

The Kikuchi Hierarchy and Tensor PCA

Alexander S. Wein^{*1}, Ahmed El Alaoui^{‡2}, and Cristopher Moore^{§3}

¹Department of Mathematics, Courant Institute of Mathematical Sciences, NYU

²Departments of Electrical Engineering and Statistics, Stanford University

³Santa Fe Institute

June 24, 2020

Abstract

For the tensor PCA (principal component analysis) problem, we propose a new hierarchy of increasingly powerful algorithms with increasing runtime. Our hierarchy is analogous to the sum-of-squares (SOS) hierarchy but is instead inspired by statistical physics and related algorithms such as belief propagation and AMP (approximate message passing). Our level- ℓ algorithm can be thought of as a linearized message-passing algorithm that keeps track of ℓ -wise dependencies among the hidden variables. Specifically, our algorithms are spectral methods based on the *Kikuchi Hessian*, which generalizes the well-studied Bethe Hessian to the higher-order Kikuchi free energies.

It is known that AMP, the flagship algorithm of statistical physics, has substantially worse performance than SOS for tensor PCA. In this work we ‘redeem’ the statistical physics approach by showing that our hierarchy gives a polynomial-time algorithm matching the performance of SOS. Our hierarchy also yields a continuum of subexponential-time algorithms, and we prove that these achieve the same (conjecturally optimal) tradeoff between runtime and statistical power as SOS. Our proofs are much simpler than prior work, and also apply to the related problem of refuting random k -XOR formulas. The results we present here apply to tensor PCA for tensors of all orders, and to k -XOR when k is even.

Our methods suggest a new avenue for systematically obtaining optimal algorithms for Bayesian inference problems, and our results constitute a step toward unifying the statistical physics and sum-of-squares approaches to algorithm design.

^{*}Email: awein@cims.nyu.edu. Partially supported by NSF grant DMS-1712730 and by the Simons Collaboration on Algorithms and Geometry.

[‡]Email: elalaoui@stanford.edu. Partially supported by IIS-1741162, and ONR N00014-18-1-2729.

[§]Email: moore@santafe.edu. Partially supported by NSF grant BIGDATA-1838251.

1 Introduction

High-dimensional Bayesian inference problems are widely studied, including planted clique [Jer92, AKS98], sparse PCA [JL04], and community detection [DKMZ11b, DKMZ11a], just to name a few. For these types of problems, two general strategies, or *meta-algorithms*, have emerged. The first is rooted in statistical physics and includes the belief propagation (BP) algorithm [Pea86, YFW03] along with variants such as approximate message passing (AMP) [DMM09], and related spectral methods such as linearized BP [KMM⁺13, BLM15], and the Bethe Hessian [SKZ14]. The second meta-algorithm is the *sum-of-squares (SOS) hierarchy* [Sho87, Par00, Las01], a hierarchy of increasingly powerful semidefinite programming relaxations to polynomial optimization problems, along with spectral methods inspired by it [HSS15, HSSS16]. Both of these meta-algorithms are known to achieve statistically-optimal performance for many problems. Furthermore, when they fail to perform a task, this is often seen as evidence that no polynomial-time algorithm can succeed. Such reasoning takes the form of free energy barriers in statistical physics [LKZ15a, LKZ15b] or SOS lower bounds (e.g., [BHK⁺16]). Thus, we generally expect both meta-algorithms to have optimal statistical performance among all *efficient* algorithms.

A fundamental question is whether we can unify statistical physics and SOS, showing that the two approaches yield, or at least predict, the same performance on a large class of problems. However, one barrier to this comes from the *tensor principal component analysis* (PCA) problem [RM14], on which the two meta-algorithms seem to have very different performance. For an integer $p \geq 2$, in the order- p tensor PCA or *spiked tensor* problem we observe a p -fold $n \times n \times \cdots \times n$ tensor

$$\mathbf{Y} = \lambda x_*^{\otimes p} + \mathbf{G}$$

where the parameter $\lambda \geq 0$ is a signal-to-noise ratio (SNR), $x_* \in \mathbb{R}^n$ is a planted signal with normalization $\|x_*\| = \sqrt{n}$ drawn from a simple prior such as the uniform distribution on $\{\pm 1\}^n$, and $\mathbf{G}$ is a symmetric noise tensor with $\mathcal{N}(0, 1)$ entries. Information-theoretically, it is possible to recover x_* given $\mathbf{Y}$ (in the limit $n \rightarrow \infty$, with p fixed) when $\lambda \gg n^{(1-p)/2}$ [RM14, LML⁺17]. (Here we ignore log factors, so $A \gg B$ can be understood to mean $A \geq B \text{ polylog}(n)$.) However, this information-theoretic threshold corresponds to exhaustive search. We would also like to understand the *computational* threshold, i.e., for what values of λ there is an efficient algorithm.

The sum-of-squares hierarchy gives a polynomial-time algorithm to recover x_* when $\lambda \gg n^{-p/4}$ [HSS15], and SOS lower bounds suggest that no polynomial-time algorithm can do better [HSS15, HKP⁺17]. However, AMP, the flagship algorithm of statistical physics, is suboptimal for $p \geq 3$ and fails unless $\lambda \gg n^{-1/2}$ [RM14]. Various other “local” algorithms such as the tensor power method, Langevin dynamics, and gradient descent also fail below this “local” threshold $\lambda \sim n^{-1/2}$ [RM14, AGJ18]. This casts serious doubts on the optimality of the statistical physics approach.

In this paper we resolve this discrepancy and “redeem” the statistical physics approach. The *Bethe free energy* associated with AMP is merely the first level of a hierarchy of *Kikuchi free energies* [Kik51, Kik94, YFW03]. From these Kikuchi free energies, we derive a hierarchy of increasingly powerful algorithms for tensor PCA, similar in spirit to generalized belief propagation [YFW03]. Roughly speaking, our level- ℓ algorithm can be thought of as an iterative message-passing algorithm that reasons about ℓ -wise dependencies among the hidden variables. As a result, it has time and space complexity $n^{O(\ell)}$. Specifically, the level- ℓ algorithm is a spectral method on a $n^{O(\ell)} \times n^{O(\ell)}$ submatrix of (a first-order approximation of) the *Kikuchi Hessian*, i.e., the matrix of second derivatives of the Kikuchi free energy. This generalizes the *Bethe Hessian* spectral method, which has been successful in the setting of community detection [SKZ14]. We note that the Ph.D. dissertation of Saade [Saa16] proposed the Kikuchi Hessian as a direction for future research.

For order- p tensor PCA with p even, we show that level $\ell = p/2$ of the Kikuchi hierarchy gives an algorithm that succeeds down to the SOS threshold $\lambda \sim n^{-p/4}$, closing the gap between SOS and statistical physics. Furthermore, by taking $\ell = n^\delta$ levels for various values of $\delta \in (0, 1)$, we obtain a continuum of subexponential-time algorithms that achieve a precise tradeoff between runtime and the signal-to-noise ratio—exactly the same tradeoff curve that SOS is known to achieve [BGG⁺16, BGL16].¹ We obtain similar results when p is odd, by combining a matrix related to the Kikuchi Hessian with a construction similar to [CGL04]; see Appendix F.2.

Our approach also applies to the problem of refuting random k -XOR formulas when k is even, showing that we can strongly refute random formulas with n variables and $m \gg n^{k/2}$ clauses in polynomial time, and with a continuum of subexponential-time algorithms that succeed at lower densities. This gives a much simpler proof of the results of [RRS17], using only the matrix Chernoff bound instead of intensive moment calculations; see Appendix F.1. We leave for future work the problem of giving a similar simplification of [RRS17] when k is odd.

Our results redeem the statistical physics approach to algorithm design and give hope that the Kikuchi hierarchy provides a systematic way to derive optimal algorithms for a large class of Bayesian inference problems. We see this as a step toward unifying the statistical physics and SOS approaches. Indeed, we propose the following informal meta-conjecture: for high-dimensional inference problems with planted solutions (and related problems such as refuting random constraint satisfaction problems) the SOS hierarchy and the Kikuchi hierarchy both achieve the optimal tradeoff between runtime and statistical power.

After the initial appearance of this paper, some related independent work has appeared. A hierarchy of algorithms similar to ours is proposed by [Has19], but with a different motivation based on a system of quantum particles. Also, [BCR19] gives an alternative “redemption” of local algorithms based on *replicated gradient descent*.

2 Preliminaries and Prior Work

2.1 Notation

Our asymptotic notation (e.g., $O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, $\omega(\cdot)$) pertains to the limit $n \rightarrow \infty$ (large dimension) and may hide constants depending on p (tensor order), which we think of as fixed. We say an event occurs *with high probability* if it occurs with probability $1 - o(1)$.

A tensor $\mathbf{T} \in (\mathbb{R}^n)^{\otimes p}$ is an $n \times n \times \cdots \times n$ (p times) multi-array with entries denoted by $\mathbf{T}_{i_1, \dots, i_p}$, where $i_k \in [n] := \{1, 2, \dots, n\}$. We call p the *order* of $\mathbf{T}$ and n the *dimension* of $\mathbf{T}$. For a vector $u \in \mathbb{R}^n$, the rank-1 tensor $u^{\otimes p}$ is defined by $(u^{\otimes p})_{i_1, \dots, i_p} = \prod_{k=1}^p u_{i_k}$. A tensor $\mathbf{T}$ is *symmetric* if $\mathbf{T}_{i_1, \dots, i_p} = \mathbf{T}_{i_{\pi(1)}, \dots, i_{\pi(p)}}$ for any permutation $\pi \in S_p$. For a symmetric tensor, if $E = \{i_1, \dots, i_p\} \subseteq [n]$, we will often write $\mathbf{T}_E := \mathbf{T}_{i_1, \dots, i_p}$.

2.2 The Spiked Tensor Model

A general formulation of the spiked tensor model is as follows. For an integer $p \geq 2$, let $\tilde{\mathbf{G}} \in (\mathbb{R}^n)^{\otimes p}$ be an asymmetric tensor with entries i.i.d. $\mathcal{N}(0, 1)$. We then symmetrize this tensor and obtain a symmetric tensor $\mathbf{G}$,

$$\mathbf{G} := \frac{1}{\sqrt{p!}} \sum_{\pi \in S_p} \tilde{\mathbf{G}}^\pi,$$

¹The strongest SOS results only apply to a variant of the spiked tensor model with Rademacher observations, but we do not expect this difference to be important; see Section 2.6.

where S_p is the symmetric group of permutations of $[p]$, and $\tilde{\mathbf{G}}_{i_1, \dots, i_p}^\pi := \tilde{\mathbf{G}}_{i_{\pi(1)}, \dots, i_{\pi(p)}}$. Note that if $i_1, \dots, i_p$ are distinct then $\mathbf{G}_{i_1, \dots, i_p} \sim \mathcal{N}(0, 1)$. We draw a ‘signal’ vector or ‘spike’ $x_* \in \mathbb{R}^n$ from a prior distribution $P_{\mathbf{x}}$ supported on the sphere $\mathcal{S}^{n-1} = \{x \in \mathbb{R}^n : \|x\| = \sqrt{n}\}$. Then we let $\mathbf{Y} \in (\mathbb{R}^n)^{\otimes p}$ be the tensor

$$\mathbf{Y} = \lambda x_*^{\otimes p} + \mathbf{G}. \quad (1)$$

We will mostly focus on the *Rademacher-spiked model* where x_* is uniform in $\{\pm 1\}^n$, i.e., $P_{\mathbf{x}} = 2^{-n} \prod_i (\delta(x_i - 1) + \delta(x_i + 1))$. We will sometimes state results without specifying the prior $P_{\mathbf{x}}$, in which case the result holds for *any* prior normalized so that $\|x_*\| = \sqrt{n}$. Let $\mathbb{P}_\lambda$ denote the law of the tensor $\mathbf{Y}$. The parameter $\lambda = \lambda(n)$ may depend on n . We will consider the limit $n \rightarrow \infty$ with p held fixed.

Our algorithms will depend only on the entries $\mathbf{Y}_{i_1, \dots, i_p}$ where the indices $i_1, \dots, i_p$ are distinct: that is, on the collection

$$\{\mathbf{Y}_E = \lambda x_*^E + \mathbf{G}_E : E \subseteq [n], |E| = p\},$$

where $\mathbf{G}_E \sim \mathcal{N}(0, 1)$ and for a vector $x \in \mathbb{R}^n$ we write $x^E = \prod_{i \in E} x_i$.

Perhaps one of the simplest statistical tasks is binary hypothesis testing. In our case this amounts to, given a tensor $\mathbf{Y}$ as input with the promise that it was sampled from $\mathbb{P}_\lambda$ with $\lambda \in \{0, \bar{\lambda}\}$, determining whether $\lambda = 0$ or $\lambda = \bar{\lambda}$. We refer to $\mathbb{P}_\lambda$ for $\lambda > 0$ as the *planted* distribution, and $\mathbb{P}_0$ as the *null* distribution.

Definition 2.1. We say that an algorithm (or test) $f : (\mathbb{R}^n)^{\otimes p} \rightarrow \{0, 1\}$ achieves *strong detection* between $\mathbb{P}_0$ and $\mathbb{P}_\lambda$ if

$$\lim_{n \rightarrow \infty} \mathbb{P}_\lambda(f(\mathbf{Y}) = 1) = 1 \quad \text{and} \quad \lim_{n \rightarrow \infty} \mathbb{P}_0(f(\mathbf{Y}) = 0) = 1.$$

Additionally we say that f achieves *weak detection* between $\mathbb{P}_0$ and $\mathbb{P}_\lambda$ if the sum of Type-I and Type-II errors remains strictly below 1:

$$\limsup_{n \rightarrow \infty} \{\mathbb{P}_0(f(\mathbf{Y}) = 1) + \mathbb{P}_\lambda(f(\mathbf{Y}) = 0)\} < 1.$$

An additional goal is to recover the planted vector x_* . Note that when p is even, x_* and $-x_*$ have the same posterior probability. Thus, our goal is to recover x_* up to a sign.

Definition 2.2. The *normalized correlation* between vectors $\hat{x}, x \in \mathbb{R}^n$ is

$$\text{corr}(\hat{x}, x) = \frac{|\langle \hat{x}, x \rangle|}{\|\hat{x}\| \|x\|}.$$

Definition 2.3. An estimator $\hat{x} = \hat{x}(\mathbf{Y})$ achieves *weak recovery* if $\text{corr}(\hat{x}, x_*)$ is lower-bounded by a strictly positive constant—and we write $\text{corr}(\hat{x}, x_*) = \Omega(1)$ —with high probability, and achieves *strong recovery* if $\text{corr}(\hat{x}, x_*) = 1 - o(1)$ with high probability.

We expect that strong detection and weak recovery are generally equally difficult, although formal implications are not known in either direction. We will see in Section 2.5 that in some regimes, weak recovery and strong recovery are equivalent.

The matrix case. When $p = 2$, the spiked tensor model reduces to the *spiked Wigner model*. We know from random matrix theory that when $\lambda = \hat{\lambda}/\sqrt{n}$ with $\hat{\lambda} > 1$, strong detection is possible by thresholding the maximum eigenvalue of $\mathbf{Y}$, and weak recovery is achieved by PCA, i.e., taking the leading eigenvector [FP07, BGN11]. For many spike priors including Rademacher, strong detection and weak recovery are statistically impossible when $\hat{\lambda} < 1$ [MRZ15, DAM15, PWBM18], so there is a sharp phase transition at $\hat{\lambda} = 1$. (Note that weak detection is still possible below $\hat{\lambda} = 1$ [EKJ18].) A more sophisticated algorithm is AMP (approximate message passing) [DMM09, FR18, LKZ15a, DAM15], which can be thought of as a modification of the matrix power method which uses certain nonlinear transformations to exploit the structure of the spike prior. For many spike priors including Rademacher, AMP is known to achieve the information-theoretically optimal correlation with the true spike [DAM15, DMK⁺16]. However, like PCA, AMP achieves zero correlation asymptotically when $\hat{\lambda} < 1$. For certain spike priors (e.g., sparse priors), statistical-computational gaps can appear in which it is information-theoretically possible to succeed for some $\hat{\lambda} < 1$ but we do not expect that polynomial-time algorithms can do so [LKZ15a, LKZ15b, BMV⁺18].

The tensor case. The tensor case $p \geq 3$ was introduced by [RM14], who also proposed various algorithms. Information-theoretically, there is a sharp phase transition similar to the matrix case: for many spike priors $P_{\mathbf{x}}$ including Rademacher, if $\lambda = \hat{\lambda}n^{(1-p)/2}$ then weak recovery is possible when $\hat{\lambda} > \lambda_c$ and impossible when $\hat{\lambda} < \lambda_c$ for a particular constant $\lambda_c = \lambda_c(p, P_{\mathbf{x}})$ depending on p and $P_{\mathbf{x}}$ [LML⁺17]. Strong detection undergoes the same transition, being possible if $\lambda > \lambda_c$ and impossible otherwise [Che17, CHL18, JLM18] (see also [PWB16]). In fact, it is shown in these works that even weak detection is impossible below λ_c , in sharp contrast with the matrix case. There are polynomial-time algorithms (e.g., SOS) that succeed at both strong detection and strong recovery when $\lambda \gg n^{-p/4}$ for any spike prior [RM14, HSS15, HSSS16], which is above the information-theoretic threshold by a factor that diverges with n . There are also SOS lower bounds suggesting that (for many priors) no polynomial-time algorithm can succeed at strong detection or weak recovery when $\lambda \ll n^{-p/4}$ [HSS15, HKP⁺17]. Thus, unlike the matrix case, we expect a large statistical-computational gap, i.e., a “possible-but-hard” regime when $n^{(1-p)/2} \ll \lambda \ll n^{-p/4}$.

2.3 The Tensor Power Method and Local Algorithms

Various algorithms have been proposed and analyzed for tensor PCA [RM14, HSS15, HSSS16, ADGM16, AGJ18]. We will present two such algorithms that are simple, representative, and will be relevant to the discussion. The first is the *tensor power method* [AGH⁺14, RM14, AGJ17].

Algorithm 2.4. (Tensor Power Method) For a vector $u \in \mathbb{R}^n$ and a tensor $\mathbf{Y} \in (\mathbb{R}^n)^{\otimes p}$, let $\mathbf{Y}\{u\} \in \mathbb{R}^n$ denote the vector

$$\mathbf{Y}\{u\}_i = \sum_{j_1, \dots, j_{p-1}} \mathbf{Y}_{i, j_1, \dots, j_{p-1}} u_{j_1} \cdots u_{j_{p-1}}.$$

The *tensor power method* begins with an initial guess $u \in \mathbb{R}^n$ (e.g., chosen at random) and repeatedly iterates the update rule $u \leftarrow \mathbf{Y}\{u\}$ until $u/\|u\|$ converges.

The tensor power method appears to only succeed when $\lambda \gg n^{-1/2}$ [RM14], which is worse than the SOS threshold $\lambda \sim n^{-p/4}$. The AMP algorithm of [RM14] is a more sophisticated variant of the tensor power method, but AMP also fails unless $\lambda \gg n^{-1/2}$ [RM14]. Two other related algorithms, gradient descent and Langevin dynamics, also fail unless $\lambda \gg n^{-1/2}$ [AGJ18]. Following [AGJ18], we refer to all of these algorithms (tensor power method, AMP, gradient descent, Langevin dynamics)

as *local algorithms*, and we refer to the corresponding threshold $\lambda \sim n^{-1/2}$ as the *local threshold*. Here “local” is not a precise notion, but roughly speaking, local algorithms keep track of a current guess for x_* and iteratively update it to a nearby vector that is more favorable in terms of e.g. the log-likelihood. This discrepancy between local algorithms and SOS is what motivated the current work.

2.4 The Tensor Unfolding Algorithm

We have seen that local algorithms do not seem able to reach the SOS threshold. Let us now describe one of the simplest algorithms that does reach this threshold: *tensor unfolding*. Tensor unfolding was first proposed by [RM14], where it was shown to succeed when $\lambda \gg n^{-\lfloor p/2 \rfloor/2}$ and conjectured to succeed when $\lambda \gg n^{-p/4}$ (the SOS threshold). For the case $p = 3$, the same algorithm was later reinterpreted as a spectral relaxation of SOS, and proven to succeed² when $\lambda \gg n^{-3/4} = n^{-p/4}$ [HSS15], confirming the conjecture of [RM14]. We now present the tensor unfolding method, restricting to the case $p = 3$ for simplicity. There is a natural extension to all p [RM14], and (a close variant of) this algorithm will in fact appear as level $\ell = \lfloor p/2 \rfloor$ in our hierarchy of algorithms (see Section 3 and Appendix C).

Algorithm 2.5. (Tensor Unfolding) Given an order-3 tensor $\mathbf{Y} \in (\mathbb{R}^n)^{\otimes 3}$, flatten it to an $n \times n^2$ matrix $\mathbf{M}$, i.e., let $M_{i,jk} = Y_{ijk}$. Compute the leading eigenvector of $\mathbf{M}\mathbf{M}^\top$.

If we use the matrix power method to compute the leading eigenvector, we can restate the tensor unfolding method as an iterative algorithm: keep track of state vectors $u \in \mathbb{R}^n$ and $v \in \mathbb{R}^{n^2}$, initialize u randomly, and alternate between applying the update steps $v \leftarrow \mathbf{M}^\top u$ and $u \leftarrow \mathbf{M}v$. We will see later (see Section 4.1) that this can be interpreted as a message-passing algorithm between singleton indices, represented by u , and *pairs* of indices, represented by v . Thus, tensor unfolding is not “local” in the sense of Section 2.3 because it keeps a state of size $O(n^2)$ (keeping track of *pairwise* information) instead of size $O(n)$. We can, however, think of it as local on a “lifted” space, and this allows it to surpass the local threshold.

Other methods have also been shown to achieve the SOS threshold $\lambda \sim n^{-p/4}$, including SOS itself and various spectral methods inspired by it [HSS15, HSSS16].

2.5 Boosting and Linearization

One fundamental difference between the matrix case ($p = 2$) and tensor case ($p \geq 3$) is the following boosting property. The following result, implicit in [RM14], shows that for $p \geq 3$, if λ is substantially above the information-theoretic threshold (i.e., $\lambda \gg n^{(1-p)/2}$) then weak recovery can be boosted to strong recovery via a single power iteration. We give a proof in Appendix D.

Proposition 2.6. *Let $\mathbf{Y} \sim \mathbb{P}_\lambda$ with any spike prior P_x supported on $\mathcal{S}^{n-1}$. Suppose we have an initial guess $u \in \mathbb{R}^n$ satisfying $\text{corr}(u, x_*) \geq \tau$. Obtain $\hat{x}$ from u via a single iteration of the tensor power method: $\hat{x} = \mathbf{Y}\{u\}$. There exists a constant $c = c(p) > 0$ such that with high probability,*

$$\text{corr}(\hat{x}, x_*) \geq 1 - c\lambda^{-1}\tau^{1-p}n^{(1-p)/2}.$$

In particular, if $\tau > 0$ is any constant and $\lambda = \omega(n^{(1-p)/2})$ then $\text{corr}(\hat{x}, x) = 1 - o(1)$.

²The analysis of [HSS15] applies to a close variant of the spiked tensor model in which the noise tensor is asymmetric. We do not expect this difference to be important.

For $p \geq 3$, since we do not expect polynomial-time algorithms to succeed when $\lambda = O(n^{(1-p)/2})$, this implies an “all-or-nothing” phenomenon: for a given $\lambda = \lambda(n)$, the optimal polynomial-time algorithm will either achieve correlation that is asymptotically 0 or asymptotically 1. This is in stark contrast to the matrix case where, for $\lambda = \hat{\lambda}/\sqrt{n}$, the optimal correlation is a constant (in $[0, 1]$) depending on both $\hat{\lambda}$ and the spike prior $P_{\mathbf{x}}$.

This boosting result substantially simplifies things when $p \geq 3$ because it implies that the only important question is identifying the threshold for weak recovery, instead of trying to achieve the optimal correlation. Heuristically, since we only want to attain the optimal threshold, statistical physics suggests that we can use a simple “linearized” spectral algorithm instead of a more sophisticated nonlinear algorithm. To illustrate this in the matrix case ($p = 2$), one needs to use AMP in order to achieve optimal correlation, but one can achieve the optimal threshold using linearized AMP, which boils down to computing the top eigenvector. In the related setting of community detection in the stochastic block model, one needs to use belief propagation to achieve optimal correlation [DKMZ11b, DKMZ11a, MNS14], but one can achieve the optimal threshold using a linearized version of belief propagation, which is a spectral method on the non-backtracking walk matrix [KMM⁺13, BLM15] or the related *Bethe Hessian* [SKZ14]. Our spectral methods for tensor PCA are based on the *Kikuchi Hessian*, which is a generalization of the Bethe Hessian.

2.6 Subexponential-time Algorithms for Tensor PCA

The degree- ℓ sum-of-squares algorithm is a large semidefinite program that requires runtime $n^{O(\ell)}$ to solve. Oftentimes the regime of interest is when ℓ is constant, so that the algorithm runs in polynomial time. However, one can also explore the power of subexponential-time algorithms by letting $\ell = n^\delta$ for $\delta \in (0, 1)$, which gives an algorithm with runtime roughly 2^{n^δ} . Results of this type are known for tensor PCA [RRS17, BGG⁺16, BGL16]. The strongest such results are for a different variant of tensor PCA, which we now define.

Definition 2.7. In the *order- p discrete spiked tensor model* with spike prior $P_{\mathbf{x}}$ (normalized so that $\|x_*\| = \sqrt{n}$) and SNR parameter $\lambda \geq 0$, we draw a spike $x_* \sim P_{\mathbf{x}}$ and then for each $1 \leq i_1 \leq \dots \leq i_p \leq n$, we independently observe a $\{\pm 1\}$ -valued random variable $\mathbf{Y}_{i_1, \dots, i_p}$ with $\mathbb{E}[\mathbf{Y}_{i_1, \dots, i_p}] = \lambda(x_*^{\otimes p})_{i_1, \dots, i_p}$.

This model differs from our usual one in that the observations are conditionally Rademacher instead of Gaussian, but we do not believe this makes an important difference. However, for technical reasons, the known SOS results are strongest in this discrete setting.

Theorem 2.8 ([BGL16, HSS15]). *For any $1 \leq \ell \leq n$, there is an algorithm with runtime $n^{O(\ell)}$ that achieves strong detection and strong recovery in the order- p discrete spiked tensor model (with any spike prior) whenever*

$$\lambda \geq \ell^{1/2-p/4} n^{-p/4} \text{polylog}(n).$$

The work of [BGL16] shows how to certify an upper bound on the injective norm of a random $\{\pm 1\}$ -valued tensor, which immediately implies the algorithm for strong detection. When combined with [HSS15], this can also be made into an algorithm for strong recovery (see Lemma 4.4 of [HSS15]). Similar (but weaker) SOS results are also known for the standard spiked tensor model (see [RRS17] and *arXiv* version 1 of [BGL16]), and we expect that Theorem 2.8 also holds for this case.

When $\ell = n^\delta$ for $\delta \in (0, 1)$, Theorem 2.8 implies that we have an algorithm with runtime $n^{O(n^\delta)} = 2^{n^{\delta+o(1)}}$ that succeeds when $\lambda \gg n^{\delta/2-p\delta/4-p/4}$. Note that this interpolates smoothly between the polynomial-time threshold ($\lambda \sim n^{-p/4}$) when $\delta = 0$, and the information-theoretic

threshold ($\lambda \sim n^{(1-p)/2}$) when $\delta = 1$. We will prove (for p even) that our algorithms achieve this same tradeoff, and we expect this tradeoff to be optimal³.

3 Main results

In this section we present our main results about detection and recovery in the spiked tensor model. We propose a hierarchy of spectral methods, which are directly derived from the hierarchy of *Kikuchi free energies*. Specifically, the symmetric difference matrix defined below appears (approximately) as a submatrix of the Hessian of the Kikuchi free energy. The full details of this derivation are given in Section 4 and Appendix E. For now we simply state the algorithms and results.

We will restrict our attention to the Rademacher-spiked tensor model, which is the setting in which we derived our algorithms. However, we show in Appendix B that the same algorithm works for a large class of priors (at least for strong detection). Furthermore, we show in Appendix F.1 that the same algorithm can also be used for refuting random k -XOR formulas (when k is even).

We will also restrict to the case where the tensor order p is even. The case of odd p is discussed in Appendix C, where we give an algorithm derived from the Kikuchi Hessian and conjecture that it achieves optimal performance. We are unable to prove this, but we are able to prove that optimal results are attained by a related algorithm (see Appendix F.2).

Our approach requires the introduction of two matrices:

The symmetric difference matrix of order ℓ . Let p be even and let $\mathbf{Y} \in (\mathbb{R}^n)^{\otimes p}$ be the observed order- p symmetric tensor. We will only use the entries $\mathbf{Y}_{i_1, \dots, i_p}$ for which the indices $i_1, \dots, i_p$ are distinct; we denote such entries by $\mathbf{Y}_E$ where $E \subseteq [n]$ with $|E| = p$. Fix an integer $\ell \in [p/2, n - p/2]$, and consider the symmetric $\binom{n}{\ell} \times \binom{n}{\ell}$ matrix $\mathbf{M}$ indexed by sets $S \subseteq [n]$ of size ℓ , having entries

$$\mathbf{M}_{S,T} = \begin{cases} \mathbf{Y}_{S \triangle T} & \text{if } |S \triangle T| = p, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Here $\triangle$ denotes the symmetric difference between sets. The leading eigenvector of $\mathbf{M}$ is intended to be an estimate of $(x_*^S)_{|S|=\ell}$ where $x_*^S := \prod_{i \in S} x_i$. The following voting matrix is a natural rounding scheme to extract an estimate of x_* from such a vector.

The voting matrix. To a vector $v \in \mathbb{R}^{\binom{n}{\ell}}$ we associate the following symmetric $n \times n$ ‘voting’ matrix $\mathbf{V}(v)$ having entries

$$\mathbf{V}_{ii}(v) = 0 \quad \forall i \in [n], \quad \text{and} \quad \mathbf{V}_{ij}(v) = \frac{1}{2} \sum_{S,T \in \binom{[n]}{\ell}} v_S v_T \mathbf{1}_{S \triangle T = \{i,j\}} \quad \forall i \neq j. \quad (3)$$

Let us define the important quantity

$$d_\ell := \binom{n-\ell}{p/2} \binom{\ell}{p/2}. \quad (4)$$

This is the number of sets T of size ℓ such that $|S \triangle T| = p$ for a given set S of size ℓ . Now we are in position to formulate our algorithms for detection and recovery.

³One form of evidence suggesting that this tradeoff is optimal is based on the *low-degree likelihood ratio*; see [KWB19].

Algorithm 3.1 (Detection for even p).

1. Compute the top eigenvalue $\lambda_{\max}(\mathbf{M})$ of the symmetric difference matrix $\mathbf{M}$.
2. Reject the null hypothesis $\mathbb{P}_0$ (i.e., return ‘1’) if $\lambda_{\max}(\mathbf{M}) \geq \lambda d_\ell/2$.

Algorithm 3.2 (Recovery for even p).

1. Compute a (unit-norm) leading eigenvector⁴ $v_{\text{top}} \in \mathbb{R}^{\binom{n}{\ell}}$ of $\mathbf{M}$.
2. Form the associated voting matrix $\mathbf{V}(v_{\text{top}})$.
3. Compute a leading eigenvector $\hat{x}$ of $\mathbf{V}(v_{\text{top}})$, and output $\hat{x}$.

The next two theorems characterize the performance of Algorithms 3.1 and 3.2 for the strong detection and recovery tasks, respectively. The proofs can be found in Appendix A.

Theorem 3.3. *Consider the Rademacher-spiked tensor model with p even. For all $\lambda \geq 0$ and $\ell \in [p/2, n - p/2]$, we have*

$$\mathbb{P}_0\left(\lambda_{\max}(\mathbf{M}) \geq \frac{\lambda d_\ell}{2}\right) \vee \mathbb{P}_\lambda\left(\lambda_{\max}(\mathbf{M}) \leq \frac{\lambda d_\ell}{2}\right) \leq 2n^\ell e^{-\lambda^2 d_\ell/8}.$$

Therefore, Algorithm 3.1 achieves strong detection between $\mathbb{P}_0$ and $\mathbb{P}_\lambda$ if $\frac{1}{8}\lambda^2 d_\ell - \ell \log n \rightarrow +\infty$ as $n \rightarrow +\infty$.

Theorem 3.4. *Consider the Rademacher-spiked tensor model with p even. Let $\hat{x} \in \mathbb{R}^n$ be the output of Algorithm 3.2. There exists an absolute constant $c_0 > 0$ such that for all $\epsilon > 0$ and $\delta \in (0, 1)$, if $\ell \leq n\epsilon^2$ and $\lambda \geq c_0\epsilon^{-4}\sqrt{\log(n^\ell/\delta)/d_\ell}$, then $\text{corr}(\hat{x}, x_*) \geq 1 - c_0\epsilon$ with probability at least $1 - \delta$.*

Remark 3.5. If $\ell = o(n)$, we have $d_\ell = \Theta(n^{p/2}\ell^{p/2})$, and so the above theorems imply that strong detection and strong recovery are possible as soon as $\lambda \gg \ell^{-(p-2)/4}n^{-p/4}\sqrt{\log n}$. Comparing with Theorem 2.8, this scaling coincides with the guarantees achieved by the level- ℓ SOS algorithm of [BGL16], up to a possible discrepancy in logarithmic factors.

Due to the particularly simple structure of the symmetric difference matrix $\mathbf{M}$ (in particular, the fact that its entries are simply entries of $\mathbf{Y}$), the proof of detection (Theorem 3.3) follows from a straightforward application of the matrix Chernoff bound. In contrast, the corresponding SOS results [RRS17, BGG⁺16, BGL16] work with more complicated matrices involving high powers of the entries of $\mathbf{Y}$, and the analysis is much more involved.

Our proof of recovery is unusual in that the signal component of $\mathbf{M}$, call it $\mathbf{X}$, is not rank-one; it even has a vanishing spectral gap when $\ell \gg 1$. Thus, the leading eigenvector of $\mathbf{M}$ does not correlate well with the leading eigenvector of $\mathbf{X}$. While this may seem to render recovery hopeless at first glance, this is not the case, due to the fact that *many* eigenvectors (actually, eigenspaces) of $\mathbf{X}$ contain non-trivial information about the spike x_* , as opposed to only the top one. We prove this by exploiting the special structure of $\mathbf{X}$ through the *Johnson scheme*, and using tools from Fourier analysis on a slice of the hypercube, in particular a Poincaré-type inequality by [Fil16].

⁴We define a leading eigenvector to be an eigenvector whose eigenvalue is maximal (although our results still hold for an eigenvector whose eigenvalue is maximal in absolute value).

Removing the logarithmic factor Both Theorem 3.3 and Theorem 3.4 involve a logarithmic factor in n in the lower bound on SNR λ . These log factors are an artifact of the matrix Chernoff bound, and we believe they can be removed. (The analysis of [HSS15] removes the log factors for the tensor unfolding algorithm, which is essentially the case $p = 3$ and $\ell = 1$ of our algorithm.) This suggests the following precise conjecture on the power of polynomial-time algorithms.

Conjecture 3.6. *Fix p and let ℓ be constant (not depending on n). There exists a constant $c_p(\ell) > 0$ with $c_p(\ell) \rightarrow 0$ as $\ell \rightarrow \infty$ (with p fixed) such that if $\lambda \geq c_p(\ell)n^{-p/4}$ then Algorithm 3.1 and Algorithm 3.2 (which run in time $n^{O(\ell)}$) achieve strong detection and strong recovery, respectively.*

Specifically, we expect $c_p(\ell) \sim \ell^{-(p-2)/4}$ for large ℓ .

4 Motivating the Symmetric Difference Matrices

In this section we motivate the symmetric difference matrices used in our algorithms. In Section 4.1 we give some high-level intuition, including an explanation of how our algorithms can be thought of as iterative message-passing procedures among subsets of size ℓ . In Section 4.2 we give a more principled derivation based on the Kikuchi Hessian, with many of the calculations deferred to Appendix E.

4.1 Intuition: Higher-Order Message-Passing and Maximum Likelihood

As stated previously, our algorithms will choose to ignore the entries $\mathbf{Y}_{i_1, \dots, i_p}$ for which $i_1, \dots, i_p$ are not distinct; these entries turn out to be unimportant asymptotically. We restrict to the Rademacher-spiked tensor model, as this yields a clean and simple derivation. The posterior distribution for the spike x_* given the observed tensor $\mathbf{Y}$ is

$$\begin{aligned} \mathbb{P}(x | \mathbf{Y}) &\propto \exp \left\{ -\frac{1}{2} \sum_{i_1 < \dots < i_p} (\mathbf{Y}_{i_1 \dots i_p} - \lambda x_{i_1} \dots x_{i_p})^2 \right\} \\ &\propto \exp \left\{ \lambda \sum_{i_1 < \dots < i_p} \mathbf{Y}_{i_1 \dots i_p} x_{i_1} \dots x_{i_p} \right\} = \exp \left\{ \lambda \sum_{|E|=p} \mathbf{Y}_E x^E \right\} \end{aligned} \quad (5)$$

over the domain $x \in \{\pm 1\}^n$. We take p to be even; a similar derivation works for odd p . Now fix $\ell \in [p/2, n - p/2]$. We can write the above as

$$\mathbb{P}(x | \mathbf{Y}) \propto \exp \left\{ \frac{\lambda}{N} \sum_{|S \Delta T|=p} \mathbf{Y}_{S \Delta T} x^S x^T \right\}, \quad (6)$$

where the sum is over ordered pairs (S, T) of sets $S, T \subseteq [n]$ with $|S| = |T| = \ell$ and $|S \Delta T| = p$, and where $N = d_\ell \binom{n}{\ell} / \binom{n}{p}$ is the number of terms (S, T) with a given symmetric difference E .

A natural message-passing algorithm to maximize the log-likelihood is the following. For each $S \subseteq [n]$ of size $|S| = \ell$, keep track of a variable $u_S \in \mathbb{R}$, which is intended to be an estimate of $x_*^S := \prod_{i \in S} (x_*)_i$. Note that there are consistency constraints that $(x_*^S)_{|S|=\ell}$ must obey, such as $x_*^S x_*^T x_*^V = 1$ when $S \Delta T \Delta V = \emptyset$; we will relax the problem and will not require our vector $u = (u_S)_{|S|=\ell}$ to obey such constraints. Instead we simply attempt to maximize

$$\frac{1}{\|u\|^2} \sum_{|S \Delta T|=p} \mathbf{Y}_{S \Delta T} u_S u_T \quad (7)$$

over all $u \in \mathbb{R}^{\binom{n}{\ell}}$. To do this, we iterate the update equations

$$u_S \leftarrow \sum_{T: |S \Delta T| = p} \mathbf{Y}_{S \Delta T} u_T. \quad (8)$$

We call S and T *neighbors* if $|S \Delta T| = p$. Intuitively, each neighbor T of S sends a message $m_{T \rightarrow S} := \mathbf{Y}_{S \Delta T} u_T$ to S , indicating T 's opinion about u_S . We update u_S to be the sum of all incoming messages.

Now note that the sum in (7) is simply $\|u\|^{-2} u^\top \mathbf{M} u$ where $\mathbf{M}$ is the symmetric difference matrix, and (8) can be written as

$$u \leftarrow \mathbf{M} u.$$

Thus our natural message-passing scheme is precisely power iteration against $\mathbf{M}$, and so we should take the leading eigenvector v_{top} of $\mathbf{M}$ as our estimate of $(x_*^S)_{|S|=\ell}$ (up to a scaling factor). Finally, defining our voting matrix $\mathbf{V}(v_{\text{top}})$ and taking its leading eigenvector is a natural method for rounding v_{top} to a vector of the form u^x where $u_S^x = x^S$, thus restoring the consistency constraints we ignored before.

Indeed, if we carry out this procedure on *all* subsets $S \subseteq [n]$ then this works as intended, and no rounding is necessary: consider the $2^n \times 2^n$ matrix $\overline{\mathbf{M}}_{S,T} = \mathbf{Y}_{S \Delta T} \mathbf{1}_{|S \Delta T|=p}$. It is easy to verify that the eigenvectors of $\overline{\mathbf{M}}$ are precisely the Fourier basis vectors on the hypercube, namely vectors of the form u^x where $u_S^x = x^S$ and $x \in \{\pm 1\}^n$. Moreover, the eigenvalue associated to u^x is

$$\frac{1}{2^n} (u^x)^\top \overline{\mathbf{M}} u^x = \frac{1}{2^n} \sum_{S,T \subseteq [n]: |S \Delta T|=p} \mathbf{Y}_{S \Delta T} x^S x^T = \sum_{|E|=p} \mathbf{Y}_E x^E.$$

This is the expression in the log-likelihood in (5). Thus the leading eigenvector of $\overline{\mathbf{M}}$ is exactly u^x where x is the maximum-likelihood estimate of x_* .

This procedure succeeds all the way down to the information-theoretic threshold $\lambda \sim n^{(1-p)/2}$, but takes exponential time. Our contribution can be viewed as showing that even when we restrict to the submatrix $\mathbf{M}$ of $\overline{\mathbf{M}}$ supported on subsets of size ℓ , the leading eigenvector still allows us to recover x_* whenever the SNR is sufficiently large. Proving this requires us to perform Fourier analysis over a slice of the hypercube rather than the simpler setting of the entire hypercube, which we do by appealing to Johnson schemes and some results of [Fil16].

4.2 Variational Inference and Kikuchi Free Energy

We now introduce the *Kikuchi approximations to the free energy* (or simply the *Kikuchi free energies*) of the above posterior (5) [Kik51, Kik94], the principle from which our algorithms are derived. For concreteness we restrict to the case of the Rademacher-spiked tensor model, but the Kikuchi free energies can be defined for general graphical models [YFW03].

The posterior distribution in (5) is a Gibbs distribution $\mathbb{P}(x | \mathbf{Y}) \propto e^{-\beta H(x)}$ with random Hamiltonian $H(x) := -\lambda \sum_{i_1 < \dots < i_p} \mathbf{Y}_{i_1 \dots i_p} x_{i_1} \dots x_{i_p}$, and inverse temperature $\beta = 1$. We let $Z_n(\beta; \mathbf{Y}) := \sum_{x \in \{\pm 1\}^n} e^{-\beta H(x)}$ be the partition function of the model, and denote by $F_n(\beta; \mathbf{Y}) := -\frac{1}{\beta} \log Z_n(\beta; \mathbf{Y})$ its *free energy*. It is a classical fact that the Gibbs distribution has the following variational characterization. Fix a finite domain Ω (e.g., $\{\pm 1\}^n$), $\beta > 0$ and $H : \Omega \rightarrow \mathbb{R}$. Consider the optimization problem

$$\inf_{\mu} F(\mu), \quad (9)$$

where the supremum is over probability distributions μ supported on Ω , and define the *free energy functional* F of μ by

$$F(\mu) := \mathbb{E}_{x \sim \mu} [H(x)] - \frac{1}{\beta} \mathcal{S}(\mu), \quad (10)$$

where $\mathcal{S}(\mu)$ is the Shannon entropy of μ , i.e., $\mathcal{S}(\mu) = -\sum_{x \in \Omega} \mu(x) \log \mu(x)$. Then the unique optimizer of (9) is the Gibbs distribution $\mu^*(x) \propto \exp(-\beta H(x))$. If we specialize this statement to our setting, $\mu^* = \mathbb{P}(\cdot | \mathbf{Y})$ and $F_n(1; \mathbf{Y}) = F(\mu^*)$. We refer to [WJ08] for more background.

In light of the above variational characterization, a natural algorithmic strategy to learn the posterior distribution is to minimize the free energy functional $F(\mu)$ over distributions μ . However, this is *a priori* intractable because (for a high-dimensional domain such as $\Omega = \{\pm 1\}^n$) an exponential number of parameters are required to represent μ . The idea underlying the *belief propagation* algorithm [Pea86, YFW03] is to work only with locally-consistent marginals, or *beliefs*, instead of a complete distribution μ . Standard belief propagation works with beliefs on singleton variables and on pairs of variables. The *Bethe free energy* is a proxy for the free energy that only depends on these beliefs, and belief propagation is a certain procedure that iteratively updates the beliefs in order to locally minimize the Bethe free energy. The level- r *Kikuchi free energy* is a generalization of the *Bethe free energy* that depends on r -wise beliefs and gives (in principle) increasingly better approximations of $F(\mu^*)$ as r increases. Our algorithms are based on the principle of locally minimizing Kikuchi free energy, which we define next.

We now define the level- r Kikuchi approximation to the free energy. We require $r \geq p$, i.e., the Kikuchi level needs to be at least as large as the interactions present in the data (although the $r < p$ case could be handled by defining a modified graphical model with auxiliary variables). The Bethe free energy is the case $r = 2$.

For $S \subseteq [n]$ with $0 < |S| \leq r$, let $b_S : \{\pm 1\}^S \rightarrow \mathbb{R}$ denote the *belief* on S , which is a probability mass function over $\{\pm 1\}^{|S|}$ representing our belief about the joint distribution of $x_S := \{x_i\}_{i \in S}$. Let $b = \{b_S : S \in \binom{[n]}{\leq r}\}$ denote the set of beliefs on s -wise interactions for all $s \leq r$. Following [YFW03], the Kikuchi free energy is a real-valued functional $\mathcal{K}$ of b having the form $\mathcal{E} - \frac{1}{\beta} \mathcal{S}$ (in our case, $\beta = 1$). Here the first term is the ‘energy’ term

$$\mathcal{E}(b) = -\lambda \sum_{|S|=p} \mathbf{Y}_S \sum_{x_S \in \{\pm 1\}^S} b_S(x_S) x^S.$$

where, recall, $x^S := \prod_{i \in S} x_i$. (This is a proxy for the term $\mathbb{E}_{x \sim \mu} [H(x)]$ in (10).) The second term in $\mathcal{K}$ is the ‘entropy’ term

$$\mathcal{S}(b) = \sum_{0 < |S| \leq r} c_S \mathcal{S}_S(b), \quad \mathcal{S}_S(b) = - \sum_{x_S \in \{\pm 1\}^S} b_S(x_S) \log b_S(x_S),$$

where the *overcounting numbers* are $c_S := \sum_{T \supseteq S, |T| \leq r} (-1)^{|T \setminus S|}$. These are defined so that for any $S \subseteq [n]$ with $0 < |S| \leq r$,

$$\sum_{T \supseteq S, |T| \leq r} c_T = 1, \quad (11)$$

which corrects for overcounting. Notice that $\mathcal{E}$ and $\mathcal{S}$ each take the form of an “expectation” with respect to the beliefs b_S ; these would be actual expectations were the beliefs the marginals of an actual probability distribution. This situation is to be contrasted with the notion of a *pseudo-expectation*, which plays a central role in the theory underlying the sum-of-squares algorithm.

Our algorithms are based on the *Kikuchi Hessian*, a generalization of the Bethe Hessian matrix that was introduced in the setting of community detection [SKZ14]. The Bethe Hessian is the

Hessian of the Bethe free energy with respect to the moments of the beliefs, evaluated at belief propagation’s so-called “uninformative fixed point.” The bottom eigenvector of the Bethe Hessian is a natural estimator for the planted signal because it represents the best direction for local improvement of the Bethe free energy, starting from belief propagation’s uninformative starting point. We generalize this method and compute the analogous Kikuchi Hessian matrix. The full derivation is given in Appendix E. The order- ℓ symmetric difference matrix (2) (approximately) appears as a submatrix of the level- r Kikuchi Hessian whenever $r \geq \ell + p/2$.

5 Conclusion

We have presented a hierarchy of spectral algorithms for tensor PCA, inspired by variational inference and statistical physics. In particular, the core idea of our approach is to locally minimize the Kikuchi free energy. We specifically implemented this via the Kikuchi Hessian, but there may be many other viable approaches to minimizing the Kikuchi free energy such as generalized belief propagation [YFW03]. Broadly speaking, we conjecture that for many average-case problems, algorithms based on Kikuchi free energy and algorithms based on sum-of-squares should both achieve the optimal tradeoff between runtime and statistical power. One direction for further work is to verify that this analogy holds for problems other than tensor PCA; in particular, we show here that it also applies to refuting random k -XOR formulas when k is even.

Perhaps one benefit of the Kikuchi hierarchy over the sum-of-squares hierarchy is that it has allowed us to *systematically* obtain spectral methods, simply by computing a certain Hessian matrix. Furthermore, the algorithms we obtained are simpler than their SOS counterparts. We are hopeful that the Kikuchi hierarchy will provide a roadmap for systematically deriving simple and optimal algorithms for a large class of problems.

Acknowledgments

We thank Alex Russell for suggesting the matrix Chernoff bound (Theorem A.4). For helpful discussions, we thank Afonso Bandeira, Sam Hopkins, Pravesh Kothari, Florent Krzakala, Tselil Schramm, Jonathan Shi, and Lenka Zdeborová. This project started during the workshop *Spin Glasses and Related Topics* held at the Banff International Research Station (BIRS) in the Fall of 2018. We thank our hosts at BIRS as well as the workshop organizers: Antonio Auffinger, Wei-Kuo Chen, Dmitry Panchenko, and Lenka Zdeborová.

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.
- [ADGM16] Anima Anandkumar, Yuan Deng, Rong Ge, and Hossein Mobahi. Homotopy analysis for tensor PCA. *arXiv preprint arXiv:1610.09322*, 2016.
- [AG89] Richard Arratia and Louis Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of mathematical biology*, 51(1):125–131, 1989.

- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research*, 15(1):2773–2832, 2014.
- [AGJ17] Animashree Anandkumar, Rong Ge, and Majid Janzamin. Analyzing tensor power method dynamics in overcomplete regime. *The Journal of Machine Learning Research*, 18(1):752–791, 2017.
- [AGJ18] Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Algorithmic thresholds for tensor PCA. *arXiv preprint arXiv:1808.00921*, 2018.
- [AKS98] Noga Alon, Michael Krivelevich, and Benny Sudakov. Finding a large hidden clique in a random graph. *Random Structures & Algorithms*, 13(3-4):457–466, 1998.
- [AOW15] Sarah R Allen, Ryan O’Donnell, and David Witmer. How to refute a random CSP. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 689–708. IEEE, 2015.
- [BCR19] Giulio Biroli, Chiara Cammarota, and Federico Ricci-Tersenghi. How to iron out rough landscapes and get optimal performances: Replicated gradient descent and its application to tensor PCA. *arXiv preprint arXiv:1905.12294*, 2019.
- [BGG⁺16] Vijay VSP Bhattiprolu, Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee, and Madhur Tulsiani. Multiplicative approximations for polynomial optimization over the unit sphere. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 23, page 1, 2016.
- [BGL16] Vijay Bhattiprolu, Venkatesan Guruswami, and Euiwoong Lee. Sum-of-squares certificates for maxima of random tensors on the sphere. *arXiv preprint arXiv:1605.00903*, 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 Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 428–437. IEEE, 2016.
- [BLM15] Charles Bordenave, Marc Lelarge, and Laurent Massoulié. Non-backtracking spectrum of random graphs: community detection and non-regular ramanujan graphs. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1347–1357. IEEE, 2015.
- [BMV⁺18] Jess Banks, Cristopher Moore, Roman Vershynin, Nicolas Verzelen, and Jiaming Xu. Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *IEEE Transactions on Information Theory*, 64(7):4872–4894, 2018.
- [Bur17] Amanda Burcroff. Johnson schemes and certain matrices with integral eigenvalues. Technical report, University of Michigan, 2017.

- [CGL04] Amin Coja-Oghlan, Andreas Goerdt, and André Lanka. Strong refutation heuristics for random k-sat. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 310–321. Springer, 2004.
- [Che17] Wei-Kuo Chen. Phase transition in the spiked random tensor with rademacher prior. *arXiv preprint arXiv:1712.01777*, 2017.
- [CHL18] Wei-Kuo Chen, Madeline Handschy, and Gilad Lerman. Phase transition in random tensors with multiple spikes. *arXiv preprint arXiv:1809.06790*, 2018.
- [DAM15] Yash Deshpande, Emmanuel Abbe, and Andrea Montanari. Asymptotic mutual information for the two-groups stochastic block model. *arXiv preprint arXiv:1507.08685*, 2015.
- [DKMZ11a] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011.
- [DKMZ11b] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Inference and phase transitions in the detection of modules in sparse networks. *Physical Review Letters*, 107(6):065701, 2011.
- [DMK⁺16] 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. In *Advances in Neural Information Processing Systems*, pages 424–432, 2016.
- [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.
- [EKJ18] Ahmed El Alaoui, Florent Krzakala, and Michael I Jordan. Fundamental limits of detection in the spiked wigner model. *arXiv preprint arXiv:1806.09588*, 2018.
- [Fil16] Yuval Filmus. An orthogonal basis for functions over a slice of the Boolean hypercube. *The Electronic Journal of Combinatorics*, 23(1):P1–23, 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.
- [FR18] Alyson K Fletcher and Sundeep Rangan. Iterative reconstruction of rank-one matrices in noise. *Information and Inference: A Journal of the IMA*, 7(3):531–562, 2018.
- [Gri01] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1-2):613–622, 2001.
- [GS10] Chris Godsil and Sung Y Song. Association schemes. *Handbook of Combinatorial Designs*,, pages 325–330, 2010.
- [Has19] Matthew B Hastings. Classical and quantum algorithms for tensor principal component analysis. *arXiv preprint arXiv:1907.12724*, 2019.

- [HKP⁺17] Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 720–731. IEEE, 2017.
- [HSS15] Samuel B Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-square proofs. In *Conference on Learning Theory*, pages 956–1006, 2015.
- [HSSS16] Samuel B Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 178–191. ACM, 2016.
- [Jer92] Mark Jerrum. Large cliques elude the metropolis process. *Random Structures & Algorithms*, 3(4):347–359, 1992.
- [JL04] Iain M Johnstone and Arthur Yu Lu. Sparse principal components analysis. *Unpublished manuscript*, 7, 2004.
- [JLM18] Aukosh Jagannath, Patrick Lopatto, and Leo Miolane. Statistical thresholds for tensor PCA. *arXiv preprint arXiv:1812.03403*, 2018.
- [Kik51] Ryoichi Kikuchi. A theory of cooperative phenomena. *Phys. Rev.*, 81:988, 1951.
- [Kik94] Ryoichi Kikuchi. Special issue in honor of R. Kikuchi. *Progr. Theor. Phys. Suppl*, 115, 1994.
- [KMM⁺13] Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, 2013.
- [KWB19] Dmitriy Kunisky, Alexander S Wein, and Afonso S Bandeira. Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio. *arXiv preprint arXiv:1907.11636*, 2019.
- [Las01] Jean B Lasserre. An explicit exact SDP relaxation for nonlinear 0-1 programs. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 293–303. Springer, 2001.
- [LKZ15a] 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.
- [LKZ15b] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Phase transitions in sparse PCA. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 1635–1639. IEEE, 2015.
- [LML⁺17] Thibault Lesieur, Léo Miolane, Marc Lelarge, Florent Krzakala, and Lenka Zdeborová. Statistical and computational phase transitions in spiked tensor estimation. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 511–515. IEEE, 2017.

- [MNS14] Elchanan Mossel, Joe Neeman, and Allan Sly. Belief propagation, robust reconstruction and optimal recovery of block models. In *Conference on Learning Theory*, pages 356–370, 2014.
- [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.
- [Oli10] Roberto Oliveira. Sums of random Hermitian matrices and an inequality by Rudelson. *Electronic Communications in Probability*, 15:203–212, 2010.
- [Par00] Pablo A Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.
- [Pea86] Judea Pearl. Fusion, propagation, and structuring in belief networks. *Artificial intelligence*, 29(3):241–288, 1986.
- [PWB16] Amelia Perry, Alexander S Wein, and Afonso S Bandeira. Statistical limits of spiked tensor models. *arXiv preprint arXiv:1612.07728*, 2016.
- [PWBM18] Amelia Perry, Alexander S Wein, Afonso S Bandeira, and Ankur Moitra. Optimality and sub-optimality of PCA I: Spiked random matrix models. *The Annals of Statistics*, 46(5):2416–2451, 2018.
- [RM14] Emile Richard and Andrea Montanari. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014.
- [RRS17] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random CSPs below the spectral threshold. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 121–131. ACM, 2017.
- [Saa16] Alaa Saade. Spectral inference methods on sparse graphs: theory and applications. *arXiv preprint arXiv:1610.04337*, 2016.
- [Sch79] Alexander Schrijver. A comparison of the Delsarte and Lovász bounds. *IEEE Trans. Information Theory*, 25(4):425–429, 1979.
- [Sch08] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 593–602. IEEE, 2008.
- [Sho87] Naum Z Shor. Class of global minimum bounds of polynomial functions. *Cybernetics*, 23(6):731–734, 1987.
- [SKZ14] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Spectral clustering of graphs with the Bethe Hessian. In *Advances in Neural Information Processing Systems*, pages 406–414, 2014.
- [Tro12] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- [Wat90] William C Waterhouse. The absolute-value estimate for symmetric multilinear forms. *Linear Algebra and its Applications*, 128:97–105, 1990.

- [WJ08] Martin J Wainwright and Michael I Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- [YFW03] Jonathan S Yedidia, William T Freeman, and Yair Weiss. Understanding belief propagation and its generalizations. *Exploring artificial intelligence in the new millennium*, 8:236–239, 2003.

A Analysis of Symmetric Difference and Voting Matrices

We adopt the notation $x^S := \prod_{i \in S} x_i$ for $x \in \{\pm 1\}^n$ and $S \subseteq [n]$. Recall the matrix $\mathbf{M}$ indexed by sets $S \subseteq [n]$ of size ℓ , having entries

$$\mathbf{M}_{S,T} = \mathbf{Y}_{S\Delta T} \mathbf{1}_{|S\Delta T|=p} \quad \text{where} \quad \mathbf{Y}_{S\Delta T} = \lambda x_*^{S\Delta T} + \mathbf{G}_{S\Delta T}. \quad (12)$$

First, observe that we can restrict our attention to the case where the spike is the all-ones vector $x_* = \mathbb{1}$ without loss of generality. To see this, conjugate $\mathbf{M}$ by a diagonal matrix $\mathbf{D}$ with diagonal entries $\mathbf{D}_{S,S} = x_*^S$ and obtain $(\mathbf{M}')_{S,T} = (\mathbf{D}^{-1} \mathbf{M} \mathbf{D})_{S,T} = \mathbf{Y}'_{S\Delta T} \mathbf{1}_{|S\Delta T|=p}$ where $\mathbf{Y}'_{S\Delta T} = x_*^S x_*^T \mathbf{Y}_{S\Delta T} = x_*^{S\Delta T} \mathbf{Y}_{S\Delta T} = \lambda + g'_{S\Delta T}$ where $g'_{S\Delta T} = x_*^S x_*^T g_{S\Delta T}$. By symmetry of the Gaussian distribution, $(g'_E)_{|E|=p}$ are i.i.d. $\mathcal{N}(0, 1)$ random variables. Therefore, the two matrices have the same spectrum and the eigenvectors of $\mathbf{M}$ can be obtained from those of $\mathbf{M}'$ by pre-multiplying by $\mathbf{D}$. So from now on we write

$$\mathbf{M} = \lambda \mathbf{X} + \mathbf{Z}, \quad (13)$$

where $\mathbf{X}_{S,T} = \mathbf{1}_{|S\Delta T|=p}$ and $\mathbf{Z}_{S,T} = g_{S\Delta T} \mathbf{1}_{|S\Delta T|=p}$, where $(g_E)_{|E|=p}$ is a collection of i.i.d. $\mathcal{N}(0, 1)$ r.v.'s.

A.1 Structure of $\mathbf{X}$

The matrix $\mathbf{X}$ is the adjacency matrix of a regular graph $J_{n,\ell,p}$ on $\binom{n}{\ell}$ vertices, where vertices are represented by sets, and two sets S and T are connected by an edge if $|S \Delta T| = p$, or equivalently $|S \cap T| = \ell - p/2$. This matrix belongs to the Bose-Mesner algebra of the (n, ℓ) -Johnson association scheme (see for instance [Sch79, GS10]). This is the algebra of $\binom{n}{\ell} \times \binom{n}{\ell}$ real- or complex-valued symmetric matrices where the entry $\mathbf{X}_{S,T}$ depends only on the size of the intersection $|S \cap T|$. In addition to this set of matrices being an algebra, it is a *commutative* algebra, which means that all such matrices are simultaneously diagonalizable and share the same eigenvectors.

Films [Fil16] provides a common eigenbasis for this algebra: for $0 \leq m \leq \ell$, let $\varphi = (a_1, b_1, \dots, a_m, b_m)$ be a sequence of $2m$ distinct elements of $[n]$. Let $|\varphi| = 2m$ denote its total length. Now define a vector $u^\varphi \in \mathbb{R}^{\binom{n}{\ell}}$ having coordinates

$$u_S^\varphi = \prod_{i=1}^m (\mathbf{1}_{a_i \in S} - \mathbf{1}_{b_i \in S}), \quad |S| = \ell.$$

In the case $m = 0$, φ is the empty sequence $\emptyset$ and we have $u^\varphi = \mathbb{1}$ (the all-ones vector).

Proposition A.1. *Each u^φ is an eigenvector of $\mathbf{X}$. Furthermore, the linear space $\mathcal{Y}_m := \text{span}\{u^\varphi : |\varphi| = 2m\}$ for $0 \leq m \leq \ell$ is an eigenspace of $\mathbf{X}$ (i.e., all vectors u^φ with sequences φ of length of $2m$ have the same eigenvalue μ_m). Lastly $\mathbb{R}^{\binom{n}{\ell}} = \bigoplus_{m=0}^{\ell} \mathcal{Y}_m$, and $\dim \mathcal{Y}_m = \binom{n}{m} - \binom{n}{m-1}$. (By convention, $\binom{n}{-1} = 0$.)*

Proof. The first two statements are the content of Lemma 4.3 in [Fil16]. The dimension of $\mathcal{Y}_m$ is given in Lemma 2.1 in [Fil16]. $\square$

We note that $(u^\varphi)_{|\varphi|=2m}$ are not linearly independent; an orthogonal basis, called the Young basis, consisting of linear combinations of the u^φ 's is given explicitly in [Fil16].

We see from the above proposition that $\mathbf{X}$ has $\ell + 1$ distinct eigenvalues $\mu_0 \geq \mu_1 \geq \dots \geq \mu_\ell$, each one corresponding to the eigenspace $\mathcal{Y}_m$. The first eigenvalue is the degree of the graph $J_{n,\ell,p}$:

$$\mu_0 = d_\ell = \binom{\ell}{p/2} \binom{n-\ell}{p/2}. \quad (14)$$

We provide an explicit formula for all the remaining eigenvalues:

Lemma A.2. *The eigenvalues of $\mathbf{X}$ are as follows:*

$$\mu_m = \sum_{s=0}^{\min(m,p/2)} (-1)^s \binom{m}{s} \binom{\ell-m}{p/2-s} \binom{n-\ell-m}{p/2-s}, \quad 0 \leq m \leq \ell. \quad (15)$$

Proof. These are the so-called Eberlein polynomials, which are polynomials in m of degree p (see, e.g., [Sch79]). We refer to [Bur17] for formulae in more general contexts, but we give a proof here for completeness. Let $A = \{a_1, \dots, a_m\}$ and $B = \{b_1, \dots, b_m\}$. Note that u_S^φ is nonzero if and only if $|S \cap \{a_i, b_i\}| = 1$ for each $1 \leq i \leq m$. By symmetry, we can assume that $A \subseteq S$ and $S \cap B = \emptyset$. Then μ_m is the sum over all sets T , such that $|S \triangle T| = p$ and $|T \cap \{a_i, b_i\}| = 1$ for each i , of $(-1)^s$ where $s = |T \cap B|$. For each s there are $\binom{m}{s}$ choices of this set of indices, giving the first binomial. Adding these b_i to T and removing these a_i from S contributes $2s$ to $|S \triangle T|$. To achieve $|S \triangle T| = p$, we also need to remove $p/2 - s$ elements of $S \setminus A$ from S , giving the second binomial. We also need to add $p/2 - s$ elements of $\overline{S \cup B}$ to T , giving the third binomial. Finally, we have $s \leq m$ and $s \leq p/2$. $\square$

As the following lemma shows, the succeeding eigenvalues decay rapidly with m .

Lemma A.3. *Let $3 \leq p \leq \sqrt{n}$ and let $\ell < n/p^2$. For all $0 \leq m \leq \ell$ it holds that*

$$\frac{|\mu_m|}{\mu_0} \leq \max \left\{ \left(1 - \frac{m}{\ell}\right)^{p/2}, \frac{p}{n} \right\}.$$

Proof. The terms in (15) have alternating signs. We will show that they decrease in absolute value beyond the first nonzero term, so that it gives a bound on μ_m . We consider two cases. First, suppose $m \leq \ell - p/2$ so that the $s = 0$ term is positive. Then the $(s+1)$ st term divided by the s th term is, in absolute value,

$$\begin{aligned} \frac{\binom{m}{s+1} \binom{\ell-m}{p/2-s-1} \binom{n-\ell-m}{p/2-s-1}}{\binom{m}{s} \binom{\ell-m}{p/2-s} \binom{n-\ell-m}{p/2-s}} &= \binom{m-s}{s+1} \left(\frac{p/2-s}{\ell-m-p/2+s+1} \right) \left(\frac{p/2-s}{n-\ell-m-p/2+s+1} \right) \\ &\leq \frac{m(p/2)(p/2)}{(\ell-m-p/2+1)(n-\ell-m-p/2+1)} \\ &\leq \frac{(\ell-p/2)(p/2)(p/2)}{n-2\ell+1} \quad \text{since } m \leq \ell - p/2 \\ &\leq \frac{\ell p^2/4}{n-2\ell} < 1. \end{aligned}$$

It follows that the $s = 0$ term is an upper bound,

$$\mu_m \leq \binom{\ell - m}{p/2} \binom{n - \ell - m}{p/2},$$

and so

$$\frac{\mu_m}{\mu_0} \leq \binom{\ell - m}{p/2} \bigg/ \binom{\ell}{p/2} \leq \left(\frac{\ell - m}{\ell} \right)^{p/2}.$$

Next we consider the case $m > \ell - p/2$, so that the first nonzero term has $s = p/2 - \ell + m \geq 1$. Intuitively, this reduces μ_m by at least a factor of n , and we will show this is the case. For the terms with $s \geq p/2 - \ell + m$, the ratio of absolute values is again bounded by

$$\begin{aligned} \frac{\binom{m}{s+1} \binom{\ell-m}{p/2-s-1} \binom{n-\ell-m}{p/2-s-1}}{\binom{m}{s} \binom{\ell-m}{p/2-s} \binom{n-\ell-m}{p/2-s}} &= \binom{m-s}{s+1} \binom{p/2-s}{\ell-m-p/2+s+1} \binom{p/2-s}{n-\ell-m-p/2+s+1} \\ &< \frac{(\ell-p/2)(p/2)(p/2)}{(p/2-\ell+m+1)(n-2\ell+1)} \leq \frac{\ell p^2/4}{n-2\ell} < 1. \end{aligned}$$

It follows that the $s = p/2 - \ell + m$ term gives a bound on the absolute value,

$$\begin{aligned} |\mu_m| &\leq \binom{m}{\ell-p/2} \binom{n-\ell-m}{\ell-m} \\ &\leq \binom{\ell}{p/2} \binom{n-\ell}{\ell-m}, \end{aligned}$$

in which case

$$\frac{|\mu_m|}{\mu_0} \leq \frac{\binom{n-\ell}{\ell-m}}{\binom{n-\ell}{p/2}} = \frac{(p/2)!(n-\ell-p/2)!}{(\ell-m)!(n-2\ell+m)!} \leq \left(\frac{p/2}{n-\ell-p/2} \right)^{p/2-\ell+m} \leq \left(\frac{p}{n} \right)^{p/2-\ell+m} \leq \frac{p}{n}.$$

Combining these two cases gives the stated result. $\square$

A.2 Proof of Strong Detection

Here we prove our strong detection result, Theorem 3.3. The proof doesn't exploit the full details of the structure exhibited above. Instead, the proof is a straightforward application of the matrix Chernoff bound for Gaussian series [Oli10] (see also Theorem 4.1.1 of [Tro12]):

Theorem A.4. *Let $\{\mathbf{A}_i\}$ be a finite sequence of fixed symmetric $d \times d$ matrices, and let $\{\xi_i\}$ be independent $\mathcal{N}(0, 1)$ random variables. Let $\mathbf{\Sigma} = \sum_i \xi_i \mathbf{A}_i$. Then, for all $t \geq 0$,*

$$\mathbb{P}(\|\mathbf{\Sigma}\|_{\text{op}} \geq t) \leq 2de^{-t^2/2\sigma^2} \quad \text{where } \sigma^2 = \|\mathbb{E}[\mathbf{\Sigma}^2]\|_{\text{op}}.$$

Let us first write $\mathbf{M}$ in the form of a Gaussian series. For a set $E \in \binom{[n]}{p}$, define the $\binom{n}{\ell} \times \binom{n}{\ell}$ matrix $\mathbf{A}_E$ as

$$(\mathbf{A}_E)_{S,T} = \mathbf{1}_{S \Delta T = E}.$$

It is immediate that for $\lambda = 0$, $\mathbf{M} = \mathbf{Z} = \sum_{|E|=p} g_E \mathbf{A}_E$ where $(g_E)_{|E|=p}$ is a collection of i.i.d. $\mathcal{N}(0, 1)$ random variables. The second moment of this random matrix is

$$\mathbb{E}[\mathbf{M}^2] = \sum_{E: |E|=p} \mathbf{A}_E^2 = d_\ell \mathbf{I},$$

since $\mathbf{A}_E^2$ is the diagonal matrix with $(\mathbf{A}_E)_{S,S} = \mathbf{1}_{|S \triangle E|=p}$, and summing over all E gives d_ℓ on the diagonal. The operator norm of the second moment is then d_ℓ . It follows that for all $t \geq 0$,

$$\mathbb{P}_0(\lambda_{\max}(\mathbf{M}) \geq t) \leq 2 \binom{n}{\ell} e^{-t^2/2d_\ell} \leq 2n^\ell e^{-t^2/2d_\ell}. \quad (16)$$

Now letting $t = \frac{\lambda d_\ell}{2}$ yields the first statement of the theorem.

As for the second statement, we have $\lambda_{\max}(\mathbf{M}) \geq \|\mathbf{X}\|_{\text{op}} - \|\mathbf{Z}\|_{\text{op}} = \lambda d_\ell - \|\mathbf{Z}\|_{\text{op}}$ where $\mathbf{Z}$ is defined in (13). Applying the same bound we have

$$\mathbb{P}_\lambda\left(\lambda_{\max}(\mathbf{M}) \leq \frac{\lambda d_\ell}{2}\right) \leq \mathbb{P}\left(\|\mathbf{Z}\|_{\text{op}} \leq \frac{\lambda d_\ell}{2}\right) \leq 2n^\ell e^{-\lambda^2 d_\ell/8}.$$

A.3 Proof of Strong Recovery

Here we prove our strong recovery result, Theorem 3.4. Let $v_0 = v_{\text{top}}(\mathbf{M})$ be a unit-norm leading eigenvector of $\mathbf{M}$. For a fixed $m \in [\ell]$ (to be determined later on), we write the orthogonal decomposition $v_0 = v^{(m)} + v^\perp$, where $v^{(m)} \in \bigoplus_{s \leq m} \mathcal{Y}_s$, and $v^\perp$ in the orthogonal complement. The goal is to first show that if m is proportional to ℓ then $v^\perp$ has small Euclidean norm, so that v_0 and $v^{(m)}$ have high inner product. The second step of the argument is to approximate the voting matrix $\mathbf{V}(v_0)$ by $\mathbf{V}(v^{(m)})$, and then use Fourier-analytic tools to reason about the latter.

Let us start with the first step.

Lemma A.5. *With $\mathbf{Z}$ defined as in (13), we have*

$$\|v^\perp\|^2 \leq 2 \frac{\|\mathbf{Z}\|_{\text{op}}}{\lambda(\mu_0 - \mu_{m+1})}.$$

Proof. Let us absorb the factor λ in the definition of the matrix $\mathbf{X}$. Let $\{u_0, \dots, u_d\}$ be a set of eigenvectors of $\mathbf{X}$ which also form an orthogonal basis for $\bigoplus_{s \leq m} \mathcal{Y}_s$, with u_0 being the top eigenvector of $\mathbf{X}$ (u_0 is the normalized all-ones vector). We start with the inequality $u_0^\top \mathbf{M} u_0 \leq v_0^\top \mathbf{M} v_0$. The left-hand side of the inequality is $\mu_0 + u_0^\top \mathbf{Z} u_0$. The right-hand side is $v_0^\top \mathbf{X} v_0 + v_0^\top \mathbf{Z} v_0$. Moreover $v_0^\top \mathbf{X} v_0 = v_0^\top \mathbf{X} v^{(m)} + v_0^\top \mathbf{X} v^\perp$. Since $v^{(m)} \in \bigoplus_{s \leq m} \mathcal{Y}_s$, by Proposition A.1, $\mathbf{X} v^{(m)}$ belongs to the space as well, and therefore $v_0^\top \mathbf{X} v^{(m)} = v^{(m)\top} \mathbf{X} v^{(m)}$. Similarly $v_0^\top \mathbf{X} v^\perp = (v^\perp)^\top \mathbf{X} v^\perp$, so $v_0^\top \mathbf{X} v_0 = v^{(m)\top} \mathbf{X} v^{(m)} + (v^\perp)^\top \mathbf{X} v^\perp$. Therefore the inequality becomes

$$\mu_0 + u_0^\top \mathbf{Z} u_0 \leq v^{(m)\top} \mathbf{X} v^{(m)} + (v^\perp)^\top \mathbf{X} v^\perp + v_0^\top \mathbf{Z} v_0.$$

Since $v^\perp$ is orthogonal to the top m eigenspaces of $\mathbf{X}$ we have $(v^\perp)^\top \mathbf{X} v^\perp \leq \mu_{m+1} \|v^\perp\|_2^2$. Moreover, $v^{(m)\top} \mathbf{X} v^{(m)} \leq \mu_0 \|v^{(m)}\|^2$, hence

$$\mu_0 + u_0^\top \mathbf{Z} u_0 \leq \mu_0 \|v^{(m)}\|^2 + \mu_{m+1} \|v^\perp\|^2 + v_0^\top \mathbf{Z} v_0.$$

By rearranging and applying the triangle inequality we get,

$$(\mu_0 - \mu_{m+1}) \|v^\perp\|^2 \leq |v_0^\top \mathbf{Z} v_0| + |u_0^\top \mathbf{Z} u_0| \leq 2 \|\mathbf{Z}\|_{\text{op}}.$$

□

Combining this fact with Lemma A.3, recalling that $\mu_0 = d_\ell$ we obtain

Lemma A.6. For any $\epsilon > 0$ and $\delta \in (0, 1)$, if $\lambda \geq \epsilon^{-1} \sqrt{2 \log(n^\ell/\delta)/d_\ell}$, then

$$\|v^\perp\|^2 \leq \epsilon \frac{\ell}{m},$$

with probability at least $1 - \delta$.

Proof. Lemma A.3 implies $\frac{1}{\mu_0 - \mu_{m+1}} \leq \frac{1}{\mu_0} \frac{1}{1 - (1 - \frac{m+1}{\ell})^{p/2}} \leq \frac{1}{\mu_0} \cdot \frac{\ell}{m}$. Therefore, Lemma A.5 implies

$$\|v^\perp\|^2 \leq \frac{\|\mathbf{Z}\|_{\text{op}}}{\lambda d_\ell} \frac{\ell}{m}.$$

The operator norm of the noise can be bounded by a matrix Chernoff bound [Oli10, Tro12], similarly to our argument in the proof of detection: for all $t \geq 0$

$$\mathbb{P}(\|\mathbf{Z}\|_{\text{op}} \geq t) \leq n^\ell e^{-t^2/2d_\ell}.$$

Therefore, letting $\lambda \geq \epsilon^{-1} \sqrt{2 \log(n^\ell/\delta)/d_\ell}$ we obtain the desired result. $\square$

A.3.1 Analysis of the Voting Matrix

Recall that the voting matrix $\mathbf{V}(v)$ of a vector $v \in \mathbb{R}^{\binom{n}{\ell}}$ has zeros on the diagonal, and off-diagonal entries

$$\begin{aligned} \mathbf{V}_{ij}(v) &= \frac{1}{2} \sum_{S, T \in \binom{[n]}{\ell}} v_S v_T \mathbf{1}_{S \Delta T = \{i, j\}} \\ &= \sum_{S \in \binom{[n]}{\ell}} v_S v_{S \Delta \{i, j\}} \mathbf{1}_{i \in S, j \notin S}, \quad 1 \leq i \neq j \leq n. \end{aligned}$$

It will be more convenient in our analysis to work with $\mathbf{V}(v^{(m)})$ instead of $\mathbf{V}(v_0)$. To this end we produce the following approximation result:

Lemma A.7. Let $u, e \in \mathbb{R}^{\binom{n}{\ell}}$ and $v = u + e$. Then

$$\|\mathbf{V}(v) - \mathbf{V}(u)\|_F^2 \leq 3\ell^2 \|e\|^2 (2\|u\|^2 + \|e\|^2).$$

In particular,

$$\|\mathbf{V}(v_0) - \mathbf{V}(v^{(m)})\|_F^2 \leq 9\ell^2 \|v^\perp\|^2.$$

Proof. Let us introduce the shorthand notation $\langle u, v \rangle_{ij} := \sum_{S \in \binom{[n]}{\ell}} v_S v_{S \Delta \{i, j\}} \mathbf{1}_{i \in S, j \notin S}$. We have

$$\begin{aligned} \|\mathbf{V}(v) - \mathbf{V}(u)\|_F^2 &= \sum_{i, j} (\mathbf{V}_{ij}(u + e) - \mathbf{V}_{ij}(u))^2 \\ &= \sum_{i, j} (\langle u + e, u + e \rangle_{ij} - \langle u, u \rangle_{ij})^2 \\ &= \sum_{i, j} (\langle u, e \rangle_{ij} + \langle e, u \rangle_{ij} + \langle e, e \rangle_{ij})^2 \\ &\leq 3 \sum_{i, j} (\langle u, e \rangle_{ij}^2 + \langle e, u \rangle_{ij}^2 + \langle e, e \rangle_{ij}^2), \end{aligned} \tag{17}$$

where the last step uses the bound $(a + b + c)^2 \leq 3(a^2 + b^2 + c^2)$. Now we expand

$$\begin{aligned}
\sum_{i,j} \langle u, e \rangle_{ij}^2 &= \sum_{i,j} \left(\sum_{S: |S|=\ell, i \in S, j \notin S} u_S e_{S \Delta \{i,j\}} \right)^2 \\
&\leq \sum_{i,j} \left(\sum_{S: |S|=\ell, i \in S, j \notin S} u_S^2 \right) \left(\sum_{S: |S|=\ell, i \in S, j \notin S} e_{S \Delta \{i,j\}}^2 \right) \quad (\text{Cauchy-Schwarz}) \\
&\leq \sum_{i,j} \left(\sum_{S: |S|=\ell, i \in S} u_S^2 \right) \left(\sum_{T: |T|=\ell, j \in T} e_T^2 \right) \\
&= \left(\sum_i \sum_{S: |S|=\ell, i \in S} u_S^2 \right) \left(\sum_j \sum_{T: |T|=\ell, j \in T} e_T^2 \right) \\
&= \left(\ell \sum_{|S|=\ell} u_S^2 \right) \left(\ell \sum_{|T|=\ell} e_T^2 \right) \\
&= \ell^2 \|u\|^2 \|e\|^2. \tag{18}
\end{aligned}$$

Plugging this back into (17) yields the desired result. To obtain $\|\mathbf{V}(v_0) - \mathbf{V}(v^{(m)})\|_F^2 \leq 9\ell^2 \|v^\perp\|^2$ we just bound $2\|v^{(m)}\|^2 + \|v^\perp\|^2$ by 3. $\square$

Let us also state the following lemma, which will be need later on:

Lemma A.8. For $u \in \mathbb{R}^{\binom{n}{\ell}}$, $\|\mathbf{V}(u)\|_F^2 \leq \ell^2 \|u\|^4$.

Proof. Note that $\|\mathbf{V}(u)\|_F^2 = \sum_{i,j} \mathbf{V}_{ij}(u)^2 = \sum_{i,j} \langle u, u \rangle_{ij}^2$ and so the desired result follows immediately from (18). $\square$

Next, in the main technical part of the proof, we show that $\mathbf{V}(v^{(m)})$ is close to a multiple of the all-ones matrix in Frobenius norm:

Proposition A.9. Let $\hat{\mathbb{1}} = \mathbb{1}/\sqrt{n}$, $\alpha = \ell \|v^{(m)}\|^2$ and $\eta = \frac{m}{\ell} + \frac{\ell}{n}$. Then

$$\|\mathbf{V}(v^{(m)}) - \alpha \hat{\mathbb{1}} \hat{\mathbb{1}}^\top\|_F^2 \leq 2\eta \alpha^2.$$

Before proving the above proposition, let us put the results together and prove our recovery result.

A.3.2 Proof of Theorem 3.4

For $\epsilon, \delta > 0$, assume $\lambda \geq \epsilon^{-1} \sqrt{2 \log(n^\ell/\delta)/d_\ell}$. By Lemma A.6 and Lemma A.7 we have

$$\|\mathbf{V}(v_0) - \mathbf{V}(v^{(m)})\|_F^2 \leq 9\ell^2 \frac{\ell}{m} \epsilon,$$

with probability at least $1 - \delta$. Moreover, by Lemma A.9, we have

$$\|\mathbf{V}(v^{(m)}) - \alpha \hat{\mathbb{1}} \hat{\mathbb{1}}^\top\|_F^2 \leq 2\eta \alpha^2,$$

With $\alpha = \ell \|v^{(m)}\|^2 \leq \ell$ and $\eta = \frac{m}{\ell} + \frac{\ell}{n}$. Therefore, by a triangle inequality we have

$$\begin{aligned}\|\mathbf{V}(v_0) - \alpha \hat{\mathbf{1}} \hat{\mathbf{1}}^\top\|_F &\leq 3\ell \sqrt{\frac{\ell\epsilon}{m}} + \sqrt{2\eta\alpha} \\ &\leq 3\ell \left(\sqrt{\frac{\ell\epsilon}{m}} + \sqrt{\frac{m}{\ell} + \frac{\ell}{n}} \right),\end{aligned}$$

with probability at least $1 - \delta$. Now let us choose a value of m that achieves a good tradeoff of the above two error terms: $m = \ell\sqrt{\epsilon}$. Let us also use the inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for positive a, b , to obtain

$$\|\mathbf{V}(v_0) - \alpha \hat{\mathbf{1}} \hat{\mathbf{1}}^\top\|_F \leq 6\ell \left(\epsilon^{1/4} + \sqrt{\frac{\ell}{n}} \right), \quad (19)$$

under the same event.

Now let $\hat{x}$ be a leading eigenvector of $\mathbf{V}(v_0)$, and let $\mathbf{R} = \mathbf{V}(v_0) - \alpha \hat{\mathbf{1}} \hat{\mathbf{1}}^\top$. Since $\hat{x}^\top \mathbf{V}(v_0) \hat{x} \geq \hat{\mathbf{1}}^\top \mathbf{V}(v_0) \hat{\mathbf{1}}$, we have

$$\alpha \langle \hat{x}, \hat{\mathbf{1}} \rangle^2 + \hat{x}^\top \mathbf{R} \hat{x} \geq \alpha + \hat{\mathbf{1}}^\top \mathbf{R} \hat{\mathbf{1}}.$$

Therefore

$$\langle \hat{x}, \hat{\mathbf{1}} \rangle^2 \geq 1 - 2 \frac{\|\mathbf{R}\|_{\text{op}}}{\alpha}.$$

Since $\alpha = \ell(1 - \|v^\perp\|^2)$, and $\|v^\perp\|^2 \leq \epsilon \frac{\ell}{m} \leq \sqrt{\epsilon}$, the bound (19) (together with $\|\mathbf{R}\|_{\text{op}} \leq \|\mathbf{R}\|_F$) implies

$$\langle \hat{x}, \hat{\mathbf{1}} \rangle^2 \geq 1 - 12 \frac{\epsilon^{1/4} + \sqrt{\ell/n}}{1 - \sqrt{\epsilon}}. \quad (20)$$

To conclude the proof of our theorem, we let $\ell \leq n\sqrt{\epsilon}$, $\epsilon < 1/16$ and then replace ϵ by ϵ^4 : we obtain $\langle \hat{x}, \hat{\mathbf{1}} \rangle^2 \geq 1 - 48\epsilon$ with probability at least $1 - \delta$ if $\lambda \geq \epsilon^{-4} \sqrt{2 \log(n^\ell/\delta)/d_\ell}$.

A.3.3 Proof of Proposition A.9

Let $\alpha = \ell \|v^{(m)}\|^2$. By Lemma A.8 we have $\|\mathbf{V}(v^{(m)})\|_F^2 \leq \alpha^2$. Therefore

$$\begin{aligned}\|\mathbf{V}(v^{(m)}) - \alpha \hat{\mathbf{1}} \hat{\mathbf{1}}^\top\|_F^2 &= \|\mathbf{V}(v^{(m)})\|_F^2 - \frac{2}{n} \alpha \sum_{i,j=1}^n \mathbf{V}_{ij}(v^{(m)}) + \alpha^2 \\ &\leq 2\alpha^2 - \frac{2}{n} \alpha \sum_{i,j=1}^n \mathbf{V}_{ij}(v^{(m)}).\end{aligned} \quad (21)$$

Now we need a lower bound on $\sum_{i,j=1}^n \mathbf{V}_{ij}(v^{(m)})$. This will crucially rely on the fact that $v^{(m)}$ lies in the span of the top m eigenspaces of $\mathbf{X}$:

Lemma A.10. *For a fixed $m \leq \ell$, let $v \in \bigoplus_{s=0}^m \mathcal{Y}_s$. Then*

$$\sum_{i,j=1}^n \mathbf{V}_{ij}(v) \geq ((\ell - m)n - \ell^2) \|v\|^2.$$

We plug the result of the above lemma in (21) to obtain

$$\begin{aligned}\|\mathbf{V}(v^{(m)}) - \alpha \hat{\mathbf{1}} \hat{\mathbf{1}}^\top\|_F^2 &\leq 2\alpha^2 \left(1 - \frac{\ell - m}{\ell} + \frac{\ell}{n} \right) \\ &= 2\alpha^2 \left(\frac{m}{\ell} + \frac{\ell}{n} \right),\end{aligned}$$

as desired.

A.3.4 A Poincaré Inequality on a Slice of the Hypercube

To prove Lemma A.10, we need some results on Fourier analysis on the slice of the hypercube $\binom{[n]}{\ell}$. Following [Fil16], we define the following. First, given a function $f : \binom{[n]}{\ell} \mapsto \mathbb{R}$, we define its expectation as its average value over all sets of size ℓ , and write

$$\mathbb{E}_{|S|=\ell} [f(S)] := \frac{1}{\binom{n}{\ell}} \sum_{|S|=\ell} f(S).$$

We also define its variance as

$$\mathbb{V}[f] := \mathbb{E}_{|S|=\ell} [f(S)^2] - \mathbb{E}_{|S|=\ell} [f(S)]^2.$$

Moreover, we identify a vector $u \in \mathbb{R}^{\binom{n}{\ell}}$ with a function on sets of size ℓ in the obvious way: $f(S) = u_S$.

Definition A.11. For $u \in \mathbb{R}^{\binom{n}{\ell}}$ and $i, j \in [n]$, let $u^{(ij)}$ denote the vector having coordinates

$$u_S^{(ij)} = \begin{cases} u_{S \triangle \{i,j\}} & \text{if } |S \cap \{i,j\}| = 1 \\ u_S & \text{otherwise.} \end{cases}$$

(The operation $u \mapsto u^{(ij)}$ exchanges the roles of i and j whenever possible.)

Definition A.12. For $u \in \mathbb{R}^{\binom{n}{\ell}}$ and $i, j \in [n]$, define the *influence* of the pair (i, j) as

$$\text{Inf}_{ij}[u] := \frac{1}{2} \mathbb{E}_{|S|=\ell} \left[(u_S^{(ij)} - u_S)^2 \right],$$

and the *total influence* as

$$\text{Inf}[u] := \frac{1}{n} \sum_{i < j} \text{Inf}_{ij}[u].$$

With this notation, our main tool is a version of Poincaré's inequality on $\binom{[n]}{\ell}$:

Lemma A.13 (Lemma 5.6 in [Fil16]). *For $v \in \bigoplus_{s=0}^m \mathcal{Y}_s$ we have*

$$\text{Inf}[v] \leq m \mathbb{V}[v].$$

Proof of Lemma A.10. We have $m \mathbb{E}_{|S|=\ell} [v_S^2] \geq m \mathbb{V}[v] \geq \text{Inf}[v] = \frac{1}{n} \sum_{i < j} \text{Inf}_{ij}[v]$, and

$$\begin{aligned} 2 \text{Inf}_{ij}[v] &= \mathbb{E}_{|S|=\ell} [(v_S^{(ij)} - v_S)^2] \\ &= \mathbb{E}_{|S|=\ell} [\mathbf{1}_{|S \cap \{i,j\}|=1} (v_S^{(ij)} - v_S)^2] \\ &= \mathbb{E}_{|S|=\ell} \left[\mathbf{1}_{|S \cap \{i,j\}|=1} \left((v_S^{(ij)})^2 - 2v_S^{(ij)} v_S + v_S^2 \right) \right]. \end{aligned}$$

Since $\mathbb{E}_{|S|=\ell} [\mathbf{1}_{|S \cap \{i,j\}|=1} (v_S^{(ij)})^2] = \mathbb{E}_{|S|=\ell} [\mathbf{1}_{|S \cap \{i,j\}|=1} v_S^2]$, the above is equal to

$$2 \mathbb{E}_{|S|=\ell} \left[\mathbf{1}_{|S \cap \{i,j\}|=1} (v_S^2 - v_S^{(ij)} v_S) \right].$$

Now,

$$\begin{aligned}
\frac{1}{n} \sum_{i < j} \text{Inf}_{ij}[v] &= \frac{1}{n} \sum_{i < j} \binom{n}{\ell}^{-1} \sum_{S: |S|=\ell, |S \cap \{i,j\}|=1} (v_S^2 - v_S^{(ij)} v_S) \\
&= \frac{1}{n} \binom{n}{\ell}^{-1} \left[\ell(n-\ell) \sum_{|S|=\ell} v_S^2 - 2 \sum_{i < j} \sum_{S: |S|=\ell, i \in S, j \notin S} v_S v_{S \Delta \{i,j\}} \right] \\
&= \frac{1}{n} \binom{n}{\ell}^{-1} \left[\ell(n-\ell) \sum_{|S|=\ell} v_S^2 - 2 \sum_{i < j} \mathbf{V}_{ij}(v) \right] \\
&= \frac{1}{n} \binom{n}{\ell}^{-1} \left[\ell(n-\ell) \|v\|^2 - \sum_{i,j} \mathbf{V}_{ij}(v) \right].
\end{aligned}$$

Therefore

$$m \|v\|^2 \geq \frac{1}{n} \left(\ell(n-\ell) \|v\|^2 - \sum_{i,j} \mathbf{V}_{ij}(v) \right),$$

and so

$$\sum_{i,j} \mathbf{V}_{ij}(v) \geq (\ell(n-\ell) - mn) \|v\|^2 \geq ((\ell-m)n - \ell^2) \|v\|^2.$$

as desired. $\square$

B Detection for General Priors

While we have mainly focused on the Rademacher-spiked tensor model, we now show that our algorithm works just as well (at least for detection) for a much larger class of spike priors.

Theorem B.1. *Let $p \geq 2$ be even. Consider the spiked tensor model with a spike prior $P_{\mathbf{x}}$ that draws the entries of x_* i.i.d. from some distribution π on $\mathbb{R}$ (which does not depend on n), normalized so that $\mathbb{E}[\pi^2] = 1$. There is a constant C (depending on p and π) such that if $\lambda \geq C \ell^{1/2} d_\ell^{-1/2} \sqrt{\log n}$ then Algorithm 3.1 achieves strong detection.*

Proof. From (16), we have $\|\mathbf{Z}\|_{\text{op}} = O(\sqrt{\ell d_\ell \log n})$ with high probability, and so it remains to give a lower bound on $\|\mathbf{X}\|_{\text{op}}$. Letting $u_S = \prod_{i \in S} \text{sgn}((x_*)_i)$ for $|S| = \ell$,

$$\|\mathbf{X}\|_{\text{op}} \geq \frac{u^\top \mathbf{X} u}{\|u\|^2} = \binom{n}{\ell}^{-1} u^\top \mathbf{X} u$$

where

$$\begin{aligned}
u^\top \mathbf{X} u &= \sum_{|S \Delta T|=p} u_S \mathbf{X}_{S,T} u_T \\
&= \sum_{|S \Delta T|=p} x_*^{S \Delta T} \prod_{i \in S} \text{sgn}((x_*)_i) \prod_{i \in T} \text{sgn}((x_*)_i) \\
&= \sum_{|S \Delta T|=p} |x_*^{S \Delta T}|.
\end{aligned}$$

We have

$$\mathbb{E}[u^\top \mathbf{X} u] = \binom{n}{\ell} d_\ell (\mathbb{E}|\pi|)^p = C(\pi, p) \binom{n}{\ell} d_\ell, \quad (22)$$

and

$$\begin{aligned} \text{Var} [u^\top \mathbf{X} u] &= \text{Var} \left[\sum_{|S \triangle T|=p} |x_*^{S \triangle T}| \right] \\ &= \sum_{|S \triangle T|=p} \sum_{|S' \triangle T'|=p} \text{Cov}(|x_*^{S \triangle T}|, |x_*^{S' \triangle T'}|). \end{aligned} \quad (23)$$

We have

$$\begin{aligned} \text{Cov}(|x_*^{S \triangle T}|, |x_*^{S' \triangle T'}|) &\leq \sqrt{\text{Var}(|x_*^{S \triangle T}|) \text{Var}(|x_*^{S' \triangle T'}|)} \\ &= \text{Var}(|x_*^{S \triangle T}|) \leq \mathbb{E}[|x_*^{S \triangle T}|^2] = (\mathbb{E}[\pi^2])^p = 1. \end{aligned}$$

Also, $\text{Cov}(|x_*^{S \triangle T}|, |x_*^{S' \triangle T'}|) = 0$ unless $S \triangle T$ and $S' \triangle T'$ have nonempty intersection. Using Lemma B.2 (below), the fraction of terms in (23) that are nonzero is at most p^2/n and so

$$\text{Var} [u^\top \mathbf{X} u] \leq \left[\binom{n}{\ell} d_\ell \right]^2 \frac{p^2}{n}. \quad (24)$$

By Chebyshev's inequality, it follows from (22) and (24) that $u^\top \mathbf{X} u \geq \frac{1}{2} C(\pi, p) \binom{n}{\ell} d_\ell$ with probability at least $1 - \frac{4p^2}{C(\pi, p)^2 n}$. This implies $\|\mathbf{X}\|_{\text{op}} \geq \frac{1}{2} C(\pi, p) d_\ell$ with the same probability, and so we have strong detection provided $\lambda \geq c_0 \ell^{1/2} d_\ell^{-1/2} \sqrt{\log n}$ for a particular constant $c_0 = c_0(\pi, p)$. $\square$

Above, we made use of the following lemma.

Lemma B.2. *Fix $A \subseteq [n]$ with $|A| = a$. Let B be chosen uniformly at random from all subsets of $[n]$ of size b . Then $\mathbb{P}(A \cap B \neq \emptyset) \leq \frac{ab}{n}$.*

Proof. Each element of A will lie in B with probability b/n , so the result follows by a union bound over the elements of A . $\square$

C The Odd- p Case

When p is odd, the Kikuchi Hessian still gives rise to a spectral algorithm. While we conjecture that this algorithm is optimal, we unfortunately only know how to prove sub-optimal results for it. (However, we can prove optimal results for a related algorithm; see Appendix F.2.) We now state the algorithm and its conjectured performance.

Let p be odd and fix an integer $\ell \in [p/2, n - p/2]$. Consider the symmetric difference matrix $\mathbf{M} \in \mathbb{R}^{\binom{n}{\ell} \times \binom{n}{\ell+1}}$ with entries

$$\mathbf{M}_{S,T} = \begin{cases} \mathbf{Y}_{S \triangle T} & \text{if } |S \triangle T| = p, \\ 0 & \text{otherwise,} \end{cases}$$

where $S, T \subseteq [n]$ with $|S| = \ell$ and $|T| = \ell + 1$.

Algorithm C.1 (Recovery for odd p). Let u be a (unit-norm) top left-singular vector of $\mathbf{M}$ and let $v = \mathbf{M}^\top u$ be the corresponding top right-singular vector. Output $\hat{x} = \hat{x}(\mathbf{Y}) \in \mathbb{R}^n$, defined by

$$\hat{x}_i = \sum_{S \in \binom{[n]}{\ell}, T \in \binom{[n]}{\ell+1}} u_S v_T \mathbf{1}_{S \Delta T = \{i\}}, \quad i \in [n].$$

Notice that the rounding step consisting in extracting an n -dimensional vector $\hat{x}$ from the singular vectors of $\mathbf{M}$ is slightly simpler than the even- p case, in that it does not require forming a voting matrix. We conjecture that, like the even case, this algorithm matches the performance of SOS.

Conjecture C.2. *Consider the Rademacher-spiked tensor model with $p \geq 3$ odd. If*

$$\lambda \gg \ell^{-(p-2)/4} n^{-p/4}$$

then (i) there is a threshold $\tau = \tau(n, p, \ell, \lambda)$ such that strong detection can be achieved by thresholding the top singular value of $\mathbf{M}$ at τ , and (ii) Algorithm C.1 achieves strong recovery.

Similarly to the proof of Theorem 3.4, the matrix Chernoff bound (Theorem A.4) can be used to show that strong recovery is achievable when $\lambda \gg \ell^{-(p-1)/4} n^{-(p-1)/4}$, which is weaker than SOS when $\ell \ll n$. We now explain the difficulties involved in improving this. We can decompose $\mathbf{M}$ into a signal part and a noise part: $\mathbf{M} = \lambda \mathbf{X} + \mathbf{Z}$. In the regime of interest, $\ell^{-(p-2)/4} n^{-p/4} \ll \lambda \ll \ell^{-(p-1)/4} n^{-(p-1)/4}$, the signal term is smaller in operator norm than the noise term, i.e., $\lambda \|\mathbf{X}\|_{\text{op}} \ll \|\mathbf{Z}\|_{\text{op}}$. While at first sight this would seem to suggest that detection and recovery are hopeless, we actually expect that $\lambda \mathbf{X}$ still affects the top singular value and singular vectors of $\mathbf{M}$. This phenomenon is already present in the analysis of tensor unfolding (the case $p = 3$, $\ell = 1$) [HSS15], but it seems that new ideas are required to extend the analysis beyond this case.

D Proof of Boosting

Definition D.1. For a tensor $\mathbf{G} \in (\mathbb{R}^n)^{\otimes p}$, the *injective tensor norm* is

$$\|\mathbf{G}\|_{\text{inj}} := \max_{\|u^{(1)}\|=\dots=\|u^{(p)}\|=1} \sum_{i_1, \dots, i_p} \mathbf{G}_{i_1, \dots, i_p} u_{i_1}^{(1)} \cdots u_{i_p}^{(p)},$$

where $u^{(j)} \in \mathbb{R}^n$. For a symmetric tensor $\mathbf{G}$, it is known [Wat90] that equivalently,

$$\|\mathbf{G}\|_{\text{inj}} = \max_{\|u\|=1} \left| \sum_{i_1, \dots, i_p} \mathbf{G}_{i_1, \dots, i_p} u_{i_1} \cdots u_{i_p} \right|.$$

Proof of Proposition 2.6. Write $\hat{x} = \lambda \langle u, x_* \rangle^{p-1} x_* + \Delta$ where $\|\Delta\| \leq \|\mathbf{G}\|_{\text{inj}} \|u\|^{p-1}$. We have

$$|\langle \hat{x}, x_* \rangle| \geq \lambda |\langle u, x_* \rangle|^{p-1} \|x_*\|^2 - \|\Delta\| \|x_*\|,$$

and

$$\|\hat{x}\| \leq \lambda |\langle u, x_* \rangle|^{p-1} \|x_*\| + \|\Delta\|,$$

and so

$$\begin{aligned}
\text{corr}(\hat{x}, x_*) &= \frac{|\langle \hat{x}, x_* \rangle|}{\|\hat{x}\| \|x_*\|} \\
&\geq \frac{\lambda |\langle u, x_* \rangle|^{p-1} \|x_*\| - \|\Delta\|}{\lambda |\langle u, x_* \rangle|^{p-1} \|x_*\| + \|\Delta\|} \\
&= 1 - \frac{2\|\Delta\|}{\lambda |\langle u, x_* \rangle|^{p-1} \|x_*\| + \|\Delta\|} \\
&\geq 1 - \frac{2\|\Delta\|}{\lambda |\langle u, x_* \rangle|^{p-1} \|x_*\|} \\
&\geq 1 - \frac{2\|\mathbf{G}\|_{\text{inj}} \|u\|^{p-1}}{\lambda |\langle u, x_* \rangle|^{p-1} \|x_*\|} \\
&\geq 1 - \frac{2\|\mathbf{G}\|_{\text{inj}}}{\lambda \tau^{p-1} \|x_*\|^p}.
\end{aligned}$$

Our prior $P_{\mathbf{x}}$ is supported on the sphere of radius $\sqrt{n}$, so $\|x_*\| = \sqrt{n}$. We need to control the injective norm of the tensor $\mathbf{G}$. To this end we use Theorem 2.12 in [ABČ13] (see also Lemma 2.1 of [RM14]): there exists a constant $c(p) > 0$ (called $E_0(p)$ in [ABČ13]) such that for all $\epsilon > 0$,

$$\mathbb{P}\left(\sqrt{\frac{p}{n}} \|\mathbf{G}\|_{\text{inj}} \geq c(p) + \epsilon\right) \xrightarrow{n \rightarrow \infty} 0.$$

Letting $\epsilon = c(p)$ we obtain

$$\text{corr}(\hat{x}, x_*) \geq 1 - 4 \frac{c(p)}{\sqrt{p}} \frac{n^{(1-p)/2}}{\lambda \tau^{p-1}},$$

with probability tending to 1 as $n \rightarrow \infty$. □

E Computing the Kikuchi Hessian

In Section 4 we defined the Kikuchi free energy and explained the high level idea of how the symmetric difference matrices are derived from the Kikuchi Hessian. We now carry out the Kikuchi Hessian computation in full detail. This is a heuristic (non-rigorous) computation, but we believe these methods are important as we hope they will be useful for systematically obtaining optimal spectral methods for a wide variety of problems.

E.1 Derivatives of Kikuchi Free Energy

Following [SKZ14], we parametrize the beliefs in terms of the moments $m_S = \mathbb{E}[x^S]$. Specifically,

$$b_S(x_S) = \frac{1}{2^{|S|}} \left(1 + \sum_{\emptyset \subset T \subseteq S} m_T x^T \right). \quad (25)$$

We imagine m_T are close enough to zero so that b_S is a positive measure. One can check that these beliefs indeed have the prescribed moments: for $T \subseteq S$,

$$\sum_{x_S} b_S(x_S) x^T = m_T.$$

Thus we can think of the Kikuchi free energy $\mathcal{K}$ as a function of the moments $\{m_S\}_{0 \leq |S| \leq r}$. This parametrization forces the beliefs to be consistent, i.e., if $T \subseteq S$ then the marginal distribution $b_S|T$ is equal to b_T .

We now compute first and second derivatives of $\mathcal{K} = \mathcal{E} - \mathcal{S}$ with respect to the moments m_S . First, the energy term:

$$\begin{aligned}\mathcal{E} &= -\lambda \sum_{|S|=p} \mathbf{Y}_S m_S \\ \frac{\partial \mathcal{E}}{\partial m_S} &= \begin{cases} -\lambda \mathbf{Y}_S & \text{if } |S| = p \\ 0 & \text{otherwise} \end{cases} \\ \frac{\partial^2 \mathcal{E}}{\partial m_S \partial m_{S'}} &= 0.\end{aligned}$$

Now the entropy term:

$$\frac{\partial \mathcal{S}_S}{\partial b_S(x_S)} = -\log b_S(x_S) - 1.$$

From (25), for $\emptyset \subset T \subseteq S$,

$$\frac{\partial b_S(x_S)}{\partial m_T} = \frac{x^T}{2^{|S|}}$$

and so

$$\begin{aligned}\frac{\partial \mathcal{S}_S}{\partial m_T} &= \sum_{x_S} \frac{\partial \mathcal{S}_S}{\partial b_S(x_S)} \cdot \frac{\partial b_S(x_S)}{\partial m_T} \\ &= -2^{-|S|} \sum_{x_S} x^T [\log b_S(x_S) + 1] \\ &= -2^{-|S|} \sum_{x_S} x^T \log b_S(x_S).\end{aligned}\tag{26}$$

For $\emptyset \subset T \subseteq S$, $\emptyset \subset T' \subseteq S$,

$$\begin{aligned}\frac{\partial^2 \mathcal{S}_S}{\partial m_T \partial m_{T'}} &= -2^{-|S|} \sum_{x_S} x^T b_S(x_S)^{-1} \cdot \frac{\partial b_S(x_S)}{\partial m_{T'}} \\ &= -2^{-2|S|} \sum_{x_S} x^T x^{T'} b_S(x_S)^{-1} \\ &= -2^{-2|S|} \sum_{x_S} x^{T \triangle T'} b_S(x_S)^{-1}.\end{aligned}$$

Finally, if $T \not\subseteq S$ then $\frac{\partial \mathcal{S}_S}{\partial m_T} = 0$.

E.2 The Case $r = p$

We first consider the simplest case, where r is as small as possible: $r = p$. (We need to require $r \geq p$ in order to express the energy term in terms of the beliefs.)

E.2.1 Trivial Stationary Point

There is a “trivial stationary point” of the Kikuchi free energy where the beliefs only depend on local information. Specifically, if $|S| < p$ then b_S is the uniform distribution over $\{\pm 1\}^{|S|}$, and if $|S| = p$ then

$$b_S(x_S) \propto \exp(\lambda \mathbf{Y}_S x^S)$$

i.e.,

$$b_S(x_S) = \frac{1}{Z_S} \exp(\lambda \mathbf{Y}_S x^S) \quad (27)$$

where

$$Z_S = \sum_{x_S} \exp(\lambda \mathbf{Y}_S x^S).$$

Note that these beliefs are consistent (if $T \subseteq S$ with $|S| \leq p$ then $b_S|_T = b_T$) and so there is a corresponding set of moments $\{m_S\}_{|S| \leq p}$.

We now check that this is indeed a stationary point of the Kikuchi free energy. Using (26) and (27) we have for $\emptyset \subset T \subseteq S$ and $|S| \leq p$,

$$\begin{aligned} \frac{\partial \mathcal{S}_S}{\partial m_T} &= -2^{-|S|} \sum_{x_S} x^T \log b_S(x_S) \\ &= -2^{-|S|} \sum_{x_S} x^T [-\log Z_S + \lambda \mathbf{1}_{|S|=p} \mathbf{Y}_S x^S] \\ &= \begin{cases} -\lambda \mathbf{Y}_T & \text{if } |T| = p \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Thus if $|T| < p$ we have $\frac{\partial \mathcal{K}}{\partial m_T} = 0$, and if $|T| = p$ we have

$$\begin{aligned} \frac{\partial \mathcal{K}}{\partial m_T} &= \frac{\partial}{\partial m_T} \left[\mathcal{E} - \sum_{\emptyset \subset S \subseteq [p]} c_S \mathcal{S}_S \right] \\ &= -\lambda \mathbf{Y}_T + c_T \lambda \mathbf{Y}_T \\ &= 0 \end{aligned}$$

This confirms that we indeed have a stationary point.

E.2.2 Hessian

We now compute the Kikuchi Hessian, the matrix indexed by subsets $\emptyset \subset T \subseteq [p]$ with entries $\mathbf{H}_{T,T'} = \frac{\partial^2 \mathcal{K}}{\partial m_T \partial m_{T'}}$, evaluated at the trivial stationary point. Similarly to the Bethe Hessian [SKZ14], we expect the bottom eigenvector of the Kikuchi Hessian to be a good estimate for the (moments of) the true signal. This is because this bottom eigenvector indicates the best local direction for improving the Kikuchi free energy, starting from the trivial stationary point. If all eigenvalues of $\mathbf{H}$ are positive then the trivial stationary point is a local minimum and so an algorithm acting locally of the beliefs should not be able to escape from it, and should not learn anything about the signal. On the other hand, a negative eigenvalue (or even an eigenvalue close to zero) indicates a (potential) direction for improvement.

Remark E.1. When p is odd, we cannot hope for a substantially negative eigenvalue because x_* and $-x_*$ are *not* equally-good solutions and so the Kikuchi free energy should be locally cubic instead of quadratic. Still, we believe that the bottom eigenvector of the Kikuchi Hessian (which will have eigenvalue close to zero) yields a good algorithm. For instance, we will see in the next section that this method yields a close variant of tensor unfolding when $r = p = 3$.

Recall that for $\emptyset \subset T \subseteq S$, and $\emptyset \subset T' \subseteq S$,

$$\frac{\partial^2 \mathcal{S}_S}{\partial m_T \partial m_{T'}} = -2^{-2|S|} \sum_{x_S} x^{T \Delta T'} b_S(x_S)^{-1}.$$

If $|S| < p$ then b_S is uniform on $\{\pm 1\}^{|S|}$ (at the trivial stationary point) and so $\frac{\partial^2 \mathcal{S}_S}{\partial m_T \partial m_{T'}} = -\mathbb{1}_{T=T'}$. If $|S| = p$ then $b_S(x_S) = \frac{1}{Z_S} \exp(\lambda \mathbf{Y}_S x^S)$ where $Z_S = \sum_{x_S} \exp(\lambda \mathbf{Y}_S x^S) = 2^{|S|} \cosh(\lambda \mathbf{Y}_S)$, and so

$$\begin{aligned} \frac{\partial^2 \mathcal{S}_S}{\partial m_T \partial m_{T'}} &= -2^{-2|S|} \sum_{x_S} x^{T \Delta T'} b_S(x_S)^{-1} \\ &= -2^{-2|S|} \sum_{x_S} x^{T \Delta T'} Z_S \exp(-\lambda \mathbf{Y}_S x^S) \\ &= -2^{-|S|} \sum_{x_S} x^{T \Delta T'} \cosh(\lambda \mathbf{Y}_S) \left(1 - \lambda \mathbf{Y}_S x^S + \frac{1}{2!} \lambda^2 \mathbf{Y}_S^2 - \frac{1}{3!} \lambda^3 \mathbf{Y}_S^3 x^S + \dots \right) \\ &= -2^{-|S|} \sum_{x_S} x^{T \Delta T'} \cosh(\lambda \mathbf{Y}_S) (\cosh(\lambda \mathbf{Y}_S) - \sinh(\lambda \mathbf{Y}_S) x^S) \end{aligned}$$

since $\cosh(x) = 1 + \frac{1}{2!}x^2 + \frac{1}{4!}x^4 + \dots$ and $\sinh(x) = x + \frac{1}{3!}x^3 + \frac{1}{5!}x^5 + \dots$

$$\begin{aligned} &= \begin{cases} -\cosh^2(\lambda \mathbf{Y}_S) & \text{if } T = T', T \subseteq S \\ \cosh(\lambda \mathbf{Y}_S) \sinh(\lambda \mathbf{Y}_S) & \text{if } T \Delta T' = S, T \cup T' \subseteq S \\ 0 & \text{otherwise} \end{cases} \\ &= \begin{cases} -\cosh^2(\lambda \mathbf{Y}_S) & \text{if } T = T', T \subseteq S \\ \cosh(\lambda \mathbf{Y}_S) \sinh(\lambda \mathbf{Y}_S) & \text{if } T \sqcup T' = S \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

where $\sqcup$ denotes disjoint union. (Note that we have replaced Δ with $\sqcup$ due to the restriction $T, T' \subseteq S$.) We can now compute the Hessian:

$$\begin{aligned} \mathbf{H}_{T, T'} &= \frac{\partial^2 \mathcal{K}}{\partial m_T \partial m_{T'}} \\ &= - \sum_{\substack{S \supseteq T \cup T' \\ |S| \leq p}} c_S \frac{\partial^2 \mathcal{S}_S}{\partial m_T \partial m_{T'}} \\ &= \begin{cases} \sum_{\substack{S \supseteq T \\ |S| < p}} c_S + \sum_{\substack{S \supseteq T \\ |S| = p}} \cosh^2(\lambda \mathbf{Y}_S) & \text{if } T = T' \\ -\cosh(\lambda \mathbf{Y}_{T \cup T'}) \sinh(\lambda \mathbf{Y}_{T \cup T'}) & \text{if } |T \sqcup T'| = p \\ 0 & \text{otherwise} \end{cases} \\ &= \begin{cases} 1 + \sum_{\substack{S \supseteq T \\ |S| = p}} [\cosh^2(\lambda \mathbf{Y}_S) - 1] & \text{if } T = T' \\ -\cosh(\lambda \mathbf{Y}_{T \cup T'}) \sinh(\lambda \mathbf{Y}_{T \cup T'}) & \text{if } |T \sqcup T'| = p \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{28}$$

where we used (11) in the last step. Suppose $\lambda \ll 1$ (since otherwise tensor PCA is very easy). If $T = T'$ then, using the cosh Taylor series, we have the leading-order approximation

$$\begin{aligned} \mathbf{H}_{T,T} &\approx 1 + \sum_{\substack{S \supset T \\ |S|=p}} \lambda^2 \mathbf{Y}_S^2 \\ &\approx 1 + \binom{n-|T|}{p-|T|} \lambda^2 \mathbb{E}[\mathbf{Y}_S^2] \\ &\approx 1 + \frac{n^{p-|T|}}{(p-|T|)!} \lambda^2. \end{aligned}$$

This means $\mathbf{H} \approx \tilde{\mathbf{H}}$ where

$$\tilde{\mathbf{H}}_{T,T'} = \begin{cases} 1 \vee \frac{n^{p-|T|}}{(p-|T|)!} \lambda^2 & \text{if } T = T' \\ -\lambda \mathbf{Y}_{T \sqcup T'} & \text{if } |T \sqcup T'| = p \\ 0 & \text{otherwise.} \end{cases}$$

E.2.3 The Case $r = p = 3$

We now restrict to the case $r = p = 3$ and show that the Kikuchi Hessian recovers (a close variant of) the tensor unfolding method. Recall that in this case the computational threshold is $\lambda \sim n^{-3/4}$ and so we can assume $\lambda \ll n^{-1/2}$ (or else the problem is easy). We have

$$\tilde{\mathbf{H}}_{T,T'} = \begin{cases} \frac{1}{2} n^2 \lambda^2 & \text{if } T = T' \text{ with } |T| = 1 \\ 1 & \text{if } T = T' \text{ with } |T| \in \{2, 3\} \\ -\lambda \mathbf{Y}_{T \sqcup T'} & \text{if } |T \sqcup T'| = 3 \\ 0 & \text{otherwise.} \end{cases}$$

This means we can write

$$\tilde{\mathbf{H}} = \begin{pmatrix} \alpha I & -\lambda \mathbf{M} & 0 \\ -\lambda \mathbf{M}^\top & I & 0 \\ 0 & 0 & I \end{pmatrix}$$

where $\alpha = \frac{1}{2} n^2 \lambda^2$ and $\mathbf{M}$ is the $n \times \binom{n}{2}$ flattening of $\mathbf{Y}$, i.e., $\mathbf{M}_{i,\{j,k\}} = \mathbf{Y}_{ijk} \mathbf{1}_{\{i,j,k \text{ distinct}\}}$.

Since we are looking for the minimum eigenvalue of $\tilde{\mathbf{H}}$, we can restrict ourselves to the submatrix $\tilde{\mathbf{H}}^{\leq 2}$ indexed by sets of size 1 and 2. We have

$$\tilde{\mathbf{H}}^{\leq 2} = \begin{pmatrix} \alpha I & -\lambda \mathbf{M} \\ -\lambda \mathbf{M}^\top & I \end{pmatrix}.$$

An eigenvector $[u \ v]^\top$ of $\tilde{\mathbf{H}}^{\leq 2}$ with eigenvalue β satisfies

$$\alpha u - \lambda \mathbf{M} v = \beta u \quad \text{and} \quad -\lambda \mathbf{M}^\top u + v = \beta v$$

which implies $(1 - \beta)v = \lambda \mathbf{M}^\top u$ and so $\lambda^2 \mathbf{M} \mathbf{M}^\top u = (\alpha - \beta)(1 - \beta)u$. This means either u is an eigenvector of $\lambda^2 \mathbf{M} \mathbf{M}^\top$ with eigenvalue $(\alpha - \beta)(1 - \beta)$, or $u = 0$ and $\beta \in \{1, \alpha\}$. Conversely, if u is an eigenvector of $\lambda^2 \mathbf{M} \mathbf{M}^\top$ with eigenvalue $(\alpha - \beta)(1 - \beta) \neq 0$, then $[u \ v]^\top$ with $v = (1 - \beta)^{-1} \lambda \mathbf{M}^\top u$ is an eigenvector of $\tilde{\mathbf{H}}^{\leq 2}$ with eigenvalue β . Letting $\mu_1 > \dots > \mu_n > 0$ be the eigenvalues of $\lambda^2 \mathbf{M} \mathbf{M}^\top$, $\tilde{\mathbf{H}}^{\leq 2}$ has $2n$ eigenvalues of the form

$$\frac{\alpha + 1 \pm \sqrt{(\alpha - 1)^2 + 4\mu_i}}{2}$$

and the remaining eigenvalues are α or 1. Thus, the u -part of the bottom eigenvector of $\tilde{\mathbf{H}}^{\leq 2}$ is precisely the leading eigenvector of $\mathbf{M}\mathbf{M}^\top$. This is a close variant of the tensor unfolding spectral method (see Section 2.4), and we expect that its performance is essentially identical.

E.2.4 The General Case: $r \geq p$

One difficulty when $r > p$ is that there is no longer a trivial stationary point that we can write down in closed form. There is, however, a natural guess for “uninformative” beliefs that only depend on the local information: for $0 < |S| \leq r$,

$$b_S(x_S) = \frac{1}{Z_S} \exp \left(\lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right)$$

for the appropriate normalizing factor Z_S . Unfortunately, these beliefs are not quite consistent, so we need separate moments for each set S :

$$m_T^{(S)} = \mathbb{E}_{x_S \sim b_S} [x^T].$$

Provided $\lambda \ll 1$, we can check that $m_T^{(S)} \approx m_T^{(S')}$ to first order, and so the above beliefs are at least approximately consistent:

$$Z_S = \sum_{x_S} \exp \left(\lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right) \approx \sum_{x_S} \left(1 + \lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right) = 2^{|S|}$$

and so

$$\begin{aligned} m_T^{(S)} &= \sum_{x_S} b_S(x_S) x^T \\ &= \frac{1}{Z_S} \sum_{x_S} x^T \exp \left(\lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right) \\ &\approx \frac{1}{Z_S} \sum_{x_S} x^T \left(1 + \lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right) \\ &= \begin{cases} \lambda \mathbf{Y}_T & \text{if } |T| = p \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

which does not depend on S . Thus we will ignore the slight inconsistencies and carry on with the

derivation. As above, the important calculation is, for $T, T' \subseteq S$,

$$\begin{aligned}
\frac{\partial^2 \mathcal{S}_S}{\partial m_T^{(S)} \partial m_{T'}^{(S)}} &= -2^{-2|S|} \sum_{x_S} x^{T \triangle T'} b_S(x_S)^{-1} \\
&= -2^{-2|S|} \sum_{x_S} x^{T \triangle T'} Z_S \exp \left(-\lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right) \\
&\approx -2^{-|S|} \sum_{x_S} x^{T \triangle T'} \exp \left(-\lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right) \\
&\approx -2^{-|S|} \sum_{x_S} x^{T \triangle T'} \left(1 - \lambda \sum_{\substack{U \subseteq S \\ |U|=p}} \mathbf{Y}_U x^U \right) \\
&= \begin{cases} -1 & \text{if } T = T' \\ \lambda \mathbf{Y}_{T \triangle T'} & \text{if } |T \triangle T'| = p \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}$$

Analogous to (28), we compute the Kikuchi Hessian

$$\mathbf{H}_{T,T'} := - \sum_{\substack{S \supseteq T \cup T' \\ |S| \leq r}} c_S \frac{\partial^2 \mathcal{S}_S}{\partial m_T^{(S)} \partial m_{T'}^{(S)}}.$$

If we fix ℓ_1, ℓ_2 , the submatrix $\mathbf{H}^{(\ell_1, \ell_2)} = (\mathbf{H}_{T,T'})_{|T|=\ell_1, |T'|=\ell_2}$ takes the form

$$\mathbf{H}^{(\ell_1, \ell_2)} \approx a(\ell_1, \ell_2) \mathbf{1}_{\ell_1=\ell_2} \mathbf{I} - b(\ell_1, \ell_2) \mathbf{M}^{(\ell_1, \ell_2)}$$

for certain scalars $a(\ell_1, \ell_2)$ and $b(\ell_1, \ell_2)$, where $\mathbf{I}$ is the identity matrix and $\mathbf{M}^{(\ell_1, \ell_2)} \in \mathbb{R}^{\binom{[n]}{\ell_1} \times \binom{[n]}{\ell_2}}$ is the symmetric difference matrix

$$\mathbf{M}_{T,T'}^{(\ell_1, \ell_2)} = \begin{cases} \mathbf{Y}_{T \triangle T'} & \text{if } |T \triangle T'| = p \\ 0 & \text{otherwise.} \end{cases}$$

Instead of working with the entire Kikuchi Hessian, we choose to work instead with $\mathbf{M}^{(\ell, \ell)}$, which (when p is even) appears as a diagonal block of the Kikuchi Hessian whenever $r \geq \ell + p/2$ (since there must exist $|T| = |T'| = \ell$ with $|T \cup T'| \leq r$ and $|T \triangle T'| = p$). Our theoretical results (see Section 3) show that indeed $\mathbf{M}^{(\ell, \ell)}$ yields algorithms matching the (conjectured optimal) performance of sum-of-squares. When p is odd, $\mathbf{M}^{(\ell, \ell)} = 0$ and so we propose to instead focus on $\mathbf{M}^{(\ell, \ell+1)}$; see Appendix C.

F Extensions

F.1 Refuting Random k -XOR Formulas for k Even

Our symmetric difference matrices can be used to give a simple algorithm and proof for a related problem: strongly refuting random k -XOR formulas (see [AOW15, RRS17] and references therein).

This is essentially a variant of the spiked tensor problem with sparse Rademacher observations instead of Gaussian ones. It is known [RRS17] that this problem exhibits a smooth tradeoff between subexponential runtime and the number of constraints required, but the proof of [RRS17] involves intensive moment calculations. When k is even, we will give a simple algorithm and a simple proof using the matrix Chernoff bound that achieves the same tradeoff. SOS lower bounds suggest that this tradeoff is optimal [Gri01, Sch08].

When k is odd, we expect that the construction in Section F.2 should achieve the optimal tradeoff, but we do not have a proof for this case.

F.1.1 Setup

Let $x_1, \dots, x_n$ be $\{\pm 1\}$ -valued variables. A k -XOR formula Φ with m constraints is specified by a sequence of subsets $U_1, \dots, U_m$ with $U_i \subseteq [n]$ and $|U_i| = k$, along with values $b_1, \dots, b_m$ with $b_i \in \{\pm 1\}$. For $1 \leq i \leq m$, constraint i is satisfied if $x^{U_i} = b_i$, where $x^{U_i} := \prod_{j \in U_i} x_j$. We write $P_\Phi(x)$ for the number of constraints satisfied by x . We will consider a uniformly random k -XOR formula in which each U_i is chosen uniformly and independently from the $\binom{n}{k}$ possible k -subsets, and each b_i is chosen uniformly and independently from $\{\pm 1\}$.

Given a formula Φ , the goal of strong refutation is to certify an upper bound on the number of constraints that can be satisfied. In other words, our algorithm should output a bound $B = B(\Phi)$ such that for *every* formula Φ , $\max_{x \in \{\pm 1\}^n} P_\Phi(x) \leq B(\Phi)$. (Note that this must be satisfied *always*, not merely with high probability.) At the same time, we want the bound B to be small with high probability over a random Φ . Since a random assignment x will satisfy roughly half the constraints, the best bound we can hope for is $B = \frac{m}{2}(1 + \varepsilon)$ with $\varepsilon > 0$ small.

F.1.2 Algorithm

Let $k \geq 2$ be even and let $\ell \geq k/2$. Given a k -XOR formula Φ , construct the order- ℓ symmetric difference matrix $\mathbf{M} \in \mathbb{R}^{\binom{[n]}{\ell} \times \binom{[n]}{\ell}}$ as follows. For $S, T \subseteq [n]$ with $|S| = |T| = \ell$, let

$$\mathbf{M}_{S,T}^{(i)} = \begin{cases} b_i & \text{if } S \triangle T = U_i \\ 0 & \text{otherwise} \end{cases}$$

and let

$$\mathbf{M} = \sum_{i=1}^m \mathbf{M}^{(i)}.$$

Define the parameter

$$d_\ell := \binom{n-\ell}{k/2} \binom{\ell}{k/2},$$

which, for any fixed $|S| = \ell$, is the number of sets $|T| = \ell$ such that $|S \triangle T| = k$. For an assignment $x \in \{\pm 1\}^n$, let $u^x \in \mathbb{R}^{\binom{[n]}{\ell}}$ be defined by $u_S^x = x^S$ for all $|S| = \ell$. We have

$$\|\mathbf{M}\| \geq \frac{(u^x)^\top \mathbf{M} u^x}{\|u^x\|^2} = \binom{n}{\ell}^{-1} \sum_{i=1}^m \sum_{S \triangle T = U_i} x^{U_i} b_i = d_\ell \binom{n}{k}^{-1} (2P_\Phi(x) - m)$$

since for any fixed U_i (of size k), the number of (S, T) pairs such that $S \triangle T = U_i$ is $\binom{n}{\ell} d_\ell \binom{n}{k}^{-1}$. Thus we can perform strong refutation by computing $\|\mathbf{M}\|$:

$$P_\Phi(x) \leq \frac{m}{2} + \frac{1}{2d_\ell} \binom{n}{k} \|\mathbf{M}\|. \quad (29)$$

Theorem F.1. *Let $k \geq 2$ be even and let $k/2 \leq \ell \leq n - k/2$. Let $\beta \in (0, 1)$. If*

$$m \geq \frac{4e^2 \binom{n}{k} \log \binom{n}{\ell}}{\beta^2 d_\ell} \quad (30)$$

then $\|\mathbf{M}\|$ certifies

$$P_\Phi(x) \leq \frac{m}{2}(1 + \beta)$$

with probability at least $1 - 3\binom{n}{\ell}^{-1}$ over a uniformly random k -XOR formula Φ with m constraints.

If k is constant and $\ell = n^\delta$ with $\delta \in (0, 1)$, the condition (30) becomes

$$m \geq O(\beta^{-2} n^{k/2} \ell^{1-k/2} \log n) = O(\beta^{-2} n^{k/2+\delta(1-k/2)} \log n),$$

matching the result of [RRS17]. In fact, our result is tighter by polylog factors.

F.1.3 Binomial Tail Bound

The main ingredients in the proof of Theorem F.1 will be the matrix Chernoff bound (Theorem A.4) and the following standard binomial tail bound.

Proposition F.2. *Let $X \sim \text{Binomial}(n, p)$. For $p < \frac{u}{n} < 1$,*

$$\mathbb{P}(X \geq u) \leq \exp \left[-u \left(\log \left(\frac{u}{pn} \right) - 1 \right) \right].$$

Proof. We begin with the standard Binomial tail bound [AG89]

$$\mathbb{P}(X \geq u) \leq \exp \left(-n D \left(\frac{u}{n} \parallel p \right) \right)$$

for $p < \frac{u}{n} < 1$, where

$$D(a \parallel p) := a \log \left(\frac{a}{p} \right) + (1 - a) \log \left(\frac{1 - a}{1 - p} \right).$$

Since $\log(x) \geq 1 - 1/x$,

$$(1 - a) \log \left(\frac{1 - a}{1 - p} \right) \geq (1 - a) \left(1 - \frac{1 - p}{1 - a} \right) = p - a \geq -a,$$

and the desired result follows. $\square$

F.1.4 Proof

Proof of Theorem F.1. We need to bound $\|\mathbf{M}\|$ with high probability over a uniformly random k -XOR formula Φ . First, fix the subsets $U_1, \dots, U_m$ and consider the randomness of the signs b_i . We can write $\mathbf{M}$ as a Rademacher series

$$\mathbf{M} = \sum_{i=1}^m b_i \mathbf{A}^{(i)}$$

where

$$\mathbf{A}_{S,T}^{(i)} = \mathbb{1}_{S \Delta T = U_i}.$$

By the matrix Chernoff bound (Theorem A.4),

$$\mathbb{P}(\|\mathbf{M}\| \geq t) \leq 2 \binom{n}{\ell} e^{-t^2/2\sigma^2} = 2 \exp \left(\log \binom{n}{\ell} - \frac{t^2}{2\sigma^2} \right)$$

where

$$\sigma^2 = \left\| \sum_{i=1}^m (\mathbf{A}^{(i)})^2 \right\|.$$

In particular,

$$\mathbb{P} \left(\|\mathbf{M}\| \geq 2 \sqrt{\sigma^2 \log \binom{n}{\ell}} \right) \leq 2 \binom{n}{\ell}^{-1}. \quad (31)$$

Now we will bound σ^2 with high probability over the random choice of $U_1, \dots, U_m$. We have

$$\sum_{i=1}^m (\mathbf{A}^{(i)})^2 = \text{diag}(D)$$

where D_S is the number of i for which $|S \triangle U_i| = \ell$. This means $\sigma^2 = \max_{|S|=\ell} D_S$. For fixed $S \subseteq [n]$ with $|S| = \ell$, the number of sets $U \subseteq [n]$ with $|U| = k$ such that $|S \triangle U| = \ell$ is d_ℓ and so $D_S \sim \text{Binomial}(m, p)$ with $p := d_\ell \binom{n}{k}^{-1}$. Using the Binomial tail bound (Proposition F.2) and a union bound over S ,

$$\mathbb{P}(\sigma^2 \geq u) \leq \binom{n}{\ell} \exp \left[-u \left(\log \left(\frac{u}{pm} \right) - 1 \right) \right] = \exp \left[\log \binom{n}{\ell} - u \left(\log \left(\frac{u}{pm} \right) - 1 \right) \right].$$

Provided

$$\frac{u}{pm} \geq e^2 \quad \text{and} \quad u \geq 2 \log \binom{n}{\ell}, \quad (32)$$

we have

$$\mathbb{P}(\sigma^2 \geq u) \leq \binom{n}{\ell}^{-1}.$$

Let $\beta \in (0, 1)$. From (29), to certify $P_\Phi(x) \leq \frac{m}{2}(1+\beta)$ it suffices to have $\|\mathbf{M}\| \leq \beta m d_\ell \binom{n}{k}^{-1} = \beta pm$. Therefore, from (31), it suffices to have

$$\sigma^2 \leq \frac{\beta^2 p^2 m^2}{4 \log \binom{n}{\ell}}.$$

From (32), this will occur provided

$$\frac{\beta^2 p^2 m^2}{4 \log \binom{n}{\ell}} \geq e^2 pm \Leftrightarrow m \geq \frac{4e^2 \log \binom{n}{\ell}}{\beta^2 p} \quad (33)$$

and

$$\frac{\beta^2 p^2 m^2}{4 \log \binom{n}{\ell}} \geq 2 \log \binom{n}{\ell} \Leftrightarrow m \geq \frac{2\sqrt{2} \log \binom{n}{\ell}}{\beta p}. \quad (34)$$

Note that (34) is subsumed by (33). This completes the proof. $\square$

F.2 Odd-Order Tensors

When the tensor order p is odd, we have given an algorithm for tensor PCA based on the Kikuchi Hessian (see Appendix C) but are unfortunately unable to give a tight analysis of it. Here we present a related algorithm for which we are able to give a better analysis, matching SOS. The idea of the algorithm is to use a construction from the SOS literature that transforms an order- p tensor (with p odd) into an order- $2(p-1)$ tensor via the Cauchy–Schwarz inequality [CGL04]. We then apply a variant of our symmetric difference matrix to the resulting even-order tensor. A similar construction was given independently in the recent work [Has19] and shown to give optimal performance for all $\ell \leq n^\delta$ for a certain constant $\delta > 0$. The proof we give here applies to the full range of ℓ values: $\ell \ll n$. Our proof uses a certain variant of the matrix Bernstein inequality combined with some fairly simple moment calculations.

F.2.1 Setup

For simplicity, we consider the following version of the problem. Let $p \geq 3$ be odd and let $\mathbf{Y} \in (\mathbb{R}^n)^{\otimes p}$ be an asymmetric tensor with i.i.d. Rademacher (uniform ± 1) entries. Our goal is to certify an upper bound on the *Rademacher injective norm*, defined as

$$\|\mathbf{Y}\|_{\pm} := \max_{x \in \{\pm 1\}^n / \sqrt{n}} |\langle \mathbf{Y}, x^{\otimes p} \rangle|.$$

The true value is $O(\sqrt{n})$ with high probability. In time n^ℓ (where $\ell = n^\delta$ with $\delta \in (0, 1)$) we will certify the bound $\|\mathbf{Y}\|_{\pm} \leq n^{p/4} \ell^{1/2-p/4} \text{polylog}(n)$, matching the results of [BGG⁺16, BGL16]. Such certification results can be turned into recovery results using sum-of-squares; see Lemma 4.4 of [HSS15]. To certify a bound on the *injective norm* instead of the Rademacher injective norm (where x is constrained to the sphere instead of the hypercube), one should use the basis-invariant version of the symmetric difference matrices given by [Has19] (but we do not do this here).

F.2.2 Algorithm

We will use a trick from [CGL04] which is often used in the sum-of-squares literature. For any $\|x\| = 1$, we have by the Cauchy–Schwarz inequality,

$$\langle \mathbf{Y}, x^{\otimes p} \rangle^2 \leq \|x\|^2 \langle \mathbf{T}, x^{\otimes 4q} \rangle = \langle \mathbf{T}, x^{\otimes 4q} \rangle$$

where $p = 2q + 1$ and $\mathbf{T}_{abcd} := \sum_{e \in [n]} \mathbf{Y}_{ace} \mathbf{Y}_{bde}$ where $a, b, c, d \in [n]^q$. We have $\mathbb{E}[\mathbf{T}]_{abcd} = n \cdot \mathbf{1}\{ac = bd\}$ and so $\langle \mathbb{E}[\mathbf{T}], x^{\otimes 4q} \rangle = n \sum_{ac} (x^a x^c)^2 = n$. Let $\tilde{\mathbf{T}} = \mathbf{T} - \mathbb{E}[\mathbf{T}]$, i.e., $\tilde{\mathbf{T}}_{abcd} = \mathbf{T}_{abcd} \cdot \mathbf{1}\{ac \neq bd\}$. Define the $n^\ell \times n^\ell$ matrix $\mathbf{M}$ as follows. For $S, T \in [n]^\ell$,

$$\mathbf{M}_{S,T} := \sum_{abcd} \tilde{\mathbf{T}}_{abcd} N_{ab,cd}^{-1} \cdot \mathbf{1}\{S \xleftrightarrow{ab,cd} T\}$$

where $S \xleftrightarrow{ab,cd} T$ roughly means that S is obtained from T by replacing ab by cd , or cd by ab ; the formal definition is given below. Also, $N_{ab,cd}$ denotes the number of (S, T) pairs for which $S \xleftrightarrow{ab,cd} T$.

Definition F.3. For $S, T \in [n]^\ell$ and $a, b, c, d \in [n]^q$, we write $S \xleftrightarrow{ab,cd} T$ if there are distinct indices $i_1, \dots, i_{2q} \in [\ell]$ such that either: (i) $S_{i_j} = (ab)_{i_j}$ and $T_{i_j} = (cd)_{i_j}$ for all $j \in [2q]$, the values in a, b, c, d do not appear anywhere else in S or T , and S, T are identical otherwise: $S_i = T_i$ for all $i \notin \{i_1, \dots, i_{2q}\}$; or (ii) the same holds but with ab and cd interchanged. (Here ab denotes concatenation.)

Note that

$$N_{ab,cd} \geq \bar{N} := \binom{\ell}{2q} (n - 4q)^{\ell-2q}. \quad (35)$$

The above construction ensures that

$$n^\ell (x^{\otimes \ell})^\top \mathbf{M} (x^{\otimes \ell}) = n^{2q} \langle \tilde{\mathbf{T}}, x^{\otimes 4q} \rangle \quad \text{for all } x \in \{\pm 1\}^n / \sqrt{n}.$$

This means we can certify an upper bound on $\|\mathbf{Y}\|_\pm$ by computing $\|\mathbf{M}\|$:

$$\|\mathbf{Y}\|_\pm \leq \sqrt{\langle \mathbf{T}, x^{\otimes 4q} \rangle} \leq \sqrt{\langle \mathbb{E}[\mathbf{T}], x^{\otimes 4q} \rangle} + \sqrt{\langle \tilde{\mathbf{T}}, x^{\otimes 4q} \rangle} \leq n^{1/2} + n^{\ell/2-q} \|\mathbf{M}\|^{1/2}.$$

Theorem F.4. *Let $p \geq 3$ be odd and let $p-1 \leq \ell \leq \min\{n - (p-1), \frac{n}{4(p-1)}, \frac{n}{8 \log n}\}$. Then $\|\mathbf{M}\|$ certifies*

$$\|\mathbf{Y}\|_\pm \leq n^{1/2} + 8p^p \ell^{1/2-p/4} n^{p/4} (\log n)^{1/4}$$

with probability at least $1 - n^{-\ell}$ over an i.i.d. Rademacher $\mathbf{Y}$.

F.2.3 Proof

We will use the following variant of the matrix Bernstein inequality; this is a special case ($\mathbf{A}_k = R \cdot \mathbf{I}$) of [Tro12], Theorem 6.2.

Theorem F.5 (Matrix Bernstein). *Consider a finite sequence $\{\mathbf{X}_i\}$ of independent random symmetric $d \times d$ matrices. Suppose $\mathbb{E}[\mathbf{X}_i] = 0$ and $\|\mathbb{E}[\mathbf{X}_i^r]\| \leq \frac{r!}{2} R^r$ for $r = 2, 3, 4, \dots$. Then*

$$\Pr \left\{ \left\| \sum_{i=1}^n \mathbf{X}_i \right\| \geq t \right\} \leq d \cdot \exp \left(\frac{-t^2/2}{nR^2 + Rt} \right).$$

For $e \in [n]$, let

$$\mathbf{M}_{S,T}^{(e)} := \sum_{abcd} \mathbf{Y}_{ace} \mathbf{Y}_{bde} N_{ab,cd}^{-1} \cdot \mathbf{1}\{S \xleftrightarrow{ab,cd} T\} \cdot \mathbf{1}\{ac \neq bd\}.$$

We will apply Theorem F.5 to the sum $\mathbf{M} = \sum_e \mathbf{M}^{(e)}$. Note that $\mathbb{E}[\mathbf{M}^{(e)}] = 0$. To bound the moments $\|\mathbb{E}[(\mathbf{M}^{(e)})^r]\|$, we will use the following basic fact.

Lemma F.6. *If $\mathbf{A}$ is a symmetric matrix,*

$$\|\mathbf{A}\| \leq \max_j \sum_i |\mathbf{A}_{ij}|.$$

Proof. Let v be the leading eigenvector of $\mathbf{A}$ so that $\mathbf{A}v = \lambda v$ where $\|\mathbf{A}\| = |\lambda|$. Normalize v so that $\|v\|_1 = 1$. Then $\|\mathbf{A}v\|_1 = |\lambda| \cdot \|v\|_1$ and so

$$\|\mathbf{A}\| = |\lambda| \cdot \|v\|_1 = \sum_i \left| \sum_j \mathbf{A}_{ij} v_j \right| \leq \sum_{ij} |\mathbf{A}_{ij}| \cdot |v_j| \leq \sum_j |v_j| \cdot \sum_i |\mathbf{A}_{ij}| \leq \|v\|_1 \cdot \max_j \sum_i |\mathbf{A}_{ij}|.$$

□

Proof of Theorem F.4. For any fixed e , we have by Lemma F.6,

$$\|\mathbb{E}[(\mathbf{M}^{(e)})^r]\| \leq \max_S \sum_T |\mathbb{E}[(\mathbf{M}^{(e)})^r]_{S,T}| =: \max_S h(r, e, S).$$

Let π denote a “path” of the form

$$\pi = (S_0, a_1, b_1, c_1, d_1, S_1, a_2, b_2, c_2, d_2, S_2, \dots, a_r, b_r, c_r, d_r, S_r)$$

such that $S_0 = S$, $(a_i, c_i) \neq (b_i, d_i)$, and $S_{i-1} \xleftrightarrow{a_i b_i, c_i d_i} S_i$. Then we have

$$h(r, e, S) = \sum_{\pi} \mathbb{E} \prod_{i=1}^r \mathbf{Y}_{a_i c_i e} \mathbf{Y}_{b_i d_i e} N_{a_i b_i, c_i d_i}^{-1}.$$

Among tuples of the form (a_i, c_i) and (b_i, d_i) , each must occur an even number of times (or else the term associated with π is 0). There are $2r$ such tuples, so there are $\binom{2r}{r} r! 2^{-r}$ ways to pair them up. Once S_{i-1} is chosen, there are at most $2(\ell n)^q$ choices for (a_i, c_i) , and the same is true for (b_i, d_i) . Once $S_{i-1}, a_i, b_i, c_i, d_i$ are chosen, there are at most $(2q)!$ possible choices for S_i . This means

$$h(r, e, S) \leq \binom{2r}{r} r! 2^{-r} [2(\ell n)^q \cdot (2q)!]^r \bar{N}^{-r}.$$

where $\bar{N}$ is defined in (35). Since $\binom{2r}{r} \leq 4^r$, we can apply Theorem F.5 with $R = 8(2q)!(\ell n)^q \bar{N}^{-1}$. This yields

$$\Pr\{\|\mathbf{M}\| \geq t\} \leq n^\ell \cdot \exp\left(\frac{-t^2/2}{nR^2 + Rt}\right).$$

Let $t = R\sqrt{8\ell n \log n}$. Provided $\ell \leq n/(8 \log n)$ we have $Rt \leq nR^2$ and so

$$\Pr\left\{\|\mathbf{M}\| \geq R\sqrt{8\ell n \log n}\right\} \leq \exp\left(\ell \log n - \frac{t^2}{4nR^2}\right) = \exp(-\ell \log n) = n^{-\ell}.$$

Thus with high probability we certify

$$\begin{aligned} \|\mathbf{Y}\|_{\pm} &\leq n^{1/2} + n^{\ell/2-q} \|\mathbf{M}\|^{1/2} \\ &\leq n^{1/2} + n^{\ell/2-q} R^{1/2} (8\ell n \log n)^{1/4} \\ &= n^{1/2} + n^{\ell/2-q} \sqrt{8(2q)!} (\ell n)^{q/2} \bar{N}^{-1/2} (8\ell n \log n)^{1/4} \\ &= n^{1/2} + 8^{3/4} \sqrt{(2q)!} n^{\ell/2-q} \bar{N}^{-1/2} (\ell n)^{1/4+q/2} (\log n)^{1/4}. \end{aligned}$$

We have the following bound on $\bar{N}$:

$$\begin{aligned}
\bar{N} &= \binom{\ell}{2q} (n - 4q)^{\ell - 2q} \\
&= n^\ell \cdot \frac{\binom{\ell}{2q} (n - 4q)^{\ell - 2q}}{n^\ell} \\
&\geq n^\ell \cdot \frac{\ell^{2q}}{(2q)^{2q}} \cdot \frac{(n - 4q)^{\ell - 2q}}{n^\ell} \\
&= \frac{n^\ell}{(2q)^{2q}} \cdot \left(\frac{n - 4q}{n} \right)^{\ell - 2q} \left(\frac{\ell}{n} \right)^{2q} \\
&= \frac{n^\ell}{(2q)^{2q}} \cdot \left(1 - \frac{4q}{n} \right)^{\ell - 2q} \left(\frac{\ell}{n} \right)^{2q} \\
&\geq \frac{n^\ell}{p^p} \cdot \left(1 - (\ell - 2q) \frac{4q}{n} \right) \left(\frac{\ell}{n} \right)^{2q} \\
&\geq \frac{n^\ell}{p^p} \cdot \left(1 - \frac{4q\ell}{n} \right) \left(\frac{\ell}{n} \right)^{2q} \\
&\geq \frac{n^\ell}{2p^p} \left(\frac{\ell}{n} \right)^{2q}
\end{aligned}$$

provided $\ell \leq n/(8q)$. Therefore we certify

$$\begin{aligned}
\|\mathbf{Y}\|_{\pm} &\leq n^{1/2} + 2^{1/2} \cdot 8^{3/4} \sqrt{(2q)!} p^{p/2} \ell^{1/4+q/2} n^{1/4-q/2} (\ell/n)^{-q} (\log n)^{1/4} \\
&\leq n^{1/2} + 8p^p \ell^{1/2-p/4} n^{p/4} (\log n)^{1/4}
\end{aligned}$$

as desired. □