\rho/2$, and 0 with probability $1 - \rho$. We expect that the information theoretic limits are the same for the i.i.d. variant, but in order to apply our lower bound argument, one should condition on the spike having close-to-typical sparsity anyway. For this reason, we consider the exact-support version in order to simplify our proof.

We will see that the sparse Rademacher prior $\mathcal{X}_{\text{spRad}}(\rho)$ has rate function given by

$$f_\rho(t) = \min_{\zeta \in [\max(\rho t, 2\rho - 1), \rho]} G(\zeta) + \zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right)$$

where

$$G(\zeta) = -H(\{\zeta, \rho - \zeta, \rho - \zeta, 1 - 2\rho + \zeta\}) + 2H(\rho)$$

and where f_{Rad} is the Rademacher rate function:

$$f_{\text{Rad}}(t) = \log 2 - H\left(\frac{1+t}{2}\right).$$

Here $H(p) = -p \log p - (1-p) \log(1-p)$ is the binary entropy and $H(\{p_1, \dots, p_k\}) = -\sum_{i=1}^k p_i \log p_i$.

The intuition behind this rate function is the following. Imagine that we first randomly choose the supports of x and x' . The probability that the support of x intersects the support of x' at exactly ζn coordinates is approximately $\exp(-nG(\zeta))$. Conditioned on the supports of x and x' , $\langle x, x' \rangle$ is a (scaled) sum of i.i.d. Rademacher ± 1 random variables, causing the f_{Rad} term to appear.

Remark 4.8. One can similarly derive the rate function for any prior with i.i.d. finitely-supported entries (say with support size s), conditioned on typical proportions of entries. As in Appendix A of [BMNN16], in general one needs to optimize over $s \times s$ matrices.

The following tail bound shows that f_ρ satisfies our definition of rate function (see Remark 2.4).

Proposition 4.9. *For every $\rho \in (0, 1]$ there exists a constant $C > 0$ such that when x, x' are drawn independently from $\mathcal{X}_{\text{spRad}}(\rho)$, we have*

$$\Pr[|\langle x, x' \rangle| \geq t] \leq Cn^{3/2} \exp(-nf_\rho(t))$$

for all $t \in [0, 1]$.

The proof is deferred to Appendix C. The next tail bound improves upon the above by eliminating the polynomial factor $n^{3/2}$; however, we only prove it for sufficiently small t . This tail bound implies the local subgaussian condition with the correct constant for $d = 2$ (see Proposition 2.8).

Proposition 4.10. *For every $\rho \in (0, 1]$ there exist constants $C, T > 0$ such that when x, x' are drawn independently from $\mathcal{X}_{\text{spRad}}(\rho)$, we have*

$$\Pr[|\langle x, x' \rangle| \geq t] \leq C \exp(-nf_\rho(t)) \quad \forall t \in [0, T].$$

The proof is deferred to Appendix C. This completes the proof of Theorem 1.5 (large d behavior), with the upper bound following from Proposition 4.1 and the lower bound following from Proposition 4.4. More explicitly, Theorem 2.7 gives the following lower bound for each d .

Theorem 4.11. *Let $d \geq 2$. Consider the sparse Rademacher prior $\mathcal{X}_{\text{spRad}}(\rho)$. Let $\lambda_{\text{spRad}(\rho),d}^*$ be the largest λ^* for which*

$$\frac{(\lambda^*)^2}{2} \frac{t^d}{1+t^d} \leq f_\rho(t) \quad \forall t \in [0, 1).$$

If $\lambda < \lambda_{\text{spRad}(\rho),d}^$ then strong detection is impossible. If furthermore $d \geq 3$, then weak detection and weak recovery are also impossible.*

We next consider the limit $\rho \rightarrow 0$, in which case the lower bound above has the following asymptotics.

Proposition 4.12. *Fix $d \geq 2$. The threshold $\lambda_{\text{spRad}(\rho),d}^*$ from Theorem 4.11 behaves as*

$$\lambda_{\text{spRad}(\rho),d}^* \geq 2\sqrt{-\rho \log \rho - O(\rho)}$$

in the limit $\rho \rightarrow 0$.

We defer the proof to Appendix D. When combined with the upper bound of Proposition 4.1, this completes the proof of Theorem 1.6 ($\rho \rightarrow 0$ limit).

5 Proof of Theorem 2.7

5.1 Overview

Throughout, our asymptotic notation (e.g. $o(1), O(1)$) pertains to the limit $n \rightarrow \infty$. The constants hidden by $o(1), O(1)$ do not depend on n but may depend on the parameters d, λ, \mathcal{X} .

Let $P_n = \text{WT}_n(d, \lambda, \mathcal{X})$ be the spiked model $T = \lambda x^{\otimes d} + W$, and let $Q_n = \text{WT}_n(d)$ be the corresponding unspiked model $T = W$. Recall that x is drawn from a prior \mathcal{X} with $\|x\| = 1$. Fix a sequence $\delta = \delta(n)$ such that $\frac{1}{\sqrt{n}} \ll \delta \ll 1$; for concreteness, say $\delta = \frac{\log n}{\sqrt{n}}$. Define the ‘good’ event $\Omega(x, T)$ by $|\langle T, x^{\otimes d} \rangle - \lambda| \leq \delta$. Note that under P_n , $\langle T, x^{\otimes d} \rangle \sim \lambda + \mathcal{N}(0, 2/n)$ and so $\Omega(x, T)$ occurs with probability $1 - o(1)$.

Let \tilde{P}_n be the conditional distribution of P_n conditioned on $\Omega(x, T)$. As discussed in Section 2.1, the idea of our proof is to compute the χ^2 -divergence and show that it is small ($o(1)$ for $d \geq 3$, $O(1)$ for $d = 2$); dropping the additive -1 for now, our goal is to compute the second moment $R \triangleq \mathbb{E}_{Q_n} \left[\left(\frac{d\tilde{P}_n}{dQ_n} \right)^2 \right]$. We have

$$\begin{aligned} \frac{d\tilde{P}_n}{dQ_n}(T) &= \frac{(1 + o(1)) \mathbb{E}_{x \sim \mathcal{X}} \mathbb{1}_{\Omega(x, T)} \exp(-\frac{n}{4} \langle T - \lambda x^{\otimes d}, T - \lambda x^{\otimes d} \rangle)}{\exp(-\frac{n}{4} \langle T, T \rangle)} \\ &= (1 + o(1)) \mathbb{E}_{x \sim \mathcal{X}} \mathbb{1}_{\Omega(x, T)} \exp \left(\frac{n}{2} \langle T, \lambda x^{\otimes d} \rangle - \frac{n\lambda^2}{4} \right) \end{aligned}$$

and so

$$R \triangleq \mathbb{E}_{Q_n} \left(\frac{d\tilde{P}_n}{dQ_n} \right)^2 = (1 + o(1)) \mathbb{E}_{x, x' \sim \mathcal{X}} r(x, x')$$

where x, x' are drawn independently from \mathcal{X} and

$$r(x, x') = \mathbb{E}_{T \sim Q_n} \mathbb{1}_{\Omega(x, T)} \mathbb{1}_{\Omega(x', T)} \exp \left(\frac{n\lambda}{2} \langle T, x^{\otimes d} + x'^{\otimes d} \rangle - \frac{n\lambda^2}{2} \right). \quad (9)$$

We will see that $r(x, x')$ only depends on $\beta \triangleq \langle x, x' \rangle^d$ and so we will write $r(\beta) \triangleq r(x, x')$. For technical reasons we will break down the computation of $R = (1 + o(1)) \mathbb{E}_\beta r(\beta)$ into two parts depending on the value of $|\beta|$:

- $[\varepsilon, 1]$ interval: $R_1 = \mathbb{E}_\beta [\mathbb{1}_{|\beta| \in [\varepsilon, 1]} r(\beta)]$
- $[0, \varepsilon)$ interval: $R_2 = \mathbb{E}_\beta [\mathbb{1}_{|\beta| \in [0, \varepsilon)} r(\beta)]$

where $\varepsilon > 0$ is a small constant (depending on λ but not n) to be chosen later. Note that $\beta \in [0, 1]$ for d even and $\beta \in [-1, 1]$ for d odd. In the subsequent subsections we bound R_1 and R_2 separately. We begin with R_1 , which is the crux of the proof.

5.2 Interval $[\varepsilon, 1]$

For x, x' fixed and $T \sim Q_n$, the joint distribution of $(y, y') \triangleq (\langle T, x^{\otimes d} \rangle, \langle T, x'^{\otimes d} \rangle)$ is the bivariate Gaussian $\mathcal{N}(0, \Sigma)$ where

$$\Sigma = \frac{2}{n} \begin{pmatrix} 1 & \beta \\ \beta & 1 \end{pmatrix}$$

where, recall, $\beta = \langle x, x' \rangle^d$. This distribution has density $\frac{1}{2\pi} |\Sigma|^{-1/2} \exp(-\frac{1}{2} y^\top \Sigma^{-1} y)$ where

$$|\Sigma| = \frac{4}{n^2} (1 - \beta^2)$$

and

$$\Sigma^{-1} = \frac{n}{2} (1 - \beta^2)^{-1} \begin{pmatrix} 1 & -\beta \\ -\beta & 1 \end{pmatrix}.$$

This means we have

$$r(\beta) = \iint \exp \left(\frac{n\lambda}{2} (y + y') - \frac{n\lambda^2}{2} \right) \frac{n}{4\pi} (1 - \beta^2)^{-1/2} \exp \left(-\frac{n}{4(1 - \beta^2)} (y^2 + y'^2 - 2\beta yy') \right) dy' dy.$$

where the integral is over the square of ‘good’ y, y' values: $|y - \lambda| \leq \delta, |y' - \lambda| \leq \delta$. Currently there are problematic terms that blow up when $|\beta| = 1$. The following change of variables is a trick to eliminate this issue; an alternative approach would be to treat the $|\beta| \in [1 - \varepsilon, \varepsilon]$ interval separately, bounding $r(\beta)$ by $r(1)$. Letting $u = \frac{1}{2}(y + y')$ and $v = \frac{1}{2}(y - y')$ yields

$$= \frac{n}{2\pi} (1 - \beta^2)^{-1/2} \iint \exp \left[n\lambda u - \frac{n\lambda^2}{2} - \frac{n}{2} \left(\frac{u^2}{1 + \beta} + \frac{v^2}{1 - \beta} \right) \right] dv du.$$

Relaxing the integration bounds for v to all of \mathbb{R} , explicitly compute the integral over v :

$$\leq \sqrt{\frac{n}{2\pi(1 + \beta)}} \int \exp \left(n\lambda u - \frac{n\lambda^2}{2} - \frac{nu^2}{2(1 + \beta)} \right) du.$$

Since $|u - \lambda| \leq \delta$ we can write $u = \lambda + \Delta$ with $|\Delta| \leq \delta$:

$$= \sqrt{\frac{n}{2\pi(1 + \beta)}} \int_{-\delta}^{\delta} \exp \left(\frac{n}{2} \cdot \frac{\beta}{1 + \beta} (\lambda^2 + 2\lambda\Delta) - \frac{n\Delta^2}{2(1 + \beta)} \right) d\Delta.$$

Upper bound this expression by removing the last term inside the exponent (which is always negative) and bounding Δ by δ :

$$\leq \sqrt{\frac{n}{2\pi(1 + \beta)}} \exp \left(\frac{n}{2} \cdot \frac{\beta}{1 + \beta} (\lambda^2 \pm 2\lambda\delta) \right)$$

where \pm has the same sign as β . By computing the derivative with respect to β we see that for sufficiently large n , the above bound is increasing in β . This will allow us to focus on the case $\beta \geq 0$. We have

$$r(\beta) \leq \sqrt{\frac{n}{2\pi(1 + |\beta|)}} \exp \left(\frac{n}{2} \cdot \frac{|\beta|}{1 + |\beta|} (\lambda^2 + 2\lambda\delta) \right) \leq \sqrt{\frac{n}{2\pi}} \exp \left(\frac{n}{2} \cdot \frac{|\beta|}{1 + |\beta|} (\lambda^2 + 2\lambda\delta) \right).$$

Letting $\tilde{\lambda} = \sqrt{\lambda^2 + 2\lambda\delta} = \lambda + o(1)$ and using $\frac{1}{\sqrt{2\pi}} \leq 1$, we now have

$$r(\beta) \leq r_1(\beta) \triangleq \sqrt{n} \exp \left(\frac{n\tilde{\lambda}^2}{2} \frac{|\beta|}{1 + |\beta|} \right).$$

Using r_1 we will now bound R_1 , the contribution of the $|\beta| \in [\varepsilon, 1]$ values to the second moment R . We have

$$\begin{aligned} R_1 &= \mathbb{E}_\beta [\mathbb{1}_{|\beta| \in [\varepsilon, 1]} r(\beta)] \\ &\leq \mathbb{E}_\beta [\mathbb{1}_{|\beta| \in [\varepsilon, 1]} r_1(\beta)] \\ &= \int_0^\infty \Pr [\mathbb{1}_{|\beta| \in [\varepsilon, 1]} r_1(\beta) \geq z] dz \\ &= \int_0^\infty \Pr \left[|\beta| \in [\varepsilon, 1] \text{ and } \sqrt{n} \exp \left(\frac{n\tilde{\lambda}^2}{2} \frac{|\beta|}{1 + |\beta|} \right) \geq z \right] dz. \end{aligned}$$

The change of variables $z = n \exp \left(\frac{n\tilde{\lambda}^2}{2} \frac{t}{1+t} \right)$ yields

$$\begin{aligned} &= \int_{-1}^\infty \Pr [|\beta| \in [\varepsilon, 1] \text{ and } |\beta| \geq t] \exp \left(\frac{n\tilde{\lambda}^2}{2} \frac{t}{1+t} \right) \frac{n^{3/2} \tilde{\lambda}^2}{2(1+t)^2} dt \\ &\leq \int_{-1}^\varepsilon \Pr [|\beta| \geq \varepsilon] \exp \left(\frac{n\tilde{\lambda}^2}{2} \frac{t}{1+t} \right) \frac{n^{3/2} \tilde{\lambda}^2}{2(1+t)^2} dt + \\ &\quad \int_\varepsilon^1 \Pr [|\beta| \geq t] \exp \left(\frac{n\tilde{\lambda}^2}{2} \frac{t}{1+t} \right) \frac{n^{3/2} \tilde{\lambda}^2}{2(1+t)^2} dt. \end{aligned}$$

Recalling the definition of rate function (Definition 2.3) and explicitly evaluating the integral in the first term,

$$\begin{aligned} &\leq \sqrt{n} \exp \left(-n f_{n,\mathcal{X}}(\varepsilon^{1/d}) + \frac{n\tilde{\lambda}^2}{2} \frac{\varepsilon}{1+\varepsilon} \right) + \\ &\quad \frac{n^{3/2}\tilde{\lambda}^2}{2} \int_{\varepsilon}^1 \exp \left(-n f_{n,\mathcal{X}}(t^{1/d}) + \frac{n\tilde{\lambda}^2}{2} \frac{t}{1+t} \right) dt. \end{aligned}$$

Recall that from the theorem hypothesis, we are guaranteed $\lambda < \lambda^*$ and $f_{\mathcal{X}}(t) \geq \frac{(\lambda^*)^2}{2} \frac{t^d}{1+t^d}$ for all $t \in (0, 1)$. Combining this with the uniform convergence $f_{n,x} \geq b_{n,\mathcal{X}} \rightarrow f_{\mathcal{X}}$ (see Definition 2.3) and the fact $\tilde{\lambda} = \lambda + o(1)$ yields the following: there exists a constant $\gamma > 0$ (depending on λ, ε but not on n) such that for sufficiently large n and for all $t \in [\varepsilon, 1]$ we have $f_{n,\mathcal{X}}(t^{1/d}) \geq \frac{\tilde{\lambda}^2}{2} \frac{t}{1+t} + \gamma$. It then follows that both terms above are bounded above by $O(n^{3/2}) \exp(-n\gamma) = o(1)$ and so we have $R_1 = o(1)$.

Note that this last step cannot be extended to all $t \in [0, 1]$ because $f_{\mathcal{X}}(t) = \frac{\lambda^2}{2} \frac{t^d}{1+t^d} = 0$ at $t = 0$; for this reason the $[0, \varepsilon)$ interval requires separate consideration.

5.3 Interval $[0, \varepsilon)$

For this case, the argument above is not enough and we will need to use the local subgaussian condition to control the deviations near zero. However, we do not need to use noise conditioning here and will revert to the standard second moment as in prior work (e.g. [MRZ15, BMVX16, PWBM16]).

Starting from the definition (9) of $r(x, x')$, we drop the conditioning on Ω and use the Gaussian moment-generating function to explicitly compute the expectation over $T \sim Q_n$:

$$\begin{aligned} r(\beta) &\leq \mathbb{E}_{T \sim Q_n} \exp \left(\frac{n\lambda}{2} \langle T, x^{\otimes d} + x'^{\otimes d} \rangle - \frac{n\lambda^2}{2} \right) \\ &= \exp \left(\frac{1}{2} \cdot \frac{2}{n} \cdot \frac{n^2\lambda^2}{4} \langle x^{\otimes d} + x'^{\otimes d}, x^{\otimes d} + x'^{\otimes d} \rangle - \frac{n\lambda^2}{2} \right) \\ &= \exp \left(\frac{n\lambda^2}{2} \langle x^{\otimes d}, x'^{\otimes d} \rangle \right) \\ &= \exp \left(\frac{n\lambda^2}{2} \beta \right) \triangleq r_2(\beta). \end{aligned}$$

A helpful trick in performing the above computation is to note that $\langle T, x^{\otimes d} \rangle = \langle T', x^{\otimes d} \rangle$ where T' is the asymmetric precursor tensor with $\mathcal{N}(0, 2/n)$ entries. Note that r_2 is an increasing function of β and so $r_2(\beta) \leq r_2(|\beta|)$. We have

$$\begin{aligned} R_2 &= \mathbb{E}_{\beta} [\mathbb{1}_{|\beta| \in [0, \varepsilon)} r(\beta)] \\ &\leq \mathbb{E}_{\beta} [\mathbb{1}_{|\beta| \in [0, \varepsilon)} r_2(|\beta|)] \\ &= \int_0^\infty \Pr [\mathbb{1}_{|\beta| \in [0, \varepsilon)} r_2(|\beta|) \geq z] dz \\ &= \int_0^\infty \Pr \left[|\beta| \in [0, \varepsilon) \text{ and } \exp \left(\frac{n\lambda^2}{2} |\beta| \right) \geq z \right] dz. \end{aligned}$$

The change of variables $z = \exp\left(\frac{n\lambda^2}{2}s\right)$ yields

$$\begin{aligned}
&= \int_{-\infty}^{\infty} \Pr[|\beta| \in [0, \varepsilon) \text{ and } |\beta| \geq s] \exp\left(\frac{n\lambda^2}{2}s\right) \frac{n\lambda^2}{2} ds \\
&= \int_{-\infty}^0 \Pr[|\beta| \in [0, \varepsilon)] \exp\left(\frac{n\lambda^2}{2}s\right) \frac{n\lambda^2}{2} ds + \int_0^{\varepsilon} \Pr[|\beta| \in [s, \varepsilon)] \exp\left(\frac{n\lambda^2}{2}s\right) \frac{n\lambda^2}{2} ds \\
&\leq \int_{-\infty}^0 \exp\left(\frac{n\lambda^2}{2}s\right) \frac{n\lambda^2}{2} ds + \int_0^{\varepsilon} \Pr[|\beta| \geq s] \exp\left(\frac{n\lambda^2}{2}s\right) \frac{n\lambda^2}{2} ds \\
&= 1 + \frac{n\lambda^2}{2} \int_0^{\varepsilon} \Pr[|\langle x, x' \rangle| \geq s^{1/d}] \exp\left(\frac{n\lambda^2}{2}s\right) ds \\
&\triangleq 1 + R'_2.
\end{aligned}$$

Letting $s = t^d$,

$$R'_2 = \frac{n\lambda^2 d}{2} \int_0^{\varepsilon^{1/d}} \Pr[|\langle x, x' \rangle| \geq t] \exp\left(\frac{n\lambda^2}{2}t^d\right) t^{d-1} dt.$$

Let $\eta > 0$ be a constant, to be chosen later. Choose ε small enough so that we can apply the local subgaussian bound, obtaining

$$\leq \frac{Cn\lambda^2 d}{2} \int_0^{\varepsilon^{1/d}} \exp\left(-\frac{n}{2\sigma^2 + \eta}t^2 + \frac{n\lambda^2}{2}t^d\right) t^{d-1} dt. \quad (10)$$

First consider the $d \geq 3$ case. Here the t^2 term dominates the t^d term, and so (by choosing ε sufficiently small) we have, for some constant $c > 0$,

$$R'_2 \leq \frac{Cn\lambda^2 d}{2} \int_0^{\varepsilon^{1/d}} \exp(-nct^2) t^{d-1} dt \leq \frac{Cn\lambda^2 d}{2} \int_{t=0}^{\infty} \exp(-nct^2) t^{d-1} dt = \frac{Cn\lambda^2 d}{4} (cn)^{-d/2} \Gamma(d/2) = o(1).$$

Now consider the $d = 2$ case. Now the t^2 and t^d terms in (10) are of the same order, so it is important that we have the correct subgaussian constant σ^2 . By hypothesis (v) of the theorem, we are guaranteed that $\lambda < 1/\sigma$ and so by choosing η small enough we again have a constant $c > 0$ such that

$$R'_2 \leq \frac{Cn\lambda^2 d}{2} \int_0^{\varepsilon^{1/d}} \exp(-nct^2) t dt \leq \frac{Cn\lambda^2 d}{2} \int_{t=0}^{\infty} \exp(-nct^2) t dt = \frac{C\lambda^2 d}{4c} = O(1).$$

In conclusion, we have shown $R_2 = 1 + o(1)$ for $d \geq 3$, and $R_2 = O(1)$ for $d = 2$.

5.4 Proof of non-detection

Combining the results from the previous subsections, we have that the χ^2 -divergence

$$\chi^2(\tilde{P}_n \parallel Q_n) \triangleq R - 1 = (1 + o(1))(R_1 + R_2) - 1$$

is $o(1)$ for $d \geq 3$ and $O(1)$ for $d = 2$. As discussed in Section 2.1, this has the following consequences. If $d \geq 3$ we have $\text{TV}(\tilde{P}_n, Q_n) = o(1)$; since $\text{TV}(P_n, \tilde{P}_n) = o(1)$ this implies $\text{TV}(P_n, Q_n) = o(1)$, and so weak detection is impossible. For all $d \geq 2$ we have $\tilde{P}_n \lhd Q_n$, implying $P_n \lhd Q_n$, and so strong detection is impossible.

5.5 Proof of non-recovery

We will show that in the $d \geq 3$ case, weak recovery is impossible. In this case we have shown $\text{TV}(P_n, Q_n) = o(1)$. Therefore, in order to show that weak recovery is impossible under P_n , it is sufficient to show that weak recovery is impossible under Q_n (i.e. we need to show that weak recovery is impossible when given only knowledge of the prior \mathcal{X} and nothing else). Assume on the contrary that weak recovery is possible under Q_n ; this means there is a deterministic unit vector v such that $\Pr_{x \sim \mathcal{X}}[|\langle v, x \rangle| \geq \varepsilon] \geq \varepsilon$ for some $\varepsilon > 0$ not depending on n . We will show that this allows us to perform weak detection, contradicting the above. Given a sample T from either P_n or Q_n , consider the statistic $m = \langle v^{\otimes d}, T \rangle$ where v is defined above. When $T \sim Q_n$ we have $m \sim \mathcal{N}(0, 2/n)$, which is $o(1)$ with probability $1 - o(1)$ as $n \rightarrow \infty$. If instead $T \sim P_n$ then $m = \lambda \langle v, x \rangle^d + \langle v^{\otimes d}, W \rangle$. The second term is $o(1)$ with probability $1 - o(1)$, but the first term exceeds $\lambda \varepsilon^d = \Theta(1)$ in absolute value with probability ε . We therefore have the following distinguisher which performs weak detection: if $|m| \geq \lambda \varepsilon^d / 2$, output ' P_n '; otherwise guess randomly.

Acknowledgements

The authors are indebted to Ankur Moitra for helpful discussions that motivated the project, for feedback on a draft of this paper, and for providing guidance throughout. We thank Gérard Ben Arous, Yash Deshpande, and Jonathan Weed for helpful conversations. We thank Yash Deshpande and Thibault Lesieur for resolving an issue in the replica predictions section of the first version of this paper. We thank the anonymous reviewer for helpful comments.

References

- [ABČ13] Antonio Auffinger, Gérard Ben Arous, and Jiří Černý. Random matrices and complexity of spin glasses. *Communications on Pure and Applied Mathematics*, 66(2):165–201, 2013.
- [ACV14] Ery Arias-Castro and Nicolas Verzelen. Community detection in dense random networks. *The Annals of Statistics*, 42(3):940–969, 2014.
- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15(1):2773–2832, 2014.
- [BBP05] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, pages 1643–1697, 2005.
- [BDM⁺16] Jean Barbier, Mohamad Dia, Nicolas Macris, Florent Krzakala, Thibault Lesieur, and Lenka Zdeborová. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. *arXiv preprint arXiv:1606.04142*, 2016.
- [BGN11] Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- [BHK⁺16] Boaz Barak, Samuel B. Hopkins, Jonathan Kelner, Pravesh K. Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *arXiv preprint arXiv:1604.03084*, 2016.
- [BM11] Mohsen Bayati and Andrea Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, 2011.

[BMNN16] Jess Banks, Christopher Moore, Joe Neeman, and Praneeth Netrapalli. Information-theoretic thresholds for community detection in sparse networks. In *29th Annual Conference on Learning Theory*, pages 383–416, june 2016.

[BMV⁺17] Jess Banks, Christopher Moore, Nicolas Verzelen, Roman Vershynin, and Jiaming Xu. Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *arXiv preprint arXiv:1607.05222*, 2017.

[BMVX16] Jess Banks, Christopher Moore, Roman Vershynin, and Jiaming Xu. Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *arXiv preprint arXiv:1607.05222 (version 1)*, 2016.

[CCK06] Patrick Chareka, Ottilia Chareka, and Sarah Kennedy. Locally sub-gaussian random variables and the strong law of large numbers. *Atlantic Electronic Journal of Mathematics*, 1(1):75–81, 2006.

[CDMF09] Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability*, pages 1–47, 2009.

[COP12] Amin Coja-Oglan and Konstantinos Panagiotou. Catching the k-naesat threshold. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 899–908. ACM, 2012.

[COP13] Amin Coja-Oghlan and Konstantinos Panagiotou. Going after the k-sat threshold. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 705–714. ACM, 2013.

[COP16] Amin Coja-Oghlan and Konstantinos Panagiotou. The asymptotic k-sat threshold. *Advances in Mathematics*, 288:985–1068, 2016.

[COV13] Amin Coja-Oghlan and Dan Vilenchik. Chasing the k-colorability threshold. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 380–389. IEEE, 2013.

[COZ12] Amin Coja-Oghlan and Lenka Zdeborová. The condensation transition in random hypergraph 2-coloring. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 241–250. SIAM, 2012.

[CS92] Andrea Crisanti and H.-J. Sommers. The spherical p -spin interaction spin glass model: the statics. *Zeitschrift für Physik B Condensed Matter*, 87(3):341–354, 1992.

[DAM16] Yash Deshpande, Emmanuel Abbe, and Andrea Montanari. Asymptotic mutual information for the binary stochastic block model. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 185–189. IEEE, 2016.

[DBKP11] Olivier Duchenne, Francis Bach, In-So Kweon, and Jean Ponce. A tensor-based algorithm for high-order graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2383–2395, 2011.

[DLMF] NIST Digital Library of Mathematical Functions. <http://dlmf.nist.gov/>, Release 1.0.11 of 2016-06-08. Online companion to [OLBC10].

[DM14] Yash Deshpande and Andrea Montanari. Information-theoretically optimal sparse PCA. In *IEEE International Symposium on Information Theory*, pages 2197–2201. IEEE, 2014.

[DMM09] David L. Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.

[DMM10] David L. Donoho, Arian Maleki, and Andrea Montanari. Message passing algorithms for compressed sensing: I. motivation and construction. *IEEE Information Theory Workshop (ITW)*, pages 115–144, 2010.

[Eld16] Sam Elder. Bayesian adaptive data analysis guarantees from subgaussianity. *arXiv preprint arXiv:1611.00065*, 2016.

[FP07] Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in mathematical physics*, 272(1):185–228, 2007.

[Gar85] Ed Gardner. Spin glasses with p -spin interactions. *Nuclear Physics B*, 257:747–765, 1985.

[GSV05] Dongning Guo, Shlomo Shamai, and Sergio Verdú. Mutual information and minimum mean-square error in gaussian channels. *IEEE Transactions on Information Theory*, 51(4):1261–1282, 2005.

[GT02] Francesco Guerra and Fabio L. Toninelli. The infinite volume limit in generalized mean field disordered models. *arXiv preprint cond-mat/0208579*, 2002.

[Gue03] Francesco Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Communications in mathematical physics*, 233(1):1–12, 2003.

[HSS15] Samuel B. Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-squares proofs. In *Proceedings of The 28th Conference on Learning Theory, COLT*, pages 3–6, 2015.

[HSSS15] Samuel B. Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. *arXiv preprint arXiv:1512.02337*, 2015.

[Jan95] Svante Janson. Random regular graphs: asymptotic distributions and contiguity. *Combinatorics, Probability and Computing*, 4(04):369–405, 1995.

[JM13] Adel Javanmard and Andrea Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference*, 2(2):115–144, 2013.

[Joh01] Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, 29(2):295–327, 2001.

[KM09] Satish Babu Korada and Nicolas Macris. Exact solution of the gauge symmetric p -spin glass model on a complete graph. *Journal of Statistical Physics*, 136(2):205–230, 2009.

[KXZ16] Florent Krzakala, Jiaming Xu, and Lenka Zdeborová. Mutual information in rank-one matrix estimation. *arXiv preprint arXiv:1603.08447*, 2016.

[LC60] Lucien Le Cam. *Locally Asymptotically Normal Families of Distributions. Certain Approximations to Families of Distributions and Their Use in the Theory of Estimation and Testing Hypotheses*. Berkeley & Los Angeles, 1960.

[LKZ15] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Mmse of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 680–687. IEEE, 2015.

[LM16] Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *arXiv preprint arXiv:1611.03888*, 2016.

[LML⁺17] Thibault Lesieur, Léo Miolane, Marc Lelarge, Florent Krzakala, and Lenka Zdeborová. Statistical and computational phase transitions in spiked tensor estimation. *arXiv preprint arXiv:1701.08010*, 2017.

[MM09] Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.

[MNS15] Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 162(3-4):431–461, 2015.

[MR16] Andrea Montanari and Emile Richard. Non-negative principal component analysis: Message passing algorithms and sharp asymptotics. *IEEE Transactions on Information Theory*, 62(3):1458–1484, 2016.

[MRRW97] Michael S. O. Molloy, Hanna Robalewska, Robert W. Robinson, and Nicholas C. Wormald. 1-factorizations of random regular graphs. *Random Structures and Algorithms*, 10(3):305–321, 1997.

[MRZ15] Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. In *Advances in Neural Information Processing Systems*, pages 217–225, 2015.

[NP33] Jerzy Neyman and Egon S Pearson. IX. On the problem of the most efficient tests of statistical hypotheses. *Phil. Trans. R. Soc. Lond. A*, 231(694-706):289–337, 1933.

[OLBC10] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, editors. *NIST Handbook of Mathematical Functions*. Cambridge University Press, New York, NY, 2010. Print companion to [DLMF].

[OMH13] Alexei Onatski, Marcelo J. Moreira, and Marc Hallin. Asymptotic power of sphericity tests for high-dimensional data. *The Annals of Statistics*, 41(3):1204–1231, 2013.

[Péc06] Sandrine Péché. The largest eigenvalue of small rank perturbations of Hermitian random matrices. *Probability Theory and Related Fields*, 134(1):127–173, 2006.

[PWBM16] Amelia Perry, Alexander S. Wein, Afonso S. Bandeira, and Ankur Moitra. Optimality and sub-optimality of PCA for spiked random matrices and synchronization. *arXiv preprint arXiv:1609.05573*, 2016.

[RM14] Emile Richard and Andrea Montanari. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014.

[RW94] Robert W. Robinson and Nicholas C. Wormald. Almost all regular graphs are Hamiltonian. *Random Structures & Algorithms*, 5(2):363–374, 1994.

[Sub15] Eliran Subag. The complexity of spherical p-spin models - a second moment approach. *arXiv preprint arXiv:1504.02251*, 2015.

[Tal06a] Michel Talagrand. Free energy of the spherical mean field model. *Probability Theory and Related Fields*, 134(3):339–382, 2006.

[Tal06b] Michel Talagrand. The parisi formula. *Annals of Mathematics*, pages 221–263, 2006.

[VAC15] Nicolas Verzelen and Ery Arias-Castro. Community detection in sparse random networks. *The Annals of Applied Probability*, 25(6):3465–3510, 2015.

[Wor99] Nicholas C. Wormald. Models of random regular graphs. *London Mathematical Society Lecture Note Series*, pages 239–298, 1999.

[ZK15] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *arXiv preprint arXiv:1511.02476*, 2015.

A Asymptotics for the spherical prior

Here we will derive the asymptotic expressions given in Section 1.2 for various bounds for the spherical prior.

A.1 Lower bound $\lambda_{\text{sph},d}^*$

Consider the threshold $\lambda \triangleq \lambda_{\text{sph},d}^*$ from Theorem 3.2, i.e. λ takes on the supremal value for which

$$\frac{\lambda^2}{2} \frac{t^d}{1+t^d} \leq -\frac{1}{2} \log(1-t^2)$$

for all $t \in [0, 1)$. This inequality is not tight as $t \rightarrow 0$, considering the derivatives at 0, nor as $t \rightarrow 1$, where the right-hand side becomes infinite while the left remains bounded. Hence there exists some point $t \in (0, 1)$ where $\lambda^2 \frac{t^d}{1+t^d}$ and $-\log(1-t^2)$ agree in value and derivative:

$$\lambda^2 \frac{t^d}{1+t^d} = -\log(1-t^2), \quad d\lambda^2 t^{d-1} (1+t^d)^{-2} = \frac{t}{1-t^2}.$$

The latter equation yields $\lambda^2 = \frac{2(1+t^d)^2}{d(1-t^2)t^{d-2}}$. Substituting into the first equation, we have

$$0 = \frac{2}{d} (1+t^d) \frac{t^2}{1-t^2} + \log(1-t^2). \quad (11)$$

This equation has a unique solution on $(0, 1)$: if we divide by $\log(1-t^2)$, we have monotonicity as the derivative is negative:

$$\begin{aligned} \frac{2t(2t^2(1+t^d) + (2+t^d(2+d(1-t^2)))\log(1-t^2))}{d(1-t^2)^2 \log^2(1-t^2)} &\leq \frac{2t(2t^2(1+t^d) + (2+t^d(2+d(1-t^2)))(-t^2))}{d(1-t^2)^2 \log^2(1-t^2)} \\ &= \frac{-t^{2+d}}{(1-t^2) \log^2(1-t^2)} < 0. \end{aligned}$$

Suppose that we evaluate (11) at the value

$$t = 1 - \frac{1}{c_1 d \log d + c_2 d \log \log d + c_3 d + \varepsilon(d)},$$

with $\varepsilon(d) = o(d)$ as $d \rightarrow \infty$. We note the expansions

$$1-t^2 = \frac{2}{c_1 d \log d + c_2 d \log \log d + c_3 d + \varepsilon(d)} + O\left(\frac{1}{d^2 \log^2 d}\right), \quad 1+t^d = 2 - \frac{1}{c_1 \log d} + o\left(\frac{1}{\log d}\right).$$

Expanding (11), we obtain the equation

$$\begin{aligned} 0 &= c_1 \log d - \frac{1}{2} \log d + c_2 \log \log d - \frac{1}{2} \log \log d + c_3 - \frac{1}{2} + \frac{1}{2} \log 2 - \frac{1}{2} \log c_1 + \varepsilon(d) + O\left(\frac{\log \log d}{\log d}\right) \\ &= \varepsilon(d) + O\left(\frac{\log \log d}{\log d}\right) \quad \text{if } c_1 = c_2 = \frac{1}{2}, c_3 = \frac{1}{2} - \log 2. \end{aligned}$$

With this choice of coefficients, the right-hand side is positive for a suitably large choice of $\varepsilon(d) = o(d)$, and negative for a suitably small choice. By the intermediate value theorem, it follows that the unique solution of this equation in $(0, 1)$ satisfies the asymptotics

$$t = 1 - \frac{2}{d \log d + d \log \log d + 1 - 2 \log 2 + o(1)},$$

which in turn implies:

$$\begin{aligned} \lambda^2 &= \frac{2(1+t^d)^2}{d(1-t^2)t^{d-2}} \\ &= \left(\frac{1}{2} \log d + \frac{1}{2} \log \log d + \frac{1}{2} + o(1) \right) \left(2 - \frac{2}{\log d} + o(1/\log d) \right)^2 \left(1 + \frac{2}{\log d} + o(1/\log d) \right) \\ &= 2 \log d + 2 \log \log d + 2 - 4 \log 2 + o(1). \end{aligned}$$

A.2 Injective norm μ_d

Recall from Theorem 3.3 that the injective norm μ_d of $\text{WT}(d)$ is given by $\mu_d = x\sqrt{2/d}$, where $x \geq 2\sqrt{d-1}$ is the unique solution to

$$0 = g(z) \triangleq \frac{2-d}{d} - \log \left(\frac{dz^2}{2} \right) + \frac{d-1}{2}z^2 - \frac{2}{d^2 z^2}, \quad z(x) = \frac{1}{(d-1)\sqrt{2d}} \left(x - \sqrt{x^2 - 4(d-1)} \right).$$

We will locate this root asymptotically. Let $x = \sqrt{d \log d + d \log \log d + dc}$, with $c = O(1)$. It is easily computed that

$$z^2 = \frac{d^2}{2(\log d + \log \log d + c - 2 + o(1))},$$

so that $g(x) = 1 - c + o(1)$. Hence the unique root is located with $c = 1 + o(1)$, so that

$$\mu_d^2 = 2 \log d + 2 \log \log d + 2 + o(1).$$

A.3 Upper bound $\Lambda_{\text{sph},d}^*$

We now develop an asymptotic understanding of the upper bound $\Lambda_{\text{sph},d}^*$ provided in Theorem 3.4 and Corollary 3.5. Recall:

Theorem 3.4. *Fix $d \geq 3$ and $\lambda \geq 0$. Let \mathcal{X} be any prior supported on the unit sphere in \mathbb{R}^d . Let $T = \lambda x^{\otimes d} + W$ be a spiked tensor drawn from $\text{WT}(d, \lambda, \mathcal{X})$. With probability $1 - o(1)$ we have*

$$\|T\| \geq L_d(\lambda) - o(1)$$

where

$$L_d(\lambda) = \max_{m \in [0,1]} m^d \left(\lambda + \sqrt{\frac{2d}{d-1}} \sqrt{M(1+M)} \right), \quad M(m) = (d-1) \frac{1-m^2}{m^2}. \quad (12)$$

Corollary 3.5. *Fix $d \geq 3$ and $\lambda \geq 0$. Let \mathcal{X} be any prior supported on the unit sphere in \mathbb{R}^d . Let μ_d be given by Theorem 3.3 and let $L_d(\lambda)$ be defined as in Theorem 3.4. If $L_d(\lambda) > \mu_d$ then strong detection and weak recovery are possible for $\text{WT}(d, \lambda, \mathcal{X})$.*

Our main result will be the following:

Claim A.1. *Assuming $\lambda = \sqrt{2 \log d} + o(\sqrt{\log d})$, we have $L_d(\lambda)^2 - \lambda^2 \geq 2 + o(1)$ as $d \rightarrow \infty$.*

Recalling from the discussion above that $\mu_d^2 = 2 \log d + 2 \log \log d + 2 + o(1)$, it follows that $L_d(\lambda)$ exceeds μ_d as soon as $\lambda > 2 \log d + 2 \log \log d + o(1)$; at this point, reliable detection is possible according to Corollary 3.5. This expansion matches the replica prediction of the next section. We prove the claim:

Proof. To produce such a bound, we can choose a value of m ; thus take $m = \exp(-c/d \log d)$, with c a constant to be chosen later. Then $M = \frac{2}{c \log d} + O(d^{-1})$, and one computes that

$$\begin{aligned} L_d(\lambda)^2 - \lambda^2 &\geq m^{2d} \left(\lambda^2 + 2\sqrt{2}\lambda \sqrt{\frac{d}{d-1}} \sqrt{M(1+M)} + 2\frac{d}{d-1}M(1+M) \right) - \lambda^2 \\ &= \left(1 - \frac{2}{c \log d} + O(\log^{-2} d) \right) \left(\lambda^2 + 2\sqrt{2}\lambda \sqrt{\frac{2}{c \log d}} + \frac{4}{c \log d} + o(1) \right) - \lambda^2 \\ &= \frac{-2\lambda^2}{c \log d} + \frac{4\lambda}{\sqrt{c \log d}} + o(1) \\ &= \frac{-4}{c} + 4\sqrt{2/c}, \end{aligned}$$

which is maximized at $c = 2$ with a value of 2. \square

We conclude that the limiting (as $n \rightarrow \infty$, in probability) injective norm of a spiked tensor exceeds that of an unspiked tensor, thus guaranteeing reliable detection, as soon as $\lambda^2 > 2 \log d + 2 \log \log d + o(1)$.

B Replica predictions

In this appendix we give a non-rigorous prediction for the exact statistical threshold for spiked tensor problems (as $n \rightarrow \infty$) with any fixed d , using the replica method from statistical physics (see [MM09] for an introduction). We expect the threshold predicted here to be correct for both strong detection and weak recovery, and when $d \geq 3$, also weak detection. An equivalent formulation of the Rademacher result has been stated previously by [KM09]. More recent independent work [LML⁺17] confirms the replica prediction for various spiked tensor problems, but only for the recovery problem (not detection).

We consider both the Rademacher prior and the spherical prior. (The sparse Rademacher prior can be approached using the same techniques; see also [LKZ15] for the solution to the $d = 2$ case using different methods.) It will be more convenient to rescale the observations as

$$T = \frac{\lambda d!}{2n^{d-1}} x^{\otimes d} + W,$$

where $\|x\|^2 = n$ and the typical (distinct-index) entries of W are distributed as $\mathcal{N}(0, \frac{d!}{2n^{d-1}})$. With respect to the normalization in previous sections, this scaling preserves the meaning of the signal-to-noise ratio λ , while matching the noise scaling conventions of [Gar85] (on which the following computation will be based).

B.1 Rademacher prior

The replica calculation for this problem in the unspiked setting can be found in [Gar85]. This section adapts the computation to the spiked case. We observe the spiked tensor T as above where each x_i is uniformly ± 1 . Ignoring lower order terms (corresponding to non-distinct indices), the posterior distribution of x given T is

$$\Pr[x \mid T] \propto \prod_{i_1 < \dots < i_d} \exp \left(-\frac{n^{d-1}}{d!} (T_{i_1 \dots i_d} - \frac{\lambda d!}{2n^{d-1}} x_{i_1} \dots x_{i_d})^2 \right) \propto \exp \left(\lambda \sum_{i_1 < \dots < i_d} T_{i_1 \dots i_d} x_{i_1} \dots x_{i_d} \right)$$

and so we are interested in the Boltzmann distribution over $\sigma \in \{\pm 1\}^n$ given by $\Pr[\sigma \mid T] \propto \exp(-\beta H(\sigma))$ with energy $H(\sigma) = -\sum_{i_1 < \dots < i_d} T_{i_1 \dots i_d} \sigma_{i_1} \dots \sigma_{i_d}$ and inverse temperature $\beta = \lambda$.

The goal is to compute the free energy density, defined as $f = -\frac{1}{\beta n} \mathbb{E} \log Z$ where

$$Z = \sum_{\sigma \in \{\pm 1\}^n} \exp(-\beta H(\sigma)).$$

The idea of the replica method is to instead compute the moments $\mathbb{E}[Z^r]$ of Z for $r \in \mathbb{N}$ and perform the (non-rigorous) analytic continuation

$$\mathbb{E}[\log Z] = \lim_{r \rightarrow 0} \frac{1}{r} \log \mathbb{E}[Z^r]. \quad (13)$$

The moment $\mathbb{E}[Z^r]$ can be expanded in terms of r ‘replicas’ $\sigma^1, \dots, \sigma^r$ with $\sigma^a \in \{\pm 1\}^n$:

$$\mathbb{E}[Z^r] = \sum_{\{\sigma^a\}} \mathbb{E} \exp \left(\beta \sum_{i_1 < \dots < i_d} T_{i_1 \dots i_d} \sum_{a=1}^r \sigma_{i_1}^a \dots \sigma_{i_d}^a \right).$$

After applying the definition of T and the Gaussian moment-generating function (to compute expectation over the noise W) we arrive at

$$\mathbb{E}[Z^r] = \sum_{\{\sigma^a\}} \exp \left[n \left(\frac{\lambda^2}{2} \sum_a c_a^d + \frac{\lambda^2}{4} \sum_{a,b} q_{ab}^d \right) \right]$$

where $q_{ab} = \frac{1}{n} \sum_i \sigma_i^a \sigma_i^b$ is the correlation between replicas a and b , and $c_a = \frac{1}{n} \sum_i \sigma_i^a x_i$ is the correlation between replica a and the truth.

Without loss of generality we can assume the true spike is $x = \mathbb{1}$ (all-ones). Let Q be the $(r+1) \times (r+1)$ matrix of overlaps (q_{ab} and c_a), including x as the zeroth replica. Note that Q is the average of n i.i.d. matrices and so by the theory of large deviations (Cramér’s Theorem in multiple dimensions), the number of configurations $\{\sigma^a\}$ corresponding to given overlap parameters q_{ab}, c_a is asymptotically

$$\inf_{\mu, \nu} \exp \left[n \left(- \sum_a \nu_a c_a - \frac{1}{2} \sum_{a \neq b} \mu_{ab} q_{ab} + \log \sum_{\sigma \in \{\pm 1\}^r} \exp \left(\sum_a \nu_a \sigma_a + \frac{1}{2} \sum_{a \neq b} \mu_{ab} \sigma_a \sigma_b \right) \right) \right]. \quad (14)$$

We now apply the saddle point method: in the large n limit, the expression for $\mathbb{E}[Z^r]$ should be dominated by a single value of the overlap parameters q_{ab}, c_a . This yields

$$\frac{1}{n} \log \mathbb{E}[Z^r] = -G(q_{ab}^*, c_a^*, \mu_{ab}^*, \nu_a^*)$$

where $(q_{ab}^*, c_a^*, \mu_{ab}^*, \nu_a^*)$ is a critical point of

$$\begin{aligned} G(q_{ab}, c_a, \mu_{ab}, \nu_a) = & -\frac{\lambda^2}{2} \sum_a c_a^d - \frac{\lambda^2}{4} \sum_{a,b} q_{ab}^d \\ & + \sum_a \nu_a c_a + \frac{1}{2} \sum_{a \neq b} \mu_{ab} q_{ab} - \log \sum_{\sigma \in \{\pm 1\}^r} \exp \left(\sum_a \nu_a \sigma_a + \frac{1}{2} \sum_{a \neq b} \mu_{ab} \sigma_a \sigma_b \right). \end{aligned}$$

We next assume that the dominant saddle point takes a particular form: the so-called replica symmetric ansatz. The validity of this assumption is justified by a phenomenon in statistical physics: there is no static replica symmetry breaking on the Nishimori line (see e.g. [ZK15]). The replica symmetric ansatz is given by $q_{aa} = 1$, $c_a = c$, $\nu_a = \nu$, and for $a \neq b$, $q_{ab} = q$ and $\mu_{ab} = \mu$ for constants q, c, μ, ν . This yields

$$\lim_{r \rightarrow 0} \frac{1}{r} G(q, c, \mu, \nu) = -\frac{\lambda^2}{2} c^d - \frac{\lambda^2}{4} + \frac{\lambda^2}{4} q^d + \nu c - \frac{1}{2} \mu(q-1) - \mathbb{E}_{z \sim \mathcal{N}(0,1)} \log(2 \cosh(\nu + \sqrt{\mu} z)) \quad (15)$$

where the last term is derived by first using the Gaussian moment-generating function to obtain the expression

$$\lim_{r \rightarrow 0} \frac{1}{r} \log \mathbb{E}_z \sum_{\sigma} \exp((\nu + \sqrt{\mu}z) \sum_a \sigma_a) = \lim_{r \rightarrow 0} \frac{1}{r} \log \mathbb{E}_z (2 \cosh(\nu + \sqrt{\mu}z))^r$$

and then applying the replica trick (13).

We next find the critical points by setting the derivatives of (15) (with respect to all four variables) to zero, which yields

$$\nu = \frac{\lambda^2}{2} d c^{d-1}, \quad \mu = \frac{\lambda^2}{2} d q^{d-1}, \quad c = \mathbb{E}_z \tanh(\nu + \sqrt{\mu}z), \quad q = \mathbb{E}_z \tanh^2(\nu + \sqrt{\mu}z).$$




Figure 3: The replica symmetric solutions for the Rademacher prior, given by (16). The left panel is $d = 2$ and the right panel is $d = 3$ (which is representative of all $d \geq 3$). The solid line is dq^{d-1} as a function of μ , where q depends on μ via (16). The dotted lines are $\frac{\mu}{\lambda^2}$ for various choices of λ (with steeper lines corresponding to smaller λ). For a given λ , the intersections between the solid and dotted line are the solutions to (16). Note that $\mu = 0$ is always a solution. For $d = 2$, a nonzero solution departs continuously from zero once λ exceeds 1. For $d \geq 3$, two nonzero solutions appear once λ exceeds a particular value (which depends on d).

Recall that the replicas are drawn from the posterior distribution $\Pr[x \mid T]$ and so the truth x behaves as if it is a replica; therefore we should have $c = q$. Using the identity $\mathbb{E}_z \tanh(\gamma + \sqrt{\gamma}z) = \mathbb{E}_z \tanh^2(\gamma + \sqrt{\gamma}z)$ (see e.g. [DAM16]), we obtain the solution $c = q$ and $\nu = \mu$ where q and μ are solutions to

$$\mu = \frac{\lambda^2}{2} d q^{d-1}, \quad q = \mathbb{E}_{z \sim \mathcal{N}(0,1)} \tanh(\mu + \sqrt{\mu}z). \quad (16)$$

The free energy density of a solution to (16) is given by

$$f = \frac{1}{\beta} \lim_{r \rightarrow 0} \frac{1}{r} G(q, c, \mu, \nu) = \frac{1}{\lambda} \left[-\frac{\lambda^2}{4} (q^d + 1) + \frac{1}{2} \mu (q + 1) - \mathbb{E}_z \log(2 \cosh(\mu + \sqrt{\mu}z)) \right].$$

Figure 3 depicts the solutions to (16). Note that (16) always admits the zero solution $q = \mu = 0$. For $d = 2$, a nonzero solution appears once λ exceeds 1, recovering the eigenvalue transition of [FP07]. For $d \geq 3$, two nonzero solutions appear once λ exceeds a particular value λ_1^* (depending on d). The smaller of the two solutions is nonsensical because the overlap q decreases with λ . In Figure 4 we plot the free energy density of these solutions as a function of λ . For each λ we expect that the correct solution is the one of minimum free energy density. Therefore we conclude that the detection threshold occurs at the point $\lambda_2^* > \lambda_1^*$ at which the




Figure 4: Free energy density of the replica symmetric solutions for the Rademacher prior (in the $n \rightarrow \infty$ limit), as a function of λ . The left panel is $d = 2$ and the right panel is $d = 3$ (which is representative of all $d \geq 3$). The dotted line is the zero solution $\mu = q = 0$. For $d = 2$, the solid line is the single nonzero solution. For $d = 3$, the solid lines are the two nonzero solutions with the bottom line corresponding to the larger (in terms of μ) solution from Figure 3. For each λ , the correct solution is the one of minimum free energy. The detection threshold is the value of λ at which a nonzero solution first drops below the zero solution.




Figure 5: Overlap $q = \frac{1}{n} |\langle x, \hat{x} \rangle|$ between the true Rademacher spike x and the MMSE estimator \hat{x} (in the $n \rightarrow \infty$ limit), given by (16). The left panel is $d = 2$ and the right panel is $d = 3$ (which is representative of all $d \geq 3$). The MMSE estimator has norm $\|\hat{x}\| = q\sqrt{n}$ or else its MSE could be improved by rescaling.

free energy density of the larger nonzero solution drops below the free energy density of the zero solution. In Figure 5 we plot the overlap q between the true spike and the optimal (MMSE) estimator; this is the same q as in (16). For $d = 2$ the overlap departs continuously from zero at the threshold, whereas for $d \geq 3$ it jumps discontinuously at the threshold.

For $d = 2$, the result of this replica calculation has been rigorously shown to give the correct strong detection threshold ($\lambda = 1$) [PWBM16] as well as the correct recovery error for each λ [DAM16].

For large d , the predicted detection threshold λ_2^* approaches the constant $2\sqrt{\log 2}$, matching the bounds in Theorem 1.4. See Figure 1 for a comparison between the replica prediction and our rigorous upper and

lower bounds.

For $d \geq 3$, it can be shown rigorously that the threshold for λ given by the replica calculation is an upper bound on the true threshold, both for strong detection and weak recovery. For recovery, this follows from Theorem 1 of [KM09] combined with the I-MMSE relation [GSV05] along with the relation between free energy and mutual information (similarly to equation (10) in [KXZ16]). For detection, this follows again from Theorem 1 of [KM09], along with the fact (stated in [KM09], based on [GT02]) that the free energy concentrates and thus provides reliable detection above the threshold. One can also deduce a rigorous lower bound on the weak recovery threshold using Theorem 3 of [KM09] but it is weaker than our rigorous lower bound (and in fact, numerically appears to coincide with the lower bound of the basic second moment method without noise conditioning).

B.2 Spherical prior

The computation in this section is similar to that of the previous section, so we will skip many of the details. This time we follow the corresponding unspiked computation in [CS92]. The problem setup and normalization is the same as in the previous section except now the spike x is drawn uniformly from vectors of norm \sqrt{n} . The only difference between this calculation and the Rademacher one is that we need to change the “entropy term” (14). We have the free energy density

$$f = \frac{1}{\beta} \lim_{r \rightarrow 0} \frac{1}{r} G(q_{ab}^*, c_a^*)$$

where, borrowing the “entropy term” from equation (3.17) in [CS92], (q_{ab}^*, c_a^*) is a critical point of

$$G(q_{ab}, c_a) = -\frac{\lambda^2}{4} \sum_{a,b} q_{ab}^d - \frac{\lambda^2}{2} \sum_a c_a^d - \frac{1}{2} \log \det(Q)$$

where Q is the $(r+1) \times (r+1)$ matrix of overlaps (q_{ab} and c_a), including the true spike as the zeroth replica. Again, since we are on the Nishimori line, the replica symmetric ansatz should be valid: $q_{aa} = 1$, $q_{ab} = c_a = q$ for $a \neq b$. This yields

$$G(q) = -\frac{\lambda^2}{4} [r + r(r+1)q^d] - \frac{1}{2} \log[(1-q)^r((r+1)q + 1 - q)]$$

and so

$$\lim_{r \rightarrow 0} \frac{1}{r} G(q) = -\frac{\lambda^2}{4}(q^d + 1) - \frac{q}{2} - \frac{1}{2} \log(1 - q).$$

Solving for stationary points in q ,

$$\frac{\lambda^2}{2} dq^{d-1}(1 - q) = q. \quad (17)$$

The free energy density of a solution to (17) is given by

$$f = \frac{1}{\beta} \lim_{r \rightarrow 0} \frac{1}{r} G(q) = \frac{1}{\lambda} \left[-\frac{\lambda^2}{4}(q^d + 1) - \frac{q}{2} - \frac{1}{2} \log(1 - q) \right].$$

Qualitatively, the results are similar to the Rademacher prior of the previous section. For $d = 2$, a single nonzero solution to (17) appears once λ exceeds 1. For $d \geq 3$, two solutions appear once λ exceeds some value λ_1^* (depending on d). The smaller of these solutions is nonsensical because q decreases with λ . The larger solution becomes correct when λ exceeds $\lambda_2^* > \lambda_1^*$, at which point its free energy drops below the free energy of the zero solution. The detection threshold occurs at λ_2^* .

For the matrix case ($d = 2$), we recover the threshold $\lambda = 1$ corresponding to the eigenvalue transition, as well as the correct overlap q between the MMSE estimator and the truth (see Theorem 1.2); the MMSE estimator is (asymptotically as $n \rightarrow \infty$) given by rescaling the top eigenvector to have norm $q\sqrt{n}$.

Following the argument style of Appendix A, one computes that as $d \rightarrow \infty$, the squared predicted detection threshold behaves as

$$(\lambda_2^*)^2 = 2 \log d + 2 \log \log d + o(1),$$

implying the cruder statement $\lambda_2^* = \sqrt{2 \log d} + o(1)$. See Figure 1 for a comparison between the replica prediction and our rigorous upper and lower bounds.

C Sparse Rademacher tail bounds

Throughout this section we fix a constant $\rho \in (0, 1]$ and consider the sparse Rademacher prior $\mathcal{X}_{\text{spRad}}(\rho)$. The notation $A \lesssim B$ means $A \leq C_\rho B$ where C_ρ is a constant that may depend on ρ but not on any other variable. Similarly, $A \simeq B$ means $C_\rho A \leq B \leq C'_\rho A$. Recall from Section 4.4 the sparse Rademacher rate function, which we will denote by f :

$$f(t) = \min_{\zeta \in [\max(\rho t, 2\rho - 1), \rho]} G(\zeta) + \zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right) \quad (18)$$

where

$$G(\zeta) = -H(\{\zeta, p - \zeta, p - \zeta, 1 - 2\rho + \zeta\}) + 2H(\rho)$$

and where f_{Rad} is the Rademacher rate function:

$$f_{\text{Rad}}(t) = \log 2 - H\left(\frac{1+t}{2}\right).$$

In this section (as discussed in Section 4.4) we prove the necessary tail bounds to show that f indeed satisfies our definition of rate function.

Proposition 4.9. *For every $\rho \in (0, 1]$ there exists a constant $C > 0$ such that when x, x' are drawn independently from $\mathcal{X}_{\text{spRad}}(\rho)$, we have for all $t \in [0, 1]$,*

$$\Pr[|\langle x, x' \rangle| \geq t] \leq Cn^{3/2} \exp(-nf(t))$$

where f is the rate function (18).

Proof. By symmetry it is sufficient to prove a one-sided tail bound on $\Pr[\langle x, x' \rangle \geq t]$. Let z be the number of indices on which the supports of x and x' overlap, and let $\zeta = z/n$. Note that the set of possible ζ values is precisely the interval $[\max(0, 2\rho - 1), \rho]$. In order to have $\Pr[\langle x, x' \rangle \geq t]$ we also need $\zeta \geq \rho t$, which restricts us to the interval $\mathcal{I} \triangleq [\max(2\rho - 1, \rho t), \rho]$ appearing in (18). Using the Rademacher Chernoff bound we can write

$$\Pr[\langle x, x' \rangle \geq t] \leq \sum_z \Pr[z] \exp\left(-n\zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right)\right). \quad (19)$$

where z ranges over all integer values for which $\zeta = z/n \in \mathcal{I}$. Note that z follows the hypergeometric distribution $z \sim \text{Hyp}(n, \rho n, \rho n)$ given by

$$\Pr[z] = \frac{\binom{\rho n}{z} \binom{(1-\rho)n}{\rho n - z}}{\binom{n}{\rho n}}.$$

By Stirling's approximation we have $n! \simeq n^{n+\frac{1}{2}} e^{-n}$ and so $\binom{n}{k} \simeq \sqrt{\frac{n}{k(n-k)}} \exp(nH(k/n))$. This means

$$\begin{aligned} \Pr[z] &\lesssim \sqrt{n} \exp\left\{n \left[\rho H\left(\frac{\zeta}{\rho}\right) + (1-\rho)H\left(\frac{\rho-\zeta}{1-\rho}\right) - H(\rho)\right]\right\} \\ &= \sqrt{n} \exp\{n [H(\{\zeta, p - \zeta, p - \zeta, 1 - 2\rho + \zeta\}) - 2H(\rho)]\} \\ &= \sqrt{n} \exp(-nG(\zeta)). \end{aligned}$$

The result now follows by bounding the sum in (19) by n times its largest term. \square

Proposition 4.10. *For every $\rho \in (0, 1]$ there exist constants $C, T > 0$ such that when x, x' are drawn independently from $\mathcal{X}_{\text{spRad}}(\rho)$, we have*

$$\Pr[|\langle x, x' \rangle| \geq t] \leq C \exp(-nf(t)) \quad \forall t \in [0, T].$$

Proof. For small t , the last term $\zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right)$ in the definition of f becomes negligible (converges to zero uniformly in $\zeta \in [\rho t, \rho]$ as $t \rightarrow 0$ by Dini's theorem) and so $f(t)$ becomes approximately the minimum value of $G(\zeta)$ over the valid range $\zeta \in [\max(\rho t, 2\rho - 1), \rho]$. The minimum of $G(\zeta)$ occurs at $\zeta = \rho^2$ and furthermore $G(\zeta)$ is strongly convex. Let $\zeta^* = \zeta^*(t)$ be the minimizer in the expression for the rate function (18) so that $f(t) = G(\zeta^*) + \zeta^* f_{\text{Rad}}\left(\frac{\rho t}{\zeta^*}\right)$. The above implies $\zeta^* \rightarrow \rho^2$ as $t \rightarrow 0$. We also have $\frac{\partial^2}{\partial \zeta^2} \left[\zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right) \right] = \frac{\rho^2 t^2}{\zeta^3 - \rho^2 t^2 \zeta}$, which converges uniformly (in ζ) to zero as $t \rightarrow 0$ in a neighborhood (for ζ) containing ρ^2 .

It follows from the above that we can choose T small enough so that there exists $\delta > 0$ (depending on ρ but not t) so that

$$G(\zeta) + \zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right) \geq f(t) + \delta(\zeta - \zeta^*)^2 \quad \forall t \in [0, T], \zeta \in \mathcal{I}.$$

Let ε be a small constant (depending on ρ but not t) to be chosen later. We first consider the sum (19) over the values $\zeta \in \mathcal{I}' \triangleq [\max(0, 2\rho - 1) + \varepsilon, \rho - \varepsilon]$, i.e. discarding a small interval at each endpoint of \mathcal{I} . For $\zeta \in \mathcal{I}'$, Stirling's approximation yields (see the proof of Proposition 4.9 above)

$$\Pr[z] \simeq \frac{1}{\sqrt{n}} \exp(-nG(\zeta)).$$

Now, starting from (19), we have

$$\begin{aligned} \sum_{z: \zeta \in \mathcal{I}'} \Pr[z] \exp\left(-n\zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right)\right) &\lesssim \frac{1}{\sqrt{n}} \sum_{z: \zeta \in \mathcal{I}'} \exp\left\{-n\left[G(\zeta) + \zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right)\right]\right\} \\ &\leq \frac{1}{\sqrt{n}} \sum_z \exp(-nf(t) - n\delta(\zeta - \zeta^*)^2) \\ &= \exp(-nf(t)) \cdot \frac{1}{\sqrt{n}} \sum_z \exp(-n\delta(\zeta - \zeta^*)^2) \\ &\lesssim \exp(-nf(t)) \end{aligned}$$

where the last step follows from (the proof of) Lemma 10 in [BMNN16].

It remains to bound the terms in (19) for which ζ lies outside \mathcal{I}' . First consider $\zeta \in [0, \varepsilon]$ (in the case $\rho < \frac{1}{2}$). Here, instead of applying Stirling's approximation to $\binom{\rho n}{z}$, we apply the simpler bound $\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$ to yield $\binom{\rho n}{z} \leq \binom{\rho n}{\varepsilon n} \leq \left(\frac{\rho e}{\varepsilon}\right)^{\varepsilon n}$. As above, we apply Stirling's approximation to the other two binomial coefficients in $\Pr[z]$, yielding

$$\frac{\binom{(1-\rho)n}{\rho n-z}}{\binom{n}{\rho n}} \simeq \exp\left\{n\left[(1-\rho)H\left(\frac{\rho-\zeta}{1-\rho}\right) - H(\rho)\right]\right\} \leq \exp(-nG(\zeta)).$$

Also, since $\zeta \in [0, \varepsilon]$ is far from the minimizer $\zeta^* \approx \rho^2$ (for small t), there exists a constant Δ (depending on ρ) such that

$$G(\zeta) + \zeta f_{\text{Rad}}\left(\frac{\rho t}{\zeta}\right) \geq f(t) + \Delta \quad \forall t \in [0, T], \zeta \in [0, \varepsilon].$$

This yields

$$\begin{aligned}
\sum_{z: \zeta \in [0, \varepsilon]} \Pr[z] \exp \left(-n\zeta f_{\text{Rad}} \left(\frac{\rho t}{\zeta} \right) \right) &\lesssim \sum_{z: \zeta \in [0, \varepsilon]} \left(\frac{\rho e}{\varepsilon} \right)^{\varepsilon n} \exp \left\{ -n \left[G(\zeta) + \zeta f_{\text{Rad}} \left(\frac{\rho t}{\zeta} \right) \right] \right\} \\
&\leq \sum_{z: \zeta \in [0, \varepsilon]} \left(\frac{\rho e}{\varepsilon} \right)^{\varepsilon n} \exp(-nf(t) - n\Delta) \\
&\leq \exp(-nf(t)) \cdot \varepsilon n \left[\left(\frac{\rho e}{\varepsilon} \right)^{\varepsilon} e^{-\Delta} \right]^n,
\end{aligned}$$

which is $o(1) \exp(-nf(t))$ provided we choose ε small enough so that the term in brackets $[\dots]$ is less than 1; this is possible because $\lim_{\varepsilon \rightarrow 0^+} \left(\frac{\rho e}{\varepsilon} \right)^{\varepsilon} = 1$.

A similar argument applies to the other boundary cases: $\zeta \in [2\rho - 1, 2\rho - 1 + \varepsilon]$ (for $\rho > \frac{1}{2}$) and $\zeta \in [\rho - \varepsilon, \rho]$. \square

D Sparse Rademacher asymptotics

In this section we analyze the $\rho \rightarrow 0$ asymptotics for our lower bound on the sparse Rademacher threshold.

Proposition 4.12. *Fix any $d \geq 2$. Let $f_\rho(t) = f_{\mathcal{X}_{\text{spRad}}(\rho)}$, the sparse Rademacher rate function. Let $\lambda^* = \lambda^*(\rho)$ be our lower bound on the detection threshold for the sparse Rademacher prior, namely*

$$(\lambda^*)^2 = 2 \inf_{t \in (0,1)} \frac{1+t^d}{t^d} f_\rho(t).$$

In the limit $\rho \rightarrow 0$, we have

$$\lambda^* \geq 2\sqrt{-\rho \log \rho - O(\rho)}.$$

Proof. Since $x \mapsto \frac{1+x}{x}$ is a decreasing function for $x > 0$, it is sufficient to prove the statement for $d = 2$. Recall the definition of f_ρ (for $\rho < \frac{1}{3}$):

$$f_\rho(t) = \min_{\zeta \in [\rho t, \rho]} G(\zeta) + \zeta f_{\text{Rad}} \left(\frac{\rho t}{\zeta} \right)$$

where

$$G(\zeta) = -H(\{\zeta, p - \zeta, p - \zeta, 1 - 2\rho + \zeta\}) + 2H(\rho).$$

We can bound the Rademacher rate function f_{Rad} by using the bound $H(a) \leq \log 2 - 2(a - \frac{1}{2})^2$ as follows:

$$f_{\text{Rad}}(t) \triangleq \log 2 - H \left(\frac{1+t}{2} \right) \geq \frac{t^2}{2}.$$

Our lower bound shows non-detection provided that

$$\lambda^2 < 2 \inf_{t \in (0,1)} \min_{\zeta \in [\rho t, \rho]} \frac{1+t^2}{t^2} \left(G(\zeta) + \frac{\rho^2 t^2}{2\zeta} \right) = 2 \inf_{\zeta \in (0, \rho]} \inf_{t \in (0, \zeta/\rho]} (1+t^2) \left(\frac{G(\zeta)}{t^2} + \frac{\rho^2}{2\zeta} \right),$$

or equivalently,

$$\lambda^2 < 2 \inf_{t \in (0, \zeta/\rho]} (1+t^2) \left(\frac{G(\zeta)}{t^2} + \frac{\rho^2}{2\zeta} \right) \quad \forall \zeta \in (0, \rho]. \quad (20)$$

We perform the optimization over t in closed form. Since $G(\zeta) \geq 0$, the function $t \mapsto (1+t^2) \left(\frac{G(\zeta)}{t^2} + \frac{\rho^2}{2\zeta} \right)$ is convex on $t \in (0, \infty)$ with its minimum at $t^2 = \frac{1}{\rho} \sqrt{2\zeta G(\zeta)}$. We will break into two cases depending on whether this minimum occurs before or after the maximum allowable value ζ/ρ for t .

First consider the ζ values satisfying

$$\frac{1}{\rho} \sqrt{2\zeta G(\zeta)} \geq \frac{\zeta^2}{\rho^2} \quad (21)$$

so that the infimum over t in (20) is optimized at $t = \zeta/\rho$. In this case we require

$$\lambda^2 < 2(1 + \zeta^2/\rho^2) \left(\frac{G(\zeta)\rho^2}{\zeta^2} + \frac{\rho^2}{2\zeta} \right) = 2 \left(\frac{\rho^2 + \zeta^2}{\zeta^2} \right) \left(G(\zeta) + \frac{\zeta}{2} \right),$$

or equivalently,

$$G(\zeta) + \frac{\zeta}{2} > \frac{\lambda^2}{2} \cdot \frac{\zeta^2}{\rho^2 + \zeta^2}.$$

If $\zeta < \frac{\rho^2}{\lambda^2}$ then this is satisfied because $G(\zeta) \geq 0$, so we can assume $\zeta \geq \frac{\rho^2}{\lambda^2}$. We will use the bound

$$G(\zeta) \geq \zeta \log \left(\frac{\zeta(1-\rho)}{ep^2} \right). \quad (22)$$

The motivation for the bound (22) comes from the argument in [BMVX16] where they compare the hypergeometric random variable to a binomial. One can prove (22) directly by showing that $G(\zeta) - \zeta \log \left(\frac{\zeta(1-\rho)}{ep^2} \right)$ is jointly convex in ζ, ρ (since the Hessian has positive trace and determinant) and has a stationary point in the limit $\zeta \rightarrow 0, \rho \rightarrow 0$ (at which point the bound holds).

Combining (22) with $\frac{\rho^2}{\lambda^2} \leq \zeta \leq \rho$ and the fact that $\zeta \mapsto \frac{\zeta}{\rho^2 + \zeta^2}$ is increasing for all $\zeta \leq \rho$, it is sufficient to show

$$\log \left(\frac{1-\rho}{e\lambda^2} \right) + \frac{1}{2} > \frac{\lambda^2}{4\rho}.$$

Rearranging this yields

$$\frac{\lambda^2}{4\rho} \exp \left(\frac{\lambda^2}{4\rho} \right) < \frac{1-\rho}{4\rho\sqrt{e}},$$

or equivalently $\frac{\lambda^2}{4\rho} < \mathcal{W} \left(\frac{1-\rho}{4\rho\sqrt{e}} \right)$ where $\mathcal{W}(y)$ (for $y \geq 0$) denotes the (unique) solution to $xe^x = y$. For $y > e$ we have $\mathcal{W}(y) < \log y$ and so for sufficiently small ρ , we are done provided that $\lambda < 2\sqrt{-\rho \log \rho - O(\rho)}$ as desired.

It remains to consider the other case: the ζ values for which (21) fails. In this case the infimum over t in (20) is optimized at $t^2 = \frac{1}{\rho} \sqrt{2\zeta G(\zeta)}$ and (20) becomes

$$\frac{\lambda^2}{2} < \left(\sqrt{G(\zeta)} + \frac{\rho}{\sqrt{2\zeta}} \right)^2$$

and so it is sufficient to show $\frac{\rho^2}{2\zeta} > -2\rho \log \rho$. Assume on the contrary that $\zeta \geq \frac{\rho}{4\log(1/\rho)}$. Combining the bound (22) with the fact that (21) fails yields

$$\frac{\zeta^2}{\rho^2} > 2 \log \left(\frac{\zeta(1-\rho)}{ep^2} \right).$$

Plugging in $\zeta \geq \frac{\rho}{4\log(1/\rho)}$ gives

$$1 \geq \frac{\zeta^2}{\rho^2} > 2 \log \left(\frac{\zeta(1-\rho)}{ep^2} \right) \geq 2 \log \left(\frac{(1-\rho)}{4e\rho \log(1/\rho)} \right) = 2 \log(1/\rho) - 2 \log \log(1/\rho) + O(1) = \Theta(\log(1/\rho))$$

as $\rho \rightarrow 0$, which is a contradiction for sufficiently small ρ . This completes the proof. \square

E Proof of MAP upper bound

In this appendix we prove the following upper bound:

Proposition 4.2. *Fix $d \geq 2$ and let \mathcal{X} be a prior on the unit sphere in \mathbb{R}^n , with bounded Rényi entropy density: there exists $\delta > 0$ with*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} H_{1-\delta}(\mathcal{X}) < \infty.$$

Suppose furthermore that the varentropy $V(\mathcal{X}) = \text{Var}_{x \sim \mathcal{X}}[-\log \Pr_{\mathcal{X}}(x)]$ is $o(n^2)$. Let $s = \limsup_{n \rightarrow \infty} \frac{1}{n} H_1(\mathcal{X})$, the Shannon entropy density. Then when $\lambda > 2\sqrt{s}$, there exists a hypothesis test distinguishing $\text{WT}(d, \lambda, \mathcal{X})$ from $\text{WT}(d)$ with probability $o(1)$ of error (i.e. strong detection is possible) and furthermore there exists an estimator that achieves nontrivial correlation with the spike (i.e. weak recovery is possible).

Proof. Given a tensor T , consider the statistic

$$m = \sup_{v \in \text{supp } \mathcal{X}} \langle T, v^{\otimes d} \rangle + \frac{2}{n\lambda} \log \Pr_{\mathcal{X}}(v). \quad (23)$$

We will show that thresholding m suitably yields a hypothesis test that distinguishes the spiked and unspiked models with $o(1)$ probability of error of either type.

Suppose first that T is drawn from the spiked model $\text{WT}(d, \lambda, \mathcal{X})$ with spike x . Then

$$\mathbb{E}m \geq \mathbb{E} \left[\langle T, x^{\otimes d} \rangle + \frac{2}{n\lambda} \log \Pr_{\mathcal{X}}(x) \right] = \lambda - \frac{2s}{\lambda},$$

and moreover, $\langle T, x^{\otimes d} \rangle + \frac{2}{n\lambda} \log \Pr_{\mathcal{X}}(x)$ concentrates to within $o(1)$ with high probability, using a Gaussian tail bound on the noise and Chebyshev on the entropy term, using the varentropy assumption. Thus for any fixed $\varepsilon > 0$, $m \geq \lambda - 2s/\lambda - \varepsilon$ with probability $1 - o(1)$.

On the other hand, suppose that T is drawn from the unspiked model $\text{WT}(d)$. Then

$$\begin{aligned} \Pr \left[m \geq \lambda - \frac{2s}{\lambda} - 2\varepsilon \right] &\leq \sum_v \Pr \left[\langle T, v^{\otimes d} \rangle \geq \lambda - 2\varepsilon - \frac{2s}{\lambda} - \frac{2}{n\lambda} \log \Pr_{\mathcal{X}}(v) \right] \\ &= \sum_v \Pr \left[\mathcal{N}(0, 2/n) \geq \lambda - 2\varepsilon - \frac{2s}{\lambda} - \frac{2}{n\lambda} \log \Pr_{\mathcal{X}}(v) \right] \\ &\leq \sum_v \exp \left(-\frac{n}{4} \left(\lambda - 2\varepsilon - \frac{2s}{\lambda} - \frac{2}{n\lambda} \log \Pr_{\mathcal{X}}(v) \right)^2 \right) \\ &= \exp \left(-\frac{n}{4} \left((\lambda - 2\varepsilon)^2 - 4s \frac{\lambda - 2\varepsilon}{\lambda} \right) \right) \sum_v \exp \left(\frac{\lambda - 2\varepsilon}{\lambda} \log \Pr_{\mathcal{X}}(v) - \frac{1}{n\lambda^2} (ns - \log \Pr_{\mathcal{X}}(v))^2 \right) \\ &\leq \exp \left(-\frac{n}{4} \left((\lambda - 2\varepsilon)^2 - 4s \frac{\lambda - 2\varepsilon}{\lambda} \right) \right) \sum_v \Pr_{\mathcal{X}}(v)^{1-2\varepsilon/\lambda} \\ &= \frac{2\varepsilon}{\lambda} \exp \left(-\frac{n}{4} \left((\lambda - 2\varepsilon)^2 - 4s \frac{\lambda - 2\varepsilon}{\lambda} - \frac{2\varepsilon}{n\lambda} H_{1-2\varepsilon/\lambda}(\mathcal{X}) \right) \right). \end{aligned}$$

As H_α is monotonically decreasing in α , we have that $\frac{1}{n} H_{1-2\varepsilon/\lambda}(\mathcal{X}) \leq \frac{1}{n} H_{1-\delta}(\mathcal{X})$ is bounded for sufficiently small ε as $n \rightarrow \infty$; thus as $\varepsilon \rightarrow 0$, we have that

$$-\frac{2\varepsilon}{n\lambda} H_{1-2\varepsilon/\lambda}(\mathcal{X}) \rightarrow 0$$

uniformly in n . Thus, so long as $\lambda > 2\sqrt{s}$, we can choose $\varepsilon > 0$ such that this probability of error is $o(1)$. Hence, thresholding the statistic m at $\lambda - \frac{2s}{\lambda} - \varepsilon$, we obtain a hypothesis test with $o(1)$ probability of error.

In order to perform weak recovery given $T = \lambda x^{\otimes d} + W$, output the v that maximizes (23). From the analysis above we have with probability $1 - o(1)$, $\langle T, v^{\otimes d} \rangle + \frac{2}{n\lambda} \log \Pr(v) \geq \lambda - \frac{2s}{\lambda} - \varepsilon$ and $\langle W, v^{\otimes d} \rangle + \frac{2}{n\lambda} \log \Pr(v) \leq \lambda - \frac{2s}{\lambda} - 2\varepsilon$. This means we have

$$\lambda - \frac{2s}{\lambda} - \varepsilon \leq \langle T, v^{\otimes d} \rangle + \frac{2}{n\lambda} \log \Pr(v) = \lambda \langle x, v \rangle^d + \langle W, v^{\otimes d} \rangle + \frac{2}{n\lambda} \log \Pr(v) \leq \lambda \langle x, v \rangle^d + \lambda - \frac{2s}{\lambda} - 2\varepsilon$$

and so $\langle x, v \rangle^d \geq \varepsilon/\lambda$, implying weak recovery. \square